# OpenShift Container Platform 4.2

## Installing on Azure

Installing OpenShift Container Platform 4.2 Azure clusters

Last Updated: 2020-05-15

# OpenShift Container Platform 4.2 Installing on Azure

Installing OpenShift Container Platform 4.2 Azure clusters

## Legal Notice

Copyright © 2020 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 ® is the registered trademark of Linus Torvalds in the United States and other countries.

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This document provides instructions for installing and uninstalling OpenShift Container Platform 4.2 clusters on Microsoft Azure.

# Table of Contents

CHAPTER 1. INSTALLING ON AZURE ..... 3
1.1. CONFIGURING AN AZURE ACCOUNT 3
1.1.1. Azure account limits 3
1.1.2. Configuring a public DNS zone in Azure 5
1.1.3. Increasing Azure account limits 6
1.1.4. Required Azure roles 7
1.1.5. Creating a service principal 7
1.1.6. Supported Azure regions 9
1.2. INSTALLING A CLUSTER QUICKLY ON AZURE 10
1.2.1. Internet and Telemetry access for OpenShift Container Platform 11
1.2.2. Generating an SSH private key and adding it to the agent 11
1.2.3. Obtaining the installation program 12
1.2.4. Deploy the cluster 13
1.2.5. Installing the CLI 15
1.2.6. Logging in to the cluster 16
1.3. INSTALLING A CLUSTER ON AZURE WITH CUSTOMIZATIONS 16
1.3.1. Internet and Telemetry access for OpenShift Container Platform 16
1.3.2. Generating an SSH private key and adding it to the agent 17
1.3.3. Obtaining the installation program 18
1.3.4. Creating the installation configuration file 19
1.3.4.1. Installation configuration parameters 20
1.3.4.2. Sample customized install-config.yaml file for Azure 24
1.3.5. Deploy the cluster 25
1.3.6. Installing the CLI 26
1.3.7. Logging in to the cluster 27
1.4. INSTALLING A CLUSTER ON AZURE WITH NETWORK CUSTOMIZATIONS 28
1.4.1. Internet and Telemetry access for OpenShift Container Platform 28
1.4.2. Generating an SSH private key and adding it to the agent 29
1.4.3. Obtaining the installation program 29
1.4.4. Creating the installation configuration file 30
1.4.4.1. Installation configuration parameters 32
1.4.4.2. Network configuration parameters 35
1.4.4.3. Sample customized install-config.yaml file for Azure 36
1.4.5. Modifying advanced network configuration parameters 38
1.4.6. Cluster Network Operator custom resource (CR) 39
1.4.6.1. Configuration parameters for OpenShift SDN 40
1.4.6.2. Configuration parameters for Open Virtual Network (OVN) SDN 41
1.4.6.3. Cluster Network Operator example CR 41
1.4.7. Deploy the cluster 41
1.4.8. Installing the CLI 42
1.4.9. Logging in to the cluster 43
1.5. UNINSTALLING A CLUSTER ON AZURE 44
1.5.1. Removing a cluster that uses installer-provisioned infrastructure 44

# CHAPTER 1. INSTALLING ON AZURE

## 1.1. CONFIGURING AN AZURE ACCOUNT

Before you can install OpenShift Container Platform, you must configure a Microsoft Azure account.

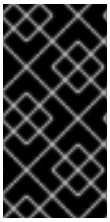

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see Resolve reserved resource name errors in the Azure documentation.

### 1.1.1. Azure account limits

The OpenShift Container Platform cluster uses a number of Microsoft Azure components, and the default Azure subscription and service limits, quotas, and constraints affect your ability to install OpenShift Container Platform clusters.

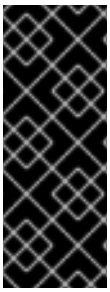

**IMPORTANT**

Default limits vary by offer category types, such as Free Trial and Pay-As-You-Go, and by series, such as Dv2, F, and G. For example, the default for Enterprise Agreement subscriptions is 350 cores.

Check the limits for your subscription type and if necessary, increase quota limits for your account before you install a default cluster on Azure.

The following table summarizes the Azure components whose limits can impact your ability to install and run OpenShift Container Platform clusters.

| Component | Number of components required by default | Default Azure limit | Description |
|---|---|---|---|

| Component | Number of components required by default | Default Azure limit | Description |
|---|---|---|---|
| vCPU | 34 | 20 per region | A default cluster requires 34 vCPUs, so you must increase the account limit.<br><br>By default, each cluster creates the following instances:<br><br>• One bootstrap machine, which is removed after installation<br>• Three control plane machines<br>• Three compute machines<br><br>Because the bootstrap machine uses **Standard_D4s_v3** machines, which use 4 vCPUS, the control plane machines use **Standard_D8s_v3** virtual machines, which use 8 vCPUs, and the worker machines use **Standard_D2s_v3** virtual machines, which use 2 vCPUs, a default cluster requires 34 vCPUs.<br><br>To deploy more worker nodes, enable autoscaling, deploy large workloads, or use a different instance type, you must further increase the vCPU limit for your account to ensure that your cluster can deploy the machines that you require.<br><br>By default, the installation program distributes control plane and compute machines across all availability zones within a region. To ensure high availability for your cluster, select a region with at least three availablity zones. If your region contains fewer than three availability zones, the installation program places more than one control plane machine in the available zones. |
| VNet | 1 | 1000 per region | Each default cluster requires one Virtual Network (VNet), which contains two subnets. |
| Network interfaces | 6 | 65,536 per region | Each default cluster requires six network interfaces. If you create more machines or your deployed workloads create load balancers, your cluster uses more network interfaces. |

