OpenShift Container Platform 4.2

Installing on OpenStack

Installing OpenShift Container Platform 4.2 OpenStack clusters

OpenShift Container Platform 4.2 Installing on OpenStack

Installing OpenShift Container Platform 4.2 OpenStack clusters

Legal Notice

Abstract

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

Table of Contents

CHAPTER 1. INSTALLING ON OPENSTACK

1.1. INSTALLING A CLUSTER ON OPENSTACK WITH CUSTOMIZATIONS

In OpenShift Container Platform version 4.2, you can install a customized cluster on Red Hat OpenStack Platform (RHOSP). To customize the installation, modify parameters in the **install-config.yaml** before you install the cluster.

Prerequisites

- Review details about the OpenShift Container Platform installation and update processes.

  - Verify that OpenShift Container Platform 4.2 is compatible with your RHOSP version in the *Available platforms* section. You can also compare platform support across different versions by viewing the OpenShift Container Platform on RHOSP support matrix .

- Have metadata service enabled in RHOSP

1.1.1. Resource guidelines for installing OpenShift Container Platform on OpenStack

Your quota must meet the following requirements to run the OpenShift Container Platform installation program in Red Hat OpenStack Platform (RHOSP).

Table 1.1. Recommended resources for a default OpenShift Container Platform cluster on RHOSP

Resource	Value
Floating IP addresses	2
Ports	15
Routers	1
Subnets	1
RAM	112 GB
vCPUs	28
Volume storage	175 GB
Instances	7
Security groups	3
Security group rules	60
Swift containers	2
Swift objects	1

Resource	Value
Swift available space	10 MB or more

NOTE

Swift space requirements vary depending on the size of the bootstrap Ignition file and image registry.

A cluster might function with fewer than recommended resources, but its performance is not guaranteed.

NOTE

By default, your security group and security group rule quotas might be low. If you encounter problems, run **openstack quota set --secgroups 3 --secgroup-rules 60 <project>** as an administrator to increase them.

An OpenShift Container Platform deployment comprises control plane machines, compute machines, and a bootstrap machine.

1.1.1.1. Control plane and compute machines

By default, the OpenShift Container Platform installation program stands up three control plane and compute machines.

Each machine requires:

- An instance from the RHOSP quota

- A port from the RHOSP quota

- A flavor with at least 16 GB memory, 4 vCPUs, and 25 GB storage space

TIP

Compute machines host the applications that you run on OpenShift Container Platform; aim to run as many as you can.

1.1.1.2. Bootstrap machine

During installation, a bootstrap machine is temporarily provisioned to stand up the control plane. After the production control plane is ready, the bootstrap machine is deprovisioned.

The bootstrap machine requires:

- An instance from the RHOSP quota

- A port from the RHOSP quota

- A flavor with at least 16 GB memory, 4 vCPUs, and 25 GB storage space

> **NOTE**
>
> The installation program cannot pass certificate authority bundles to Ignition on control plane machines. Therefore, the bootstrap machine cannot retrieve Ignition configurations from Swift if your endpoint uses self-signed certificates.

1.1.2. 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.

1.1.3. Enabling Swift on OpenStack

OpenShift Container Platform on Red Hat OpenStack Platform (RHOSP) uses OpenStack Object Storage (Swift) to store and serve user configuration files.

Swift is operated by a user account with the **swiftoperator** role and **temp-url** support.

Prerequisites

- A RHOSP administrator account on the target environment

- On Ceph RGW, the **account in url** option must be enabled

Procedure

To enable Swift on RHOSP:

1. As an administrator in the RHOSP CLI, add the **swiftoperator** role to the account that will access Swift:

   ```
   $ openstack role add --user <user> --project <project> swiftoperator
   ```

2. As the account with the **swiftoperator** role, set a temporary URL property for the account:

```
$ openstack object store account set --property Temp-URL-Key=superkey
```

Your RHOSP deployment can now use Swift to store and serve files.

1.1.4. Creating the Red Hat Enterprise Linux CoreOS (RHCOS) image

The OpenShift Container Platform installation program requires that a Red Hat Enterprise Linux CoreOS (RHCOS) image be present in the Red Hat OpenStack Platform (RHOSP) cluster. Retrieve the latest RHCOS image, then upload it using the RHOSP CLI.

Prerequisites

- The RHOSP CLI is installed.

