



Red Hat

OpenShift Container Platform 4.19

Installing on OCI

Installing OpenShift Container Platform on Oracle Cloud Infrastructure

OpenShift Container Platform 4.19 Installing on OCI

Installing OpenShift Container Platform on Oracle Cloud Infrastructure

Legal Notice

Copyright © Red Hat.

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document describes how to install OpenShift Container Platform on Oracle Cloud Infrastructure.

Table of Contents

CHAPTER 1. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE ASSISTED INSTALLER	4
1.1. ABOUT THE ASSISTED INSTALLER AND OCI INTEGRATION	4
1.1.1. Preinstallation considerations	4
1.1.2. Workflow	5
1.2. PREPARING THE OCI ENVIRONMENT	6
1.3. USING THE ASSISTED INSTALLER TO GENERATE AN OCI-COMPATIBLE DISCOVERY ISO IMAGE	7
1.3.1. Creating the cluster	7
1.3.2. Generating the Discovery ISO image	8
1.4. PROVISIONING OCI INFRASTRUCTURE FOR YOUR CLUSTER	8
1.5. COMPLETING THE REMAINING ASSISTED INSTALLER STEPS	10
1.5.1. Assigning node roles	10
1.5.2. Adding custom manifests	10
1.6. VERIFYING A SUCCESSFUL CLUSTER INSTALLATION ON OCI	11
1.7. ADDING HOSTS TO THE CLUSTER FOLLOWING THE INSTALLATION	12
1.8. TROUBLESHOOTING THE INSTALLATION OF A CLUSTER ON OCI	12
1.8.1. The Ingress Load Balancer in OCI is not at a healthy status	12
1.8.2. OCI create stack operation fails with an Error: 400-InvalidParameter message	13
CHAPTER 2. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE AGENT-BASED INSTALLER	14
2.1. THE AGENT-BASED INSTALLER AND OCI OVERVIEW	14
2.2. INSTALLATION PROCESS WORKFLOW	16
2.3. CREATING OCI INFRASTRUCTURE RESOURCES AND SERVICES	17
2.4. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON OCI	17
2.5. CONFIGURING YOUR FIREWALL FOR OPENSHIFT CONTAINER PLATFORM	22
2.6. RUNNING A CLUSTER ON OCI	24
2.7. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON OCI	25
2.8. ADDITIONAL RESOURCES	26
CHAPTER 3. INSTALLING A CLUSTER ON ORACLE COMPUTE CLOUD@CUSTOMER BY USING THE AGENT-BASED INSTALLER	28
3.1. INSTALLATION PROCESS WORKFLOW	28
3.2. CREATING ORACLE COMPUTE CLOUD@CUSTOMER INFRASTRUCTURE RESOURCES AND SERVICES	28
3.3. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON COMPUTE CLOUD@CUSTOMER	29
3.4. CONFIGURING YOUR FIREWALL FOR OPENSHIFT CONTAINER PLATFORM	33
3.5. RUNNING A CLUSTER ON COMPUTE CLOUD@CUSTOMER	39
3.6. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON COMPUTE CLOUD@CUSTOMER	39
3.7. ADDITIONAL RESOURCES	41
CHAPTER 4. INSTALLING A CLUSTER ON ORACLE PRIVATE CLOUD APPLIANCE BY USING THE AGENT-BASED INSTALLER	42
4.1. INSTALLATION PROCESS WORKFLOW	42
4.2. CREATING ORACLE PRIVATE CLOUD APPLIANCE INFRASTRUCTURE RESOURCES AND SERVICES	42
4.3. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON PRIVATE CLOUD APPLIANCE	43
4.4. CONFIGURING YOUR FIREWALL FOR OPENSHIFT CONTAINER PLATFORM	47
4.5. RUNNING A CLUSTER ON PRIVATE CLOUD APPLIANCE	53
4.6. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON PRIVATE CLOUD APPLIANCE	53
4.7. ADDITIONAL RESOURCES	55

CHAPTER 5. INSTALLING A CLUSTER ON ORACLE COMPUTE CLOUD@CUSTOMER BY USING THE ASSISTED INSTALLER	56
5.1. OVERVIEW	56
5.2. PREPARING THE OCI BASTION SERVER	57
5.3. RUNNING THE TERRAFORM SCRIPT VIA THE HOME REGION	57
5.4. PREPARING THE OCI IMAGE	58
5.4.1. Generating the image in the Assisted Installer	58
5.4.2. Converting and uploading the image to Oracle Compute Cloud@Customer	59
5.5. RUNNING THE TERRAFORM SCRIPT VIA THE C3 REGION	59
5.6. COMPLETING THE INSTALLATION BY USING THE ASSISTED INSTALLER WEB CONSOLE	60
5.6.1. Assigning node roles	60
5.6.2. Configuring networking	61
5.6.3. Adding custom manifests	61
5.7. OPENING OPENSHIFT CONTAINER PLATFORM FROM THE ORACLE COMPUTE CLOUD@CUSTOMER WEB CONSOLE	62

CHAPTER 1. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE ASSISTED INSTALLER

You can use the Assisted Installer to install a cluster on Oracle® Cloud Infrastructure (OCI). This method is recommended for most users, and requires an internet connection.

If you want to set up the cluster manually or using other automation tools, or if you are working in a disconnected environment, you can use the Red Hat Agent-based Installer for the installation. For details, see [Installing a cluster on Oracle Cloud Infrastructure \(OCI\) by using the Agent-based Installer](#).



NOTE

You can deploy OpenShift Container Platform on a [Dedicated Region](#) (Oracle documentation) the same as any region from Oracle Cloud Infrastructure (OCI).

1.1. ABOUT THE ASSISTED INSTALLER AND OCI INTEGRATION

You can run cluster workloads on Oracle® Cloud Infrastructure (OCI) infrastructure that supports dedicated, hybrid, public, and multiple cloud environments. Both Red Hat and Oracle test, validate, and support running OCI in an OpenShift Container Platform cluster on OCI.

This section explains how to use the Assisted Installer to install an OpenShift Container Platform cluster on the OCI platform. The installation deploys cloud-native components such as Oracle Cloud Controller Manager (CCM) and Oracle Container Storage Interface (CSI), and integrates your cluster with OCI API resources such as instance node, load balancer, and storage.

The installation process uses the OpenShift Container Platform discovery ISO image provided by Red Hat, together with the scripts and manifests provided and maintained by OCI.

1.1.1. Preinstallation considerations

Before installing OpenShift Container Platform on Oracle Cloud Infrastructure (OCI), you must consider the following configuration choices.

Deployment platforms

The integration between OpenShift Container Platform and Oracle Cloud Infrastructure (OCI) is certified on both virtual machines (VMs) and bare-metal (BM) machines. Bare-metal installations using iSCSI boot drives require a secondary vNIC that is automatically created in the Terraform stack provided by Oracle.

Before you create a virtual machine (VM) or bare-metal (BM) machine, you must identify the relevant OCI shape. For details, see the following resource:

- [Cloud instance types \(Red Hat Ecosystem Catalog portal\)](#).

VPU sizing recommendations

To ensure the best performance conditions for your cluster workloads that operate on OCI, ensure that volume performance units (VPUs) for your block volume are sized for your workloads. The following list provides guidance for selecting the VPUs needed for specific performance needs:

- Test or proof of concept environment: 100 GB, and 20 to 30 VPUs.

- Basic environment: 500 GB, and 60 VPUs.
- Heavy production environment: More than 500 GB, and 100 or more VPUs.


Consider reserving additional VPUs to provide sufficient capacity for updates and scaling activities. For more information about VPUs, see [Volume Performance Units \(Oracle documentation\)](#).

Instance sizing recommendations

Find recommended values for compute instance CPU, memory, VPU, and volume size for OpenShift Container Platform nodes. For details, see [Instance Sizing Recommendations for OpenShift Container Platform on OCI Nodes \(Oracle documentation\)](#).

1.1.2. Workflow

Figure 1.1. High-level workflow for using the Assisted Installer in a connected environment to install a cluster on OCI



The procedure for using the Assisted Installer in a connected environment to install a cluster on OCI is outlined below:

1. In the OCI console, configure an OCI account to host the cluster:
 - a. Create a new child compartment under an existing compartment.
 - b. Create a new object storage bucket or use one provided by OCI.
 - c. Download the stack file template stored locally.
2. In the Assisted Installer console, set up a cluster:
 - a. Enter the cluster configurations.
 - b. Generate and download the discovery ISO image.
3. In the OCI console, create the infrastructure:
 - a. Upload the discovery ISO image to the OCI bucket.
 - b. Create a Pre-Authenticated Request (PAR) for the ISO image.
 - c. Upload the stack file template, and use it to create and apply the stack.
 - d. Copy the custom manifest YAML file from the stack.
4. In the Assisted Installer console, complete the cluster installation:

- a. Set roles for the cluster nodes.
- b. Upload the manifests provided by Oracle.
- c. Install the cluster.



IMPORTANT

The steps for provisioning OCI resources are provided as an example only. You can also choose to create the required resources through other methods; the scripts are just an example. Installing a cluster with infrastructure that you provide requires knowledge of the cloud provider and the installation process on OpenShift Container Platform. You can access OCI configurations to complete these steps, or use the configurations to model your own custom script.

Additional resources

- [Assisted Installer for OpenShift Container Platform](#)
- [Installing a Cluster with Red Hat's Assisted Installer \(Oracle documentation\)](#)
- [Internet access for OpenShift Container Platform](#)

1.2. PREPARING THE OCI ENVIRONMENT

Before installing OpenShift Container Platform using Assisted Installer, create the necessary resources and download the configuration file in the OCI environment.

Prerequisites

- You have an OCI account to host the cluster.
- If you use a firewall and you plan to use a Telemetry service, you configured your firewall to allow OpenShift Container Platform to access the sites required.

Procedure

1. Log in to your [Oracle Cloud Infrastructure \(OCI\)](#) account with administrator privileges.
2. Configure the account by defining the [Cloud Accounts and Resources \(Oracle documentation\)](#). Ensure that you create the following resources:
 - a. Create a child compartment for organizing, restricting access, and setting usage limits to OCI resources. For the full procedure, see [Creating a Compartment \(Oracle documentation\)](#).
 - b. Create a new object storage bucket into which you will upload the discovery ISO image. For the full procedure, see [Creating an Object Storage Bucket \(Oracle documentation\)](#).
3. Download the latest version of the **create-cluster-vX.X.X.zip** configuration file from the [oracle-quickstart/oci-openshift](#) repository. This file provides the infrastructure for the cluster and contains configurations for the following:
 - **Terraform Stacks:** The Terraform stack code for provisioning OCI resources to create and manage OpenShift Container Platform clusters on OCI.

- **Custom Manifests:** The manifest files needed for the installation of OpenShift Container Platform clusters on OCI.



NOTE

To make any changes to the manifests, you can clone the entire Oracle GitHub repository and access the **custom_manifests** and **terraform-stacks** directories directly.

For details, see [Configuration Files \(Oracle documentation\)](#).

1.3. USING THE ASSISTED INSTALLER TO GENERATE AN OCI-COMPATIBLE DISCOVERY ISO IMAGE

Create the cluster configuration and generate the discovery ISO image in the Assisted Installer web console.

Prerequisites

- You created a child compartment and an object storage bucket on OCI. For details, see [Preparing the OCI environment](#).
- You reviewed details about the OpenShift Container Platform installation and update processes.

1.3.1. Creating the cluster

Set the cluster details.

Procedure

1. Log into [Assisted Installer web console](#) with your credentials.
2. In the **Red Hat OpenShift** tile, select **OpenShift**.
3. In the **Red Hat OpenShift Container Platform** tile, select **Create Cluster**.
4. On the **Cluster Type** page, scroll down to the end of the **Cloud** tab, and select **Oracle Cloud Infrastructure (virtual machines)**.
5. On the **Create an OpenShift Cluster** page, select the **Interactive** tile.
6. On the **Cluster Details** page, complete the following fields:

Field	Action required
Cluster name	Specify the name of your cluster, such as oci . This is the same value as the cluster name in OCI.
Base domain	Specify the base domain of the cluster, such as openshift-demo.devcluster.openshift.com . This must be the same value as the zone DNS server in OCI.