| Component | Number of components required by default | Default Azure limit | Description |
|---|---|---|---|
| Network security groups | 2 | 5000 | Each default cluster Each cluster creates network security groups for each subnet in the VNet. The default cluster creates network security groups for the control plane and for the compute node subnets:<br>**controlplane** Allows the control plane machines to be reached on port 6443 from anywhere<br>**node** Allows worker nodes to be reached from the internet on ports 80 and 443 |
| Network load balancers | 3 | 1000 per region | Each cluster creates the following load balancers:<br>**default** Public IP address that load balances requests to ports 80 and 443 across worker machines<br>**internal** Private IP address that load balances requests to ports 6443 and 22623 across control plane machines<br>**external** Public IP address that load balances requests to port 6443 across control plane machines<br>If your applications create more Kubernetes LoadBalancer Service objects, your cluster uses more load balancers. |
| Public IP addresses | 3 | | Each of the two public load balancers uses a public IP address. The bootstrap machine also uses a public IP address so that you can SSH into the machine to troubleshoot issues during installation. The IP address for the bootstrap node is used only during installation. |
| Private IP addresses | 7 | | The internal loadbalancer, each of the three control plane machines, and each of the three worker machines each use a private IP address. |

### 1.1.2. Configuring a public DNS zone in Azure

To install OpenShift Container Platform, the Microsoft Azure account you use must have a dedicated public hosted DNS zone in your account. This zone must be authoritative for the domain. This service provides cluster DNS resolution and name lookup for external connections to the cluster.

**Procedure**

1. Identify your domain, or subdomain, and registrar. You can transfer an existing domain and registrar or obtain a new one through Azure or another source.

**NOTE**

For more information about purchasing domains through Azure, see Buy a custom domain name for Azure App Service in the Azure documentation.

2. If you are using an existing domain and registrar, migrate its DNS to Azure. See Migrate an active DNS name to Azure App Service in the Azure documentation.

3. Configure DNS for your domain. Follow the steps in the Tutorial: Host your domain in Azure DNS in the Azure documentation to create a public hosted zone for your domain or subdomain, extract the new authoritative name servers, and update the registrar records for the name servers that your domain uses.
   Use an appropriate root domain, such as **openshiftcorp.com**, or subdomain, such as **clusters.openshiftcorp.com**.

4. If you use a subdomain, follow your company's procedures to add its delegation records to the parent domain.

### 1.1.3. Increasing Azure account limits

To increase an account limit, file a support request on the Azure portal.

**NOTE**

You can increase only one type of quota per support request.

**Procedure**

1. From the Azure portal, click **Help + support** in the lower left corner.

2. Click **New support request** and then select the required values:

   a. From the **Issue type** list, select **Service and subscription limits (quotas)**.

   b. From the **Subscription** list, select the subscription to modify.

   c. From the **Quota type** list, select the quota to increase. For example, select **Compute-VM (cores-vCPUs) subscription limit increases** to increase the number of vCPUs, which is required to install a cluster.

   d. Click **Next: Solutions**.

3. On the PROBLEM DETAILS page, provide the required information for your quota increase:

   a. Click **Provide details** and provide the required details in the "Quota details" window.

b. In the SUPPORT METHOD and CONTACT INFO sections, provide the issue severity and your contact details.

4. Click **Next: Review + create** and then click **Create**.

### 1.1.4. Required Azure roles

Your Microsoft Azure account must have the following roles for the subscription that you use:

- **User Access Administrator**

To set roles on the Azure portal, see the Manage access to Azure resources using RBAC and the Azure portal in the Azure documentation.

### 1.1.5. Creating a service principal

Because OpenShift Container Platform and its installation program must create Microsoft Azure resources through Azure Resource Manager, you must create a service principal to represent it.

#### Prerequisites

- Install or update the Azure CLI.
- Install the **jq** package.
- Your Azure account has the required roles for the subscription that you use.

#### Procedure

1. Log in to the Azure CLI:

```
$ az login
```

Log in to Azure in the web console by using your credentials.

2. If your Azure account uses subscriptions, ensure that you are using the right subscription.

a. View the list of available accounts and record the **tenantId** value for the subscription you want to use for your cluster:

```
$ az account list --refresh
[
  {
    "cloudName": "AzureCloud",
    "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
    "isDefault": true,
    "name": "Subscription Name",
    "state": "Enabled",
    "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee",
    "user": {
      "name": "you@example.com",
      "type": "user"
    }
  }
]
```

b. View your active account details and confirm that the **tenantId** matches the subscription you want to use:

```
$ az account show
{
  "environmentName": "AzureCloud",
  "id": "9bab1460-96d5-40b3-a78e-17b15e978a80",
  "isDefault": true,
  "name": "Subscription Name",
  "state": "Enabled",
  "tenantId": "6057c7e9-b3ae-489d-a54e-de3f6bf6a8ee", (1)
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

(1) Ensure that the value of the **tenantId** parameter is the UUID of the correct subscription.

c. If you are not using the right subscription, change the active subscription:

```
$ az account set -s <id> (1)
```

(1) Substitute the value of the **id** for the subscription that you want to use for **<id>**.

d. If you changed the active subscription, display your account information again:

```
$ az account show

{
  "environmentName": "AzureCloud",
  "id": "33212d16-bdf6-45cb-b038-f6565b61edda",
  "isDefault": true,
  "name": "Subscription Name",
  "state": "Enabled",
  "tenantId": "8049c7e9-c3de-762d-a54e-dc3f6be6a7ee",
  "user": {
    "name": "you@example.com",
    "type": "user"
  }
}
```

3. Record the values of the **tenantId** and **id** parameters from the previous output. You need these values during OpenShift Container Platform installation.

4. Create the service principal for your account:

```
$ az ad sp create-for-rbac --role Contributor --name <service_principal> (1)
Changing "<service_principal>" to a valid URI of "http://<service_principal>", which is the
required format used for service principal names
Retrying role assignment creation: 1/36
Retrying role assignment creation: 2/36
```

