



OpenShift Container Platform 4.20

Installing on Oracle Edge Cloud

Installing OpenShift Container Platform on Oracle Cloud Edge

OpenShift Container Platform 4.20 Installing on Oracle Edge Cloud

Installing OpenShift Container Platform on Oracle Cloud Edge

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 Edge Cloud Infrastructure.

Table of Contents

CHAPTER 1. INSTALLING A CLUSTER ON ORACLE EDGE CLOUD BY USING THE ASSISTED INSTALLER	3
1.1. SUPPORTED ORACLE EDGE CLOUD INFRASTRUCTURES	3
1.2. OVERVIEW	3
1.3. PREPARING THE OCI BASTION SERVER	4
1.4. RUNNING THE TERRAFORM SCRIPT VIA THE HOME REGION	4
1.5. PREPARING THE OCI IMAGE	5
1.5.1. Generating the image in the Assisted Installer	5
1.5.2. Converting and uploading the image to Oracle Edge Cloud	7
1.6. RUNNING THE TERRAFORM SCRIPT VIA THE C3 REGION	7
1.7. COMPLETING THE INSTALLATION BY USING THE ASSISTED INSTALLER WEB CONSOLE	7
1.7.1. Assigning node roles	8
1.7.2. Configuring networking	8
1.7.3. Adding custom manifests	8
1.8. OPENING OPENSIFT CONTAINER PLATFORM FROM THE ORACLE EDGE CLOUD WEB CONSOLE	9
CHAPTER 2. INSTALLING A CLUSTER ON ORACLE EDGE CLOUD BY USING THE AGENT-BASED INSTALLER	10
2.1. SUPPORTED ORACLE EDGE CLOUD INFRASTRUCTURES	10
2.2. INSTALLATION PROCESS WORKFLOW	10
2.3. CREATING OCI INFRASTRUCTURE RESOURCES AND SERVICES	10
2.4. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON ORACLE EDGE CLOUD	11
2.5. CONFIGURING YOUR FIREWALL FOR OPENSIFT CONTAINER PLATFORM	15
2.6. RUNNING A CLUSTER ON ORACLE EDGE CLOUD	21
2.7. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON ORACLE EDGE CLOUD	22
2.8. ADDITIONAL RESOURCES	23

CHAPTER 1. INSTALLING A CLUSTER ON ORACLE EDGE CLOUD BY USING THE ASSISTED INSTALLER

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

The following procedures describe a cluster installation on Oracle® Compute Cloud@Customer as an example.

1.1. SUPPORTED ORACLE EDGE CLOUD INFRASTRUCTURES

The following table describes the support status of each Oracle® Edge Cloud infrastructure offering:

Table 1.1. Oracle Edge Cloud infrastructure support statuses

Infrastructure type	Support status
Private Cloud Appliance	General Availability
Oracle Compute Cloud@Customer	General Availability
Roving Edge	Technology Preview

1.2. OVERVIEW

You can install OpenShift Container Platform on Oracle Edge Cloud by using the Assisted Installer.

For an alternative installation method, see "Installing a cluster on Oracle® Edge Cloud 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 Oracle® Edge Cloud system (two Dynamic Groups and one Policy).
- The second script builds the infrastructure resources on the Oracle® Edge Cloud 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 Oracle® Edge Cloud DRG Subnet or Internet GW Subnet within the Oracle® Edge Cloud 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 Edge Cloud.

Main steps

The main steps of the procedure are as follows:

1. Preparing the Oracle® Edge Cloud bastion server.
2. Running the Terraform script via the Home region.
3. Preparing the OpenShift Container Platform image for Oracle Edge Cloud.
4. Running the Terraform script via the Oracle® Edge Cloud region.
5. Installing the cluster by using the Assisted Installer web console.

1.3. PREPARING THE OCI BASTION SERVER

By implementing a bastion host, you can securely and efficiently manage access to your Oracle Cloud Infrastructure (OCI) 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 \(Oracle documentation\)](#).

1.4. 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 Oracle Cloud Infrastructure (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](#).

1.5. 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 Edge Cloud compatible image and upload it to the **Custom Images** page of your Oracle Edge Cloud 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.

1.5.1. 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.

Field	Action required
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 Oracle® Edge Cloud. 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 Edge Cloud, 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**.

1.5.2. Converting and uploading the image to Oracle Edge Cloud

Convert the ISO image to an Oracle Cloud Infrastructure (OCI) image and upload it to your Oracle Edge Cloud 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 Oracle® Edge Cloud 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](#).

1.6. RUNNING THE TERRAFORM SCRIPT VIA THE C3 REGION

Run the **terraform.tfvars** Terraform script to create all infrastructure resources on Oracle® Edge Cloud. 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 Oracle® Edge Cloud portal.

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

1.7. 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.

1.7.1. 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**.

1.7.2. 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**.

1.7.3. 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_manifest_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.

1.8. OPENING OPENSIFT CONTAINER PLATFORM FROM THE ORACLE EDGE CLOUD WEB CONSOLE

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

CHAPTER 2. INSTALLING A CLUSTER ON ORACLE EDGE CLOUD BY USING THE AGENT-BASED INSTALLER

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

The following procedures describe a cluster installation on Oracle® Compute Cloud@Customer as an example.

