



# OpenShift Container Platform 4.12

## Installing on Nutanix

Installing OpenShift Container Platform on Nutanix



# OpenShift Container Platform 4.12 Installing on Nutanix

---

Installing OpenShift Container Platform on Nutanix

## Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This document describes how to install OpenShift Container Platform on Nutanix.

## Table of Contents

<b>CHAPTER 1. PREPARING TO INSTALL ON NUTANIX</b> .....	<b>4</b>
1.1. NUTANIX VERSION REQUIREMENTS	4
1.2. ENVIRONMENT REQUIREMENTS	4
1.2.1. Required account privileges	4
1.2.2. Cluster limits	5
1.2.3. Cluster resources	5
1.2.4. Networking requirements	6
1.2.4.1. Required IP Addresses	6
1.2.4.2. DNS records	6
1.3. CONFIGURING THE CLOUD CREDENTIAL OPERATOR UTILITY	7
<b>CHAPTER 2. INSTALLING A CLUSTER ON NUTANIX</b> .....	<b>9</b>
2.1. PREREQUISITES	9
2.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM	9
2.3. INTERNET ACCESS FOR PRISM CENTRAL	10
2.4. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	10
2.5. OBTAINING THE INSTALLATION PROGRAM	12
2.6. ADDING NUTANIX ROOT CA CERTIFICATES TO YOUR SYSTEM TRUST	13
2.7. CREATING THE INSTALLATION CONFIGURATION FILE	13
2.7.1. Installation configuration parameters	14
2.7.1.1. Required configuration parameters	15
2.7.1.2. Network configuration parameters	16
2.7.1.3. Optional configuration parameters	18
2.7.1.4. Additional Nutanix configuration parameters	23
2.7.2. Sample customized install-config.yaml file for Nutanix	25
2.7.3. Configuring the cluster-wide proxy during installation	27
2.8. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	28
Installing the OpenShift CLI on Linux	29
Installing the OpenShift CLI on Windows	29
Installing the OpenShift CLI on macOS	30
2.9. CONFIGURING IAM FOR NUTANIX	30
2.10. DEPLOYING THE CLUSTER	33
2.11. CONFIGURING THE DEFAULT STORAGE CONTAINER	35
2.12. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	35
2.13. ADDITIONAL RESOURCES	35
2.14. NEXT STEPS	35
<b>CHAPTER 3. INSTALLING A CLUSTER ON NUTANIX IN A RESTRICTED NETWORK</b> .....	<b>36</b>
3.1. PREREQUISITES	36
3.2. ABOUT INSTALLATIONS IN RESTRICTED NETWORKS	36
3.2.1. Additional limits	37
3.3. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS	37
3.4. ADDING NUTANIX ROOT CA CERTIFICATES TO YOUR SYSTEM TRUST	39
3.5. DOWNLOADING THE RHCOS CLUSTER IMAGE	39
3.6. CREATING THE INSTALLATION CONFIGURATION FILE	40
3.6.1. Installation configuration parameters	42
3.6.1.1. Required configuration parameters	43
3.6.1.2. Network configuration parameters	44
3.6.1.3. Optional configuration parameters	46
3.6.1.4. Additional Nutanix configuration parameters	51
3.6.2. Sample customized install-config.yaml file for Nutanix	53

3.6.3. Configuring the cluster-wide proxy during installation	55
3.7. INSTALLING THE OPENSIFT CLI BY DOWNLOADING THE BINARY	57
Installing the OpenShift CLI on Linux	57
Installing the OpenShift CLI on Windows	58
Installing the OpenShift CLI on macOS	58
3.8. CONFIGURING IAM FOR NUTANIX	59
3.9. DEPLOYING THE CLUSTER	62
3.10. POST INSTALLATION	63
3.10.1. Disabling the default OperatorHub catalog sources	63
3.10.2. Installing the policy resources into the cluster	63
3.10.3. Configuring the default storage container	64
3.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM	64
3.12. ADDITIONAL RESOURCES	65
3.13. NEXT STEPS	65
<b>CHAPTER 4. UNINSTALLING A CLUSTER ON NUTANIX</b> .....	<b>66</b>
4.1. REMOVING A CLUSTER THAT USES INSTALLER-PROVISIONED INFRASTRUCTURE	66



# CHAPTER 1. PREPARING TO INSTALL ON NUTANIX

Before you install an OpenShift Container Platform cluster, be sure that your Nutanix environment meets the following requirements.

## 1.1. NUTANIX VERSION REQUIREMENTS

You must install the OpenShift Container Platform cluster to a Nutanix environment that meets the following requirements.

**Table 1.1. Version requirements for Nutanix virtual environments**

Component	Required version
Nutanix AOS	5.20.4+ or 6.5.1+
Prism Central	2022.4+

## 1.2. ENVIRONMENT REQUIREMENTS

Before you install an OpenShift Container Platform cluster, review the following Nutanix AOS environment requirements.

### 1.2.1. Required account privileges

The installation program requires access to a Nutanix account with the necessary permissions to deploy the cluster and to maintain the daily operation of it. The following options are available to you:

- You can use a local Prism Central user account with administrative privileges. Using a local account is the quickest way to grant access to an account with the required permissions.
- If your organization's security policies require that you use a more restrictive set of permissions, use the permissions that are listed in the following table to create a custom Cloud Native role in Prism Central. You can then assign the role to a user account that is a member of a Prism Central authentication directory. When assigning entities to the role, ensure that the user can access only the Prism Element and subnet that are required to deploy the virtual machines. For more information, see the Nutanix documentation about creating a [Custom Cloud Native role](#) and [assigning a role](#).

#### Example 1.1. Required permissions for creating a Custom Cloud Native role

Nutanix Object	Required permissions in Nutanix API	Description
----------------	-------------------------------------	-------------



Nutanix Object	Required permissions in Nutanix API	Description
Categories	<b>Create_Category_Mapping</b> <b>Create_Or_Update_Name_Category</b> <b>Create_Or_Update_Value_Category</b> <b>Delete_Category_Mapping</b> <b>Delete_Name_Category</b> <b>Delete_Value_Category</b> <b>View_Category_Mapping</b> <b>View_Name_Category</b> <b>View_Value_Category</b>	Create, read, and delete categories that are assigned to the OpenShift Container Platform machines.
Images	<b>Create_Image</b> <b>Delete_Image</b> <b>View_Image</b>	Create, read, and delete the operating system images used for the OpenShift Container Platform machines.
Virtual Machines	<b>Create_Virtual_Machine</b> <b>Delete_Virtual_Machine</b> <b>View_Virtual_Machine</b>	Create, read, and delete the OpenShift Container Platform machines.
Clusters	<b>View_Cluster</b>	View the Prism Element clusters that host the OpenShift Container Platform machines.
Subnets	<b>View_Subnet</b>	View the subnets that host the OpenShift Container Platform machines.

### 1.2.2. Cluster limits

Available resources vary between clusters. The number of possible clusters within a Nutanix environment is limited primarily by available storage space and any limitations associated with the resources that the cluster creates, and resources that you require to deploy the cluster, such as IP addresses and networks.

### 1.2.3. Cluster resources

A minimum of 800 GB of storage is required to use a standard cluster.

When you deploy a OpenShift Container Platform cluster that uses installer-provisioned infrastructure, the installation program must be able to create several resources in your Nutanix instance. Although these resources use 856 GB of storage, the bootstrap node is destroyed as part of the installation process.


A standard OpenShift Container Platform installation creates the following resources:

- 1 label
- Virtual machines:
  - 1 disk image
  - 1 temporary bootstrap node
  - 3 control plane nodes
  - 3 compute machines

## 1.2.4. Networking requirements

You must use AHV IP Address Management (IPAM) for the network and ensure that it is configured to provide persistent IP addresses to the cluster machines. Additionally, create the following networking resources before you install the OpenShift Container Platform cluster:

- IP addresses
- DNS records



### NOTE

It is recommended that each OpenShift Container Platform node in the cluster have access to a Network Time Protocol (NTP) server that is discoverable via DHCP. Installation is possible without an NTP server. However, an NTP server prevents errors typically associated with asynchronous server clocks.

### 1.2.4.1. Required IP Addresses

An installer-provisioned installation requires two static virtual IP (VIP) addresses:

- A VIP address for the API is required. This address is used to access the cluster API.
- A VIP address for ingress is required. This address is used for cluster ingress traffic.

You specify these IP addresses when you install the OpenShift Container Platform cluster.

### 1.2.4.2. DNS records

You must create DNS records for two static IP addresses in the appropriate DNS server for the Nutanix instance that hosts your OpenShift Container Platform cluster. In each record, **<cluster\_name>** is the cluster name and **<base\_domain>** is the cluster base domain that you specify when you install the cluster.

If you use your own DNS or DHCP server, you must also create records for each node, including the bootstrap, control plane, and compute nodes.

