



# OpenShift Container Platform 4.12

## Installing on AWS

Installing OpenShift Container Platform on Amazon Web Services



# OpenShift Container Platform 4.12 Installing on AWS

---

Installing OpenShift Container Platform on Amazon Web Services

## 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> 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<sup>®</sup> 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 describes how to install OpenShift Container Platform on Amazon Web Services.

## Table of Contents

<b>CHAPTER 1. PREPARING TO INSTALL ON AWS</b> .....	<b>12</b>
1.1. PREREQUISITES	12
1.2. REQUIREMENTS FOR INSTALLING OPENSIFT CONTAINER PLATFORM ON AWS	12
1.3. CHOOSING A METHOD TO INSTALL OPENSIFT CONTAINER PLATFORM ON AWS	12
1.3.1. Installing a cluster on installer-provisioned infrastructure	12
1.3.2. Installing a cluster on user-provisioned infrastructure	13
1.4. NEXT STEPS	13
<b>CHAPTER 2. CONFIGURING AN AWS ACCOUNT</b> .....	<b>14</b>
2.1. CONFIGURING ROUTE 53	14
2.1.1. Ingress Operator endpoint configuration for AWS Route 53	14
2.2. AWS ACCOUNT LIMITS	15
2.3. REQUIRED AWS PERMISSIONS FOR THE IAM USER	17
2.4. CREATING AN IAM USER	25
2.5. IAM POLICIES AND AWS AUTHENTICATION	26
2.5.1. Default permissions for IAM instance profiles	27
2.5.2. Specifying an existing IAM role	28
2.5.3. Using AWS IAM Analyzer to create policy templates	29
2.6. SUPPORTED AWS MARKETPLACE REGIONS	30
2.7. SUPPORTED AWS REGIONS	30
2.7.1. AWS public regions	30
2.7.2. AWS GovCloud regions	31
2.7.3. AWS SC2S and C2S secret regions	31
2.7.4. AWS China regions	32
2.8. NEXT STEPS	32
<b>CHAPTER 3. MANUALLY CREATING IAM FOR AWS</b> .....	<b>33</b>
3.1. ALTERNATIVES TO STORING ADMINISTRATOR-LEVEL SECRETS IN THE KUBE-SYSTEM PROJECT	33
3.2. MANUALLY CREATE IAM	34
3.3. MINT MODE	37
3.4. MINT MODE WITH REMOVAL OR ROTATION OF THE ADMINISTRATOR-LEVEL CREDENTIAL	37
3.5. NEXT STEPS	37
<b>CHAPTER 4. INSTALLING A CLUSTER QUICKLY ON AWS</b> .....	<b>39</b>
4.1. PREREQUISITES	39
4.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	39
4.3. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	40
4.4. OBTAINING THE INSTALLATION PROGRAM	41
4.5. DEPLOYING THE CLUSTER	42
4.6. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	45
Installing the OpenShift CLI on Linux	45
Installing the OpenShift CLI on Windows	46
Installing the OpenShift CLI on macOS	46
4.7. LOGGING IN TO THE CLUSTER BY USING THE CLI	47
4.8. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	47
4.9. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	48
4.10. NEXT STEPS	48
<b>CHAPTER 5. INSTALLING A CLUSTER ON AWS WITH CUSTOMIZATIONS</b> .....	<b>50</b>
5.1. PREREQUISITES	50
5.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	50
5.3. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	51

5.4. OBTAINING AN AWS MARKETPLACE IMAGE	52
5.5. OBTAINING THE INSTALLATION PROGRAM	53
5.6. CREATING THE INSTALLATION CONFIGURATION FILE	54
5.6.1. Installation configuration parameters	55
5.6.1.1. Required configuration parameters	56
5.6.1.2. Network configuration parameters	57
5.6.1.3. Optional configuration parameters	59
5.6.1.4. Optional AWS configuration parameters	65
5.6.2. Minimum resource requirements for cluster installation	69
5.6.3. Tested instance types for AWS	70
5.6.4. Tested instance types for AWS on 64-bit ARM infrastructures	71
5.6.5. Sample customized install-config.yaml file for AWS	71
5.6.6. Configuring the cluster-wide proxy during installation	74
5.7. DEPLOYING THE CLUSTER	76
5.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	77
Installing the OpenShift CLI on Linux	78
Installing the OpenShift CLI on Windows	78
Installing the OpenShift CLI on macOS	79
5.9. LOGGING IN TO THE CLUSTER BY USING THE CLI	79
5.10. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	80
5.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	81
5.12. NEXT STEPS	81
<b>CHAPTER 6. INSTALLING A CLUSTER ON AWS WITH NETWORK CUSTOMIZATIONS</b>	<b>82</b>
6.1. PREREQUISITES	82
6.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	82
6.3. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	83
6.4. OBTAINING THE INSTALLATION PROGRAM	84
6.5. NETWORK CONFIGURATION PHASES	85
6.6. CREATING THE INSTALLATION CONFIGURATION FILE	86
6.6.1. Installation configuration parameters	87
6.6.1.1. Required configuration parameters	87
6.6.1.2. Network configuration parameters	89
6.6.1.3. Optional configuration parameters	91
6.6.1.4. Optional AWS configuration parameters	97
6.6.2. Minimum resource requirements for cluster installation	101
6.6.3. Tested instance types for AWS	102
6.6.4. Tested instance types for AWS on 64-bit ARM infrastructures	102
6.6.5. Sample customized install-config.yaml file for AWS	103
6.6.6. Configuring the cluster-wide proxy during installation	106
6.7. CLUSTER NETWORK OPERATOR CONFIGURATION	107
6.7.1. Cluster Network Operator configuration object	108
defaultNetwork object configuration	108
Configuration for the OpenShift SDN network plugin	109
Configuration for the OVN-Kubernetes network plugin	110
kubeProxyConfig object configuration	114
6.8. SPECIFYING ADVANCED NETWORK CONFIGURATION	115
6.9. CONFIGURING AN INGRESS CONTROLLER NETWORK LOAD BALANCER ON A NEW AWS CLUSTER	116
6.10. CONFIGURING HYBRID NETWORKING WITH OVN-KUBERNETES	117
6.11. DEPLOYING THE CLUSTER	119
6.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	121
Installing the OpenShift CLI on Linux	121

Installing the OpenShift CLI on Windows	122
Installing the OpenShift CLI on macOS	122
6.13. LOGGING IN TO THE CLUSTER BY USING THE CLI	123
6.14. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	123
6.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	124
6.16. NEXT STEPS	124
<b>CHAPTER 7. INSTALLING A CLUSTER ON AWS IN A RESTRICTED NETWORK</b> .....	<b>125</b>
7.1. PREREQUISITES	125
7.2. ABOUT INSTALLATIONS IN RESTRICTED NETWORKS	126
7.2.1. Additional limits	126
7.3. ABOUT USING A CUSTOM VPC	126
7.3.1. Requirements for using your VPC	126
Option 1: Create VPC endpoints	127
Option 2: Create a proxy without VPC endpoints	128
Option 3: Create a proxy with VPC endpoints	128
7.3.2. VPC validation	129
7.3.3. Division of permissions	130
7.3.4. Isolation between clusters	130
7.4. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	130
7.5. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	131
7.6. CREATING THE INSTALLATION CONFIGURATION FILE	132
7.6.1. Installation configuration parameters	135
7.6.1.1. Required configuration parameters	135
7.6.1.2. Network configuration parameters	136
7.6.1.3. Optional configuration parameters	138
7.6.1.4. Optional AWS configuration parameters	144
7.6.2. Minimum resource requirements for cluster installation	148
7.6.3. Sample customized install-config.yaml file for AWS	149
7.6.4. Configuring the cluster-wide proxy during installation	152
7.7. DEPLOYING THE CLUSTER	154
7.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	156
Installing the OpenShift CLI on Linux	156
Installing the OpenShift CLI on Windows	157
Installing the OpenShift CLI on macOS	157
7.9. LOGGING IN TO THE CLUSTER BY USING THE CLI	158
7.10. DISABLING THE DEFAULT OPERATORHUB CATALOG SOURCES	158
7.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	159
7.12. NEXT STEPS	159
<b>CHAPTER 8. INSTALLING A CLUSTER ON AWS INTO AN EXISTING VPC</b> .....	<b>160</b>
8.1. PREREQUISITES	160
8.2. ABOUT USING A CUSTOM VPC	160
8.2.1. Requirements for using your VPC	160
Option 1: Create VPC endpoints	162
Option 2: Create a proxy without VPC endpoints	162
Option 3: Create a proxy with VPC endpoints	162
8.2.2. VPC validation	164
8.2.3. Division of permissions	164
8.2.4. Isolation between clusters	164
8.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	165
8.4. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	165
8.5. OBTAINING THE INSTALLATION PROGRAM	167

8.6. CREATING THE INSTALLATION CONFIGURATION FILE	168
8.6.1. Installation configuration parameters	169
8.6.1.1. Required configuration parameters	169
8.6.1.2. Network configuration parameters	171
8.6.1.3. Optional configuration parameters	173
8.6.1.4. Optional AWS configuration parameters	179
8.6.2. Minimum resource requirements for cluster installation	183
8.6.3. Tested instance types for AWS	184
8.6.4. Tested instance types for AWS on 64-bit ARM infrastructures	184
8.6.5. Sample customized install-config.yaml file for AWS	185
8.6.6. Configuring the cluster-wide proxy during installation	188
8.7. DEPLOYING THE CLUSTER	190
8.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	191
Installing the OpenShift CLI on Linux	191
Installing the OpenShift CLI on Windows	192
Installing the OpenShift CLI on macOS	192
8.9. LOGGING IN TO THE CLUSTER BY USING THE CLI	193
8.10. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	194
8.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	194
8.12. NEXT STEPS	195
<b>CHAPTER 9. INSTALLING A PRIVATE CLUSTER ON AWS</b>	<b>196</b>
9.1. PREREQUISITES	196
9.2. PRIVATE CLUSTERS	196
9.2.1. Private clusters in AWS	197
9.2.1.1. Limitations	197
9.3. ABOUT USING A CUSTOM VPC	197
9.3.1. Requirements for using your VPC	198
Option 1: Create VPC endpoints	199
Option 2: Create a proxy without VPC endpoints	199
Option 3: Create a proxy with VPC endpoints	199
9.3.2. VPC validation	201
9.3.3. Division of permissions	201
9.3.4. Isolation between clusters	201
9.4. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	201
9.5. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	202
9.6. OBTAINING THE INSTALLATION PROGRAM	204
9.7. MANUALLY CREATING THE INSTALLATION CONFIGURATION FILE	205
9.7.1. Installation configuration parameters	205
9.7.1.1. Required configuration parameters	206
9.7.1.2. Network configuration parameters	207
9.7.1.3. Optional configuration parameters	209
9.7.1.4. Optional AWS configuration parameters	215
9.7.2. Minimum resource requirements for cluster installation	219
9.7.3. Tested instance types for AWS	220
9.7.4. Tested instance types for AWS on 64-bit ARM infrastructures	221
9.7.5. Sample customized install-config.yaml file for AWS	221
9.7.6. Configuring the cluster-wide proxy during installation	224
9.8. DEPLOYING THE CLUSTER	226
9.9. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	228
Installing the OpenShift CLI on Linux	228
Installing the OpenShift CLI on Windows	229
Installing the OpenShift CLI on macOS	229



---

9.10. LOGGING IN TO THE CLUSTER BY USING THE CLI	230
9.11. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	230
9.12. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	231
9.13. NEXT STEPS	231
<b>CHAPTER 10. INSTALLING A CLUSTER ON AWS INTO A GOVERNMENT REGION</b> .....	<b>232</b>
10.1. PREREQUISITES	232
10.2. AWS GOVERNMENT REGIONS	232
10.3. INSTALLATION REQUIREMENTS	232
10.4. PRIVATE CLUSTERS	233
10.4.1. Private clusters in AWS	233
10.4.1.1. Limitations	234
10.5. ABOUT USING A CUSTOM VPC	234
10.5.1. Requirements for using your VPC	234
Option 1: Create VPC endpoints	235
Option 2: Create a proxy without VPC endpoints	235
Option 3: Create a proxy with VPC endpoints	235
10.5.2. VPC validation	237
10.5.3. Division of permissions	237
10.5.4. Isolation between clusters	238
10.6. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	238
10.7. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	238
10.8. OBTAINING AN AWS MARKETPLACE IMAGE	240
10.9. OBTAINING THE INSTALLATION PROGRAM	241
10.10. MANUALLY CREATING THE INSTALLATION CONFIGURATION FILE	242
10.10.1. Installation configuration parameters	243
10.10.1.1. Required configuration parameters	243
10.10.1.2. Network configuration parameters	244
10.10.1.3. Optional configuration parameters	246
10.10.1.4. Optional AWS configuration parameters	252
10.10.2. Minimum resource requirements for cluster installation	256
10.10.3. Tested instance types for AWS	257
10.10.4. Tested instance types for AWS on 64-bit ARM infrastructures	258
10.10.5. Sample customized install-config.yaml file for AWS	258
10.10.6. Configuring the cluster-wide proxy during installation	261
10.11. DEPLOYING THE CLUSTER	263
10.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	265
Installing the OpenShift CLI on Linux	265
Installing the OpenShift CLI on Windows	266
Installing the OpenShift CLI on macOS	266
10.13. LOGGING IN TO THE CLUSTER BY USING THE CLI	267
10.14. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	267
10.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	268
10.16. NEXT STEPS	268
<b>CHAPTER 11. INSTALLING A CLUSTER ON AWS INTO A SECRET OR TOP SECRET REGION</b> .....	<b>269</b>
11.1. PREREQUISITES	269
11.2. AWS SECRET REGIONS	269
11.3. INSTALLATION REQUIREMENTS	270
11.4. PRIVATE CLUSTERS	270
11.4.1. Private clusters in AWS	271
11.4.1.1. Limitations	271
11.5. ABOUT USING A CUSTOM VPC	271

---

11.5.1. Requirements for using your VPC	272
Option 1: Create VPC endpoints	273
Option 2: Create a proxy without VPC endpoints	273
Option 3: Create a proxy with VPC endpoints	273
11.5.2. VPC validation	275
11.5.3. Division of permissions	275
11.5.4. Isolation between clusters	275
11.6. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	276
11.7. UPLOADING A CUSTOM RHCOS AMI IN AWS	276
11.8. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	278
11.9. OBTAINING THE INSTALLATION PROGRAM	280
11.10. MANUALLY CREATING THE INSTALLATION CONFIGURATION FILE	281
11.10.1. Installation configuration parameters	282
11.10.1.1. Required configuration parameters	282
11.10.1.2. Network configuration parameters	284
11.10.1.3. Optional configuration parameters	286
11.10.1.4. Optional AWS configuration parameters	292
11.10.2. Tested instance types for AWS	296
11.10.3. Sample customized install-config.yaml file for AWS	296
11.10.4. Configuring the cluster-wide proxy during installation	299
11.11. DEPLOYING THE CLUSTER	301
11.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	303
Installing the OpenShift CLI on Linux	303
Installing the OpenShift CLI on Windows	304
Installing the OpenShift CLI on macOS	304
11.13. LOGGING IN TO THE CLUSTER BY USING THE CLI	305
11.14. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	305
11.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	306
11.16. NEXT STEPS	306
<b>CHAPTER 12. INSTALLING A CLUSTER ON AWS CHINA</b> .....	<b>307</b>
12.1. PREREQUISITES	307
12.2. INSTALLATION REQUIREMENTS	307
12.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	307
12.4. PRIVATE CLUSTERS	308
12.4.1. Private clusters in AWS	309
12.4.1.1. Limitations	309
12.5. ABOUT USING A CUSTOM VPC	309
12.5.1. Requirements for using your VPC	309
Option 1: Create VPC endpoints	311
Option 2: Create a proxy without VPC endpoints	311
Option 3: Create a proxy with VPC endpoints	311
12.5.2. VPC validation	312
12.5.3. Division of permissions	313
12.5.4. Isolation between clusters	313
12.6. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	313
12.7. UPLOADING A CUSTOM RHCOS AMI IN AWS	315
12.8. OBTAINING THE INSTALLATION PROGRAM	317
12.9. MANUALLY CREATING THE INSTALLATION CONFIGURATION FILE	318
12.9.1. Installation configuration parameters	319
12.9.1.1. Required configuration parameters	319
12.9.1.2. Network configuration parameters	321
12.9.1.3. Optional configuration parameters	323

12.9.2. Sample customized install-config.yaml file for AWS	328
12.9.3. Minimum resource requirements for cluster installation	331
12.9.4. Tested instance types for AWS	332
12.9.5. Tested instance types for AWS on 64-bit ARM infrastructures	333
12.9.6. Configuring the cluster-wide proxy during installation	333
12.10. DEPLOYING THE CLUSTER	335
12.11. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	337
Installing the OpenShift CLI on Linux	337
Installing the OpenShift CLI on Windows	338
Installing the OpenShift CLI on macOS	338
12.12. LOGGING IN TO THE CLUSTER BY USING THE CLI	339
12.13. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	339
12.14. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	340
12.15. NEXT STEPS	340
<b>CHAPTER 13. INSTALLING A CLUSTER ON USER-PROVISIONED INFRASTRUCTURE IN AWS BY USING CLOUDFORMATION TEMPLATES</b>	<b>341</b>
13.1. PREREQUISITES	341
13.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	342
13.3. REQUIREMENTS FOR A CLUSTER WITH USER-PROVISIONED INFRASTRUCTURE	342
13.3.1. Required machines for cluster installation	342
13.3.2. Minimum resource requirements for cluster installation	343
13.3.3. Tested instance types for AWS	344
13.3.4. Tested instance types for AWS on 64-bit ARM infrastructures	344
13.3.5. Certificate signing requests management	345
13.4. REQUIRED AWS INFRASTRUCTURE COMPONENTS	345
13.4.1. Other infrastructure components	346
Option 1: Create VPC endpoints	346
Option 2: Create a proxy without VPC endpoints	346
Option 3: Create a proxy with VPC endpoints	346
13.4.2. Cluster machines	353
13.4.3. Required AWS permissions for the IAM user	354
13.5. OBTAINING AN AWS MARKETPLACE IMAGE	362
13.6. OBTAINING THE INSTALLATION PROGRAM	362
13.7. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	363
13.8. CREATING THE INSTALLATION FILES FOR AWS	365
13.8.1. Optional: Creating a separate /var partition	365
13.8.2. Creating the installation configuration file	367
13.8.3. Configuring the cluster-wide proxy during installation	369
13.8.4. Creating the Kubernetes manifest and Ignition config files	370
13.9. EXTRACTING THE INFRASTRUCTURE NAME	373
13.10. CREATING A VPC IN AWS	374
13.10.1. CloudFormation template for the VPC	375
13.11. CREATING NETWORKING AND LOAD BALANCING COMPONENTS IN AWS	381
13.11.1. CloudFormation template for the network and load balancers	385
13.12. CREATING SECURITY GROUP AND ROLES IN AWS	393
13.12.1. CloudFormation template for security objects	395
13.13. ACCESSING RHCOS AMIS WITH STREAM METADATA	406
13.14. RHCOS AMIS FOR THE AWS INFRASTRUCTURE	407
13.14.1. AWS regions without a published RHCOS AMI	409
13.14.2. Uploading a custom RHCOS AMI in AWS	410
13.15. CREATING THE BOOTSTRAP NODE IN AWS	412
13.15.1. CloudFormation template for the bootstrap machine	417

13.16. CREATING THE CONTROL PLANE MACHINES IN AWS	421
13.16.1. CloudFormation template for control plane machines	425
13.17. CREATING THE WORKER NODES IN AWS	431
13.17.1. CloudFormation template for worker machines	434
13.18. INITIALIZING THE BOOTSTRAP SEQUENCE ON AWS WITH USER-PROVISIONED INFRASTRUCTURE	437
13.19. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	438
Installing the OpenShift CLI on Linux	438
Installing the OpenShift CLI on Windows	439
Installing the OpenShift CLI on macOS	439
13.20. LOGGING IN TO THE CLUSTER BY USING THE CLI	440
13.21. APPROVING THE CERTIFICATE SIGNING REQUESTS FOR YOUR MACHINES	440
13.22. INITIAL OPERATOR CONFIGURATION	443
13.22.1. Image registry storage configuration	444
13.22.1.1. Configuring registry storage for AWS with user-provisioned infrastructure	444
13.22.1.2. Configuring storage for the image registry in non-production clusters	445
13.23. DELETING THE BOOTSTRAP RESOURCES	446
13.24. CREATING THE INGRESS DNS RECORDS	446
13.25. COMPLETING AN AWS INSTALLATION ON USER-PROVISIONED INFRASTRUCTURE	449
13.26. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	450
13.27. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	451
13.28. ADDITIONAL RESOURCES	451
13.29. NEXT STEPS	451
<b>CHAPTER 14. INSTALLING A CLUSTER USING AWS LOCAL ZONES</b>	<b>452</b>
14.1. PREREQUISITES	452
14.2. CLUSTER LIMITATIONS IN AWS LOCAL ZONES	453
14.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	453
14.4. OPTING INTO AWS LOCAL ZONES	454
14.5. OBTAINING AN AWS MARKETPLACE IMAGE	455
14.6. CREATING A VPC THAT USES AWS LOCAL ZONES	456
14.6.1. CloudFormation template for the VPC that uses AWS Local Zones	458
14.7. CREATING A SUBNET IN AWS LOCAL ZONES	464
14.7.1. CloudFormation template for the subnet that uses AWS Local Zones	467
14.8. OBTAINING THE INSTALLATION PROGRAM	468
14.9. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	469
14.10. CREATING THE INSTALLATION FILES FOR AWS	471
14.10.1. Minimum resource requirements for cluster installation	471
14.10.2. Tested instance types for AWS	471
14.10.3. Creating the installation configuration file	472
14.10.4. Creating the Kubernetes manifest files	474
14.11. DEPLOYING THE CLUSTER	478
14.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	479
Installing the OpenShift CLI on Linux	480
Installing the OpenShift CLI on Windows	480
Installing the OpenShift CLI on macOS	481
14.13. LOGGING IN TO THE CLUSTER BY USING THE CLI	481
14.14. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	482
14.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	483
14.16. NEXT STEPS	483
<b>CHAPTER 15. INSTALLING A CLUSTER ON AWS IN A RESTRICTED NETWORK WITH USER-PROVISIONED INFRASTRUCTURE</b>	<b>484</b>

