



# OpenShift Container Platform 4.4

## Installing on RHV

Installing OpenShift Container Platform RHV clusters



# OpenShift Container Platform 4.4 Installing on RHV

---

Installing OpenShift Container Platform RHV clusters

## Legal Notice

Copyright © 2021 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 provides instructions for installing and uninstalling OpenShift Container Platform clusters on Red Hat Virtualization.

# Table of Contents

<b>CHAPTER 1. INSTALLING ON RHV</b> .....	<b>4</b>
1.1. CREATING A CUSTOM VIRTUAL MACHINE TEMPLATE ON RHV	4
1.1.1. Downloading a specific RHCOS image to the RHV Manager machine	4
1.1.2. Using an Ansible playbook to upload an RHCOS image to a data storage domain	5
1.1.3. Determining the resources available to customize your virtual machine template	6
1.1.4. Attaching the virtual disk with the RHCOS image to a virtual machine	7
1.1.5. Configuring and creating of the virtual machine	8
1.1.6. Creating a virtual machine template	8
1.1.7. Exporting some environment variables	8
1.2. INSTALLING A CLUSTER QUICKLY ON RHV	9
1.2.1. Prerequisites	10
1.2.2. Internet and Telemetry access for OpenShift Container Platform	10
1.2.3. Requirements for the RHV environment	11
1.2.4. Verifying the requirements for the RHV environment	12
1.2.5. Preparing the network environment on RHV	14
1.2.6. Setting up the CA certificate for RHV	15
1.2.7. Generating an SSH private key and adding it to the agent	16
1.2.8. Obtaining the installation program	16
1.2.9. Deploying the cluster	17
<b>CHAPTER 2. INSTALLING THE CLI BY DOWNLOADING THE BINARY</b> .....	<b>21</b>
2.1. INSTALLING THE CLI ON LINUX	21
2.2. INSTALLING THE CLI ON WINDOWS	21
2.3. INSTALLING THE CLI ON MACOS	22
<b>CHAPTER 3. LOGGING IN TO THE CLUSTER</b> .....	<b>23</b>
3.1. VERIFYING CLUSTER STATUS	23
3.2. ACCESSING THE OPENSIFT CONTAINER PLATFORM WEB CONSOLE ON RHV	24
3.3. TROUBLESHOOTING COMMON ISSUES WITH INSTALLING ON RED HAT VIRTUALIZATION (RHV)	24
3.3.1. CPU load increases and nodes go into a Not Ready state	24
3.3.2. Trouble connecting the OpenShift Container Platform cluster API	25
3.4. POST-INSTALLATION TASKS	25
3.5. INSTALLING A CLUSTER ON RHV WITH CUSTOMIZATIONS	25
3.5.1. Prerequisites	26
3.5.2. Internet and Telemetry access for OpenShift Container Platform	27
3.5.3. Requirements for the RHV environment	27
3.5.4. Verifying the requirements for the RHV environment	28
3.5.5. Preparing the network environment on RHV	30
3.5.6. Setting up the CA certificate for RHV	31
3.5.7. Generating an SSH private key and adding it to the agent	32
3.5.8. Obtaining the installation program	33
3.5.9. Creating the installation configuration file	33
3.5.10. Installation configuration parameters for RHV	36
3.5.11. Deploying the cluster	37
3.5.12. Installing the CLI by downloading the binary	38
3.5.12.1. Installing the CLI on Linux	38
3.5.12.2. Installing the CLI on Windows	39
3.5.12.3. Installing the CLI on macOS	39
3.5.13. Logging in to the cluster	40
3.5.14. Verifying cluster status	40
3.5.15. Accessing the OpenShift Container Platform web console on RHV	41

3.5.16. Troubleshooting common issues with installing on Red Hat Virtualization (RHV)	42
3.5.16.1. CPU load increases and nodes go into a Not Ready state	42
3.5.16.2. Trouble connecting the OpenShift Container Platform cluster API	42
3.5.17. Post-installation tasks	43
3.5.18. Next steps	43
3.6. UNINSTALLING A CLUSTER ON RHV	43
3.6.1. Removing a cluster that uses installer-provisioned infrastructure	43



# CHAPTER 1. INSTALLING ON RHV

## 1.1. CREATING A CUSTOM VIRTUAL MACHINE TEMPLATE ON RHV

Before installing a OpenShift Container Platform cluster on your Red Hat Virtualization (RHV) environment, you create a custom virtual machine template and configure your environment so the installation program uses the template.



### IMPORTANT

Creating a custom virtual machine template for RHV is a workaround for a known issue ([BZ#1818577](#)). If you do not create a custom virtual machine, the OpenShift Container Platform cluster you install will fail.

### Prerequisites

- Review details about the [OpenShift Container Platform installation and update](#) processes.
- If you use a firewall, [configure it to allow the sites](#) that your cluster requires access to.

### 1.1.1. Downloading a specific RHCOS image to the RHV Manager machine

Download a specific Red Hat Enterprise Linux CoreOS (RHCOS) image to the Red Hat Virtualization (RHV) Manager machine.



### IMPORTANT

Do not substitute the RHCOS image specified by these instructions with another image.

### Procedure

1. Open a terminal session with the RHV Manager machine. For example, if you run the Manager as a self-hosted engine:
  - a. Go to the RHV Administration Portal and go to the **Compute** → **Virtual Machines** page.
  - b. Select the **HostedEngine** virtual machine and click **Console**.
2. On the Manager's command line, create a working directory and change to it.

```
$ mkdir rhcos  
$ cd rhcos
```

3. In a browser, go to <https://github.com/openshift/installer/blob/release-4.4/data/data/rhcos.json>.
4. In **rhcos.json**, find **baseURI** and copy its value.
5. On the Manager, start a **wget** command and paste the value of **baseURI** but **do not press Enter**. For example:

```
$ wget https://releases-art-rhcos.svc.ci.openshift.org/art/storage/releases/rhcos-4.4/44.81.202003062006-0/x86_64/
```

6. In `rhcos.json` document, find **openstack** and copy the value of **path**.
7. On the Manager, paste the value of **path**. For example:

```
$ wget https://releases-art-rhcos.svc.ci.openshift.org/art/storage/releases/rhcos-4.4/44.81.202003062006-0/x86_64/rhcos-44.81.202003062006-0-openstack.x86_64.qcow2.gz
```

8. Press **Enter** and wait for **wget** to finish downloading the RHCOS image.
9. Unzip the RHCOS image. For example:

```
$ gunzip rhcos*
```

### 1.1.2. Using an Ansible playbook to upload an RHCOS image to a data storage domain

Use an Ansible playbook to upload an Red Hat Enterprise Linux CoreOS (RHCOS) image to a data storage domain.

#### Prerequisites

- [Red Hat Ansible Engine, version 2.9 or later](#)
- [Python oVirt Engine SDK, version 4.3 or later](#)

#### Procedure

1. Create an Ansible playbook file **upload\_rhcos\_disk.yaml** on the Red Hat Virtualization (RHV) Manager. For example:

```
$ vi upload_rhcos_disk.yaml
```

2. Go to [this file on the Red Hat Customer Portal](#) .
3. Paste its contents into your playbook.
4. In your playbook, update the parameter values that have callouts.

```
- hosts: localhost
gather_facts: no
tasks:
  - name: Authenticate with engine
    ovirt_auth:
      url: https://<virtlab.example.com>/ovirt-engine/api 1
      username: <username@profile> 2
      password: <password> 3
      insecure: yes

  - name: Upload RHCOS image
    ovirt_disk:
      auth: "{{ ovirt_auth }}"
      name: <rhcos_44-81_img-diskname> 4
```

```

size: 120GiB 5
interface: virtio_scsi
storage_domain: <SSD_RAID_10> 6
bootable: yes
timeout: 3600 7
upload_image_path: </custom/rhcos-44.81.202003110027-0-openstack.x86_64.qcow2>
8
wait: yes

```

- 1** For **<virtlab.example.com>**, specify the fully-qualified domain name (FQDN) of your RHV Manager.
- 2** For **<username@profile>**, specify the admin user name and profile.
- 3** For **<password>**, specify the admin password.
- 4** For **<rhcos\_44-81\_img-diskname>**, specify a disk name.
- 5** Specify **120GiB** or more for production environments. For resource-constrained or non-production environments, specify **32GiB** or more.
- 6** For **<SSD\_RAID\_10>**, specify a data storage domain name.
- 7** Specify a timeout period, in seconds, that gives your RHV environment enough time to upload a 2.4 GB image to storage. The default value, **3600** seconds, gives one hour.
- 8** For **</custom/rhcos-44.81.202003110027-0-openstack.x86\_64.qcow2>**, specify the image path and file name.

