



Quickstart Guide:

# Installing OpenShift Container Platform on Red Hat Virtualization

Revision 1.4, March 27th, 2020



## Table of Contents

<b>Introduction</b>	<b>3</b>
What is Red Hat OpenShift Container Platform on Red Hat Virtualization?	3
Getting help	3
Reporting bugs	3
About the examples in this document	3
<b>Verifying RHV Environment Prerequisites</b>	<b>4</b>
Prerequisites	4
Checking the RHV version	5
Inspecting the data center, cluster, and storage	5
Inspecting the RHV host resources	5
Verifying the virtual network for installing RHOCP	6
<b>Preparing the Network Environment</b>	<b>7</b>
Reserving static IP addresses	7
Configuring DNS	7
<b>Installing RHOCP on RHV</b>	<b>8</b>
Set up the installer software	8
Setting up the SSH keys and CA certificate	8
Running the installer	9
Customizing RHCOS template	12
Post-installation tasks	13
<b>Appendix: Customizing your installation</b>	<b>14</b>
Creating an installation configuration file	14
The installation configuration file	15
Copying install-config.yaml and run the installer	16
<b>Troubleshooting Installation Problems</b>	<b>17</b>
Troubleshooting resources	17
Getting debug information for installation failures	17
Common Issues	18
CPU load increases and nodes go into a Not Ready state	18
Trouble connecting the RHOCP cluster API	19
<b>Revision history</b>	<b>20</b>





## Introduction

# What is Red Hat OpenShift Container Platform on Red Hat Virtualization?

Red Hat OpenShift Container Platform (RHOC) can run master and worker nodes on virtual machines in Red Hat Virtualization (RHV).

The installer that deploys RHOC on RHV automates the process using “installer-provisioned infrastructure.” This installer is available for Linux and macOS only.

## Getting help

For help, contact [Red Hat Customer Experience & Engagement](#).

## Reporting bugs

For issues with the [Developer Preview version](#) of the RHOC on RHV installer, you can create a bug report in [Bugzilla](#) and set its parameters as shown here:

- **Classification:** Red Hat
- **Product:** OpenShift Container Platform
- **Component:** Installer | Cloud compute (for cluster-api-provider-ovirt) | MCO
- **Sub Component:** OpenShift on RHV
- **Title:** Begin with “[ovirt]”

## About the examples in this document

This document uses a consistent set of examples to help illustrate how the information in one section relates to another. You should adapt these examples to your own environment and preferences.



### Verifying RHV Environment Prerequisites

Check that the RHV environment meets the requirements for installing RHOCP on RHV. You can avoid installation failures by fixing unmet requirements. This verification process also helps you gather key information you will use later when you run the installer.

## Prerequisites

The RHV environment must meet the following requirements:

- An RHV administrator account with ClusterAdmin or SuperUser privileges
- RHV 4.3.7.2 or later for the Developer Preview version of the installer software
- RHV 4.3.9.x or later for the RHOCP 4.4 GA version of the installer software
- Has at least one datacenter and cluster
- The RHV cluster must have the following resources exclusively for the RHOCP cluster:
  - Minimum 24 vCPUs (4 vCPUs for each of the 6 virtual machines created during installation)
  - 64 GiB RAM or more
    - 8 GiB required (16 GiB recommended) for the Bootstrap node, which provides the temporary control plane
    - 8 GiB required (16 GiB recommended) for each of the three master nodes which provide the control plane
    - 16 GiB for each of the two worker nodes, which run the application workloads
- The storage domain for the RHOCP cluster must have 175 GiB or more
- An Internet connection to download images from the Red Hat Ecosystem Catalog during installation.
- A virtual network that has access to the REST API on the RHV Manager

**Note:** All together, the hosts must have the required memory and CPU resources *aside from what they use for non-RHOCP operations.*

## Checking the RHV version

1. In the RHV Administration Portal, click the “?” help icon in the upper-right corner and select **About**.
2. In the dialog that opens, confirm that the RHV Software Version meets the requirement.