15.1. PREREQUISITES	484
15.2. ABOUT INSTALLATIONS IN RESTRICTED NETWORKS	485
15.2.1. Additional limits	485
15.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	486
15.4. REQUIREMENTS FOR A CLUSTER WITH USER-PROVISIONED INFRASTRUCTURE	486
15.4.1. Required machines for cluster installation	486
15.4.2. Minimum resource requirements for cluster installation	487
15.4.3. Tested instance types for AWS	488
15.4.4. Tested instance types for AWS on 64-bit ARM infrastructures	488
15.4.5. Certificate signing requests management	489
15.5. REQUIRED AWS INFRASTRUCTURE COMPONENTS	489
15.5.1. Other infrastructure components	490
Option 1: Create VPC endpoints	490
Option 2: Create a proxy without VPC endpoints	490
Option 3: Create a proxy with VPC endpoints	490
15.5.2. Cluster machines	497
15.5.3. Required AWS permissions for the IAM user	498
15.6. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	506
15.7. CREATING THE INSTALLATION FILES FOR AWS	507
15.7.1. Optional: Creating a separate /var partition	508
15.7.2. Creating the installation configuration file	510
15.7.3. Configuring the cluster-wide proxy during installation	512
15.7.4. Creating the Kubernetes manifest and Ignition config files	514
15.8. EXTRACTING THE INFRASTRUCTURE NAME	517
15.9. CREATING A VPC IN AWS	517
15.9.1. CloudFormation template for the VPC	519
15.10. CREATING NETWORKING AND LOAD BALANCING COMPONENTS IN AWS	525
15.10.1. CloudFormation template for the network and load balancers	528
15.11. CREATING SECURITY GROUP AND ROLES IN AWS	536
15.11.1. CloudFormation template for security objects	539
15.12. ACCESSING RHCOS AMIS WITH STREAM METADATA	550
15.13. RHCOS AMIS FOR THE AWS INFRASTRUCTURE	551
15.14. CREATING THE BOOTSTRAP NODE IN AWS	553
15.14.1. CloudFormation template for the bootstrap machine	558
15.15. CREATING THE CONTROL PLANE MACHINES IN AWS	562
15.15.1. CloudFormation template for control plane machines	567
15.16. CREATING THE WORKER NODES IN AWS	572
15.16.1. CloudFormation template for worker machines	576
15.17. INITIALIZING THE BOOTSTRAP SEQUENCE ON AWS WITH USER-PROVISIONED INFRASTRUCTURE	578
15.18. LOGGING IN TO THE CLUSTER BY USING THE CLI	579
15.19. APPROVING THE CERTIFICATE SIGNING REQUESTS FOR YOUR MACHINES	580
15.20. INITIAL OPERATOR CONFIGURATION	582
15.20.1. Disabling the default OperatorHub catalog sources	583
15.20.2. Image registry storage configuration	584
15.20.2.1. Configuring registry storage for AWS with user-provisioned infrastructure	584
15.20.2.2. Configuring storage for the image registry in non-production clusters	585
15.21. DELETING THE BOOTSTRAP RESOURCES	585
15.22. CREATING THE INGRESS DNS RECORDS	586
15.23. COMPLETING AN AWS INSTALLATION ON USER-PROVISIONED INFRASTRUCTURE	588
15.24. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	589
15.25. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	590
15.26. ADDITIONAL RESOURCES	591

15.27. NEXT STEPS	591
<b>CHAPTER 16. INSTALLING A CLUSTER ON AWS WITH REMOTE WORKERS ON AWS OUTPOSTS</b> . . . . .	<b>592</b>
16.1. PREREQUISITES	592
16.2. ABOUT USING A CUSTOM VPC	593
16.2.1. Requirements for using your VPC	593
Option 1: Create VPC endpoints	595
Option 2: Create a proxy without VPC endpoints	595
Option 3: Create a proxy with VPC endpoints	595
16.2.2. VPC validation	597
16.2.3. Division of permissions	597
16.2.4. Isolation between clusters	598
16.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	598
16.4. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	598
16.5. OBTAINING THE INSTALLATION PROGRAM	600
16.6. MINIMUM RESOURCE REQUIREMENTS FOR CLUSTER INSTALLATION	601
16.7. IDENTIFYING YOUR AWS OUTPOSTS INSTANCE TYPES	602
16.8. CREATING THE INSTALLATION CONFIGURATION FILE	602
16.8.1. Installation configuration parameters	604
16.8.1.1. Required configuration parameters	604
16.8.1.2. Network configuration parameters	606
16.8.1.3. Optional configuration parameters	608
16.8.1.4. Optional AWS configuration parameters	614
16.8.2. Sample customized install-config.yaml file for AWS	618
16.9. GENERATING MANIFEST FILES	620
16.9.1. Modifying manifest files	621
16.10. DEPLOYING THE CLUSTER	622
16.11. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	624
Installing the OpenShift CLI on Linux	624
Installing the OpenShift CLI on Windows	624
Installing the OpenShift CLI on macOS	625
16.12. LOGGING IN TO THE CLUSTER BY USING THE CLI	625
16.13. LOGGING IN TO THE CLUSTER BY USING THE WEB CONSOLE	626
16.14. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	627
16.15. CLUSTER LIMITATIONS	627
16.16. NEXT STEPS	628
<b>CHAPTER 17. UNINSTALLING A CLUSTER ON AWS</b> . . . . .	<b>629</b>
17.1. REMOVING A CLUSTER THAT USES INSTALLER-PROVISIONED INFRASTRUCTURE	629
17.2. DELETING AWS RESOURCES WITH THE CLOUD CREDENTIAL OPERATOR UTILITY	629
17.3. DELETING A CLUSTER WITH A CONFIGURED AWS LOCAL ZONE INFRASTRUCTURE	631



# CHAPTER 1. PREPARING TO INSTALL ON AWS

## 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](#).

## 1.2. REQUIREMENTS FOR INSTALLING OPENSIFT CONTAINER PLATFORM ON AWS

Before installing OpenShift Container Platform on Amazon Web Services (AWS), you must create an AWS account. See [Configuring an AWS account](#) for details about configuring an account, account limits, account permissions, IAM user setup, and supported AWS 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 IAM for AWS](#) for other options, including [configuring the Cloud Credential Operator \(CCO\) to use the Amazon Web Services Security Token Service \(AWS STS\)](#).

## 1.3. CHOOSING A METHOD TO INSTALL OPENSIFT CONTAINER PLATFORM ON AWS

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.

### 1.3.1. 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.

### 1.3.2. 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.

## 1.4. NEXT STEPS

- [Configuring an AWS account](#)

## CHAPTER 2. CONFIGURING AN AWS ACCOUNT

Before you can install OpenShift Container Platform, you must configure an Amazon Web Services (AWS) account.

### 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.

#### 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.

## 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

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"> <li>• One bootstrap machine, which is removed after installation</li> <li>• Three control plane nodes</li> <li>• Three worker nodes</li> </ul> <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 <b>m6i.large</b> instance and the bootstrap and control plane machines use <b>m6i.xlarge</b> instances. In some regions, including all regions that do not support these instance types, <b>m5.large</b> and <b>m5.xlarge</b> instances are used instead.</p>
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 <a href="#">availability zone within a region</a>. Each private subnet requires a <a href="#">NAT Gateway</a>, and each NAT gateway requires a separate <a href="#">elastic IP</a>. Review the <a href="#">AWS region map</a> 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><b>IMPORTANT</b></p> <p>To use the <b>us-east-1</b> 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.

Component	Number of clusters available by default	Default AWS limit	Description
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 <b>Service</b> objects with type <b>LoadBalancer</b> will create additional <a href="#">load balancers</a> .
NAT Gateways	5	5 per availability zone	The cluster deploys one NAT gateway in each availability zone.
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 <b>us-east-1</b> region contains six availability zones, so a cluster that is deployed in that zone uses 27 ENIs. Review the <a href="#">AWS region map</a> 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.

## 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 2.1. Required EC2 permissions for installation


- **ec2:AuthorizeSecurityGroupEgress**
- **ec2:AuthorizeSecurityGroupIngress**
- **ec2:CopyImage**
- **ec2:CreateNetworkInterface**
- **ec2:AttachNetworkInterface**
- **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:DescribeRegions**
- **ec2:DescribeRouteTables**
- **ec2:DescribeSecurityGroups**
- **ec2:DescribeSubnets**
- **ec2:DescribeTags**
- **ec2:DescribeVolumes**
- **ec2:DescribeVpcAttribute**
- **ec2:DescribeVpcClassicLink**
- **ec2:DescribeVpcClassicLinkDnsSupport**
- **ec2:DescribeVpcEndpoints**
- **ec2:DescribeVpcs**
- **ec2:GetEbsDefaultKmsKeyId**
- **ec2:ModifyInstanceAttribute**
- **ec2:ModifyNetworkInterfaceAttribute**
- **ec2:RevokeSecurityGroupEgress**
- **ec2:RevokeSecurityGroupIngress**
- **ec2:RunInstances**
- **ec2:TerminateInstances**

Example 2.2. 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 VPC, your account does not require these permissions for creating network resources.

#### Example 2.3. Required Elastic Load Balancing permissions (ELB) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:ApplySecurityGroupsToLoadBalancer**
- **elasticloadbalancing:AttachLoadBalancerToSubnets**
- **elasticloadbalancing:ConfigureHealthCheck**
- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateLoadBalancerListeners**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing:DeregisterInstancesFromLoadBalancer**
- **elasticloadbalancing:DescribeInstanceHealth**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTags**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:RegisterInstancesWithLoadBalancer**
- **elasticloadbalancing:SetLoadBalancerPoliciesOfListener**

#### Example 2.4. Required Elastic Load Balancing permissions (ELBv2) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:CreateListener**




- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateTargetGroup**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing:DeregisterTargets**
- **elasticloadbalancing:DescribeListeners**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTargetGroupAttributes**
- **elasticloadbalancing:DescribeTargetHealth**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:ModifyTargetGroup**
- **elasticloadbalancing:ModifyTargetGroupAttributes**
- **elasticloadbalancing:RegisterTargets**

#### Example 2.5. 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:TagRole**

**NOTE**

If you have not created a load balancer in your AWS account, the IAM user also requires the **iam:CreateServiceLinkedRole** permission.

**Example 2.6. 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 2.7. Required S3 permissions for installation**

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3:GetAccelerateConfiguration**
- **s3:GetBucketAcl**
- **s3:GetBucketCors**
- **s3:GetBucketLocation**
- **s3:GetBucketLogging**
- **s3:GetBucketPolicy**
- **s3:GetBucketObjectLockConfiguration**

- **s3:GetBucketReplication**
- **s3:GetBucketRequestPayment**
- **s3:GetBucketTagging**
- **s3:GetBucketVersioning**
- **s3:GetBucketWebsite**
- **s3:GetEncryptionConfiguration**
- **s3:GetLifecycleConfiguration**
- **s3:GetReplicationConfiguration**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketTagging**
- **s3:PutEncryptionConfiguration**

Example 2.8. S3 permissions that cluster Operators require

- **s3>DeleteObject**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**
- **s3:GetObjectVersion**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Example 2.9. Required permissions to delete base cluster resources

- **autoscaling:DescribeAutoScalingGroups**
- **ec2>DeletePlacementGroup**
- **ec2>DeleteNetworkInterface**
- **ec2>DeleteVolume**
- **elasticloadbalancing>DeleteTargetGroup**
- **elasticloadbalancing:DescribeTargetGroups**

- **iam:DeleteAccessKey**
- **iam:DeleteUser**
- **iam:ListAttachedRolePolicies**
- **iam:ListInstanceProfiles**
- **iam:ListRolePolicies**
- **iam:ListUserPolicies**
- **s3:DeleteObject**
- **s3:ListBucketVersions**
- **tag:GetResources**

Example 2.10. 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 2.11. Required permissions to delete a cluster with shared instance roles

- **iam:UntagRole**

### Example 2.12. Additional IAM and S3 permissions that are required to create manifests

- **iam:DeleteAccessKey**
- **iam:DeleteUser**
- **iam:DeleteUserPolicy**
- **iam:GetUserPolicy**
- **iam:ListAccessKeys**
- **iam:PutUserPolicy**
- **iam:TagUser**
- **s3:PutBucketPublicAccessBlock**
- **s3:GetBucketPublicAccessBlock**
- **s3:PutLifecycleConfiguration**
- **s3:ListBucket**
- **s3:ListBucketMultipartUploads**
- **s3:AbortMultipartUpload**



#### 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 2.13. Optional permissions for instance and quota checks for installation

- **ec2:DescribeInstanceTypeOfferings**
- **servicequotas:ListAWSDefaultServiceQuotas**

## 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-lived credentials.

#### Additional resources

- See [Manually creating IAM for AWS](#) for steps to set the Cloud Credential Operator (CCO) to manual mode prior to installation. Use this mode in environments where the cloud identity and access management (IAM) APIs are not reachable, or if you prefer not to store an administrator-level credential secret in the cluster **kube-system** project.

## 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".

### 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 2.14. 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 2.15. Default IAM role permissions for compute instance profiles

- **ec2:DescribeInstances**
- **ec2:DescribeRegions**

## 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 [RHOC P 4 AWS cloud-credentials access key is expired](#) (Red Hat Knowledgebase).

#### Additional resources

- See [Deploying the cluster](#).

### 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.

## 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.

## 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.

### 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)
- **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)
- **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)

### 2.7.2. AWS GovCloud regions

The following AWS GovCloud regions are supported:

- **us-gov-west-1**
- **us-gov-east-1**

### 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)

#### 2.7.4. AWS China regions

The following AWS China regions are supported:

- **cn-north-1** (Beijing)
- **cn-northwest-1** (Ningxia)

### 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](#)
  - [Installing a cluster on AWS with remote workers on AWS Outposts](#)

## CHAPTER 3. MANUALLY CREATING IAM FOR AWS

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, you can put the Cloud Credential Operator (CCO) into manual mode before you install the cluster.


### 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.

If you prefer not to store an administrator-level credential secret in the cluster **kube-system** project, you can choose one of the following options when installing OpenShift Container Platform:

- **Use the Amazon Web Services Security Token Service**

You can use the CCO utility (**ccoctl**) to configure the cluster to use the Amazon Web Services Security Token Service (AWS STS). When the CCO utility is used to configure the cluster for STS, it assigns IAM roles that provide short-term, limited-privilege security credentials to components.



#### NOTE

This credentials strategy is supported for only new OpenShift Container Platform clusters and must be configured during installation. You cannot reconfigure an existing cluster that uses a different credentials strategy to use this feature.

- **Manage cloud credentials manually.**

You can set the **credentialsMode** parameter for the CCO to **Manual** to manage 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.

- **Remove the administrator-level credential secret after installing OpenShift Container Platform with mint mode:**

If you are using the CCO with the **credentialsMode** parameter set to **Mint**, you can remove or rotate the administrator-level credential after installing OpenShift Container Platform. Mint mode is the default configuration for the CCO. This option requires the presence of the administrator-level credential during an installation. The administrator-level credential is used during the installation to mint other credentials with some permissions granted. The original credential secret is not stored in the cluster permanently.



#### NOTE

Prior to a non z-stream upgrade, you must reinstate the credential secret with the administrator-level credential. If the credential is not present, the upgrade might be blocked.

## Additional resources

- To learn how to use the CCO utility (**ccoctl**) to configure the CCO to use the AWS STS, see [Using manual mode with STS](#).
- To learn how to rotate or remove the administrator-level credential secret after installing OpenShift Container Platform, see [Rotating or removing cloud provider credentials](#).
- For a detailed description of all available CCO credential modes and their supported platforms, see [About the Cloud Credential Operator](#).

## 3.2. MANUALLY CREATE IAM

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. Change to the directory that contains the installation program and create the **install-config.yaml** file by running the following command:

```
$ openshift-install create install-config --dir <installation_directory>
```

where **<installation\_directory>** is the directory in which the installation program creates files.

2. 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**.

3. Generate the manifests by running the following command from the directory that contains the installation program:

```
$ openshift-install create manifests --dir <installation_directory>
```

where **<installation\_directory>** is the directory in which the installation program creates files.

4. 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 by running the following command:

```
$ openshift-install version
```

### Example output

```
release image quay.io/openshift-release-dev/ocp-release:4.y.z-x86_64
```

5. Locate all **CredentialsRequest** objects in this release image that target the cloud you are deploying on by running the following command:

```
$ oc adm release extract quay.io/openshift-release-dev/ocp-release:4.y.z-x86_64 \
  --credentials-requests \
  --cloud=aws
```

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: "*"
    ...
```

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
    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

The release image includes **CredentialsRequest** objects for Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set. You can identify these objects by their use of the **release.openshift.io/feature-set: TechPreviewNoUpgrade** annotation.

- If you are not using any of these features, do not create secrets for these objects. Creating secrets for Technology Preview features that you are not using can cause the installation to fail.
- If you are using any of these features, you must create secrets for the corresponding objects.

- To find **CredentialsRequest** objects with the **TechPreviewNoUpgrade** annotation, run the following command:

```
$ grep "release.openshift.io/feature-set" *
```

#### Example output

```
0000_30_capi-operator_00_credentials-request.yaml: release.openshift.io/feature-set:
TechPreviewNoUpgrade
```

7. From the directory that contains the installation program, proceed with your cluster creation:

```
$ openshift-install create cluster --dir <installation_directory>
```



#### 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](#)

## 3.3. MINT MODE

Mint mode is the default Cloud Credential Operator (CCO) credentials mode for OpenShift Container Platform on platforms that support it. In this mode, the CCO uses the provided administrator-level cloud credential to run the cluster. Mint mode is supported for AWS and GCP.

In mint mode, the **admin** credential is stored in the **kube-system** namespace and then used by the CCO to process the **CredentialsRequest** objects in the cluster and create users for each with specific permissions.

The benefits of mint mode include:

- Each cluster component has only the permissions it requires
- Automatic, on-going reconciliation for cloud credentials, including additional credentials or permissions that might be required for upgrades

One drawback is that mint mode requires **admin** credential storage in a cluster **kube-system** secret.

## 3.4. MINT MODE WITH REMOVAL OR ROTATION OF THE ADMINISTRATOR-LEVEL CREDENTIAL

Currently, this mode is only supported on AWS and GCP.

In this mode, a user installs OpenShift Container Platform with an administrator-level credential just like the normal mint mode. However, this process removes the administrator-level credential secret from the cluster post-installation.

The administrator can have the Cloud Credential Operator make its own request for a read-only credential that allows it to verify if all **CredentialsRequest** objects have their required permissions, thus the administrator-level credential is not required unless something needs to be changed. After the associated credential is removed, it can be deleted or deactivated on the underlying cloud, if desired.



### NOTE

Prior to a non z-stream upgrade, you must reinstate the credential secret with the administrator-level credential. If the credential is not present, the upgrade might be blocked.

The administrator-level credential is not stored in the cluster permanently.

Following these steps still requires the administrator-level credential in the cluster for brief periods of time. It also requires manually re-instating the secret with administrator-level credentials for each upgrade.

## 3.5. NEXT STEPS

- Install an OpenShift Container Platform cluster:


- [Installing a cluster quickly on AWS](#) 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](#)

## CHAPTER 4. INSTALLING A CLUSTER QUICKLY ON AWS

In OpenShift Container Platform version 4.12, you can install a cluster on Amazon Web Services (AWS) that uses the default configuration options.

### 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-lived 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.
- 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 IAM credentials](#). Manual mode can also be used in environments where the cloud IAM APIs are not reachable.

### 4.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

### 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 FIPS validated or Modules In Process cryptographic libraries on 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.

## 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 the 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.

## 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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 the Red Hat OpenShift Cluster Manager](#) .

**NOTE**

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 4.6. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 4.7. 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
```

## 4.8. 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.

## 4.9. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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

## 4.10. 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](#).

## CHAPTER 5. INSTALLING A CLUSTER ON AWS WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.12, 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.

### 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 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-lived 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.
- 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 IAM credentials](#).

### 5.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

**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 FIPS validated or Modules In Progress cryptographic libraries on 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.

## 5.4. 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 worker 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 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 worker 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.

## 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 the 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.

## 5.6. CREATING THE INSTALLATION CONFIGURATION FILE

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Obtain service principal permissions at the subscription level.

### 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.
  - vii. Paste the [pull secret from the Red Hat OpenShift Cluster Manager](#) .
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.

### 5.6.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for

the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.



## NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

### 5.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 5.1. Required parameters

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;</b> . <b>&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}</b> . <b>{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .

Parameter	Description	Values
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud</b> , <b>aws</b> , <b>baremetal</b> , <b>azure</b> , <b>gcp</b> , <b>ibmcloud</b> , <b>nutanix</b> , <b>openstack</b> , <b>ovirt</b> , <b>vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 5.6.1.2. 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 5.2. Network parameters

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object  <b>NOTE</b> You cannot modify parameters specified by the <b>networking</b> object after installation.
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:   - cidr: 10.128.0.0/14     hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:   - 172.30.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:   - cidr: 10.0.0.0/16</pre>
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation.  For example, <b>10.0.0.0/16</b> .   <p><b>NOTE</b></p> <p>Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.</p>

### 5.6.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:


Table 5.3. Optional parameters



Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String

Parameter	Description	Values
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>




Parameter	Description	Values
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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

Parameter	Description	Values
<b>controlPlane.hyperthreading</b>	<p>Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b>, 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><b>IMPORTANT</b></p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.

Parameter	Description	Values
<b>credentialsMode</b>	<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> <div data-bbox="485 510 595 891" style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">  </div> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> <div data-bbox="485 936 595 1288" style="border: 1px solid black; padding: 5px;">  </div> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<p><b>fips</b></p>	<p>Enable or disable FIPS mode. The default is <b>false</b> (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 1391" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p><b>IMPORTANT</b></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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> </div> <div data-bbox="486 1440 592 1637" style="background-color: #f0f0f0; padding: 5px; margin-bottom: 10px;"> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p><b>false</b> or <b>true</b></p>


Parameter	Description	Values
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

#### 5.6.1.4. Optional AWS configuration parameters

Optional AWS configuration parameters are described in the following table:


Table 5.4. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID or the key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .

Parameter	Description	Values
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	<p>Any valid <a href="#">AWS region</a>, such as <b>us-east-1</b>. 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; margin-top: 10px;">  <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability map</a> in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.



Parameter	Description	Values
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.   <b>NOTE</b> You can add up to 25 user defined tags during installation. The remaining 25 tags are reserved for OpenShift Container Platform.
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

## 5.6.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 5.5. 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

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.


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](#)

### 5.6.3. 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 "Minimum resource requirements for cluster installation".