5. On the Manager, run the Ansible playbook. For example:

```
$ ansible-playbook rhcos_image.yaml -vvv
```

6. In the Administration Portal, go to the **Storage → Disks** page and find the disk name you specified in the playbook. For example, **rhcos\_44-81\_img-diskname**.
7. When the status of the disk is **OK**, the playbook has finished uploading the RHCOS image to the storage domain.

### 1.1.3. Determining the resources available to customize your virtual machine template

If you plan to [install a cluster quickly on RHV using the default configuration options](#), skip to the next task, [Attaching the virtual disk with the RHCOS image to the virtual machine](#).

Otherwise, if you plan to [install a cluster on RHV with customizations](#), consider the following information.

You can customize the virtual machine template the installation program uses to create the control and compute machines in your OpenShift Container Platform cluster. You can customize the template so these machines have more than the default virtual CPU, memory, and storage resources. Your customizations apply equally to both control and compute machines; they cannot be customized differently.

Over-allocating resources can cause the installation or operation of your cluster to fail. To avoid over-

allocating resources, you must determine how much of each resource is available for each virtual machine: Inspect the resources in your Red Hat Virtualization (RHV) cluster and divide them by the number of virtual machines. The resulting quotients are the maximum values you can allocate for each resource.

### Procedure

1. Inspect the available resources in your RHV cluster.
  - a. Go to [Verifying the requirements for the RHV environment](#)
  - b. Record the values of the storage **Free Space**, **Logical CPU Cores** and **Max free Memory for scheduling new VMs**.
2. Record the number of virtual machines in your OpenShift Container Platform cluster.
  - By default, your cluster contains seven machines.
  - You can also customize the number of machines, as described in [Installation configuration parameters for RHV](#). To account for the bootstrap machine, add one machine to the number in the configuration file.
3. Divide each resource by the number of virtual machines.
4. Record these values for use later on.

### 1.1.4. Attaching the virtual disk with the RHCOS image to a virtual machine

Attach the virtual disk with the Red Hat Enterprise Linux CoreOS (RHCOS) image to a virtual machine.

### Procedure

1. Go to the **Compute** → **Virtual Machines** page and click **New**.
2. In the **New Virtual Machine** → **General** window that opens, use **Cluster** to select the RHV cluster where you plan to create the OpenShift Container Platform cluster.
3. Leave **Template** unchanged, with **Blank** selected.
4. Set **Operating System** to **Red Hat Enterprise Linux CoreOS**
5. Leave **Optimized for** unchanged, with **Desktop** selected.
6. For **Name**, specify the name of the virtual machine. For example, **vmname**.
7. For **Instance Images**, click **Attach**.
8. In the **Attach Virtual Disks** window that opens, find the disk you specified in the playbook. For example, **rhcos\_44-81\_img-diskname**.
9. Click the radio button for this disk.
10. Select the **OS** checkbox for this disk. This makes the disk bootable.
11. Click **OK** to save and close the **Attach Virtual Disks** window.

## 1.1.5. Configuring and creating of the virtual machine

Continue configuring the virtual machine and then create it.

### Procedure

1. In the **New Virtual Machine** → **General** window, for each NIC under **Instantiate VM network interfaces by picking a vNIC profile**, select a vNIC profile.
2. Go to the **New Virtual Machine** → **System** window.
3. Set **Memory Size** to **16384 MB** or more. This default value is equivalent to 16 GiB. You can adjust this value according to the workload you expect on the compute machines and the amount of memory resources that are available.
4. Due to a known issue ([bz#1821215](#)), set **Physical Memory Guaranteed** to **8192 MB**. This value is equivalent to 8 GiB.
5. Set **Total Virtual CPUs** to **4** or more. You can adjust this value according to the workload you expect on the compute machines and the resources that are available.
6. Expand **Advanced Parameters**.
7. Adjust **Cores per Virtual Socket** to **4** or more so it matches the **Total Virtual CPUs** setting.
8. Press **OK** to save and close the **New Virtual Machine**.  
The new virtual machine appears in the **Compute** → **Virtual Machines** window. Because the virtual machine is not starting, it appears quickly.

## 1.1.6. Creating a virtual machine template

Create the virtual machine template.

### Procedure

1. Select the new virtual machine.
2. In the upper-right corner of the Administration Portal window, click the **More actions** menu, and select **Make Template**.
3. In the **New Template** window, specify a value for the **Name**. For example, **vm-tmplname**.
4. Verify that **Target** and the other parameter values are correct and reflect your previous choices.
5. Verify that **Seal Template (Linux only)** is **not** selected.
6. Click **OK**.
7. Go to the **Compute** → **Templates** page and wait for the template to be created.

## 1.1.7. Exporting some environment variables



## IMPORTANT

Using these environment variables to customize the install is a temporary workaround for a known issue ([BZ#1818577](#)). This content might be removed in a future release.

1. Determine the values you specified for **Memory** and **Total Virtual CPUs** in the virtual machine settings.
2. Run the following environment variables on the command line of your installation machine.

```
$ export OPENSIFT_INSTALL_OS_IMAGE_OVERRIDE=<vm-tmplname> 1
$ export TF_VAR_ovirt_template_mem=<mem-value> 2
$ export TF_VAR_ovirt_template_cpu=<cpu-value> 3
$ export TF_VAR_ovirt_master_mem=<mem-value> 4
$ export TF_VAR_ovirt_master_cpu=<cpu-value> 5
```

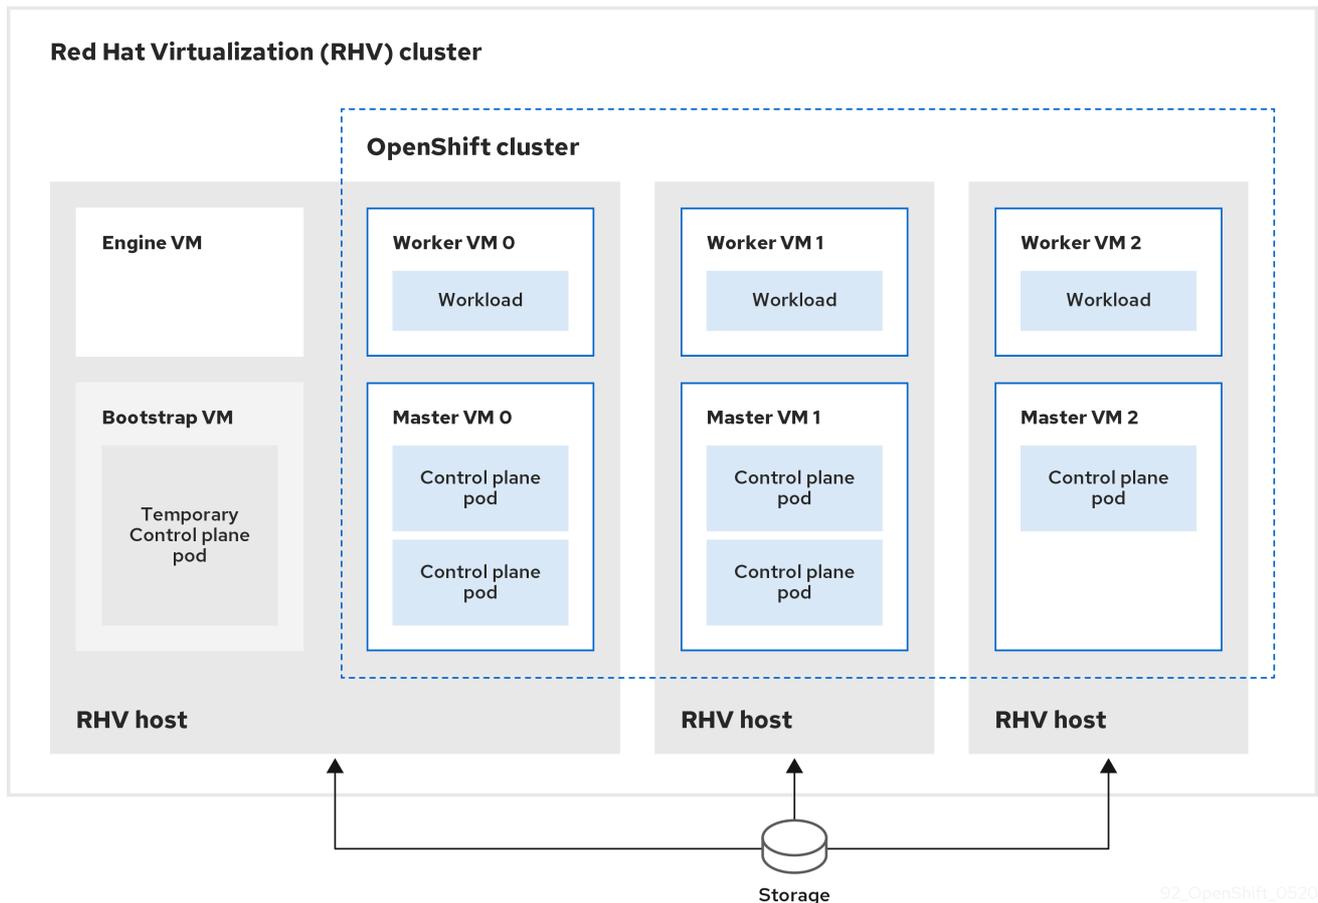
- 1 For **<vm-tmplname>**, specify the name of the template that you created.
- 2 4 For **<mem-value>**, specify the value of **Memory** in MB, in the virtual machine. For example, **16384**.
- 3 5 For **<cpu-value>**, specify value of **Total Virtual CPUs** in the virtual machine. For example, **4**.