A complete DNS record takes the form: **<component>.<cluster\_name>.<base\_domain>..**

#### Table 1.2. Required DNS records

Component	Record	Description
API VIP	<b>api.&lt;cluster_name&gt;.&lt;base_domain&gt;.</b>	This DNS A/AAAA or CNAME record must point to the load balancer for the control plane machines. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.
Ingress VIP	<b>*.apps.&lt;cluster_name&gt;.&lt;base_domain&gt;.</b>	A wildcard DNS A/AAAA or CNAME record that points to the load balancer that targets the machines that run the Ingress router pods, which are the worker nodes by default. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

### 1.3. CONFIGURING THE CLOUD CREDENTIAL OPERATOR UTILITY

The Cloud Credential Operator (CCO) manages cloud provider credentials as Kubernetes custom resource definitions (CRDs). To install a cluster on Nutanix, you must set the CCO to **manual** mode as part of the installation process.

To create and manage cloud credentials from outside of the cluster when the Cloud Credential Operator (CCO) is operating in manual mode, extract and prepare the CCO utility (**ccoctl**) binary.



#### NOTE

The **ccoctl** utility is a Linux binary that must run in a Linux environment.

#### Prerequisites

- You have access to an OpenShift Container Platform account with cluster administrator access.
- You have installed the OpenShift CLI (**oc**).


#### Procedure

1. Obtain the OpenShift Container Platform release image by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

2. Obtain the CCO container image from the OpenShift Container Platform release image by running the following command:

```
$ CCO_IMAGE=$(oc adm release info --image-for='cloud-credential-operator'
$RELEASE_IMAGE -a ~/.pull-secret)
```

**NOTE**

Ensure that the architecture of the **\$RELEASE\_IMAGE** matches the architecture of the environment in which you will use the **ccoctl** tool.

3. Extract the **ccoctl** binary from the CCO container image within the OpenShift Container Platform release image by running the following command:

```
$ oc image extract $CCO_IMAGE --file="/usr/bin/ccoctl" -a ~/.pull-secret
```

4. Change the permissions to make **ccoctl** executable by running the following command:

```
$ chmod 775 ccoctl
```

**Verification**

- To verify that **ccoctl** is ready to use, display the help file by running the following command:

```
$ ccoctl --help
```

**Output of ccoctl --help**

```
OpenShift credentials provisioning tool
```

```
Usage:
```

```
ccoctl [command]
```

```
Available Commands:
```

```
alibabacloud Manage credentials objects for alibaba cloud
aws          Manage credentials objects for AWS cloud
gcp         Manage credentials objects for Google cloud
help       Help about any command
ibmcloud    Manage credentials objects for IBM Cloud
nutanix     Manage credentials objects for Nutanix
```

```
Flags:
```

```
-h, --help  help for ccoctl
```

```
Use "ccoctl [command] --help" for more information about a command.
```

**Additional resources**

- [Preparing to update a cluster with manually maintained credentials](#)

## CHAPTER 2. INSTALLING A CLUSTER ON NUTANIX

In OpenShift Container Platform version 4.12, you can choose one of the following options to install a cluster on your Nutanix instance:

**Using installer-provisioned infrastructure:** Use the procedures in the following sections to use installer-provisioned infrastructure. Installer-provisioned infrastructure is ideal for installing in connected or disconnected network environments. The installer-provisioned infrastructure includes an installation program that provisions the underlying infrastructure for the cluster.

**Using the Assisted Installer:** The [Assisted Installer](#) hosted at [console.redhat.com](https://console.redhat.com). The Assisted Installer cannot be used in disconnected environments. The Assisted Installer does not provision the underlying infrastructure for the cluster, so you must provision the infrastructure before the running the Assisted Installer. Installing with the Assisted Installer also provides integration with Nutanix, enabling autoscaling. See [Installing an on-premise cluster using the Assisted Installer](#) for additional details.

**Using user-provisioned infrastructure:** Complete the relevant steps outlined in the [Installing a cluster on any platform](#) documentation.

### 2.1. PREREQUISITES

- You have reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- The installation program requires access to port 9440 on Prism Central and Prism Element. You verified that port 9440 is accessible.
- If you use a firewall, you have met these prerequisites:
  - You confirmed that port 9440 is accessible. Control plane nodes must be able to reach Prism Central and Prism Element on port 9440 for the installation to succeed.
  - You configured the firewall to [grant access](#) to the sites that OpenShift Container Platform requires. This includes the use of Telemetry.
- If your Nutanix environment is using the default self-signed SSL certificate, replace it with a certificate that is signed by a CA. The installation program requires a valid CA-signed certificate to access to the Prism Central API. For more information about replacing the self-signed certificate, see the [Nutanix AOS Security Guide](#).

If your Nutanix environment uses an internal CA to issue certificates, you must configure a cluster-wide proxy as part of the installation process. For more information, see [Configuring a custom PKI](#).



#### IMPORTANT


Use 2048-bit certificates. The installation fails if you use 4096-bit certificates with Prism Central 2022.x.

### 2.2. INTERNET ACCESS FOR OPENSIFT CONTAINER PLATFORM

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

You must have internet access to:

- Access [OpenShift Cluster Manager Hybrid Cloud Console](#) to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



### IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the required content and use it to populate a mirror registry with the installation packages. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

## 2.3. INTERNET ACCESS FOR PRISM CENTRAL

Prism Central requires internet access to obtain the Red Hat Enterprise Linux CoreOS (RHCOS) image that is required to install the cluster. The RHCOS image for Nutanix is available at [rhcos.mirror.openshift.com](https://rhcos.mirror.openshift.com).

## 2.4. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



### IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



### NOTE

You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

### Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



#### NOTE

If you plan to install an OpenShift Container Platform cluster that uses FIPS validated or Modules In Process cryptographic libraries on the **x86\_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



#### NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

#### Example output

```
Agent pid 31874
```



#### NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

## Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

## Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

## 2.5. OBTAINING THE INSTALLATION PROGRAM

Before you install OpenShift Container Platform, download the installation file on the host you are using for installation.

## Prerequisites

- You have a computer that runs Linux or macOS, with 500 MB of local disk space.

## Procedure

1. Access the [Infrastructure Provider](#) page on the OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Select your infrastructure provider.
3. Navigate to the page for your installation type, download the installation program that corresponds with your host operating system and architecture, and place the file in the directory where you will store the installation configuration files.



### IMPORTANT

The installation program creates several files on the computer that you use to install your cluster. You must keep the installation program and the files that the installation program creates after you finish installing the cluster. Both files are required to delete the cluster.



### IMPORTANT

Deleting the files created by the installation program does not remove your cluster, even if the cluster failed during installation. To remove your cluster, complete the OpenShift Container Platform uninstallation procedures for your specific cloud provider.

4. Extract the installation program. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -xvf openshift-install-linux.tar.gz
```

5. Download your installation [pull secret from the Red Hat OpenShift Cluster Manager](#) . This pull secret allows you to authenticate with the services that are provided by the included authorities, including Quay.io, which serves the container images for OpenShift Container Platform components.



## 2.6. ADDING NUTANIX ROOT CA CERTIFICATES TO YOUR SYSTEM TRUST

Because the installation program requires access to the Prism Central API, you must add your Nutanix trusted root CA certificates to your system trust before you install an OpenShift Container Platform cluster.

### Procedure

1. From the Prism Central web console, download the Nutanix root CA certificates.
2. Extract the compressed file that contains the Nutanix root CA certificates.
3. Add the files for your operating system to the system trust. For example, on a Fedora operating system, run the following command:

```
# cp certs/lin/* /etc/pki/ca-trust/source/anchors
```

4. Update your system trust. For example, on a Fedora operating system, run the following command:

```
# update-ca-trust extract
```

## 2.7. CREATING THE INSTALLATION CONFIGURATION FILE

You can customize the OpenShift Container Platform cluster you install on Nutanix.

### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify that you have met the Nutanix networking requirements. For more information, see "Preparing to install on Nutanix".

### Procedure

1. Create the **install-config.yaml** file.
  - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1** For **<installation\_directory>**, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
- Use an empty directory. Some installation assets, such as bootstrap X.509 certificates,

have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.


- b. At the prompts, provide the configuration details for your cloud:
  - i. Optional: Select an SSH key to use to access your cluster machines.



#### NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **nutanix** as the platform to target.
  - iii. Enter the Prism Central domain name or IP address.
  - iv. Enter the port that is used to log into Prism Central.
  - v. Enter the credentials that are used to log into Prism Central.  
The installation program connects to Prism Central.
  - vi. Select the Prism Element that will manage the OpenShift Container Platform cluster.
  - vii. Select the network subnet to use.
  - viii. Enter the virtual IP address that you configured for control plane API access.
  - ix. Enter the virtual IP address that you configured for cluster ingress.
  - x. Enter the base domain. This base domain must be the same one that you configured in the DNS records.
  - xi. Enter a descriptive name for your cluster. The cluster name you enter must match the cluster name you specified when configuring the DNS records.
  - xii. Paste the [pull secret from the Red Hat OpenShift Cluster Manager](#) .
2. Optional: Update one or more of the default configuration parameters in the **install.config.yaml** file to customize the installation.  
For more information about the parameters, see "Installation configuration parameters".
  3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



#### IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

### 2.7.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's

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



## NOTE

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

### 2.7.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 2.1. Required parameters

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

Parameter	Description	Values
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <code>{}</code> . For additional information about <b>platform</b> . <code>&lt;platform&gt;</code> parameters, consult the table for your specific platform that follows.	Object
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{   "auths":{     "cloud.openshift.com":{       "auth":"b3Blb=",       "email":"you@example.com"     },     "quay.io":{       "auth":"b3Blb=",       "email":"you@example.com"     }   } }</pre>

### 2.7.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

Only IPv4 addresses are supported.




#### NOTE

Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 2.2. Network parameters

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

Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object  <b>NOTE</b> You cannot modify parameters specified by the <b>networking</b> object after installation.
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . An IP address block.  An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>


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


### 2.7.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 2.3. Optional parameters

Parameter	Description	Values
<b>additionalTrustBundle</b>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<b>capabilities</b>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String



Parameter	Description	Values
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.
<b>compute.architecture</b>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are <b>amd64</b> (the default).	String
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b>  If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .

Parameter	Description	Values
<b>featureSet</b>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are <b>amd64</b> (the default).	String
<b>controlPlane.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b> If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.



Parameter	Description	Values
<b>credentialsMode</b>	<p>The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.</p> <div data-bbox="485 510 595 891" style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> <div data-bbox="485 936 595 1288" style="border: 1px solid black; padding: 5px;"> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<b>fips</b>	<p>Enable or disable FIPS mode. The default is <b>false</b> (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 20px; height: 100px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, black 2px, black 4px); margin-right: 10px;"></div> <div> <p><b>IMPORTANT</b></p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="width: 20px; height: 50px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, black 2px, black 4px); margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div> </div>	<b>false</b> or <b>true</b>
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String

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

#### 2.7.1.4. Additional Nutanix configuration parameters

Additional Nutanix configuration parameters are described in the following table:

Table 2.4. Additional Nutanix cluster parameters

Parameter	Description	Values
<b>platform.nutanix.apiVIP</b>	The virtual IP (VIP) address that you configured for control plane API access.	IP address
<b>platform.nutanix.ingressVIP</b>	The virtual IP (VIP) address that you configured for cluster ingress.	IP address

Parameter	Description	Values
<b>platform.nutanix.pris mCentral.endpoint.a ddress</b>	The Prism Central domain name or IP address.	String
<b>platform.nutanix.pris mCentral.endpoint.p ort</b>	The port that is used to log into Prism Central.	String
<b>platform.nutanix.pris mCentral.password</b>	The password for the Prism Central user name.	String
<b>platform.nutanix.pris mCentral.username</b>	The user name that is used to log into Prism Central.	String
<b>platform.nutanix.pris mElements.endpoint. address</b>	The Prism Element domain name or IP address. [1]	String
<b>platform.nutanix.pris mElements.endpoint. port</b>	The port that is used to log into Prism Element.	String
<b>platform.nutanix.pris mElements.uuid</b>	The universally unique identifier (UUID) for Prism Element.	String
<b>platform.nutanix.sub netUUIDs</b>	The UUID of the Prism Element network that contains the virtual IP addresses and DNS records that you configured. [2]	String
<b>platform.nutanix.clu sterOSImage</b>	Optional: By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image. If Prism Central does not have internet access, you can override the default behavior by hosting the RHCOS image on any HTTP server and pointing the installation program to the image.	An HTTP or HTTPS URL, optionally with a SHA-256 checksum. For example, <code>http://example.com/images/rhcos-47.83.202103221318-0-nutanix.x86_64.qcow2</code>

1. The **prismElements** section holds a list of Prism Elements (clusters). A Prism Element encompasses all of the Nutanix resources, for example virtual machines and subnets, that are used to host the OpenShift Container Platform cluster. Only a single Prism Element is supported.
2. Only one subnet per OpenShift Container Platform cluster is supported.

## 2.7.2. Sample customized install-config.yaml file for Nutanix

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



### IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 3
  platform:
    nutanix: 4
      cpus: 2
      coresPerSocket: 2
      memoryMiB: 8196
      osDisk:
        diskSizeGiB: 120
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3
  platform:
    nutanix: 7
      cpus: 4
      coresPerSocket: 2
      memoryMiB: 16384
      osDisk:
        diskSizeGiB: 120
metadata:
  creationTimestamp: null
  name: test-cluster 8
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 9
  serviceNetwork:
    - 172.30.0.0/16
platform:
  nutanix:
    apiVIP: 10.40.142.7 10
    ingressVIP: 10.40.142.8 11
  prismCentral:
    endpoint:
      address: your.prismcentral.domainname 12

```

```

port: 9440 13
password: <password> 14
username: <username> 15
prismElements:
- endpoint:
  address: your.prismelement.domainname
  port: 9440
  uuid: 0005b0f1-8f43-a0f2-02b7-3cecef193712
subnetUUIDs:
- c7938dc6-7659-453e-a688-e26020c68e43
clusterOSImage: http://example.com/images/rhcos-47.83.202103221318-0-nutanix.x86_64.qcow2
16
credentialsMode: Manual
publish: External
pullSecret: '{"auths": ...}' 17
fips: false 18
sshKey: ssh-ed25519 AAAA... 19

```

**1 8 10 11 12 13 14 15 17** Required. The installation program prompts you for this value.

**2 5** The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, -, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.

**3 6** Whether to enable or disable simultaneous multithreading, or **hyperthreading**. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores. You can disable it by setting the parameter value to **Disabled**. If you disable simultaneous multithreading in some cluster machines, you must disable it in all cluster machines.



### IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.

**4 7** Optional: Provide additional configuration for the machine pool parameters for the compute and control plane machines.

**9** The cluster network plugin to install. The supported values are **OVNKubernetes** and **OpenShiftSDN**. The default value is **OVNKubernetes**.

**16** Optional: By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image. If Prism Central does not have internet access, you can override the default behavior by hosting the RHCOS image on any HTTP server and pointing the installation program to the image.

**18** Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.

**IMPORTANT**

The use of FIPS Validated or Modules in Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the **x86\_64**, **ppc64le**, and **s390x** architectures.

**19**

Optional: You can provide the **sshKey** value that you use to access the machines in your cluster.

**NOTE**


For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

### 2.7.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

#### Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.

**NOTE**

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

#### Procedure


1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
additionalTrustBundle: | 4
-----BEGIN CERTIFICATE-----
```

```
<MY_TRUSTED_CA_CERT>
-----END CERTIFICATE-----
```

```
additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1** A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2** A proxy URL to use for creating HTTPS connections outside the cluster.
- 3** A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use **\*** to bypass the proxy for all destinations.
- 4** If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5** Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



#### NOTE

The installation program does not support the proxy **readinessEndpoints** field.



#### NOTE

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

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

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.




#### NOTE

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

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



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



### IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.12. Download and install the new version of **oc**.

## Installing the OpenShift CLI on Linux

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

### Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.12 Linux Client** entry and save the file.
5. Unpack the archive:

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

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

```
$ echo $PATH
```

### Verification

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

```
$ oc <command>
```

## Installing the OpenShift CLI on Windows

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

### Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.12 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.  
To check your **PATH**, open the command prompt and execute the following command:

—

```
C:\> path
```

### Verification

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

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

### Installing the OpenShift CLI on macOS

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

#### Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.12 macOS Client** entry and save the file.



#### NOTE

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

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

```
$ echo $PATH
```

### Verification

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

```
$ oc <command>
```

## 2.9. CONFIGURING IAM FOR NUTANIX

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets.

### Prerequisites

- You have configured the **ccoctl** binary.
- You have an **install-config.yaml** file.

### Procedure

1. Create a YAML file that contains the credentials data in the following format:

## Credentials data format

```
credentials:
- type: basic_auth 1
  data:
    prismCentral: 2
      username: <username_for_prism_central>
      password: <password_for_prism_central>
    prismElements: 3
      - name: <name_of_prism_element>
        username: <username_for_prism_element>
        password: <password_for_prism_element>
```

**1** Specify the authentication type. Only basic authentication is supported.

**2** Specify the Prism Central credentials.

**3** Optional: Specify the Prism Element credentials.

2. Set a **\$RELEASE\_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

3. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
  --cloud=nutanix \
  --to=<path_to_directory_with_list_of_credentials_requests>/credrequests 1
```

**1** Specify the path to the directory that contains the files for the component **CredentialsRequests** objects. If the specified directory does not exist, this command creates it.

## Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  annotations:
    include.release.openshift.io/self-managed-high-availability: "true"
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-machine-api-nutanix
  namespace: openshift-cloud-credential-operator
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: NutanixProviderSpec
```

```
secretRef:
  name: nutanix-credentials
  namespace: openshift-machine-api
```

- If your cluster uses cluster capabilities to disable one or more optional components, delete the **CredentialsRequest** custom resources for any disabled components.

### Example credrequests directory contents for OpenShift Container Platform 4.12 on Nutanix

```
0000_30_machine-api-operator_00_credentials-request.yaml 1
```

- The Machine API Operator CR is required.