Field	Action required
OpenShift version	<p>* For installations on virtual machines only, specify OpenShift 4.14 or a later version.</p> <p>* For installations that include bare metal machines, specify OpenShift 4.16 or a later version.</p>
CPU architecture	Specify x86_64 or Arm64 .
Integrate with external partner platforms	<p>Specify Oracle Cloud Infrastructure.</p> <p>After you specify this value, the Include custom manifests checkbox is selected by default and the Custom manifests page is added to the wizard.</p>

7. Leave the default settings for the remaining fields, and click **Next**.

8. On the **Operators** page, click **Next**.

1.3.2. Generating the Discovery ISO image

Generate and download the Discovery ISO image.

Procedure

1. On the **Host Discovery** page, click **Add hosts** and complete the following steps:
 - a. For the **Provisioning type** field, select **Minimal image file**
 - b. For the **SSH public key** field, add the SSH public key from your local system, by copying the output of the following command:


```
$ cat ~/.ssh/id_rsa.pub
```

 The SSH public key will be installed on all OpenShift Container Platform control plane and compute nodes.
 - c. Click **Generate Discovery ISO** to generate the discovery ISO image file.
 - d. Click **Download Discovery ISO** to save the file to your local system.

Additional resources

- [Installation and update](#)
- [Configuring your firewall](#)

1.4. PROVISIONING OCI INFRASTRUCTURE FOR YOUR CLUSTER

When using the Assisted Installer to create details for your OpenShift Container Platform cluster, you

specify these details in a Terraform stack. A stack is an OCI feature that automates the provisioning of all necessary OCI infrastructure resources that are required for installing an OpenShift Container Platform cluster on OCI.

Prerequisites

- You downloaded the discovery ISO image to a local directory. For details, see "Using the Assisted Installer to generate an OCI-compatible discovery ISO image".
- You downloaded the Terraform stack template to a local directory. For details, see "Preparing the OCI environment".

Procedure

1. Log in to your [Oracle Cloud Infrastructure \(OCI\)](#) account.
 2. Upload the discovery ISO image from your local drive to the new object storage bucket you created. For the full procedure, see [Uploading an Object Storage Object to a Bucket \(Oracle documentation\)](#).
 3. Locate the uploaded discovery ISO, and complete the following steps:
 - a. Create a Pre-Authenticated Request (PAR) for the ISO from the adjacent options menu.
 - b. Copy the generated URL to use as the OpenShift Image Source URI in the next step.
- For the full procedure, see [Creating a Pre-Authenticated Requests in Object Storage \(Oracle documentation\)](#).
4. Create and apply the Terraform stack:



IMPORTANT

The Terraform stack includes files for creating cluster resources and custom manifests. The stack also includes a script, and when you apply the stack, the script creates OCI resources, such as DNS records, an instance, and other resources. For a list of the resources, see the **terraform-stacks** folder in [OpenShift on OCI \(OSO\)](#).

- a. Upload the Terraform stacks template [terraform-stacks](#) to the new object storage bucket.
- b. Complete the stack information and click **Next**.



IMPORTANT

- Make sure that **Cluster Name** matches **Cluster Name** in Assisted Installer, and **Zone DNS** matches **Base Domain** in Assisted Installer.
- In the **OpenShift Image Source URI** field, paste the Pre-Authenticated Request URL link that you generated in the previous step.
- Ensure that the correct **Compute Shape** field value is defined, depending on whether you are installing on bare metal or a virtual machine. If not, select a different shape from the list. For details, see [Compute Shapes \(Oracle documentation\)](#).

- Click **Apply** to apply the stack.

For the full procedure, see [Creating OpenShift Container Platform Infrastructure Using Resource Manager \(Oracle documentation\)](#).

- Copy the **dynamic_custom_manifest.yml** file from the **Outputs** page of the Terraform stack.



NOTE

The YAML file contains all the required manifests, concatenated and preformatted with the configuration values. For details, see the [Custom Manifests README file](#).

For the full procedure, see [Getting the OpenShift Container Platform Custom Manifests for Installation \(Oracle documentation\)](#).

1.5. COMPLETING THE REMAINING ASSISTED INSTALLER STEPS

After you provision Oracle® Cloud Infrastructure (OCI) resources and upload OpenShift Container Platform custom manifest configuration files to OCI, you must complete the remaining cluster installation steps on the Assisted Installer before you can create an instance OCI. These steps include assigning node roles and adding custom manifests.

1.5.1. Assigning node roles

Following host discovery, the role of all nodes appears as **Auto-assign** by default. Change each of the node roles to either **Control Plane node** or **Worker**.

Prerequisites

- You created and applied the Terraform stack in OCI. For details, see "Provisioning OCI infrastructure for your cluster".

Procedure

- From the Assisted Installer user interface, go to the **Host discovery** page.
- Under the **Role** column, select either **Control plane node** or **Worker** for each targeted hostname. Then click **Next**.



NOTE

- Before continuing to the next step, wait for each node to reach **Ready** status.
- Expand the node to verify that the hardware type is bare metal.

- Accept the default settings for the **Storage** and **Networking** pages. Then click **Next**.

1.5.2. Adding custom manifests

Add the mandatory custom manifests provided by Oracle. For details, see [Custom Manifests \(Oracle documentation\)](#).

Prerequisites

- You copied the **dynamic_custom_manifest.yml** file from the Terraform stack in OCI. For details, see "Provisioning OCI infrastructure for your cluster".

Procedure

1. On the **Custom manifests** page, in the **Folder** field, select **manifests**. This is the Assisted Installer folder where you want to save the custom manifest file.
2. In the **File name** field, enter a filename, for example, **dynamic_custom_manifest.yml**.
3. Paste the contents of the **dynamic_custom_manifest.yml** file that you copied from OCI:
 - a. In the **Content** section, click the **Paste content** icon.
 - b. If you are using Firefox, click **OK** to close the dialog box, and then press **Ctrl+V**. Otherwise, skip this step.
4. Click **Next** to save the custom manifest.
5. From the **Review and create** page, click **Install cluster** to create your OpenShift Container Platform cluster on OCI.

After the cluster installation and initialization operations, the Assisted Installer indicates the completion of the cluster installation operation. For more information, see "Completing the installation" section in the Assisted Installer for OpenShift Container Platform document.

Additional resources

- [Assisted Installer for OpenShift Container Platform](#)

1.6. VERIFYING A SUCCESSFUL CLUSTER INSTALLATION ON OCI

Verify that your cluster was installed and is running effectively on Oracle® Cloud Infrastructure (OCI).

Procedure

1. From the [Red Hat Hybrid Cloud Console](#), go to **Clusters > Assisted Clusters** and select your cluster's name.
2. On the **Installation Progress** page, check that the Installation progress bar is at 100% and a message displays indicating **Installation completed successfully**.
3. Under **Host inventory**, confirm that the status of all control plane and compute nodes is **Installed**.



NOTE

OpenShift Container Platform designates one of the control plane nodes as the bootstrap virtual machine, eliminating the need for a separate bootstrap machine.

4. Click the Web Console URL, to access the OpenShift Container Platform web console.

5. From the menu, select **Compute > Nodes**
6. Locate your node from the **Nodes** table.
7. From the **Terminal** tab, verify that iSCSI appears next to the serial number.
8. From the **Overview** tab, check that your node has a **Ready** status.
9. Select the **YAML** tab.
10. Check the **labels** parameter, and verify that the listed labels apply to your configuration. For example, the **topology.kubernetes.io/region=us-sanjose-1** label indicates in what OCI region the node was deployed.

1.7. ADDING HOSTS TO THE CLUSTER FOLLOWING THE INSTALLATION

After creating a cluster with the Assisted Installer, you can use the Red Hat Hybrid Cloud Console to add new host nodes to the cluster and approve their certificate signing requests (CRSs).

For details, see [Adding Nodes to a Cluster \(Oracle documentation\)](#) .

1.8. TROUBLESHOOTING THE INSTALLATION OF A CLUSTER ON OCI

If you experience issues with using the Assisted Installer to install an OpenShift Container Platform cluster on Oracle® Cloud Infrastructure (OCI), read the following sections to troubleshoot common problems.

1.8.1. The Ingress Load Balancer in OCI is not at a healthy status

This issue is classed as a **Warning** because by using OCI to create a stack, you created a pool of compute nodes, 3 by default, that are automatically added as backend listeners for the Ingress Load Balancer. By default, the OpenShift Container Platform deploys 2 router pods, which are based on the default values from the OpenShift Container Platform manifest files. The **Warning** is expected because a mismatch exists with the number of router pods available, two, to run on the three compute nodes.

Figure 1.2. Example of a **Warning** message that is under the Backend set information tab on OCI

openshift_cluster_ingress_http_backend

Backend set information

Overall health

Backend health

Critical	Warning	Unknown	OK
2	0	0	2

Traffic between this network load balancer and its backends is subject to the governing security lists and network security groups.

Learn more about [network load balancers](#) and [security lists](#).

You do not need to modify the Ingress Load Balancer configuration. Instead, you can point the Ingress

Load Balancer to specific compute nodes that operate in your cluster on OpenShift Container Platform. To do this, use placement mechanisms, such as annotations, on OpenShift Container Platform to ensure router pods only run on the compute nodes that you originally configured on the Ingress Load Balancer as backend listeners.

1.8.2. OCI create stack operation fails with an Error: 400-InvalidParameter message

On attempting to create a stack on OCI, you identified that the **Logs** section of the job outputs an error message. For example:

Error: 400-InvalidParameter, DNS Label oci-demo does not follow Oracle requirements
Suggestion: Please update the parameter(s) in the Terraform config as per error message DNS Label oci-demo does not follow Oracle requirements
Documentation: https://registry.terraform.io/providers/oracle/oci/latest/docs/resources/core_vcn

Go to the [Install OpenShift with the Assisted Installer](#) page on the Hybrid Cloud Console, and check the **Cluster name** field on the **Cluster Details** step. Remove any special characters, such as a hyphen (-), from the name, because these special characters are not compatible with the OCI naming conventions. For example, change **oci-demo** to **ocidemo**.

Additional resources

- [Troubleshooting OpenShift Container Platform on OCI \(Oracle documentation\)](#)
- [Installing an on-premise cluster using the Assisted Installer](#)

CHAPTER 2. INSTALLING A CLUSTER ON ORACLE CLOUD INFRASTRUCTURE (OCI) BY USING THE AGENT-BASED INSTALLER

In OpenShift Container Platform 4.19, you can use the Agent-based Installer to install a cluster on Oracle® Cloud Infrastructure (OCI), so that you can run cluster workloads on infrastructure that supports dedicated, hybrid, public, and multiple cloud environments.

Installing a cluster on OCI is supported for virtual machines (VMs) and bare-metal machines.



NOTE

You can deploy OpenShift Container Platform on a [Dedicated Region](#) (Oracle documentation) the same as any region from Oracle Cloud Infrastructure (OCI).

2.1. THE AGENT-BASED INSTALLER AND OCI OVERVIEW

You can install an OpenShift Container Platform cluster on Oracle® Cloud Infrastructure (OCI) by using the Agent-based Installer. Red Hat and Oracle test, validate, and support running OCI workloads in an OpenShift Container Platform cluster.

The Agent-based Installer provides the ease of use of the Assisted Installation service, but with the capability to install a cluster in either a connected or disconnected environment.

The following diagrams show workflows for connected and disconnected environments:

Figure 2.1. Workflow for using the Agent-based installer in a connected environment to install a cluster on OCI





Figure 2.2. Workflow for using the Agent-based installer in a disconnected environment to install a cluster on OCI



684_OpenShift_0624

OCI provides services that can meet your regulatory compliance, performance, and cost-effectiveness needs. OCI supports 64-bit **x86** instances and 64-bit **ARM** instances.



NOTE

Consider selecting a nonvolatile memory express (NVMe) drive or a solid-state drive (SSD) for your boot disk, because these drives offer low latency and high throughput capabilities for your boot disk.