For example:

```
$ export OPENSIFT_INSTALL_OS_IMAGE_OVERRIDE=vm-tmplname
$ export TF_VAR_ovirt_template_mem=16384
$ export TF_VAR_ovirt_template_cpu=4
$ export TF_VAR_ovirt_master_mem=16384
$ export TF_VAR_ovirt_master_cpu=4
```

## 1.2. INSTALLING A CLUSTER QUICKLY ON RHV

You can quickly install a default, non-customized, OpenShift Container Platform cluster on a Red Hat Virtualization (RHV) cluster, similar to the one shown in the following diagram.

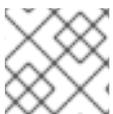


92\_OpenShift\_0520

The installation program uses installer-provisioned infrastructure to automate creating and deploying the cluster.

To install a default cluster, you prepare the environment, run the installation program and answer its prompts. Then, the installation program creates the OpenShift Container Platform cluster.

For an alternative to installing a default cluster, see [Installing a cluster with customizations](#).



#### NOTE

This installation program is available for Linux and macOS only.

### 1.2.1. Prerequisites

- Review details about the [OpenShift Container Platform installation and update](#) processes.
- If you use a firewall, [configure it to allow the sites](#) that your cluster requires access to.

### 1.2.2. Internet and Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.4, you require access to the Internet to install your cluster. The Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, also requires Internet access. If your cluster is connected to the Internet, Telemetry runs automatically, and your cluster is registered to the [Red Hat OpenShift Cluster Manager \(OCM\)](#).

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

You must have Internet access to:

- Access the [Red Hat OpenShift Cluster Manager](#) page 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 content that is required and use it to populate a mirror registry with the packages that you need to install a cluster and generate the installation program. 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.

### 1.2.3. Requirements for the RHV environment

To install and run an OpenShift Container Platform cluster, the RHV environment must meet the following requirements. Not meeting these requirements can cause failures.

The following requirements for CPUs, memory, and storage are based on *default* values multiplied by the default number of virtual machines the installation program creates.

By default, the installation program creates seven machines during the installation process, which includes one bootstrap machine. When the installation program finishes, it deletes the bootstrap machine and frees up its resources. If you perform a custom installation, you can increase the number of virtual machines the installation program creates.



### IMPORTANT

If you increase the resource usage or cluster size in the `install_config.yaml` file, increase these requirements accordingly.

#### Requirements

- The RHV version is 4.3.9 or later.
- The RHV environment has one data center whose state is **Up**.
- The RHV data center contains an RHV cluster.
- The RHV cluster has the following resources exclusively for the OpenShift Container Platform cluster:
  - Minimum 28 vCPUs, which is 4 vCPUs for each of the seven virtual machines created during installation.
  - 112 GiB RAM or more, including:
    - 16 GiB or more for the bootstrap machine, which provides the temporary control plane.

- 16 GiB or more for each of the three control plane machines which provide the control plane.
  - 16 GiB or more for each of the three compute machines, which run the application workloads.
- The RHV storage domain must meet [these etcd backend performance requirements](#).
- In production environments, each virtual machine must have 120 GiB or more, so the storage domain must have 840 GiB or more for the OpenShift Container Platform cluster. In resource-constrained or non-production environments, each virtual machine must have 32 GiB or more, so the storage domain must have 230 GiB or more for the OpenShift Container Platform cluster.
- The RHV cluster must have access to an Internet connection to download images from the Red Hat Ecosystem Catalog during installation and updates, and for the Telemetry service to simplify the subscription and entitlement process.
- The RHV cluster has a virtual network with access to the REST API on the RHV Manager.



#### NOTE

- All together, the hosts must have the required memory and CPU resources **in addition to and aside from** what they use to operate or provide to non-OpenShift Container Platform operations.
- The release cycles of OpenShift Container Platform and RHV are different and versions tested might vary in the future depending on the release dates of both products.
- The bootstrap machine provides a temporary control plane while the installation program creates the OpenShift Container Platform cluster. After it creates the cluster, the installation program removes the bootstrap machine and releases its resources.

### 1.2.4. Verifying the requirements for the RHV environment

Verify that the RHV environment meets the requirements to install and run an OpenShift Container Platform cluster. Not meeting these requirements can cause failures.



#### IMPORTANT

These requirements are based on the default resources the installation program uses to create control plane and compute machines. These resources include vCPUs, memory, and storage. If you change these resources or increase the number of OpenShift Container Platform machines, adjust these requirements accordingly.

#### Procedure

1. Check the RHV version.
  - a. In the RHV Administration Portal, click the ? help icon in the upper-right corner and select **About**.
  - b. In the window that opens, confirm that the **RHV Software Version** is **4.3.9** or higher.

2. Inspect the data center, cluster, and storage.
  - a. In the RHV Administration Portal, click **Compute → Data Centers**.
  - b. Confirm the data center where you plan to install OpenShift Container Platform displays a green up arrow, meaning it is "Up".
  - c. Click the name of that data center.
  - d. In the data center details, on the **Storage** tab, confirm the storage domain where you plan to install OpenShift Container Platform is **Active**.
  - e. Record the **Domain Name** for use later on.
  - f. Confirm **Free Space** has at least 230 GiB.
  - g. Confirm that the storage domain meets [these etcd backend performance requirements](#), which can be [measured using the fio performance benchmarking tool](#).
  - h. In the data center details, click the **Clusters** tab.
    - i. Find the RHV cluster where you plan to install OpenShift Container Platform. Record the cluster name for use later on.
3. Inspect the RHV host resources.
  - a. In the RHV Administration Portal, click **Compute > Clusters**.
  - b. Click the cluster where you plan to install OpenShift Container Platform.
  - c. In the cluster details, click the **Hosts** tab.
  - d. Inspect the hosts and confirm they have a combined total of at least 28 **Logical CPU Cores** available *exclusively* for the OpenShift Container Platform cluster.
  - e. Record the number of available **Logical CPU Cores** for use later on.
  - f. Confirm that these CPU cores are distributed so each of the seven virtual machines created during installation can have four cores.
  - g. Confirm that, all together, the hosts have 112 GiB of **Max free Memory for scheduling new VMs** distributed to meet the requirements for each of the following OpenShift Container Platform machines:
    - 16 GiB required for the bootstrap machine
    - 16 GiB required for each of the three control plane machines
    - 16 GiB for each of the three compute machines
  - h. Record the amount of **Max free Memory for scheduling new VMs** for use later on.
4. Verify that the virtual network for installing OpenShift Container Platform has access to the RHV Manager's REST API. From a virtual machine on this network, use a curl command with the RHV Manager's REST API. Use the following format:

```
curl -k -u <username>@<profile>:<password> \ 1
https://<engine-fqdn>/ovirt-engine/api 2
```

- 
- 1 For **<username>**, specify the user name of an RHV administrator. For **<profile>**, specify the login profile, which you can get by going to the RHV Administration Portal login page and reviewing the **Profile** dropdown list. For **<password>**, specify the admin password.
- 2 For **<engine-fqdn>**, specify the fully qualified domain name of the RHV environment.