## Inspecting the data center, cluster, and storage

1. In the RHV Administration Portal, click **Compute > Data Centers**.
2. Click the data center for installing your RHOCP cluster.
3. From the data center details, record the name of the storage domain for installing RHOCP.
4. Confirm that the storage domain has at least 175 GiB of free space. (120 GiB per RHOCP node is recommended for production environments.)
5. Click **Compute > Clusters**.
6. Identify the RHV cluster for installing RHOCP and record its name.
7. For this cluster, confirm that the data center column contains the data center for installing RHOCP.

Note: To customize the disk size, see [Customizing RHCOS template](#).

## Inspecting the RHV host resources

1. In the RHV Administration Portal, click **Compute > Clusters**.
2. Click the cluster for installing RHOCP to display its details.
3. In the cluster details, click the **Hosts** tab.
4. Confirm that the hosts have a combined total of at least 24 CPU cores available *exclusively* for the RHOCP cluster. These CPUs must be distributed so that each of the 6 virtual machines created during installation can have 4 CPU cores.
5. Confirm that, all together, the hosts have 64 GiB free memory for the following RHOCP nodes:
  - 8 GiB required (16 GiB recommended) for the Bootstrap node, which provides the temporary control plane



- 8 GiB required (16 GiB recommended) for each of the three master nodes which provide the control plane
- 16 GiB for each of the two worker nodes, which run the application workloads

**Note:** The installer removes the bootstrap node when it is no longer needed and releases its resources.

## Verifying the virtual network for installing RHOCP

Verify that the virtual network for installing RHOCP 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 <ovirt-engine username>:<password>  
https://<engine-fqdn>/ovirt-engine/api
```

For example:

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



## Preparing the Network Environment

Configure three static IP addresses for the RHOCP cluster and create DNS entries for two of them.

## Reserving static IP addresses

1. Identify three static IP addresses on the network for installing RHOCP that are outside the DHCP lease pool. This document uses three example IP addresses:
  - 10.35.1.19 for the RHOCP cluster's REST API
  - 10.35.1.20 for the RHOCP cluster's apps (ingress/load balancer)
  - 10.35.1.21 for the RHOCP cluster's internal DNS
2. Connect to a host on this network and verify that each of the IP addresses is not in use. For example, use `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
```

3. Reserve three static IP addresses following the standard practices for your network environment.

## Configuring DNS

Create DNS entries for RHOCP REST API and apps domain names using this format:

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

For example:

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

The third static IP address (10.35.1.21 in our examples) does not require a DNS entry. The RHOCP cluster uses that address for its internal DNS service.





Installing RHOCP on RHV

## Set up the installer software

Select the system where you will run the installer. Often, this system is a bastion host. This document calls the system “installer machine” from here on.

### Prerequisites

An installer machine that:

- Runs Linux or macOS.
- Has an Internet connection and browser.
- Has a fast network connection to the RHV Manager.

### Procedure

1. On the installer machine, [download the installer](#) for your Linux or macOS operating system.

2. Extract the installer from the archive:

```
$ tar -xf ~/Downloads/<openshift-install-your_os>.tar.gz
```

3. Optionally, [download and extract the client software](#) for your operating system.

```
$ tar -xf ~/Downloads/<openshift-client-your_os>.tar.gz
```

## Setting up the SSH keys and CA certificate

Identify an existing SSH key pair or create a new key pair. The installer uses this key when it creates the new RHOCP cluster. Only the key owner can connect to the RHOCP cluster.

### Procedure

1. To recreate a new key pair, enter:

```
$ cd ~/.ssh/  
$ mkdir key_backup  
$ copy id_rsa* key_backup  
$ ssh-keygen -t rsa -C "your\_email@example.com"
```



2. Download the CA certificate from the RHV Manager to the installer machine:

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

For example:

```
curl -k  
'https://rhv-env.virtlab.example.com/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA' -o /tmp/ca.pem
```

3. Copy the CA certificate:

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

4. Update the CA trust:

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

**Note:** If you're using your own certificate authority, make sure the system trusts it.

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

## Running the installer

Run the installer and respond to the interactive prompts. When you've the prompts the installer creates the RHOCP cluster.