By running your OpenShift Container Platform cluster on OCI, you can access the following capabilities:

- Compute flexible shapes, where you can customize the number of Oracle® CPUs (OCPUs) and memory resources for your VM. With access to this capability, a cluster's workload can perform operations in a resource-balanced environment. You can find all RHEL-certified OCI shapes by going to the Oracle page on the Red Hat Ecosystem Catalog portal.
- Block Volume storage, where you can configure scaling and auto-tuning settings for your storage volume, so that the Block Volume service automatically adjusts the performance level to optimize performance.



IMPORTANT

To ensure the best performance conditions for your cluster workloads that operate on OCI and on the OCVS service, ensure volume performance units (VPUs) for your block volume is sized for your workloads. The following list provides some guidance in selecting the VPUs needed for specific performance needs:

- Test or proof of concept environment: 100 GB, and 20 to 30 VPUs.
- Basic environment: 500 GB, and 60 VPUs.
- Heavy production environment: More than 500 GB, and 100 or more VPUs.

Consider reserving additional VPUs to provide sufficient capacity for updates and scaling activities. For more information about VPUs, see Volume Performance Units (Oracle documentation).

Additional resources

- [Installation process](#)
- [Internet access for OpenShift Container Platform](#)
- [Understanding the Agent-based Installer](#)
- [Overview of the Compute Service \(Oracle documentation\)](#)
- [Volume Performance Units \(Oracle documentation\)](#)
- [Instance Sizing Recommendations for OpenShift Container Platform on OCI Nodes \(Oracle documentation\)](#)

2.2. INSTALLATION PROCESS WORKFLOW

The following workflow describes a high-level outline for the process of installing an OpenShift Container Platform cluster on OCI using the Agent-based Installer:

1. Create OCI resources and services (Oracle).
2. Disconnected environments: Prepare a web server that is accessible by OCI instances (Red Hat).
3. Prepare configuration files for the Agent-based Installer (Red Hat).
4. Generate the agent ISO image (Red Hat).
5. Disconnected environments: Upload the rootfs image to the web server (Red Hat).
6. Configure your firewall for OpenShift Container Platform (Red Hat).
7. Upload the agent ISO image to a storage bucket (Oracle).
8. Create a custom image from the uploaded agent ISO image (Oracle).
9. Create compute instances on OCI (Oracle).
10. Verify that your cluster runs on OCI (Oracle).

2.3. CREATING OCI INFRASTRUCTURE RESOURCES AND SERVICES

You must create an OCI environment on your virtual machine (VM) or bare-metal shape. By creating this environment, you can install OpenShift Container Platform and deploy a cluster on an infrastructure that supports a wide range of cloud options and strong security policies. Having prior knowledge of OCI components can help you with understanding the concept of OCI resources and how you can configure them to meet your organizational needs.

The Agent-based Installer method for installing an OpenShift Container Platform cluster on OCI requires that you manually create OCI resources and services.



IMPORTANT

To ensure compatibility with OpenShift Container Platform, you must set **A** as the record type for each DNS record and name records as follows:

- **api.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer
- **api-int.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer
- ***.apps.<cluster_name>.<base_domain>**, which targets the **ingressVIP** parameter of the Ingress load balancer

The **api.*** and **api-int.*** DNS records relate to control plane machines, so you must ensure that all nodes in your installed OpenShift Container Platform cluster can access these DNS records.

Prerequisites

- You configured an OCI account to host the OpenShift Container Platform cluster. See [Prerequisites \(Oracle documentation\)](#).

Procedure

- Create the required OCI resources and services.

For installations in a connected environment, see [Provisioning Cluster Infrastructure Using Terraform \(Oracle documentation\)](#).

For installations in a disconnected environment, see [Provisioning OCI Resources for the Agent-based Installer in Disconnected Environments \(Oracle documentation\)](#).

Additional resources

- [Learn About Oracle Cloud Basics \(Oracle documentation\)](#)

2.4. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON OCI

You must create the **install-config.yaml** and the **agent-config.yaml** configuration files so that you can use the Agent-based Installer to generate a bootable ISO image. The Agent-based installation comprises a bootable ISO that has the Assisted discovery agent and the Assisted Service. Both of these

components are required to perform the cluster installation, but the latter component runs on only one of the hosts.



NOTE

You can also use the Agent-based Installer to generate or accept Zero Touch Provisioning (ZTP) custom resources.

Prerequisites

- You reviewed details about the OpenShift Container Platform installation and update processes.
- You read the documentation on selecting a cluster installation method and preparing the method for users.
- You have read the "Preparing to install with the Agent-based Installer" documentation.
- You downloaded the Agent-Based Installer and the command-line interface (CLI) from the [Red Hat Hybrid Cloud Console](#).
- If you are installing in a disconnected environment, you have prepared a mirror registry in your environment and mirrored release images to the registry.



IMPORTANT

Check that your **openshift-install** binary version relates to your local image container registry and not a shared registry, such as Red Hat Quay, by running the following command:

```
$ ./openshift-install version
```

Example output for a shared registry binary

```
./openshift-install 4.19.0
built from commit ae7977b7d1ca908674a0d45c5c243c766fa4b2ca
release image registry.ci.openshift.org/origin/release:4.19ocp-
release@sha256:0da6316466d60a3a4535d5fed3589feb0391989982fba59d47d
4c729912d6363
release architecture amd64
```

- You have logged in to the OpenShift Container Platform with administrator privileges.

Procedure

1. Create an installation directory to store configuration files in by running the following command:

```
$ mkdir ~/<directory_name>
```

2. Configure the **install-config.yaml** configuration file to meet the needs of your organization and save the file in the directory you created.

install-config.yaml file that sets an external platform

```

# install-config.yaml
apiVersion: v1
baseDomain: <base_domain> ①
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  networkType: OVNKubernetes
  machineNetwork:
    - cidr: <ip_address_from_cidr> ②
  serviceNetwork:
    - 172.30.0.0/16
compute:
  - architecture: amd64 ③
    hyperthreading: Enabled
    name: worker
    replicas: 0
controlPlane:
  architecture: amd64 ④
  hyperthreading: Enabled
  name: master
  replicas: 3
platform:
  external:
    platformName: oci ⑤
    cloudControllerManager: External
    sshKey: <public_ssh_key> ⑥
    pullSecret: '<pull_secret>' ⑦
# ...

```

- ① The base domain of your cloud provider.
- ② The IP address from the virtual cloud network (VCN) that the CIDR allocates to resources and components that operate on your network.
- ③ ④ Depending on your infrastructure, you can select either **arm64** or **amd64**.
- ⑤ Set **OCI** as the external platform, so that OpenShift Container Platform can integrate with OCI.
- ⑥ Specify your SSH public key.
- ⑦ The pull secret that you need for authenticate purposes when downloading container images for OpenShift Container Platform components and services, such as Quay.io. See [Install OpenShift Container Platform 4](#) from the Red Hat Hybrid Cloud Console.

3. Create a directory on your local system named **openshift**. This must be a subdirectory of the installation directory.



IMPORTANT

Do not move the **install-config.yaml** or **agent-config.yaml** configuration files to the **openshift** directory.

4. If you used a stack to provision OCI infrastructure resources: Copy and paste the **dynamic_custom_manifest** output of the OCI stack into a file titled **manifest.yaml** and save the file in the **openshift** directory.
5. If you did not use a stack to provision OCI infrastructure resources: Download and prepare custom manifests to create an Agent ISO image:
 - a. Go to [Configuration Files](#) (Oracle documentation) and follow the link to the custom manifests directory on GitHub.
 - b. Copy the contents of the **condensed-manifest.yml** file and save it locally to a file in the **openshift** directory.
 - c. In the **condensed-manifest.yml** file, update the sections marked with **TODO** to specify the compartment Oracle® Cloud Identifier (OCID), VCN OCID, subnet OCID from the load balancer, and the security lists OCID.
6. Configure the **agent-config.yaml** configuration file to meet your organization's requirements.

Sample **agent-config.yaml** file for an IPv4 network.

```
apiVersion: v1beta1
metadata:
  name: <cluster_name> ①
  namespace: <cluster_namespace> ②
  rendezvousIP: <ip_address_from_CIDR> ③
  bootArtifactsBaseUrl: <server_URL> ④
# ...
```

- ① The cluster name that you specified in your DNS record.
- ② The namespace of your cluster on OpenShift Container Platform.
- ③ If you use IPv4 as the network IP address format, ensure that you set the **rendezvousIP** parameter to an IPv4 address that the VCN's Classless Inter-Domain Routing (CIDR) method allocates on your network. Also ensure that at least one instance from the pool of instances that you booted with the ISO matches the IP address value you set for the **rendezvousIP** parameter.
- ④ The URL of the server where you want to upload the rootfs image. This parameter is required only for disconnected environments.

7. Generate a minimal ISO image, which excludes the rootfs image, by entering the following command in your installation directory:

```
$ ./openshift-install agent create image --log-level debug
```

The command also completes the following actions:

- Creates a subdirectory, **./<installation_directory>/auth directory**, and places **kubeadmin-password** and **kubeconfig** files in the subdirectory.
- Creates a **rendezvousIP** file based on the IP address that you specified in the **agent-config.yaml** configuration file.

- Optional: Any modifications you made to **agent-config.yaml** and **install-config.yaml** configuration files get imported to the Zero Touch Provisioning (ZTP) custom resources.



IMPORTANT

The Agent-based Installer uses Red Hat Enterprise Linux CoreOS (RHCOS). The rootfs image, which is mentioned in a later step, is required for booting, recovering, and repairing your operating system.

8. Disconnected environments only: Upload the rootfs image to a web server.
 - a. Go to the `./<installation_directory>/boot-artifacts` directory that was generated when you created the minimal ISO image.
 - b. Use your preferred web server, such as any Hypertext Transfer Protocol daemon (**httpd**), to upload the rootfs image to the location specified in the **bootArtifactsBaseUrl** parameter of the **agent-config.yaml** file.

For example, if the **bootArtifactsBaseUrl** parameter states **http://192.168.122.20**, you would upload the generated rootfs image to this location so that the Agent-based installer can access the image from **http://192.168.122.20/agent.x86_64-rootfs.img**. After the Agent-based installer boots the minimal ISO for the external platform, the Agent-based Installer downloads the rootfs image from the **http://192.168.122.20/agent.x86_64-rootfs.img** location into the system memory.



NOTE

The Agent-based Installer also adds the value of the **bootArtifactsBaseUrl** to the minimal ISO Image's configuration, so that when the Operator boots a cluster's node, the Agent-based Installer downloads the rootfs image into system memory.



IMPORTANT

Consider that the full ISO image, which is in excess of **1 GB**, includes the rootfs image. The image is larger than the minimal ISO Image, which is typically less than **150 MB**.


Additional resources

- [About OpenShift Container Platform installation](#)
- [Selecting a cluster installation type](#)
- [Preparing to install with the Agent-based Installer](#)
- [Downloading the Agent-based Installer](#)
- [Creating a mirror registry with mirror registry for Red Hat OpenShift](#)
- [Mirroring the OpenShift Container Platform image repository](#)
- [Optional: Using ZTP manifests](#)

2.5. CONFIGURING YOUR FIREWALL FOR OPENSHIFT CONTAINER PLATFORM

Before you install OpenShift Container Platform, you must configure your firewall to grant access to the sites that OpenShift Container Platform requires. When using a firewall, make additional configurations to the firewall so that OpenShift Container Platform can access the sites that it requires to function.

For a disconnected environment, you must mirror content from both Red Hat and Oracle. This environment requires that you create firewall rules to expose your firewall to specific ports and registries.



NOTE

If your environment has a dedicated load balancer in front of your OpenShift Container Platform cluster, review the allowlists between your firewall and load balancer to prevent unwanted network restrictions to your cluster.