For example:

```
$ curl -k -u rhvadmin@internal:pw123 \
https://rhv-env.virtlab.example.com/ovirt-engine/api
```

## 1.2.5. Preparing the network environment on RHV

Configure three static IP addresses for the OpenShift Container Platform cluster and create DNS entries using two of these addresses.

### Procedure

1. Reserve three static IP addresses
  - a. On the network where you plan to install OpenShift Container Platform, identify three static IP addresses that are outside the DHCP lease pool.
  - b. Connect to a host on this network and verify that each of the IP addresses is not in use. For example, use Address Resolution Protocol (ARP) to check that none of the IP addresses have entries:
 

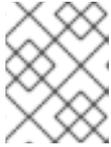
```
$ arp 10.35.1.19
10.35.1.19 (10.35.1.19) -- no entry
```
  - c. Reserve three static IP addresses following the standard practices for your network environment.
  - d. Record these IP addresses for future reference.
2. Create DNS entries for the OpenShift Container Platform REST API and apps domain names using this format:

```
api.<cluster-name>.<base-domain> <ip-address> 1
*.apps.<cluster-name>.<base-domain> <ip-address> 2
```

- 1 For **<cluster-name>**, **<base-domain>**, and **<ip-address>**, specify the cluster name, base domain, and static IP address of your OpenShift Container Platform API.
- 2 Specify the cluster name, base domain, and static IP address of your OpenShift Container Platform apps for Ingress and the load balancer.

For example:

```
api.my-cluster.virtlab.example.com 10.35.1.19
*.apps.my-cluster.virtlab.example.com 10.35.1.20
```

**NOTE**

The third static IP address does not require a DNS entry. The OpenShift Container Platform cluster uses that address for its internal DNS service.

## 1.2.6. Setting up the CA certificate for RHV

Download the CA certificate from the Red Hat Virtualization (RHV) Manager and set it up on the installation machine.

You can download the certificate from a webpage on the RHV Manager or by using a **curl** command.

Later, you provide the certificate to the installation program.

### Procedure

1. Use either of these two methods to download the CA certificate:

- Go to the Manager's webpage, <https://<engine-fqdn>/ovirt-engine/>. Then, under **Downloads**, click the **CA Certificate** link.
- Run the following command:

```
$ curl -k 'https://<engine-fqdn>/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA' -o /tmp/ca.pem 1
```

- 1 For **<engine-fqdn>**, specify the fully qualified domain name of the RHV Manager, such as **rhv-env.virtlab.example.com**.

2. Configure the CA file to grant rootless user access to the Manager. Set the CA file permissions to have an octal value of **0644** (symbolic value: **-rw-r--r--**):

```
$ sudo chmod 0644 /tmp/ca.pem
```

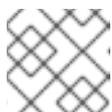
3. For Linux, copy the CA certificate to the directory for server certificates. Use **-p** to preserve the permissions:

```
$ sudo cp -p /tmp/ca.pem /etc/pki/ca-trust/source/anchors/ca.pem
```

4. Add the certificate to the certificate manager for your operating system:

- For macOS, double-click the certificate file and use the **Keychain Access** utility to add the file to the **System** keychain.
- For Linux, update the CA trust:

```
$ sudo update-ca-trust
```

**NOTE**

If you use your own certificate authority, make sure the system trusts it.

### Additional Resources

To learn more, see [Authentication and Security](#) in the RHV documentation.

## 1.2.7. Generating an SSH private key and adding it to the agent

If you want to perform installation debugging or disaster recovery on your cluster, you must provide an SSH key to both your **ssh-agent** and the installation program. You can use this key to access the bootstrap machine in a public cluster to troubleshoot installation issues.



### NOTE

In a production environment, you require disaster recovery and debugging.

You can use this key to SSH into the master nodes as the user **core**. When you deploy the cluster, the key is added to the **core** user's `~/.ssh/authorized_keys` list.

### Procedure

1. If you do not have an SSH key that is configured for password-less authentication on your computer, 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_rsa`, of the SSH key. Do not specify an existing SSH key, as it will be overwritten.

Running this command generates an SSH key that does not require a password in the location that you specified.

2. Start the **ssh-agent** process as a background task:

```
$ eval "$(ssh-agent -s)"
Agent pid 31874
```

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

```
$ ssh-add <path>/<file_name> 1
Identity added: /home/<you>/<path>/<file_name> (<computer_name>)
```

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

### Next steps

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

## 1.2.8. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on a local computer.

### Prerequisites

- You must install the cluster from a computer that uses Linux or macOS.
- You need 500 MB of local disk space to download the installation program.

### Procedure

1. Access the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Navigate to the page for your installation type, download the installation program for your operating system, 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 both the installation program and the files that the installation program creates after you finish installing the cluster.



#### IMPORTANT

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

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

```
$ tar xvf <installation_program>.tar.gz
```

4. From the [Pull Secret](#) page on the Red Hat OpenShift Cluster Manager site, download your installation pull secret as a **.txt** file. 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.

### 1.2.9. 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.

## Procedure

1. Run the installation program:

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

- 1** For **<installation\_directory>**, specify the directory name to store the files that the installation program creates.
- 2** To view different installation details, specify **warn**, **debug**, or **error** instead of **info**.



### IMPORTANT

Specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

Respond to the installation program prompts.

- a. Optional: For **SSH Public Key**, select a password-less public key, such as `~/.ssh/id_rsa.pub`. This key authenticates connections with the new OpenShift Container Platform cluster.



### NOTE

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

- b. For **Platform**, select **ovirt**.
- c. For **Enter oVirt's API endpoint URL**, enter the URL of the RHV API using this format:

```
https://<engine-fqdn>/ovirt-engine/api 1
```

- 1** For **<engine-fqdn>**, specify the fully qualified domain name of the RHV environment.

For example:

```
https://rhv-env.virtlab.example.com/ovirt-engine/api
```

- d. For **Is the installed oVirt certificate trusted?**, enter **Yes** since you have already set up a CA certificate. Otherwise, enter **No**.
- e. For **oVirt's CA bundle**, if you entered **Yes** for the preceding question, copy the certificate content from `/etc/pki/ca-trust/source/anchors/ca.pem` and paste it here. Then, press **Enter** twice. Otherwise, if you entered **No** for the preceding question, this question does not appear.

- f. For **Enter the oVirt engine username**, enter the username and profile of the RHV administrator using this format:

```
<username>@<profile> 1
```

- 1 For **<username>**, specify the username of an RHV administrator. For **<profile>**, specify the login profile, which you can get by going to the RHV Administration Portal login page and reviewing the **Profile** dropdown list. Together, the user name and profile should look similar to this example:

```
admin@internal
```

- g. For **Enter password**, enter the RHV admin password.
- h. For **Select the oVirt cluster**, select the cluster for installing OpenShift Container Platform.
- i. For **Select the oVirt storage domain**, select the storage domain for installing OpenShift Container Platform.
- j. For **Select the oVirt network**, select a virtual network that has access to the RHV Manager REST API.
- k. For **Enter the internal API Virtual IP**, enter the static IP address you set aside for the cluster's REST API.
- l. For **Enter the internal DNS Virtual IP**, enter the static IP address you set aside for the cluster's internal DNS service.
- m. For **Enter the ingress IP**, enter the static IP address you reserved for the wildcard apps domain.
- n. For **Base domain**, enter the base domain of the OpenShift Container Platform cluster. If this cluster is exposed to the outside world, this must be a valid domain recognized by DNS infrastructure. For example, enter: **virtlab.example.com**
- o. For **Cluster name**, enter the name of the cluster. For example, **my-cluster**. Use cluster name from the externally registered/resolvable DNS entries you created for the OpenShift Container Platform REST API and apps domain names. The installation program also gives this name to the cluster in the RHV environment.
- p. For **Pull secret**, copy the pull secret from the **pull-secret.txt** file you downloaded earlier and paste it here. You can also get a copy of the same pull secret from the [Pull Secret](#) page on the Red Hat OpenShift Cluster Manager site.



#### 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.

When the cluster deployment completes, directions for accessing your cluster, including a link to its web console and credentials for the **kubeadmin** user, display in your terminal.



### 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.



### IMPORTANT

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



### IMPORTANT

You have completed the steps required to install the cluster. The remaining steps show you how to verify the cluster and troubleshoot the installation.