- Use the **ccoctl** tool to process all of the **CredentialsRequest** objects in the **credrequests** directory by running the following command:

```
$ ccoctl nutanix create-shared-secrets \
  --credentials-requests-dir=
  <path_to_directory_with_list_of_credentials_requests>/credrequests 1
  --output-dir=<ccoctl_output_dir> \ 2
  --credentials-source-filepath=<path_to_credentials_file> 3
```

- Specify the path to the directory that contains the files for the component **CredentialsRequests** objects.
- Specify the directory that contains the files of the component credentials secrets, under the **manifests** directory. By default, the **ccoctl** tool creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag.
- Optional: Specify the directory that contains the credentials data YAML file. By default, **ccoctl** expects this file to be in **<home\_directory>/.**nutanix/credentials****. To specify a different directory, use the **--credentials-source-filepath** flag.

- Edit the **install-config.yaml** configuration file so that the **credentialsMode** parameter is set to **Manual**.

### Example install-config.yaml configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
...
```

- Add this line to set the **credentialsMode** parameter to **Manual**.

- Create the installation manifests by running the following command:

```
$ openshift-install create manifests --dir <installation_directory> 1
```

- 1 Specify the path to the directory that contains the **install-config.yaml** file for your cluster.
8. Copy the generated credential files to the target manifests directory by running the following command:

```
$ cp <ccoctl_output_dir>/manifests/*credentials.yaml ./<installation_directory>/manifests
```

### Verification

- Ensure that the appropriate secrets exist in the **manifests** directory.

```
$ ls ./<installation_directory>/manifests
```

### Example output

```
total 64
-rw-r----- 1 <user> <user> 2335 Jul  8 12:22 cluster-config.yaml
-rw-r----- 1 <user> <user>  161 Jul  8 12:22 cluster-dns-02-config.yml
-rw-r----- 1 <user> <user>  864 Jul  8 12:22 cluster-infrastructure-02-config.yml
-rw-r----- 1 <user> <user>  191 Jul  8 12:22 cluster-ingress-02-config.yml
-rw-r----- 1 <user> <user> 9607 Jul  8 12:22 cluster-network-01-crd.yml
-rw-r----- 1 <user> <user>  272 Jul  8 12:22 cluster-network-02-config.yml
-rw-r----- 1 <user> <user>  142 Jul  8 12:22 cluster-proxy-01-config.yaml
-rw-r----- 1 <user> <user>  171 Jul  8 12:22 cluster-scheduler-02-config.yml
-rw-r----- 1 <user> <user>  200 Jul  8 12:22 cvo-overrides.yaml
-rw-r----- 1 <user> <user>  118 Jul  8 12:22 kube-cloud-config.yaml
-rw-r----- 1 <user> <user> 1304 Jul  8 12:22 kube-system-configmap-root-ca.yaml
-rw-r----- 1 <user> <user> 4090 Jul  8 12:22 machine-config-server-tls-secret.yaml
-rw-r----- 1 <user> <user> 3961 Jul  8 12:22 openshift-config-secret-pull-secret.yaml
-rw----- 1 <user> <user>  283 Jul  8 12:24 openshift-machine-api-nutanix-credentials-credentials.yaml
```

## 2.10. DEPLOYING THE CLUSTER

You can install OpenShift Container Platform on a compatible cloud platform.



### IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

### Prerequisites


- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

### Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation\_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.



## NOTE

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

## Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation\_directory>/openshift\_install.log**.



## IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

## Example output

```
...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s
```



## IMPORTANT

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

## 2.11. CONFIGURING THE DEFAULT STORAGE CONTAINER

After you install the cluster, you must install the Nutanix CSI Operator and configure the default storage container for the cluster.

For more information, see the Nutanix documentation for [installing the CSI Operator](#) and [configuring registry storage](#).

## 2.12. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

## 2.13. ADDITIONAL RESOURCES

- [About remote health monitoring](#)

## 2.14. NEXT STEPS

- [Opt out of remote health reporting](#)
- [Customize your cluster](#)

## CHAPTER 3. INSTALLING A CLUSTER ON NUTANIX IN A RESTRICTED NETWORK

In OpenShift Container Platform 4.12, you can install a cluster on Nutanix infrastructure in a restricted network by creating an internal mirror of the installation release content.

### 3.1. PREREQUISITES

- You have reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- The installation program requires access to port 9440 on Prism Central and Prism Element. You verified that port 9440 is accessible.
- If you use a firewall, you have met these prerequisites:
  - You confirmed that port 9440 is accessible. Control plane nodes must be able to reach Prism Central and Prism Element on port 9440 for the installation to succeed.
  - You configured the firewall to [grant access](#) to the sites that OpenShift Container Platform requires. This includes the use of Telemetry.
- If your Nutanix environment is using the default self-signed SSL/TLS certificate, replace it with a certificate that is signed by a CA. The installation program requires a valid CA-signed certificate to access to the Prism Central API. For more information about replacing the self-signed certificate, see the [Nutanix AOS Security Guide](#).  
If your Nutanix environment uses an internal CA to issue certificates, you must configure a cluster-wide proxy as part of the installation process. For more information, see [Configuring a custom PKI](#).



#### IMPORTANT

Use 2048-bit certificates. The installation fails if you use 4096-bit certificates with Prism Central 2022.x.

- You have a container image registry, such as Red Hat Quay. If you do not already have a registry, you can create a mirror registry using [mirror registry for Red Hat OpenShift](#).
- You have used the [oc-mirror OpenShift CLI \(oc\) plugin](#) to mirror all of the required OpenShift Container Platform content and other images, including the Nutanix CSI Operator, to your mirror registry.



#### IMPORTANT

Because the installation media is on the mirror host, you can use that computer to complete all installation steps.

### 3.2. ABOUT INSTALLATIONS IN RESTRICTED NETWORKS

In OpenShift Container Platform 4.12, you can perform an installation that does not require an active connection to the internet to obtain software components. Restricted network installations can be completed using installer-provisioned infrastructure or user-provisioned infrastructure, depending on the cloud platform to which you are installing the cluster.



If you choose to perform a restricted network installation on a cloud platform, you still require access to its cloud APIs. Some cloud functions, like Amazon Web Service's Route 53 DNS and IAM services, require internet access. Depending on your network, you might require less internet access for an installation on bare metal hardware, Nutanix, or on VMware vSphere.

To complete a restricted network installation, you must create a registry that mirrors the contents of the OpenShift image registry and contains the installation media. You can create this registry on a mirror host, which can access both the internet and your closed network, or by using other methods that meet your restrictions.

### 3.2.1. Additional limits

Clusters in restricted networks have the following additional limitations and restrictions:

- The **ClusterVersion** status includes an **Unable to retrieve available updates** error.
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required image stream tags.

## 3.3. GENERATING A KEY PAIR FOR CLUSTER NODE SSH ACCESS

During an OpenShift Container Platform installation, you can provide an SSH public key to the installation program. The key is passed to the Red Hat Enterprise Linux CoreOS (RHCOS) nodes through their Ignition config files and is used to authenticate SSH access to the nodes. The key is added to the `~/.ssh/authorized_keys` list for the **core** user on each node, which enables password-less authentication.

After the key is passed to the nodes, you can use the key pair to SSH in to the RHCOS nodes as the user **core**. To access the nodes through SSH, the private key identity must be managed by SSH for your local user.

If you want to SSH in to your cluster nodes to perform installation debugging or disaster recovery, you must provide the SSH public key during the installation process. The `./openshift-install gather` command also requires the SSH public key to be in place on the cluster nodes.



#### IMPORTANT

Do not skip this procedure in production environments, where disaster recovery and debugging is required.



#### NOTE


You must use a local key, not one that you configured with platform-specific approaches such as [AWS key pairs](#).

### Procedure

1. If you do not have an existing SSH key pair on your local machine to use for authentication onto your cluster nodes, create one. For example, on a computer that uses a Linux operating system, run the following command:

```
$ ssh-keygen -t ed25519 -N "" -f <path>/<file_name> 1
```

- 1 Specify the path and file name, such as `~/.ssh/id_ed25519`, of the new SSH key. If you have an existing key pair, ensure your public key is in the your `~/.ssh` directory.



#### NOTE

If you plan to install an OpenShift Container Platform cluster that uses FIPS validated or Modules In Process cryptographic libraries on the **x86\_64**, **ppc64le**, and **s390x** architectures, do not create a key that uses the **ed25519** algorithm. Instead, create a key that uses the **rsa** or **ecdsa** algorithm.

2. View the public SSH key:

```
$ cat <path>/<file_name>.pub
```

For example, run the following to view the `~/.ssh/id_ed25519.pub` public key:

```
$ cat ~/.ssh/id_ed25519.pub
```

3. Add the SSH private key identity to the SSH agent for your local user, if it has not already been added. SSH agent management of the key is required for password-less SSH authentication onto your cluster nodes, or if you want to use the `./openshift-install gather` command.



#### NOTE

On some distributions, default SSH private key identities such as `~/.ssh/id_rsa` and `~/.ssh/id_dsa` are managed automatically.

- a. If the **ssh-agent** process is not already running for your local user, start it as a background task:

```
$ eval "$(ssh-agent -s)"
```

#### Example output

```
Agent pid 31874
```



#### NOTE

If your cluster is in FIPS mode, only use FIPS-compliant algorithms to generate the SSH key. The key must be either RSA or ECDSA.

4. Add your SSH private key to the **ssh-agent**:

```
$ ssh-add <path>/<file_name> 1
```

- 1 Specify the path and file name for your SSH private key, such as `~/.ssh/id_ed25519`

#### Example output

```
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

### Next steps

- When you install OpenShift Container Platform, provide the SSH public key to the installation program.

