



OpenShift Container Platform 4.2

Installing on IBM Z

Installing OpenShift Container Platform 4.2 IBM Z clusters

OpenShift Container Platform 4.2 Installing on IBM Z

Installing OpenShift Container Platform 4.2 IBM Z 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 OpenShift Container Platform 4.2 clusters on IBM Z.

Table of Contents

CHAPTER 1. INSTALLING ON IBM Z	3
1.1. INSTALLING A CLUSTER ON IBM Z AND LINUXONE	3
1.1.1. Internet and Telemetry access for OpenShift Container Platform	3
1.1.2. Machine requirements for a cluster with user-provisioned infrastructure	4
1.1.2.1. Required machines	4
1.1.2.2. Network connectivity requirements	4
1.1.2.3. IBM Z network connectivity requirements	5
1.1.2.4. Minimum resource requirements	5
1.1.2.5. Minimum IBM Z system requirements	5
Hardware requirements	5
Operating system requirements	5
Disk storage for the z/VM guest virtual machines	5
Storage / Main Memory	6
1.1.2.6. Preferred IBM Z system requirements	6
Hardware requirements	6
Operating system requirements	6
Disk storage for the z/VM guest virtual machines	6
Storage / Main Memory	6
1.1.2.7. Certificate signing requests management	7
1.1.3. Creating the user-provisioned infrastructure	7
1.1.3.1. Networking requirements for user-provisioned infrastructure	7
Network topology requirements	8
1.1.3.2. User-provisioned DNS requirements	9
1.1.4. Generating an SSH private key and adding it to the agent	11
1.1.5. Obtaining the installation program	12
1.1.6. Installing the CLI	13
1.1.7. Manually creating the installation configuration file	13
1.1.7.1. Sample install-config.yaml file for IBM Z	14
1.1.8. Creating the Kubernetes manifest and Ignition config files	16
1.1.9. Creating Red Hat Enterprise Linux CoreOS (RHCOS) machines	17
1.1.10. Creating the cluster	19
1.1.11. Logging in to the cluster	20
1.1.12. Approving the CSRs for your machines	20
1.1.13. Initial Operator configuration	22
1.1.13.1. Image registry storage configuration	22
1.1.13.1.1. Configuring registry storage for IBM Z	23
1.1.13.1.2. Configuring storage for the image registry in non-production clusters	24
1.1.14. Completing installation on user-provisioned infrastructure	24
1.1.15. Collecting debugging information	26

CHAPTER 1. INSTALLING ON IBM Z


1.1. INSTALLING A CLUSTER ON IBM Z AND LINUXONE

In OpenShift Container Platform version 4.2, you can install a cluster on IBM Z or LinuxONE infrastructure that you provision.



NOTE

While this document refers only to IBM Z, all information in it also applies to LinuxONE.



IMPORTANT

Additional considerations exist for non-bare metal platforms. Review the information in the [guidelines for deploying OpenShift Container Platform on non-tested platforms](#) before you install an OpenShift Container Platform cluster.



WARNING

Restricted network installations are not tested or supported on IBM Z in OpenShift Container Platform 4.2.

Prerequisites

- Provision [persistent storage using NFS](#) for your cluster. To deploy a private image registry, your storage must provide ReadWriteMany access modes.
- Review details about the [OpenShift Container Platform installation and update](#) processes.
- If you use a firewall, you must [configure it to allow the sites](#) that your cluster requires access to.



NOTE

Be sure to also review this [site list](#) if you are configuring a proxy.

1.1.1. Internet and Telemetry access for OpenShift Container Platform

In OpenShift Container Platform 4.2, you require access to the internet to install your cluster. The Telemetry service, which runs by default to provide metrics about cluster health and the success of updates, also requires internet access. If your cluster is connected to the internet, Telemetry runs automatically, and your cluster is registered to the [Red Hat OpenShift Cluster Manager \(OCM\)](#).

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

You must have internet access to:

- Access the [Red Hat OpenShift Cluster Manager](#) page to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
- Access [Quay.io](#) to obtain the packages that are required to install your cluster.
- Obtain the packages that are required to perform cluster updates.



IMPORTANT

If your cluster cannot have direct internet access, you can perform a restricted network installation on some types of infrastructure that you provision. During that process, you download the content that is required and use it to populate a mirror registry with the packages that you need to install a cluster and generate the installation program. With some installation types, the environment that you install your cluster in will not require internet access. Before you update the cluster, you update the content of the mirror registry.