```
Retrying role assignment creation: 3/36
Retrying role assignment creation: 4/36
{
  "appId": "8bd0d04d-0ac2-43a8-928d-705c598c6956",
  "displayName": "<service_principal>",
  "name": "http://<service_principal>",
  "password": "ac461d78-bf4b-4387-ad16-7e32e328aec6",
  "tenant": "6048c7e9-b2ad-488d-a54e-dc3f6be6a7ee"
}
```

1 Replace **<service_principal>** with the name to assign to the service principal.

5. Record the values of the **appId** and **password** parameters from the previous output. You need these values during OpenShift Container Platform installation.

6. Grant additional permissions to the service principal. The service principal requires the legacy **Azure Active Directory Graph → Application.ReadWrite.OwnedBy** permission and the **User Access Administrator** role for the cluster to assign credentials for its components.

   a. To assign the **User Access Administrator** role, run the following command:

   ```
   $ az role assignment create --role "User Access Administrator" \
       --assignee-object-id $(az ad sp list --filter "appId eq '<appId>'" \ 1
         | jq '.[0].objectId' -r)
   ```

   1 Replace **<appId>** with the **appId** parameter value for your service principal.

   b. To assign the **Azure Active Directory Graph** permission, run the following command:

   ```
   $ az ad app permission add --id <appId> \ 1
       --api 00000002-0000-0000-c000-000000000000 \
       --api-permissions 824c81eb-e3f8-4ee6-8f6d-de7f50d565b7=Role

       Invoking "az ad app permission grant --id 46d33abc-b8a3-46d8-8c84-f0fd58177435 --
   api 00000002-0000-0000-c000-000000000000" is needed to make the change effective
   ```

   1 Replace **<appId>** with the **appId** parameter value for your service principal.

   For more information about the specific permissions that you grant with this command, see the GUID Table for Windows Azure Active Directory Permissions .

   c. Approve the permissions request. If your account does not have the Azure Active Directory tenant administrator role, follow the guidelines for your organization to request that the tenant administrator approve your permissions request.

   ```
   $ az ad app permission grant --id <appId> \ 1
       --api 00000002-0000-0000-c000-000000000000
   ```

   1 Replace **<appId>** with the **appId** parameter value for your service principal.

### 1.1.6. Supported Azure regions

The installation program dynamically generates the list of available Microsoft Azure regions based on your subscription. The following Azure regions were tested and validated in OpenShift Container Platform version 4.2.0:

- centralus (Central US)
- eastus (East US)
- eastus2 (East US 2)
- northcentralus (North Central US)
- southcentralus (South Central US)
- westcentralus (West Central US)
- westus (West US)
- westus2 (West US 2)
- uksouth (UK South)
- ukwest (UK West)
- francecentral (France Central)
- northeurope (North Europe)
- westeurope (West Europe)
- japaneast (Japan East)
- japanwest (Japan West)
- koreacentral (Korea Central)
- koreasouth (Korea South)
- eastasia (East Asia)
- southeastasia (Southeast Asia)
- southindia (South India)
- centralindia (Central India)
- westindia (West India)
- uaenorth (UAE North)

**Next steps**

- Install an OpenShift Container Platform cluster on Azure. You can install a customized cluster or quickly install a cluster with default options.

## 1.2. INSTALLING A CLUSTER QUICKLY ON AZURE

In OpenShift Container Platform version 4.2, you can install a cluster on Microsoft Azure that uses the default configuration options.

**Prerequisites**

- Review details about the OpenShift Container Platform installation and update processes.
- Configure an Azure account to host the cluster and determine the tested and validated region to deploy the cluster to.
- If you use a firewall, you must configure it to allow the sites that your cluster requires access to.

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

In OpenShift Container Platform 4.2, you require access to the internet to install and entitle 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 . From there, you can allocate entitlements to your cluster.

You must have internet access to:

- Access the Red Hat OpenShift Cluster Manager page to download the installation program and perform subscription management and entitlement. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster. If the Telemetry service cannot entitle your cluster, you must manually entitle it on the Cluster registration page.
- 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.

### 1.2.2. 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 to the installation program.

**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.

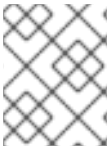

**NOTE**

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

**Procedure**

1. If you do not have an 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 rsa -b 4096 -N '' \
    -f <path>/<file_name> 1
```

1 Specify the path and file name, such as ~/**.ssh/id_rsa**, of the SSH key.

   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.

### 1.2.3. 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.

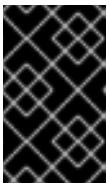

**IMPORTANT**

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.

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.

### 1.2.4. Deploy 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**

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

**Procedure**

1. Run the installation program:

```
$ ./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**.

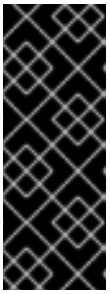

**IMPORTANT**

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.

Provide values at the prompts:

a. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

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

b. Select **azure** as the platform to target.

c. If you do not have a Microsoft Azure profile stored on your computer, specify the following Azure parameter values for your subscription and service principal:

- **azure subscription id** The subscription ID to use for the cluster. Specify the **id** value in your account output.
- **azure tenant id** The tenant ID. Specify the **tenantId** value in your account output.
- **azure service principal client id** The value of the **appId** parameter for the service principal.
- **azure service principal client secret** The value of the **password** parameter for the service principal.

d. Select the region to deploy the cluster to.

e. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.

f. Enter a descriptive name for your cluster.

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see Resolve reserved resource name errors in the Azure documentation.

g. Paste the pull secret that you obtained from the Pull Secret page on the Red Hat OpenShift Cluster Manager site.

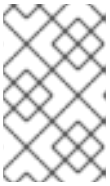

NOTE

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.

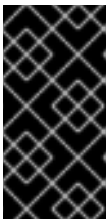

IMPORTANT

The Ignition config files that the installation program generates contain certificates that expire after 24 hours. You must keep the cluster running for 24 hours in a non-degraded state to ensure that the first certificate rotation has finished.

IMPORTANT

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

### 1.2.5. Installing the CLI

You can install the CLI in order to interact with OpenShift Container Platform using a command-line interface.

IMPORTANT

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.2. Download and install the new version of **oc**.

**Procedure**

1. From the Infrastructure Provider page on the Red Hat OpenShift Cluster Manager site, navigate to the page for your installation type and click **Download Command-line Tools**.
2. Click the folder for your operating system and architecture and click the compressed file.

NOTE

You can install **oc** on Linux, Windows, or macOS.

3. Save the file to your file system.
4. Extract the compressed file.
5. Place it in a directory that is on your **PATH**.

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

```
$ oc <command>
```

### 1.2.6. 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
```

#### Next steps

- Customize your cluster.
- If necessary, you can opt out of remote health reporting .