Procedure

1. Set the following registry URLs for your firewall's allowlist:

URL	Port	Function
registry.redhat.io	443	Provides core container images
access.redhat.com	443	Hosts a signature store that a container client requires for verifying images pulled from registry.access.redhat.com . In a firewall environment, ensure that this resource is on the allowlist.
registry.access.redhat.co m	443	Hosts all the container images that are stored on the Red Hat Ecosystem Catalog, including core container images.
quay.io	443	Provides core container images
cdn.quay.io	443	Provides core container images
cdn01.quay.io	443	Provides core container images
cdn02.quay.io	443	Provides core container images
cdn03.quay.io	443	Provides core container images
cdn04.quay.io	443	Provides core container images
cdn05.quay.io	443	Provides core container images
cdn06.quay.io	443	Provides core container images

URL	Port	Function
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com
icr.io	443	Provides IBM Cloud Pak container images. This domain is only required if you use IBM Cloud Paks.
cp.icr.io	443	Provides IBM Cloud Pak container images. This domain is only required if you use IBM Cloud Paks.

- You can use the wildcard ***.quay.io** instead of **cdn.quay.io** and **cdn0[1-6].quay.io** in your allowlist.
 - You can use the wildcard ***.access.redhat.com** to simplify the configuration and ensure that all subdomains, including **registry.access.redhat.com**, are allowed.
 - When you add a site, such as **quay.io**, to your allowlist, do not add a wildcard entry, such as ***.quay.io**, to your denylist. In most cases, image registries use a content delivery network (CDN) to serve images. If a firewall blocks access, image downloads are denied when the initial download request redirects to a hostname such as **cdn01.quay.io**.
2. Set your firewall's allowlist to include any site that provides resources for a language or framework that your builds require.
 3. If you do not disable Telemetry, you must grant access to the following URLs to access Red Hat Insights:

URL	Port	Function
cert-api.access.redhat.com	443	Required for Telemetry
api.access.redhat.com	443	Required for Telemetry
infogw.api.openshift.com	443	Required for Telemetry
console.redhat.com	443	Required for Telemetry and for insights-operator

4. Set your firewall's allowlist to include the following registry URLs:

URL	Port	Function
-----	------	----------

URL	Port	Function
api.openshift.com	443	Required both for your cluster token and to check if updates are available for the cluster.
rhcosh.mirror.openshift.com	443	Required to download Red Hat Enterprise Linux CoreOS (RHCOS) images.

- Set your firewall's allowlist to include the following external URLs. Each repository URL hosts OCI containers. Consider mirroring images to as few repositories as possible to reduce any performance issues.

URL	Port	Function
k8s.gcr.io	port	A Kubernetes registry that hosts container images for a community-based image registry. This image registry is hosted on a custom Google Container Registry (GCR) domain.
ghcr.io	port	A GitHub image registry where you can store and manage Open Container Initiative images. Requires an access token to publish, install, and delete private, internal, and public packages.
storage.googleapis.com	443	A source of release image signatures, although the Cluster Version Operator needs only a single functioning source.
registry.k8s.io	port	Replaces the k8s.gcr.io image registry because the k8s.gcr.io image registry does not support other platforms and vendors.

2.6. RUNNING A CLUSTER ON OCI

To run a cluster on Oracle® Cloud Infrastructure (OCI), you must upload the generated agent ISO image to the default Object Storage bucket on OCI. Additionally, you must create a compute instance from the supplied base image, so that your OpenShift Container Platform and OCI can communicate with each other for the purposes of running the cluster on OCI.



NOTE

OCI supports the following OpenShift Container Platform cluster topologies:

- Installing an OpenShift Container Platform cluster on a single node.
- A highly available cluster that has a minimum of three control plane instances and two compute instances.
- A compact three-node cluster that has a minimum of three control plane instances.

Prerequisites

- You generated an agent ISO image. See the "Creating configuration files for installing a cluster on OCI" section.

Procedure

1. Upload the agent ISO image to Oracle's default Object Storage bucket and import the agent ISO image as a custom image to this bucket. Ensure you that you configure the custom image to boot in Unified Extensible Firmware Interface (UEFI) mode. For more information, see [Creating the OpenShift Container Platform ISO Image \(Oracle documentation\)](#) .
2. Create a compute instance from the supplied base image for your cluster topology. See [Creating the OpenShift Container Platform cluster on OCI \(Oracle documentation\)](#) .



IMPORTANT

Before you create the compute instance, check that you have enough memory and disk resources for your cluster. Additionally, ensure that at least one compute instance has the same IP address as the address stated under `rendezvousIP` in the `agent-config.yaml` file.

Additional resources

- [Recommended resources for topologies](#)
- [Instance Sizing Recommendations for OpenShift Container Platform on OCI Nodes \(Oracle documentation\)](#)
- [Troubleshooting OpenShift Container Platform on OCI \(Oracle documentation\)](#)

2.7. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON OCI

Verify that your cluster was installed and is running effectively on Oracle® Cloud Infrastructure (OCI).

Prerequisites

- You created all the required OCI resources and services. See the "Creating OCI infrastructure resources and services" section.

- You created **install-config.yaml** and **agent-config.yaml** configuration files. See the "Creating configuration files for installing a cluster on OCI" section.
- You uploaded the agent ISO image to a default Oracle Object Storage bucket, and you created a compute instance on OCI. For more information, see "Running a cluster on OCI".

Procedure

After you deploy the compute instance on a self-managed node in your OpenShift Container Platform cluster, you can monitor the cluster's status by choosing one of the following options:

- From the OpenShift Container Platform CLI, enter the following command:

```
$ ./openshift-install agent wait-for install-complete --log-level debug
```

Check the status of the **rendezvous** host node that runs the bootstrap node. After the host reboots, the host forms part of the cluster.

- Use the **kubeconfig** API to check the status of various OpenShift Container Platform components. For the **KUBECONFIG** environment variable, set the relative path of the cluster's **kubeconfig** configuration file:

```
$ export KUBECONFIG=~/auth/kubeconfig
```

Check the status of each of the cluster's self-managed nodes. CCM applies a label to each node to designate the node as running in a cluster on OCI.

```
$ oc get nodes -A
```

Output example

NAME	STATUS	ROLES	AGE	VERSION
main-0.private.agenttest.oraclevcn.com	Ready	control-plane, master	7m	v1.27.4+6eeca63
main-1.private.agenttest.oraclevcn.com	Ready	control-plane, master	15m	v1.27.4+d7fa83f
main-2.private.agenttest.oraclevcn.com	Ready	control-plane, master	15m	v1.27.4+d7fa83f

Check the status of each of the cluster's Operators, with the CCM Operator status being a good indicator that your cluster is running.

```
$ oc get co
```

Truncated output example

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
MESSAGE					
authentication	4.19.0-0	True	False	False	6m18s
baremetal	4.19.0-0	True	False	False	2m42s
network	4.19.0-0	True	True	False	5m58s
					Progressing: ...
					...

2.8. ADDITIONAL RESOURCES

- [Gathering log data from a failed Agent-based installation](#)

- Adding worker nodes to an on-premise cluster

CHAPTER 3. INSTALLING A CLUSTER ON ORACLE COMPUTE CLOUD@CUSTOMER BY USING THE AGENT-BASED INSTALLER

You can use the Agent-based Installer to install a cluster on Oracle® Compute Cloud@Customer, so that you can run cluster workloads on on-premise infrastructure while still using Oracle® Cloud Infrastructure (OCI) services.

3.1. INSTALLATION PROCESS WORKFLOW

The following workflow describes a high-level outline for the process of installing an OpenShift Container Platform cluster on Compute Cloud@Customer using the Agent-based Installer:

1. Create Compute Cloud@Customer resources and services (Oracle).
2. Prepare configuration files for the Agent-based Installer (Red Hat).
3. Generate the agent ISO image (Red Hat).
4. Convert the ISO image to an Oracle Cloud Infrastructure (OCI) image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Compute Cloud@Customer system (Oracle).
5. Disconnected environments: Prepare a web server that is accessible by OCI instances (Red Hat).
6. Disconnected environments: Upload the rootfs image to the web server (Red Hat).
7. Configure your firewall for OpenShift Container Platform (Red Hat).
8. Create control plane nodes and configure load balancers (Oracle).
9. Create compute nodes and configure load balancers (Oracle).
10. Verify that your cluster runs on OCI (Oracle).

3.2. CREATING ORACLE COMPUTE CLOUD@CUSTOMER INFRASTRUCTURE RESOURCES AND SERVICES

You must create an Compute Cloud@Customer environment on your virtual machine (VM) shape. By creating this environment, you can install OpenShift Container Platform and deploy a cluster on an infrastructure that supports a wide range of cloud options and strong security policies. Having prior knowledge of OCI components can help you with understanding the concept of OCI resources and how you can configure them to meet your organizational needs.



IMPORTANT

To ensure compatibility with OpenShift Container Platform, you must set **A** as the record type for each DNS record and name records as follows:

- **api.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer
- **api-int.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer
- ***.apps.<cluster_name>.<base_domain>**, which targets the **ingressVIP** parameter of the Ingress load balancer

The **api.*** and **api-int.*** DNS records relate to control plane machines, so you must ensure that all nodes in your installed OpenShift Container Platform cluster can access these DNS records.

Prerequisites

- You configured an OCI account to host the OpenShift Container Platform cluster. See "Access and Considerations" in [OpenShift Cluster Setup with Agent Based Installer on Compute Cloud@Customer](#) (Oracle documentation).

Procedure

- Create the required Compute Cloud@Customer resources and services. For more information, see "Terraform Script Execution" in [OpenShift Cluster Setup with Agent Based Installer on Compute Cloud@Customer](#) (Oracle documentation).

Additional resources

- [Learn About Oracle Cloud Basics](#) (Oracle documentation)

3.3. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON COMPUTE CLOUD@CUSTOMER

You must create the **install-config.yaml** and the **agent-config.yaml** configuration files so that you can use the Agent-based Installer to generate a bootable ISO image. The Agent-based installation comprises a bootable ISO that has the Assisted discovery agent and the Assisted Service. Both of these components are required to perform the cluster installation, but the latter component runs on only one of the hosts.



NOTE

You can also use the Agent-based Installer to generate or accept Zero Touch Provisioning (ZTP) custom resources.

Prerequisites

- You reviewed details about the OpenShift Container Platform installation and update processes.

- You read the documentation on selecting a cluster installation method and preparing the method for users.
- You have read the "Preparing to install with the Agent-based Installer" documentation.
- You downloaded the Agent-Based Installer and the command-line interface (CLI) from the [Red Hat Hybrid Cloud Console](#).
- If you are installing in a disconnected environment, you have prepared a mirror registry in your environment and mirrored release images to the registry.



IMPORTANT

Check that your **openshift-install** binary version relates to your local image container registry and not a shared registry, such as Red Hat Quay, by running the following command:

```
$ ./openshift-install version
```

Example output for a shared registry binary

```
./openshift-install 4.19.0
built from commit ae7977b7d1ca908674a0d45c5c243c766fa4b2ca
release image registry.ci.openshift.org/origin/release:4.19ocp-
release@sha256:0da6316466d60a3a4535d5fed3589feb0391989982fba59d47d
4c729912d6363
release architecture amd64
```

- You have logged in to the OpenShift Container Platform with administrator privileges.

Procedure

1. Create an installation directory to store configuration files in by running the following command:

```
$ mkdir ~/<directory_name>
```

2. Configure the **install-config.yaml** configuration file to meet the needs of your organization and save the file in the directory you created.

install-config.yaml file that sets an external platform

```
# install-config.yaml
apiVersion: v1
baseDomain: <base_domain> ①
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  network type: OVNKubernetes
  machineNetwork:
    - cidr: <ip_address_from_cidr> ②
  serviceNetwork:
    - 172.30.0.0/16
```

```

compute:
  - architecture: amd64 3
    hyperthreading: Enabled
    name: worker
    replicas: 0
controlPlane:
  architecture: amd64 4
  hyperthreading: Enabled
  name: master
  replicas: 3
platform:
  external:
    platformName: oci 5
    cloudControllerManager: External
    sshKey: <public_ssh_key> 6
    pullSecret: '<pull_secret>' 7
# ...