1.1.2. Machine requirements for a cluster with user-provisioned infrastructure

For a cluster that contains user-provisioned infrastructure, you must deploy all of the required machines.

1.1.2.1. Required machines

The smallest OpenShift Container Platform clusters require the following hosts:

- One temporary bootstrap machine
- Three control plane, or master, machines
- At least two compute machines, which are also known as worker machines



NOTE

The cluster requires the bootstrap machine to deploy the OpenShift Container Platform cluster on the three control plane machines. You can remove the bootstrap machine after you install the cluster.



IMPORTANT

To improve high availability of your cluster, distribute the control plane machines over different z/VM instances. These can, but need not, run on the same Z or LinuxONE hardware.

The bootstrap, control plane, and compute machines must use the Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system.

Note that RHCOS is based on Red Hat Enterprise Linux 8 and inherits all of its hardware certifications and requirements. See [Red Hat Enterprise Linux technology capabilities and limits](#) .

1.1.2.2. Network connectivity requirements

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require network in **inittamfs** during boot to fetch Ignition config files from the Machine Config Server. The machines are configured with static IP addresses. No DHCP server is required.

1.1.2.3. IBM Z network connectivity requirements

To install on IBM Z under z/VM, you require a single z/VM virtual NIC in layer 2 mode. You also need:

- A direct-attached OSA or RoCE network adapter
- A z/VM VSwitch set up. For a preferred setup, use OSA link aggregation.

1.1.2.4. Minimum resource requirements

Each cluster machine must meet the following minimum requirements:

Machine	Operating System	vCPU	Virtual RAM	Storage
Bootstrap	RHCOS	4	16 GB	120 GB
Control plane	RHCOS	4	16 GB	120 GB
Compute	RHCOS	2	8 GB	120 GB

1.1.2.5. Minimum IBM Z system requirements

You can install OpenShift Container Platform version 4.2 on the following IBM hardware:

- IBM Z: z13, z13s, all z14 models, all z15 models
- LinuxONE: all models

Hardware requirements

- 1 LPAR with 3 IFLs that supports SMT2
- 1 OSA or RoCE network adapter

Operating system requirements

- One instance of z/VM 7.1

On your z/VM instance, set up:

- 3 guest virtual machines for OpenShift Container Platform control plane machines
- 2 guest virtual machines for OpenShift Container Platform compute machines
- 1 guest virtual machine for the temporary OpenShift Container Platform bootstrap machine

Disk storage for the z/VM guest virtual machines

- FICON attached disk storage (DASDs). These can be z/VM minidisks, fullpack minidisks, or dedicated DASDs. To reach the minimum required DASD size for RHCOS installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.

**NOTE**

In OpenShift Container Platform 4.2, you must use DASD disks if you require multi-path support.

- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

1.1.2.6. Preferred IBM Z system requirements**Hardware requirements**

- 3 LPARs with 6 IFLs that support SMT2
- 1 or 2 OSA or RoCE network adapters, or both

Operating system requirements

- 2 or 3 instances of z/VM 7.1 for high availability

On your z/VM instances, set up:

- 3 guest virtual machines for OpenShift Container Platform control plane machines, one per z/VM instance
- At least 6 guest virtual machines for OpenShift Container Platform compute machines, distributed across the z/VM instances
- 1 guest virtual machine for the temporary OpenShift Container Platform bootstrap machine

Disk storage for the z/VM guest virtual machines

- FICON attached disk storage (DASDs). These can be z/VM minidisks, fullpack minidisks, or dedicated DASDs. To reach the minimum required DASD size for RHCOS installations, you need extended address volumes (EAV). If available, use HyperPAV to ensure optimal performance.

**NOTE**

In OpenShift Container Platform 4.2, you must use DASD disks if you require multi-path support.

- FCP attached disk storage

Storage / Main Memory

- 16 GB for OpenShift Container Platform control plane machines
- 8 GB for OpenShift Container Platform compute machines
- 16 GB for the temporary OpenShift Container Platform bootstrap machine

1.1.2.7. Certificate signing requests management

Because your cluster has limited access to automatic machine management when you use infrastructure that you provision, you must provide a mechanism for approving cluster certificate signing requests (CSRs) after installation. The **kube-controller-manager** only approves the kubelet client CSRs. The **machine-approver** cannot guarantee the validity of a serving certificate that is requested by using kubelet credentials because it cannot confirm that the correct machine issued the request. You must determine and implement a method of verifying the validity of the kubelet serving certificate requests and approving them.