## 1.3. INSTALLING A CLUSTER ON AZURE WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.2, you can install a customized cluster on infrastructure that the installation program provisions on Microsoft Azure. To customize the installation, you modify parameters in the **install-config.yaml** file before you install the cluster.

#### Prerequisites

- Review details about the OpenShift Container Platform installation and update processes.
- Configure an Azure account to host the cluster and determine the tested and validated region to deploy the cluster to.
- If you use a firewall, you must configure it to allow the sites that your cluster requires access to.

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

In OpenShift Container Platform 4.2, you require access to the internet to install and entitle 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 . From there, you can allocate entitlements to your cluster.

You must have internet access to:

- Access the Red Hat OpenShift Cluster Manager page to download the installation program and perform subscription management and entitlement. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster. If the Telemetry service cannot entitle your cluster, you must manually entitle it on the Cluster registration page.
- 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.

### 1.3.2. 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 to the installation program.

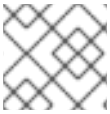

**NOTE**

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.

**NOTE**

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

**Procedure**

1. If you do not have an 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 rsa -b 4096 -N '' \
    -f <path>/<file_name> (1)
```

(1) Specify the path and file name, such as **~/.ssh/id_rsa**, of the SSH key.

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.

## 1.3.3. 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.

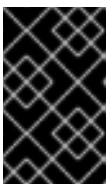

**IMPORTANT**

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.

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.

### 1.3.4. Creating the installation configuration file

You can customize your installation of OpenShift Container Platform on Microsoft Azure.

#### Prerequisites

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

#### Procedure

1. Create the **install-config.yaml** file.
   a. Run the following command:

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

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

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

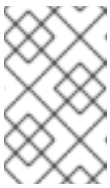

**NOTE**

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

      ii. Select **azure** as the platform to target.
      iii. If you do not have a Microsoft Azure profile stored on your computer, specify the following Azure parameter values for your subscription and service principal:
         - **azure subscription id** The subscription ID to use for the cluster. Specify the **id** value in your account output.

- **azure tenant id** The tenant ID. Specify the **tenantId** value in your account output.
- **azure service principal client id** The value of the **appId** parameter for the service principal.
- **azure service principal client secret** The value of the **password** parameter for the service principal.

iv. Select the region to deploy the cluster to.

v. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.

vi. Enter a descriptive name for your cluster.

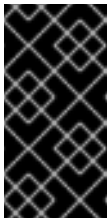

IMPORTANT

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see Resolve reserved resource name errors in the Azure documentation.

vii. Paste the pull secret that you obtained 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.

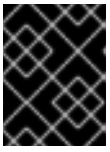

IMPORTANT

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

#### 1.3.4.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.

NOTE

You cannot modify these parameters in the **install-config.yaml** file after installation.

**Table 1.1. Required parameters**

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

| Parameter | Description | Values |
|---|---|---|
| **baseDomain** | The base domain of your cloud provider. This value is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the **baseDomain** and **metadata.name** parameter values that uses the **<metadata.name>.<baseDomain>** format. | A fully-qualified domain or subdomain name, such as **example.com**. |
| **controlPlane.platform** | The cloud provider to host the control plane machines. This parameter value must match the **compute.platform** parameter value. | **aws**, **azure**, **gcp**, **openstack**, or **{}** |
| **compute.platform** | The cloud provider to host the worker machines. This parameter value must match the **controlPlane.platform** parameter value. | **aws**, **azure**, **gcp**, **openstack**, or **{}** |
| **metadata.name** | The name of your cluster. | A string that contains uppercase or lowercase letters, such as **dev**. |
| **platform.<platform>.region** | The region to deploy your cluster in. | A valid region for your cloud, such as **us-east-1** for AWS, **centralus** for Azure, or **region1** for Red Hat OpenStack Platform (RHOSP). |
| **pullSecret** | The pull secret that you obtained from the Pull Secret page on the Red Hat OpenShift Cluster Manager site. You use this pull secret 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. | {<br>"auths":{<br>"cloud.openshift.com":{<br>"auth":"b3Blb=",<br>"email":"you@example.com"<br>},<br>"quay.io":{<br>"auth":"b3Blb=",<br>"email":"you@example.com"<br>}<br>}<br>} |

Table 1.2. Optional parameters

| Parameter | Description | Values |
| --- | --- | --- |
| **sshKey** | The SSH key to use to access your cluster machines.<br><br> <br>**NOTE**<br><br>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery on, specify an SSH key that your **ssh-agent** process uses. | A valid, local public SSH key that you added to the **ssh-agent** process. |
| **compute.hyperthreading** | Whether to enable or disable simultaneous multithreading, or **hyperthreading**, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.<br><br> 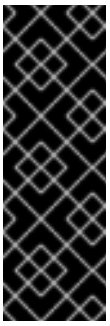 <br>**IMPORTANT**<br><br>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. | **Enabled** or **Disabled** |
| **compute.replicas** | The number of compute machines, which are also known as worker machines, to provision. | A positive integer greater than or equal to **2**. The default value is **3**. |

| Parameter | Description | Values |
|---|---|---|
| **controlPlane.hyperthreading** | Whether to enable or disable simultaneous multithreading, or **hyperthreading**, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.<br><br> 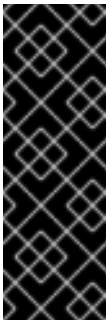 <br>**IMPORTANT**<br>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. | **Enabled** or **Disabled** |
| **controlPlane.replicas** | The number of control plane machines to provision. | A positive integer greater than or equal to **3**. The default value is **3**. |

Table 1.3. Additional Azure parameters

| Parameter | Description | Values |
|---|---|---|
| **machines.platform.azure.type** | The Azure VM instance type. | VMs that use Windows or Linux as the operating system. See the Guest operating systems supported on Azure Stack in the Azure documentation. |
| **machines.platform.azure.osDisk.diskSizeGB** | The Azure disk size for the VM. | Integer that represents the size of the disk in GB, for example **512**. The minimum supported disk size is **120**. |
| **platform.azure.baseDomainResourceGroupName** | The name of the resource group that contains the DNS zone for your base domain. | String, for example **production_cluster**. |
| **platform.azure.region** | The name of the Azure region that hosts your cluster. | Any valid region name. |

| Parameter | Description | Values |
| --- | --- | --- |
| **platform.azure.zone** | List of availability zones to place machines in. For high availability, specify at least two zones. | List of zones, for example **["1", "2", "3"]** |

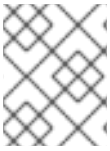

NOTE

You cannot customize Azure Availability Zones or Use tags to organize your Azure resources with an Azure cluster.

#### 1.3.4.2. Sample customized install-config.yaml file for Azure

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

IMPORTANT

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