## CHAPTER 2. INSTALLING THE CLI BY DOWNLOADING THE BINARY

You can install the OpenShift CLI (**oc**) in order 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 4.4. Download and install the new version of **oc**.

### 2.1. INSTALLING THE CLI ON LINUX

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

#### Procedure

1. Navigate to the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site.
2. Select your infrastructure provider, and, if applicable, your installation type.
3. In the **Command-line interface** section, select **Linux** from the drop-down menu and click **Download command-line tools**.
4. Unpack the archive:

```
$ tar xvzf <file>
```

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

```
$ echo $PATH
```

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

```
$ oc <command>
```

### 2.2. INSTALLING THE CLI ON WINDOWS

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

#### Procedure

1. Navigate to the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site.
2. Select your infrastructure provider, and, if applicable, your installation type.
3. In the **Command-line interface** section, select **Windows** from the drop-down menu and click **Download command-line tools**.
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
```

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

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

## 2.3. INSTALLING THE CLI ON MACOS

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

### Procedure

1. Navigate to the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site.
2. Select your infrastructure provider, and, if applicable, your installation type.
3. In the **Command-line interface** section, select **MacOS** from the drop-down menu and click **Download command-line tools**.
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
```

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

```
$ oc <command>
```

To learn more, see [Getting started with the CLI](#).

## CHAPTER 3. LOGGING IN TO THE CLUSTER

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

### Prerequisites

- Deploy an OpenShift Container Platform cluster.
- Install the **oc** CLI.

### Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation\_directory>**, specify the path to the directory that you stored the installation files in.

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

```
$ oc whoami  
system:admin
```

## 3.1. VERIFYING CLUSTER STATUS

You can verify your OpenShift Container Platform cluster's status during or after installation.

### Procedure

1. In the cluster environment, export the administrator's kubeconfig file:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation\_directory>**, specify the path to the directory that you stored the installation files in.

The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server.

2. View the control plane and compute machines created after a deployment:

```
$ oc get nodes
```

3. View your cluster's version:

```
$ oc get clusterversion
```

4. View your Operators' status:

```
$ oc get clusteroperator
```

5. View all running pods in the cluster:

```
$ oc get pods -A
```

## Troubleshooting

If the installation fails, the installation program times out and displays an error message. To learn more, see [Troubleshooting installation issues](#).

## 3.2. ACCESSING THE OPENSIFT CONTAINER PLATFORM WEB CONSOLE ON RHV

After the OpenShift Container Platform cluster initializes, you can log into the OpenShift Container Platform web console.

### Procedure

1. Optional: In the Red Hat Virtualization (RHV) Administration Portal, open **Compute** → **Cluster**.
2. Verify that the installation program creates the virtual machines.
3. Return to the command line where the installation program is running. When the installation program finishes, it displays the user name and temporary password for logging into the OpenShift Container Platform web console.
4. In a browser, open the URL of the OpenShift Container Platform web console. The URL uses this format:

```
console-openshift-console.apps.<clustername>.<basedomain> 1
```

- 1** For **<clustername>.<basedomain>**, specify the cluster name and base domain.

For example:

```
console-openshift-console.apps.my-cluster.virtlab.example.com
```

## 3.3. TROUBLESHOOTING COMMON ISSUES WITH INSTALLING ON RED HAT VIRTUALIZATION (RHV)

Here are some common issues you might encounter, along with proposed causes and solutions.

### 3.3.1. CPU load increases and nodes go into a Not Ready state

- **Symptom:** CPU load increases significantly and nodes start going into a **Not Ready** state.
- **Cause:** The storage domain latency might be too high, especially for master nodes.
- **Solution:**

Make the nodes ready again by restarting the kubelet service. Enter:

```
$ systemctl restart kubelet
```

Inspect the OpenShift Container Platform metrics service, which automatically gathers and reports on some valuable data such as the etcd disk sync duration. If the cluster is operational, use this data to help determine whether storage latency or throughput is the root issue. If so, consider using a storage resource that has lower latency and higher throughput.

To get raw metrics, enter the following command as kubeadmin or user with cluster-admin privileges:

```
$ oc get --insecure-skip-tls-verify --server=https://localhost:<port> --raw=/metrics`
```

To learn more, see [Exploring Application Endpoints for the purposes of Debugging with OpenShift 4.x](#)

### 3.3.2. Trouble connecting the OpenShift Container Platform cluster API

- **Symptom:** The installation program completes but the OpenShift Container Platform cluster API is not available. The bootstrap virtual machine remains up after the bootstrap process is complete. When you enter the following command, the response will time out.

```
$ oc login -u kubeadmin -p *** <apiurl>
```

- **Cause:** The bootstrap VM was not deleted by the installation program and has not released the cluster's API IP address.
- **Solution:** Use the **wait-for** subcommand to be notified when the bootstrap process is complete:

```
$ ./openshift-install wait-for bootstrap-complete
```

When the bootstrap process is complete, delete the bootstrap virtual machine:

```
$ ./openshift-install destroy bootstrap
```

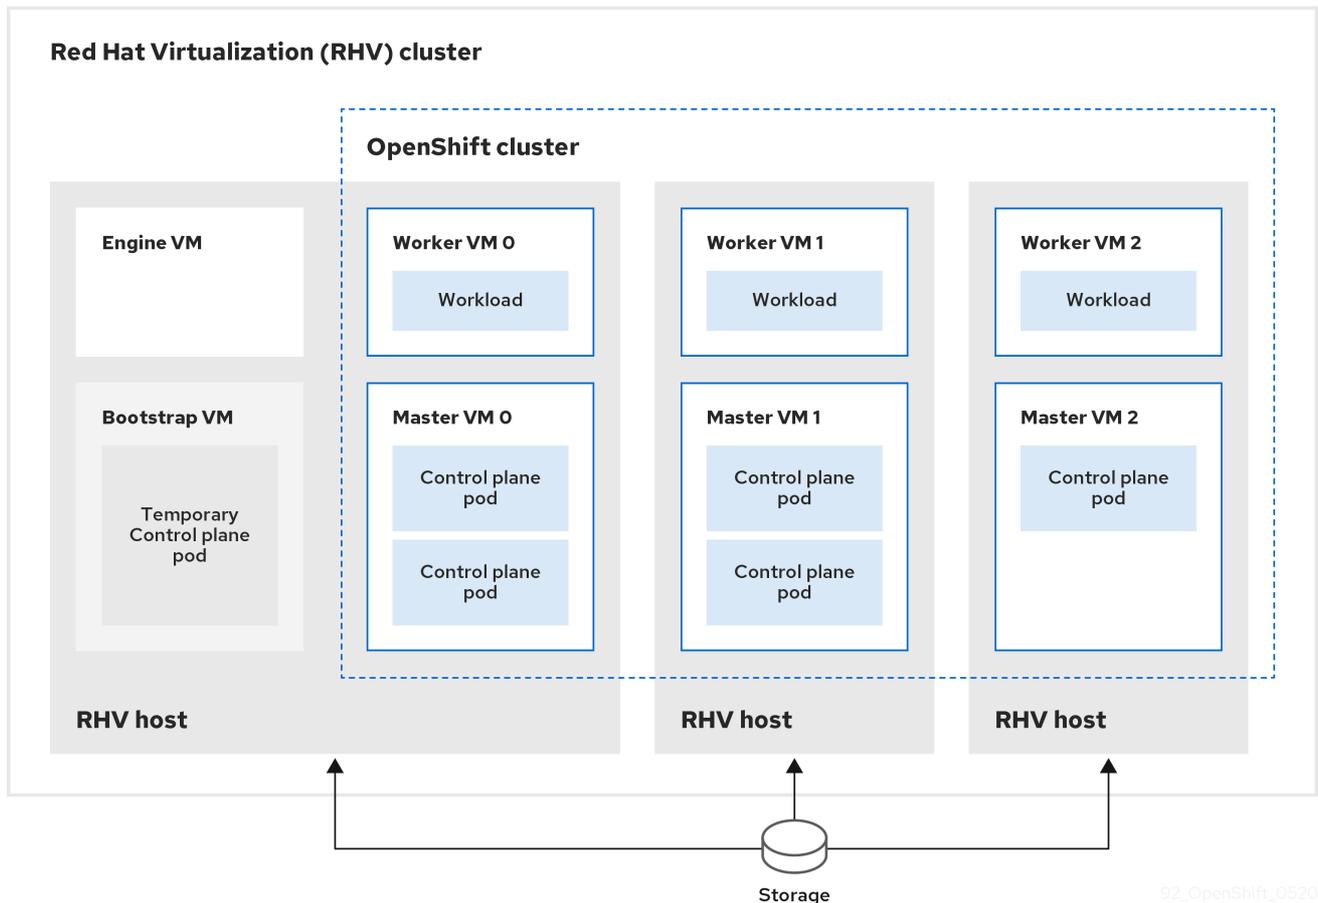
## 3.4. POST-INSTALLATION TASKS

After the OpenShift Container Platform cluster initializes, you can perform the following tasks.