Additional resources

- See [Bridging a HiperSockets LAN with a z/VM Virtual Switch](#) in the IBM Knowledge Center.
- See [Scaling HyperPAV alias devices on Linux guests on z/VM](#) for performance optimization.

1.1.3. Creating the user-provisioned infrastructure

Before you deploy an OpenShift Container Platform cluster that uses user-provisioned infrastructure, you must create the underlying infrastructure.

Prerequisites

- Review the [OpenShift Container Platform 4.x Tested Integrations](#) page before you create the supporting infrastructure for your cluster.

Procedure

1. Set up static IP addresses.
2. Set up an FTP server.
3. Provision the required load balancers.
4. Configure the ports for your machines.
5. Configure DNS.
6. Ensure network connectivity.

1.1.3.1. Networking requirements for user-provisioned infrastructure

All the Red Hat Enterprise Linux CoreOS (RHCOS) machines require network in **initramfs** during boot to fetch Ignition config from the Machine Config Server.

During the initial boot, the machines require an FTP server in order to establish a network connection to download their Ignition config files.

Ensure that the machines have persistent IP addresses and host names.

The Kubernetes API server, which runs on each master node after a successful cluster installation, must be able to resolve the node names of the cluster machines. If the API servers and worker nodes are in different zones, you can configure a default DNS search zone to allow the API server to resolve the node names. Another supported approach is to always refer to hosts by their fully-qualified domain names in both the node objects and all DNS requests.

You must configure the network connectivity between machines to allow cluster components to communicate. Each machine must be able to resolve the host names of all other machines in the cluster.

Table 1.1. All machines to all machines

Protocol	Port	Description
ICMP	N/A	Network reachability tests
TCP	9000-9999	Host level services, including the node exporter on ports 9100-9101 and the Cluster Version Operator on port 9099 .
	10250-10259	The default ports that Kubernetes reserves
	10256	openshift-sdn
UDP	4789	VXLAN and GENEVE
	6081	VXLAN and GENEVE
	9000-9999	Host level services, including the node exporter on ports 9100-9101 .
TCP/UDP	30000-32767	Kubernetes NodePort

Table 1.2. All machines to control plane

Protocol	Port	Description
TCP	2379-2380	etcd server, peer, and metrics ports
	6443	Kubernetes API

Network topology requirements

The infrastructure that you provision for your cluster must meet the following network topology requirements.




IMPORTANT

OpenShift Container Platform requires all nodes to have internet access to pull images for platform containers and provide telemetry data to Red Hat.

Load balancers

Before you install OpenShift Container Platform, you must provision two layer-4 load balancers. The API requires one load balancer and the default Ingress Controller needs the second load balancer to provide ingress to applications.

Port	Machines	Internal	External	Description
6443	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	x	x	Kubernetes API server
22623	Bootstrap and control plane. You remove the bootstrap machine from the load balancer after the bootstrap machine initializes the cluster control plane.	x		Machine Config server
443	The machines that run the Ingress router pods, compute, or worker, by default.	x	x	HTTPS traffic
80	The machines that run the Ingress router pods, compute, or worker by default.	x	x	HTTP traffic



NOTE


A working configuration for the Ingress router is required for an OpenShift Container Platform cluster. You must configure the Ingress router after the control plane initializes.

1.1.3.2. User-provisioned DNS requirements

The following DNS records are required for an OpenShift Container Platform cluster that uses user-provisioned infrastructure. In each record, **<cluster_name>** is the cluster name and **<base_domain>** is the cluster base domain that you specify in the **install-config.yaml** file. A complete DNS record takes the form: **<component>.<cluster_name>.<base_domain>.**

Table 1.3. Required DNS records


Component	Record	Description
Kubernetes API	api.<cluster_name>.<base_domain>.	This DNS A/AAAA or CNAME record must point to the load balancer for the control plane machines. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.