```

- 1** The base domain of your cloud provider.
- 2** The IP address from the virtual cloud network (VCN) that the CIDR allocates to resources and components that operate on your network.
- 3** **4** Depending on your infrastructure, you can select either **arm64** or **amd64**.
- 5** Set **OCI** as the external platform, so that OpenShift Container Platform can integrate with OCI.
- 6** Specify your SSH public key.
- 7** The pull secret that you need for authenticate purposes when downloading container images for OpenShift Container Platform components and services, such as Quay.io. See [Install OpenShift Container Platform 4](#) from the Red Hat Hybrid Cloud Console.

3. Create a directory on your local system named **openshift**. This must be a subdirectory of the installation directory.



IMPORTANT

Do not move the **install-config.yaml** or **agent-config.yaml** configuration files to the **openshift** directory.

4. Configure the Oracle custom manifest files.
 - a. Go to "Prepare the OpenShift Master Images" in [OpenShift Cluster Setup with Agent Based Installer on Compute Cloud@Customer](#) (Oracle documentation).
 - b. Copy and paste the **oci-ccm.yaml**, **oci-csi.yaml**, and **machineconfig-ccm.yaml** files into your **openshift** directory.
 - c. Edit the **oci-ccm.yaml** and **oci-csi.yaml** files to specify the compartment Oracle® Cloud Identifier (OCID), VCN OCID, subnet OCID from the load balancer, the security lists OCID, and the **c3-cert.pem** section.

5. Configure the **agent-config.yaml** configuration file to meet your organization's requirements.

Sample agent-config.yaml file for an IPv4 network.

```
apiVersion: v1beta1
metadata:
  name: <cluster_name> 1
  namespace: <cluster_namespace> 2
  rendezvousIP: <ip_address_from_CIDR>
  bootArtifactsBaseUrl: <server_URL> 4
# ...
```

- 1 The cluster name that you specified in your DNS record.
 - 2 The namespace of your cluster on OpenShift Container Platform.
 - 3 If you use IPv4 as the network IP address format, ensure that you set the **rendezvousIP** parameter to an IPv4 address that the VCN's Classless Inter-Domain Routing (CIDR) method allocates on your network. Also ensure that at least one instance from the pool of instances that you booted with the ISO matches the IP address value you set for the **rendezvousIP** parameter.
 - 4 The URL of the server where you want to upload the rootfs image. This parameter is required only for disconnected environments.

6. Generate a minimal ISO image, which excludes the rootfs image, by entering the following command in your installation directory:

```
$ ./openshift-install agent create image --log-level debug
```

The command also completes the following actions:

- Creates a subdirectory, `./<installation_directory>/auth directory`, and places **kubeadmin-password** and **kubeconfig** files in the subdirectory.
 - Creates a **rendezvousIP** file based on the IP address that you specified in the **agent-config.yaml** configuration file.
 - Optional: Any modifications you made to **agent-config.yaml** and **install-config.yaml** configuration files get imported to the Zero Touch Provisioning (ZTP) custom resources.



IMPORTANT

The Agent-based Installer uses Red Hat Enterprise Linux CoreOS (RHCOS). The rootfs image, which is mentioned in a later step, is required for booting, recovering, and repairing your operating system.

7. Disconnected environments only: Upload the rootfs image to a web server.

- a. Go to the `./<installation_directory>/boot-artifacts` directory that was generated when you created the minimal ISO image.

- b. Use your preferred web server, such as any Hypertext Transfer Protocol daemon (**httpd**), to upload the rootfs image to the location specified in the **bootArtifactsBaseUrl** parameter of the **agent-config.yaml** file.

For example, if the **bootArtifactsBaseUrl** parameter states **http://192.168.122.20**, you would upload the generated rootfs image to this location so that the Agent-based installer can access the image from **http://192.168.122.20/agent.x86_64-rootfs.img**. After the Agent-based installer boots the minimal ISO for the external platform, the Agent-based Installer downloads the rootfs image from the **http://192.168.122.20/agent.x86_64-rootfs.img** location into the system memory.



NOTE

The Agent-based Installer also adds the value of the **bootArtifactsBaseUrl** to the minimal ISO Image's configuration, so that when the Operator boots a cluster's node, the Agent-based Installer downloads the rootfs image into system memory.



IMPORTANT

Consider that the full ISO image, which is in excess of **1** GB, includes the rootfs image. The image is larger than the minimal ISO Image, which is typically less than **150** MB.

Additional resources

- [About OpenShift Container Platform installation](#)
- [Selecting a cluster installation type](#)
- [Preparing to install with the Agent-based Installer](#)
- [Downloading the Agent-based Installer](#)
- [Creating a mirror registry with mirror registry for Red Hat OpenShift](#)
- [Mirroring the OpenShift Container Platform image repository](#)
- [Optional: Using ZTP manifests](#)

3.4. CONFIGURING YOUR FIREWALL FOR OPENSHIFT CONTAINER PLATFORM

Before you install OpenShift Container Platform, you must configure your firewall to grant access to the sites that OpenShift Container Platform requires. When using a firewall, make additional configurations to the firewall so that OpenShift Container Platform can access the sites that it requires to function.

There are no special configuration considerations for services running on only controller nodes compared to worker nodes.



NOTE

If your environment has a dedicated load balancer in front of your OpenShift Container Platform cluster, review the allowlists between your firewall and load balancer to prevent unwanted network restrictions to your cluster.

Procedure

1. Set the following registry URLs for your firewall's allowlist:

URL	Port	Function
registry.redhat.io	443	Provides core container images
access.redhat.com	443	Hosts a signature store that a container client requires for verifying images pulled from registry.access.redhat.com . In a firewall environment, ensure that this resource is on the allowlist.
registry.access.redhat.com	443	Hosts all the container images that are stored on the Red Hat Ecosystem Catalog, including core container images.
quay.io	443	Provides core container images
cdn.quay.io	443	Provides core container images
cdn01.quay.io	443	Provides core container images
cdn02.quay.io	443	Provides core container images
cdn03.quay.io	443	Provides core container images
cdn04.quay.io	443	Provides core container images
cdn05.quay.io	443	Provides core container images
cdn06.quay.io	443	Provides core container images
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com
icr.io	443	Provides IBM Cloud Pak container images. This domain is only required if you use IBM Cloud Paks.
cp.icr.io	443	Provides IBM Cloud Pak container images. This domain is only required if you use IBM Cloud Paks.

- You can use the wildcard ***.quay.io** instead of **cdn.quay.io** and **cdn0[1-6].quay.io** in your allowlist.

- You can use the wildcard ***.access.redhat.com** to simplify the configuration and ensure that all subdomains, including **registry.access.redhat.com**, are allowed.
 - When you add a site, such as **quay.io**, to your allowlist, do not add a wildcard entry, such as ***.quay.io**, to your denylist. In most cases, image registries use a content delivery network (CDN) to serve images. If a firewall blocks access, image downloads are denied when the initial download request redirects to a hostname such as **cdn01.quay.io**.
2. Set your firewall's allowlist to include any site that provides resources for a language or framework that your builds require.
 3. If you do not disable Telemetry, you must grant access to the following URLs to access Red Hat Insights:

URL	Port	Function
cert-api.access.redhat.com	443	Required for Telemetry
api.access.redhat.com	443	Required for Telemetry
infogw.api.openshift.com	443	Required for Telemetry
console.redhat.com	443	Required for Telemetry and for insights-operator

4. If you use Alibaba Cloud, Amazon Web Services (AWS), Microsoft Azure, or Google Cloud to host your cluster, you must grant access to the URLs that offer the cloud provider API and DNS for that cloud:

Cloud	URL	Port	Function
Alibaba	*.aliyuncs.com	443	Required to access Alibaba Cloud services and resources. Review the Alibaba endpoints_config.go file to find the exact endpoints to allow for the regions that you use.
AWS	aws.amazon.com	443	Used to install and manage clusters in an AWS environment.
	*.amazonaws.com	443	Required to access AWS services and resources. Review the AWS Service Endpoints in the AWS documentation to find the exact endpoints to allow for the regions that you use.
	Alternatively, if you choose to not use a wildcard for AWS APIs, you must include the following URLs in your allowlist:		
	ec2.amazonaws.com	443	Used to install and manage clusters in an AWS environment.

Cloud	URL	Port	Function
	events.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	iam.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	route53.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.dualstack. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	sts.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	sts. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	tagging.us-east-1.amazonaws.com	443	Used to install and manage clusters in an AWS environment. This endpoint is always us-east-1 , regardless of the region the cluster is deployed in.
	ec2. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	elasticloadbalancing. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	servicequotas. <aws_region>.amazonaws.com	443	Required. Used to confirm quotas for deploying the service.
	tagging. <aws_region>.amazonaws.com	443	Allows the assignment of metadata about AWS resources in the form of tags.

Cloud	URL	Port	Function
	*.cloudfront.net	443	Used to provide access to CloudFront. If you use the AWS Security Token Service (STS) and the private S3 bucket, you must provide access to CloudFront.
GCP	*.googleapis.com	443	Required to access Google Cloud services and resources. Review Cloud Endpoints in the Google Cloud documentation to find the endpoints to allow for your APIs.
	accounts.google.com	443	Required to access your Google Cloud account.
Microsoft Azure	management.azure.com	443	Required to access Microsoft Azure services and resources. Review the Microsoft Azure REST API reference in the Microsoft Azure documentation to find the endpoints to allow for your APIs.
	*.blob.core.windows.net	443	Required to download Ignition files.
	login.microsoftonline.com	443	Required to access Microsoft Azure services and resources. Review the Azure REST API reference in the Microsoft Azure documentation to find the endpoints to allow for your APIs.

5. Allowlist the following URLs:

URL	Port	Function
*.apps.<cluster_name>.<base_domain>	443	Required to access the default cluster routes unless you set an ingress wildcard during installation.
api.openshift.com	443	Required both for your cluster token and to check if updates are available for the cluster.
console.redhat.com	443	Required for your cluster token.

URL	Port	Function
mirror.openshift.com	443	Required to access mirrored installation content and images. This site is also a source of release image signatures, although the Cluster Version Operator needs only a single functioning source.
quayio-production-s3.s3.amazonaws.com	443	Required to access Quay image content in AWS.
rhcosh.mirror.openshift.com	443	Required to download Red Hat Enterprise Linux CoreOS (RHCOS) images.
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com
storage.googleapis.com/openshift-release	443	A source of release image signatures, although the Cluster Version Operator needs only a single functioning source.

Operators require route access to perform health checks. Specifically, the authentication and web console Operators connect to two routes to verify that the routes work. If you are the cluster administrator and do not want to allow `*.apps.<cluster_name>.<base_domain>`, then allow these routes:

- `oauth-openshift.apps.<cluster_name>.<base_domain>`
- `canary-openshift-ingress-canary.apps.<cluster_name>.<base_domain>`
- `console-openshift-console.apps.<cluster_name>.<base_domain>`, or the hostname that is specified in the `spec.route.hostname` field of the `consoles.operator/cluster` object if the field is not empty.

6. Allowlist the following URL for optional third-party content:

URL	Port	Function
registry.connect.redhat.com	443	Required for all third-party images and certified operators.

7. If you use a default Red Hat Network Time Protocol (NTP) server allow the following URLs:

- `1.rhel.pool.ntp.org`
- `2.rhel.pool.ntp.org`
- `3.rhel.pool.ntp.org`



NOTE

If you do not use a default Red Hat NTP server, verify the NTP server for your platform and allow it in your firewall.

3.5. RUNNING A CLUSTER ON COMPUTE CLOUD@CUSTOMER

To run a cluster on Oracle® Compute Cloud@Customer, you must first convert your generated Agent ISO image into an OCI image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Compute Cloud@Customer system.



NOTE

Compute Cloud@Customer supports the following OpenShift Container Platform cluster topologies:

- Installing an OpenShift Container Platform cluster on a single node.
- A highly available cluster that has a minimum of three control plane instances and two compute instances.
- A compact three-node cluster that has a minimum of three control plane instances.

Prerequisites

- You generated an Agent ISO image. See the "Creating configuration files for installing a cluster on Compute Cloud@Customer" section.