Procedure

1. Log in to the Red Hat customer portal's Product Downloads page .

2. Under **Version**, select version 4.2.0 for RHEL 8.

3. Download the *Red Hat Enterprise Linux CoreOS – OpenStack Image (QCOW)* .

   > **IMPORTANT**
   >
   > The RHCOS images might not change with every release of OpenShift Container Platform. You must download images with the highest version that is less than or equal to the OpenShift Container Platform version that you install. Use the image versions that match your OpenShift Container Platform version if they are available.

4. Decompress the image.

   > **NOTE**
   >
   > You must decompress the OpenStack image before the cluster can use it.

5. From the image that you downloaded, create an image that is named **rhcos** in your cluster by using the RHOSP CLI:

```
$ openstack image create --container-format=bare --disk-format=qcow2 --file rhcos-${RHCOS_VERSION}-openstack.qcow2 rhcos
```

   CAUTION

   If the installation program finds multiple images with the same name, it chooses one of them at random. To avoid this behavior, create unique names for resources in RHOSP.

After you upload the image to RHOSP, it is available to the installation program.

1.1.5. Verifying external network access

The OpenShift Container Platform installer requires external network access. You must provide an external network value to it, or deployment fails. Before you run the installer, verify that a network with the External router type exists in Red Hat OpenStack Platform (RHOSP).

Prerequisites

- On RHOSP, the **NeutronDhcpAgentDnsmasqDnsServers** parameter must be configured to allow DHCP agents to forward instances' DNS queries. One way to set this parameter is to:

  a. Create a new environment file in the template directory.

  b. Provide parameter values in the file. For example:

     Sample neutron-dhcp-agent-dnsmasq-dns-servers.yaml file

     ```
     parameter_defaults:
       NeutronDhcpAgentDnsmasqDnsServers:
     ['<DNS_server_address_1>','<DNS_server_address_2']
     ```

  c. Include the environment file in your Overcloud deploy command. For example:

     ```
     $ openstack overcloud deploy --templates -e neutron-dhcp-agent-dnsmasq-dns-servers.yaml ...
     ```

Procedure

1. Using the RHOSP CLI, verify the name and ID of the 'External' network:

   ```
   $ openstack network list --long -c ID -c Name -c "Router Type"

   +--------------------------------------+----------------+-------------+
   | ID                                   | Name           | Router Type |
   +--------------------------------------+----------------+-------------+
   | 148a8023-62a7-4672-b018-003462f8d7dc | public_network | External    |
   +--------------------------------------+----------------+-------------+
   ```

A network with an External router type appears in the network list. If at least one does not, see Create an external network.
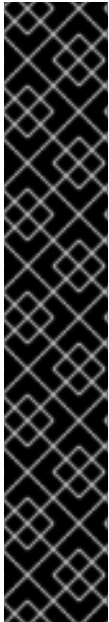
IMPORTANT

If the external network's CIDR range overlaps one of the default network ranges, you must change the matching network ranges in the **install-config.yaml** file before you run the installation program.

The default network ranges are:

Network	Range
machineCIDR	10.0.0.0/16
serviceNetwork	172.30.0.0/16
clusterNetwork	10.128.0.0/14

CAUTION

If the installation program finds multiple networks with the same name, it sets one of them at random. To avoid this behavior, create unique names for resources in RHOSP.
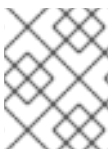
NOTE

If the Neutron trunk service plug-in is enabled, a trunk port is created by default. For more information, see Neutron trunk port .

1.1.6. Defining parameters for the installation program

The OpenShift Container Platform installation program relies on a file called **clouds.yaml**. The file describes Red Hat OpenStack Platform (RHOSP) configuration parameters, including the project name, log in information, and authorization service URLs.

Procedure

1. Create the **clouds.yaml** file:

   - If your OpenStack distribution includes the Horizon web UI, generate a **clouds.yaml** file in it.

     IMPORTANT

     Remember to add a password to the **auth** field. You can also keep secrets in a separate file from **clouds.yaml**.

   - If your OpenStack distribution does not include the Horizon web UI, or you do not want to use Horizon, create the file yourself. For detailed information about **clouds.yaml**, see Config files in the RHOSP documentation.

     ```
     clouds:
       shiftstack:
         auth:
           auth_url: http://10.10.14.42:5000/v3
     ```

```
        project_name: shiftstack
        username: shiftstack_user
        password: XXX
        user_domain_name: Default
        project_domain_name: Default
    dev-env:
      region_name: RegionOne
      auth:
        username: 'devuser'
        password: XXX
        project_name: 'devonly'
        auth_url: 'https://10.10.14.22:5001/v2.0'
```

2. Place the file that you generate in one of the following locations:

   a. The value of the **OS_CLIENT_CONFIG_FILE** environment variable

   b. The current directory

   c. A Unix-specific user configuration directory, for example ~/**.config**/**openstack**/**clouds.yaml**

   d. A Unix-specific site configuration directory, for example /**etc**/**openstack**/**clouds.yaml**
      The installation program searches for **clouds.yaml** in that order.