- Optional: After deployment, add or replace SSH keys using the Machine Config Operator (MCO) in OpenShift Container Platform.
- Optional: Remove the **kubeadmin** user. Instead, use the authentication provider to create a user with cluster-admin privileges.

## 3.5. INSTALLING A CLUSTER ON RHV WITH CUSTOMIZATIONS

You can customize and install an OpenShift Container Platform cluster on Red Hat Virtualization (RHV), similar to the one shown in the following diagram.



92\_OpenShift\_0520

The installation program uses installer-provisioned infrastructure to automate creating and deploying the cluster.

To install a customized cluster, you prepare the environment and perform the following steps:

1. Create an installation configuration file, the **install-config.yaml** file, by running the installation program and answering its prompts.
2. Inspect and modify parameters in the **install-config.yaml** file.
3. Make a working copy of the **install-config.yaml** file.
4. Run the installation program with a copy of the **install-config.yaml** file.

Then, the installation program creates the OpenShift Container Platform cluster.

For an alternative to installing a customized cluster, see [Installing a default cluster](#).



#### NOTE

This installation program is available for Linux and macOS only.

### 3.5.1. Prerequisites

- Review details about the [OpenShift Container Platform installation and update](#) processes.
- If you use a firewall, [configure it to allow the sites](#) that your cluster requires access to.

### 3.5.2. Internet and Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.4, you require access to the Internet to install your cluster. The Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, also requires Internet access. If your cluster is connected to the Internet, Telemetry runs automatically, and your cluster is registered to the [Red Hat OpenShift Cluster Manager \(OCM\)](#).

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

You must have Internet access to:

- Access the [Red Hat OpenShift Cluster Manager](#) page 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 content that is required and use it to populate a mirror registry with the packages that you need to install a cluster and generate the installation program. 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.

### 3.5.3. Requirements for the RHV environment

To install and run an OpenShift Container Platform cluster, the RHV environment must meet the following requirements. Not meeting these requirements can cause failures.

The following requirements for CPUs, memory, and storage are based on *default* values multiplied by the default number of virtual machines the installation program creates.

By default, the installation program creates seven machines during the installation process, which includes one bootstrap machine. When the installation program finishes, it deletes the bootstrap machine and frees up its resources. If you perform a custom installation, you can increase the number of virtual machines the installation program creates.



#### IMPORTANT

If you increase the resource usage or cluster size in the `install_config.yaml` file, increase these requirements accordingly.

#### Requirements

- The RHV version is 4.3.9 or later.
- The RHV environment has one data center whose state is **Up**.
- The RHV data center contains an RHV cluster.

- The RHV cluster has the following resources exclusively for the OpenShift Container Platform cluster:
  - Minimum 28 vCPUs, which is 4 vCPUs for each of the seven virtual machines created during installation.
  - 112 GiB RAM or more, including:
    - 16 GiB or more for the bootstrap machine, which provides the temporary control plane.
    - 16 GiB or more for each of the three control plane machines which provide the control plane.
    - 16 GiB or more for each of the three compute machines, which run the application workloads.
- The RHV storage domain must meet [these etcd backend performance requirements](#).
- In production environments, each virtual machine must have 120 GiB or more, so the storage domain must have 840 GiB or more for the OpenShift Container Platform cluster. In resource-constrained or non-production environments, each virtual machine must have 32 GiB or more, so the storage domain must have 230 GiB or more for the OpenShift Container Platform cluster.
- The RHV cluster must have access to an Internet connection to download images from the Red Hat Ecosystem Catalog during installation and updates, and for the Telemetry service to simplify the subscription and entitlement process.
- The RHV cluster has a virtual network with access to the REST API on the RHV Manager.



#### NOTE

- All together, the hosts must have the required memory and CPU resources **in addition to and aside from** what they use to operate or provide to non-OpenShift Container Platform operations.
- The release cycles of OpenShift Container Platform and RHV are different and versions tested might vary in the future depending on the release dates of both products.
- The bootstrap machine provides a temporary control plane while the installation program creates the OpenShift Container Platform cluster. After it creates the cluster, the installation program removes the bootstrap machine and releases its resources.

### 3.5.4. Verifying the requirements for the RHV environment

Verify that the RHV environment meets the requirements to install and run an OpenShift Container Platform cluster. Not meeting these requirements can cause failures.



#### IMPORTANT

These requirements are based on the default resources the installation program uses to create control plane and compute machines. These resources include vCPUs, memory, and storage. If you change these resources or increase the number of OpenShift Container Platform machines, adjust these requirements accordingly.

## Procedure

1. Check the RHV version.
  - a. In the RHV Administration Portal, click the ? help icon in the upper-right corner and select **About**.
  - b. In the window that opens, confirm that the **RHV Software Version** is **4.3.9** or higher.
2. Inspect the data center, cluster, and storage.
  - a. In the RHV Administration Portal, click **Compute → Data Centers**.
  - b. Confirm the data center where you plan to install OpenShift Container Platform displays a green up arrow, meaning it is "Up".
  - c. Click the name of that data center.
  - d. In the data center details, on the **Storage** tab, confirm the storage domain where you plan to install OpenShift Container Platform is **Active**.
  - e. Record the **Domain Name** for use later on.
  - f. Confirm **Free Space** has at least 230 GiB.
  - g. Confirm that the storage domain meets [these etcd backend performance requirements](#), which can be [measured using the fio performance benchmarking tool](#).
  - h. In the data center details, click the **Clusters** tab.
    - i. Find the RHV cluster where you plan to install OpenShift Container Platform. Record the cluster name for use later on.
3. Inspect the RHV host resources.
  - a. In the RHV Administration Portal, click **Compute > Clusters**.
  - b. Click the cluster where you plan to install OpenShift Container Platform.
  - c. In the cluster details, click the **Hosts** tab.
  - d. Inspect the hosts and confirm they have a combined total of at least 28 **Logical CPU Cores** available *exclusively* for the OpenShift Container Platform cluster.
  - e. Record the number of available **Logical CPU Cores** for use later on.
  - f. Confirm that these CPU cores are distributed so each of the seven virtual machines created during installation can have four cores.
  - g. Confirm that, all together, the hosts have 112 GiB of **Max free Memory for scheduling new VMs** distributed to meet the requirements for each of the following OpenShift Container Platform machines:
    - 16 GiB required for the bootstrap machine
    - 16 GiB required for each of the three control plane machines
    - 16 GiB for each of the three compute machines

- h. Record the amount of **Max free Memory for scheduling new VMs** for use later on.
4. Verify that the virtual network for installing OpenShift Container Platform has access to the RHV Manager's REST API. From a virtual machine on this network, use a curl command with the RHV Manager's REST API. Use the following format:

```
curl -k -u <username>@<profile>:<password> \ 1
https://<engine-fqdn>/ovirt-engine/api 2
```

1 For **<username>**, specify the user name of an RHV administrator. For **<profile>**, specify the login profile, which you can get by going to the RHV Administration Portal login page and reviewing the **Profile** dropdown list. For **<password>**, specify the admin password.

2 For **<engine-fqdn>**, specify the fully qualified domain name of the RHV environment.

For example:

```
$ curl -k -u rhvadmin@internal:pw123 \
https://rhv-env.virtlab.example.com/ovirt-engine/api
```

### 3.5.5. Preparing the network environment on RHV

Configure three static IP addresses for the OpenShift Container Platform cluster and create DNS entries using two of these addresses.