2.1. SUPPORTED ORACLE EDGE CLOUD INFRASTRUCTURES

The following table describes the support status of each Oracle® Edge Cloud infrastructure offering:

Table 2.1. Oracle Edge Cloud infrastructure support statuses

Infrastructure type	Support status
Private Cloud Appliance	General Availability
Oracle Compute Cloud@Customer	General Availability
Roving Edge	Technology Preview

2.2. INSTALLATION PROCESS WORKFLOW

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

1. Create Oracle Cloud Infrastructure (OCI) 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 OCI image, upload it to an OCI Home Region Bucket, and then import the uploaded image to the Oracle Edge Cloud system (Oracle).
5. Disconnected environments: Prepare a web server that is accessible by Oracle Edge Cloud 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 Oracle Edge Cloud (Oracle).

2.3. CREATING OCI INFRASTRUCTURE RESOURCES AND SERVICES

You must create an Oracle Edge Cloud 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 Oracle Cloud Infrastructure (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 OCI 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)

2.4. CREATING CONFIGURATION FILES FOR INSTALLING A CLUSTER ON ORACLE EDGE CLOUD

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 4.20.0
built from commit ae7977b7d1ca908674a0d45c5c243c766fa4b2ca
release image registry.ci.openshift.org/origin/release:4.20ocp-
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> ❷
serviceNetwork:
- 172.30.0.0/16
compute:
- architecture: amd64 ❸
hyperthreading: Enabled
name: worker
replicas: 0
controlPlane:
architecture: amd64 ❹
hyperthreading: Enabled
name: master
replicas: 3
platform:
external:
platformName: oci ❺
cloudControllerManager: External
sshKey: <public_ssh_key> ❻
pullSecret: '<pull_secret>' ❼
# ...

```

- ❶ The base domain of your cloud provider.
- ❷ The IP address from the virtual cloud network (VCN) that the CIDR allocates to resources and components that operate on your network.
- ❸ ❹ Depending on your infrastructure, you can select either **arm64** or **amd64**.
- ❺ Set **OCI** as the external platform, so that OpenShift Container Platform can integrate with OCI.
- ❻ Specify your SSH public key.
- ❼ 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.yml** and **oci-csi.yml** 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> ❶
  namespace: <cluster_namespace> ❷
rendezvousIP: <ip_address_from_CIDR> ❸
bootArtifactsBaseURL: <server_URL> ❹
# ...
```

- ❶ The cluster name that you specified in your DNS record.
- ❷ The namespace of your cluster on OpenShift Container Platform.
- ❸ 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.
- ❹ 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](#)

2.5. CONFIGURING YOUR FIREWALL FOR OPENSIFT 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.

URL	Port	Function
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 Lightspeed:

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.

Cloud	URL	Port	Function
AWS	aws.amazon.com	443	Used to install and manage clusters in an AWS environment.
	*.amazonaws.com Alternatively, if you choose to not use a wildcard for AWS APIs, you must include the following URLs in your allowlist:	443	Required to access AWS services and resources. Review the AWS Service Endpoints in the AWS documentation to find the exact endpoints to allow for the regions that you use.
	ec2.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	events.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	iam.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	route53.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	*.s3.dualstack. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	sts.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	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.

Cloud	URL	Port	Function
	elasticloadbalancing. <aws_region>.amazonaws.com	443	Used to install and manage clusters in an AWS environment.
	servicequotas. <aws_region>.amazonaws.com	443	Required. Used to confirm quotas for deploying the service.
	tagging. <aws_region>.amazonaws.com	443	Allows the assignment of metadata about AWS resources in the form of tags.
	*.cloudfront.net	443	Used to provide access to CloudFront. If you use the AWS Security Token Service (STS) and the private S3 bucket, you must provide access to CloudFront.
GCP	*.googleapis.com	443	Required to access 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
-----	------	----------

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.
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.
rhcos.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.

2.6. RUNNING A CLUSTER ON ORACLE EDGE CLOUD

To run a cluster on Oracle® Edge Cloud, 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 Oracle Edge Cloud system.



NOTE

Oracle Edge Cloud 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 Oracle Edge Cloud" 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 Oracle Edge Cloud 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 Oracle Edge Cloud. 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.

2.7. VERIFYING THAT YOUR AGENT-BASED CLUSTER INSTALLATION RUNS ON ORACLE EDGE CLOUD

Verify that your cluster was installed and is running effectively on Oracle Edge Cloud.

Prerequisites

- You created all the required Oracle Cloud Infrastructure (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 Oracle Edge Cloud" section.
- You uploaded the agent ISO image to a default Oracle Object Storage bucket, and you created a compute instance on Oracle Edge Cloud. For more information, see "Running a cluster on Oracle Edge Cloud".

Procedure

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

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

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

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

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

```
$ export KUBECONFIG=~/.kube/kubeconfig
```

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

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

Output example

NAME	STATUS	ROLES	AGE	VERSION
------	--------	-------	-----	---------

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

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

```
$ oc get co
```

Truncated output example

```
NAME          VERSION   AVAILABLE PROGRESSING  DEGRADED  SINCE
MESSAGE
authentication 4.20.0-0 True      False        False     6m18s
baremetal      4.20.0-0 True      False        False     2m42s
network        4.20.0-0 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](#)