1.1.7. 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.

1.1.8. Creating the installation configuration file

You can customize your installation of OpenShift Container Platform on OpenStack.

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 **openstack** as the platform to target.

   iii. Specify the Red Hat OpenStack Platform (RHOSP) external network name to use for installing the cluster.

   iv. Specify the Floating IP address to use for external access to the OpenShift API.

v. Specify a RHOSP flavor with at least 16 GB RAM to use for control plane and compute nodes.

vi. Select the base domain to deploy the cluster to. All DNS records will be sub-domains of this base and will also include the cluster name.

vii. Enter a name for your cluster. The name must be 14 or fewer characters long.

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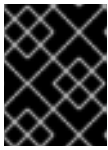
> **IMPORTANT**
>
> The **install-config.yaml** file is consumed during the installation process. If you want to reuse the file, you must back it up now.

1.1.9. 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.
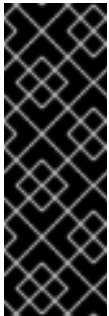
Table 1.2. Required parameters

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**.

Parameter	Description	Values
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**. The string must be 14 characters or fewer long.
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.	``` { "auths":{ "cloud.openshift.com":{ "auth":"b3Blb=", "email":"you@example.com" }, "quay.io":{ "auth":"b3Blb=", "email":"you@example.com" } } } ```

Table 1.3. Optional parameters

Parameter	Description	Values

Parameter	Description	Values
sshKey	The 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.	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. **IMPORTANT** 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. **IMPORTANT** 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 1.4. Additional Red Hat OpenStack Platform (RHOSP) parameters

Parameter	Description	Values
compute.platform.openstack.rootVolume.size	For compute machines, the size in gigabytes of the root volume. If you do not set this value, machines use ephemeral storage.	Integer, for example **30**.
compute.platform.openstack.rootVolume.type	For compute machines, the root volume's type.	String, for example **performance**.
controlPlane.platform.openstack.rootVolume.size	For control plane machines, the size in gigabytes of the root volume. If you do not set this value, machines use ephemeral storage.	Integer, for example **30**.
controlPlane.platform.openstack.rootVolume.type	For control plane machines, the root volume's type.	String, for example **performance**.
platform.openstack.region	The region where the RHOSP cluster is created.	String, for example **region1**.

Parameter	Description	Values
platform.openstack. cloud	The name of the RHOSP cloud to use from the list of clouds in the **clouds.yaml** file.	String, for example **MyCloud**.
platform.openstack. externalDNS	*Optional.* IP addresses for external DNS servers that cluster instances use for DNS resolution.	A list of IP addresses as strings, for example **["8.8.8.8", "192.168.1.12"]**.
platform.openstack. externalNetwork	The RHOSP external network name to be used for installation.	String, for example **external**.
platform.openstack. computeFlavor	The RHOSP flavor to use for control plane and compute machines.	String, for example **m1.xlarge**.
platform.openstack.l bFloatingIP	An existing floating IP address to associate with the load balancer API.	An IP address, for example **128.0.0.1**.

platform.openstack. defaultMachinePlatf orm	*Optional.* The default machine pool platform configuration.	{
		"type": "ml.large",
		"rootVolume": {
		"size": 30,
		"type": "performance"
		}
		}

1.1.9.1. Sample customized **install-config.yaml** file for OpenStack

This sample **install-config.yaml** demonstrates all of the possible Red Hat OpenStack Platform (RHOSP) customization options.

> **IMPORTANT**
>
> This sample file is provided for reference only. You must obtain your **install-config.yaml** file by using the installation program.

```
apiVersion: v1
baseDomain: example.com
clusterID: os-test
controlPlane:
  name: master
  platform: {}
  replicas: 3
compute:
- name: worker
  platform:
```

```
    openstack:
      type: ml.large
    replicas: 3
metadata:
  name: example
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineCIDR: 10.0.0.0/16
  serviceNetwork:
  - 172.30.0.0/16
  networkType: OpenShiftSDN
platform:
  openstack:
    region: region1
    cloud: mycloud
    externalNetwork: external
    computeFlavor: m1.xlarge
    lbFloatingIP: 128.0.0.1
pullSecret: '{"auths": ...}'
sshKey: ssh-ed25519 AAAA...
```

1.1.10. 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.
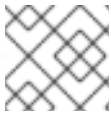
> **NOTE**
>
> In a production environment, you require disaster recovery and debugging.