#### Example 5.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***

- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

#### 5.6.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 5.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

#### 5.6.5. 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
  amiID: ami-96c6f8f7 16
  serviceEndpoints: 17
    - name: ec2
      url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
  fips: false 18
  sshKey: ssh-ed25519 AAAA... 19
  pullSecret: '{"auths": ...}' 20

```

- 1 12 14 20** 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.
- 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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.

- 16 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.
- 17 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.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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.

## 5.6.6. 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

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.

## 5.7. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



#### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

**5.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY**

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 5.9. 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.

- Verify you can run **oc** commands successfully using the exported configuration:

```
$ oc whoami
```

#### Example output

```
system:admin
```

## 5.10. 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

- 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.

## 5.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 5.12. 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](#) .


## CHAPTER 6. INSTALLING A CLUSTER ON AWS WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform version 4.12, 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.

### 6.1. 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-lived 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.
- 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 IAM credentials](#).

### 6.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 6.3. 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 FIPS validated or Modules In Progress cryptographic libraries on 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.

## 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 the 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.

## 6.5. 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.

**6.6. CREATING THE INSTALLATION CONFIGURATION FILE**

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

**Prerequisites**

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Obtain service principal permissions at the subscription level.

**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.
  - vii. Paste the [pull secret from the Red Hat OpenShift Cluster Manager](#) .
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.

### 6.6.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.



#### NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

#### 6.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 6.1. Required parameters

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;</b> . <b>&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}</b> . <b>{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud</b> , <b>aws</b> , <b>baremetal</b> , <b>azure</b> , <b>gcp</b> , <b>ibmcloud</b> , <b>nutanix</b> , <b>openstack</b> , <b>ovirt</b> , <b>vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 6.6.1.2. 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 6.2. Network parameters

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;"> <div> <p><b>NOTE</b></p> <p>You cannot modify parameters specified by the <b>networking</b> object after installation.</p> </div> </div>

Parameter	Description	Values
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:     - cidr: 10.0.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation. For example, <b>10.0.0.0/16</b> .  <b>NOTE</b> Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.

### 6.6.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 6.3. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.


Parameter	Description	Values
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .



Parameter	Description	Values
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>controlPlane.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>

Parameter	Description	Values
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.
<b>credentialsMode</b>	<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> <div style="display: flex; align-items: flex-start;"> <div style="width: 20px; height: 100px; border: 1px solid black; margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="width: 20px; height: 100px; border: 1px solid black; margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ( <b>""</b> ).
<b>fips</b>	Enable or disable FIPS mode. The default is <b>false</b> (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.	<b>false</b> or <b>true</b>

Parameter	Description	Values
	<p data-bbox="671 107 863 138"><b>IMPORTANT</b></p> <p data-bbox="671 181 932 909">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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> <p data-bbox="671 965 762 996"><b>NOTE</b></p> <p data-bbox="671 1032 922 1151">If you are using Azure File storage, you cannot enable FIPS mode.</p>	
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.


Parameter	Description	Values
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

### 6.6.1.4. Optional AWS configuration parameters


Optional AWS configuration parameters are described in the following table:

Table 6.4. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> or the <a href="#">key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	Any valid <a href="#">AWS region</a> , such as <b>us-east-1</b> . You can use the AWS CLI to access the regions available based on your selected instance type. For example: <pre>aws ec2 describe-instance-type-offerings -- filters Name=instance- type,Values=c7g.xlarge</pre> <div style="display: flex; align-items: flex-start; margin-top: 10px;">  <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability map</a> in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.

Parameter	Description	Values
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .

Parameter	Description	Values
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	<p>Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></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>
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .



Parameter	Description	Values
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

## 6.6.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

**Table 6.5. 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:  $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$ .
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.

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](#)

### 6.6.3. 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 "Minimum resource requirements for cluster installation".

#### Example 6.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***
- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

### 6.6.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 6.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

## 6.6.5. 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
      amiID: ami-96c6f8f7 17
      serviceEndpoints: 18
        - name: ec2
          url: https://vpce-id.ec2.us-west-2.vpce.amazonaws.com
    fips: false 19
    sshKey: ssh-ed25519 AAAA... 20
    pullSecret: '{"auths": ...}' 21

```

**1 12 15 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.

**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



### 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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the **x86\_64**, **ppc64le**, and **s390x** architectures.

- 20 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.

## 6.6.6. 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

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.

## 6.7. 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 and these fields cannot be changed:

#### **clusterNetwork**

IP address pools from which pod IP addresses are allocated.

#### **serviceNetwork**

IP address pool for services.

#### **defaultNetwork.type**

Cluster network plugin, such as OpenShift SDN or OVN-Kubernetes.

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**.

### 6.7.1. Cluster Network Operator configuration object

The fields for the Cluster Network Operator (CNO) are described in the following table:

**Table 6.6. Cluster Network Operator configuration object**


Field	Type	Description
<b>metadata.name</b>	<b>string</b>	The name of the CNO object. This name is always <b>cluster</b> .
<b>spec.clusterNetwork</b>	<b>array</b>	<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> <p>You can customize this field only in the <b>install-config.yaml</b> file before you create the manifests. The value is read-only in the manifest file.</p>
<b>spec.serviceNetwork</b>	<b>array</b>	<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 <b>install-config.yaml</b> file before you create the manifests. The value is read-only in the manifest file.</p>
<b>spec.defaultNetwork</b>	<b>object</b>	Configures the network plugin for the cluster network.
<b>spec.kubeProxyConfig</b>	<b>object</b>	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 6.7. defaultNetwork object**



Field	Type	Description
<b>type</b>	<b>string</b>	<p>Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b>. 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><b>NOTE</b></p> <p>OpenShift Container Platform uses the OVN-Kubernetes network plugin by default.</p> </div>
<b>openshiftSDNConfig</b>	<b>object</b>	This object is only valid for the OpenShift SDN network plugin.
<b>ovnKubernetesConfig</b>	<b>object</b>	This object is only valid for the OVN-Kubernetes network plugin.

### Configuration for the OpenShift SDN network plugin

The following table describes the configuration fields for the OpenShift SDN network plugin:

**Table 6.8. openshiftSDNConfig object**

Field	Type	Description
<b>mode</b>	<b>string</b>	<p>Configures the network isolation mode for OpenShift SDN. The default value is <b>NetworkPolicy</b>.</p> <p>The values <b>Multitenant</b> and <b>Subnet</b> are available for backwards compatibility with OpenShift Container Platform 3.x but are not recommended. This value cannot be changed after cluster installation.</p>
<b>mtu</b>	<b>integer</b>	<p>The maximum transmission unit (MTU) for the VXLAN 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 <b>50</b> less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of <b>9001</b>, and some have an MTU of <b>1500</b>, you must set this value to <b>1450</b>.</p> <p>This value cannot be changed after cluster installation.</p>

Field	Type	Description
<b>vxlanPort</b>	<b>integer</b>	<p>The port to use for all VXLAN packets. The default value is <b>4789</b>. This value cannot be changed after cluster installation.</p> <p>If you are running in a virtualized environment with existing nodes that are part of another VXLAN network, then you might be required to change this. For example, when running an OpenShift SDN overlay on top of VMware NSX-T, you must select an alternate port for the VXLAN, because both SDNs use the same default VXLAN port number.</p> <p>On Amazon Web Services (AWS), you can select an alternate port for the VXLAN between port <b>9000</b> and port <b>9999</b>.</p>

### Example OpenShift SDN configuration


```
defaultNetwork:
  type: OpenShiftSDN
  openshiftSDNConfig:
    mode: NetworkPolicy
    mtu: 1450
    vxlanPort: 4789
```

### Configuration for the OVN-Kubernetes network plugin

The following table describes the configuration fields for the OVN-Kubernetes network plugin:

**Table 6.9. ovnKubernetesConfig object**

Field	Type	Description
<b>mtu</b>	<b>integer</b>	<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 <b>100</b> less than the lowest MTU value in your cluster. For example, if some nodes in your cluster have an MTU of <b>9001</b>, and some have an MTU of <b>1500</b>, you must set this value to <b>1400</b>.</p>
<b>genevePort</b>	<b>integer</b>	The port to use for all Geneve packets. The default value is <b>6081</b> . This value cannot be changed after cluster installation.
<b>ipsecConfig</b>	<b>object</b>	Specify an empty object to enable IPsec encryption.

Field	Type	Description
<b>policyAuditConfig</b>	<b>object</b>	Specify a configuration object for customizing network policy audit logging. If unset, the defaults audit log settings are used.
<b>gatewayConfig</b>	<b>object</b>	<p>Optional: Specify a configuration object for customizing how egress traffic is sent to the node gateway.</p> <p> <b>NOTE</b></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>

Field	Type	Description
<b>v4InternalSubnet</b>	<p>If your existing network infrastructure overlaps with the <b>100.64.0.0/16</b> 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 <b>clusterNetwork.cidr</b> value is <b>10.128.0.0/14</b> and the <b>clusterNetwork.hostPrefix</b> value is <b>/23</b>, then the maximum number of nodes is <b><math>2^{(23-14)}=512</math></b>.</p> <p>This field cannot be changed after installation.</p>	The default value is <b>100.64.0.0/16</b> .

Field	Type	Description
<b>v6InternalSubnet</b>	<p>If your existing network infrastructure overlaps with the <b>fd98::/48</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>This field cannot be changed after installation.</p>	The default value is <b>fd98::/48</b> .

Table 6.10. `policyAuditConfig` object

Field	Type	Description
<b>rateLimit</b>	integer	The maximum number of messages to generate every second per node. The default value is <b>20</b> messages per second.
<b>maxFileSize</b>	integer	The maximum size for the audit log in bytes. The default value is <b>50000000</b> or 50 MB.

Field	Type	Description
<b>destination</b>	string	<p>One of the following additional audit log targets:</p> <p><b>libc</b> The libc <b>syslog()</b> function of the journald process on the host.</p> <p><b>udp:&lt;host&gt;:&lt;port&gt;</b> A syslog server. Replace <b>&lt;host&gt;:&lt;port&gt;</b> with the host and port of the syslog server.</p> <p><b>unix:&lt;file&gt;</b> A Unix Domain Socket file specified by <b>&lt;file&gt;</b>.</p> <p><b>null</b> Do not send the audit logs to any additional target.</p>
<b>syslogFacility</b>	string	The syslog facility, such as <b>kern</b> , as defined by RFC5424. The default value is <b>local0</b> .

Table 6.11. gatewayConfig object

Field	Type	Description
<b>routingViaHost</b>	<b>boolean</b>	<p>Set this field to <b>true</b> 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 <b>false</b>.</p> <p>This field has an interaction with the Open vSwitch hardware offloading feature. If you set this field to <b>true</b>, you do not receive the performance benefits of the offloading because egress traffic is processed by the host networking stack.</p>

### Example OVN-Kubernetes configuration with IPsec enabled

```


defaultNetwork:
  type: OVNKubernetes
  ovnKubernetesConfig:
    mtu: 1400
    genevePort: 6081
    ipsecConfig: {}

```

### kubeProxyConfig object configuration

The values for the **kubeProxyConfig** object are defined in the following table:

Table 6.12. kubeProxyConfig object

Field	Type	Description
<b>iptablesSyncPeriod</b>	<b>string</b>	<p>The refresh period for <b>iptables</b> rules. The default value is <b>30s</b>. Valid suffixes include <b>s</b>, <b>m</b>, and <b>h</b> and are described in the <a href="#">Go time package</a> documentation.</p> <div style="display: flex; align-items: flex-start;">  <div> <p><b>NOTE</b></p> <p>Because of performance improvements introduced in OpenShift Container Platform 4.3 and greater, adjusting the <b>iptablesSyncPeriod</b> parameter is no longer necessary.</p> </div> </div>
<b>proxyArguments.iptables-min-sync-period</b>	<b>array</b>	<p>The minimum duration before refreshing <b>iptables</b> rules. This field ensures that the refresh does not happen too frequently. Valid suffixes include <b>s</b>, <b>m</b>, and <b>h</b> and are described in the <a href="#">Go time package</a>. The default value is:</p> <pre>kubeProxyConfig:   proxyArguments:     iptables-min-sync-period:       - 0s</pre>

## 6.8. 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.

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 examples:

#### Specify a different VXLAN port for the OpenShift SDN network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    openshiftSDNConfig:
      vxlanPort: 4800
```

#### Enable IPsec for the OVN-Kubernetes network provider

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      ipsecConfig: {}
```

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.



#### 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](#).

## 6.9. 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.

1. Change to the directory that contains the installation program and create the manifests:

```
$ ./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:

```
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
```

4. Save the **cluster-ingress-default-ingresscontroller.yaml** file and quit the text editor.
5. Optional: Back up the **manifests/cluster-ingress-default-ingresscontroller.yaml** file. The installation program deletes the **manifests/** directory when creating the cluster.

## 6.10. CONFIGURING HYBRID NETWORKING WITH OVN-KUBERNETES

You can configure your cluster to use hybrid networking with OVN-Kubernetes. This allows a hybrid cluster that supports different node networking configurations. For example, this is necessary to run both Linux and Windows nodes in a cluster.



## IMPORTANT

You must configure hybrid networking with OVN-Kubernetes during the installation of your cluster. You cannot switch to hybrid networking after the installation process.

### 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.yaml** in the **<installation\_directory>/manifests/** directory:

```
$ cat <<EOF > <installation_directory>/manifests/cluster-network-03-config.yaml
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.yaml** 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
    hybridOverlayVXLANPort: 9898 2

```

- 1 Specify the CIDR configuration used for nodes on the additional overlay network. The **hybridClusterNetwork** CIDR cannot overlap with the **clusterNetwork** CIDR.
- 2 Specify a custom VXLAN port for the additional overlay network. This is required for running Windows nodes in a cluster installed on vSphere, and must not be configured for any other cloud provider. The custom port can be any open port excluding the default **4789** port. For more information on this requirement, see the Microsoft documentation on [Pod-to-pod connectivity between hosts is broken](#).



#### NOTE

Windows Server Long-Term Servicing Channel (LTSC): Windows Server 2019 is not supported on clusters with a custom **hybridOverlayVXLANPort** value because this Windows server version does not support selecting a custom VXLAN port.

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.



#### NOTE

For more information on using Linux and Windows nodes in the same cluster, see [Understanding Windows container workloads](#).

## 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 6.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 6.13. 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
```

## 6.14. 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**

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

## 6.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 6.16. 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](#) .



## CHAPTER 7. INSTALLING A CLUSTER ON AWS IN A RESTRICTED NETWORK

In OpenShift Container Platform version 4.12, 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).

### 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 [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-lived 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.

- 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 IAM credentials](#).

## 7.2. ABOUT INSTALLATIONS IN RESTRICTED NETWORKS

In OpenShift Container Platform 4.12, 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.

### 7.2.1. 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.

## 7.3. ABOUT USING A CUSTOM VPC

In OpenShift Container Platform 4.12, 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.

### 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.
- 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** field in the **install-config.yaml** file.

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"> <li>• <b>AWS::EC2::VPC</b></li> <li>• <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>• <b>AWS::EC2::Subnet</b></li> <li>• <b>AWS::EC2::SubnetNetworkACLAssociation</b></li> </ul>	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"> <li>• <b>AWS::EC2::InternetGateway</b></li> <li>• <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>• <b>AWS::EC2::RouteTable</b></li> <li>• <b>AWS::EC2::Route</b></li> <li>• <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>• <b>AWS::EC2::NatGateway</b></li> <li>• <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkAclEntry</b></li> </ul>	You must allow the VPC to access the following ports:	
		<b>Port</b>	<b>Reason</b>
		<b>80</b>	Inbound HTTP traffic
		<b>443</b>	Inbound HTTPS traffic
		<b>22</b>	Inbound SSH traffic
		<b>1024 - 65535</b>	Inbound ephemeral traffic
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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.	
		<b>0 - 65535</b>	Outbound ephemeral traffic

### 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.

### 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.

### 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.

## 7.4. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#) 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. 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 FIPS validated or Modules In Progress cryptographic libraries on 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. CREATING THE INSTALLATION CONFIGURATION FILE

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).



## 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.
- Have the **imageContentSources** values that were generated during mirror registry creation.
- Obtain the contents of the certificate for your mirror registry.
- Obtain service principal permissions at the subscription level.

## 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.
  - vii. Paste the [pull secret from the 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":{"<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.

3. Make any other modifications to the **install-config.yaml** file that you require. 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.

### 7.6.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.

**NOTE**

After installation, you cannot modify these parameters in the **install-config.yaml** file.

#### 7.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 7.1. Required parameters

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;</b> . <b>&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object

Parameter	Description	Values
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}.{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 7.6.1.2. 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 7.2. Network parameters

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object  <b>NOTE</b> You cannot modify parameters specified by the <b>networking</b> object after installation.
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:   - cidr: 10.0.0/16</pre>
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation.  For example, <b>10.0.0/16</b> .   <p><b>NOTE</b></p> <p>Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.</p>

### 7.6.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:


Table 7.3. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String

Parameter	Description	Values
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>

Parameter	Description	Values
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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



Parameter	Description	Values
<b>controlPlane.hyperthreading</b>	<p>Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b>, 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><b>IMPORTANT</b></p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.

Parameter	Description	Values
<b>credentialsMode</b>	<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> <div data-bbox="485 510 595 891" style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> <div data-bbox="485 936 595 1288" style="border: 1px solid black; padding: 5px;"> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<b>fips</b>	<p>Enable or disable FIPS mode. The default is <b>false</b> (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 1393" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p><b>IMPORTANT</b></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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> </div> <div data-bbox="486 1440 592 1637" style="background-color: #f0f0f0; padding: 5px;"> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<b>false</b> or <b>true</b>


Parameter	Description	Values
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

#### 7.6.1.4. Optional AWS configuration parameters


Optional AWS configuration parameters are described in the following table:

Table 7.4. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> or the <a href="#">key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .

Parameter	Description	Values
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	<p>Any valid <a href="#">AWS region</a>, such as <b>us-east-1</b>. 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; margin-top: 10px;">  <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability</a> map in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.

Parameter	Description	Values
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.   <b>NOTE</b> You can add up to 25 user defined tags during installation. The remaining 25 tags are reserved for OpenShift Container Platform.
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

## 7.6.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 7.5. 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



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: (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.

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](#)

### 7.6.3. 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-96c6f8f7 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.
- 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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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.

## 7.6.4. 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: 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 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.

## 7.7. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.

**NOTE**

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 7.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



### NOTE

For macOS arm64, choose the **OpenShift v4.12 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. 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. 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.

## 7.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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. 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](#).

## CHAPTER 8. INSTALLING A CLUSTER ON AWS INTO AN EXISTING VPC

In OpenShift Container Platform version 4.12, 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.

### 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-lived 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.
- 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 IAM credentials](#).

### 8.2. ABOUT USING A CUSTOM VPC

In OpenShift Container Platform 4.12, 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.

#### 8.2.1. 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.

- 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** field in the **install-config.yaml** file.

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

Component	AWS type	Description												
VPC	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::VPC</b></li> <li>● <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::SubnetNetworkAclAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkAclEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> <tr> <td><b>443</b></td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td><b>22</b></td> <td>Inbound SSH traffic</td> </tr> <tr> <td><b>1024 - 65535</b></td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td><b>0 - 65535</b></td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic	<b>443</b>	Inbound HTTPS traffic	<b>22</b>	Inbound SSH traffic	<b>1024 - 65535</b>	Inbound ephemeral traffic	<b>0 - 65535</b>	Outbound ephemeral traffic
Port	Reason													
<b>80</b>	Inbound HTTP traffic													
<b>443</b>	Inbound HTTPS traffic													
<b>22</b>	Inbound SSH traffic													
<b>1024 - 65535</b>	Inbound ephemeral traffic													
<b>0 - 65535</b>	Outbound ephemeral traffic													

Component	AWS type	Description
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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.

### 8.2.2. 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.

### 8.2.3. 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.

### 8.2.4. 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.

### 8.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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. 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 FIPS validated or Modules In Process cryptographic libraries on 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. 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 the 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.

## 8.6. CREATING THE INSTALLATION CONFIGURATION FILE

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Obtain service principal permissions at the subscription level.

### 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.
  - vii. Paste the [pull secret from the Red Hat OpenShift Cluster Manager](#) .
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.

### 8.6.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.



#### 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 8.1. Required parameters

Parameter	Description	Values
-----------	-------------	--------


Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;.&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}</b> . <b>{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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 8.2. Network parameters

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;">  <div> <p><b>NOTE</b></p> <p>You cannot modify parameters specified by the <b>networking</b> object after installation.</p> </div> </div>

Parameter	Description	Values
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:     - cidr: 10.0.0.0/16</pre>




Parameter	Description	Values
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation. For example, <b>10.0.0.0/16</b> .  <b>NOTE</b> Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.


### 8.6.1.3. Optional configuration parameters



Optional installation configuration parameters are described in the following table:

Table 8.3. Optional parameters


Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.

Parameter	Description	Values
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .

Parameter	Description	Values
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>controlPlane.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>

Parameter	Description	Values
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.
<b>credentialsMode</b>	<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> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ( <b>""</b> ).
<b>fips</b>	Enable or disable FIPS mode. The default is <b>false</b> (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.	<b>false</b> or <b>true</b>

Parameter	Description	Values
	<p data-bbox="488 192 592 909"></p> <p data-bbox="671 114 863 143"><b>IMPORTANT</b></p> <p data-bbox="671 181 932 909">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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> <p data-bbox="488 958 592 1155"></p> <p data-bbox="671 965 762 994"><b>NOTE</b></p> <p data-bbox="671 1032 922 1151">If you are using Azure File storage, you cannot enable FIPS mode.</p>	
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.

Parameter	Description	Values
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

### 8.6.1.4. Optional AWS configuration parameters

Optional AWS configuration parameters are described in the following table:


Table 8.4. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> or the <a href="#">key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	Any valid <a href="#">AWS region</a> , such as <b>us-east-1</b> . You can use the AWS CLI to access the regions available based on your selected instance type. For example: <pre>aws ec2 describe-instance-type-offerings -- filters Name=instance- type,Values=c7g.xlarge</pre> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="background-color: black; width: 20px; height: 20px; margin-right: 5px;"></div> <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability</a> map in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.



Parameter	Description	Values
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .

Parameter	Description	Values
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	<p>Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></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>
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .

Parameter	Description	Values
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

## 8.6.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

**Table 8.5. 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:  $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$ .
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.

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](#)

### 8.6.3. 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 "Minimum resource requirements for cluster installation".