Component	Record	Description
	api-int.<cluster_name>.<base_domain>.	<p>This DNS A/AAAA or CNAME record must point to the load balancer for the control plane machines. This record must be resolvable from all the nodes within the cluster.</p> <div data-bbox="1040 517 1147 987" style="background-color: black; color: white; padding: 10px; width: fit-content;">  </div> <p>IMPORTANT</p> <p>The API server must be able to resolve the worker nodes by the host names that are recorded in Kubernetes. If it cannot resolve the node names, proxied API calls can fail, and you cannot retrieve logs from Pods.</p>
Routes	*.apps.<cluster_name>.<base_domain>.	<p>A wildcard DNS A/AAAA or CNAME record that points to the load balancer that targets the machines that run the Ingress router pods, which are the worker nodes by default. This record must be resolvable by both clients external to the cluster and from all the nodes within the cluster.</p>
etcd	etcd-<index>.<cluster_name>.<base_domain>.	<p>OpenShift Container Platform requires DNS A/AAAA records for each etcd instance to point to the control plane machines that host the instances. The etcd instances are differentiated by <index> values, which start with 0 and end with n-1, where n is the number of control plane machines in the cluster. The DNS record must resolve to an unicast IPv4 address for the control plane machine, and the records must be resolvable from all the nodes in the cluster.</p>

Component	Record	Description
	<code>_etcd-server-ssl._tcp.<cluster_name>.<base_domain>.</code>	<p>For each control plane machine, OpenShift Container Platform also requires a SRV DNS record for etcd server on that machine with priority 0, weight 10 and port 2380. A cluster that uses three control plane machines requires the following records:</p> <pre># _service._proto.name. TTL class SRV priority weight port target. _etcd-server-ssl._tcp. <cluster_name>. <base_domain>. 86400 IN SRV 0 10 2380 etcd- 0.<cluster_name>. <base_domain> _etcd-server-ssl._tcp. <cluster_name>. <base_domain>. 86400 IN SRV 0 10 2380 etcd- 1.<cluster_name>. <base_domain> _etcd-server-ssl._tcp. <cluster_name>. <base_domain>. 86400 IN SRV 0 10 2380 etcd- 2.<cluster_name>. <base_domain></pre>


1.1.4. 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.



IMPORTANT

Do not skip this procedure in production environments where disaster recovery and debugging is required.

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> 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.

1.1.5. Obtaining the installation program

Before you install OpenShift Container Platform, download the installation file on your provisioning machine.


Prerequisites

- You must install the cluster from a machine that runs Linux, for example Red Hat Enterprise Linux 8.
- 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.

1.1.6. 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>
```

1.1.7. Manually creating the installation configuration file

For installations of OpenShift Container Platform that use user-provisioned infrastructure, you must manually generate your installation configuration file.


Prerequisites

- Obtain the OpenShift Container Platform installation program and the access token for your cluster.

Procedure

1. Create an installation directory to store your required installation assets in:

```
$ mkdir <installation_directory>
```



IMPORTANT

You must create a 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.


2. Customize the following **install-config.yaml** file template and save it in the **<installation_directory>**.



NOTE

You must name this configuration file **install-config.yaml**.

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 next step of the installation process. You must back it up now.

1.1.7.1. Sample **install-config.yaml** file for IBM Z

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.

```
apiVersion: v1
baseDomain: example.com 1
compute:
- hyperthreading: Enabled 2 3
  name: worker
  replicas: 0 4
controlPlane:
  hyperthreading: Enabled 5 6
  name: master 7
  replicas: 3 8
```