In the background, the installer creates a configuration file from your responses, installs the RHOCP cluster, and deletes the configuration file.

**Note:** There is an alternative workflow that separates the installation process into distinct steps. It enables you to troubleshoot issues, customize the installation, and create templates for multiple installations. To learn more, see [Appendix: Customizing your installation](#).

### Prerequisites

You have [verified the prerequisites](#) and [prepared the installer machine](#), so you have:

- An SSH public key, `~/.ssh/id_rsa.pub`
- A CA certificate, `/etc/pki/ca-trust/source/anchors/ca.pem`

## Procedure

1. On the installer machine, go to the directory that has the installer and run it.

```
$ sudo ./openshift-install create cluster --dir=<assets-dir>
```

*<assets-dir>* is the name of a new directory the installer creates for storing installation information, such as log files. If you specify an existing directory, the installer will look for the installation configuration file described in [Appendix: Customizing your installation](#).

2. Respond to the following installer prompts for information. For more information about a prompt, enter “?”.

**Note:** The order of the prompts may be different from the order shown here.

- **SSH Public Key:** Select the public key on the installer machine, such as `~/.ssh/id_rsa.pub`, that will be used to authenticate connections with the new OpenShift cluster. Only the owner of this public key can ssh to the new OpenShift nodes. Optionally, after the installation, you can add and replace keys using the Machine Config Operator (MCO).

- **Platform: ovirt**

- **Enter oVirt's API endpoint URL:** Enter the URL of the oVirt API using this format:

```
<ovirt-admin>:<password> https://<engine-fqdn>/ovirt-engine/api
```

For example:

```
admin:pw123 https://rhv-env.virtlab.example.com/ovirt-engine/api
```

- **Is the installed oVirt certificate trusted?:** Enter **Yes**. (You provided a CA certificate, `ca.pem`.)
- **oVirt's CA bundle:** If you entered **Yes** for the preceding question, copy the certificate 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.
- **Enter the oVirt engine username:** Enter the RHV admin user name and profile using this format:

```
<ovirt-admin>@<ovirt-login-profile>
```



For example:

```
admin@internal
```

- **Enter password:** Enter the RHV admin password.
- **Select the oVirt cluster:** Select the cluster for installing RHOCP.
- **Select the oVirt storage domain:** Select the storage domain for installing RHOCP.
- **Select the oVirt network:** Select a virtual network that has access to the RHV Manager REST API.
- **Enter the internal API Virtual IP:** Enter the static IP address you set aside for the RHOCP cluster's REST API. Such as `10.35.1.19` from the example in which you configured DNS for `api.ocp-cluster.virtlab.example.com`.
- **Enter the internal DNS Virtual IP:** Enter the static IP address you set aside for the RHOCP cluster's internal DNS service, such as `10.35.1.21` from the previous example in which you reserved the IP addresses.
- **Enter the ingress IP:** Enter the static IP address you reserved for the wildcard apps domain. Such as `10.35.1.20` from the example in which you configured DNS for `*.apps.ocp-cluster.virtlab.example.com`).
- **Base domain:** Enter the base domain of the RHOCP 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`
- **Cluster name:** Enter the name of the cluster. For example, `ocp-cluster`. This value should match the cluster name you used in the externally registered/resolvable API and ingress DNS entries. The installer also uses this to name the cluster in the RHV environment.
- **Pull secret:** [Copy the pull secret](#) and paste it here. The installer uses this pull secret to get images from the Red Hat Container Registry.

**Note:** At this point in the sequence, the installer begins creating the cluster. Optionally, in the RHV Administration Portal, watch the installer create the virtual machines in **Compute > Cluster**.