#### Example 8.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***
- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

### 8.6.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 8.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

## 8.6.5. 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-96c6f8f7 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.

**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 **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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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.

### 8.6.6. 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

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.

## 8.7. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 8.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.

**NOTE**

For macOS arm64, choose the **OpenShift v4.12 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>
```

**8.9. 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
```

## 8.10. 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.

## 8.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 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](#).

## CHAPTER 9. INSTALLING A PRIVATE CLUSTER ON AWS

In OpenShift Container Platform version 4.12, 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.

### 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 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-lived 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.
- 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 IAM credentials](#).

### 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.

### 9.2.1. 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.

#### 9.2.1.1. 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.

## 9.3. ABOUT USING A CUSTOM VPC

In OpenShift Container Platform 4.12, 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.

### 9.3.1. 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.

- 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** field in the **install-config.yaml** file.

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"> <li>• <b>AWS::EC2::VPC</b></li> <li>• <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::SubnetNetworkACLAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkACLEntry</b></li> </ul>	You must allow the VPC to access the following ports:	
		<b>Port</b>	<b>Reason</b>
		<b>80</b>	Inbound HTTP traffic
		<b>443</b>	Inbound HTTPS traffic
		<b>22</b>	Inbound SSH traffic
		<b>1024 - 65535</b>	Inbound ephemeral traffic
	<b>0 - 65535</b>	Outbound ephemeral traffic	
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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.	

### 9.3.2. 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.

### 9.3.3. 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.

### 9.3.4. 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.

## 9.4. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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. 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 FIPS validated or Modules In Process cryptographic libraries on 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. 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 the 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. 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.

### 9.7.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.

**NOTE**

After installation, you cannot modify these parameters in the **install-config.yaml** file.

**9.7.1.1. Required configuration parameters**

Required installation configuration parameters are described in the following table:

**Table 9.1. Required parameters**

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;.&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}.{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .

Parameter	Description	Values
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <code>{}</code> . For additional information about <b>platform</b> . <code>&lt;platform&gt;</code> parameters, consult the table for your specific platform that follows.	Object
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 9.7.1.2. 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 9.2. Network parameters

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object  <b>NOTE</b> You cannot modify parameters specified by the <b>networking</b> object after installation.
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:   - cidr: 10.0.0.0/16</pre>
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation.  For example, <b>10.0.0.0/16</b> .   <p><b>NOTE</b></p> <p>Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.</p>

### 9.7.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 9.3. Optional parameters



Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String

Parameter	Description	Values
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>


Parameter	Description	Values
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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

Parameter	Description	Values
<b>controlPlane.hyperthreading</b>	<p>Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b>, 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><b>IMPORTANT</b></p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.



Parameter	Description	Values
<b>credentialsMode</b>	<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> <div data-bbox="485 510 595 891" style="border: 1px solid black; padding: 5px; margin-bottom: 10px;">  </div> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> <div data-bbox="485 936 595 1285" style="border: 1px solid black; padding: 5px;">  </div> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<p><b>fips</b></p>	<p>Enable or disable FIPS mode. The default is <b>false</b> (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 1391" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p><b>IMPORTANT</b></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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> </div> <div data-bbox="486 1440 592 1637" style="background-color: #f0f0f0; padding: 5px;"> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<p><b>false</b> or <b>true</b></p>


Parameter	Description	Values
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

#### 9.7.1.4. Optional AWS configuration parameters


Optional AWS configuration parameters are described in the following table:

Table 9.4. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID or the key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .

Parameter	Description	Values
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	<p>Any valid <a href="#">AWS region</a>, such as <b>us-east-1</b>. 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; margin-top: 10px;">  <div style="flex-grow: 1;"> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability</a> map in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.

Parameter	Description	Values
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.   <b>NOTE</b> You can add up to 25 user defined tags during installation. The remaining 25 tags are reserved for OpenShift Container Platform.
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

## 9.7.2. 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

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.

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](#)

### 9.7.3. 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 "Minimum resource requirements for cluster installation".

#### Example 9.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***



- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

#### 9.7.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 9.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

#### 9.7.5. 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-96c6f8f7 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.
- 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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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**.

### 9.7.6. 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

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

**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. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



#### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 9.9. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



#### NOTE

For macOS arm64, choose the **OpenShift v4.12 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. 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. 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**

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

## 9.12. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 9.13. 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](#) .

## CHAPTER 10. INSTALLING A CLUSTER ON AWS INTO A GOVERNMENT REGION

In OpenShift Container Platform version 4.12, 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.

### 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](#).
- 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-lived 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.
- 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 IAM credentials](#).

### 10.2. 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**

### 10.3. 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**).

## 10.4. 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.

### 10.4.1. 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.

#### 10.4.1.1. 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.

## 10.5. ABOUT USING A CUSTOM VPC

In OpenShift Container Platform 4.12, 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.

### 10.5.1. 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.
- 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** field in the **install-config.yaml** file.

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"> <li>• <b>AWS::EC2::VPC</b></li> <li>• <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>• <b>AWS::EC2::Subnet</b></li> <li>• <b>AWS::EC2::SubnetNetworkACLAssociation</b></li> </ul>	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"> <li>• <b>AWS::EC2::InternetGateway</b></li> <li>• <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>• <b>AWS::EC2::RouteTable</b></li> <li>• <b>AWS::EC2::Route</b></li> <li>• <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>• <b>AWS::EC2::NatGateway</b></li> <li>• <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>• <b>AWS::EC2::NetworkACL</b></li> <li>• <b>AWS::EC2::NetworkACLEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic
Port	Reason					
<b>80</b>	Inbound HTTP traffic					



Component	AWS type	Description	
		<b>443</b>	Inbound HTTPS traffic
		<b>22</b>	Inbound SSH traffic
		<b>1024 - 65535</b>	Inbound ephemeral traffic
		<b>0 - 65535</b>	Outbound ephemeral traffic
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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.	

### 10.5.2. 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.

### 10.5.3. 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.

#### 10.5.4. 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.

### 10.6. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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. 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 FIPS validated or Modules In Process cryptographic libraries on 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. 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 worker 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 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 worker 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.

## 10.9. 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 the 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.

## 10.10. 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.

## 10.10.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.



### 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 10.1. Required parameters

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;. &lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object

Parameter	Description	Values
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}.{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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 10.2. Network parameters



Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object  <b>NOTE</b> You cannot modify parameters specified by the <b>networking</b> object after installation.
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:   - cidr: 10.0.0.0/16</pre>
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation.  For example, <b>10.0.0.0/16</b> .   <p><b>NOTE</b></p> <p>Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.</p>

### 10.10.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:



Table 10.3. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String

Parameter	Description	Values
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>

Parameter	Description	Values
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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

Parameter	Description	Values
<b>controlPlane.hyperthreading</b>	<p>Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b>, 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><b>IMPORTANT</b></p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.

Parameter	Description	Values
<b>credentialsMode</b>	<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> <p> <b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> <p> <b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<b>fips</b>	<p>Enable or disable FIPS mode. The default is <b>false</b> (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 1391" style="background-color: black; color: white; padding: 5px; margin-bottom: 10px;"> <p><b>IMPORTANT</b></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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> </div> <div data-bbox="486 1435 592 1637" style="background-color: #f0f0f0; padding: 5px; margin-bottom: 10px;"> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div>	<b>false</b> or <b>true</b>

Parameter	Description	Values
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: center;">  <div> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>


#### 10.10.1.4. Optional AWS configuration parameters

Optional AWS configuration parameters are described in the following table:


Table 10.4. Optional AWS parameters



Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> or the <a href="#">key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .

Parameter	Description	Values
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	<p>Any valid <a href="#">AWS region</a>, such as <b>us-east-1</b>. 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; margin-top: 10px;">  <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability</a> map in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.

Parameter	Description	Values
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.   <b>NOTE</b> You can add up to 25 user defined tags during installation. The remaining 25 tags are reserved for OpenShift Container Platform.
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

### 10.10.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 10.5. 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

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: (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.

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](#)

### 10.10.3. 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 "Minimum resource requirements for cluster installation".

#### Example 10.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***

- **m4.\***
- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

#### 10.10.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 10.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

#### 10.10.5. 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-96c6f8f7 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.

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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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**.

### 10.10.6. 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

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

**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.

## 10.11. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



#### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 10.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



#### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 10.13. 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
```

## 10.14. 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**

- See [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.

## 10.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 10.16. 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](#) .



## CHAPTER 11. INSTALLING A CLUSTER ON AWS INTO A SECRET OR TOP SECRET REGION

In OpenShift Container Platform version 4.12, 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.

### 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 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-lived 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.
- 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 IAM credentials](#).

### 11.2. 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*


## 11.3. 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.

## 11.4. 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.

### 11.4.1. 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.

#### 11.4.1.1. 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.

## 11.5. ABOUT USING A CUSTOM VPC

In OpenShift Container Platform 4.12, 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.

## 11.5.1. 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.
- 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** field in the **install-config.yaml** file.

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"> <li>● <b>AWS::EC2::VPC</b></li> <li>● <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::SubnetNetworkACLAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkACL</b></li> <li>● <b>AWS::EC2::NetworkACLEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> <tr> <td><b>443</b></td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td><b>22</b></td> <td>Inbound SSH traffic</td> </tr> <tr> <td><b>1024 - 65535</b></td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td><b>0 - 65535</b></td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic	<b>443</b>	Inbound HTTPS traffic	<b>22</b>	Inbound SSH traffic	<b>1024 - 65535</b>	Inbound ephemeral traffic	<b>0 - 65535</b>	Outbound ephemeral traffic
Port	Reason													
<b>80</b>	Inbound HTTP traffic													
<b>443</b>	Inbound HTTPS traffic													
<b>22</b>	Inbound SSH traffic													
<b>1024 - 65535</b>	Inbound ephemeral traffic													
<b>0 - 65535</b>	Outbound ephemeral traffic													

Component	AWS type	Description
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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.

### 11.5.2. 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.

### 11.5.3. 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.

### 11.5.4. 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.

## 11.6. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 11.7. 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> ❶
```

- Export the region to associate with your custom AMI as an environment variable:

```
$ export AWS_DEFAULT_REGION=<aws_region> ❶
```

- Export the version of RHCOS you uploaded to Amazon S3 as an environment variable:

```
$ export RHCOS_VERSION=<version> ❶
```

❶ ❶ ❶ The RHCOS VMDK version, like **4.12.0**.

- Export the Amazon S3 bucket name as an environment variable:

```
$ export VMIMPORT_BUCKET_NAME=<s3_bucket_name>
```

- 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
```

- 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.

- 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](#).

## 11.8. 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 FIPS validated or Modules In Process cryptographic libraries on 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.

## 11.9. 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 the 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.

## 11.10. 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.

### 11.10.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.

**NOTE**

After installation, you cannot modify these parameters in the **install-config.yaml** file.

#### 11.10.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

**Table 11.1. Required parameters**

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String

Parameter	Description	Values
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;</b> . <b>&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}</b> . <b>{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 11.10.1.2. 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 11.2. Network parameters

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;"> <div> <p><b>NOTE</b></p> <p>You cannot modify parameters specified by the <b>networking</b> object after installation.</p> </div> </div>



Parameter	Description	Values
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:     - cidr: 10.0.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation. For example, <b>10.0.0.0/16</b> .  <b>NOTE</b> Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.


### 11.10.1.3. Optional configuration parameters



Optional installation configuration parameters are described in the following table:

Table 11.3. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.

Parameter	Description	Values
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .

Parameter	Description	Values
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>controlPlane.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>

Parameter	Description	Values
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.
<b>credentialsMode</b>	<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> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ( <b>""</b> ).
<b>fips</b>	Enable or disable FIPS mode. The default is <b>false</b> (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.	<b>false</b> or <b>true</b>

Parameter	Description	Values
	<p data-bbox="671 107 863 141"><b>IMPORTANT</b></p> <p data-bbox="671 181 932 909">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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> <p data-bbox="671 965 762 999"><b>NOTE</b></p> <p data-bbox="671 1032 922 1151">If you are using Azure File storage, you cannot enable FIPS mode.</p>	
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.

Parameter	Description	Values
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

### 11.10.1.4. Optional AWS configuration parameters

Optional AWS configuration parameters are described in the following table:

Table 11.4. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> or the <a href="#">key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.



Parameter	Description	Values
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	Any valid <a href="#">AWS region</a> , such as <b>us-east-1</b> . You can use the AWS CLI to access the regions available based on your selected instance type. For example: <pre>aws ec2 describe-instance-type-offerings -- filters Name=instance- type,Values=c7g.xlarge</pre> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="width: 20px; height: 20px; background-color: black; margin-right: 5px;"></div> <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability</a> map in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.

Parameter	Description	Values
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .

Parameter	Description	Values
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	<p>Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></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>
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .

Parameter	Description	Values
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

### 11.10.2. 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 "Minimum resource requirements for cluster installation".

#### Example 11.1. Machine types based on 64-bit x86 architecture for secret regions

- c4.\*
- c5.\*
- i3.\*
- m4.\*
- m5.\*
- r4.\*
- r5.\*
- t3.\*

### 11.10.3. 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.

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.

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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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.

### 11.10.4. 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: 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.

## 11.11. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

## 11.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



#### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 11.13. 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
```

## 11.14. 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](#)

**11.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM**

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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](#)

**11.16. 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](#) .

## CHAPTER 12. INSTALLING A CLUSTER ON AWS CHINA

In OpenShift Container Platform version 4.12, you can install a cluster to the following Amazon Web Services (AWS) China regions:

- **cn-north-1** (Beijing)
- **cn-northwest-1** (Ningxia)

### 12.1. 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.
- 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 IAM credentials](#).



#### 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-lived 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.

### 12.2. 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.

### 12.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 12.4. 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.

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.

### 12.4.1. 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.

#### 12.4.1.1. 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.

## 12.5. ABOUT USING A CUSTOM VPC

In OpenShift Container Platform 4.12, 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.

### 12.5.1. 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.
- 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** field in the **install-config.yaml** file.

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"> <li>• <b>AWS::EC2::VPC</b></li> <li>• <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>• <b>AWS::EC2::Subnet</b></li> <li>• <b>AWS::EC2::SubnetNetworkACLAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	<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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkAclEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> <tr> <td><b>443</b></td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td><b>22</b></td> <td>Inbound SSH traffic</td> </tr> <tr> <td><b>1024 - 65535</b></td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td><b>0 - 65535</b></td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic	<b>443</b>	Inbound HTTPS traffic	<b>22</b>	Inbound SSH traffic	<b>1024 - 65535</b>	Inbound ephemeral traffic	<b>0 - 65535</b>	Outbound ephemeral traffic
Port	Reason													
<b>80</b>	Inbound HTTP traffic													
<b>443</b>	Inbound HTTPS traffic													
<b>22</b>	Inbound SSH traffic													
<b>1024 - 65535</b>	Inbound ephemeral traffic													
<b>0 - 65535</b>	Outbound ephemeral traffic													
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	<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>												

### 12.5.2. 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.

### 12.5.3. 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.

### 12.5.4. 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.

## 12.6. 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 FIPS validated or Modules In Process cryptographic libraries on 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.

## 12.7. 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 **4.12.0**.

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](#).

## 12.8. 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 the 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.

## 12.9. 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.

## 12.9.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.



### NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

### 12.9.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

**Table 12.1. Required parameters**

Parameter	Description	Values
-----------	-------------	--------

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;</b> . <b>&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}</b> . <b>{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud</b> , <b>aws</b> , <b>baremetal</b> , <b>azure</b> , <b>gcp</b> , <b>ibmcloud</b> , <b>nutanix</b> , <b>openstack</b> , <b>ovirt</b> , <b>vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 12.9.1.2. 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 12.2. Network parameters

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object <div style="display: flex; align-items: center; margin-top: 10px;">  <div> <p><b>NOTE</b></p> <p>You cannot modify parameters specified by the <b>networking</b> object after installation.</p> </div> </div>

Parameter	Description	Values
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:     - cidr: 10.0.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation. For example, <b>10.0.0.0/16</b> .  <b>NOTE</b> Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.

### 12.9.1.3. Optional configuration parameters


Optional installation configuration parameters are described in the following table:

Table 12.3. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.



Parameter	Description	Values
<b>compute.architecture</b>	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 <b>amd64</b> (the default).	String
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .



Parameter	Description	Values
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> (the default).	String
<b>controlPlane.hypertreading</b>	<p>Whether to enable or disable simultaneous multithreading, or <b>hypertreading</b>, 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><b>IMPORTANT</b></p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div> </div>	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere, or {}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.


Parameter	Description	Values
<b>credentialsMode</b>	<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> <div data-bbox="485 510 595 891" style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> <div data-bbox="485 936 595 1288" style="border: 1px solid black; padding: 5px;"> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<b>fips</b>	<p>Enable or disable FIPS mode. The default is <b>false</b> (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="488 584 592 1391" style="background-color: black; width: 65px; height: 360px; margin-bottom: 10px;"></div> <p><b>IMPORTANT</b></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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> <div data-bbox="488 1440 592 1637" style="background-color: black; width: 65px; height: 88px; margin-bottom: 10px;"></div> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p>	<b>false</b> or <b>true</b>
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String

Parameter	Description	Values
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<p><b>Internal</b> or <b>External</b>. The default value is <b>External</b>.</p> <p>Setting this field to <b>Internal</b> is not supported on non-cloud platforms.</p> <div style="display: flex; align-items: flex-start;">  <div> <p><b>IMPORTANT</b></p> <p>If the value of the field is set to <b>Internal</b>, the cluster will become non-functional. For more information, refer to <a href="#">BZ#1953035</a>.</p> </div> </div>
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;">  <div> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

### 12.9.2. 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:
      - 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.
- 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 supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 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](#). The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on 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**.

### 12.9.3. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

Table 12.4. 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.


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](#)

## 12.9.4. 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 "Minimum resource requirements for cluster installation".

#### Example 12.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***
- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***



- r5.\*
- r5a.\*
- r6i.\*
- t3.\*
- t3a.\*

### 12.9.5. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 12.2. Machine types based on 64-bit ARM architecture

- c6g.\*
- c7g.\*
- m6g.\*
- m7g.\*
- r8g.\*

### 12.9.6. 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

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.

## 12.10. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



#### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

- 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.