```
apiVersion: v1
baseDomain: example.com (1)
controlPlane: (2)
  hyperthreading: Enabled (3) (4)
  name: master
  platform:
    azure:
      osDisk:
        diskSizeGB: 512 (5)
      type: Standard_D8s_v3
  replicas: 3
compute: (6)
- hyperthreading: Enabled (7)
  name: worker
  platform:
    azure:
      type: Standard_D2s_v3
      osDisk:
        diskSizeGB: 512 (8)
      zones: (9)
      - "1"
      - "2"
      - "3"
  replicas: 5
metadata:
  name: test-cluster (10)
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
```

```
  machineCIDR: 10.0.0.0/16
  networkType: OpenShiftSDN
  serviceNetwork:
  - 172.30.0.0/16
platform:
  azure:
    region: centralus 11
    baseDomainResourceGroupName: resource-group 12
pullSecret: '{"auths": ...}' 13
sshKey: ssh-ed25519 AAAA... 14
```

1 10 11 13 Required. The installation program prompts you for this value.

2 6 If you do not provide these parameters and values, the installation program provides the default value.

3 7 The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, **-**, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.

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

IMPORTANT

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

5 8 You can specify the size of the disk to use in GB.

9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.

12 Specify the name of the resource group that contains the DNS zone for your base domain.

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

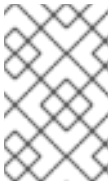

NOTE

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

### 1.3.5. Deploy the cluster

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

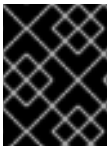

IMPORTANT

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

**Prerequisites**

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

**Procedure**

1. Run the installation program:

```
$ ./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**.

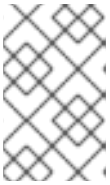

NOTE

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. You must keep the cluster running for 24 hours in a non-degraded state to ensure that the first certificate rotation has finished.

IMPORTANT

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

### 1.3.6. Installing the CLI

You can install the CLI in order to interact with OpenShift Container Platform using a command-line interface.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.2. Download and install the new version of **oc**.

**Procedure**

1. From the Infrastructure Provider page on the Red Hat OpenShift Cluster Manager site, navigate to the page for your installation type and click **Download Command-line Tools**.
2. Click the folder for your operating system and architecture and click the compressed file.

**NOTE**

You can install **oc** on Linux, Windows, or macOS.

3. Save the file to your file system.
4. Extract the compressed file.
5. Place it in a directory that is on your **PATH**.

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

```
$ oc <command>
```

### 1.3.7. 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
```

**Next steps**

- Customize your cluster.
- If necessary, you can opt out of remote health reporting .

## 1.4. INSTALLING A CLUSTER ON AZURE WITH NETWORK CUSTOMIZATIONS

In OpenShift Container Platform version 4.2, you can install a cluster with a customized network configuration on infrastructure that the installation program provisions on Microsoft Azure. By customizing your network configuration, your cluster can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations.

You must set most of the network configuration parameters during installation, and you can modify only **kubeProxy** configuration parameters in a running cluster.

**Prerequisites**

- Review details about the OpenShift Container Platform installation and update processes.
- Configure an Azure account to host the cluster and determine the tested and validated region to deploy the cluster to.
- If you use a firewall, you must configure it to allow the sites that your cluster requires access to.

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

In OpenShift Container Platform 4.2, you require access to the internet to install and entitle 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 . From there, you can allocate entitlements to your cluster.

You must have internet access to:

- Access the Red Hat OpenShift Cluster Manager page to download the installation program and perform subscription management and entitlement. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster. If the Telemetry service cannot entitle your cluster, you must manually entitle it on the Cluster registration page.
- 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.

### 1.4.2. 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 to the installation program.

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.

NOTE

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

**Procedure**

1. If you do not have an 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 rsa -b 4096 -N '' \
    -f <path>/<file_name> 1
```

1 Specify the path and file name, such as ~/**.ssh/id_rsa**, of the SSH key.

   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.

### 1.4.3. 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.

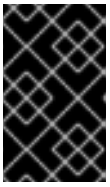

**IMPORTANT**

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.

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.

## 1.4.4. Creating the installation configuration file

You can customize your installation of OpenShift Container Platform on Microsoft Azure.