```

metadata:
  name: test 9
networking:
  clusterNetwork:
    - cidr: 10.128.0.0/14 10
      hostPrefix: 23 11
  networkType: OpenShiftSDN
  serviceNetwork: 12
    - 172.30.0.0/16
platform:
  none: {} 13
pullSecret: '{"auths": ...}' 14
sshKey: 'ssh-ed25519 AAAA...' 15

```

- 1 The base domain of the cluster. All DNS records must be sub-domains of this base and include the cluster name.
- 2 5 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.
- 3 6 7 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.

- 4 You must set the value of the **replicas** parameter to **0**. This parameter controls the number of workers that the cluster creates and manages for you, which are functions that the cluster does not perform when you use user-provisioned infrastructure. You must manually deploy worker machines for the cluster to use before you finish installing OpenShift Container Platform.
- 8 The number of control plane machines that you add to the cluster. Because the cluster uses this values as the number of etcd endpoints in the cluster, the value must match the number of control plane machines that you deploy.
- 9 The cluster name that you specified in your DNS records.
- 10 A block of IP addresses from which Pod IP addresses are allocated. This block must not overlap with existing physical networks. These IP addresses are used for the Pod network, and if you need to access the Pods from an external network, configure load balancers and routers to manage the traffic.
- 11 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**, which allows for 510 ($2^{(32 - 23)} - 2$) pod IPs addresses. If you are required to provide access to nodes from an external network, configure load balancers and routers to manage the traffic.

- 12 The IP address pool to use for service IP addresses. You can enter only one IP address pool. If you need to access the services from an external network, configure load balancers and routers to
- 13 You must set the platform to **none**. You cannot provide additional platform configuration variables for IBM Z infrastructure.
- 14 The pull secret that you obtained from the [Pull Secret](#) page on the Red Hat OpenShift Cluster Manager site. 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.
- 15 The public portion of the default SSH key for the **core** user in Red Hat Enterprise Linux CoreOS (RHCOS).




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.

1.1.8. Creating the Kubernetes manifest and Ignition config files

Because you must modify some cluster definition files and manually start the cluster machines, you must generate the Kubernetes manifest and Ignition config files that the cluster needs to make its machines.



IMPORTANT

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

Prerequisites

- Obtain the OpenShift Container Platform installation program.
- Create the **install-config.yaml** installation configuration file.

Procedure

1. Generate the Kubernetes manifests for the cluster:

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


```
WARNING There are no compute nodes specified. The cluster will not fully initialize without compute nodes.
```

```
INFO Consuming "Install Config" from target directory
```

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

Because you create your own compute machines later in the installation process, you can safely ignore this warning.

2. Modify the **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. This step might not be required in a future minor version of OpenShift Container Platform.

3. Obtain the Ignition config files:

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

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

The following files are generated in the directory:

```
.
├── auth
│   ├── kubeadmin-password
│   └── kubeconfig
├── bootstrap.ign
├── master.ign
├── metadata.json
└── worker.ign
```

1.1.9. Creating Red Hat Enterprise Linux CoreOS (RHCOS) machines


Before you install a cluster on IBM Z infrastructure that you provision, you must install RHCOS on z/VM guest virtual machines for the cluster to use. Complete the following steps to create the machines.

Prerequisites

- An FTP server running on your provisioning machine that is accessible to the machines you create.

Procedure

1. Log in to Linux on your provisioning machine.
2. Download the Red Hat Enterprise Linux CoreOS installation files from the [RHCOS image mirror](#).



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.

Download the following files:

- The initramfs: **rhcos-<version>-installer-initramfs.img**
- The kernel: **rhcos-<version>-installer-kernel**
- The operating system image for the disk on which you want to install RHCOS. This type can differ by virtual machine:
rhcos<version>-dasd.s390x.raw.gz for DASD
rhcos<version>-metal.s390x.raw.gz for FCP

3. Create parameter files. The following parameters are specific for a particular virtual machine:

- For **coreos.inst.install_dev=**, specify **dasda** for a DASD installation, or **sda** for FCP. Note that FCP requires **zfcplib.allow_lun_scan=0**.
- For **rd.dasd=**, specifies the DASD where RHCOS is to be installed.
- **rd.zfcplib=<adapter>,<wwpn>,<lun>** specifies the FCP disk to install RHCOS on.
- For **ip=**, specify the following seven entries:
 - i. The IP address for the machine.
 - ii. An empty string.
 - iii. The gateway.
 - iv. The netmask.
 - v. The machine host and domain name in the form **hostname.domainname**. Omit this value to let RHCOS decide set it.
 - vi. The network interface name. Omit this value to let RHCOS decide set it.
 - vii. If you use static IP addresses, an empty string.
- For **coreos.inst.ignition_url=**, specify the Ignition file for the machine role. Use **bootstrap.ign**, **master.ign**, or **worker.ign**.
- All other parameters can stay as they are.
 Example parameter file, **bootstrap-0.parm**, for the bootstrap machine:

```
rd.neednet=1 coreos.inst=yes coreos.inst.install_dev=dasda coreos.inst.image_url=ftp://
cl1.provide.example.com:8080/assets/rhcos-42.80.20191105.0-metal-dasd.raw.gz
coreos.inst.ignition_url=ftp://cl1.provide.example.com:8080/ignition-bootstrap-0
ip=172.18.78.2::172.18.78.1:255.255.255.0::none nameserver=172.18.78.1
```