## 12.11. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



#### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 12.12. 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
```

## 12.13. 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.

## 12.14. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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 [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.
- See [About remote health monitoring](#) for more information about the Telemetry service.

## 12.15. 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](#).



## CHAPTER 13. INSTALLING A CLUSTER ON USER-PROVISIONED INFRASTRUCTURE IN AWS BY USING CLOUDFORMATION TEMPLATES

In OpenShift Container Platform version 4.12, 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.

### 13.1. 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-lived 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 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 IAM credentials](#).

## 13.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 13.3. 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.

### 13.3.1. Required machines for cluster installation

The smallest OpenShift Container Platform clusters require the following hosts:

**Table 13.1. 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) 8 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

### 13.3.2. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

**Table 13.2. 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.

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](#)

### 13.3.3. 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 "Minimum resource requirements for cluster installation".

#### Example 13.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***
- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

### 13.3.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 13.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

### 13.3.5. 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.

## 13.4. 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.

### 13.4.1. 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"> <li>● <b>AWS::EC2::VPC</b></li> <li>● <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::SubnetNetworkAclAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkAclEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> <tr> <td><b>443</b></td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td><b>22</b></td> <td>Inbound SSH traffic</td> </tr> <tr> <td><b>1024 - 65535</b></td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td><b>0 - 65535</b></td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic	<b>443</b>	Inbound HTTPS traffic	<b>22</b>	Inbound SSH traffic	<b>1024 - 65535</b>	Inbound ephemeral traffic	<b>0 - 65535</b>	Outbound ephemeral traffic
Port	Reason													
<b>80</b>	Inbound HTTP traffic													
<b>443</b>	Inbound HTTPS traffic													
<b>22</b>	Inbound SSH traffic													
<b>1024 - 65535</b>	Inbound ephemeral traffic													
<b>0 - 65535</b>	Outbound ephemeral traffic													

Component	AWS type	Description
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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	<b>AWS::Route53::HostedZone</b>	The hosted zone for your internal DNS.
Public load balancer	<b>AWS::ElasticLoadBalancingV2::LoadBalancer</b>	The load balancer for your public subnets.
External API server record	<b>AWS::Route53::RecordSetGroup</b>	Alias records for the external API server.
External listener	<b>AWS::ElasticLoadBalancingV2::Listener</b>	A listener on port 6443 for the external load balancer.
External target group	<b>AWS::ElasticLoadBalancingV2::TargetGroup</b>	The target group for the external load balancer.



Component	AWS type	Description
Private load balancer	<b>AWS::ElasticLoadBalancingV2::LoadBalancer</b>	The load balancer for your private subnets.
Internal API server record	<b>AWS::Route53::RecordSetGroup</b>	Alias records for the internal API server.
Internal listener	<b>AWS::ElasticLoadBalancingV2::Listener</b>	A listener on port 22623 for the internal load balancer.
Internal target group	<b>AWS::ElasticLoadBalancingV2::TargetGroup</b>	The target group for the internal load balancer.
Internal listener	<b>AWS::ElasticLoadBalancingV2::Listener</b>	A listener on port 6443 for the internal load balancer.
Internal target group	<b>AWS::ElasticLoadBalancingV2::TargetGroup</b>	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
<b>MasterSecurityGroup</b>	<b>AWS::EC2::SecurityGroup</b>	<b>icmp</b>	<b>0</b>
		<b>tcp</b>	<b>22</b>
		<b>tcp</b>	<b>6443</b>
		<b>tcp</b>	<b>22623</b>
<b>WorkerSecurityGroup</b>	<b>AWS::EC2::SecurityGroup</b>	<b>icmp</b>	<b>0</b>
		<b>tcp</b>	<b>22</b>

Group	Type	IP Protocol	Port range
<b>BootstrapSecurityGroup</b>	<b>AWS::EC2::SecurityGroup</b>	<b>tcp</b>	<b>22</b>
		<b>tcp</b>	<b>19531</b>

## 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
<b>MasterIngressEtc</b>	etcd	<b>tcp</b>	<b>2379- 2380</b>
<b>MasterIngressVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>
<b>MasterIngressWorkerVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>
<b>MasterIngressInternal</b>	Internal cluster communication and Kubernetes proxy metrics	<b>tcp</b>	<b>9000 - 9999</b>
<b>MasterIngressWorkerInternal</b>	Internal cluster communication	<b>tcp</b>	<b>9000 - 9999</b>
<b>MasterIngressKube</b>	Kubernetes kubelet, scheduler and controller manager	<b>tcp</b>	<b>10250 - 10259</b>
<b>MasterIngressWorkerKube</b>	Kubernetes kubelet, scheduler and controller manager	<b>tcp</b>	<b>10250 - 10259</b>
<b>MasterIngressIngressServices</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>MasterIngressWorkerIngressServices</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>MasterIngressGeneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>
<b>MasterIngressWorkerGeneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>

Ingress group	Description	IP protocol	Port range
<b>MasterIngressIpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>MasterIngressWorkerIpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>MasterIngressIpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>MasterIngressWorkerIpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>MasterIngressIpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>
<b>MasterIngressWorkerIpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>
<b>MasterIngressInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>MasterIngressWorkerInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>MasterIngressIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>
<b>MasterIngressWorkerIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>

## 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
<b>WorkerIngressVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>

Ingress group	Description	IP protocol	Port range
<b>WorkerIngress WorkerVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>
<b>WorkerIngress Internal</b>	Internal cluster communication	<b>tcp</b>	<b>9000 - 9999</b>
<b>WorkerIngress WorkerInternal</b>	Internal cluster communication	<b>tcp</b>	<b>9000 - 9999</b>
<b>WorkerIngress Kube</b>	Kubernetes kubelet, scheduler, and controller manager	<b>tcp</b>	<b>10250</b>
<b>WorkerIngress WorkerKube</b>	Kubernetes kubelet, scheduler, and controller manager	<b>tcp</b>	<b>10250</b>
<b>WorkerIngress IngressServices</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>WorkerIngress WorkerIngress Services</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>WorkerIngress Geneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>
<b>WorkerIngress MasterGeneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>
<b>WorkerIngress IpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>WorkerIngress MasterIpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>WorkerIngress IpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>WorkerIngress MasterIpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>WorkerIngress IpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>

Ingress group	Description	IP protocol	Port range
<b>WorkerIngressMasterIpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>
<b>WorkerIngressInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>WorkerIngressMasterInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>WorkerIngressIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>
<b>WorkerIngressMasterIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>

### 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	<b>Allow</b>	<b>ec2:*</b>	<b>*</b>
	<b>Allow</b>	<b>elasticloadbalancing:*</b>	<b>*</b>
	<b>Allow</b>	<b>iam:PassRole</b>	<b>*</b>
	<b>Allow</b>	<b>s3:GetObject</b>	<b>*</b>
Worker	<b>Allow</b>	<b>ec2:Describe*</b>	<b>*</b>
Bootstrap	<b>Allow</b>	<b>ec2:Describe*</b>	<b>*</b>
	<b>Allow</b>	<b>ec2:AttachVolume</b>	<b>*</b>
	<b>Allow</b>	<b>ec2:DetachVolume</b>	<b>*</b>

### 13.4.2. 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.

### 13.4.3. 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 13.3. Required EC2 permissions for installation


- **ec2:AuthorizeSecurityGroupEgress**
- **ec2:AuthorizeSecurityGroupIngress**
- **ec2:CopyImage**
- **ec2:CreateNetworkInterface**
- **ec2:AttachNetworkInterface**
- **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:DescribeRegions**
- **ec2:DescribeRouteTables**
- **ec2:DescribeSecurityGroups**
- **ec2:DescribeSubnets**
- **ec2:DescribeTags**
- **ec2:DescribeVolumes**
- **ec2:DescribeVpcAttribute**
- **ec2:DescribeVpcClassicLink**
- **ec2:DescribeVpcClassicLinkDnsSupport**
- **ec2:DescribeVpcEndpoints**
- **ec2:DescribeVpcs**
- **ec2:GetEbsDefaultKmsKeyId**
- **ec2:ModifyInstanceAttribute**
- **ec2:ModifyNetworkInterfaceAttribute**
- **ec2:RevokeSecurityGroupEgress**
- **ec2:RevokeSecurityGroupIngress**
- **ec2:RunInstances**

- **ec2:TerminateInstances**

#### Example 13.4. 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 VPC, your account does not require these permissions for creating network resources.

#### Example 13.5. Required Elastic Load Balancing permissions (ELB) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:ApplySecurityGroupsToLoadBalancer**
- **elasticloadbalancing:AttachLoadBalancerToSubnets**
- **elasticloadbalancing:ConfigureHealthCheck**
- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateLoadBalancerListeners**
- **elasticloadbalancing>DeleteLoadBalancer**



- **elasticloadbalancing:DeregisterInstancesFromLoadBalancer**
- **elasticloadbalancing:DescribeInstanceHealth**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTags**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:RegisterInstancesWithLoadBalancer**
- **elasticloadbalancing:SetLoadBalancerPoliciesOfListener**

Example 13.6. Required Elastic Load Balancing permissions (ELBv2) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:CreateListener**
- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateTargetGroup**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing:DeregisterTargets**
- **elasticloadbalancing:DescribeListeners**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTargetGroupAttributes**
- **elasticloadbalancing:DescribeTargetHealth**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:ModifyTargetGroup**
- **elasticloadbalancing:ModifyTargetGroupAttributes**
- **elasticloadbalancing:RegisterTargets**

Example 13.7. 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:TagRole**



#### NOTE

If you have not created a load balancer in your AWS account, the IAM user also requires the **iam:CreateServiceLinkedRole** permission.

#### Example 13.8. 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 13.9. Required S3 permissions for installation

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3:GetAccelerateConfiguration**
- **s3:GetBucketAcl**
- **s3:GetBucketCors**
- **s3:GetBucketLocation**
- **s3:GetBucketLogging**
- **s3:GetBucketPolicy**
- **s3:GetBucketObjectLockConfiguration**
- **s3:GetBucketReplication**
- **s3:GetBucketRequestPayment**
- **s3:GetBucketTagging**
- **s3:GetBucketVersioning**
- **s3:GetBucketWebsite**
- **s3:GetEncryptionConfiguration**
- **s3:GetLifecycleConfiguration**
- **s3:GetReplicationConfiguration**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketTagging**
- **s3:PutEncryptionConfiguration**

#### Example 13.10. S3 permissions that cluster Operators require

- **s3>DeleteObject**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**

- **s3:GetObjectVersion**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Example 13.11. Required permissions to delete base cluster resources

- **autoscaling:DescribeAutoScalingGroups**
- **ec2:DeletePlacementGroup**
- **ec2:DeleteNetworkInterface**
- **ec2:DeleteVolume**
- **elasticloadbalancing:DeleteTargetGroup**
- **elasticloadbalancing:DescribeTargetGroups**
- **iam:DeleteAccessKey**
- **iam:DeleteUser**
- **iam>ListAttachedRolePolicies**
- **iam>ListInstanceProfiles**
- **iam>ListRolePolicies**
- **iam>ListUserPolicies**
- **s3>DeleteObject**
- **s3>ListBucketVersions**
- **tag:GetResources**

Example 13.12. 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 13.13. Required permissions to delete a cluster with shared instance roles

- **iam:UntagRole**

#### Example 13.14. Additional IAM and S3 permissions that are required to create manifests

- **iam>DeleteAccessKey**
- **iam>DeleteUser**
- **iam>DeleteUserPolicy**
- **iam:GetUserPolicy**
- **iam>ListAccessKeys**
- **iam:PutUserPolicy**
- **iam:TagUser**
- **s3:PutBucketPublicAccessBlock**
- **s3:GetBucketPublicAccessBlock**
- **s3:PutLifecycleConfiguration**
- **s3>ListBucket**
- **s3>ListBucketMultipartUploads**
- **s3:AbortMultipartUpload**



#### 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 13.15. Optional permissions for instance and quota checks for installation

- `ec2:DescribeInstanceTypeOfferings`
- `servicequotas:ListAWSDefaultServiceQuotas`

## 13.5. 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 worker 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 region. If you use the CloudFormation template to deploy your worker nodes, you must update the `worker0.type.properties.ImageID` parameter with this value.

## 13.6. 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 the 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.

## 13.7. 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 FIPS validated or Modules In Process cryptographic libraries on 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. If you install a cluster on infrastructure that you provision, you must provide the key to the installation program.

## 13.8. 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.

### 13.8.1. 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: 4.12.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/<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

- 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.

### 13.8.2. 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 a region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to a 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 the 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.

## Additional resources

- See [Configuration and credential file settings](#) in the AWS documentation for more information about AWS profile and credential configuration.

### 13.8.3. 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> ❶
  httpsProxy: https://<username>:<pswd>@<ip>:<port> ❷
  noProxy: ec2.<aws_region>.amazonaws.com,elasticloadbalancing.
<aws_region>.amazonaws.com,s3.<aws_region>.amazonaws.com ❸
  additionalTrustBundle: | ❹
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> ❺
```

❶ A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.

❷ A proxy URL to use for creating HTTPS connections outside the cluster.

❸

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.

### 13.8.4. 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
```

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

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:

```
1. Open the <installation_directory>/manifests/cluster-scheduler-02-config.yml file
```

- 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. Optional: If you manually created a cloud identity and access management (IAM) role, locate any `CredentialsRequest` objects with the `TechPreviewNoUpgrade` annotation in the release image by running the following command:

```
$ oc adm release extract quay.io/openshift-release-dev/ocp-release:4.y.z-x86_64 --
credentials-requests --cloud=<platform_name>
```

### Example output

```
0000_30_capi-operator_00_credentials-request.yaml: release.openshift.io/feature-set:
TechPreviewNoUpgrade
```



### IMPORTANT

The release image includes `CredentialsRequest` objects for Technology Preview features that are enabled by the `TechPreviewNoUpgrade` feature set. You can identify these objects by their use of the `release.openshift.io/feature-set: TechPreviewNoUpgrade` annotation.

- If you are not using any of these features, do not create secrets for these objects. Creating secrets for Technology Preview features that you are not using can cause the installation to fail.
- If you are using any of these features, you must create secrets for the corresponding objects.

- a. Delete all `CredentialsRequest` objects that have the `TechPreviewNoUpgrade` annotation.



- 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
```

## 13.9. 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
```

- 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.

## 13.10. 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> ❶
  --template-body file://<template>.yaml ❷
  --parameters file://<parameters>.json ❸
```

- <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:

<b>VpcId</b>	The ID of your VPC.
<b>PublicSubnetIds</b>	The IDs of the new public subnets.
<b>PrivateSubnetIds</b>	The IDs of the new private subnets.

### 13.10.1. 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 13.16. 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[0-9][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:
    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"]]
    ]

```

### Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

## 13.11. 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:

<b>PrivateHostedZoneId</b>	Hosted zone ID for the private DNS.
<b>ExternalApiLoadBalancerName</b>	Full name of the external API load balancer.
<b>InternalApiLoadBalancerName</b>	Full name of the internal API load balancer.
<b>ApiServerDnsName</b>	Full hostname of the API server.
<b>RegisterNlbTargetLambda</b>	Lambda ARN useful to help register/deregister IP targets for these load balancers.
<b>ExternalApiTargetGroupArn</b>	ARN of external API target group.

<b>InternalApiTargetGroupArn</b>	ARN of internal API target group.
<b>InternalServiceTargetGroupArn</b>	ARN of internal service target group.

### 13.11.1. 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 13.17. 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  
Vpclid:  
Ref: Vpclid  
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  
Vpclid:  
Ref: Vpclid



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  
 Vpclid:  
 Ref: Vpclid  
 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:
      - "RegisterTargetLambdRole"
      - "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

```

#### RegisterSubnetTagsLambdRole:

```

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:
      - "RegisterSubnetTagsLambdalaRole"
      - "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.

## 13.12. 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
```

- 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:

<b>MasterSecurityGroupID</b>	Master Security Group ID
<b>WorkerSecurityGroupID</b>	Worker Security Group ID
<b>MasterInstanceProfile</b>	Master IAM Instance Profile
<b>WorkerInstanceProfile</b>	Worker IAM Instance Profile

### 13.12.1. 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 13.18. 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: [10.0.0.0/16](#)

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&gt;DeleteSecurityGroup"

- "ec2&gt;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](#).

## 13.13. 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.

## 13.14. 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 13.3. x86\_64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-073850a7021953a5c
ap-east-1	ami-0f8800a05c09be42d
ap-northeast-1	ami-0a226dbcc9a561c40
ap-northeast-2	ami-041ae0537e2eddec1
ap-northeast-3	ami-0bb8d9b69dc5b7670
ap-south-1	ami-0e9c18058fc5f94fd
ap-southeast-1	ami-03022d358ba2168be
ap-southeast-2	ami-09ffdc5be9b973be0

AWS zone	AWS AMI
ap-southeast-3	ami-0facf1a0edeb20314
ca-central-1	ami-028cea206c2d03317
eu-central-1	ami-002eb441f329ccb0f
eu-north-1	ami-0b1a1fb68b3b9fee7
eu-south-1	ami-0bd0fd41a1d3f799a
eu-west-1	ami-04504e8799057980c
eu-west-2	ami-0cc9297ddb3bce971
eu-west-3	ami-06f98f607a50937c6
me-south-1	ami-0fe39da7871a5b2a5
sa-east-1	ami-08265cc3226697767
us-east-1	ami-0fe05b1aa8dacfa90
us-east-2	ami-0ff64f495c7e977cf
us-gov-east-1	ami-0c99658076c41872a
us-gov-west-1	ami-0ca4acd5b8ba1cb1d
us-west-1	ami-01dc5d8e6bb6f23f4
us-west-2	ami-0404a109adfd00019

Table 13.4. aarch64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-0574bcc5f80b0ad9a
ap-east-1	ami-0a65e79822ae2d235
ap-northeast-1	ami-0f7ef19d48e22353b
ap-northeast-2	ami-051dc6de359975e3c

AWS zone	AWS AMI
ap-northeast-3	ami-0fd0b4222595650ac
ap-south-1	ami-05f9d14fe4a90ed6f
ap-southeast-1	ami-0afdb9133d22fba5f
ap-southeast-2	ami-0ef979abe82d07d44
ap-southeast-3	ami-025f9103ac4310e7f
ca-central-1	ami-0588cdf59e5c14847
eu-central-1	ami-0ef24c0e18f93fa42
eu-north-1	ami-0439e2a3bf315df1a
eu-south-1	ami-0714e7c2e0106cdd3
eu-west-1	ami-0b960e76764ccd0c3
eu-west-2	ami-02621f50de62b3b89
eu-west-3	ami-0933ce7f5e2bfb50e
me-south-1	ami-074bde61a2ab740ee
sa-east-1	ami-03b4f97cfc8033ae0
us-east-1	ami-02a574449d4f4d280
us-east-2	ami-020e5600ef28c60ae
us-gov-east-1	ami-069f60e1dcf766d24
us-gov-west-1	ami-0db3cda4dbaccda02
us-west-1	ami-0c90cabeb5dee3178
us-west-2	ami-0f96437a23aeae53f

### 13.14.1. 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.

### 13.14.2. 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 **4.12.0**.

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](#).

## 13.15. 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": "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>**.

- 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.

- 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.
- 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.
- 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
```

- <name>** is the name for the CloudFormation stack, such as **cluster-bootstrap**. 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.
- 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
```

- 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:

<b>Bootstrap InstanceId</b>	The bootstrap Instance ID.
-----------------------------	----------------------------

<b>Bootstrap PublicIp</b>	The bootstrap node public IP address.
<b>Bootstrap PrivateIp</b>	The bootstrap node private IP address.

### 13.15.1. 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 13.19. 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]{1-9}|[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}([0-9]{1-9}|[0-9]{2}|2[0-
    4][0-9]|25[0-5])\V([0-9]{1-9}|2[0-9]|3[0-2]))$
    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:

```

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

- 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.

## 13.16. 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": "RhcoshAmi", 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"
  },
  {
    "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](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>
```

## 13.16.1. 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 13.20. 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:
```

```
- AssociatePublicIp: "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](#).

## 13.17. 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](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.

- 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
```

- 1** **<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.
- 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-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.

**13.17.1. 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 13.21. 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 master nodes into.

Type: AWS::EC2::Subnet::Id

**WorkerSecurityGroupId:**

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/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 master 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](#).

## 13.18. 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 v1.25.0 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](#).

## 13.19. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



#### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 13.20. 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
```

## 13.21. 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   v1.25.0
master-1  Ready    master   63m   v1.25.0
master-2  Ready    master   64m   v1.25.0
```

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}}' | 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
```

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  v1.25.0
master-1  Ready   master 73m  v1.25.0
master-2  Ready   master 74m  v1.25.0
worker-0  Ready   worker 11m  v1.25.0
worker-1  Ready   worker 11m  v1.25.0
```



### 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](#).

## 13.22. 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	4.12.0	True	False	False	19m
baremetal	4.12.0	True	False	False	37m
cloud-credential	4.12.0	True	False	False	40m
cluster-autoscaler	4.12.0	True	False	False	37m
config-operator	4.12.0	True	False	False	38m

console	4.12.0	True	False	False	26m
csi-snapshot-controller	4.12.0	True	False	False	37m
dns	4.12.0	True	False	False	37m
etcd	4.12.0	True	False	False	36m
image-registry	4.12.0	True	False	False	31m
ingress	4.12.0	True	False	False	30m
insights	4.12.0	True	False	False	31m
kube-apiserver	4.12.0	True	False	False	26m
kube-controller-manager	4.12.0	True	False	False	36m
kube-scheduler	4.12.0	True	False	False	36m
kube-storage-version-migrator	4.12.0	True	False	False	37m
machine-api	4.12.0	True	False	False	29m
machine-approver	4.12.0	True	False	False	37m
machine-config	4.12.0	True	False	False	36m
marketplace	4.12.0	True	False	False	37m
monitoring	4.12.0	True	False	False	29m
network	4.12.0	True	False	False	38m
node-tuning	4.12.0	True	False	False	37m
openshift-apiserver	4.12.0	True	False	False	32m
openshift-controller-manager	4.12.0	True	False	False	30m
openshift-samples	4.12.0	True	False	False	32m
operator-lifecycle-manager	4.12.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.12.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.12.0	True	False	False	32m
service-ca	4.12.0	True	False	False	38m
storage	4.12.0	True	False	False	37m

2. Configure the Operators that are not available.

### 13.22.1. 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.

#### 13.22.1.1. 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.

## 13.22.1.2. 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.

## 13.23. 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](#).

## 13.24. 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

1. 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
```

- 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.

## 13.25. 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.

## 13.26. 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.

## 13.27. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 13.28. ADDITIONAL RESOURCES

- See [Working with stacks](#) in the AWS documentation for more information about AWS CloudFormation stacks.

## 13.29. 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](#).

## CHAPTER 14. INSTALLING A CLUSTER USING AWS LOCAL ZONES

In OpenShift Container Platform version 4.12, you can install a cluster on Amazon Web Services (AWS) into an existing VPC, extending workers to the edge of the Cloud Infrastructure using AWS Local Zones.

After you create an Amazon Web Service (AWS) Local Zone environment, and you deploy your cluster, you can use edge worker nodes to create user workloads in Local Zone subnets.

AWS Local Zones are a type of infrastructure that place Cloud Resources close to the metropolitan regions. For more information, see the [AWS Local Zones Documentation](#).

OpenShift Container Platform can be installed in existing VPCs with Local Zone subnets. The Local Zone subnets can be used to extend the regular workers' nodes to the edge networks. The edge worker nodes are dedicated to running user workloads.

One way to create the VPC and subnets 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 an installer-provisioned infrastructure installation are provided as an example only. Installing a cluster with VPC you provide requires knowledge of the cloud provider and the installation process of OpenShift Container Platform. The 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.

### 14.1. 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-lived 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 noted the region and supported [AWS Local Zones locations](#) to create the network resources in.
- You read the [Features](#) for each AWS Local Zones location.



- 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 IAM credentials](#).

## 14.2. CLUSTER LIMITATIONS IN AWS LOCAL ZONES

Some limitations exist when you attempt to deploy a cluster with a default installation configuration in Amazon Web Services (AWS) Local Zones.



#### IMPORTANT

The following list details limitations when deploying a cluster in AWS Local Zones:

- The Maximum Transmission Unit (MTU) between an Amazon EC2 instance in a Local 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 on the deployment.
- Network resources such as Network Load Balancer (NLB), Classic Load Balancer, and Network Address Translation (NAT) Gateways are not supported in AWS Local Zones.
- 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 Local Zone locations. By default, the nodes running in Local Zones are deployed with the **gp2** EBS volume. The **gp2-csi StorageClass** must be set when creating workloads on Local Zone nodes.

#### Additional resources

- [Storage classes](#)

## 14.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 14.4. OPTING INTO AWS LOCAL ZONES

If you plan to create the subnets in AWS Local Zones, you must opt in to each zone group separately.

### Prerequisites

- You have installed the AWS CLI.
- You have determined into which region you will deploy your OpenShift Container Platform cluster.

### Procedure

1. Export a variable to contain the name of the region in which you plan to deploy your OpenShift Container Platform cluster by running the following command:

```
$ export CLUSTER_REGION="<region_name>" 1
```

- 1** For **<region\_name>**, specify a valid AWS region name, such as **us-east-1**.