### Prerequisites

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

### Procedure

1. Create the **install-config.yaml** file.
   a. Run the following command:

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

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

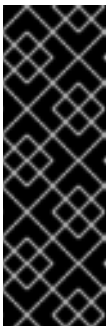

**IMPORTANT**

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. At the prompts, provide the configuration details for your cloud:

i. Optional: Select an SSH key to use to access your cluster machines.

**NOTE**

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

ii. Select **azure** as the platform to target.

iii. If you do not have a Microsoft Azure profile stored on your computer, specify the following Azure parameter values for your subscription and service principal:

- **azure subscription id** The subscription ID to use for the cluster. Specify the **id** value in your account output.
- **azure tenant id** The tenant ID. Specify the **tenantId** value in your account output.
- **azure service principal client id** The value of the **appId** parameter for the service principal.
- **azure service principal client secret** The value of the **password** parameter for the service principal.

iv. Select the region to deploy the cluster to.

v. Select the base domain to deploy the cluster to. The base domain corresponds to the Azure DNS Zone that you created for your cluster.

vi. Enter a descriptive name for your cluster.

**IMPORTANT**

All Azure resources that are available through public endpoints are subject to resource name restrictions, and you cannot create resources that use certain terms. For a list of terms that Azure restricts, see Resolve reserved resource name errors in the Azure documentation.

vii. Paste the pull secret that you obtained 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.

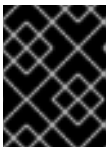

IMPORTANT

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

#### 1.4.4.1. Installation configuration parameters

Before you deploy an OpenShift Container Platform cluster, you provide parameter values to describe your account on the cloud platform that hosts your cluster and optionally customize your cluster's platform. When you create the **install-config.yaml** installation configuration file, you provide values for the required parameters through the command line. If you customize your cluster, you can modify the **install-config.yaml** file to provide more details about the platform.

NOTE

You cannot modify these parameters in the **install-config.yaml** file after installation.

Table 1.4. Required parameters

| Parameter | Description | Values |
|---|---|---|
| **baseDomain** | The base domain of your cloud provider. This value is used to create routes to your OpenShift Container Platform cluster components. The full DNS name for your cluster is a combination of the **baseDomain** and **metadata.name** parameter values that uses the **<metadata.name>.<baseDomain>** format. | A fully-qualified domain or subdomain name, such as **example.com**. |
| **controlPlane.platform** | The cloud provider to host the control plane machines. This parameter value must match the **compute.platform** parameter value. | **aws**, **azure**, **gcp**, **openstack**, or **{}** |
| **compute.platform** | The cloud provider to host the worker machines. This parameter value must match the **controlPlane.platform** parameter value. | **aws**, **azure**, **gcp**, **openstack**, or **{}** |
| **metadata.name** | The name of your cluster. | A string that contains uppercase or lowercase letters, such as **dev**. |

| Parameter | Description | Values |
|---|---|---|
| **platform.<platform>.region** | The region to deploy your cluster in. | A valid region for your cloud, such as **us-east-1** for AWS, **centralus** for Azure, or **region1** for Red Hat OpenStack Platform (RHOSP). |
| **pullSecret** | The pull secret that you obtained from the Pull Secret page on the Red Hat OpenShift Cluster Manager site. You use this pull secret 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. | <pre>{<br>  "auths":{<br>  "cloud.openshift.com":{<br>  "auth":"b3Blb=",<br>  "email":"you@example.com"<br>  },<br>  "quay.io":{<br>  "auth":"b3Blb=",<br>  "email":"you@example.com"<br>  }<br>  }<br>}</pre> |

Table 1.5. Optional parameters

| Parameter | Description | Values |
|---|---|---|
| **sshKey** | The SSH key to use to access your cluster machines.<br><br>**NOTE**<br><br>For production OpenShift Container Platform clusters on which you want to perform installation debugging or disaster recovery on, specify an SSH key that your **ssh-agent** process uses. | A valid, local public SSH key that you added to the **ssh-agent** process. |

| Parameter | Description | Values |
| --- | --- | --- |
| **compute.hyperthreading** | Whether to enable or disable simultaneous multithreading, or **hyperthreading**, on compute machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.<br><br> <br>**IMPORTANT**<br><br>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. | **Enabled** or **Disabled** |
| **compute.replicas** | The number of compute machines, which are also known as worker machines, to provision. | A positive integer greater than or equal to **2**. The default value is **3**. |
| **controlPlane.hyperthreading** | Whether to enable or disable simultaneous multithreading, or **hyperthreading**, on control plane machines. By default, simultaneous multithreading is enabled to increase the performance of your machines' cores.<br><br> <br>**IMPORTANT**<br><br>If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. | **Enabled** or **Disabled** |
| **controlPlane.replicas** | The number of control plane machines to provision. | A positive integer greater than or equal to **3**. The default value is **3**. |

Table 1.6. Additional Azure parameters

| Parameter | Description | Values |
|---|---|---|
| **machines.platform.azure.type** | The Azure VM instance type. | VMs that use Windows or Linux as the operating system. See the Guest operating systems supported on Azure Stack in the Azure documentation. |
| **machines.platform.azure.osDisk.diskSizeGB** | The Azure disk size for the VM. | Integer that represents the size of the disk in GB, for example **512**. The minimum supported disk size is **120**. |
| **platform.azure.baseDomainResourceGroupName** | The name of the resource group that contains the DNS zone for your base domain. | String, for example **production_cluster**. |
| **platform.azure.region** | The name of the Azure region that hosts your cluster. | Any valid region name. |
| **platform.azure.zone** | List of availability zones to place machines in. For high availability, specify at least two zones. | List of zones, for example **["1", "2", "3"]** |

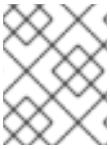

NOTE

You cannot customize Azure Availability Zones or Use tags to organize your Azure resources with an Azure cluster.