## 3.4. ADDING NUTANIX ROOT CA CERTIFICATES TO YOUR SYSTEM TRUST

Because the installation program requires access to the Prism Central API, you must add your Nutanix trusted root CA certificates to your system trust before you install an OpenShift Container Platform cluster.

### Procedure

1. From the Prism Central web console, download the Nutanix root CA certificates.
2. Extract the compressed file that contains the Nutanix root CA certificates.
3. Add the files for your operating system to the system trust. For example, on a Fedora operating system, run the following command:

```
# cp certs/lin/* /etc/pki/ca-trust/source/anchors
```

4. Update your system trust. For example, on a Fedora operating system, run the following command:

```
# update-ca-trust extract
```

## 3.5. DOWNLOADING THE RHCOS CLUSTER IMAGE

Prism Central requires access to the Red Hat Enterprise Linux CoreOS (RHCOS) image to install the cluster. You can use the installation program to locate and download the RHCOS image and make it available through an internal HTTP server or Nutanix Objects.

### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.

### Procedure

1. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install coreos print-stream-json
```

2. Use the output of the command to find the location of the Nutanix image, and click the link to download it.

### Example output

```
"nutanix": {
  "release": "411.86.202210041459-0",
  "formats": {
    "qcow2": {
      "disk": {
        "location": "https://rhcos.mirror.openshift.com/art/storage/releases/rhcos-4.11/411.86.202210041459-0/x86_64/rhcos-411.86.202210041459-0-nutanix.x86_64.qcow2",
        "sha256":
"42e227cac6f11ac37ee8a2f9528bb3665146566890577fd55f9b950949e5a54b"
```

3. Make the image available through an internal HTTP server or Nutanix Objects.
4. Note the location of the downloaded image. You update the **platform** section in the installation configuration file (**install-config.yaml**) with the image's location before deploying the cluster.

### Snippet of an **install-config.yaml** file that specifies the RHCOS image

```
platform:
  nutanix:
    clusterOSImage: http://example.com/images/rhcos-411.86.202210041459-0-
    nutanix.x86_64.qcow2
```

## 3.6. CREATING THE INSTALLATION CONFIGURATION FILE

You can customize the OpenShift Container Platform cluster you install on Nutanix.

### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster. For a restricted network installation, these files are on your mirror host.
- Have the **imageContentSourcePolicy.yaml** file that was created when you mirrored your registry.
- Have the location of the Red Hat Enterprise Linux CoreOS (RHCOS) image you download.
- Obtain the contents of the certificate for your mirror registry.
- Retrieve a Red Hat Enterprise Linux CoreOS (RHCOS) image and upload it to an accessible location.
- Verify that you have met the Nutanix networking requirements. For more information, see "Preparing to install on Nutanix".

### Procedure

1. Create the **install-config.yaml** file.
  - a. Change to the directory that contains the installation program and run the following command:

```
$ ./openshift-install create install-config --dir <installation_directory> 1
```

- 1 For `<installation_directory>`, specify the directory name to store the files that the installation program creates.

When specifying the directory:

- Verify that the directory has the **execute** permission. This permission is required to run Terraform binaries under the installation directory.
  - Use an empty directory. Some installation assets, such as bootstrap X.509 certificates, have short expiration intervals, therefore you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.
- b. At the prompts, provide the configuration details for your cloud:
    - i. Optional: Select an SSH key to use to access your cluster machines.



#### NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.

- ii. Select **nutanix** as the platform to target.
  - iii. Enter the Prism Central domain name or IP address.
  - iv. Enter the port that is used to log into Prism Central.
  - v. Enter the credentials that are used to log into Prism Central.  
The installation program connects to Prism Central.
  - vi. Select the Prism Element that will manage the OpenShift Container Platform cluster.
  - vii. Select the network subnet to use.
  - viii. Enter the virtual IP address that you configured for control plane API access.
  - ix. Enter the virtual IP address that you configured for cluster ingress.
  - x. Enter the base domain. This base domain must be the same one that you configured in the DNS records.
  - xi. Enter a descriptive name for your cluster. The cluster name you enter must match the cluster name you specified when configuring the DNS records.
  - xii. Paste the [pull secret from the Red Hat OpenShift Cluster Manager](#) .
2. In the **install-config.yaml** file, set the value of **platform.nutanix.clusterOSImage** to the image location or name. For example:

```
platform:
  nutanix:
    clusterOSImage: http://mirror.example.com/images/rhcos-47.83.202103221318-0-
```

nutanix.x86\_64.qcow2

- 3. Edit the **install-config.yaml** file to give the additional information that is required for an installation in a restricted network.
  - a. Update the **pullSecret** value to contain the authentication information for your registry:

```
pullSecret: '{"auths":{"<mirror_host_name>:5000": {"auth": "<credentials>","email": "you@example.com"}}}'
```

For **<mirror\_host\_name>**, specify the registry domain name that you specified in the certificate for your mirror registry, and for **<credentials>**, specify the base64-encoded user name and password for your mirror registry.

- b. Add the **additionalTrustBundle** parameter and value.

```
additionalTrustBundle: |
  -----BEGIN CERTIFICATE-----

  /-----/
  -----END CERTIFICATE-----
```


The value must be the contents of the certificate file that you used for your mirror registry. The certificate file can be an existing, trusted certificate authority, or the self-signed certificate that you generated for the mirror registry.

- c. Add the image content resources, which resemble the following YAML excerpt:

```
imageContentSources:
  - mirrors:
    - <mirror_host_name>:5000/<repo_name>/release
      source: quay.io/openshift-release-dev/ocp-release
  - mirrors:
    - <mirror_host_name>:5000/<repo_name>/release
      source: registry.redhat.io/ocp/release
```

For these values, use the **imageContentSourcePolicy.yaml** file that was created when you mirrored the registry.

- 4. Optional: Update one or more of the default configuration parameters in the **install.config.yaml** file to customize the installation. For more information about the parameters, see "Installation configuration parameters".
- 5. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



### IMPORTANT

The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

### 3.6.1. Installation configuration parameters

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

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



## NOTE

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

### 3.6.1.1. Required configuration parameters

Required installation configuration parameters are described in the following table:

Table 3.1. Required parameters

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

Parameter	Description	Values
<b>platform</b>	The configuration for the specific platform upon which to perform the installation: <b>alibabacloud, aws, baremetal, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere,</b> or <code>{}</code> . For additional information about <b>platform</b> . <b>&lt;platform&gt;</b> parameters, consult the table for your specific platform that follows.	Object
<b>pullSecret</b>	Get a <a href="#">pull secret from the Red Hat OpenShift Cluster Manager</a> to authenticate downloading container images for OpenShift Container Platform components from services such as Quay.io.	<pre>{   "auths":{     "cloud.openshift.com":{       "auth":"b3Blb=",       "email":"you@example.com"     },     "quay.io":{       "auth":"b3Blb=",       "email":"you@example.com"     }   } }</pre>

### 3.6.1.2. Network configuration parameters

You can customize your installation configuration based on the requirements of your existing network infrastructure. For example, you can expand the IP address block for the cluster network or provide different IP address blocks than the defaults.

Only IPv4 addresses are supported.




#### NOTE


Globalnet is not supported with Red Hat OpenShift Data Foundation disaster recovery solutions. For regional disaster recovery scenarios, ensure that you use a nonoverlapping range of private IP addresses for the cluster and service networks in each cluster.

Table 3.2. Network parameters

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



Parameter	Description	Values
<b>networking</b>	The configuration for the cluster network.	Object  <b>NOTE</b> You cannot modify parameters specified by the <b>networking</b> object after installation.
<b>networking.networkType</b>	The Red Hat OpenShift Networking network plugin to install.	Either <b>OpenShiftSDN</b> or <b>OVNKubernetes</b> . <b>OpenShiftSDN</b> is a CNI plugin for all-Linux networks. <b>OVNKubernetes</b> is a CNI plugin for Linux networks and hybrid networks that contain both Linux and Windows servers. The default value is <b>OVNKubernetes</b> .
<b>networking.clusterNetwork</b>	The IP address blocks for pods.  The default value is <b>10.128.0.0/14</b> with a host prefix of <b>/23</b> .  If you specify multiple IP address blocks, the blocks must not overlap.	An array of objects. For example:  <pre>networking:   clusterNetwork:     - cidr: 10.128.0.0/14       hostPrefix: 23</pre>
<b>networking.clusterNetwork.cidr</b>	Required if you use <b>networking.clusterNetwork</b> . An IP address block.  An IPv4 network.	An IP address block in Classless Inter-Domain Routing (CIDR) notation. The prefix length for an IPv4 block is between <b>0</b> and <b>32</b> .
<b>networking.clusterNetwork.hostPrefix</b>	The subnet prefix length to assign to each individual node. For example, if <b>hostPrefix</b> is set to <b>23</b> then each node is assigned a <b>/23</b> subnet out of the given <b>cidr</b> . A <b>hostPrefix</b> value of <b>23</b> provides 510 ( $2^{(32 - 23)} - 2$ ) pod IP addresses.	A subnet prefix.  The default value is <b>23</b> .
<b>networking.serviceNetwork</b>	The IP address block for services. The default value is <b>172.30.0.0/16</b> .  The OpenShift SDN and OVN-Kubernetes network plugins support only a single IP address block for the service network.	An array with an IP address block in CIDR format. For example:  <pre>networking:   serviceNetwork:     - 172.30.0.0/16</pre>


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


### 3.6.1.3. Optional configuration parameters

Optional installation configuration parameters are described in the following table:

Table 3.3. Optional parameters



Parameter	Description	Values
<b>additionalTrustBundle</b>	A PEM-encoded X.509 certificate bundle that is added to the nodes' trusted certificate store. This trust bundle may also be used when a proxy has been configured.	String
<b>capabilities</b>	Controls the installation of optional core cluster components. You can reduce the footprint of your OpenShift Container Platform cluster by disabling optional components. For more information, see the "Cluster capabilities" page in <i>Installing</i> .	String array
<b>capabilities.baselineCapabilitySet</b>	Selects an initial set of optional capabilities to enable. Valid values are <b>None</b> , <b>v4.11</b> , <b>v4.12</b> and <b>vCurrent</b> . The default value is <b>vCurrent</b> .	String

Parameter	Description	Values
<b>capabilities.additionalEnabledCapabilities</b>	Extends the set of optional capabilities beyond what you specify in <b>baselineCapabilitySet</b> . You may specify multiple capabilities in this parameter.	String array
<b>compute</b>	The configuration for the machines that comprise the compute nodes.	Array of <b>MachinePool</b> objects.
<b>compute.architecture</b>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are <b>amd64</b> (the default).	String
<b>compute.hyperthreading</b>	Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b> , on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.   <b>IMPORTANT</b>  If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.	<b>Enabled</b> or <b>Disabled</b>
<b>compute.name</b>	Required if you use <b>compute</b> . The name of the machine pool.	<b>worker</b>
<b>compute.platform</b>	Required if you use <b>compute</b> . Use this parameter to specify the cloud provider to host the worker machines. This parameter value must match the <b>controlPlane.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>compute.replicas</b>	The number of compute machines, which are also known as worker machines, to provision.	A positive integer greater than or equal to <b>2</b> . The default value is <b>3</b> .

Parameter	Description	Values
<b>featureSet</b>	Enables the cluster for a feature set. A feature set is a collection of OpenShift Container Platform features that are not enabled by default. For more information about enabling a feature set during installation, see "Enabling features using feature gates".	String. The name of the feature set to enable, such as <b>TechPreviewNoUpgrade</b> .
<b>controlPlane</b>	The configuration for the machines that comprise the control plane.	Array of <b>MachinePool</b> objects.
<b>controlPlane.architecture</b>	Determines the instruction set architecture of the machines in the pool. Currently, clusters with varied architectures are not supported. All pools must specify the same architecture. Valid values are <b>amd64</b> (the default).	String
<b>controlPlane.hyperthreading</b>	<p>Whether to enable or disable simultaneous multithreading, or <b>hyperthreading</b>, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.</p> <div style="display: flex; align-items: center;">  <p><b>IMPORTANT</b></p> <p>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance.</p> </div>	<b>Enabled</b> or <b>Disabled</b>
<b>controlPlane.name</b>	Required if you use <b>controlPlane</b> . The name of the machine pool.	<b>master</b>
<b>controlPlane.platform</b>	Required if you use <b>controlPlane</b> . Use this parameter to specify the cloud provider that hosts the control plane machines. This parameter value must match the <b>compute.platform</b> parameter value.	<b>alibabacloud, aws, azure, gcp, ibmcloud, nutanix, openstack, ovirt, vsphere</b> , or <b>{}</b>
<b>controlPlane.replicas</b>	The number of control plane machines to provision.	The only supported value is <b>3</b> , which is the default value.

Parameter	Description	Values
<b>credentialsMode</b>	<p>The Cloud Credential Operator (CCO) mode. If no mode is specified, the CCO dynamically tries to determine the capabilities of the provided credentials, with a preference for mint mode on the platforms where multiple modes are supported.</p> <div data-bbox="486 510 595 891" style="border: 1px solid black; padding: 5px; margin-bottom: 10px;"> <p><b>NOTE</b></p> <p>Not all CCO modes are supported for all cloud providers. For more information about CCO modes, see the <i>Cloud Credential Operator</i> entry in the <i>Cluster Operators reference</i> content.</p> </div> <div data-bbox="486 936 595 1288" style="border: 1px solid black; padding: 5px;"> <p><b>NOTE</b></p> <p>If your AWS account has service control policies (SCP) enabled, you must configure the <b>credentialsMode</b> parameter to <b>Mint</b>, <b>Passthrough</b> or <b>Manual</b>.</p> </div>	<b>Mint, Passthrough, Manual</b> or an empty string ("").

Parameter	Description	Values
<b>fips</b>	<p>Enable or disable FIPS mode. The default is <b>false</b> (disabled). If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 20px; height: 100px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, black 2px, black 4px); margin-right: 10px;"></div> <div> <p><b>IMPORTANT</b></p> <p>To enable FIPS mode for your cluster, you must run the installation program from a Red Hat Enterprise Linux (RHEL) computer configured to operate in FIPS mode. For more information about configuring FIPS mode on RHEL, see <a href="#">Installing the system in FIPS mode</a>. The use of FIPS validated or Modules In Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the <b>x86_64</b>, <b>ppc64le</b>, and <b>s390x</b> architectures.</p> </div> </div> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="width: 20px; height: 40px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, black 2px, black 4px); margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>If you are using Azure File storage, you cannot enable FIPS mode.</p> </div> </div>	<b>false</b> or <b>true</b>
<b>imageContentSources</b>	Sources and repositories for the release-image content.	Array of objects. Includes a <b>source</b> and, optionally, <b>mirrors</b> , as described in the following rows of this table.
<b>imageContentSources.source</b>	Required if you use <b>imageContentSources</b> . Specify the repository that users refer to, for example, in image pull specifications.	String

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

#### 3.6.1.4. Additional Nutanix configuration parameters

Additional Nutanix configuration parameters are described in the following table:

**Table 3.4. Additional Nutanix cluster parameters**

Parameter	Description	Values
<b>platform.nutanix.apiVIP</b>	The virtual IP (VIP) address that you configured for control plane API access.	IP address
<b>platform.nutanix.ingressVIP</b>	The virtual IP (VIP) address that you configured for cluster ingress.	IP address

Parameter	Description	Values
<b>platform.nutanix.pris mCentral.endpoint.a ddress</b>	The Prism Central domain name or IP address.	String
<b>platform.nutanix.pris mCentral.endpoint.p ort</b>	The port that is used to log into Prism Central.	String
<b>platform.nutanix.pris mCentral.password</b>	The password for the Prism Central user name.	String
<b>platform.nutanix.pris mCentral.username</b>	The user name that is used to log into Prism Central.	String
<b>platform.nutanix.pris mElements.endpoint. address</b>	The Prism Element domain name or IP address. [1]	String
<b>platform.nutanix.pris mElements.endpoint. port</b>	The port that is used to log into Prism Element.	String
<b>platform.nutanix.pris mElements.uuid</b>	The universally unique identifier (UUID) for Prism Element.	String
<b>platform.nutanix.sub netUUIDs</b>	The UUID of the Prism Element network that contains the virtual IP addresses and DNS records that you configured. [2]	String
<b>platform.nutanix.clu sterOSImage</b>	Optional: By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image. If Prism Central does not have internet access, you can override the default behavior by hosting the RHCOS image on any HTTP server and pointing the installation program to the image.	An HTTP or HTTPS URL, optionally with a SHA-256 checksum. For example, <code>http://example.com/images/rhcos-47.83.202103221318-0-nutanix.x86_64.qcow2</code>

1. The **prismElements** section holds a list of Prism Elements (clusters). A Prism Element encompasses all of the Nutanix resources, for example virtual machines and subnets, that are used to host the OpenShift Container Platform cluster. Only a single Prism Element is supported.
2. Only one subnet per OpenShift Container Platform cluster is supported.



### 3.6.2. Sample customized install-config.yaml file for Nutanix

You can customize the **install-config.yaml** file to specify more details about your OpenShift Container Platform cluster's platform or modify the values of the required parameters.



#### IMPORTANT

This sample YAML file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program and modify it.