```
rd.znet=qeth,0.0.bdf0,0.0.bdf1,0.0.bdf2,layer2=1,portno=0 zfcplib.allow_lun_scan=0
cio_ignore=all,
!condev rd.dasd=0.0.3490
```

4. Transfer the initramfs, kernel, parameter files, and RHCOS images to z/VM, for example with FTP. For details about how to transfer the files with FTP and boot from the virtual reader, see [Installing under Z/VM](#).
5. Punch the files to the virtual reader of the z/VM guest virtual machine that is to become your bootstrap node.
See [PUNCH](#) in the IBM Knowledge Center.

TIP

You can use the CP PUNCH command or, if you use Linux, the **vmur** command to transfer files between two z/VM guest virtual machines.

6. Log in to CMS on the bootstrap machine.
7. IPL the bootstrap machine from the reader:

```
$ ipl c
```

See [IPL](#) in the IBM Knowledge Center.

8. Repeat this procedure for the other machines in the cluster.

1.1.10. Creating the cluster

To create the OpenShift Container Platform cluster, you wait for the bootstrap process to complete on the machines that you provisioned by using the Ignition config files that you generated with the installation program.

Prerequisites

- Create the required infrastructure for the cluster.
- You obtained the installation program and generated the Ignition config files for your cluster.
- You used the Ignition config files to create RHCOS machines for your cluster.
- Your machines have direct internet access.

Procedure

1. Monitor the bootstrap process:

```
$ ./openshift-install --dir=<installation_directory> wait-for bootstrap-complete \ 1
--log-level=info 2
INFO Waiting up to 30m0s for the Kubernetes API at https://api.test.example.com:6443...
INFO API v1.14.6+c4799753c up
INFO Waiting up to 30m0s for the bootstrap-complete event...
```


- 1 For **<installation_directory>**, specify the path to the directory that you stored the

installation files in.

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

The command succeeds when the Kubernetes API server signals that it has been bootstrapped on the control plane machines.

2. After bootstrap process is complete, remove the bootstrap machine from the load balancer.



IMPORTANT

You must remove the bootstrap machine from the load balancer at this point. You can also remove or reformat the machine itself.

1.1.11. 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
```

1.1.12. Approving the CSRs for your machines

When you add machines to a cluster, two pending certificates signing request (CSRs) are generated for each machine that you added. You must confirm that these CSRs are approved or, if necessary, approve them yourself.

Prerequisites

- You added machines to your cluster.

Procedure

1. Confirm that the cluster recognizes the machines:

```
$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
master-0.cl1mstr0.example.com	Ready	master	20h	v1.14.6+888f9c630
master-1.cl1mstr1.example.com	Ready	master	20h	v1.14.6+888f9c630
master-2.cl1mstr2.example.com	Ready	master	20h	v1.14.6+888f9c630
worker-0.cl1wrk00.example.com	Ready	worker	20h	v1.14.6+888f9c630
worker-1.cl1wrk01.example.com	Ready	worker	20h	v1.14.6+888f9c630

The output lists all of the machines that you created.

2. Review the pending certificate signing requests (CSRs) and ensure that you see a client and server request with **Pending** or **Approved** status for each machine that you added to the cluster:


```
$ oc get csr
```

NAME	AGE	REQUESTOR	CONDITION
csr-8b2br	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending 1
csr-8vnps	15m	system:serviceaccount:openshift-machine-config-operator:node-bootstrapper	Pending
csr-bfd72	5m26s	system:node:ip-10-0-50-126.us-east-2.compute.internal	Pending 2
csr-c57lv	5m26s	system:node:ip-10-0-95-157.us-east-2.compute.internal	Pending
...			

- 1** A client request CSR.
- 2** A server request CSR.

In this example, two machines are joining the cluster. You might see more approved CSRs in the list.

3. If the CSRs were not approved, after all of the pending CSRs for the machines you added are in **Pending** status, approve the CSRs for your cluster machines:



NOTE

Because the CSRs rotate automatically, approve your CSRs within an hour of adding the machines to the cluster. If you do not approve them within an hour, the certificates will rotate, and more than two certificates will be present for each node. You must approve all of these certificates. After you approve the initial CSRs, the subsequent node client CSRs are automatically approved by the cluster **kube-controller-manager**. You must implement a method of automatically approving the kubelet serving certificate requests.

- To approve them individually, run the following command for each valid CSR:

```
$ oc adm certificate approve <csr_name> 1
```

1 **<csr_name>** is the name of a CSR from the list of current CSRs.

- To approve all pending CSRs, run the following command:

```
$ oc get csr -o go-template='{{range .items}}{{if not .status}}{{.metadata.name}}{\n"}\n{{end}}' | xargs oc adm certificate approve
```

1.1.13. Initial Operator configuration

After the control plane initializes, you must immediately configure some Operators so that they all become available.

Prerequisites

- Your control plane has initialized.

Procedure

- Watch the cluster components come online:

```
$ watch -n5 oc get clusteroperators
```

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.2.0	True	False	False	69s
cloud-credential	4.2.0	True	False	False	12m
cluster-autoscaler	4.2.0	True	False	False	11m
console	4.2.0	True	False	False	46s
dns	4.2.0	True	False	False	11m
image-registry	4.2.0	False	True	False	5m26s
ingress	4.2.0	True	False	False	5m36s
kube-apiserver	4.2.0	True	False	False	8m53s
kube-controller-manager	4.2.0	True	False	False	7m24s
kube-scheduler	4.2.0	True	False	False	12m
machine-api	4.2.0	True	False	False	12m
machine-config	4.2.0	True	False	False	7m36s
marketplace	4.2.0	True	False	False	7m54m
monitoring	4.2.0	True	False	False	7h54s
network	4.2.0	True	False	False	5m9s
node-tuning	4.2.0	True	False	False	11m
openshift-apiserver	4.2.0	True	False	False	11m
openshift-controller-manager	4.2.0	True	False	False	5m943s
openshift-samples	4.2.0	True	False	False	3m55s
operator-lifecycle-manager	4.2.0	True	False	False	11m
operator-lifecycle-manager-catalog	4.2.0	True	False	False	11m
service-ca	4.2.0	True	False	False	11m
service-catalog-apiserver	4.2.0	True	False	False	5m26s
service-catalog-controller-manager	4.2.0	True	False	False	5m25s
storage	4.2.0	True	False	False	5m30s

- Configure the Operators that are not available.

1.1.13.1. Image registry storage configuration

If the **image-registry** Operator is not available, you must configure storage for it. Instructions for both configuring a PersistentVolume, which is required for production clusters, and for configuring an empty directory as the storage location, which is available for only non-production clusters, are shown.

1.1.13.1.1. Configuring registry storage for IBM Z

As a cluster administrator, following installation you must configure your registry to use storage.


Prerequisites

- Cluster administrator permissions.
- A cluster on IBM Z.
- Provision persistent storage for your cluster, such as NFS. To deploy a private image registry, your storage must provide ReadWriteMany access mode.
- Must have "100Gi" capacity.

Procedure

1. To configure your registry to use storage, change the **spec.storage.pvc** in the **configs.imageregistry/cluster** resource.
2. Verify you do not have a registry Pod:

```
$ oc get pod -n openshift-image-registry
```



NOTE

If the storage type is **emptyDIR**, the replica number cannot be greater than **1**. If the storage type is **NFS**, and you want to scale up the registry Pod by setting **replica>1** you must enable the **no_wdelay** mount option. For example:

```
# cat /etc/exports
/mnt/data *(rw,sync,no_wdelay,no_root_squash,insecure,fsid=0)
sh-4.3# exportfs -rv
exporting */mnt/data
```

3. Check the registry configuration:

```
$ oc edit configs.imageregistry.operator.openshift.io

storage:
  pvc:
    claim:
```

Leave the **claim** field blank to allow the automatic creation of an **image-registry-storage** PVC.

4. Check the **clusteroperator** status:

```
$ oc get clusteroperator image-registry
```

1.1.13.1.2. Configuring storage for the image registry in non-production clusters

You must configure storage for the image registry Operator. For non-production clusters, you can set the image registry to an empty directory. If you do so, all images are lost if you restart the registry.

Procedure

- To set the image registry storage to an empty directory:

```
$ oc patch configs.imageregistry.operator.openshift.io cluster --type merge --patch '{"spec": {"storage":{"emptyDir":{}}}'
```



WARNING

Configure this option for only non-production clusters.

If you run this command before the Image Registry Operator initializes its components, the **oc patch** command fails with the following error:

```
Error from server (NotFound): configs.imageregistry.operator.openshift.io "cluster" not found
```

Wait a few minutes and run the command again.

1.1.14. Completing installation on user-provisioned infrastructure

After you complete the Operator configuration, you can finish installing the cluster on infrastructure that you provide.

Prerequisites

- Your control plane has initialized.
- You have completed the initial Operator configuration.

Procedure

- Confirm that all the cluster components are online:

```
$ watch -n5 oc get clusteroperators
```

NAME	VERSION	AVAILABLE	PROGRESSING	DEGRADED	SINCE
authentication	4.2.0	True	False	False	10m
cloud-credential	4.2.0	True	False	False	22m
cluster-autoscaler	4.2.0	True	False	False	21m
console	4.2.0	True	False	False	10m
dns	4.2.0	True	False	False	21m
image-registry	4.2.0	True	False	False	16m
ingress	4.2.0	True	False	False	16m

kube-apiserver	4.2.0	True	False	False	19m
kube-controller-manager	4.2.0	True	False	False	18m
kube-scheduler	4.2.0	True	False	False	22m
machine-api	4.2.0	True	False	False	22m
machine-config	4.2.0	True	False	False	18m
marketplace	4.2.0	True	False	False	18m
monitoring	4.2.0	True	False	False	18m
network	4.2.0	True	False	False	16m
node-tuning	4.2.0	True	False	False	21m
openshift-apiserver	4.2.0	True	False	False	21m
openshift-controller-manager	4.2.0	True	False	False	17m
openshift-samples	4.2.0	True	False	False	14m
operator-lifecycle-manager	4.2.0	True	False	False	21m
operator-lifecycle-manager-catalog	4.2.0	True	False	False	21m
service-ca	4.2.0	True	False	False	21m
service-catalog-apiserver	4.2.0	True	False	False	16m
service-catalog-controller-manager	4.2.0	True	False	False	16m
storage	4.2.0	True	False	False	16m

When all of the cluster Operators are **AVAILABLE**, you can complete the installation.

2. Monitor for cluster completion:

```
$ ./openshift-install --dir=<installation_directory> wait-for install-complete 1
INFO Waiting up to 30m0s for the cluster to initialize...
```

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

The command succeeds when the Cluster Version Operator finishes deploying the OpenShift Container Platform cluster from Kubernetes API server.



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.

3. Confirm that the Kubernetes API server is communicating with the Pods.

- a. To view a list of all Pods, use the following command:

```
$ oc get pods --all-namespaces
```

NAMESPACE	NAME	READY	STATUS
RESTARTS	AGE		
openshift-apiserver-operator	openshift-apiserver-operator-85cb746d55-zqhs8	1/1	Running
1	9m		
openshift-apiserver	apiserver-67b9g	1/1	Running
3m			
openshift-apiserver	apiserver-ljcmx	1/1	Running
1m			
openshift-apiserver	apiserver-z25h4	1/1	Running
			0

```

2m
openshift-authentication-operator authentication-operator-69d5d8bf84-vh2n8 1/1
Running 0 5m
...

```

- b. View the logs for a Pod that is listed in the output of the previous command by using the following command:

```
$ oc logs <pod_name> -n <namespace> 1
```

- 1 Specify the Pod name and namespace, as shown in the output of the previous command.

If the Pod logs display, the Kubernetes API server can communicate with the cluster machines.

1.1.15. Collecting debugging information

You can gather debugging information that might help you to troubleshoot and debug certain issues with an OpenShift Container Platform installation on IBM Z.

Prerequisites

- The **oc** CLI tool installed.

Procedure

1. Log in to the cluster:

```
$ oc login
```

2. On the node you want to gather hardware information about, start a debugging container:

```
$ oc debug node/<nodename>
```

3. Change to the **/host** file system and start **toolbox**:

```
$ chroot /host
$ toolbox
```

4. Collect the **dbginfo** data:

```
$ dbginfo.sh
```

5. You can then retrieve the data, for example, using **scp**.

Additional resources

- See also [How to generate SOSREPORT within OpenShift4 nodes without SSH](#) .

Next steps

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