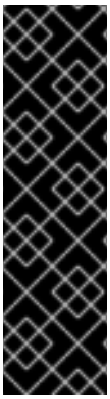

IMPORTANT

The Open Virtual Networking (OVN) Kubernetes network plug-in is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of the OVN Technology Preview, see https://access.redhat.com/articles/4380121.

### 1.4.4.2. Network configuration parameters

You can modify your cluster network configuration parameters in the **install-config.yaml** configuration file. The following table describes the parameters.

NOTE

You cannot modify these parameters in the **install-config.yaml** file after installation.

**Table 1.7. Required network parameters**

| Parameter | Description | Value |
|---|---|---|
| **networking.networkType** | The network plug-in to deploy. The **OpenShiftSDN** plug-in is the only plug-in supported in OpenShift Container Platform 4.2. The **OVNKubernetes** plug-in is available as Technology Preview in OpenShift Container Platform 4.2. | Either **OpenShiftSDN** or **OVNKubernetes**. The default value is **OpenShiftSDN**. |
| **networking.clusterNetwork.cidr** | A block of IP addresses from which Pod IP addresses are allocated. The **OpenShiftSDN** network plug-in supports multiple cluster networks. The address blocks for multiple cluster networks must not overlap. Select address pools large enough to fit your anticipated workload. | An IP address allocation in CIDR format. The default value is **10.128.0.0/14**. |
| **networking.clusterNetwork.hostPrefix** | The subnet prefix length to assign to each individual node. For example, if **hostPrefix** is set to **23**, then each node is assigned a **/23** subnet out of the given **cidr**, allowing for 510 (2^(32 - 23) - 2) Pod IP addresses. | A subnet prefix. The default value is **23**. |
| **networking.serviceNetwork** | A block of IP addresses for services. **OpenShiftSDN** allows only one **serviceNetwork** block. The address block must not overlap with any other network block. | An IP address allocation in CIDR format. The default value is **172.30.0.0/16**. |
| **networking.machineCIDR** | A block of IP addresses used by the OpenShift Container Platform installation program while installing the cluster. The address block must not overlap with any other network block. | An IP address allocation in CIDR format. The default value is **10.0.0.0/16**. |

### 1.4.4.3. Sample customized install-config.yaml file for Azure

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

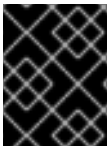

IMPORTANT

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

```
apiVersion: v1
baseDomain: example.com (1)
controlPlane: (2)
  hyperthreading: Enabled (3) (4)
  name: master
  platform:
    azure:
      osDisk:
        diskSizeGB: 512 (5)
      type: Standard_D8s_v3
  replicas: 3
```

```
compute: 6
- hyperthreading: Enabled 7
  name: worker
  platform:
    azure:
      type: Standard_D2s_v3
      osDisk:
        diskSizeGB: 512 8
      zones: 9
      - "1"
      - "2"
      - "3"
  replicas: 5
metadata:
  name: test-cluster 10
networking: 11
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineCIDR: 10.0.0.0/16
  networkType: OpenShiftSDN
  serviceNetwork:
  - 172.30.0.0/16
platform:
  azure:
    region: centralus 12
    baseDomainResourceGroupName: resource-group 13
pullSecret: '{"auths": ...}' 14
sshKey: ssh-ed25519 AAAA... 15
```

1 10 12 14 Required. The installation program prompts you for this value.

2 6 11 If you do not provide these parameters and values, the installation program provides the default value.

3 7 The **controlPlane** section is a single mapping, but the compute section is a sequence of mappings. To meet the requirements of the different data structures, the first line of the **compute** section must begin with a hyphen, **-**, and the first line of the **controlPlane** section must not. Although both sections currently define a single machine pool, it is possible that future versions of OpenShift Container Platform will support defining multiple compute pools during installation. Only one control plane pool is used.

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

**IMPORTANT**

If you disable simultaneous multithreading, ensure that your capacity planning accounts for the dramatically decreased machine performance. Use larger virtual machine types, such as **Standard_D8s_v3**, for your machines if you disable simultaneous multithreading.

5 8 You can specify the size of the disk to use in GB.

9 Specify a list of zones to deploy your machines to. For high availability, specify at least two zones.

13 Specify the name of the resource group that contains the DNS zone for your base domain.

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

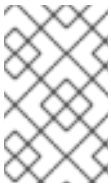

NOTE

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

### 1.4.5. Modifying advanced network configuration parameters

You can modify the advanced network configuration parameters only before you install the cluster. Advanced configuration customization lets you integrate your cluster into your existing network environment by specifying an MTU or VXLAN port, by allowing customization of kube-proxy settings, and by specifying a different **mode** for the **openshiftSDNConfig** parameter.

IMPORTANT

Modifying the OpenShift Container Platform manifest files directly is not supported.

**Prerequisites**

- Create the **install-config.yaml** file and complete any modifications to it.

**Procedure**

1. Use the following command to create manifests:

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

1 For **<installation_directory>**, specify the name of the directory that contains the **install-config.yaml** file for your cluster.

2. Modify the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file to prevent Pods from being scheduled on the control plane machines:

   a. Open the **manifests/cluster-scheduler-02-config.yml** file.

   b. Locate the **mastersSchedulable** parameter and set its value to **False**.

   c. Save and exit the file.

NOTE

Currently, due to a Kubernetes limitation, router Pods running on control plane machines will not be reachable by the ingress load balancer.

3. Create a file that is named **cluster-network-03-config.yml** in the **<installation_directory>/manifests/** directory:

```
$ touch <installation_directory>/manifests/cluster-network-03-config.yml 1
```

1 For **<installation_directory>**, specify the directory name that contains the **manifests/** directory for your cluster.

