



OpenShift Container Platform 4.17

Installing on a single node

Installing OpenShift Container Platform on a single node

OpenShift Container Platform 4.17 Installing on a single node

Installing OpenShift Container Platform on a single node

Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

Linux[®] 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 describes how to install OpenShift Container Platform on a single node.

Table of Contents

CHAPTER 1. PREPARING TO INSTALL ON A SINGLE NODE	3
1.1. PREREQUISITES	3
1.2. ABOUT OPENSIFT ON A SINGLE NODE	3
1.3. REQUIREMENTS FOR INSTALLING OPENSIFT ON A SINGLE NODE	3
CHAPTER 2. INSTALLING OPENSIFT ON A SINGLE NODE	6
2.1. INSTALLING SINGLE-NODE OPENSIFT USING THE ASSISTED INSTALLER	6
2.1.1. Generating the discovery ISO with the Assisted Installer	6
2.1.2. Installing single-node OpenShift with the Assisted Installer	7
2.2. INSTALLING SINGLE-NODE OPENSIFT MANUALLY	7
2.2.1. Generating the installation ISO with coreos-installer	8
2.2.2. Monitoring the cluster installation using openshift-install	10
2.3. INSTALLING SINGLE-NODE OPENSIFT ON CLOUD PROVIDERS	11
2.3.1. Additional requirements for installing single-node OpenShift on a cloud provider	11
2.3.2. Supported cloud providers for single-node OpenShift	12
2.3.3. Installing single-node OpenShift on AWS	12
2.3.4. Installing single-node OpenShift on Azure	12
2.3.5. Installing single-node OpenShift on GCP	12
2.4. CREATING A BOOTABLE ISO IMAGE ON A USB DRIVE	12
2.5. BOOTING FROM AN HTTP-HOSTED ISO IMAGE USING THE REDFISH API	13
2.6. CREATING A CUSTOM LIVE RHCOS ISO FOR REMOTE SERVER ACCESS	14
2.7. INSTALLING SINGLE-NODE OPENSIFT WITH IBM Z AND IBM LINUXONE	16
Hardware requirements	16
2.7.1. Installing single-node OpenShift with z/VM on IBM Z and IBM LinuxONE	16
2.7.2. Installing single-node OpenShift with RHEL KVM on IBM Z and IBM LinuxONE	21
2.7.3. Installing single-node OpenShift in an LPAR on IBM Z and IBM LinuxONE	24
2.8. INSTALLING SINGLE-NODE OPENSIFT WITH IBM POWER	29
Hardware requirements	29
2.8.1. Setting up basion for single-node OpenShift with IBM Power	29
2.8.2. Installing single-node OpenShift with IBM Power	32

CHAPTER 1. PREPARING TO INSTALL ON A SINGLE NODE

1.1. PREREQUISITES

- You reviewed details about the [OpenShift Container Platform installation and update](#) processes.
- You have read the documentation on [selecting a cluster installation method and preparing it for users](#).

1.2. ABOUT OPENSIFT ON A SINGLE NODE

You can create a single-node cluster with standard installation methods. OpenShift Container Platform on a single node is a specialized installation that requires the creation of a special Ignition configuration file. The primary use case is for edge computing workloads, including intermittent connectivity, portable clouds, and 5G radio access networks (RAN) close to a base station. The major tradeoff with an installation on a single node is the lack of high availability.



IMPORTANT

The use of OpenShiftSDN with single-node OpenShift is not supported. OVN-Kubernetes is the default network plugin for single-node OpenShift deployments.

1.3. REQUIREMENTS FOR INSTALLING OPENSIFT ON A SINGLE NODE

Installing OpenShift Container Platform on a single node alleviates some of the requirements for high availability and large scale clusters. However, you must address the following requirements:

- **Administration host:** You must have a computer to prepare the ISO, to create the USB boot drive, and to monitor the installation.



NOTE

For the **ppc64le** platform, the host should prepare the ISO, but does not need to create the USB boot drive. The ISO can be mounted to PowerVM directly.



NOTE

ISO is not required for IBM Z[®] installations.

- **CPU Architecture:** Installing OpenShift Container Platform on a single node supports **x86_64**, **arm64**, **ppc64le**, and **s390x** CPU architectures.
- **Supported platforms:** Installing OpenShift Container Platform on a single node is supported on bare metal and [Certified third-party hypervisors](#). In most cases, you must specify the **platform.none: {}** parameter in the **install-config.yaml** configuration file. The following list shows the only exceptions and the corresponding parameter to specify in the **install-config.yaml** configuration file:
 - Amazon Web Services (AWS), where you use **platform=aws**

- Google Cloud Platform (GCP), where you use **platform=gcp**
- Microsoft Azure, where you use **platform=azure**
- **Production-grade server:** Installing OpenShift Container Platform on a single node requires a server with sufficient resources to run OpenShift Container Platform services and a production workload.

Table 1.1. Minimum resource requirements

Profile	vCPU	Memory	Storage
Minimum	8 vCPUs	16 GB of RAM	120 GB

**NOTE**

One vCPU equals one physical core. However, if you enable simultaneous multithreading (SMT), or Hyper-Threading, use the following formula to calculate the number of vCPUs that represent one physical core:

- $(\text{threads per core} \times \text{cores}) \times \text{sockets} = \text{vCPUs}$
- Adding Operators during the installation process might increase the minimum resource requirements.

The server must have a Baseboard Management Controller (BMC) when booting with virtual media.

**NOTE**

BMC is not supported on IBM Z® and IBM Power®.