Procedure

1. Convert the agent ISO image to an OCI image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Compute Cloud@Customer system. See "Prepare the OpenShift Master Images" in [OpenShift Cluster Setup with Agent Based Installer on Compute Cloud@Customer \(Oracle documentation\)](#) for instructions.
2. Create control plane instances on Compute Cloud@Customer. See "Create control plane instances on C3 and Master Node LB Backend Sets" in [OpenShift Cluster Setup with Agent Based Installer on Compute Cloud@Customer \(Oracle documentation\)](#) for instructions.
3. Create a compute instance from the supplied base image for your cluster topology. See "Add worker nodes" in [OpenShift Cluster Setup with Agent Based Installer on Compute Cloud@Customer \(Oracle documentation\)](#) for instructions.



IMPORTANT

Before you create the compute instance, check that you have enough memory and disk resources for your cluster. Additionally, ensure that at least one compute instance has the same IP address as the address stated under `rendezvousIP` in the `agent-config.yaml` file.

3.6. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON COMPUTE CLOUD@CUSTOMER

Verify that your cluster was installed and is running effectively on Compute Cloud@Customer.

Prerequisites

- You created all the required Oracle® Compute Cloud@Customer resources and services. See the "Creating Oracle Compute Cloud@Customer infrastructure resources and services" section.
- You created **install-config.yaml** and **agent-config.yaml** configuration files. See the "Creating configuration files for installing a cluster on Compute Cloud@Customer" section.
- You uploaded the agent ISO image to a default Oracle Object Storage bucket, and you created a compute instance on Compute Cloud@Customer. For more information, see "Running a cluster on Compute Cloud@Customer".

Procedure

After you deploy the compute instance on a self-managed node in your OpenShift Container Platform cluster, you can monitor the cluster's status by choosing one of the following options:

- From the OpenShift Container Platform CLI, enter the following command:

```
$ ./openshift-install agent wait-for install-complete --log-level debug
```

Check the status of the **rendezvous** host node that runs the bootstrap node. After the host reboots, the host forms part of the cluster.

- Use the **kubeconfig** API to check the status of various OpenShift Container Platform components. For the **KUBECONFIG** environment variable, set the relative path of the cluster's **kubeconfig** configuration file:

```
$ export KUBECONFIG=~/auth/kubeconfig
```

Check the status of each of the cluster's self-managed nodes. CCM applies a label to each node to designate the node as running in a cluster on OCI.

```
$ oc get nodes -A
```

Output example

NAME	STATUS	ROLES	AGE	VERSION
main-0.private.agenttest.oraclevcn.com	Ready	control-plane, master	7m	v1.27.4+6eeca63
main-1.private.agenttest.oraclevcn.com	Ready	control-plane, master	15m	v1.27.4+d7fa83f
main-2.private.agenttest.oraclevcn.com	Ready	control-plane, master	15m	v1.27.4+d7fa83f

Check the status of each of the cluster's Operators, with the CCM Operator status being a good indicator that your cluster is running.

```
$ oc get co
```

Truncated output example

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.19.0-0	True	False	False	6m18s

baremetal	4.19.0-0	True	False	False	2m42s
network	4.19.0-0	True	True	False	5m58s Progressing: ...
...					

3.7. ADDITIONAL RESOURCES

- Gathering log data from a failed Agent-based installation
- Adding worker nodes to an on-premise cluster

CHAPTER 4. INSTALLING A CLUSTER ON ORACLE PRIVATE CLOUD APPLIANCE BY USING THE AGENT-BASED INSTALLER

You can use the Agent-based Installer to install a cluster on Oracle® Private Cloud Appliance, so that you can run cluster workloads on on-premise infrastructure while still using Oracle® Cloud Infrastructure (OCI) services.


4.1. INSTALLATION PROCESS WORKFLOW

The following workflow describes a high-level outline for the process of installing an OpenShift Container Platform cluster on Private Cloud Appliance using the Agent-based Installer:

1. Create Private Cloud Appliance resources and services (Oracle).
2. Prepare configuration files for the Agent-based Installer (Red Hat).
3. Generate the agent ISO image (Red Hat).
4. Convert the ISO image to an Oracle Cloud Infrastructure (OCI) image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Private Cloud Appliance system (Oracle).
5. Disconnected environments: Prepare a web server that is accessible by OCI instances (Red Hat).
6. Disconnected environments: Upload the rootfs image to the web server (Red Hat).
7. Configure your firewall for OpenShift Container Platform (Red Hat).
8. Create control plane nodes and configure load balancers (Oracle).
9. Create compute nodes and configure load balancers (Oracle).
10. Verify that your cluster runs on OCI (Oracle).

4.2. CREATING ORACLE PRIVATE CLOUD APPLIANCE INFRASTRUCTURE RESOURCES AND SERVICES

You must create an Private Cloud Appliance environment on your virtual machine (VM) shape. By creating this environment, you can install OpenShift Container Platform and deploy a cluster on an infrastructure that supports a wide range of cloud options and strong security policies. Having prior knowledge of OCI components can help you with understanding the concept of OCI resources and how you can configure them to meet your organizational needs.



IMPORTANT

To ensure compatibility with OpenShift Container Platform, you must set **A** as the record type for each DNS record and name records as follows:

- **api.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer
- **api-int.<cluster_name>.<base_domain>**, which targets the **apiVIP** parameter of the API load balancer
- ***.apps.<cluster_name>.<base_domain>**, which targets the **ingressVIP** parameter of the Ingress load balancer

The **api.*** and **api-int.*** DNS records relate to control plane machines, so you must ensure that all nodes in your installed OpenShift Container Platform cluster can access these DNS records.

Prerequisites

- You configured an OCI account to host the OpenShift Container Platform cluster. See "Access and Considerations" in [OpenShift Cluster Setup with Agent Based Installer on Private Cloud Appliance](#) (Oracle documentation).

Procedure

- Create the required Private Cloud Appliance resources and services. For more information, see "Terraform Script Execution" in [OpenShift Cluster Setup with Agent Based Installer on Private Cloud Appliance](#) (Oracle documentation).

Additional resources

- [Learn About Oracle Cloud Basics](#) (Oracle documentation)

4.3. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON PRIVATE CLOUD APPLIANCE

You must create the **install-config.yaml** and the **agent-config.yaml** configuration files so that you can use the Agent-based Installer to generate a bootable ISO image. The Agent-based installation comprises a bootable ISO that has the Assisted discovery agent and the Assisted Service. Both of these components are required to perform the cluster installation, but the latter component runs on only one of the hosts.



NOTE

You can also use the Agent-based Installer to generate or accept Zero Touch Provisioning (ZTP) custom resources.

Prerequisites

- You reviewed details about the OpenShift Container Platform installation and update processes.

- You read the documentation on selecting a cluster installation method and preparing the method for users.
- You have read the "Preparing to install with the Agent-based Installer" documentation.
- You downloaded the Agent-Based Installer and the command-line interface (CLI) from the [Red Hat Hybrid Cloud Console](#).
- If you are installing in a disconnected environment, you have prepared a mirror registry in your environment and mirrored release images to the registry.



IMPORTANT

Check that your **openshift-install** binary version relates to your local image container registry and not a shared registry, such as Red Hat Quay, by running the following command:

```
$ ./openshift-install version
```

Example output for a shared registry binary

```
./openshift-install 4.19.0
built from commit ae7977b7d1ca908674a0d45c5c243c766fa4b2ca
release image registry.ci.openshift.org/origin/release:4.19ocp-
release@sha256:0da6316466d60a3a4535d5fed3589feb0391989982fba59d47d
4c729912d6363
release architecture amd64
```

- You have logged in to the OpenShift Container Platform with administrator privileges.

Procedure

1. Create an installation directory to store configuration files in by running the following command:

```
$ mkdir ~/<directory_name>
```

2. Configure the **install-config.yaml** configuration file to meet the needs of your organization and save the file in the directory you created.

install-config.yaml file that sets an external platform

```
# install-config.yaml
apiVersion: v1
baseDomain: <base_domain> ①
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  network type: OVNKubernetes
  machineNetwork:
    - cidr: <ip_address_from_cidr> ②
  serviceNetwork:
    - 172.30.0.0/16
```

```

compute:
  - architecture: amd64 3
  hyperthreading: Enabled
  name: worker
  replicas: 0
controlPlane:
  architecture: amd64 4
  hyperthreading: Enabled
  name: master
  replicas: 3
platform:
  external:
    platformName: oci 5
    cloudControllerManager: External
  sshKey: <public_ssh_key> 6
  pullSecret: '<pull_secret>' 7
# ...

```

- 1** The base domain of your cloud provider.
- 2** The IP address from the virtual cloud network (VCN) that the CIDR allocates to resources and components that operate on your network.
- 3** **4** Depending on your infrastructure, you can select either **arm64** or **amd64**.
- 5** Set **OCI** as the external platform, so that OpenShift Container Platform can integrate with OCI.
- 6** Specify your SSH public key.
- 7** The pull secret that you need for authenticate purposes when downloading container images for OpenShift Container Platform components and services, such as Quay.io. See [Install OpenShift Container Platform 4](#) from the Red Hat Hybrid Cloud Console.

3. Create a directory on your local system named **openshift**. This must be a subdirectory of the installation directory.



IMPORTANT

Do not move the **install-config.yaml** or **agent-config.yaml** configuration files to the **openshift** directory.

4. Configure the Oracle custom manifest files.
 - a. Go to "Prepare the OpenShift Master Images" in [OpenShift Cluster Setup with Agent Based Installer on Private Cloud Appliance](#) (Oracle documentation).
 - b. Copy and paste the **oci-ccm.yaml**, **oci-csi.yaml**, and **machineconfig-ccm.yaml** files into your **openshift** directory.
 - c. Edit the **oci-ccm.yaml** and **oci-csi.yaml** files to specify the compartment Oracle® Cloud Identifier (OCID), VCN OCID, subnet OCID from the load balancer, the security lists OCID, and the **c3-cert.pem** section.

5. Configure the **agent-config.yaml** configuration file to meet your organization's requirements.

Sample agent-config.yaml file for an IPv4 network.

```
apiVersion: v1beta1
metadata:
  name: <cluster_name> 1
  namespace: <cluster_namespace> 2
  rendezvousIP: <ip_address_from_CIDR>
  bootArtifactsBaseUrl: <server_URL> 4
# ...
```

- 1 The cluster name that you specified in your DNS record.
 - 2 The namespace of your cluster on OpenShift Container Platform.
 - 3 If you use IPv4 as the network IP address format, ensure that you set the **rendezvousIP** parameter to an IPv4 address that the VCN's Classless Inter-Domain Routing (CIDR) method allocates on your network. Also ensure that at least one instance from the pool of instances that you booted with the ISO matches the IP address value you set for the **rendezvousIP** parameter.
 - 4 The URL of the server where you want to upload the rootfs image. This parameter is required only for disconnected environments.

6. Generate a minimal ISO image, which excludes the rootfs image, by entering the following command in your installation directory:

```
$ ./openshift-install agent create image --log-level debug
```

The command also completes the following actions:

- Creates a subdirectory, `./<installation_directory>/auth` directory, and places **kubeadmin-password** and **kubeconfig** files in the subdirectory.
 - Creates a **rendezvousIP** file based on the IP address that you specified in the **agent-config.yaml** configuration file.
 - Optional: Any modifications you made to **agent-config.yaml** and **install-config.yaml** configuration files get imported to the Zero Touch Provisioning (ZTP) custom resources.



IMPORTANT

The Agent-based Installer uses Red Hat Enterprise Linux CoreOS (RHCOS). The rootfs image, which is mentioned in a later step, is required for booting, recovering, and repairing your operating system.

7. Disconnected environments only: Upload the rootfs image to a web server.

- a. Go to the `./<installation_directory>/boot-artifacts` directory that was generated when you created the minimal ISO image.

- b. Use your preferred web server, such as any Hypertext Transfer Protocol daemon (**httpd**), to upload the rootfs image to the location specified in the **bootArtifactsBaseUrl** parameter of the **agent-config.yaml** file.