2. List the zones that are available in your region by running the following command:

```
$ aws --region ${CLUSTER_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 region, the list of available zones can be long. The command will return the following fields:

#### ZoneName

The name of the Local Zone.

#### GroupName

The group that the zone is part of. You need to save this name to opt in.

#### Status

The status of the Local Zone group. If the status is **not-opted-in**, you must opt in the **GroupName** by running the commands that follow.

3. Export a variable to contain the name of the Local Zone to host your VPC by running the following command:

```
$ export ZONE_GROUP_NAME="<value_of_GroupName>" 1
```

- 1** The **<value\_of\_GroupName>** specifies the name of the group of the Local Zone you want to create subnets on. For example, specify **us-east-1-nyc-1** to use the zone **us-east-1-nyc-1a**, US East (New York).

4. Opt in to the zone group on your AWS account by running the following command:

```
$ aws ec2 modify-availability-zone-group \
  --group-name "${ZONE_GROUP_NAME}" \
  --opt-in-status opted-in
```

## 14.5. 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 worker 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

- Complete the OpenShift Container Platform subscription from the [AWS Marketplace](#).
- Record the AMI ID for your specific 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 worker 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

## 14.6. CREATING A VPC THAT USES AWS LOCAL ZONES

You must create a Virtual Private Cloud (VPC), and subnets for each Local Zone location, in Amazon Web Services (AWS) for your OpenShift Container Platform cluster to extend worker nodes to the edge locations. You can further customize the VPC to meet your requirements, including VPN, route tables, and add new Local Zone subnets that are 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 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 template requires:

```
[
  {
    "ParameterKey": "ClusterName", 1
    "ParameterValue": "mycluster" 2
  },
  {
    "ParameterKey": "VpcCidr", 3
    "ParameterValue": "10.0.0.0/16" 4
  },
  {
    "ParameterKey": "AvailabilityZoneCount", 5
    "ParameterValue": "3" 6
  },
  {
    "ParameterKey": "SubnetBits", 7
    "ParameterValue": "12" 8
  }
]
```

- 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 CIDR block for the VPC.
  - 4 Specify a CIDR block in the format **x.x.x.x/16-24**.
  - 5 The number of availability zones to deploy the VPC in.
  - 6 Specify an integer between **1** and **3**.
  - 7 The size of each subnet in each availability zone.
  - 8 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 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 to and 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:

<b>VpcId</b>	The ID of your VPC.
<b>PublicSubnetIds</b>	The IDs of the new public subnets.
<b>PrivateSubnetIds</b>	The IDs of the new private subnets.
<b>PublicRouteTableId</b>	The ID of the new public route table ID.

### 14.6.1. CloudFormation template for the VPC that uses AWS Local Zones

You can use the following CloudFormation template to deploy the VPC that you need for your OpenShift Container Platform cluster that uses AWS Local Zones.

#### Example 14.1. CloudFormation template for the VPC

```

AWSTemplateFormatVersion: 2010-09-09
Description: Template for Best Practice VPC with 1-3 AZs

Parameters:
  ClusterName:
    Type: String
    Description: ClusterName used to prefix resource names
  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:
ClusterName:
  default: ""
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
  Tags:
    - Key: Name
      Value: !Join [ "", [ !Ref ClusterName, "-vpc" ] ]
    - Key: !Join [ "", [ "kubernetes.io/cluster/unmanaged" ] ]
      Value: "shared"

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"
  Tags:
    - Key: Name
      Value: !Join [ "", [ !Ref ClusterName, "-public-1" ] ]

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"
  Tags:
    - Key: Name

```

```

    Value: !Join [ "", [ !Ref ClusterName, "-public-2" ] ]
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"
  Tags:
    - Key: Name
      Value: !Join [ "", [ !Ref ClusterName, "-public-3" ] ]

InternetGateway:
  Type: "AWS::EC2::InternetGateway"
  Properties:
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-igw" ] ]
GatewayToInternet:
  Type: "AWS::EC2::VPCGatewayAttachment"
  Properties:
    VpId: !Ref VPC
    InternetGatewayId: !Ref InternetGateway

PublicRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpId: !Ref VPC
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-rtb-public" ] ]
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"
  Properties:
    SubnetId: !Ref PublicSubnet2
    RouteTableId: !Ref PublicRouteTable
PublicSubnetRouteTableAssociation3:
  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"
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-private-1" ] ]
PrivateRouteTable:
  Type: "AWS::EC2::RouteTable"
  Properties:
    VpcId: !Ref VPC
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-rtb-private-1" ] ]
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
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-natgw-private-1" ] ]
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"
Tags:
- Key: Name
  Value: !Join [ "", [ !Ref ClusterName, "-private-2" ] ]
PrivateRouteTable2:
Type: "AWS::EC2::RouteTable"
Condition: DoAz2
Properties:
  VpcId: !Ref VPC
Tags:
- Key: Name
  Value: !Join [ "", [ !Ref ClusterName, "-rtb-private-2" ] ]
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
Tags:
- Key: Name
  Value: !Join [ "", [ !Ref ClusterName, "-natgw-private-2" ] ]
EIP2:
Type: "AWS::EC2::EIP"
Condition: DoAz2
Properties:
  Domain: vpc
Tags:
- Key: Name
  Value: !Join [ "", [ !Ref ClusterName, "-eip-private-2" ] ]
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"
Tags:
- Key: Name
  Value: !Join [ "", [ !Ref ClusterName, "-private-3" ] ]
PrivateRouteTable3:
  Type: "AWS::EC2::RouteTable"
  Condition: DoAz3
  Properties:
    Vpclid: !Ref VPC
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-rtb-private-3" ] ]
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
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-natgw-private-3" ] ]
EIP3:
  Type: "AWS::EC2::EIP"
  Condition: DoAz3
  Properties:
    Domain: vpc
    Tags:
      - Key: Name
        Value: !Join [ "", [ !Ref ClusterName, "-eip-private-3" ] ]
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

PrivateRouteTableId:

Description: Private Route table ID

Value: !Ref PrivateRouteTable

## 14.7. CREATING A SUBNET IN AWS LOCAL ZONES

You must create a subnet in AWS Local Zones before you configure a worker machineset for your OpenShift Container Platform cluster.

You must repeat the following process for each Local Zone you want to deploy worker nodes to.

You can use the provided CloudFormation template and a custom parameter file to create a stack of AWS resources that represent the 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 Zone group.


## Procedure

1. Create a JSON file that contains the parameter values that the template requires:

```
[
  {
    "ParameterKey": "ClusterName", 1
    "ParameterValue": "mycluster" 2
  },
  {
    "ParameterKey": "VpcId", 3
    "ParameterValue": "vpc-<random_string>" 4
  },
  {
    "ParameterKey": "PublicRouteTableId", 5
    "ParameterValue": "<vpc_rtb_pub>" 6
  },
  {
    "ParameterKey": "LocalZoneName", 7
    "ParameterValue": "<cluster_region_name>-<location_identifier>-<zone_identifier>" 8
  },
  {
    "ParameterKey": "LocalZoneNameShort", 9
    "ParameterValue": "<lz_zone_shortname>" 10
  },
  {
    "ParameterKey": "PublicSubnetCidr", 11
    "ParameterValue": "10.0.128.0/20" 12
  }
]
```

- 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 VPC ID in which the Local Zone's subnet will be created.
  - 4 Specify the **VpcId** value from the output of the CloudFormation template for the VPC.
  - 5 The Public Route Table ID for the VPC.
  - 6 Specify the **PublicRouteTableId** value from the output of the CloudFormation template for the VPC.
  - 7 The Local Zone name that the VPC belongs to.
  - 8 Specify the Local Zone that you opted your AWS account into, such as **us-east-1-nyc-1a**.
  - 9 The shortname of the AWS Local Zone that the VPC belongs to.
  - 10 Specify a short name for the AWS Local Zone that you opted your AWS account into, such as **<zone\_group\_identified><zone\_identifier>**. For example, **us-east-1-nyc-1a** is shortened to **nyc-1a**.
  - 11 The CIDR block to allow access to the Local Zone.
  - 12 Specify a CIDR block in the format **x.x.x.x/16-24**.
2. Copy the template from the **CloudFormation template for the subnet** 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 by running the following command:



### IMPORTANT

You must enter the command on a single line.

```
$ aws cloudformation create-stack --stack-name <subnet_stack_name> \ 1
  --template-body file://<template>.yaml \ 2
  --parameters file://<parameters>.json 3
```

- 1 **<subnet\_stack\_name>** is the name for the CloudFormation stack, such as **cluster-lz-<local\_zone\_shortname>**. 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-1:123456789012:stack/cluster-lz-nyc1/dbedae40-2fd3-11eb-820e-12a48460849f
```

4. Confirm that the template components exist by running the following command:

```
$ aws cloudformation describe-stacks --stack-name <subnet_stack_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:

<b>PublicSubnetIds</b>	The IDs of the new public subnets.
------------------------	------------------------------------

### 14.7.1. CloudFormation template for the subnet that uses AWS Local Zones

You can use the following CloudFormation template to deploy the subnet that you need for your OpenShift Container Platform cluster that uses AWS Local Zones.

#### Example 14.2. CloudFormation template for the subnet

```
# CloudFormation template used to create Local Zone subnets and dependencies
AWSTemplateFormatVersion: 2010-09-09
Description: Template for Best Practice VPC with 1-3 AZs

Parameters:
  ClusterName:
    Description: ClusterName used to prefix resource names
    Type: String
  VpcId:
    Description: VPC Id
    Type: String
  LocalZoneName:
    Description: Local Zone Name (Example us-east-1-bos-1)
    Type: String
  LocalZoneNameShort:
    Description: Short name for Local Zone used on tag Name (Example bos1)
    Type: String
  PublicRouteTableId:
    Description: Public Route Table ID to associate the Local Zone subnet
    Type: String
  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

Resources:
  PublicSubnet:
    Type: "AWS::EC2::Subnet"
    Properties:
      VpcId: !Ref VpcId
      CidrBlock: !Ref PublicSubnetCidr
      AvailabilityZone: !Ref LocalZoneName
    Tags:
      - Key: Name
        Value: !Join
```

```

- ""
- [ !Ref ClusterName, "-public-", !Ref LocalZoneNameShort, "-1" ]
- Key: kubernetes.io/cluster/unmanaged
  Value: "true"

```

```

PublicSubnetRouteTableAssociation:
  Type: "AWS::EC2::SubnetRouteTableAssociation"
  Properties:
    SubnetId: !Ref PublicSubnet
    RouteTableId: !Ref PublicRouteTableId

```

```

Outputs:
PublicSubnetIds:
  Description: Subnet IDs of the public subnets.
  Value:
    !Join [
      "",
      [!Ref PublicSubnet]
    ]

```

### Additional resources

- You can view details about the CloudFormation stacks that you create by navigating to the [AWS CloudFormation console](#).

## 14.8. 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

- 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.
- Select your infrastructure provider.
- 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 the 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.

## 14.9. 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 FIPS validated or Modules In Process cryptographic libraries on 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.

## 14.10. CREATING THE INSTALLATION FILES FOR AWS

To install OpenShift Container Platform on Amazon Web Services (AWS) and use AWS Local Zones, 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 and Kubernetes manifests.

### 14.10.1. Minimum resource requirements for cluster installation

Each cluster machine must meet the following minimum requirements:

**Table 14.1. 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.


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](#)

### 14.10.2. 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 "Minimum resource requirements for cluster installation".

### Example 14.3. 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 for more information about AWS Local Zones and the supported instances types and services.

## 14.10.3. 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 a region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to a 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. The region that you specify must be the same region that contains the Local Zone that you opted into for your AWS account.
  - 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 the Red Hat OpenShift Cluster Manager](#) .
2. Edit the **install-config.yaml** file to provide the subnets for the availability zones that your VPC uses:

```
platform:
aws:
subnets: 1
```

- publicSubnetId-1
- publicSubnetId-2
- publicSubnetId-3
- privateSubnetId-1
- privateSubnetId-2
- privateSubnetId-3

- 1 Add the **subnets** section and specify the **PrivateSubnetIds** and **PublicSubnetIds** values from the outputs of the CloudFormation template for the VPC. Do not include the Local Zone subnets here.

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.

## 14.10.4. Creating the Kubernetes manifest files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest files that the cluster needs to configure the machines.

### Prerequisites

- You obtained the OpenShift Container Platform installation program.
- You created the **install-config.yaml** installation configuration file.
- You installed the **jq** package.

### Procedure

1. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster by running the following command:

```
$ ./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. Set the default Maximum Transmission Unit (MTU) according to the network plugin:



## IMPORTANT

Generally, the Maximum Transmission Unit (MTU) between an Amazon EC2 instance in a Local Zone and an Amazon EC2 instance in the Region is 1300. See [How Local Zones work](#) in the AWS documentation. The cluster network MTU must be always less than the EC2 MTU to account for the overhead. The specific overhead is determined by your network plugin, for example:

- OVN-Kubernetes: **100 bytes**
- OpenShift SDN: **50 bytes**

The network plugin could provide additional features, like IPsec, that also must be decreased the MTU. Check the documentation for additional information.

- a. If you are using the **OVN-Kubernetes** network plugin, enter the following command:

```
$ cat <<EOF > <installation_directory>/manifests/cluster-network-03-config.yml
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      mtu: 1200
EOF
```

- b. If you are using the **OpenShift SDN** network plugin, enter the following command:

```
$ cat <<EOF > <installation_directory>/manifests/cluster-network-03-config.yml
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    openshiftSDNConfig:
      mtu: 1250
EOF
```

3. Create the machine set manifests for the worker nodes in your Local Zone.

- a. Export a local variable that contains the name of the Local Zone that you opted your AWS account into by running the following command:

```
$ export LZ_ZONE_NAME="<local_zone_name>" 1
```

- 1** For **<local\_zone\_name>**, specify the Local Zone that you opted your AWS account into, such as **us-east-1-nyc-1a**.

- b. Review the instance types for the location that you will deploy to by running the following command:

```
$ aws ec2 describe-instance-type-offerings \
  --location-type availability-zone \
  --filters Name=location,Values=${LZ_ZONE_NAME} \
  --region <region> 1
```

- 1** For **<region>**, specify the name of the region that you will deploy to, such as **us-east-1**.

- c. Export a variable to define the instance type for the worker machines to deploy on the Local Zone subnet by running the following command:

```
$ export INSTANCE_TYPE="<instance_type>" 1
```

- 1** Set **<instance\_type>** to a tested instance type, such as **c5d.2xlarge**.

- d. Store the AMI ID as a local variable by running the following command:

```
$ export AMI_ID=$(grep ami
  <installation_directory>/openshift/99_openshift-cluster-api_worker-machineset-0.yaml \
  | tail -n1 | awk '{print$2}')
```

- e. Store the subnet ID as a local variable by running the following command:

```
$ export SUBNET_ID=$(aws cloudformation describe-stacks --stack-name "
  <subnet_stack_name>" \
  | jq -r '.Stacks[0].Outputs[0].OutputValue') 1
```

- 1** For **<subnet\_stack\_name>**, specify the name of the subnet stack that you created.

- f. Store the cluster ID as local variable by running the following command:

```
$ export CLUSTER_ID="$(awk '/infrastructureName: / {print $2}'
  <installation_directory>/manifests/cluster-infrastructure-02-config.yml)"
```

- g. Create the worker manifest file for the Local Zone that your VPC uses by running the following command:

```
$ cat <<EOF > <installation_directory>/openshift/99_openshift-cluster-api_worker-
  machineset-nyc1.yaml
  apiVersion: machine.openshift.io/v1beta1
  kind: MachineSet
  metadata:
    labels:
      machine.openshift.io/cluster-api-cluster: ${CLUSTER_ID}
      name: ${CLUSTER_ID}-edge-${LZ_ZONE_NAME}
      namespace: openshift-machine-api
  spec:
    replicas: 1
    selector:
      matchLabels:
        machine.openshift.io/cluster-api-cluster: ${CLUSTER_ID}
```



```

    machine.openshift.io/cluster-api-machineset: ${CLUSTER_ID}-edge-
    ${LZ_ZONE_NAME}
  template:
    metadata:
      labels:
        machine.openshift.io/cluster-api-cluster: ${CLUSTER_ID}
        machine.openshift.io/cluster-api-machine-role: edge
        machine.openshift.io/cluster-api-machine-type: edge
        machine.openshift.io/cluster-api-machineset: ${CLUSTER_ID}-edge-
        ${LZ_ZONE_NAME}
    spec:
      metadata:
        labels:
          machine.openshift.com/zone-type: local-zone
          machine.openshift.com/zone-group: ${ZONE_GROUP_NAME}
          node-role.kubernetes.io/edge: ""
      taints:
        - key: node-role.kubernetes.io/edge
          effect: NoSchedule
      providerSpec:
        value:
          ami:
            id: ${AMI_ID}
          apiVersion: machine.openshift.io/v1beta1
          blockDevices:
            - ebs:
                volumeSize: 120
                volumeType: gp2
          credentialsSecret:
            name: aws-cloud-credentials
          deviceIndex: 0
          iamInstanceProfile:
            id: ${CLUSTER_ID}-worker-profile
          instanceType: ${INSTANCE_TYPE}
          kind: AWSMachineProviderConfig
          placement:
            availabilityZone: ${LZ_ZONE_NAME}
            region: ${CLUSTER_REGION}
          securityGroups:
            - filters:
                - name: tag:Name
                  values:
                    - ${CLUSTER_ID}-worker-sg
          subnet:
            id: ${SUBNET_ID}
          publicIp: true
          tags:
            - name: kubernetes.io/cluster/${CLUSTER_ID}
              value: owned
          userDataSecret:
            name: worker-user-data
EOF

```

### Additional resources

- [Changing the MTU for the cluster network](#)

- [Enabling IPsec encryption](#)

## 14.11. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.



### NOTE

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

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.

### Next steps

- [Creating user workloads in AWS Local Zones](#)

## 14.12. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 14.13. 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
```

## 14.14. 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.

## 14.15. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 14.16. 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](#).

## CHAPTER 15. INSTALLING A CLUSTER ON AWS IN A RESTRICTED NETWORK WITH USER-PROVISIONED INFRASTRUCTURE

In OpenShift Container Platform version 4.12, 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.

### 15.1. 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-lived 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

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 IAM credentials](#).

## 15.2. ABOUT INSTALLATIONS IN RESTRICTED NETWORKS

In OpenShift Container Platform 4.12, 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.

### 15.2.1. 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.

## 15.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#) 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.

## 15.4. 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.

### 15.4.1. 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.

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 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) 8 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#).

**15.4.2. Minimum resource requirements for cluster installation**

Each cluster machine must meet the following minimum requirements:

**Table 15.2. Minimum resource requirements**

Machine	Operating System	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.

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](#)

### 15.4.3. 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 "Minimum resource requirements for cluster installation".

#### Example 15.1. Machine types based on 64-bit x86 architecture

- **c4.\***
- **c5.\***
- **c5a.\***
- **i3.\***
- **m4.\***
- **m5.\***
- **m5a.\***
- **m6a.\***
- **m6i.\***
- **r4.\***
- **r5.\***
- **r5a.\***
- **r6i.\***
- **t3.\***
- **t3a.\***

### 15.4.4. Tested instance types for AWS on 64-bit ARM infrastructures

The following Amazon Web Services (AWS) ARM64 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 15.2. Machine types based on 64-bit ARM architecture

- **c6g.\***
- **c7g.\***
- **m6g.\***
- **m7g.\***
- **r8g.\***

## 15.4.5. 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.

## 15.5. 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.

### 15.5.1. 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"> <li>● <b>AWS::EC2::VPC</b></li> <li>● <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::SubnetNetworkAclAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkAclEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> <tr> <td><b>443</b></td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td><b>22</b></td> <td>Inbound SSH traffic</td> </tr> <tr> <td><b>1024 - 65535</b></td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td><b>0 - 65535</b></td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic	<b>443</b>	Inbound HTTPS traffic	<b>22</b>	Inbound SSH traffic	<b>1024 - 65535</b>	Inbound ephemeral traffic	<b>0 - 65535</b>	Outbound ephemeral traffic
Port	Reason													
<b>80</b>	Inbound HTTP traffic													
<b>443</b>	Inbound HTTPS traffic													
<b>22</b>	Inbound SSH traffic													
<b>1024 - 65535</b>	Inbound ephemeral traffic													
<b>0 - 65535</b>	Outbound ephemeral traffic													

Component	AWS type	Description
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	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	<b>AWS::Route53::HostedZone</b>	The hosted zone for your internal DNS.
Public load balancer	<b>AWS::ElasticLoadBalancingV2::LoadBalancer</b>	The load balancer for your public subnets.
External API server record	<b>AWS::Route53::RecordSetGroup</b>	Alias records for the external API server.
External listener	<b>AWS::ElasticLoadBalancingV2::Listener</b>	A listener on port 6443 for the external load balancer.
External target group	<b>AWS::ElasticLoadBalancingV2::TargetGroup</b>	The target group for the external load balancer.