- **Networking:** The server must have access to the internet or access to a local registry if it is not connected to a routable network. The server must have a DHCP reservation or a static IP address for the Kubernetes API, ingress route, and cluster node domain names. You must configure the DNS to resolve the IP address to each of the following fully qualified domain names (FQDN):

Table 1.2. Required DNS records

Usage	FQDN	Description
Kubernetes API	api.<cluster_name>.<base_domain>	Add a DNS A/AAAA or CNAME record. This record must be resolvable by both clients external to the cluster and within the cluster.

Usage	FQDN	Description
Internal API	api-int.<cluster_name>.<base_domain>	Add a DNS A/AAAA or CNAME record when creating the ISO manually. This record must be resolvable by nodes within the cluster.
Ingress route	*.apps.<cluster_name>.<base_domain>	Add a wildcard DNS A/AAAA or CNAME record that targets the node. This record must be resolvable by both clients external to the cluster and within the cluster.



IMPORTANT

Without persistent IP addresses, communications between the **apiserver** and **etcd** might fail.

CHAPTER 2. INSTALLING OPENSIFT ON A SINGLE NODE

You can install single-node OpenShift by using either the web-based Assisted Installer or the **coreos-installer** tool to generate a discovery ISO image. The discovery ISO image writes the Red Hat Enterprise Linux CoreOS (RHCOS) system configuration to the target installation disk, so that you can run a single-cluster node to meet your needs.

Consider using single-node OpenShift when you want to run a cluster in a low-resource or an isolated environment for testing, troubleshooting, training, or small-scale project purposes.

2.1. INSTALLING SINGLE-NODE OPENSIFT USING THE ASSISTED INSTALLER

To install OpenShift Container Platform on a single node, use the web-based Assisted Installer wizard to guide you through the process and manage the installation.

See the [Assisted Installer for OpenShift Container Platform](#) documentation for details and configuration options.

2.1.1. Generating the discovery ISO with the Assisted Installer

Installing OpenShift Container Platform on a single node requires a discovery ISO, which the Assisted Installer can generate.

Procedure

1. On the administration host, open a browser and navigate to [Red Hat OpenShift Cluster Manager](#).
2. Click **Create New Cluster** to create a new cluster.
3. In the **Cluster name** field, enter a name for the cluster.
4. In the **Base domain** field, enter a base domain. For example:

```
example.com
```

All DNS records must be subdomains of this base domain and include the cluster name, for example:

```
<cluster_name>.example.com
```



NOTE

You cannot change the base domain or cluster name after cluster installation.

5. Select **Install single node OpenShift (SNO)** and complete the rest of the wizard steps. Download the discovery ISO.
6. Complete the remaining Assisted Installer wizard steps.



IMPORTANT

Ensure that you take note of the discovery ISO URL for installing with virtual media.

If you enable OpenShift Virtualization during this process, you must have a second local storage device of at least 50GiB for your virtual machines.

Additional resources

- [Persistent storage using logical volume manager storage](#)
- [What you can do with OpenShift Virtualization](#)

2.1.2. Installing single-node OpenShift with the Assisted Installer

Use the Assisted Installer to install the single-node cluster.

Prerequisites

- Ensure that the boot drive order in the server BIOS settings defaults to booting the server from the target installation disk.

Procedure

1. Attach the discovery ISO image to the target host.
2. Boot the server from the discovery ISO image. The discovery ISO image writes the system configuration to the target installation disk and automatically triggers a server restart.
3. On the administration host, return to the browser. Wait for the host to appear in the list of discovered hosts. If necessary, reload the [Assisted Clusters](#) page and select the cluster name.
4. Complete the install wizard steps. Add networking details, including a subnet from the available subnets. Add the SSH public key if necessary.
5. Monitor the installation's progress. Watch the cluster events. After the installation process finishes writing the operating system image to the server's hard disk, the server restarts.
6. Optional: Remove the discovery ISO image.
The server restarts several times automatically, deploying the control plane.

Additional resources

- [Creating a bootable ISO image on a USB drive](#)
- [Booting from an HTTP-hosted ISO image using the Redfish API](#)
- [Adding worker nodes to single-node OpenShift clusters](#)

2.2. INSTALLING SINGLE-NODE OPENSIFT MANUALLY

To install OpenShift Container Platform on a single node, first generate the installation ISO, and then boot the server from the ISO. You can monitor the installation using the **openshift-install** installation program.

Additional resources

- [Networking requirements for user-provisioned infrastructure](#)
- [User-provisioned DNS requirements](#)
- [Configuring DHCP or static IP addresses](#)

2.2.1. Generating the installation ISO with coreos-installer

Installing OpenShift Container Platform on a single node requires an installation ISO, which you can generate with the following procedure.

Prerequisites

- Install **podman**.



NOTE

See "Requirements for installing OpenShift on a single node" for networking requirements, including DNS records.