3. While the installer is running, visit the RHV Administration Portal and verify that the installer creates and runs the virtual machines for the RHOCP nodes.
4. Get the user name and temporary password the installer displays when the new OpenShift cluster is up and running.
5. [Access the OpenShift web console](#) to verify the cluster is running.

The web console URL uses the following format:

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

For example:

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

You have finished installing the RHOCP cluster on RHV.

**Note:** If something goes wrong with the installation, the installer times out and displays an error message. To learn more, see [Troubleshooting Installation Problems](#).

## Customizing RHCOS template

**Note:** This is a temporary hack for the Developer Preview version of the installer. A better solution should be available for a later version of RHOCP.

Optionally, you can customize a virtual machine template to create virtual machines with non-default characteristics, such as 120 GiB disks.

1. From the installer machine, open the RHV Administration Portal and go to **Compute > Virtual machines**.
2. Check that the cluster has sufficient resources for the customizations you intend to make.
3. Run the installer from the command line. For example:

```
openshift-install create cluster --dir template-customization
```

4. Watch the Administration Portal for new virtual machines being created by the installer. When you see them, enter **ctrl+c** on the command line to interrupt the installer.
5. Connect to the created virtual machines and locate the template that was just created. For example, with a cluster named **test**, the virtual machines and template would be named something like **test-vjhb-master-0** and a template named something like **test-vjhb-rhcos**.



6. Rename the template to something like **rhcos-ocp-4.4**.
7. Create a virtual machine from the template.
8. Change the virtual machine characteristics. For example, extend its disk to 120 GiB.
9. Make a new template from the updated virtual machine and name it something like **rhcos-ocp-4.4-120gb**.
10. On the command line, set the image override shell variable to use this new template. For example:

```
$ export OPENSIFT_INSTALL_OS_IMAGE_OVERRIDE=rhcos-ocp-4.4-120gb
```

11. Rerun the installer.

```
$ openshift-install create cluster --dir ocp-production
```

12. Verify that the virtual machines in the new cluster contain the expected customizations.

## Post-installation tasks

- Optionally, after deployment, you can add or replace SSH keys using the Machine Config Operator (MCO) in OpenShift.
- Optionally, [remove the kubeadmin user](#). Instead, use the authentication provider to create a user with cluster-admin privileges.



## Appendix: Customizing your installation

You can customize your installation of RHOSP. Here's an overview of the steps:

1. Creating an installation configuration file, `install-config.yaml`, by answering the interactive prompts.
2. Inspecting and modifying parameters in `install-config.yaml`.
3. Copying `install-config.yaml`.
4. Running the installer with a copy of `install-config.yaml` (no prompts).

## Creating an installation configuration file

### Prerequisites

You have [verified the prerequisites](#) and [prepared the installer machine](#), so you have:

- An SSH public key, `~/.ssh/id_rsa.pub`
- A CA certificate, `/etc/pki/ca-trust/source/anchors/ca.pem`

### Procedure

Use the `create install-config` subcommand to create a configuration file without installing an RHOC cluster.

1. Change to the directory with the installer and enter:

```
$ sudo ./openshift-install create install-config --dir=assets-dir
```

*assets-dir* is the assets directory where the installer creates the new installation configuration file, `install-config.yaml`. For example, name the asset directory `my_custom_config`.

2. Respond to the prompts detailed in [Run the Installer](#).

When you finish, the installer uses your responses to create the install configuration file, `install-config.yaml`, in the assets directory. For example:

```
~/my_custom_config/install-config.yaml
```



## The installation configuration file

The installer uses your entries to create `install-config.yaml` in the `assets` directory. Here's an example of that file. You can edit the values of parameters in bold text.