Component	AWS type	Description
Private load balancer	<b>AWS::ElasticLoadBalancingV2::LoadBalancer</b>	The load balancer for your private subnets.
Internal API server record	<b>AWS::Route53::RecordSetGroup</b>	Alias records for the internal API server.
Internal listener	<b>AWS::ElasticLoadBalancingV2::Listener</b>	A listener on port 22623 for the internal load balancer.
Internal target group	<b>AWS::ElasticLoadBalancingV2::TargetGroup</b>	The target group for the internal load balancer.
Internal listener	<b>AWS::ElasticLoadBalancingV2::Listener</b>	A listener on port 6443 for the internal load balancer.
Internal target group	<b>AWS::ElasticLoadBalancingV2::TargetGroup</b>	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
<b>MasterSecurityGroup</b>	<b>AWS::EC2::SecurityGroup</b>	<b>icmp</b>	<b>0</b>
		<b>tcp</b>	<b>22</b>
		<b>tcp</b>	<b>6443</b>
		<b>tcp</b>	<b>22623</b>
<b>WorkerSecurityGroup</b>	<b>AWS::EC2::SecurityGroup</b>	<b>icmp</b>	<b>0</b>
		<b>tcp</b>	<b>22</b>

Group	Type	IP Protocol	Port range
<b>BootstrapSecurityGroup</b>	<b>AWS::EC2::SecurityGroup</b>	<b>tcp</b>	<b>22</b>
		<b>tcp</b>	<b>19531</b>

## 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
<b>MasterIngressEtcd</b>	etcd	<b>tcp</b>	<b>2379- 2380</b>
<b>MasterIngressVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>
<b>MasterIngressWorkerVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>
<b>MasterIngressInternal</b>	Internal cluster communication and Kubernetes proxy metrics	<b>tcp</b>	<b>9000 - 9999</b>
<b>MasterIngressWorkerInternal</b>	Internal cluster communication	<b>tcp</b>	<b>9000 - 9999</b>
<b>MasterIngressKube</b>	Kubernetes kubelet, scheduler and controller manager	<b>tcp</b>	<b>10250 - 10259</b>
<b>MasterIngressWorkerKube</b>	Kubernetes kubelet, scheduler and controller manager	<b>tcp</b>	<b>10250 - 10259</b>
<b>MasterIngressIngressServices</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>MasterIngressWorkerIngressServices</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>MasterIngressGeneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>
<b>MasterIngressWorkerGeneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>

Ingress group	Description	IP protocol	Port range
<b>MasterIngressIpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>MasterIngressWorkerIpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>MasterIngressIpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>MasterIngressWorkerIpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>MasterIngressIpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>
<b>MasterIngressWorkerIpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>
<b>MasterIngressInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>MasterIngressWorkerInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>MasterIngressIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>
<b>MasterIngressWorkerIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>

### 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
<b>WorkerIngressVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>

Ingress group	Description	IP protocol	Port range
<b>WorkerIngress WorkerVxlan</b>	Vxlan packets	<b>udp</b>	<b>4789</b>
<b>WorkerIngress Internal</b>	Internal cluster communication	<b>tcp</b>	<b>9000 - 9999</b>
<b>WorkerIngress WorkerInternal</b>	Internal cluster communication	<b>tcp</b>	<b>9000 - 9999</b>
<b>WorkerIngress Kube</b>	Kubernetes kubelet, scheduler, and controller manager	<b>tcp</b>	<b>10250</b>
<b>WorkerIngress WorkerKube</b>	Kubernetes kubelet, scheduler, and controller manager	<b>tcp</b>	<b>10250</b>
<b>WorkerIngress IngressServices</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>WorkerIngress WorkerIngress Services</b>	Kubernetes Ingress services	<b>tcp</b>	<b>30000 - 32767</b>
<b>WorkerIngress Geneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>
<b>WorkerIngress MasterGeneve</b>	Geneve packets	<b>udp</b>	<b>6081</b>
<b>WorkerIngress IpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>WorkerIngress MasterIpsecIke</b>	IPsec IKE packets	<b>udp</b>	<b>500</b>
<b>WorkerIngress IpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>WorkerIngress MasterIpsecNat</b>	IPsec NAT-T packets	<b>udp</b>	<b>4500</b>
<b>WorkerIngress IpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>

Ingress group	Description	IP protocol	Port range
<b>WorkerIngressMasterIpsecEsp</b>	IPsec ESP packets	<b>50</b>	<b>All</b>
<b>WorkerIngressInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>WorkerIngressMasterInternalUDP</b>	Internal cluster communication	<b>udp</b>	<b>9000 - 9999</b>
<b>WorkerIngressIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>
<b>WorkerIngressMasterIngressServicesUDP</b>	Kubernetes Ingress services	<b>udp</b>	<b>30000 - 32767</b>

### 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	<b>Allow</b>	<b>ec2:*</b>	<b>*</b>
	<b>Allow</b>	<b>elasticloadbalancing:*</b>	<b>*</b>
	<b>Allow</b>	<b>iam:PassRole</b>	<b>*</b>
	<b>Allow</b>	<b>s3:GetObject</b>	<b>*</b>
Worker	<b>Allow</b>	<b>ec2:Describe*</b>	<b>*</b>
Bootstrap	<b>Allow</b>	<b>ec2:Describe*</b>	<b>*</b>
	<b>Allow</b>	<b>ec2:AttachVolume</b>	<b>*</b>
	<b>Allow</b>	<b>ec2:DetachVolume</b>	<b>*</b>

### 15.5.2. 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.

### 15.5.3. 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 15.3. Required EC2 permissions for installation

- **ec2:AuthorizeSecurityGroupEgress**
- **ec2:AuthorizeSecurityGroupIngress**
- **ec2:CopyImage**
- **ec2:CreateNetworkInterface**
- **ec2:AttachNetworkInterface**
- **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:DescribeRegions**
- **ec2:DescribeRouteTables**
- **ec2:DescribeSecurityGroups**
- **ec2:DescribeSubnets**
- **ec2:DescribeTags**
- **ec2:DescribeVolumes**
- **ec2:DescribeVpcAttribute**
- **ec2:DescribeVpcClassicLink**
- **ec2:DescribeVpcClassicLinkDnsSupport**
- **ec2:DescribeVpcEndpoints**
- **ec2:DescribeVpcs**
- **ec2:GetEbsDefaultKmsKeyId**
- **ec2:ModifyInstanceAttribute**
- **ec2:ModifyNetworkInterfaceAttribute**
- **ec2:RevokeSecurityGroupEgress**
- **ec2:RevokeSecurityGroupIngress**
- **ec2:RunInstances**

- **ec2:TerminateInstances**

#### Example 15.4. 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 VPC, your account does not require these permissions for creating network resources.

#### Example 15.5. Required Elastic Load Balancing permissions (ELB) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:ApplySecurityGroupsToLoadBalancer**
- **elasticloadbalancing:AttachLoadBalancerToSubnets**
- **elasticloadbalancing:ConfigureHealthCheck**
- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateLoadBalancerListeners**
- **elasticloadbalancing>DeleteLoadBalancer**



- **elasticloadbalancing:DeregisterInstancesFromLoadBalancer**
- **elasticloadbalancing:DescribeInstanceHealth**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTags**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:RegisterInstancesWithLoadBalancer**
- **elasticloadbalancing:SetLoadBalancerPoliciesOfListener**


Example 15.6. Required Elastic Load Balancing permissions (ELBv2) for installation

- **elasticloadbalancing:AddTags**
- **elasticloadbalancing:CreateListener**
- **elasticloadbalancing:CreateLoadBalancer**
- **elasticloadbalancing:CreateTargetGroup**
- **elasticloadbalancing>DeleteLoadBalancer**
- **elasticloadbalancing:DeregisterTargets**
- **elasticloadbalancing:DescribeListeners**
- **elasticloadbalancing:DescribeLoadBalancerAttributes**
- **elasticloadbalancing:DescribeLoadBalancers**
- **elasticloadbalancing:DescribeTargetGroupAttributes**
- **elasticloadbalancing:DescribeTargetHealth**
- **elasticloadbalancing:ModifyLoadBalancerAttributes**
- **elasticloadbalancing:ModifyTargetGroup**
- **elasticloadbalancing:ModifyTargetGroupAttributes**
- **elasticloadbalancing:RegisterTargets**

Example 15.7. 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:TagRole**



#### NOTE

If you have not created a load balancer in your AWS account, the IAM user also requires the **iam:CreateServiceLinkedRole** permission.

#### Example 15.8. 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 15.9. Required S3 permissions for installation

- **s3:CreateBucket**
- **s3>DeleteBucket**
- **s3:GetAccelerateConfiguration**
- **s3:GetBucketAcl**
- **s3:GetBucketCors**
- **s3:GetBucketLocation**
- **s3:GetBucketLogging**
- **s3:GetBucketPolicy**
- **s3:GetBucketObjectLockConfiguration**
- **s3:GetBucketReplication**
- **s3:GetBucketRequestPayment**
- **s3:GetBucketTagging**
- **s3:GetBucketVersioning**
- **s3:GetBucketWebsite**
- **s3:GetEncryptionConfiguration**
- **s3:GetLifecycleConfiguration**
- **s3:GetReplicationConfiguration**
- **s3:ListBucket**
- **s3:PutBucketAcl**
- **s3:PutBucketTagging**
- **s3:PutEncryptionConfiguration**

Example 15.10. S3 permissions that cluster Operators require

- **s3>DeleteObject**
- **s3:GetObject**
- **s3:GetObjectAcl**
- **s3:GetObjectTagging**

- **s3:GetObjectVersion**
- **s3:PutObject**
- **s3:PutObjectAcl**
- **s3:PutObjectTagging**

Example 15.11. Required permissions to delete base cluster resources

- **autoscaling:DescribeAutoScalingGroups**
- **ec2:DeletePlacementGroup**
- **ec2:DeleteNetworkInterface**
- **ec2:DeleteVolume**
- **elasticloadbalancing:DeleteTargetGroup**
- **elasticloadbalancing:DescribeTargetGroups**
- **iam:DeleteAccessKey**
- **iam:DeleteUser**
- **iam>ListAttachedRolePolicies**
- **iam>ListInstanceProfiles**
- **iam>ListRolePolicies**
- **iam>ListUserPolicies**
- **s3>DeleteObject**
- **s3>ListBucketVersions**
- **tag:GetResources**

Example 15.12. 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 15.13. Required permissions to delete a cluster with shared instance roles

- **iam:UntagRole**

#### Example 15.14. Additional IAM and S3 permissions that are required to create manifests

- **iam>DeleteAccessKey**
- **iam>DeleteUser**
- **iam>DeleteUserPolicy**
- **iam:GetUserPolicy**
- **iam>ListAccessKeys**
- **iam:PutUserPolicy**
- **iam:TagUser**
- **s3:PutBucketPublicAccessBlock**
- **s3:GetBucketPublicAccessBlock**
- **s3:PutLifecycleConfiguration**
- **s3>ListBucket**
- **s3>ListBucketMultipartUploads**
- **s3:AbortMultipartUpload**



#### 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 15.15. Optional permissions for instance and quota checks for installation**

- **ec2:DescribeInstanceTypeOfferings**
- **servicequotas:ListAWSDefaultServiceQuotas**

## 15.6. 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 FIPS validated or Modules In Process cryptographic libraries on 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.

## 15.7. 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.

### 15.7.1. 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: 4.12.0
metadata:
  labels:
    machineconfiguration.openshift.io/role: worker
  name: 98-var-partition
storage:
  disks:
    - device: /dev/<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 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.

### 15.7.2. 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 a region with an accompanying Red Hat Enterprise Linux CoreOS (RHCOS) AMI published by Red Hat. If you are deploying to a 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 the 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.

## 15.7.3. 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.

### 15.7.4. 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.

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
```

Because you create and manage the worker machines yourself, you do not need to initialize these machines.

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: 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.

6. Optional: If you manually created a cloud identity and access management (IAM) role, locate any **CredentialsRequest** objects with the **TechPreviewNoUpgrade** annotation in the release image by running the following command:

```
$ oc adm release extract quay.io/openshift-release-dev/ocp-release:4.y.z-x86_64 --
credentials-requests --cloud=<platform_name>
```

### Example output

```
0000_30_capi-operator_00_credentials-request.yaml: release.openshift.io/feature-set:
TechPreviewNoUpgrade
```



### IMPORTANT

The release image includes **CredentialsRequest** objects for Technology Preview features that are enabled by the **TechPreviewNoUpgrade** feature set. You can identify these objects by their use of the **release.openshift.io/feature-set: TechPreviewNoUpgrade** annotation.

- If you are not using any of these features, do not create secrets for these objects. Creating secrets for Technology Preview features that you are not using can cause the installation to fail.
- If you are using any of these features, you must create secrets for the corresponding objects.

- a. Delete all **CredentialsRequest** objects that have the **TechPreviewNoUpgrade** annotation.
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> 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
```





### Additional resources

- [Manually creating IAM](#)

## 15.8. 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.

## 15.9. 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:

<b>VpcId</b>	The ID of your VPC.
<b>PublicSubnetIds</b>	The IDs of the new public subnets.
<b>PrivateSubnetIds</b>	The IDs of the new private subnets.

### 15.9.1. 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 15.16. 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:
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::VPCGatewayAttachment"
Properties:
VpcId: !Ref VPC
InternetGatewayId: !Ref InternetGateway
PublicRouteTable:
Type: "AWS::EC2::RouteTable"
Properties:
VpcId: !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"]]
]
```

## 15.10. 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 has 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
```

- 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:

<b>PrivateHostedZoneId</b>	Hosted zone ID for the private DNS.
<b>ExternalApiLoadBalancerName</b>	Full name of the external API load balancer.
<b>InternalApiLoadBalancerName</b>	Full name of the internal API load balancer.
<b>ApiServerDnsName</b>	Full hostname of the API server.
<b>RegisterNlbIpTargetsLambda</b>	Lambda ARN useful to help register/deregister IP targets for these load balancers.
<b>ExternalApiTargetGroupArn</b>	ARN of external API target group.
<b>InternalApiTargetGroupArn</b>	ARN of internal API target group.
<b>InternalServiceTargetGroupArn</b>	ARN of internal service target group.

### 15.10.1. 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 15.17. 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 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

```

#### RegisterNlbPTargets:

```

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]]`

RegisterNlbTargetsLambda:

Description: Lambda ARN useful to help register or deregister IP targets for these load balancers.

Value: `!GetAtt RegisterNlbTargets.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.

## 15.11. 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:

<b>MasterSecurityGroupID</b>	Master Security Group ID
------------------------------	--------------------------

<b>WorkerSecurityGroupID</b>	Worker Security Group ID
<b>MasterInstanceProfile</b>	Master IAM Instance Profile
<b>WorkerInstanceProfile</b>	Worker IAM Instance Profile

### 15.11.1. 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 15.18. 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}(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/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
        FromPort: 6443
        ToPort: 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&gt;DeleteSecurityGroup"

- "ec2&gt;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&gt;DeleteListener"

- "elasticloadbalancing&gt;DeleteLoadBalancer"

- "elasticloadbalancing&gt;DeleteLoadBalancerListeners"

- "elasticloadbalancing&gt;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

```

## 15.12. 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.

### 15.13. 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 15.3. x86\_64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-073850a7021953a5c
ap-east-1	ami-0f8800a05c09be42d
ap-northeast-1	ami-0a226dbcc9a561c40
ap-northeast-2	ami-041ae0537e2eddec1
ap-northeast-3	ami-0bb8d9b69dc5b7670
ap-south-1	ami-0e9c18058fc5f94fd
ap-southeast-1	ami-03022d358ba2168be
ap-southeast-2	ami-09ffdc5be9b973be0
ap-southeast-3	ami-0facf1a0edeb20314
ca-central-1	ami-028cea206c2d03317
eu-central-1	ami-002eb441f329ccb0f
eu-north-1	ami-0b1a1fb68b3b9fee7
eu-south-1	ami-0bd0fd41a1d3f799a
eu-west-1	ami-04504e8799057980c
eu-west-2	ami-0cc9297ddb3bce971

AWS zone	AWS AMI
eu-west-3	ami-06f98f607a50937c6
me-south-1	ami-0fe39da7871a5b2a5
sa-east-1	ami-08265cc3226697767
us-east-1	ami-0fe05b1aa8dacfa90
us-east-2	ami-0ff64f495c7e977cf
us-gov-east-1	ami-0c99658076c41872a
us-gov-west-1	ami-0ca4acd5b8ba1cb1d
us-west-1	ami-01dc5d8e6bb6f23f4
us-west-2	ami-0404a109adfd00019

Table 15.4. aarch64 RHCOS AMIs

AWS zone	AWS AMI
af-south-1	ami-0574bcc5f80b0ad9a
ap-east-1	ami-0a65e79822ae2d235
ap-northeast-1	ami-0f7ef19d48e22353b
ap-northeast-2	ami-051dc6de359975e3c
ap-northeast-3	ami-0fd0b4222595650ac
ap-south-1	ami-05f9d14fe4a90ed6f
ap-southeast-1	ami-0afdb9133d22fba5f
ap-southeast-2	ami-0ef979abe82d07d44
ap-southeast-3	ami-025f9103ac4310e7f
ca-central-1	ami-0588cdf59e5c14847
eu-central-1	ami-0ef24c0e18f93fa42

AWS zone	AWS AMI
eu-north-1	ami-0439e2a3bf315df1a
eu-south-1	ami-0714e7c2e0106cdd3
eu-west-1	ami-0b960e76764ccd0c3
eu-west-2	ami-02621f50de62b3b89
eu-west-3	ami-0933ce7f5e2bfb50e
me-south-1	ami-074bde61a2ab740ee
sa-east-1	ami-03b4f97cfc8033ae0
us-east-1	ami-02a574449d4f4d280
us-east-2	ami-020e5600ef28c60ae
us-gov-east-1	ami-069f60e1dcf766d24
us-gov-west-1	ami-0db3cda4dbaccda02
us-west-1	ami-0c90cabeb5dee3178
us-west-2	ami-0f96437a23aeae53f

## 15.14. 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": "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>**.
- 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** 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:

<b>Bootstrap InstanceId</b>	The bootstrap Instance ID.
<b>Bootstrap PublicIp</b>	The bootstrap node public IP address.
<b>Bootstrap PrivateIp</b>	The bootstrap node private IP address.

### 15.14.1. 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 15.19. 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]{1-9}[0-9]{1-9}|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){3}([0-9]{1-9}|1[0-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.)\.[0-9]{1-9}|1[0-9]|2[0-9]|3[0-2]))$
    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
- RegisterNLbIpTargetsLambdaArn
- 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 "MasterSecurityGroupId"

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.

## 15.15. 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": "MasterSecurityGroupID", 17
    "ParameterValue": "sg-<random_string>" 18
  },
  {
```

```

    "ParameterKey": "IgnitionLocation", 19
    "ParameterValue": "https://api-int.<cluster_name>.<domain_name>:22623/config/master"
  },
  {
    "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](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
  - 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>
```

### 15.15.1. 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 15.20. 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:
      - AssociatePublicIp: "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 RegisterNlbIpTargetsLambdaArn
  TargetArn: !Ref ExternalApiTargetGroupArn
  TargetIp: !GetAtt Master2.PrivateIp

```

**RegisterMaster2InternalApiTarget:**

```

Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
  ServiceToken: !Ref RegisterNlbIpTargetsLambdaArn
  TargetArn: !Ref InternalApiTargetGroupArn
  TargetIp: !GetAtt Master2.PrivateIp

```

**RegisterMaster2InternalServiceTarget:**

```

Condition: DoRegistration
Type: Custom::NLBRegister
Properties:
  ServiceToken: !Ref RegisterNlbIpTargetsLambdaArn
  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]
]

```

## 15.16. 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": "RhcosAmi", 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"
  },
  {
    "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](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> 1
  --template-body file://<template>.yaml \ 2
  --parameters file://<parameters>.json 3
```

- 1** **<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.
- 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-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.

**15.16.1. 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 15.21. 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 master nodes into.
    Type: AWS::EC2::Subnet::Id
  WorkerSecurityGroupId:
    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/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 master 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

## 15.17. 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 v1.25.0 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.

## 15.18. 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
```

## 15.19. 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

- Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

#### Example output

```
NAME      STATUS  ROLES  AGE  VERSION
master-0  Ready   master 63m  v1.25.0
master-1  Ready   master 63m  v1.25.0
master-2  Ready   master 64m  v1.25.0
```

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}}' | 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  v1.25.0
master-1  Ready   master 73m  v1.25.0
master-2  Ready   master 74m  v1.25.0
worker-0  Ready   worker 11m  v1.25.0
worker-1  Ready   worker 11m  v1.25.0
```



### 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](#).

## 15.20. 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	4.12.0	True	False	False	19m
baremetal	4.12.0	True	False	False	37m
cloud-credential	4.12.0	True	False	False	40m
cluster-autoscaler	4.12.0	True	False	False	37m
config-operator	4.12.0	True	False	False	38m
console	4.12.0	True	False	False	26m
csi-snapshot-controller	4.12.0	True	False	False	37m
dns	4.12.0	True	False	False	37m
etcd	4.12.0	True	False	False	36m
image-registry	4.12.0	True	False	False	31m
ingress	4.12.0	True	False	False	30m
insights	4.12.0	True	False	False	31m
kube-apiserver	4.12.0	True	False	False	26m
kube-controller-manager	4.12.0	True	False	False	36m
kube-scheduler	4.12.0	True	False	False	36m
kube-storage-version-migrator	4.12.0	True	False	False	37m
machine-api	4.12.0	True	False	False	29m
machine-approver	4.12.0	True	False	False	37m
machine-config	4.12.0	True	False	False	36m
marketplace	4.12.0	True	False	False	37m
monitoring	4.12.0	True	False	False	29m
network	4.12.0	True	False	False	38m
node-tuning	4.12.0	True	False	False	37m
openshift-apiserver	4.12.0	True	False	False	32m
openshift-controller-manager	4.12.0	True	False	False	30m
openshift-samples	4.12.0	True	False	False	32m
operator-lifecycle-manager	4.12.0	True	False	False	37m
operator-lifecycle-manager-catalog	4.12.0	True	False	False	37m
operator-lifecycle-manager-packageserver	4.12.0	True	False	False	32m
service-ca	4.12.0	True	False	False	38m
storage	4.12.0	True	False	False	37m

2. Configure the Operators that are not available.

### 15.20.1. 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.

### 15.20.2. 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.

#### 15.20.2.1. 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.

**15.20.2.2. 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.

**15.21. 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](#).

## 15.22. 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

1. 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>
```

2. 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
```

3. Locate the hosted zone ID for the load balancer:

```
$ aws elb describe-load-balancers | jq -r '.LoadBalancerDescriptions[] | select(.DNSName ==
"<external_ip>").CanonicalHostedZoneNameID' 1
```

- 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.

4. 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 '{
> "Changes": [
> {
>   "Action": "CREATE",
>   "ResourceRecordSet": {
>     "Name": "\\052.apps.<cluster_domain>",
>     "Type": "A",
>     "AliasTarget": {
>       "HostedZoneId": "<hosted_zone_id>",
>       "DNSName": "<external_ip>.",
>       "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.

## 15.23. 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-bootstrap** 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.

## 15.24. 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.

## 15.25. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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

## 15.26. ADDITIONAL RESOURCES

- See [Working with stacks](#) in the AWS documentation for more information about AWS CloudFormation stacks.

## 15.27. 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](#).

## CHAPTER 16. INSTALLING A CLUSTER ON AWS WITH REMOTE WORKERS ON AWS OUTPOSTS

In OpenShift Container Platform version 4.12, you can install a cluster on Amazon Web Services (AWS) with remote workers running in AWS Outposts. This can be achieved by customizing the default AWS installation and performing some manual steps.



### IMPORTANT

Installing a cluster on AWS with remote workers on AWS Outposts 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](#).

For more info about AWS Outposts see [AWS Outposts Documentation](#).



### IMPORTANT

In order to install a cluster with remote workers in AWS Outposts, all worker instances must be located within the same Outpost instance and cannot be located in an AWS region. It is not possible for the cluster to have instances in both AWS Outposts and AWS region. In addition, it also follows that control plane nodes mustn't be schedulable.

## 16.1. 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.
- You are familiar with the instance types are supported in the AWS Outpost instance you use. This can be validated with [get-outpost-instance-types AWS CLI command](#)
- You are familiar with the AWS Outpost instance details, such as OutpostArn and AvailabilityZone. This can be validated with [list-outposts AWS CLI command](#)



### IMPORTANT

Since the cluster uses the provided AWS credentials to create AWS resources for its entire life cycle, the credentials must be key-based and long-lived. So, If you have an AWS profile stored on your computer, it must not use a temporary session token, generated while using a multi-factor authentication device. For more information about generating the appropriate keys, see [Managing Access Keys for IAM Users](#) in the AWS documentation. You may supply the keys when you run the installation program.

- You have access to an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). See the section "About using a custom VPC" for more information.
- If a firewall is used, it was [configured 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 IAM credentials](#).

## 16.2. ABOUT USING A CUSTOM VPC

OpenShift Container Platform 4.12 installer cannot automatically deploy AWS Subnets on AWS Outposts, so you will need to manually configure the VPC. Therefore, you have to deploy the cluster into existing subnets in an existing Amazon Virtual Private Cloud (VPC) in Amazon Web Services (AWS). In addition, 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.

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.