Procedure

1. Set the OpenShift Container Platform version:

```
$ export OCP_VERSION=<ocp_version> 1
```

- 1 Replace **<ocp_version>** with the current version, for example, **latest-4.17**

2. Set the host architecture:

```
$ export ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture, for example, **aarch64** or **x86_64**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-  
v4/clients/ocp/$OCP_VERSION/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar zxf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/clients/ocp/$OCP_VERSION/openshift-install-linux.tar.gz -o openshift-install-linux.tar.gz
```

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

```
$ chmod +x openshift-install
```

- Retrieve the RHCOS ISO URL by running the following command:

```
$ export ISO_URL=$(./openshift-install coreos print-stream-json | grep location | grep $ARCH
| grep iso | cut -d\" -f4)
```

- Download the RHCOS ISO:

```
$ curl -L $ISO_URL -o rhcos-live.iso
```

- Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9
```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.

- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
- 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 9 Add the public SSH key from the administration host so that you can log in to the cluster after installation.

8. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

```
$ ./openshift-install --dir=ocp create single-node-ignition-config
```

9. Embed the ignition data into the RHCOS ISO by running the following commands:

```
$ alias coreos-installer='podman run --privileged --pull always --rm \
-v /dev:/dev -v /run/udev:/run/udev -v $PWD:/data \
-w /data quay.io/coreos/coreos-installer:release'
```

```
$ coreos-installer iso ignition embed -fi ocp/bootstrap-in-place-for-live-iso.ign rhcos-live.iso
```

Additional resources

- See [Requirements for installing OpenShift on a single node](#) for more information about installing OpenShift Container Platform on a single node.
- See [Cluster capabilities](#) for more information about enabling cluster capabilities that were disabled before installation.
- See [Optional cluster capabilities in OpenShift Container Platform 4.17](#) for more information about the features provided by each capability.

2.2.2. Monitoring the cluster installation using openshift-install

Use **openshift-install** to monitor the progress of the single-node cluster installation.

Prerequisites

- Ensure that the boot drive order in the server BIOS settings defaults to booting the server from the target installation disk.

Procedure

1. Attach the discovery ISO image to the target host.
2. Boot the server from the discovery ISO image. The discovery ISO image writes the system configuration to the target installation disk and automatically triggers a server restart.
3. On the administration host, monitor the installation by running the following command:

```
$ ./openshift-install --dir=ocp wait-for install-complete
```

4. Optional: Remove the discovery ISO image.
The server restarts several times while deploying the control plane.

Verification

- After the installation is complete, check the environment by running the following command:

```
$ export KUBECONFIG=ocp/auth/kubeconfig
```

```
$ oc get nodes
```

Example output

```
NAME                STATUS ROLES    AGE   VERSION
control-plane.example.com Ready  master,worker 10m   v1.30.3
```

Additional resources

- [Creating a bootable ISO image on a USB drive](#)
- [Booting from an HTTP-hosted ISO image using the Redfish API](#)
- [Adding worker nodes to single-node OpenShift clusters](#)

2.3. INSTALLING SINGLE-NODE OPENSIFT ON CLOUD PROVIDERS

2.3.1. Additional requirements for installing single-node OpenShift on a cloud provider

The documentation for installer-provisioned installation on cloud providers is based on a high availability cluster consisting of three control plane nodes. When referring to the documentation, consider the differences between the requirements for a single-node OpenShift cluster and a high availability cluster.

- A high availability cluster requires a temporary bootstrap machine, three control plane machines, and at least two compute machines. For a single-node OpenShift cluster, you need only a temporary bootstrap machine and one cloud instance for the control plane node and no compute nodes.
- The minimum resource requirements for high availability cluster installation include a control plane node with 4 vCPUs and 100GB of storage. For a single-node OpenShift cluster, you must have a minimum of 8 vCPUs and 120GB of storage.

- The **controlPlane.replicas** setting in the **install-config.yaml** file should be set to **1**.
- The **compute.replicas** setting in the **install-config.yaml** file should be set to **0**. This makes the control plane node schedulable.

2.3.2. Supported cloud providers for single-node OpenShift

The following table contains a list of supported cloud providers and CPU architectures.

Table 2.1. Supported cloud providers

Cloud provider	CPU architecture
Amazon Web Service (AWS)	x86_64 and AArch64
Microsoft Azure	x86_64
Google Cloud Platform (GCP)	x86_64 and AArch64

2.3.3. Installing single-node OpenShift on AWS

Installing a single-node cluster on AWS requires installer-provisioned installation using the "Installing a cluster on AWS with customizations" procedure.

Additional resources

- [Installing a cluster on AWS with customizations](#)

2.3.4. Installing single-node OpenShift on Azure

Installing a single node cluster on Azure requires installer-provisioned installation using the "Installing a cluster on Azure with customizations" procedure.

Additional resources

- [Installing a cluster on Azure with customizations](#)

2.3.5. Installing single-node OpenShift on GCP

Installing a single node cluster on GCP requires installer-provisioned installation using the "Installing a cluster on GCP with customizations" procedure.

Additional resources

- [Installing a cluster on GCP with customizations](#)

2.4. CREATING A BOOTABLE ISO IMAGE ON A USB DRIVE

You can install software using a bootable USB drive that contains an ISO image. Booting the server with the USB drive prepares the server for the software installation.

Procedure

1. On the administration host, insert a USB drive into a USB port.
2. Create a bootable USB drive, for example:

```
# dd if=<path_to_iso> of=<path_to_usb> status=progress
```

where:

<path_to_iso>

is the relative path to the downloaded ISO file, for example, **rhcos-live.iso**.

<path_to_usb>

is the location of the connected USB drive, for example, **/dev/sdb**.

After the ISO is copied to the USB drive, you can use the USB drive to install software on the server.

2.5. BOOTING FROM AN HTTP-HOSTED ISO IMAGE USING THE REDFISH API

You can provision hosts in your network using ISOs that you install using the Redfish Baseboard Management Controller (BMC) API.



NOTE

This example procedure demonstrates the steps on a Dell server.



IMPORTANT

Ensure that you have the latest firmware version of iDRAC that is compatible with your hardware. If you have any issues with the hardware or firmware, you must contact the provider.

Prerequisites

- Download the installation Red Hat Enterprise Linux CoreOS (RHCOS) ISO.
- Use a Dell PowerEdge server that is compatible with iDRAC9.

Procedure

1. Copy the ISO file to an HTTP server accessible in your network.
2. Boot the host from the hosted ISO file, for example:
 - a. Call the Redfish API to set the hosted ISO as the **VirtualMedia** boot media by running the following command:

```
$ curl -k -u <bmc_username>:<bmc_password> -d '{"Image": "<hosted_iso_file>",
  "Inserted": true}' -H "Content-Type: application/json" -X POST
<host_bmc_address>/redfish/v1/Managers/iDRAC.Embedded.1/VirtualMedia/CD/Actions/Vi
rtualMedia.InsertMedia
```

Where:

<bmc_username>:<bmc_password>

Is the username and password for the target host BMC.

<hosted_iso_file>

Is the URL for the hosted installation ISO, for example:

<http://webserver.example.com/rhcos-live-minimal.iso>. The ISO must be accessible from the target host machine.

<host_bmc_address>

Is the BMC IP address of the target host machine.

- b. Set the host to boot from the **VirtualMedia** device by running the following command:

```
$ curl -k -u <bmc_username>:<bmc_password> -X PATCH -H 'Content-Type: application/json' -d '{"Boot": {"BootSourceOverrideTarget": "Cd", "BootSourceOverrideMode": "UEFI", "BootSourceOverrideEnabled": "Once"}}' <host_bmc_address>/redfish/v1/Systems/System.Embedded.1
```

- c. Reboot the host:

```
$ curl -k -u <bmc_username>:<bmc_password> -d '{"ResetType": "ForceRestart"}' -H 'Content-type: application/json' -X POST <host_bmc_address>/redfish/v1/Systems/System.Embedded.1/Actions/ComputerSystem.Reset
```

- d. Optional: If the host is powered off, you can boot it using the **{"ResetType": "On"}** switch. Run the following command:

```
$ curl -k -u <bmc_username>:<bmc_password> -d '{"ResetType": "On"}' -H 'Content-type: application/json' -X POST <host_bmc_address>/redfish/v1/Systems/System.Embedded.1/Actions/ComputerSystem.Reset
```

2.6. CREATING A CUSTOM LIVE RHCOS ISO FOR REMOTE SERVER ACCESS

In some cases, you cannot attach an external disk drive to a server, however, you need to access the server remotely to provision a node. It is recommended to enable SSH access to the server. You can create a live RHCOS ISO with SSHd enabled and with predefined credentials so that you can access the server after it boots.

Prerequisites

- You installed the **butane** utility.

Procedure

- Download the **coreos-installer** binary from the **coreos-installer** image [mirror](#) page.
- Download the latest live RHCOS ISO from mirror.openshift.com.
- Create the **embedded.yaml** file that the **butane** utility uses to create the Ignition file:

```

variant: openshift
version: 4.17.0
metadata:
  name: sshd
  labels:
    machineconfiguration.openshift.io/role: worker
passwd:
users:
  - name: core 1
  ssh_authorized_keys:
    - '<ssh_key>'

```

1 The **core** user has sudo privileges.

- Run the **butane** utility to create the Ignition file using the following command:

```
$ butane -pr embedded.yaml -o embedded.ign
```

- After the Ignition file is created, you can include the configuration in a new live RHCOS ISO, which is named **rhcos-sshd-4.17.0-x86_64-live.x86_64.iso**, with the **coreos-installer** utility:

```
$ coreos-installer iso ignition embed -i embedded.ign rhcos-4.17.0-x86_64-live.x86_64.iso -o rhcos-sshd-4.17.0-x86_64-live.x86_64.iso
```

Verification

- Check that the custom live ISO can be used to boot the server by running the following command:

```
# coreos-installer iso ignition show rhcos-sshd-4.17.0-x86_64-live.x86_64.iso
```

Example output

```

{
  "ignition": {
    "version": "3.2.0"
  },
  "passwd": {
    "users": [
      {
        "name": "core",
        "sshAuthorizedKeys": [
          "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQAZnG8AlzIDAhpyENpK2qKiTT8EbRWOrz7NXjRzo
pbPu215mocaJgjjwJjh1cYhgPhpAp6M/ttTk7I4OI7g4588Apx4bwJep6oWTU35LkY8ZxkGVPJL
8kVITdKQviDv3XX12l4QfnDom4tm4gVbRH0gNT1wzhnLP+LKYM2Ohr9D7p9NBnAdro6k++X
WgkDeijLRUTwdEyWunldW1f8G0Mg8Y1Xzr13BUo3+8aey7HLKJMDtobkz/C8ESYA/f7HJc5Fx
F0XbapWWovSSDJrr9OmIL9f4TfE+cQk3s+eoKiz2bgNPRgEEwihVbGsCN4grA+RzLCAOpec+
2dTJrQvFqsD alosadag@sonnelicht.local"
        ]
      }
    ]
  }
}

```



2.7. INSTALLING SINGLE-NODE OPENSIFT WITH IBM Z AND IBM LINUXONE

Installing a single-node cluster on IBM Z® and IBM® LinuxONE requires user-provisioned installation using one of the following procedures:

- [Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE](#)
- [Installing a cluster with RHEL KVM on IBM Z® and IBM® LinuxONE](#)
- [Installing a cluster in an LPAR on IBM Z® and IBM® LinuxONE](#)




NOTE

Installing a single-node cluster on IBM Z® simplifies installation for development and test environments and requires less resource requirements at entry level.

Hardware requirements

- The equivalent of two Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to both connect to the **LoadBalancer** service and to serve data for traffic outside the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Z®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.

2.7.1. Installing single-node OpenShift with z/VM on IBM Z and IBM LinuxONE

Prerequisites

- You have installed **podman**.

Procedure

1. Set the OpenShift Container Platform version by running the following command:

```
$ OCP_VERSION=<ocp_version> 1
```

- 1** Replace **<ocp_version>** with the current version. For example, **latest-4.17**.

2. Set the host architecture by running the following command:

```
$ ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture **s390x**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar zxf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-install-linux.tar.gz -o openshift-install-
linux.tar.gz
```

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

```
$ chmod +x openshift-install
```

5. Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
```

```
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9
```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
- 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 9 Add the public SSH key from the administration host so that you can log in to the cluster after installation.


6. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

```
$ ./openshift-install --dir=ocp create single-node-ignition-config
```

7. Obtain the RHEL **kernel**, **initramfs**, and **rootfs** artifacts from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



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. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

kernel

```
rhcos-<version>-live-kernel-<architecture>
```

initramfs

rhcos-<version>-live-initramfs.<architecture>.img

rootfs

rhcos-<version>-live-rootfs.<architecture>.img



NOTE

The **rootfs** image is the same for FCP and DASD.

8. Move the following artifacts and files to an HTTP or HTTPS server:

- Downloaded RHEL live **kernel**, **initramfs**, and **rootfs** artifacts
- Ignition files

9. Create parameter files for a particular virtual machine:

Example parameter file

```

cio_ignore=all,!condev rd.neednet=1 \
console=ttysclp0 \
ignition.firstboot ignition.platform.id=metal \
ignition.config.url=http://<http_server>:8080/ignition/bootstrap-in-place-for-live-iso.ign \ 1
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 2
ip=<ip>::<gateway>:<mask>:<hostname>::none nameserver=<dns> \ 3
rd.znet=qeth,0.0.bdd0,0.0.bdd1,0.0.bdd2,layer2=1 \
rd.dasd=0.0.4411 \ 4
rd.zfcp=0.0.8001,0x50050763040051e3,0x4000406300000000 \ 5
zfcp.allow_lun_scan=0

```

- 1 For the **ignition.config.url=** parameter, specify the Ignition file for the machine role. Only HTTP and HTTPS protocols are supported.
- 2 For the **coreos.live.rootfs_url=** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 3 For the **ip=** parameter, assign the IP address automatically using DHCP or manually as described in "Installing a cluster with z/VM on IBM Z® and IBM® LinuxONE".
- 4 For installations on DASD-type disks, use **rd.dasd=** to specify the DASD where RHCOS is to be installed. Omit this entry for FCP-type disks.
- 5 For installations on FCP-type disks, use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. Omit this entry for DASD-type disks.

Leave all other parameters unchanged.

10. Transfer the following artifacts, files, and images to z/VM. For example by using FTP:

- **kernel** and **initramfs** artifacts
- Parameter files
- RHCOS images

For details about how to transfer the files with FTP and boot from the virtual reader, see [Installing under Z/VM](#).

11. Punch the files to the virtual reader of the z/VM guest virtual machine that is to become your bootstrap node.
12. Log in to CMS on the bootstrap machine.
13. IPL the bootstrap machine from the reader by running the following command:

```
$ cp ipl c
```

14. After the first reboot of the virtual machine, run the following commands directly after one another:

- a. To boot a DASD device after first reboot, run the following commands:

```
$ cp i <devno> clear loadparm prompt
```

where:

<devno>

Specifies the device number of the boot device as seen by the guest.

```
$ cp vi vmmsg 0 <kernel_parameters>
```

where:

<kernel_parameters>

Specifies a set of kernel parameters to be stored as system control program data (SCPDATA). When booting Linux, these kernel parameters are concatenated to the end of the existing kernel parameters that are used by your boot configuration. The combined parameter string must not exceed 896 characters.

- b. To boot an FCP device after first reboot, run the following commands:

```
$ cp set loaddev portname <wwpn> lun <lun>
```

where:

<wwpn>

Specifies the target port and **<lun>** the logical unit in hexadecimal format.

```
$ cp set loaddev bootprog <n>
```

where:

<n>

Specifies the kernel to be booted.

```
$ cp set loaddev scpdata {APPEND|NEW} '<kernel_parameters>'
```

where:

<kernel_parameters>

Specifies a set of kernel parameters to be stored as system control program data (SCPDATA). When booting Linux, these kernel parameters are concatenated to the end of the existing kernel parameters that are used by your boot configuration. The combined parameter string must not exceed 896 characters.

<APPEND|NEW>

Optional: Specify **APPEND** to append kernel parameters to existing SCPDATA. This is the default. Specify **NEW** to replace existing SCPDATA.

Example

```
$ cp set loaddev scpdata
'rd.zfcp=0.0.8001,0x500507630a0350a4,0x4000409D00000000
ip=enbddd0:dhcp::02:00:00:02:34:02 rd.neednet=1'
```

To start the IPL and boot process, run the following command:

```
$ cp i <devno>
```

where:

<devno>

Specifies the device number of the boot device as seen by the guest.

2.7.2. Installing single-node OpenShift with RHEL KVM on IBM Z and IBM LinuxONE

Prerequisites

- You have installed **podman**.

Procedure

- Set the OpenShift Container Platform version by running the following command:

```
$ OCP_VERSION=<ocp_version> 1
```

- Replace **<ocp_version>** with the current version. For example, **latest-4.17**.