For example, if the **bootArtifactsBaseUrl** parameter states **http://192.168.122.20**, you would upload the generated rootfs image to this location so that the Agent-based installer can access the image from **http://192.168.122.20/agent.x86_64-rootfs.img**. After the Agent-based installer boots the minimal ISO for the external platform, the Agent-based Installer downloads the rootfs image from the **http://192.168.122.20/agent.x86_64-rootfs.img** location into the system memory.



NOTE

The Agent-based Installer also adds the value of the **bootArtifactsBaseUrl** to the minimal ISO Image's configuration, so that when the Operator boots a cluster's node, the Agent-based Installer downloads the rootfs image into system memory.



IMPORTANT

Consider that the full ISO image, which is in excess of **1** GB, includes the rootfs image. The image is larger than the minimal ISO Image, which is typically less than **150** MB.

Additional resources

- [About OpenShift Container Platform installation](#)
- [Selecting a cluster installation type](#)
- [Preparing to install with the Agent-based Installer](#)
- [Downloading the Agent-based Installer](#)
- [Creating a mirror registry with mirror registry for Red Hat OpenShift](#)
- [Mirroring the OpenShift Container Platform image repository](#)
- [Optional: Using ZTP manifests](#)

4.4. CONFIGURING YOUR FIREWALL FOR OPENSHIFT CONTAINER PLATFORM

Before you install OpenShift Container Platform, you must configure your firewall to grant access to the sites that OpenShift Container Platform requires. When using a firewall, make additional configurations to the firewall so that OpenShift Container Platform can access the sites that it requires to function.

There are no special configuration considerations for services running on only controller nodes compared to worker nodes.



NOTE

If your environment has a dedicated load balancer in front of your OpenShift Container Platform cluster, review the allowlists between your firewall and load balancer to prevent unwanted network restrictions to your cluster.

Procedure

1. Set the following registry URLs for your firewall's allowlist:

URL	Port	Function
registry.redhat.io	443	Provides core container images
access.redhat.com	443	Hosts a signature store that a container client requires for verifying images pulled from registry.access.redhat.com . In a firewall environment, ensure that this resource is on the allowlist.
registry.access.redhat.co m	443	Hosts all the container images that are stored on the Red Hat Ecosystem Catalog, including core container images.
quay.io	443	Provides core container images
cdn.quay.io	443	Provides core container images
cdn01.quay.io	443	Provides core container images
cdn02.quay.io	443	Provides core container images
cdn03.quay.io	443	Provides core container images
cdn04.quay.io	443	Provides core container images
cdn05.quay.io	443	Provides core container images
cdn06.quay.io	443	Provides core container images
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com
icr.io	443	Provides IBM Cloud Pak container images. This domain is only required if you use IBM Cloud Paks.
cp.icr.io	443	Provides IBM Cloud Pak container images. This domain is only required if you use IBM Cloud Paks.

- You can use the wildcard ***.quay.io** instead of **cdn.quay.io** and **cdn0[1-6].quay.io** in your allowlist.

- You can use the wildcard ***.access.redhat.com** to simplify the configuration and ensure that all subdomains, including **registry.access.redhat.com**, are allowed.
 - When you add a site, such as **quay.io**, to your allowlist, do not add a wildcard entry, such as ***.quay.io**, to your denylist. In most cases, image registries use a content delivery network (CDN) to serve images. If a firewall blocks access, image downloads are denied when the initial download request redirects to a hostname such as **cdn01.quay.io**.
2. Set your firewall's allowlist to include any site that provides resources for a language or framework that your builds require.
 3. If you do not disable Telemetry, you must grant access to the following URLs to access Red Hat Insights:

URL	Port	Function
cert-api.access.redhat.com	443	Required for Telemetry
api.access.redhat.com	443	Required for Telemetry
infogw.api.openshift.com	443	Required for Telemetry
console.redhat.com	443	Required for Telemetry and for insights-operator

4. If you use Alibaba Cloud, Amazon Web Services (AWS), Microsoft Azure, or Google Cloud to host your cluster, you must grant access to the URLs that offer the cloud provider API and DNS for that cloud:

Cloud	URL	Port	Function
Alibaba	*.aliyuncs.com	443	Required to access Alibaba Cloud services and resources. Review the Alibaba endpoints_config.go file to find the exact endpoints to allow for the regions that you use.
AWS	aws.amazon.com	443	Used to install and manage clusters in an AWS environment.
	*.amazonaws.com	443	Required to access AWS services and resources. Review the AWS Service Endpoints in the AWS documentation to find the exact endpoints to allow for the regions that you use.
	Alternatively, if you choose to not use a wildcard for AWS APIs, you must include the following URLs in your allowlist:		
	ec2.amazonaws.com	443	Used to install and manage clusters in an AWS environment.

Cloud	URL	Port	Function
	events.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	iam.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	route53.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.dualstack. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	sts.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	sts. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	tagging.us-east-1.amazonaws.com	443	Used to install and manage clusters in an AWS environment. This endpoint is always us-east-1 , regardless of the region the cluster is deployed in.
	ec2. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	elasticloadbalancing. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	servicequotas. <aws_region>.amazonaws.com	443	Required. Used to confirm quotas for deploying the service.
	tagging. <aws_region>.amazonaws.com	443	Allows the assignment of metadata about AWS resources in the form of tags.

Cloud	URL	Port	Function
	*.cloudfront.net	443	Used to provide access to CloudFront. If you use the AWS Security Token Service (STS) and the private S3 bucket, you must provide access to CloudFront.
GCP	*.googleapis.com	443	Required to access Google Cloud services and resources. Review Cloud Endpoints in the Google Cloud documentation to find the endpoints to allow for your APIs.
	accounts.google.com	443	Required to access your Google Cloud account.
Microsoft Azure	management.azure.com	443	Required to access Microsoft Azure services and resources. Review the Microsoft Azure REST API reference in the Microsoft Azure documentation to find the endpoints to allow for your APIs.
	*.blob.core.windows.net	443	Required to download Ignition files.
	login.microsoftonline.com	443	Required to access Microsoft Azure services and resources. Review the Azure REST API reference in the Microsoft Azure documentation to find the endpoints to allow for your APIs.

5. Allowlist the following URLs:

URL	Port	Function
*.apps.<cluster_name>.<base_domain>	443	Required to access the default cluster routes unless you set an ingress wildcard during installation.
api.openshift.com	443	Required both for your cluster token and to check if updates are available for the cluster.
console.redhat.com	443	Required for your cluster token.

URL	Port	Function
mirror.openshift.com	443	Required to access mirrored installation content and images. This site is also a source of release image signatures, although the Cluster Version Operator needs only a single functioning source.
quayio-production-s3.s3.amazonaws.com	443	Required to access Quay image content in AWS.
rhcosh.mirror.openshift.com	443	Required to download Red Hat Enterprise Linux CoreOS (RHCOS) images.
sso.redhat.com	443	The https://console.redhat.com site uses authentication from sso.redhat.com
storage.googleapis.com/openshift-release	443	A source of release image signatures, although the Cluster Version Operator needs only a single functioning source.

Operators require route access to perform health checks. Specifically, the authentication and web console Operators connect to two routes to verify that the routes work. If you are the cluster administrator and do not want to allow `*.apps.<cluster_name>.<base_domain>`, then allow these routes:

- `oauth-openshift.apps.<cluster_name>.<base_domain>`
- `canary-openshift-ingress-canary.apps.<cluster_name>.<base_domain>`
- `console-openshift-console.apps.<cluster_name>.<base_domain>`, or the hostname that is specified in the `spec.route.hostname` field of the `consoles.operator/cluster` object if the field is not empty.

6. Allowlist the following URL for optional third-party content:

URL	Port	Function
registry.connect.redhat.com	443	Required for all third-party images and certified operators.

7. If you use a default Red Hat Network Time Protocol (NTP) server allow the following URLs:

- `1.rhel.pool.ntp.org`
- `2.rhel.pool.ntp.org`
- `3.rhel.pool.ntp.org`



NOTE

If you do not use a default Red Hat NTP server, verify the NTP server for your platform and allow it in your firewall.

4.5. RUNNING A CLUSTER ON PRIVATE CLOUD APPLIANCE

To run a cluster on Oracle® Private Cloud Appliance, you must first convert your generated Agent ISO image into an OCI image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Private Cloud Appliance system.



NOTE

Private Cloud Appliance supports the following OpenShift Container Platform cluster topologies:

- Installing an OpenShift Container Platform cluster on a single node.
- A highly available cluster that has a minimum of three control plane instances and two compute instances.
- A compact three-node cluster that has a minimum of three control plane instances.

Prerequisites

- You generated an Agent ISO image. See the "Creating configuration files for installing a cluster on Private Cloud Appliance" section.

Procedure

1. Convert the agent ISO image to an OCI image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Private Cloud Appliance system. See "Prepare the OpenShift Master Images" in [OpenShift Cluster Setup with Agent Based Installer on Private Cloud Appliance \(Oracle documentation\)](#) for instructions.
2. Create control plane instances on Private Cloud Appliance. See "Create control plane instances on PCA and Master Node LB Backend Sets" in [OpenShift Cluster Setup with Agent Based Installer on Private Cloud Appliance \(Oracle documentation\)](#) for instructions.
3. Create a compute instance from the supplied base image for your cluster topology. See "Add worker nodes" in [OpenShift Cluster Setup with Agent Based Installer on Private Cloud Appliance \(Oracle documentation\)](#) for instructions.



IMPORTANT

Before you create the compute instance, check that you have enough memory and disk resources for your cluster. Additionally, ensure that at least one compute instance has the same IP address as the address stated under `rendezvousIP` in the `agent-config.yaml` file.

4.6. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON PRIVATE CLOUD APPLIANCE

Verify that your cluster was installed and is running effectively on Private Cloud Appliance.

Prerequisites

- You created all the required Oracle® Private Cloud Appliance resources and services. See the "Creating Oracle Private Cloud Appliance infrastructure resources and services" section.
- You created **install-config.yaml** and **agent-config.yaml** configuration files. See the "Creating configuration files for installing a cluster on Private Cloud Appliance" section.
- You uploaded the agent ISO image to a default Oracle Object Storage bucket, and you created a compute instance on Private Cloud Appliance. For more information, see "Running a cluster on Private Cloud Appliance".

Procedure

After you deploy the compute instance on a self-managed node in your OpenShift Container Platform cluster, you can monitor the cluster's status by choosing one of the following options:

- From the OpenShift Container Platform CLI, enter the following command:

```
$ ./openshift-install agent wait-for install-complete --log-level debug
```

Check the status of the **rendezvous** host node that runs the bootstrap node. After the host reboots, the host forms part of the cluster.

- Use the **kubeconfig** API to check the status of various OpenShift Container Platform components. For the **KUBECONFIG** environment variable, set the relative path of the cluster's **kubeconfig** configuration file:

```
$ export KUBECONFIG=~/auth/kubeconfig
```

Check the status of each of the cluster's self-managed nodes. CCM applies a label to each node to designate the node as running in a cluster on OCI.

```
$ oc get nodes -A
```

Output example

NAME	STATUS	ROLES	AGE	VERSION
main-0.private.agenttest.oraclevcn.com	Ready	control-plane, master	7m	v1.27.4+6eeca63
main-1.private.agenttest.oraclevcn.com	Ready	control-plane, master	15m	v1.27.4+d7fa83f
main-2.private.agenttest.oraclevcn.com	Ready	control-plane, master	15m	v1.27.4+d7fa83f

Check the status of each of the cluster's Operators, with the CCM Operator status being a good indicator that your cluster is running.

```
$ oc get co
```

Truncated output example

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.19.0-0	True	False	False	6m18s

baremetal	4.19.0-0	True	False	False	2m42s
network	4.19.0-0	True	True	False	5m58s Progressing: ...
...					

4.7. ADDITIONAL RESOURCES

- Gathering log data from a failed Agent-based installation
- Adding worker nodes to an on-premise cluster

CHAPTER 5. INSTALLING A CLUSTER ON ORACLE COMPUTE CLOUD@CUSTOMER BY USING THE ASSISTED INSTALLER

With Oracle® Compute Cloud@Customer (C3), you can run applications and middleware by using Oracle® Cloud Infrastructure (OCI) services on high performance cloud infrastructure in your data center.