#### Procedure

1. Reserve three static IP addresses
  - a. On the network where you plan to install OpenShift Container Platform, identify three static IP addresses that are outside the DHCP lease pool.
  - b. Connect to a host on this network and verify that each of the IP addresses is not in use. For example, use Address Resolution Protocol (ARP) to check that none of the IP addresses have entries:

```
$ arp 10.35.1.19
10.35.1.19 (10.35.1.19) -- no entry
```

- c. Reserve three static IP addresses following the standard practices for your network environment.
  - d. Record these IP addresses for future reference.
2. Create DNS entries for the OpenShift Container Platform REST API and apps domain names using this format:

```
api.<cluster-name>.<base-domain> <ip-address> 1
*.apps.<cluster-name>.<base-domain> <ip-address> 2
```

1 For **<cluster-name>**, **<base-domain>**, and **<ip-address>**, specify the cluster name, base domain, and static IP address of your OpenShift Container Platform API.

- 2 Specify the cluster name, base domain, and static IP address of your OpenShift Container Platform apps for Ingress and the load balancer.

For example:

```
api.my-cluster.virtlab.example.com 10.35.1.19
*.apps.my-cluster.virtlab.example.com 10.35.1.20
```



#### NOTE

The third static IP address does not require a DNS entry. The OpenShift Container Platform cluster uses that address for its internal DNS service.

### 3.5.6. Setting up the CA certificate for RHV

Download the CA certificate from the Red Hat Virtualization (RHV) Manager and set it up on the installation machine.

You can download the certificate from a webpage on the RHV Manager or by using a **curl** command.

Later, you provide the certificate to the installation program.

#### Procedure

1. Use either of these two methods to download the CA certificate:

- Go to the Manager's webpage, <https://<engine-fqdn>/ovirt-engine/>. Then, under **Downloads**, click the **CA Certificate** link.
- Run the following command:

```
$ curl -k 'https://<engine-fqdn>/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA' -o /tmp/ca.pem 1
```

- 1 For **<engine-fqdn>**, specify the fully qualified domain name of the RHV Manager, such as **rhv-env.virtlab.example.com**.

2. Configure the CA file to grant rootless user access to the Manager. Set the CA file permissions to have an octal value of **0644** (symbolic value: **-rw-r--r--**):

```
$ sudo chmod 0644 /tmp/ca.pem
```

3. For Linux, copy the CA certificate to the directory for server certificates. Use **-p** to preserve the permissions:

```
$ sudo cp -p /tmp/ca.pem /etc/pki/ca-trust/source/anchors/ca.pem
```

4. Add the certificate to the certificate manager for your operating system:

- For macOS, double-click the certificate file and use the **Keychain Access** utility to add the file to the **System** keychain.

- For Linux, update the CA trust:

```
$ sudo update-ca-trust
```



#### NOTE

If you use your own certificate authority, make sure the system trusts it.

### Additional Resources

To learn more, see [Authentication and Security](#) in the RHV documentation.

### 3.5.7. Generating an SSH private key and adding it to the agent

If you want to perform installation debugging or disaster recovery on your cluster, you must provide an SSH key to both your **ssh-agent** and the installation program. You can use this key to access the bootstrap machine in a public cluster to troubleshoot installation issues.



#### NOTE

In a production environment, you require disaster recovery and debugging.

You can use this key to SSH into the master nodes as the user **core**. When you deploy the cluster, the key is added to the **core** user's `~/.ssh/authorized_keys` list.

#### Procedure

1. If you do not have an SSH key that is configured for password-less authentication on your computer, 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_rsa`, of the SSH key. Do not specify an existing SSH key, as it will be overwritten.

Running this command generates an SSH key that does not require a password in the location that you specified.

2. Start the **ssh-agent** process as a background task:

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

Agent pid 31874
```

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

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

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

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

### Next steps

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

## 3.5.8. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on a local computer.

### Prerequisites

- You must install the cluster from a computer that uses Linux or macOS.
- You need 500 MB of local disk space to download the installation program.

### Procedure

1. Access the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site. If you have a Red Hat account, log in with your credentials. If you do not, create an account.
2. Navigate to the page for your installation type, download the installation program for your operating system, 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 both the installation program and the files that the installation program creates after you finish installing the cluster.



#### IMPORTANT

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

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

```
$ tar xvf <installation_program>.tar.gz
```

4. From the [Pull Secret](#) page on the Red Hat OpenShift Cluster Manager site, download your installation pull secret as a `.txt` file. 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.

## 3.5.9. Creating the installation configuration file

You can customize the OpenShift Container Platform cluster you install on Red Hat Virtualization (RHV).

## Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

## Procedure

1. Create the **install-config.yaml** file.
  - a. For Red Hat Virtualization (RHV), run the installation program with **sudo**:

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

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



### IMPORTANT

Specify an empty directory. Some installation assets, like bootstrap X.509 certificates have short expiration intervals, so you must not reuse an installation directory. If you want to reuse individual files from another cluster installation, you can copy them into your directory. However, the file names for the installation assets might change between releases. Use caution when copying installation files from an earlier OpenShift Container Platform version.

- b. Respond to the installation program prompts.
  - i. For **SSH Public Key**, select a password-less public key, such as `~/.ssh/id_rsa.pub`. This key authenticates connections with the new OpenShift Container Platform cluster.



### NOTE

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

- ii. For **Platform**, select **ovirt**.
- iii. For **Enter oVirt's API endpoint URL**, enter the URL of the RHV API using this format:

```
https://<engine-fqdn>/ovirt-engine/api 1
```

- 1** For **<engine-fqdn>**, specify the fully qualified domain name of the RHV environment.

For example:

```
https://rhv-env.virtlab.example.com/ovirt-engine/api
```

- iv. For **Is the installed oVirt certificate trusted?**, enter **Yes** since you have already set up a CA certificate. Otherwise, enter **No**.
- v. For **oVirt's CA bundle**, if you entered **Yes** for the preceding question, copy the certificate content from `/etc/pki/ca-trust/source/anchors/ca.pem` and paste it here. Then, press **Enter** twice. Otherwise, if you entered **No** for the preceding question, this question does not appear.
- vi. For **Enter the oVirt engine username**, enter the username and profile of the RHV administrator using this format:

```
<username>@<profile> 1
```

- 1 For **<username>**, specify the username of an RHV administrator. For **<profile>**, specify the login profile, which you can get by going to the RHV Administration Portal login page and reviewing the **Profile** dropdown list. Together, the user name and profile should look similar to this example:

```
admin@internal
```

- vii. For **Enter password**, enter the RHV admin password.
  - viii. For **Select the oVirt cluster**, select the cluster for installing OpenShift Container Platform.
  - ix. For **Select the oVirt storage domain**, select the storage domain for installing OpenShift Container Platform.
  - x. For **Select the oVirt network**, select a virtual network that has access to the RHV Manager REST API.
  - xi. For **Enter the internal API Virtual IP**, enter the static IP address you set aside for the cluster's REST API.
  - xii. For **Enter the internal DNS Virtual IP**, enter the static IP address you set aside for the cluster's internal DNS service.
  - xiii. For **Enter the ingress IP**, enter the static IP address you reserved for the wildcard apps domain.
  - xiv. For **Base domain**, enter the base domain of the OpenShift Container Platform cluster. If this cluster is exposed to the outside world, this must be a valid domain recognized by DNS infrastructure. For example, enter: **virtlab.example.com**
  - xv. For **Cluster name**, enter the name of the cluster. For example, **my-cluster**. Use cluster name from the externally registered/resolvable DNS entries you created for the OpenShift Container Platform REST API and apps domain names. The installation program also gives this name to the cluster in the RHV environment.
  - xvi. For **Pull secret**, copy the pull secret from the **pull-secret.txt** file you downloaded earlier and paste it here. You can also get a copy of the same pull secret from the [Pull Secret](#) page on the Red Hat OpenShift Cluster Manager site.
2. Modify the **install-config.yaml** file. You can find more information about the available parameters in the **Installation configuration parameters** section.

3. Back up the **install-config.yaml** file so that you can use it to install multiple clusters.



### IMPORTANT

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

## 3.5.10. Installation configuration parameters for RHV

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.

The following example is specific to installing OpenShift Container Platform on RHV. It uses numbered callouts to show which parameter values you can edit. Do not modify the parameters values without callouts.

This file is located in the **<installation directory>** you specified when you ran the following command.

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



### NOTE

- Do not copy the following example. Instead, run the installation program to create one.
- You cannot modify these parameters in the **install-config.yaml** file after installation.



### IMPORTANT

If you make customizations that require additional resources, such as adding control plane and compute machines, verify that your RHV environment has enough resources. Otherwise, these customizations might cause the installation to fail.

## Example install-config.yaml configuration file