**Note:** Do not copy and paste this example. Its details may become outdated. Instead, use your file from [Create an installation configuration file](#).

```
apiVersion: v1
baseDomain: <base domain of the RHOCP cluster, such as "virtlab.example.com">
compute:
- hyperthreading: Enabled
  name: worker
  platform: {}
  replicas: <2 or more>
controlPlane:
  hyperthreading: Enabled
  name: master
  platform: {}
  replicas: <3 or more>
metadata:
  name: <mycluster>
platform:
  ovirt:
    api_vip: <Virtual IP for internal API>
    dns_vip: <Virtual IP for OpenShift DNS>
    ingress_vip: <Ingress virtual IP>
    ovirt_cluster_id: <UUID of RHV cluster that will host your RHOCP cluster>
    ovirt_storage_domain_id: <UUID of RHV storage domain that will host your
RHOCP cluster>
publish: External
pullSecret: |
  <Your pull secret in JSON format goes here>
sshKey: |
  <Public SSH keys that are to be stored on OpenShift RHCOS nodes go here>
```





## Copying install-config.yaml and run the installer

**Important:** When you run the installer, it deletes the `install-config.yaml` installer file. To preserve your customized file, make a working copy of the file in a separate directory.

1. When you have finished customizing `install-config.yaml`, make a working copy of the asset directory. For example:

```
$ cp -r ~/my_custom_config/ ~/my_working_dir/
```

2. Run the installer and specify the working directory. For example:

```
$ sudo ./openshift-install create cluster --my_working_dir
```

The installer runs, using the information in `install-config.yaml` instead of prompting you for information. It also deletes `install-config.yaml` from the working directory.

You have finished installing the RHOC cluster on RHV.



Troubleshooting Installation Problems

## Troubleshooting resources

See [Debugging OpenShift 4.x: Debugging an Install](#).

Some installation failures are caused by installing in an environment that does not meet the prerequisites. To find and fix these issues, see [Verifying Prerequisites and Gathering Information](#).

## Getting debug information for installation failures

You can do the following to debug a failed installation:

- Run the installer with `--log-level=debug`:

```
$ sudo ./openshift-install create cluster --dir=assets-dir
--log-level=debug
```

- Inspect the debug messages. To see the log from a past installation, look in the assets directory for the hidden file, `.openshift_install.log`. For example, enter:

```
$ cat ~/assets-dir/.openshift_install.log
```

- Gather the most relevant error messages from the installer. For example:

```
$ sudo ./openshift-install gather bootstrap
```

```
FATAL bootstrap host address and at least one control plane host address
must be provided
```

- Check for timeout errors.

Sometimes, the installer times out while the RHOCPC cluster is spinning up in the background. The installer has three stages where it waits for cluster events. As the events happen, the installer displays status messages:

- **Bootstrap:** After 20 minutes, "Waiting up to %v for bootstrapping to complete..."
- **Initialization:** After 30 minutes, "Waiting up to %v for the cluster at %s to initialize..."
- **Console route created:** After 10 minutes, "Waiting up to %v for the openshift-console route to be created..."



Then use the `wait-for` subcommands to troubleshoot issues related to these stages.

To wait until cluster bootstrapping has completed:

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

To wait until the cluster is ready:

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

## Common Issues

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 may be too high, especially for master nodes.
- **Solution:**

- Make the nodes Ready again by restarting the kubelet service. Enter:

```
$ systemctl restart kubelet
```

- Inspect the RHOCP 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`
```

See [Exploring Application Endpoints for the purposes of Debugging with OpenShift 4.x](#)



## Trouble connecting the RHOCP cluster API

- **Symptom:** The installer completes but the RHOCP 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 installer 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:

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

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

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



#### Revision history

- Revision 1.0, March 18th, 2020, 2020: Initial publication
- Revision 1.1, March 19th, 2020: Copy edits and minor additions.
- Revision 1.2, March 20th, 2020: Copy edits.
- Revision 1.3, March 22nd, 2020: CA certificate issue fixed. Changed “Is the installed oVirt certificate trusted?” to “Yes.”
- Revision 1.4, March 27th, 2020: Added the [Customizing RHCOS template](#) topic.