- Set the host architecture by running the following command:

```
$ ARCH=<architecture> 1
```

- Replace **<architecture>** with the target host architecture **s390x**.

- Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar xzf oc.tar.gz
```

```
$ chmod +x oc
```

- Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-install-linux.tar.gz -o openshift-install-
linux.tar.gz
```

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

```
$ chmod +x openshift-install
```

- Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9
```

- Add the cluster domain name.
- Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- Set the **metadata** name to the cluster name.

- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
- 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 9 Add the public SSH key from the administration host so that you can log in to the cluster after installation.

6. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

```
$ ./openshift-install --dir=ocp create single-node-ignition-config
```

7. Obtain the RHEL **kernel**, **initramfs**, and **rootfs** artifacts from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



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. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

kernel

```
rhcos-<version>-live-kernel-<architecture>
```

initramfs

```
rhcos-<version>-live-initramfs.<architecture>.img
```

rootfs

```
rhcos-<version>-live-rootfs.<architecture>.img
```

8. Before you launch **virt-install**, move the following files and artifacts to an HTTP or HTTPS server:
 - Downloaded RHEL live **kernel**, **initramfs**, and **rootfs** artifacts
 - Ignition files
9. Create the KVM guest nodes by using the following components:
 - RHEL **kernel** and **initramfs** artifacts

- Ignition files
- The new disk image
- Adjusted parm line arguments

```
$ virt-install \
  --name <vm_name> \
  --autostart \
  --memory=<memory_mb> \
  --cpu host \
  --vcpus <vcpus> \
  --location <media_location>,kernel=<rhcos_kernel>,initrd=<rhcos_initrd> \ 1
  --disk size=100 \
  --network network=<virt_network_parm> \
  --graphics none \
  --noautoconsole \
  --extra-args "rd.neednet=1 ignition.platform.id=metal ignition.firstboot" \
  --extra-args "ignition.config.url=http://<http_server>/bootstrap.ign" \ 2
  --extra-args "coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.
  <architecture>.img" \ 3
  --extra-args "ip=<ip>::<gateway>:<mask>:<hostname>::none" \ 4
  --extra-args "nameserver=<dns>" \
  --extra-args "console=ttysclp0" \
  --wait
```

- 1 For the **--location** parameter, specify the location of the kernel/initrd on the HTTP or HTTPS server.
- 2 Specify the location of the **bootstrap.ign** config file. Only HTTP and HTTPS protocols are supported.
- 3 For the **coreos.live.rootfs_url=** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 4 For the **ip=** parameter, assign the IP address manually as described in "Installing a cluster with RHEL KVM on IBM Z® and IBM® LinuxONE".

2.7.3. Installing single-node OpenShift in an LPAR on IBM Z and IBM LinuxONE

Prerequisites

- If you are deploying a single-node cluster there are zero compute nodes, the Ingress Controller pods run on the control plane nodes. In single-node cluster deployments, you must configure your application ingress load balancer to route HTTP and HTTPS traffic to the control plane nodes. See the *Load balancing requirements for user-provisioned infrastructure* section for more information.

Procedure

1. Set the OpenShift Container Platform version by running the following command:

```
$ OCP_VERSION=<ocp_version> 1
```

- 1 Replace **<ocp_version>** with the current version. For example, **latest-4.17**.

2. Set the host architecture by running the following command:

```
$ ARCH=<architecture> 1
```

- 1 Replace **<architecture>** with the target host architecture **s390x**.

3. Download the OpenShift Container Platform client (**oc**) and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-client-linux.tar.gz -o oc.tar.gz
```

```
$ tar zxvf oc.tar.gz
```

```
$ chmod +x oc
```

4. Download the OpenShift Container Platform installer and make it available for use by entering the following commands:

```
$ curl -k https://mirror.openshift.com/pub/openshift-
v4/${ARCH}/clients/ocp/${OCP_VERSION}/openshift-install-linux.tar.gz -o openshift-install-
linux.tar.gz
```

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

```
$ chmod +x openshift-install
```

5. Prepare the **install-config.yaml** file:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
```

```

serviceNetwork:
- 172.30.0.0/16
platform:
  none: {}
pullSecret: '<pull_secret>' 7
sshKey: |
  <ssh_key> 8

```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures the cluster runs on a single node.
- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.
- 7 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
- 8 Add the public SSH key from the administration host so that you can log in to the cluster after installation.

6. Generate OpenShift Container Platform assets by running the following commands:

```
$ mkdir ocp
```

```
$ cp install-config.yaml ocp
```

7. Change to the directory that contains the OpenShift Container Platform installation program and generate the Kubernetes manifests for the cluster:

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

- 1 For **<installation_directory>**, specify the installation directory that contains the **install-config.yaml** file you created.

8. Check that the **mastersSchedulable** parameter in the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** Kubernetes manifest file is set to **true**.

- a. Open the **<installation_directory>/manifests/cluster-scheduler-02-config.yml** file.
- b. Locate the **mastersSchedulable** parameter and ensure that it is set to **true** as shown in the following **spec** stanza:

```
spec:
  mastersSchedulable: true
  status: {}
```

c. Save and exit the file.

9. Create the Ignition configuration files by running the following command from the directory that contains the installation program:

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

1 For **<installation_directory>**, specify the same installation directory.

10. Obtain the RHEL **kernel**, **initramfs**, and **rootfs** artifacts from the [Product Downloads](#) page on the Red Hat Customer Portal or from the [RHCOS image mirror](#) page.



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. Only use the appropriate **kernel**, **initramfs**, and **rootfs** artifacts described in the following procedure.

The file names contain the OpenShift Container Platform version number. They resemble the following examples:

kernel