### 16.2.1. 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:



#### NOTE

To allow the creation of OpenShift Container Platform with remote workers in the AWS Outposts, you must create at least one private subnet in the AWS Outpost instance for the workload instances creation and one private subnet in an AWS region for the control plane instances creation. If you specify more than one private subnet in the region, the control plane instances will be distributed across these subnets. You will also need to create a public subnet in each of the availability zones used for private subnets, including the Outpost private subnet, as Network Load Balancers will be created in the AWS region for the API server and Ingress network as part of the cluster installation. It is possible to create an AWS region private subnet in the same Availability zone as an Outpost private subnet.

- Create a public and private subnet in the AWS Region for each availability zone that your control plane uses. Each availability zone can contain no more than one public and one private subnet in the AWS region. For an example of this type of configuration, see [VPC with public and private subnets \(NAT\)](#) in the AWS documentation.

To create a private subnet in the AWS Outposts, you need to first ensure that the Outpost instance is located in the desired availability zone. Then, you can create the private subnet within that availability zone within the Outpost instance, by adding the Outpost ARN. Make sure there is another public subnet in the AWS Region created in the same availability zone.

Record each subnet ID. Completing the installation requires that you enter all the subnets IDs, created in the AWS Region, in the **platform** section of the **install-config.yaml** file and changing the workers **machineset** to use the private subnet ID created in the Outpost. See [Finding a subnet ID](#) in the AWS documentation.



#### IMPORTANT

In case you need to create a public subnet in the AWS Outposts, verify that this subnet is not used for the Network or Classic LoadBalancer, otherwise the LoadBalancer creation fails. To achieve that, the **kubernetes.io/cluster/.\*-outposts: owned** special tag must be included in the subnet.

- 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.

**NOTE**

To access your local cluster over your local network, the VPC must be associated with your Outpost's local gateway route table. For more information, see [VPC associations](#) in the AWS Outposts User Guide.

- 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.
- 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** field in the **install-config.yaml** file.

**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"> <li>● <b>AWS::EC2::VPC</b></li> <li>● <b>AWS::EC2::VPCEndpoint</b></li> </ul>	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"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::SubnetNetworkAclAssociation</b></li> </ul>	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"> <li>● <b>AWS::EC2::InternetGateway</b></li> <li>● <b>AWS::EC2::VPCGatewayAttachment</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::Route</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> <li>● <b>AWS::EC2::NatGateway</b></li> <li>● <b>AWS::EC2::EIP</b></li> </ul>	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"> <li>● <b>AWS::EC2::NetworkAcl</b></li> <li>● <b>AWS::EC2::NetworkAclEntry</b></li> </ul>	<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><b>80</b></td> <td>Inbound HTTP traffic</td> </tr> <tr> <td><b>443</b></td> <td>Inbound HTTPS traffic</td> </tr> <tr> <td><b>22</b></td> <td>Inbound SSH traffic</td> </tr> <tr> <td><b>1024 - 65535</b></td> <td>Inbound ephemeral traffic</td> </tr> <tr> <td><b>0 - 65535</b></td> <td>Outbound ephemeral traffic</td> </tr> </tbody> </table>	Port	Reason	<b>80</b>	Inbound HTTP traffic	<b>443</b>	Inbound HTTPS traffic	<b>22</b>	Inbound SSH traffic	<b>1024 - 65535</b>	Inbound ephemeral traffic	<b>0 - 65535</b>	Outbound ephemeral traffic
Port	Reason													
<b>80</b>	Inbound HTTP traffic													
<b>443</b>	Inbound HTTPS traffic													
<b>22</b>	Inbound SSH traffic													
<b>1024 - 65535</b>	Inbound ephemeral traffic													
<b>0 - 65535</b>	Outbound ephemeral traffic													



Component	AWS type	Description
Private subnets	<ul style="list-style-type: none"> <li>● <b>AWS::EC2::Subnet</b></li> <li>● <b>AWS::EC2::RouteTable</b></li> <li>● <b>AWS::EC2::SubnetRouteTableAssociation</b></li> </ul>	Your VPC can have private subnets. The provided CloudFormation templates can create private subnets for between 1 and 3 availability zones. To enable remote workers running in the Outpost, the VPC must include a private subnet located within the Outpost instance, in addition to the private subnets located within the corresponding AWS region. If you use private subnets, you must provide appropriate routes and tables for them.

### 16.2.2. 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 exactly one public and one private subnet in the AWS region (not created in the Outpost instance). The availability zone in which the Outpost instance is installed should include one additional private subnet in the Outpost instance.
- 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.

### 16.2.3. 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.

## 16.2.4. 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.

## 16.3. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, you require access to the internet to install your cluster.

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) 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.

## 16.4. 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 FIPS validated or Modules In Process cryptographic libraries on 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.

## 16.5. 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 the 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.

## 16.6. MINIMUM RESOURCE REQUIREMENTS FOR CLUSTER INSTALLATION

Each cluster machine must meet the following minimum requirements:

**Table 16.1. 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.

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](#)

## 16.7. IDENTIFYING YOUR AWS OUTPOSTS INSTANCE TYPES

AWS Outposts rack catalog includes options supporting the latest generation Intel powered EC2 instance types with or without local instance storage. Identify which instance types are configured in your AWS Outpost instance. As part of the installation process, you must update the **install-config.yaml** file with the instance type that the installation program will use to deploy worker nodes.

### Procedure

Use the AWS CLI to get the list of supported instance types by running the following command:

```
$ aws outposts get-outpost-instance-types --outpost-id <outpost_id> 1
```

- 1 For **<outpost\_id>**, specify the Outpost ID, used in the AWS account for the worker instances



### IMPORTANT

When you purchase capacity for your AWS Outpost instance, you specify an EC2 capacity layout that each server provides. Each server supports a single family of instance types. A layout can offer a single instance type or multiple instance types. Dedicated Hosts allows you to alter whatever you chose for that initial layout. If you allocate a host to support a single instance type for the entire capacity, you can only start a single instance type from that host.

Supported instance types in AWS Outposts might be changed. For more information, you can check the [Compute and Storage](#) page in AWS Outposts documents.

## 16.8. CREATING THE INSTALLATION CONFIGURATION FILE

You can customize the OpenShift Container Platform cluster you install on Amazon Web Services (AWS).

### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Obtain service principal permissions at the subscription level.

### 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.
  - vii. Paste the [pull secret from the Red Hat OpenShift Cluster Manager](#) .
2. Modify the **install-config.yaml** file. The AWS Outposts installation has the following limitations which require manual modification of the **install-config.yaml** file:
    - Unlike AWS Regions, which offer near-infinite scale, AWS Outposts are limited by their provisioned capacity, EC2 family and generations, configured instance sizes, and availability of compute capacity that is not already consumed by other workloads. Therefore, when creating new OpenShift Container Platform cluster, you need to provide the supported instance type in the **compute.platform.aws.type** section in the configuration file.
    - When deploying OpenShift Container Platform cluster with remote workers running in AWS Outposts, only one Availability Zone can be used for the compute instances - the Availability Zone in which the Outpost instance was created in. Therefore, when creating

new OpenShift Container Platform cluster, it is recommended to provide the relevant Availability Zone in the **compute.platform.aws.zones** section in the configuration file, in order to limit the compute instances to this Availability Zone.

- Amazon Elastic Block Store (EBS) gp3 volumes aren't supported by the AWS Outposts service. This volume type is the default type used by the OpenShift Container Platform cluster. Therefore, when creating new OpenShift Container Platform cluster, you must change the volume type in the **compute.platform.aws.rootVolume.type** section to gp2. You will find more information about how to change these values below.
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.

## 16.8.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.



### NOTE

After installation, you cannot modify these parameters in the **install-config.yaml** file.

### 16.8.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 16.2. Required parameters

Parameter	Description	Values
<b>apiVersion</b>	The API version for the <b>install-config.yaml</b> content. The current version is <b>v1</b> . The installation program may also support older API versions.	String



Parameter	Description	Values
<b>baseDomain</b>	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 <b>baseDomain</b> and <b>metadata.name</b> parameter values that uses the <b>&lt;metadata.name&gt;</b> . <b>&lt;baseDomain&gt;</b> format.	A fully-qualified domain or subdomain name, such as <b>example.com</b> .
<b>metadata</b>	Kubernetes resource <b>ObjectMeta</b> , from which only the <b>name</b> parameter is consumed.	Object
<b>metadata.name</b>	The name of the cluster. DNS records for the cluster are all subdomains of <b>{{.metadata.name}}</b> . <b>{{.baseDomain}}</b> .	String of lowercase letters, hyphens (-), and periods (.), such as <b>dev</b> .
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object

Parameter	Description	Values
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> 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>

### 16.8.1.2. 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.


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 16.3. Network parameters

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	<p>Object</p>  <p><b>NOTE</b></p> <p>You cannot modify parameters specified by the <b>networking</b> object after installation.</p>

Parameter	Description	Values
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . 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 <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>
<b>networking.machineNetwork</b>	The IP address blocks for machines.  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   machineNetwork:     - cidr: 10.0.0.0/16</pre>


Parameter	Description	Values
<b>networking.machineNetwork.cidr</b>	Required if you use <b>networking.machineNetwork</b> . An IP address block. The default value is <b>10.0.0.0/16</b> for all platforms other than libvirt. For libvirt, the default value is <b>192.168.126.0/24</b> .	An IP network block in CIDR notation. For example, <b>10.0.0.0/16</b> .  <b>NOTE</b> Set the <b>networking.machineNetwork</b> to match the CIDR that the preferred NIC resides in.


### 16.8.1.3. Optional configuration parameters



Optional installation configuration parameters are described in the following table:



Table 16.4. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	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
<b>capabilities</b>	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
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.

Parameter	Description	Values
<b>compute.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .

Parameter	Description	Values
<b>featureSet</b>	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 <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	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 <b>amd64</b> and <b>arm64</b> . 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
<b>controlPlane.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>

Parameter	Description	Values
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.
<b>credentialsMode</b>	<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> <div style="display: flex; align-items: flex-start;"> <div style="border: 1px solid black; padding: 5px; margin-right: 10px;">  </div> <div> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="border: 1px solid black; padding: 5px; margin-right: 10px;">  </div> <div> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ( <b>""</b> ).
<b>fips</b>	Enable or disable FIPS mode. The default is <b>false</b> (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.	<b>false</b> or <b>true</b>

Parameter	Description	Values
	<p style="text-align: center;"><b>IMPORTANT</b></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 <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> <p style="text-align: center;"><b>NOTE</b></p>  <p>If you are using Azure File storage, you cannot enable FIPS mode.</p>	
<p><b>imageContentSources</b></p>	<p>Sources and repositories for the release-image content.</p>	<p>Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b>, as described in the following rows of this table.</p>



Parameter	Description	Values
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String
<b>imageContentSources.mirrors</b>	Specify one or more repositories that may also contain the same images.	Array of strings
<b>platform.aws.lbType</b>	Required to set the NLB load balancer type in AWS. Valid values are <b>Classic</b> or <b>NLB</b> . If no value is specified, the installation program defaults to <b>Classic</b> . 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.	<b>Classic</b> or <b>NLB</b> . The default value is <b>Classic</b> .
<b>publish</b>	How to publish or expose the user-facing endpoints of your cluster, such as the Kubernetes API, OpenShift routes.	<b>Internal</b> or <b>External</b> . To deploy a private cluster, which cannot be accessed from the internet, set <b>publish</b> to <b>Internal</b> . The default value is <b>External</b> .
<b>sshKey</b>	<p>The SSH key to authenticate access to your cluster machines.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></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 <b>ssh-agent</b> process uses.</p> </div> </div>	For example, <b>sshKey: ssh-ed25519 AAAA...</b>

### 16.8.1.4. Optional AWS configuration parameters


Optional AWS configuration parameters are described in the following table:

Table 16.5. Optional AWS parameters

Parameter	Description	Values
<b>compute.platform.aws.amiID</b>	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.
<b>compute.platform.aws.iamRole</b>	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.
<b>compute.platform.aws.rootVolume.iops</b>	The Input/Output Operations Per Second (IOPS) that is reserved for the root volume.	Integer, for example <b>4000</b> .
<b>compute.platform.aws.rootVolume.size</b>	The size in GiB of the root volume.	Integer, for example <b>500</b> .
<b>compute.platform.aws.rootVolume.type</b>	The type of the root volume.	Valid <a href="#">AWS EBS volume type</a> , such as <b>io1</b> .
<b>compute.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> or the <a href="#">key ARN</a>
<b>compute.platform.aws.type</b>	The EC2 instance type for the compute machines.	Valid AWS instance type, such as <b>m4.2xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.

Parameter	Description	Values
<b>compute.platform.aws.zones</b>	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 <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>compute.aws.region</b>	The AWS region that the installation program creates compute resources in.	Any valid <a href="#">AWS region</a> , such as <b>us-east-1</b> . You can use the AWS CLI to access the regions available based on your selected instance type. For example: <pre>aws ec2 describe-instance-type-offerings -- filters Name=instance- type,Values=c7g.xlarge</pre> <div style="display: flex; align-items: flex-start; margin-top: 10px;">  <div> <p><b>IMPORTANT</b></p> <p>When running on ARM based AWS instances, ensure that you enter a region where AWS Graviton processors are available. See <a href="#">Global availability map</a> in the AWS documentation. Currently, AWS Graviton3 processors are only available in some regions.</p> </div> </div>
<b>controlPlane.platform.aws.amiID</b>	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.
<b>controlPlane.platform.aws.iamRole</b>	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.


Parameter	Description	Values
<b>controlPlane.platform.aws.rootVolume.kmsKeyARN</b>	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 <a href="#">key ID</a> and the <a href="#">key ARN</a>
<b>controlPlane.platform.aws.type</b>	The EC2 instance type for the control plane machines.	Valid AWS instance type, such as <b>m6i.xlarge</b> . See the <b>Supported AWS machine types</b> table that follows.
<b>controlPlane.platform.aws.zones</b>	The availability zones where the installation program creates machines for the control plane machine pool.	A list of valid AWS availability zones, such as <b>us-east-1c</b> , in a <a href="#">YAML sequence</a> .
<b>controlPlane.aws.region</b>	The AWS region that the installation program creates control plane resources in.	Valid <a href="#">AWS region</a> , such as <b>us-east-1</b> .
<b>platform.aws.amiID</b>	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.
<b>platform.aws.hostedZone</b>	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 <b>Z3URY6TWQ91KVV</b> .

Parameter	Description	Values
<b>platform.aws.serviceEndpoints.name</b>	The AWS service endpoint name. 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 <a href="#">AWS service endpoint</a> name.
<b>platform.aws.serviceEndpoints.url</b>	The AWS service endpoint URL. The URL must use the <b>https</b> protocol and the host must trust the certificate.	Valid <a href="#">AWS service endpoint</a> URL.
<b>platform.aws.userTags</b>	A map of keys and values that the installation program adds as tags to all resources that it creates.	<p>Any valid YAML map, such as key value pairs in the <b>&lt;key&gt;: &lt;value&gt;</b> format. For more information about AWS tags, see <a href="#">Tagging Your Amazon EC2 Resources</a> in the AWS documentation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></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>
<b>platform.aws.propagateUserTags</b>	A flag that directs in-cluster Operators to include the specified user tags in the tags of the AWS resources that the Operators create.	Boolean values, for example <b>true</b> or <b>false</b> .

Parameter	Description	Values
<b>platform.aws.subnets</b>	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 <b>machineNetwork[].cidr</b> ranges that you specify. For a standard cluster, specify a public and a private subnet for each availability zone. For a private cluster, specify a private subnet for each availability zone.	Valid subnet IDs.

## 16.8.2. 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: {}
  replicas: 3
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    aws:
      type: m5.large 8
      zones:
        - us-east-1a 9
      rootVolume:
        type: gp2 10
        size: 120
      replicas: 3
  metadata:
    name: test-cluster 11
networking:

```

```

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:
aws:
  region: us-west-2 13
  propagateUserTags: true 14
  userTags:
    adminContact: jdoe
    costCenter: 7536
  subnets: 15
- subnet-1
- subnet-2
- subnet-3
sshKey: ssh-ed25519 AAAA... 16
pullSecret: '{"auths": ...}' 17

```

- 1 11 13 17 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, instead of having the CCO dynamically try to determine the capabilities of the credentials. For details about CCO modes, see the *Cloud Credential Operator* entry in the *Red Hat Operators reference* content.
- 3 6 14 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 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. Use larger instance types, such as **m4.2xlarge** or **m5.2xlarge**, for your machines if you disable simultaneous multithreading.

- 8 For compute instances running in an AWS Outpost instance, specify a supported instance type in the AWS Outpost instance.
- 9 For compute instances running in AWS Outpost instance, specify the Availability Zone where the Outpost instance is located.

- 10 For compute instances running in AWS Outpost instance, specify volume type gp2, to avoid using gp3 volume type which is not supported.
- 12 The cluster network plugin to install. The supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.
- 15 If you provide your own VPC, specify subnets for each availability zone that your cluster uses.
- 16 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.

## 16.9. GENERATING MANIFEST FILES

Use the installation program to generate a set of manifest files in the assets directory. Manifest files are required to specify the AWS Outposts subnets to use for worker machines, and to specify settings required by the network provider.

If you plan to reuse the **install-config.yaml** file, create a backup file before you generate the manifest files.

### Procedure

- Optional: Create a backup copy of the **install-config.yaml** file:

```
$ cp install-config.yaml install-config.yaml.backup
```

- Generate a set of manifests in your assets directory:

```
$ 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
```

The command generates the following manifest files:

### Example output

```
$ tree
.
├── manifests
│   ├── 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
    ├── 99_cloud-creds-secret.yaml
    ├── 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
    ├── 99_openshift-cluster-api_master-user-data-secret.yaml
    ├── 99_openshift-cluster-api_worker-machineset-0.yaml
    ├── 99_openshift-cluster-api_worker-user-data-secret.yaml
    ├── 99_openshift-machineconfig_99-master-ssh.yaml
    ├── 99_openshift-machineconfig_99-worker-ssh.yaml
    ├── 99_role-cloud-creds-secret-reader.yaml
    └── openshift-install-manifests.yaml

```

### 16.9.1. Modifying manifest files



#### NOTE

The AWS Outposts environments has the following limitations which require manual modification in the manifest generated files:

- The maximum transmission unit (MTU) of a network connection is the size, in bytes, of the largest permissible packet that can be passed over the connection. The Outpost service link supports a maximum packet size of 1300 bytes. For more information about the service link, see [Outpost connectivity to AWS Regions](#)

You will find more information about how to change these values below.

- Use Outpost Subnet for workers **machineset**  
Modify the following file: <installation\_directory>/openshift/99\_openshift-cluster-api\_worker-machineset-0.yaml Find the subnet ID and replace it with the ID of the private subnet created in the Outpost. As a result, all the worker machines will be created in the Outpost.
- Specify MTU value for the Network Provider  
Outpost service links support a maximum packet size of 1300 bytes. It's required to modify the MTU of the Network Provider to follow this requirement. Create a new file under manifests directory, named cluster-network-03-config.yaml

If OpenShift SDN network provider is used, set the MTU value to 1250

```

apiVersion: operator.openshift.io/v1
kind: Network

```

```

metadata:
  name: cluster
spec:
  defaultNetwork:
    openshiftSDNConfig:
      mtu: 1250

```

If OVN-Kubernetes network provider is used, set the MTU value to 1200

```

apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  defaultNetwork:
    ovnKubernetesConfig:
      mtu: 1200

```

## 16.10. 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

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify 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**.

**NOTE**

If the cloud provider account that you configured on your host does not have sufficient permissions to deploy the cluster, the installation process stops, and the missing permissions are displayed.

- 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.

## 16.11. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY

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 4.12. 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 v4.12 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 v4.12 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 v4.12 macOS Client** entry and save the file.



#### NOTE

For macOS arm64, choose the **OpenShift v4.12 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>
```

## 16.12. 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
```

## 16.13. 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.

## 16.14. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM


In OpenShift Container Platform 4.12, 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 Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) 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 [Accessing the web console](#) for more details about accessing and understanding the OpenShift Container Platform web console.
- See [About remote health monitoring](#) for more information about the Telemetry service.

## 16.15. CLUSTER LIMITATIONS




### IMPORTANT

Network Load Balancer (NLB) and Classic Load Balancer are not supported on AWS Outposts. After the cluster is created, all the Load Balancers are created in the AWS region. In order to use Load Balancers created inside the Outpost instances, Application Load Balancer should be used. The AWS Load Balancer Operator can be used in order to achieve that goal.

If you want to use a public subnet located in the outpost instance for the ALB, you need to remove the special tag (**kubernetes.io/cluster/.\*-outposts: owned**) that was added earlier during the VPC creation. This will prevent you from creating new Services of type LoadBalancer (Network Load Balancer).

See [Understanding the AWS Load Balancer Operator](#) for more information



## IMPORTANT

### Persistent storage using AWS Elastic Block Store limitations

- AWS Outposts does not support Amazon Elastic Block Store (EBS) gp3 volumes. After installation, the cluster includes two storage classes - gp3-csi and gp2-csi, with gp3-csi being the default storage class. It is important to always use gp2-csi. You can change the default storage class using the following OpenShift CLI (oc) commands:

```
$ oc annotate --overwrite storageclass gp3-csi storageclass.kubernetes.io/is-  
default-class=false  
$ oc annotate --overwrite storageclass gp2-csi storageclass.kubernetes.io/is-  
default-class=true
```

- To create a Volume in the Outpost instance, the CSI driver determines the Outpost ARN based on the topology keys stored on the CSINode objects. To ensure that the CSI driver uses the correct topology values, it is necessary to use the **WaitForConsumer** volume binding mode and avoid setting allowed topologies on any new storage class created.

## 16.16. 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](#) .



## CHAPTER 17. UNINSTALLING A CLUSTER ON AWS

You can remove a cluster that you deployed to Amazon Web Services (AWS).

### 17.1. 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 clusters. There might be resources that the installation program did not create or that the installation program 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. On the computer that you used to install the cluster, go to the directory that contains the installation program, and 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.

### 17.2. DELETING AWS RESOURCES WITH THE CLOUD CREDENTIAL OPERATOR UTILITY

To clean up resources after uninstalling an OpenShift Container Platform cluster with the Cloud Credential Operator (CCO) in manual mode with STS, you can use the CCO utility (**ccoctl**) to remove the AWS resources that **ccoctl** created during installation.

## Prerequisites

- Extract and prepare the **ccoctl** binary.
- Install an OpenShift Container Platform cluster with the CCO in manual mode with STS.

## Procedure

- Delete the AWS resources that **ccoctl** created:

```
$ 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.

## 17.3. 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](#)