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.

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> ❶
   ```

   ❶ 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>  ❶

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

❶   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.

1.1.11. Enabling access to the environment

At deployment, all OpenShift Container Platform machines are created in a Red Hat OpenStack Platform (RHOSP)–tenant network. Therefore, they are not accessible directly in most RHOSP deployments.

You can configure the OpenShift Container Platform API to be accessible either with or without floating IP addresses.

1.1.11.1. Enabling access with floating IP addresses

Make OpenShift Container Platform API endpoints accessible by attaching two floating IP (FIP) addresses to them: one for the API load balancer (**lb FIP**), and one for OpenShift Container Platform applications (**apps FIP**).

> **IMPORTANT**
>
> The load balancer FIP is also used in the **install-config.yaml** file.

Procedure

1. Using the Red Hat OpenStack Platform (RHOSP) CLI, create a new external network:

```
$ openstack floating ip create <external network>
```

2. Add a record that follows this pattern to your DNS server:

```
api.<cluster name>.<base domain>  IN  A  <lb FIP>
```

> **NOTE**
>
> If you do not control the DNS server you can add the record to your **/etc/hosts** file instead. This action makes the API accessible to you only, which is not suitable for production deployment but does allow installation for development and testing.

TIP

You can make OpenShift Container Platform resources available outside of the cluster by assigning a floating IP address and updating your firewall configuration.

1.1.11.2. Enabling access without floating IP addresses

If you cannot use floating IP addresses, the OpenShift Container Platform installation might still finish. However, the installation program fails after it times out waiting for API access.

After the installation program times out, the cluster might still initialize. After the bootstrapping processing begins, it must complete. You must edit the cluster's networking configuration after it is deployed, however.

1.1.12. 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

- 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> \ ❶
       --log-level=info ❷
   ```

   ❶ For **<installation_directory>**, specify the location of your customized **./install-config.yaml** file.

   ❷ 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.
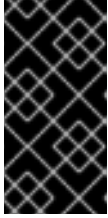
 IMPORTANT

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.

1.1.13. Verifying cluster status

To verify your OpenShift Container Platform cluster's status during or after installation:

Procedure

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

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

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

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

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

   ```
   $ oc get nodes
   ```

3. View your cluster's version:

   ```
   $ oc get clusterversion
   ```

4. View your operators' status:

   ```
   $ oc get clusteroperator
   ```

5. View all running Pods in the cluster:

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

1.1.14. 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 ❶
   ```

   ❶   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
   ```

1.1.15. Configuring application access with floating IP addresses

After you install OpenShift Container Platform, configure Red Hat OpenStack Platform (RHOSP) to allow application network traffic.

Prerequisites

- OpenShift Container Platform cluster must be installed

- Floating IP addresses are enabled as described in *Enabling access to the environment*.

Procedure

After you install the OpenShift Container Platform cluster, attach a floating IP address to the ingress port:

1. Show the port:
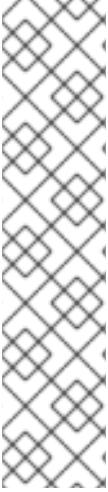
   ```
   $ openstack port show <cluster name>-<clusterID>-ingress-port
   ```

2. Attach the port to the IP address:

   ```
   $ openstack floating ip set --port <ingress port ID> <apps FIP>
   ```

3. Add a wildcard **A** record for ***apps.** to your DNS file:

   ```
   *.apps.<cluster name>.<base domain>  IN  A  <apps FIP>
   ```

> **NOTE**
>
> If you do not control the DNS server but want to enable application access for non-production purposes, you can add these hostnames to /**etc**/**hosts**:
>
> > <apps FIP> console-openshift-console.apps.<cluster name>.<base domain>
> > <apps FIP> integrated-oauth-server-openshift-authentication.apps.<cluster name>.<base domain>
> > <apps FIP> oauth-openshift.apps.<cluster name>.<base domain>
> > <apps FIP> prometheus-k8s-openshift-monitoring.apps.<cluster name>.<base domain>
> > <apps FIP> grafana-openshift-monitoring.apps.<cluster name>.<base domain>
> > <apps FIP> <app name>.apps.<cluster name>.<base domain>

Next steps

- [Customize your cluster](#).

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

1.2. UNINSTALLING A CLUSTER ON OPENSTACK

You can remove a cluster that you deployed to Red Hat OpenStack Platform (RHOSP).

1.2.1. 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.