```
rhcos-<version>-live-kernel-<architecture>
```

initramfs

```
rhcos-<version>-live-initramfs.<architecture>.img
```

rootfs

```
rhcos-<version>-live-rootfs.<architecture>.img
```



NOTE

The **rootfs** image is the same for FCP and DASD.

11. Move the following artifacts and files to an HTTP or HTTPS server:
 - Downloaded RHEL live **kernel**, **initramfs**, and **rootfs** artifacts
 - Ignition files
12. Create a parameter file for the bootstrap in an LPAR:

Example parameter file for the bootstrap machine

```
cio_ignore=all,!condev rd.neednet=1 \
console=ttySCLP0 \
```

```

coreos.inst.install_dev=/dev/<block_device> \ 1
coreos.inst.ignition_url=http://<http_server>/bootstrap.ign \ 2
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \ 3
ip=<ip>::<gateway>:<netmask>:<hostname>::none nameserver=<dns> \ 4
rd.znet=qeth,0.0.1140,0.0.1141,0.0.1142,layer2=1,portno=0 \
rd.dasd=0.0.4411 \ 5
rd.zfcp=0.0.8001,0x50050763040051e3,0x4000406300000000 \ 6
zfcp.allow_lun_scan=0

```

- 1 Specify the block device on the system to install to. For installations on DASD-type disk use **dasda**, for installations on FCP-type disks use **sda**.
- 2 Specify the location of the **bootstrap.ign** config file. Only HTTP and HTTPS protocols are supported.
- 3 For the **coreos.live.rootfs_url=** artifact, specify the matching **rootfs** artifact for the **kernel** and **initramfs** you are booting. Only HTTP and HTTPS protocols are supported.
- 4 For the **ip=** parameter, assign the IP address manually as described in "Installing a cluster in an LPAR on IBM Z® and IBM® LinuxONE".
- 5 For installations on DASD-type disks, use **rd.dasd=** to specify the DASD where RHCOS is to be installed. Omit this entry for FCP-type disks.
- 6 For installations on FCP-type disks, use **rd.zfcp=<adapter>,<wwpn>,<lun>** to specify the FCP disk where RHCOS is to be installed. Omit this entry for DASD-type disks.

You can adjust further parameters if required.

13. Create a parameter file for the control plane in an LPAR:

Example parameter file for the control plane machine

```

cio_ignore=all,!condev rd.neednet=1 \
console=ttysclp0 \
coreos.inst.install_dev=/dev/<block_device> \
coreos.inst.ignition_url=http://<http_server>/master.ign \ 1
coreos.live.rootfs_url=http://<http_server>/rhcos-<version>-live-rootfs.<architecture>.img \
ip=<ip>::<gateway>:<netmask>:<hostname>::none nameserver=<dns> \
rd.znet=qeth,0.0.1140,0.0.1141,0.0.1142,layer2=1,portno=0 \
rd.dasd=0.0.4411 \
rd.zfcp=0.0.8001,0x50050763040051e3,0x4000406300000000 \
zfcp.allow_lun_scan=0

```

- 1 Specify the location of the **master.ign** config file. Only HTTP and HTTPS protocols are supported.

14. Transfer the following artifacts, files, and images to the LPAR. For example by using FTP:

- **kernel** and **initramfs** artifacts
- Parameter files
- RHCOS images

For details about how to transfer the files with FTP and boot, see [Installing in an LPAR](#).

15. Boot the bootstrap machine.
16. Boot the control plane machine.

2.8. INSTALLING SINGLE-NODE OPENSIFT WITH IBM POWER

Installing a single-node cluster on IBM Power® requires user-provisioned installation using the "Installing a cluster with IBM Power®" procedure.



NOTE

Installing a single-node cluster on IBM Power® simplifies installation for development and test environments and requires less resource requirements at entry level.

Hardware requirements

- The equivalent of two Integrated Facilities for Linux (IFL), which are SMT2 enabled, for each cluster.
- At least one network connection to connect to the **LoadBalancer** service and to serve data for traffic outside of the cluster.



NOTE

You can use dedicated or shared IFLs to assign sufficient compute resources. Resource sharing is one of the key strengths of IBM Power®. However, you must adjust capacity correctly on each hypervisor layer and ensure sufficient resources for every OpenShift Container Platform cluster.

Additional resources

- [Installing a cluster on IBM Power®](#)

2.8.1. Setting up basion for single-node OpenShift with IBM Power

Prior to installing single-node OpenShift on IBM Power®, you must set up bastion. Setting up a bastion server for single-node OpenShift on IBM Power® requires the configuration of the following services:

- PXE is used for the single-node OpenShift cluster installation. PXE requires the following services to be configured and run:
 - DNS to define api, api-int, and *.apps
 - DHCP service to enable PXE and assign an IP address to single-node OpenShift node
 - HTTP to provide ignition and RHCOS rootfs image
 - TFTP to enable PXE
- You must install **dnsmasq** to support DNS, DHCP and PXE, httpd for HTTP.

Use the following procedure to configure a bastion server that meets these requirements.

Procedure

1. Use the following command to install **grub2**, which is required to enable PXE for PowerVM:

```
grub2-mknetdir --net-directory=/var/lib/tftpboot
```

Example `/var/lib/tftpboot/boot/grub2/grub.cfg` file