After creating the file, several network configuration files are in the **manifests/** directory, as shown:

```
$ ls <installation_directory>/manifests/cluster-network-*
cluster-network-01-crd.yml
cluster-network-02-config.yml
cluster-network-03-config.yml
```

4. Open the **cluster-network-03-config.yml** file in an editor and enter a CR that describes the Operator configuration you want:

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec: 1
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  serviceNetwork:
  - 172.30.0.0/16
  defaultNetwork:
    type: OpenShiftSDN
    openshiftSDNConfig:
      mode: NetworkPolicy
      mtu: 1450
      vxlanPort: 4789
```

1 The parameters for the **spec** parameter are only an example. Specify your configuration for the Cluster Network Operator in the CR.

The CNO provides default values for the parameters in the CR, so you must specify only the parameters that you want to change.

5. Save the **cluster-network-03-config.yml** file and quit the text editor.

6. Optional: Back up the **manifests/cluster-network-03-config.yml** file. The installation program deletes the **manifests/** directory when creating the cluster.

### 1.4.6. Cluster Network Operator custom resource (CR)

The cluster network configuration in the **Network.operator.openshift.io** custom resource (CR) stores the configuration settings for the Cluster Network Operator (CNO). The Operator manages the cluster network.

You can specify the cluster network configuration for your OpenShift Container Platform cluster by setting the parameters for the **defaultNetwork** parameter in the CNO CR. The following CR displays the default configuration for the CNO and explains both the parameters you can configure and valid parameter values:

#### Cluster Network Operator CR

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  clusterNetwork: 1
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  serviceNetwork: 2
  - 172.30.0.0/16
  defaultNetwork: 3
    ...
  kubeProxyConfig: 4
    iptablesSyncPeriod: 30s 5
    proxyArguments:
      iptables-min-sync-period: 6
      - 30s
```

1 2 Specified in the **install-config.yaml** file.

3 Configures the software-defined networking (SDN) for the cluster network.

4 The parameters for this object specify the **kube-proxy** configuration. If you do not specify the parameter values, the Network Operator applies the displayed default parameter values.

5 The refresh period for **iptables** rules. The default value is **30s**. Valid suffixes include **s**, **m**, and **h** and are described in the Go time package documentation.

6 The minimum duration before refreshing **iptables** rules. This parameter ensures that the refresh does not happen too frequently. Valid suffixes include **s**, **m**, and **h** and are described in the Go time package

#### 1.4.6.1. Configuration parameters for OpenShift SDN

The following YAML object describes the configuration parameters for OpenShift SDN:

```
defaultNetwork:
  type: OpenShiftSDN 1
  openshiftSDNConfig: 2
    mode: NetworkPolicy 3
    mtu: 1450 4
    vxlanPort: 4789 5
```

1 Specified in the **install-config.yaml** file.

2 Specify only if you want to override part of the OpenShift SDN configuration.

3 Configures the network isolation mode for **OpenShiftSDN**. The allowed values are **Multitenant**, **Subnet**, or **NetworkPolicy**. The default value is **NetworkPolicy**.

4 MTU for the VXLAN overlay network. This value is normally configured automatically, but if the nodes in your cluster do not all use the same MTU, then you must set this explicitly to 50 less than the smallest node MTU value.

5 The port to use for all VXLAN packets. The default value is **4789**. If you are running in a virtualized environment with existing nodes that are part of another VXLAN network, then you might be required to change this. For example, when running an OpenShift SDN overlay on top of VMware NSX-T, you must select an alternate port for VXLAN, since both SDNs use the same default VXLAN port number.

On Amazon Web Services (AWS), you can select an alternate port for the VXLAN between port **9000** and port **9999**.

#### 1.4.6.2. Configuration parameters for Open Virtual Network (OVN) SDN

The OVN SDN does not have any configuration parameters in OpenShift Container Platform 4.2.

#### 1.4.6.3. Cluster Network Operator example CR

A complete CR for the CNO is displayed in the following example:

**Cluster Network Operator example CR**

```
apiVersion: operator.openshift.io/v1
kind: Network
metadata:
  name: cluster
spec:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  serviceNetwork:
  - 172.30.0.0/16
  defaultNetwork:
    type: OpenShiftSDN
    openshiftSDNConfig:
      mode: NetworkPolicy
      mtu: 1450
      vxlanPort: 4789
  kubeProxyConfig:
    iptablesSyncPeriod: 30s
    proxyArguments:
      iptables-min-sync-period:
      - 30s
```

### 1.4.7. Deploy 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**

- Configure an account with the cloud platform that hosts your cluster.
- Obtain the OpenShift Container Platform installation program and the pull secret for your cluster.

**Procedure**

1. Run the installation program:

```
$ ./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**.

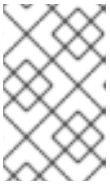

NOTE

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. You must keep the cluster running for 24 hours in a non-degraded state to ensure that the first certificate rotation has finished.

IMPORTANT

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

### 1.4.8. Installing the CLI

You can install the CLI in order to interact with OpenShift Container Platform using a command-line interface.

**IMPORTANT**

If you installed an earlier version of **oc**, you cannot use it to complete all of the commands in OpenShift Container Platform 4.2. Download and install the new version of **oc**.

**Procedure**

1. From the Infrastructure Provider page on the Red Hat OpenShift Cluster Manager site, navigate to the page for your installation type and click **Download Command-line Tools**.
2. Click the folder for your operating system and architecture and click the compressed file.

**NOTE**

You can install **oc** on Linux, Windows, or macOS.

3. Save the file to your file system.
4. Extract the compressed file.
5. Place it in a directory that is on your **PATH**.

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

```
$ oc <command>
```

### 1.4.9. 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
```

### Next steps

- Customize your cluster.
- If necessary, you can opt out of remote health reporting .

## 1.5. UNINSTALLING A CLUSTER ON AZURE

You can remove a cluster that you deployed to Microsoft Azure.

### 1.5.1. 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.