```

apiVersion: v1
baseDomain: example.com 1
compute: 2
- hyperthreading: Enabled 3
  name: worker
  replicas: 3
  platform:
    nutanix: 4
      cpus: 2
      coresPerSocket: 2
      memoryMiB: 8196
      osDisk:
        diskSizeGiB: 120
controlPlane: 5
  hyperthreading: Enabled 6
  name: master
  replicas: 3
  platform:
    nutanix: 7
      cpus: 4
      coresPerSocket: 2
      memoryMiB: 16384
      osDisk:
        diskSizeGiB: 120
metadata:
  creationTimestamp: null
  name: test-cluster 8
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14
      hostPrefix: 23
  machineNetwork:
    - cidr: 10.0.0.0/16
  networkType: OVNKubernetes 9
  serviceNetwork:
    - 172.30.0.0/16
platform:
  nutanix:
    apiVIP: 10.40.142.7 10
    ingressVIP: 10.40.142.8 11
  prismCentral:
    endpoint:
      address: your.prismcentral.domainname 12

```



- 16 Optional: By default, the installation program downloads and installs the Red Hat Enterprise Linux CoreOS (RHCOS) image. If Prism Central does not have internet access, you can override the
- 17 For **<local\_registry>**, specify the registry domain name, and optionally the port, that your mirror registry uses to serve content. For example **registry.example.com** or **registry.example.com:5000**. For **<credentials>**, specify the base64-encoded user name and password for your mirror registry.
- 18 Whether to enable or disable FIPS mode. By default, FIPS mode is not enabled. If FIPS mode is enabled, the Red Hat Enterprise Linux CoreOS (RHCOS) machines that OpenShift Container Platform runs on bypass the default Kubernetes cryptography suite and use the cryptography modules that are provided with RHCOS instead.



### IMPORTANT

The use of FIPS Validated or Modules in Process cryptographic libraries is only supported on OpenShift Container Platform deployments on the **x86\_64**, **ppc64le**, and **s390x** architectures.

- 19 Optional: You can provide the **sshKey** value that you use to access the machines in your cluster.



### NOTE

For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery, specify an SSH key that your **ssh-agent** process uses.


- 20 Provide the contents of the certificate file that you used for your mirror registry.
- 21 Provide these values from the **metadata.name: release-0** section of the **imageContentSourcePolicy.yaml** file that was created when you mirrored the registry.

### 3.6.3. Configuring the cluster-wide proxy during installation

Production environments can deny direct access to the internet and instead have an HTTP or HTTPS proxy available. You can configure a new OpenShift Container Platform cluster to use a proxy by configuring the proxy settings in the **install-config.yaml** file.

#### Prerequisites

- You have an existing **install-config.yaml** file.
- You reviewed the sites that your cluster requires access to and determined whether any of them need to bypass the proxy. By default, all cluster egress traffic is proxied, including calls to hosting cloud provider APIs. You added sites to the **Proxy** object's **spec.noProxy** field to bypass the proxy if necessary.



## NOTE

The **Proxy** object **status.noProxy** field is populated with the values of the **networking.machineNetwork[].cidr**, **networking.clusterNetwork[].cidr**, and **networking.serviceNetwork[]** fields from your installation configuration.

For installations on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, and Red Hat OpenStack Platform (RHOSP), the **Proxy** object **status.noProxy** field is also populated with the instance metadata endpoint (**169.254.169.254**).

## Procedure

1. Edit your **install-config.yaml** file and add the proxy settings. For example:

```
apiVersion: v1
baseDomain: my.domain.com
proxy:
  httpProxy: http://<username>:<pswd>@<ip>:<port> 1
  httpsProxy: https://<username>:<pswd>@<ip>:<port> 2
  noProxy: example.com 3
  additionalTrustBundle: | 4
    -----BEGIN CERTIFICATE-----
    <MY_TRUSTED_CA_CERT>
    -----END CERTIFICATE-----
  additionalTrustBundlePolicy: <policy_to_add_additionalTrustBundle> 5
```

- 1 A proxy URL to use for creating HTTP connections outside the cluster. The URL scheme must be **http**.
- 2 A proxy URL to use for creating HTTPS connections outside the cluster.
- 3 A comma-separated list of destination domain names, IP addresses, or other network CIDRs to exclude from proxying. Preface a domain with **.** to match subdomains only. For example, **.y.com** matches **x.y.com**, but not **y.com**. Use **\*** to bypass the proxy for all destinations.
- 4 If provided, the installation program generates a config map that is named **user-ca-bundle** in the **openshift-config** namespace that contains one or more additional CA certificates that are required for proxying HTTPS connections. The Cluster Network Operator then creates a **trusted-ca-bundle** config map that merges these contents with the Red Hat Enterprise Linux CoreOS (RHCOS) trust bundle, and this config map is referenced in the **trustedCA** field of the **Proxy** object. The **additionalTrustBundle** field is required unless the proxy's identity certificate is signed by an authority from the RHCOS trust bundle.
- 5 Optional: The policy to determine the configuration of the **Proxy** object to reference the **user-ca-bundle** config map in the **trustedCA** field. The allowed values are **Proxyonly** and **Always**. Use **Proxyonly** to reference the **user-ca-bundle** config map only when **http/https** proxy is configured. Use **Always** to always reference the **user-ca-bundle** config map. The default value is **Proxyonly**.



## NOTE

The installation program does not support the proxy **readinessEndpoints** field.

**NOTE**

If the installer times out, restart and then complete the deployment by using the **wait-for** command of the installer. For example:

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

2. Save the file and reference it when installing OpenShift Container Platform.

The installation program creates a cluster-wide proxy that is named **cluster** that uses the proxy settings in the provided **install-config.yaml** file. If no proxy settings are provided, a **cluster Proxy** object is still created, but it will have a nil **spec**.

**NOTE**

Only the **Proxy** object named **cluster** is supported, and no additional proxies can be created.

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

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

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.12. Download and install the new version of **oc**.

#### Installing the OpenShift CLI on Linux

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

##### Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the architecture from the **Product Variant** drop-down list.
3. Select the appropriate version from the **Version** drop-down list.
4. Click **Download Now** next to the **OpenShift v4.12 Linux Client** entry and save the file.
5. Unpack the archive:

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

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

```
$ echo $PATH
```

## Verification

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

```
$ oc <command>
```

## Installing the OpenShift CLI on Windows

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

### Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.12 Windows Client** entry and save the file.
4. Unzip the archive with a ZIP program.
5. Move the **oc** binary to a directory that is on your **PATH**.

To check your **PATH**, open the command prompt and execute the following command:

```
C:\> path
```

## Verification

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


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

## Installing the OpenShift CLI on macOS

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

### Procedure

1. Navigate to the [OpenShift Container Platform downloads page](#) on the Red Hat Customer Portal.
2. Select the appropriate version from the **Version** drop-down list.
3. Click **Download Now** next to the **OpenShift v4.12 macOS Client** entry and save the file.



#### NOTE

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

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

```
$ echo $PATH
```

## Verification

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

```
$ oc <command>
```

## 3.8. CONFIGURING IAM FOR NUTANIX

Installing the cluster requires that the Cloud Credential Operator (CCO) operate in manual mode. While the installation program configures the CCO for manual mode, you must specify the identity and access management secrets.

### Prerequisites

- You have configured the **ccoctl** binary.
- You have an **install-config.yaml** file.

### Procedure

1. Create a YAML file that contains the credentials data in the following format:

#### Credentials data format

```
credentials:
- type: basic_auth 1
  data:
    prismCentral: 2
      username: <username_for_prism_central>
      password: <password_for_prism_central>
    prismElements: 3
      - name: <name_of_prism_element>
        username: <username_for_prism_element>
        password: <password_for_prism_element>
```

- 1 Specify the authentication type. Only basic authentication is supported.
- 2 Specify the Prism Central credentials.
- 3 Optional: Specify the Prism Element credentials.

2. Set a **\$RELEASE\_IMAGE** variable with the release image from your installation file by running the following command:

```
$ RELEASE_IMAGE=$(./openshift-install version | awk '/release image/ {print $3}')
```

3. Extract the list of **CredentialsRequest** custom resources (CRs) from the OpenShift Container Platform release image by running the following command:

```
$ oc adm release extract \
  --from=$RELEASE_IMAGE \
  --credentials-requests \
```

```
--cloud=nutanix \
--to=<path_to_directory_with_list_of_credentials_requests>/credrequests 1
```

- 1** Specify the path to the directory that contains the files for the component **CredentialsRequests** objects. If the specified directory does not exist, this command creates it.

### Sample CredentialsRequest object

```
apiVersion: cloudcredential.openshift.io/v1
kind: CredentialsRequest
metadata:
  annotations:
    include.release.openshift.io/self-managed-high-availability: "true"
  labels:
    controller-tools.k8s.io: "1.0"
  name: openshift-machine-api-nutanix
  namespace: openshift-cloud-credential-operator
spec:
  providerSpec:
    apiVersion: cloudcredential.openshift.io/v1
    kind: NutanixProviderSpec
  secretRef:
    name: nutanix-credentials
    namespace: openshift-machine-api
```

- If your cluster uses cluster capabilities to disable one or more optional components, delete the **CredentialsRequest** custom resources for any disabled components.

### Example credrequests directory contents for OpenShift Container Platform 4.12 on Nutanix