```
apiVersion: v1
baseDomain: <virtlab.example.com> ❶
compute:
- hyperthreading: Enabled
  name: worker
  platform: {}
  replicas: 3 ❷
controlPlane:
  hyperthreading: Enabled
  name: master
  platform: {}
  replicas: 3 ❸
metadata:
  name: <my-cluster> ❹
```

```

platform:
  ovirt:
    api_vip: <ip-address> 5
    dns_vip: <ip-address> 6
    ingress_vip: <ip-address> 7
    ovirt_cluster_id: <rhv-cluster-id> 8
    ovirt_storage_domain_id: <rhv-storage-domain-id> 9
  publish: External
  pullSecret: |
    <pull-secret> 10
  sshKey: |
    <ssh-public-key> 11

```

- 1 For **<virtlab.example.com>**, specify the base domain of the OpenShift Container Platform cluster.
- 2 Specify **3** or more compute machines. The default value is **3**.
- 3 Specify **3** or more control plane machines. The default value is **3**.
- 4 For **<my-cluster>**, specify the name of the new OpenShift Container Platform cluster.
- 5 For **<ip-address>**, specify the static IP address of the API for which you created the **api**. DNS entry.
- 6 For **<ip-address>**, specify the static IP address of the internal DNS of the OpenShift Container Platform cluster.
- 7 For **<ip-address>**, specify the static IP address of the cluster applications for which you created the **\*.apps**. DNS entry.
- 8 For **<rhv-cluster-id>**, specify an RHV cluster ID.
- 9 For **<rhv-storage-domain-id>**, specify an RHV storage domain ID.
- 10 For **<pull-secret>**, specify your pull secret in JSON format.
- 11 For **<ssh-public-key>**, specify your public SSH key.

### 3.5.11. Deploying the cluster

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



#### IMPORTANT

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

#### Prerequisites

- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

#### Procedure

1. Run the installation program:

```
$ sudo ./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.

When the cluster deployment completes, directions for accessing your cluster, including a link to its web console and credentials for the **kubeadmin** user, display in your terminal.



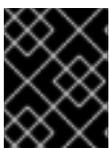
#### IMPORTANT

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.



#### IMPORTANT

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

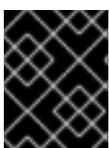


#### IMPORTANT

You have completed the steps required to install the cluster. The remaining steps show you how to verify the cluster and troubleshoot the installation.

### 3.5.12. Installing the CLI by downloading the binary

You can install the OpenShift CLI (**oc**) in order 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 4.4. Download and install the new version of **oc**.

#### 3.5.12.1. Installing the CLI on Linux

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

### Procedure

1. Navigate to the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site.
2. Select your infrastructure provider, and, if applicable, your installation type.
3. In the **Command-line interface** section, select **Linux** from the drop-down menu and click **Download command-line tools**.
4. Unpack the archive:

```
$ tar xvzf <file>
```

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

```
$ echo $PATH
```

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

```
$ oc <command>
```

### 3.5.12.2. Installing the CLI on Windows

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

### Procedure

1. Navigate to the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site.
2. Select your infrastructure provider, and, if applicable, your installation type.
3. In the **Command-line interface** section, select **Windows** from the drop-down menu and click **Download command-line tools**.
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
```

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

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

### 3.5.12.3. Installing the CLI on macOS

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

## Procedure

1. Navigate to the [Infrastructure Provider](#) page on the Red Hat OpenShift Cluster Manager site.
2. Select your infrastructure provider, and, if applicable, your installation type.
3. In the **Command-line interface** section, select **MacOS** from the drop-down menu and click **Download command-line tools**.
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
```

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

```
$ oc <command>
```

### 3.5.13. Logging in to the cluster

You can log in to your cluster as a default system user by exporting the cluster **kubeconfig** file. The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during OpenShift Container Platform installation.

#### Prerequisites

- Deploy an OpenShift Container Platform cluster.
- Install the **oc** CLI.

#### Procedure

1. Export the **kubeadmin** credentials:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation\_directory>**, specify the path to the directory that you stored the installation files in.

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

```
$ oc whoami  
system:admin
```

To learn more, see [Getting started with the CLI](#).

### 3.5.14. Verifying cluster status

You can verify your OpenShift Container Platform cluster's status during or after installation.

## Procedure

1. In the cluster environment, export the administrator's kubeconfig file:

```
$ export KUBECONFIG=<installation_directory>/auth/kubeconfig 1
```

- 1** For **<installation\_directory>**, specify the path to the directory that you stored the installation files in.

The **kubeconfig** file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server.

2. View the control plane and compute machines created after a deployment:

```
$ oc get nodes
```

3. View your cluster's version:

```
$ oc get clusterversion
```

4. View your Operators' status:

```
$ oc get clusteroperator
```

5. View all running pods in the cluster:

```
$ oc get pods -A
```

## Troubleshooting

If the installation fails, the installation program times out and displays an error message. To learn more, see [Troubleshooting installation issues](#).

### 3.5.15. Accessing the OpenShift Container Platform web console on RHV

After the OpenShift Container Platform cluster initializes, you can log into the OpenShift Container Platform web console.

## Procedure

1. Optional: In the Red Hat Virtualization (RHV) Administration Portal, open **Compute → Cluster**.
2. Verify that the installation program creates the virtual machines.
3. Return to the command line where the installation program is running. When the installation program finishes, it displays the user name and temporary password for logging into the OpenShift Container Platform web console.
4. In a browser, open the URL of the OpenShift Container Platform web console. The URL uses this format:

```
console-openshift-console.apps.<clustername>.<basedomain> 1
```

- 1 For `<clustername>.<basedomain>`, specify the cluster name and base domain.

For example:

```
console-openshift-console.apps.my-cluster.virtlab.example.com
```

### 3.5.16. Troubleshooting common issues with installing on Red Hat Virtualization (RHV)

Here are some common issues you might encounter, along with proposed causes and solutions.

#### 3.5.16.1. CPU load increases and nodes go into a **Not Ready** state

- **Symptom:** CPU load increases significantly and nodes start going into a **Not Ready** state.
- **Cause:** The storage domain latency might be too high, especially for master nodes.
- **Solution:**  
Make the nodes ready again by restarting the kubelet service. Enter:

```
$ systemctl restart kubelet
```

Inspect the OpenShift Container Platform metrics service, which automatically gathers and reports on some valuable data such as the etcd disk sync duration. If the cluster is operational, use this data to help determine whether storage latency or throughput is the root issue. If so, consider using a storage resource that has lower latency and higher throughput.

To get raw metrics, enter the following command as kubeadmin or user with cluster-admin privileges:

```
$ oc get --insecure-skip-tls-verify --server=https://localhost:<port> --raw=/metrics`
```

To learn more, see [Exploring Application Endpoints for the purposes of Debugging with OpenShift 4.x](#)

#### 3.5.16.2. Trouble connecting the OpenShift Container Platform cluster API

- **Symptom:** The installation program completes but the OpenShift Container Platform cluster API is not available. The bootstrap virtual machine remains up after the bootstrap process is complete. When you enter the following command, the response will time out.

```
$ oc login -u kubeadmin -p *** <apiurl>
```

- **Cause:** The bootstrap VM was not deleted by the installation program and has not released the cluster's API IP address.
- **Solution:** Use the **wait-for** subcommand to be notified when the bootstrap process is complete:

```
$ ./openshift-install wait-for bootstrap-complete
```

When the bootstrap process is complete, delete the bootstrap virtual machine:

```
$ ./openshift-install destroy bootstrap
```

### 3.5.17. Post-installation tasks

After the OpenShift Container Platform cluster initializes, you can perform the following tasks.

- Optional: After deployment, add or replace SSH keys using the Machine Config Operator (MCO) in OpenShift Container Platform.
- Optional: Remove the **kubeadmin** user. Instead, use the authentication provider to create a user with cluster-admin privileges.

### 3.5.18. Next steps

- [Customize your cluster](#).
- If necessary, you can [opt out of remote health reporting](#).

## 3.6. UNINSTALLING A CLUSTER ON RHV

You can remove an OpenShift Container Platform cluster from Red Hat Virtualization (RHV).

### 3.6.1. Removing a cluster that uses installer-provisioned infrastructure

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

#### Prerequisites

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

#### Procedure

1. From the computer that you used to install the cluster, run the following command:

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

**1** For **<installation\_directory>**, specify the path to the directory that you stored the installation files in.

**2** To view different details, specify **warn**, **debug**, or **error** instead of **info**.



#### NOTE

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

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