5.1. OVERVIEW

You can install OpenShift Container Platform on Oracle Compute Cloud@Customer by using the Assisted Installer.

For an alternative installation method, see "Installing a cluster on Oracle® Compute Cloud@Customer by using the Agent-based Installer".

Preinstallation considerations

- Ensure that your installation meets the prerequisites specified for Oracle. For details, see the "Access and Considerations" section in the [Oracle documentation](#).
- Ensure that your infrastructure is certified and uses a compatible cloud instance type. For details, see [Oracle Cloud Infrastructure](#).
- Ensure that you are performing the installation on a virtual machine.

Installation process

The installation process builds a bastion host within the designated compartment of the OpenShift Container Platform cluster. The bastion host is used to run two Terraform scripts:

- The first script builds IAM Resources in the OCI Home region of the Compute Cloud@Customer system (two Dynamic Groups and one Policy).
- The second script builds the infrastructure resources on the Compute Cloud@Customer system to support the OpenShift Container Platform cluster, including the OpenShift Container Platform VCN, public and private subnets, load balancers, Internet GW, NAT GW, and DNS server. The script includes all the resources needed to activate the control plane nodes and compute nodes that form a cluster.

The bastion host is installed in the designated OpenShift Container Platform Compartment and configured to communicate through a designated Compute Cloud@Customer DRG Subnet or Internet GW Subnet within the Compute Cloud@Customer parent tenancy.

The installation process subsequently provisions three control plane (master) nodes and three compute (worker) nodes, together with the external and internal Load Balancers that form the cluster. This is the standard implementation for Oracle Cloud Infrastructure (OCI).

Main steps

The main steps of the procedure are as follows:

1. Preparing the Compute Cloud@Customer bastion server.
2. Running the Terraform script via the Home region.
3. Preparing the OpenShift Container Platform image for Oracle Cloud Infrastructure (OCI).

4. Running the Terraform script via the Compute Cloud@Customer region.
5. Installing the cluster by using the Assisted Installer web console.

5.2. PREPARING THE OCI BASTION SERVER

By implementing a bastion host, you can securely and efficiently manage access to your Oracle Compute Cloud@Customer resources, ensuring that your private instances remain protected and accessible only through a secure, controlled entry point.

Prerequisites

- See the "Bastion server - prerequisites" section in the [Oracle documentation](#).

Procedure

1. Install the bastion server. For details, see the "Bastion Installation" section in the [Oracle documentation](#).
2. Install the Terraform application which is used to run the Terraform script. For details, see the "Terraform Installation" section in the [Oracle documentation](#).
3. Install and configure the OCI command-line interface (CLI). For details, see the "Installing and Configuring the OCI CLI" section in the [Oracle documentation](#).

Additional resources

- [Quick start - Installing the CLI \(Oracle documentation\)](#).

5.3. RUNNING THE TERRAFORM SCRIPT VIA THE HOME REGION

Copy the Terraform scripts **createInfraResources.tf** and **terraform.tfvars** onto the bastion server. Then run the **createInfraResources.tf** script to create the Dynamic Group Identity resources on your Compute Cloud@Customer OCI Home Region. These resources include dynamic groups, policies, and tags.

Prerequisites

- You have tenancy privileges to create Dynamic Groups and Policies. If not, you can manually provision them during this procedure.

Procedure

1. Connect to the bastion server via SSH.
2. Create **OpenShift\createResourceOnHomeRegion** folders.
3. Copy the **createInfraResources.tf** and **terraform.tfvars** files from the C3_PCA GitHub repository into the **createResourceOnHomeRegion** folder.
4. Ensure that you have access to the source environment, and that your C3 certificate has been exported.
5. Run the **createInfraResources.tf** Terraform script.

For the full procedure, see the "Terraform Script Execution Part-1 (Run Script via Home Region)" section in the [Oracle documentation](#).

5.4. PREPARING THE OCI IMAGE

Generate the OpenShift Container Platform ISO image in the Assisted Installer on the Red Hat portal. Then, convert the image to an Oracle Cloud Infrastructure (OCI) compatible image and upload it to the **Custom Images** page of your Oracle Compute Cloud@Customer environment.

You can generate, convert and upload the image on your laptop and not on the bastion server or within environments such as Oracle Solution Center.

5.4.1. Generating the image in the Assisted Installer

Create a cluster and download the discovery ISO image.

Procedure

1. Log in to [Assisted Installer web console](#) with your credentials.
2. In the **Red Hat OpenShift** tile, select **OpenShift**.
3. In the **Red Hat OpenShift Container Platform** tile, select **Create Cluster**.
4. On the **Cluster Type** page, scroll to the end of the **Cloud** tab, and select **Oracle Cloud Infrastructure (virtual machines)**.
5. On the **Create an OpenShift Cluster** page, select the **Interactive** tile.
6. On the **Cluster Details** page, complete the following fields:

Field	Action required
Cluster name	Specify the name of your OpenShift Container Platform cluster. This name is the same name you used to create the resource via the Terraform scripts. The name must be between 1-54 characters. It can use lowercase alphanumeric characters or hyphen (-), but must start and end with a lowercase letter or a number.
Base domain	Specify the base domain of the cluster. This is the value used for the zone_dns variables in the Terraform scripts that run on Compute Cloud@Customer. Make a note of the value.
OpenShift version	Select OpenShift 4.16.20 . If it is not immediately visible, scroll to the end of the dropdown menu, select Show all available versions , and type the version in the search box.
Integrate with external partner platforms	Select Oracle Cloud Infrastructure . After you specify this value, the Include custom manifests checkbox is selected by default and the Custom manifests page is added to the wizard.

7. Leave the default settings for the remaining fields, and click **Next**.
8. On the **Operators** page, click **Next**.
9. On the **Host Discovery** page, click **Add hosts** and complete the following steps:



NOTE

The minimal ISO image is the mandatory **Provisioning type** for the Oracle Cloud Infrastructure (OCI), and cannot be changed.

- a. In the **SSH public key** field, add the SSH public key by copying the output of the following command:

```
$ cat ~/.ssh/id_rsa.pub
```

The SSH public key will be installed on all OpenShift Container Platform control plane and compute nodes.

- b. Click the **Show proxy settings** checkbox.
- c. Add the proxy variables from the **/etc/environment** file of the bastion server that you configured earlier:

```
http_proxy=http://www-proxy.<your_domain>.com:80
https_proxy=http://www-proxy.<your_domain>.com:80
no_proxy=localhost,127.0.0.1,1,2,3,4,5,6,7,8,9,0,.<your_domain>.com
#(ie.oracle.com,,oraclecorp.com)
```

- d. Click **Generate Discovery ISO** to generate the discovery ISO image file.
10. Click **Download Discovery ISO** to save the file to your local system. After you download the ISO file, you can rename it as required, for example **discovery_image_<your_cluster_name>.iso**.

5.4.2. Converting and uploading the image to Oracle Compute Cloud@Customer

Convert the ISO image to an OCI image and upload it to your Compute Cloud@Customer system from your OCI Home Region Object Store.

Procedure

1. Convert the image from ISO to OCI.
2. Upload the OCI image to an OCI bucket, and generate a Pre-Authenticated Request (PAR) URL.
3. Import the OCI image to the Compute Cloud@Customer portal.
4. Copy the Oracle Cloud Identifier (OCID) of the image for use in the next procedure.

For the full procedure, see step 6 - 8 in the "OpenShift Image Preparation" section of the [Oracle documentation](#).

5.5. RUNNING THE TERRAFORM SCRIPT VIA THE C3 REGION

Run the **terraform.tfvars** Terraform script to create all infrastructure resources on Compute Cloud@Customer. These resources include the OpenShift Container Platform VCN, public and private subnets, load balancers, internet GW, NAT GW, and DNS server.

This procedure deploys a cluster consisting of three control plane (master) and three compute (worker) nodes. After deployment, you must rename and reboot the nodes. This process temporarily duplicates nodes, requiring manual cleanup in the next procedure.

Procedure

1. Connect to the bastion server via SSH.
2. Set the C3 Certificate location and export the certificate.
3. Run the **terraform.tfvars** script to create three control plane nodes and three compute nodes.
4. Update the labels for the control plane and compute nodes.
5. Stop and restart the instances one by one on the Compute Cloud@Customer portal.

For the full procedure, see the "Terraform Script Execution - Part 2" section in the [Oracle documentation](#).

5.6. COMPLETING THE INSTALLATION BY USING THE ASSISTED INSTALLER WEB CONSOLE

After you configure the infrastructure, the instances are now running and are ready to be registered with Red Hat.

5.6.1. Assigning node roles

If the Terraform scripts completed successfully, twelve hosts are now listed for the cluster. Three control plane hosts and three compute hosts have the status "Disconnected". Three control plane hosts and three compute hosts have the status "Insufficient".

Delete the disconnected hosts and assign roles to the remaining hosts.

Procedure

1. From the [Assisted Installer web console](#), select the cluster and navigate to the **Host discovery** page.
2. Delete the six hosts with a "Disconnected" status, by clicking the option button for each host and selecting **Remove host**. The status of the remaining hosts changes from "Insufficient" to "Ready". This process can take up to three minutes.
3. From the **Role** column, assign the **Control plane** role to the three nodes with a boot size of 1.10 TB. Assign the **Worker** role to the three nodes with boot size of 100 GB.
4. Rename any hosts with a name shorter than 63 characters, by clicking the option button for the host and selecting **Change hostname**. Otherwise the cluster installation will fail.
5. Click **Next**.
6. On the **Storage** page, click **Next**.

5.6.2. Configuring networking

On the **Networking** page, add the NTP sources for any hosts that display the **Some validations failed** status.

Procedure

1. In the **Host inventory** table, click the **Some validations failed** link for each host displaying this status.
2. Click **Add NTP sources**, and then add the IP address **169.254.169.254** for one of the nodes.
3. Wait for 2 - 3 minutes until all the **Some validations failed** indicators disappear.
4. Select **Next**.

5.6.3. Adding custom manifests

Create, modify, and upload the four mandatory custom manifests provided by Oracle.

- In the **C3/custom_manifests_C3/manifests** folder, the following manifests are mandatory:
 - **oci-ccm.yml**
 - **oci-csi.yml**
- In the **C3/custom_manifests_C3/openshift** folder, the following manifests are mandatory:
 - **machineconfig-ccm.yml**
 - **machineconfig-csi.yml**

Prerequisites

- Prepare the custom manifests. For details, see step 8 in the "Install the Cluster using the RH Assisted Installer UI" section of the [Oracle documentation](#).

Procedure

1. Navigate to the **Custom manifests** page.
2. Upload and save the **oci-ccm.yml** and **oci-csi.yml** manifest files:
 - a. In the **Folder** field, select **manifests**.
 - b. In the **File name** field, enter **oci-ccm.yml**.
 - c. In the **Content** section, click **Browse**.
 - d. Select the **oci-ccm.yml** file from the **C3/custom_manifests_C3/manifests** folder.
 - e. Click **Add another manifest** and repeat the previous substeps for the **oci-csi.yml** file.
3. Upload and save the **machineconfig-ccm.yml** and **machineconfig-csi.yml** manifest files:
 - a. Click **Add another manifest**.

- b. In the **Folder** field, select **openshift**.
 - c. In the **File name** field, enter **machineconfig-ccm.yml**.
 - d. In the **Content** section, click **Browse**.
 - e. Select the **machineconfig-ccm.yml** file from the **C3/custom_ manifest_C3/openshift** folder.
 - f. Click **Add another manifest** and repeat the previous substeps for the **machineconfig-csi.yml** file.
4. Click **Next** to save the custom manifests.
 5. From the **Review and create** page, click **Install cluster** to create your OpenShift Container Platform cluster. This process takes approximately thirty minutes.

5.7. OPENING OPENSHIFT CONTAINER PLATFORM FROM THE ORACLE COMPUTE CLOUD@CUSTOMER WEB CONSOLE

For instructions to access the OpenShift Container Platform console from Oracle Compute Cloud@Customer, see step 15 - 17 in the "Install the Cluster using the RH Assisted Installer UI" section of the [Oracle documentation](#).