```
0000_30_machine-api-operator_00_credentials-request.yaml 1
```

- 1** The Machine API Operator CR is required.
- Use the **ccoctl** tool to process all of the **CredentialsRequest** objects in the **credrequests** directory by running the following command:

```
$ ccoctl nutanix create-shared-secrets \
--credentials-requests-dir=
<path_to_directory_with_list_of_credentials_requests>/credrequests 1
--output-dir=<ccoctl_output_dir> 2
--credentials-source-filepath=<path_to_credentials_file> 3
```

- 1** Specify the path to the directory that contains the files for the component **CredentialsRequests** objects.
- 2** Specify the directory that contains the files of the component credentials secrets, under the **manifests** directory. By default, the **ccoctl** tool creates objects in the directory in which the commands are run. To create the objects in a different directory, use the **--output-dir** flag.



- 3 Optional: Specify the directory that contains the credentials data YAML file. By default, **ccoctl** expects this file to be in **<home\_directory>/nutanix/credentials**. To specify a
6. Edit the **install-config.yaml** configuration file so that the **credentialsMode** parameter is set to **Manual**.

#### Example install-config.yaml configuration file

```
apiVersion: v1
baseDomain: cluster1.example.com
credentialsMode: Manual 1
...
```

- 1 Add this line to set the **credentialsMode** parameter to **Manual**.

7. Create the installation manifests by running the following command:

```
$ openshift-install create manifests --dir <installation_directory> 1
```

- 1 Specify the path to the directory that contains the **install-config.yaml** file for your cluster.

8. Copy the generated credential files to the target manifests directory by running the following command:

```
$ cp <ccoctl_output_dir>/manifests/*credentials.yaml ./<installation_directory>/manifests
```

#### Verification

- Ensure that the appropriate secrets exist in the **manifests** directory.

```
$ ls ./<installation_directory>/manifests
```

#### Example output

```
total 64
-rw-r----- 1 <user> <user> 2335 Jul  8 12:22 cluster-config.yaml
-rw-r----- 1 <user> <user>  161 Jul  8 12:22 cluster-dns-02-config.yml
-rw-r----- 1 <user> <user>  864 Jul  8 12:22 cluster-infrastructure-02-config.yml
-rw-r----- 1 <user> <user>  191 Jul  8 12:22 cluster-ingress-02-config.yml
-rw-r----- 1 <user> <user> 9607 Jul  8 12:22 cluster-network-01-crd.yml
-rw-r----- 1 <user> <user>  272 Jul  8 12:22 cluster-network-02-config.yml
-rw-r----- 1 <user> <user>  142 Jul  8 12:22 cluster-proxy-01-config.yaml
-rw-r----- 1 <user> <user>  171 Jul  8 12:22 cluster-scheduler-02-config.yml
-rw-r----- 1 <user> <user>  200 Jul  8 12:22 cvo-overrides.yaml
-rw-r----- 1 <user> <user>  118 Jul  8 12:22 kube-cloud-config.yaml
-rw-r----- 1 <user> <user> 1304 Jul  8 12:22 kube-system-configmap-root-ca.yaml
-rw-r----- 1 <user> <user> 4090 Jul  8 12:22 machine-config-server-tls-secret.yaml
-rw-r----- 1 <user> <user> 3961 Jul  8 12:22 openshift-config-secret-pull-secret.yaml
-rw----- 1 <user> <user>  283 Jul  8 12:24 openshift-machine-api-nutanix-credentials-credentials.yaml
```

## 3.9. DEPLOYING THE CLUSTER

You can install OpenShift Container Platform on a compatible cloud platform.



### IMPORTANT

You can run the **create cluster** command of the installation program only once, during initial installation.

### Prerequisites


- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.
- Verify the cloud provider account on your host has the correct permissions to deploy the cluster. An account with incorrect permissions causes the installation process to fail with an error message that displays the missing permissions.

### Procedure

- Change to the directory that contains the installation program and initialize the cluster deployment:

```
$ ./openshift-install create cluster --dir <installation_directory> \ 1
--log-level=info 2
```

- 1** For **<installation\_directory>**, specify the location of your customized **./install-config.yaml** file.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.




### NOTE

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

### Verification

When the cluster deployment completes successfully:

- The terminal displays directions for accessing your cluster, including a link to the web console and credentials for the **kubeadmin** user.
- Credential information also outputs to **<installation\_directory>/openshift\_install.log**.



### IMPORTANT

Do not delete the installation program or the files that the installation program creates. Both are required to delete the cluster.

### Example output

■

```

...
INFO Install complete!
INFO To access the cluster as the system:admin user when using 'oc', run 'export
KUBECONFIG=/home/myuser/install_dir/auth/kubeconfig'
INFO Access the OpenShift web-console here: https://console-openshift-
console.apps.mycluster.example.com
INFO Login to the console with user: "kubeadmin", and password: "password"
INFO Time elapsed: 36m22s

```

## IMPORTANT

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

## 3.10. POST INSTALLATION

Complete the following steps to complete the configuration of your cluster.

### 3.10.1. Disabling the default OperatorHub catalog sources

Operator catalogs that source content provided by Red Hat and community projects are configured for OperatorHub by default during an OpenShift Container Platform installation. In a restricted network environment, you must disable the default catalogs as a cluster administrator.

#### Procedure

- Disable the sources for the default catalogs by adding **disableAllDefaultSources: true** to the **OperatorHub** object:

```

$ oc patch OperatorHub cluster --type json \
  -p '[{"op": "add", "path": "/spec/disableAllDefaultSources", "value": true}]'

```

#### TIP

Alternatively, you can use the web console to manage catalog sources. From the **Administration** → **Cluster Settings** → **Configuration** → **OperatorHub** page, click the **Sources** tab, where you can create, update, delete, disable, and enable individual sources.

### 3.10.2. Installing the policy resources into the cluster

Mirroring the OpenShift Container Platform content using the `oc-mirror` OpenShift CLI (`oc`) plugin creates resources, which include **catalogSource-certified-operator-index.yaml** and **imageContentSourcePolicy.yaml**.

- The **ImageContentSourcePolicy** resource associates the mirror registry with the source registry and redirects image pull requests from the online registries to the mirror registry.
- The **CatalogSource** resource is used by Operator Lifecycle Manager (OLM) to retrieve information about the available Operators in the mirror registry, which lets users discover and install Operators.

After you install the cluster, you must install these resources into the cluster.

### Prerequisites

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

### Procedure

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

```
$ oc apply -f ./oc-mirror-workspace/results-<id>/
```

### Verification

1. Verify that the **ImageContentSourcePolicy** resources were successfully installed:

```
$ oc get imagecontentsourcepolicy
```

2. Verify that the **CatalogSource** resources were successfully installed:

```
$ oc get catalogsource --all-namespaces
```

### 3.10.3. Configuring the default storage container

After you install the cluster, you must install the Nutanix CSI Operator and configure the default storage container for the cluster.

For more information, see the Nutanix documentation for [installing the CSI Operator](#) and [configuring registry storage](#).

## 3.11. TELEMETRY ACCESS FOR OPENSIFT CONTAINER PLATFORM

In OpenShift Container Platform 4.12, the Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to [OpenShift Cluster Manager Hybrid Cloud Console](#).

After you confirm that your [OpenShift Cluster Manager Hybrid Cloud Console](#) inventory is correct, either maintained automatically by Telemetry or manually by using OpenShift Cluster Manager, [use](#)

[subscription watch](#) to track your OpenShift Container Platform subscriptions at the account or multi-cluster level.

## 3.12. ADDITIONAL RESOURCES

- [About remote health monitoring](#)

## 3.13. NEXT STEPS

- If necessary, see [Opt out of remote health reporting](#)
- If necessary, see [Registering your disconnected cluster](#)
- [Customize your cluster](#)

## CHAPTER 4. UNINSTALLING A CLUSTER ON NUTANIX

You can remove a cluster that you deployed to Nutanix.

### 4.1. REMOVING A CLUSTER THAT USES INSTALLER-PROVISIONED INFRASTRUCTURE

You can remove a cluster that uses installer-provisioned infrastructure from your cloud.



#### NOTE

After uninstallation, check your cloud provider for any resources not removed properly, especially with user-provisioned infrastructure clusters. There might be resources that the installation program did not create or that the installation program is unable to access.

#### Prerequisites

- You have a copy of the installation program that you used to deploy the cluster.
- You have the files that the installation program generated when you created your cluster.

#### Procedure

1. On the computer that you used to install the cluster, go to the directory that contains the installation program, and run the following command:

```
$. /openshift-install destroy cluster \  
--dir <installation_directory> --log-level info 1 2
```

- 1** For **<installation\_directory>**, specify the path to the directory that you stored the installation files in.
- 2** To view different details, specify **warn**, **debug**, or **error** instead of **info**.



#### NOTE

You must specify the directory that contains the cluster definition files for your cluster. The installation program requires the **metadata.json** file in this directory to delete the cluster.

2. Optional: Delete the **<installation\_directory>** directory and the OpenShift Container Platform installation program.