```
default=0
fallback=1
timeout=1
if [ ${net_default_mac} == fa:b0:45:27:43:20 ]; then
menuentry "CoreOS (BIOS)" {
    echo "Loading kernel"
    linux "/rhcos/kernel" ip=dhcp rd.neednet=1 ignition.platform.id=metal ignition.firstboot
    coreos.live.rootfs_url=http://192.168.10.5:8000/install/rootfs.img
    ignition.config.url=http://192.168.10.5:8000/ignition/sno.ign
    echo "Loading initrd"
    initrd "/rhcos/initramfs.img"
}
fi
```

2. Use the following commands to download RHCOS image files from the mirror repo for PXE.

- a. Enter the following command to assign the **RHCOS_URL** variable the follow 4.12 URL:

```
$ export RHCOS_URL=https://mirror.openshift.com/pub/openshift-
v4/ppc64le/dependencies/rhcos/4.12/latest/
```

- b. Enter the following command to navigate to the `/var/lib/tftpboot/rhcos` directory:

```
$ cd /var/lib/tftpboot/rhcos
```

- c. Enter the following command to download the specified RHCOS kernel file from the URL stored in the **RHCOS_URL** variable:

```
$ wget ${RHCOS_URL}/rhcos-live-kernel-ppc64le -o kernel
```

- d. Enter the following command to download the RHCOS **initramfs** file from the URL stored in the **RHCOS_URL** variable:

```
$ wget ${RHCOS_URL}/rhcos-live-initramfs.ppc64le.img -o initramfs.img
```

- e. Enter the following command to navigate to the `/var//var/www/html/install/` directory:

```
$ cd /var//var/www/html/install/
```

- f. Enter the following command to download, and save, the RHCOS **root filesystem** image file from the URL stored in the **RHCOS_URL** variable:

```
$ wget ${RHCOS_URL}/rhcos-live-rootfs.ppc64le.img -o rootfs.img
```

3. To create the ignition file for a single-node OpenShift cluster, you must create the **install-config.yaml** file.

- a. Enter the following command to create the work directory that holds the file:

```
$ mkdir -p ~/sno-work
```

- b. Enter the following command to navigate to the **~/sno-work** directory:

```
$ cd ~/sno-work
```

- c. Use the following sample file can to create the required **install-config.yaml** in the **~/sno-work** directory:

```
apiVersion: v1
baseDomain: <domain> 1
compute:
- name: worker
  replicas: 0 2
controlPlane:
  name: master
  replicas: 1 3
metadata:
  name: <name> 4
networking: 5
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16 6
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  none: {}
bootstrapInPlace:
  installationDisk: /dev/disk/by-id/<disk_id> 7
pullSecret: '<pull_secret>' 8
sshKey: |
  <ssh_key> 9
```

- 1 Add the cluster domain name.
- 2 Set the **compute** replicas to **0**. This makes the control plane node schedulable.
- 3 Set the **controlPlane** replicas to **1**. In conjunction with the previous **compute** setting, this setting ensures that the cluster runs on a single node.
- 4 Set the **metadata** name to the cluster name.
- 5 Set the **networking** details. OVN-Kubernetes is the only allowed network plugin type for single-node clusters.
- 6 Set the **cidr** value to match the subnet of the single-node OpenShift cluster.

- 7 Set the path to the installation disk drive, for example, **/dev/disk/by-id/wwn-0x64cd98f04fde100024684cf3034da5c2**.
 - 8 Copy the [pull secret from Red Hat OpenShift Cluster Manager](#) and add the contents to this configuration setting.
 - 9 Add the public SSH key from the administration host so that you can log in to the cluster after installation.
4. Download the **openshift-install** image to create the ignition file and copy it to the **http** directory.

- a. Enter the following command to download the **openshift-install-linux-4.12.0** .tar file:

```
$ wget https://mirror.openshift.com/pub/openshift-
v4/ppc64le/clients/ocp/4.12.0/openshift-install-linux-4.12.0.tar.gz
```

- b. Enter the following command to unpack the **openshift-install-linux-4.12.0.tar.gz** archive:

```
$ tar xzvf openshift-install-linux-4.12.0.tar.gz
```

- c. Enter the following command to

```
$ ./openshift-install --dir=~/.sno-work create create single-node-ignition-config
```

- d. Enter the following command to create the ignition file:

```
$ cp ~/.sno-work/single-node-ignition-config.ign /var/www/html/ignition/sno.ign
```

- e. Enter the following command to restore SELinux file for the **/var/www/html** directory:

```
$ restorecon -vR /var/www/html || true
```

Bastion now has all the required files and is properly configured in order to install single-node OpenShift.

2.8.2. Installing single-node OpenShift with IBM Power

Prerequisites

- You have set up bastion.

Procedure

There are two steps for the single-node OpenShift cluster installation. First the single-node OpenShift logical partition (LPAR) needs to boot up with PXE, then you need to monitor the installation progress.

1. Use the following command to boot powerVM with netboot:

```
$ lpar_netboot -i -D -f -t ent -m <sno_mac> -s auto -d auto -S <server_ip> -C <sno_ip> -G
<gateway> <lpar_name> default_profile <cec_name>
```

where:

sno_mac

Specifies the MAC address of the single-node OpenShift cluster.

sno_ip

Specifies the IP address of the single-node OpenShift cluster.

server_ip

Specifies the IP address of bastion (PXE server).

gateway

Specifies the Network's gateway IP.

lpar_name

Specifies the single-node OpenShift lpar name in HMC.

cec_name

Specifies the System name where the sno_lpar resides

2. After the single-node OpenShift LPAR boots up with PXE, use the **openshift-install** command to monitor the progress of installation:

- a. Run the following command after the bootstrap is complete:

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

- b. Run the following command after it returns successfully:

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