



# OpenShift Container Platform 4.13

## Installation overview

Overview content for installing OpenShift Container Platform



## OpenShift Container Platform 4.13 Installation overview

---

Overview content for installing OpenShift Container Platform

## Legal Notice

Copyright © 2025 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

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

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

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

Node.js<sup>®</sup> 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<sup>®</sup> 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 an overview on how to install OpenShift Container Platform.

## Table of Contents

<b>CHAPTER 1. OPENSIFT CONTAINER PLATFORM INSTALLATION OVERVIEW .....</b>	<b>3</b>
1.1. ABOUT OPENSIFT CONTAINER PLATFORM INSTALLATION	3
1.1.1. About the installation program	3
1.1.2. About Red Hat Enterprise Linux CoreOS (RHCOS)	4
1.1.3. Glossary of common terms for OpenShift Container Platform installing	4
1.1.4. Installation process	6
The installation process with the Assisted Installer	7
The installation process with Agent-based infrastructure	7
The installation process with installer-provisioned infrastructure	7
The installation process with user-provisioned infrastructure	8
Installation process details	8
1.1.5. Verifying node state after installation	10
Installation scope	11
1.1.6. OpenShift Local overview	11
1.2. SUPPORTED PLATFORMS FOR OPENSIFT CONTAINER PLATFORM CLUSTERS	12
<b>CHAPTER 2. SELECTING A CLUSTER INSTALLATION METHOD AND PREPARING IT FOR USERS .....</b>	<b>14</b>
2.1. SELECTING A CLUSTER INSTALLATION TYPE	14
2.1.1. Do you want to install and manage an OpenShift Container Platform cluster yourself?	14
2.1.2. Have you used OpenShift Container Platform 3 and want to use OpenShift Container Platform 4?	15
2.1.3. Do you want to use existing components in your cluster?	15
2.1.4. Do you need extra security for your cluster?	16
2.2. PREPARING YOUR CLUSTER FOR USERS AFTER INSTALLATION	16
2.3. PREPARING YOUR CLUSTER FOR WORKLOADS	17
2.4. SUPPORTED INSTALLATION METHODS FOR DIFFERENT PLATFORMS	17
<b>CHAPTER 3. CLUSTER CAPABILITIES .....</b>	<b>23</b>
3.1. SELECTING CLUSTER CAPABILITIES	23
3.2. OPTIONAL CLUSTER CAPABILITIES IN OPENSIFT CONTAINER PLATFORM 4.13	24
3.2.1. Bare-metal capability	25
3.2.2. Cluster storage capability	25
Notes	25
3.2.3. Console capability	26
3.2.4. CSI snapshot controller capability	26
3.2.5. Insights capability	26
Notes	26
3.2.6. Marketplace capability	26
3.2.7. Node Tuning capability	27
3.2.8. OpenShift samples capability	27
3.3. ADDITIONAL RESOURCES	27
<b>CHAPTER 4. SUPPORT FOR FIPS CRYPTOGRAPHY .....</b>	<b>29</b>
4.1. FIPS VALIDATION IN OPENSIFT CONTAINER PLATFORM	29
4.2. FIPS SUPPORT IN COMPONENTS THAT THE CLUSTER USES	29
4.2.1. etcd	30
4.2.2. Storage	30
4.2.3. Runtimes	30
4.3. INSTALLING A CLUSTER IN FIPS MODE	30



# CHAPTER 1. OPENSIFT CONTAINER PLATFORM INSTALLATION OVERVIEW

## 1.1. ABOUT OPENSIFT CONTAINER PLATFORM INSTALLATION

The OpenShift Container Platform installation program offers four methods for deploying a cluster which are detailed in the following list:

- **Interactive:** You can deploy a cluster with the web-based [Assisted Installer](#). This is an ideal approach for clusters with networks connected to the internet. The Assisted Installer is the easiest way to install OpenShift Container Platform, it provides smart defaults, and it performs pre-flight validations before installing the cluster. It also provides a RESTful API for automation and advanced configuration scenarios.
- **Local Agent-based:** You can deploy a cluster locally with the Agent-based Installer for disconnected environments or restricted networks. It provides many of the benefits of the Assisted Installer, but you must download and configure the [Agent-based Installer](#) first. Configuration is done with a command-line interface. This approach is ideal for disconnected environments.
- **Automated:** You can deploy a cluster on installer-provisioned infrastructure. The installation program uses each cluster host's baseboard management controller (BMC) for provisioning. You can deploy clusters in connected or disconnected environments.
- **Full control:** You can deploy a cluster on infrastructure that you prepare and maintain, which provides maximum customizability. You can deploy clusters in connected or disconnected environments.

Each method deploys a cluster with the following characteristics:


- Highly available infrastructure with no single points of failure, which is available by default.
- Administrators can control what updates are applied and when.

### 1.1.1. About the installation program

You can use the installation program to deploy each type of cluster. The installation program generates the main assets, such as Ignition config files for the bootstrap, control plane, and compute machines. You can start an OpenShift Container Platform cluster with these three machine configurations, provided you correctly configured the infrastructure.

The OpenShift Container Platform installation program uses a set of targets and dependencies to manage cluster installations. The installation program has a set of targets that it must achieve, and each target has a set of dependencies. Because each target is only concerned with its own dependencies, the installation program can act to achieve multiple targets in parallel with the ultimate target being a running cluster. The installation program recognizes and uses existing components instead of running commands to create them again because the program meets the dependencies.

Figure 1.1. OpenShift Container Platform installation targets and dependencies



### 1.1.2. About Red Hat Enterprise Linux CoreOS (RHCOS)

Post-installation, each cluster machine uses Red Hat Enterprise Linux CoreOS (RHCOS) as the operating system. RHCOS is the immutable container host version of Red Hat Enterprise Linux (RHEL) and features a RHEL kernel with SELinux enabled by default. RHCOS includes the **kubelet**, which is the Kubernetes node agent, and the CRI-O container runtime, which is optimized for Kubernetes.

Every control plane machine in an OpenShift Container Platform 4.13 cluster must use RHCOS, which includes a critical first-boot provisioning tool called Ignition. This tool enables the cluster to configure the machines. Operating system updates are delivered as a bootable container image, using **OSTree** as a backend, that is deployed across the cluster by the Machine Config Operator. Actual operating system changes are made in-place on each machine as an atomic operation by using **rpm-ostree**. Together, these technologies enable OpenShift Container Platform to manage the operating system like it manages any other application on the cluster, by in-place upgrades that keep the entire platform up to date. These in-place updates can reduce the burden on operations teams.

If you use RHCOS as the operating system for all cluster machines, the cluster manages all aspects of its components and machines, including the operating system. Because of this, only the installation program and the Machine Config Operator can change machines. The installation program uses Ignition config files to set the exact state of each machine, and the Machine Config Operator completes more changes to the machines, such as the application of new certificates or keys, after installation.

### 1.1.3. Glossary of common terms for OpenShift Container Platform installing

The glossary defines common terms that relate to the installation content. Read the following list of terms to better understand the installation process.

#### Assisted Installer



An installer hosted at [console.redhat.com](https://console.redhat.com) that provides a web-based user interface or a RESTful API for creating a cluster configuration. The [Assisted Installer](#) generates a discovery image. Cluster machines boot with the discovery image, which installs RHCOS and an agent. Together, the Assisted Installer and agent provide preinstallation validation and installation for the cluster.

**Agent-based Installer**

An installer similar to the Assisted Installer, but you must download the [Agent-based Installer](#) first. The Agent-based Installer is ideal for disconnected environments.

**Bootstrap node**

A temporary machine that runs a minimal Kubernetes configuration required to deploy the OpenShift Container Platform control plane.

**Control plane**

A container orchestration layer that exposes the API and interfaces to define, deploy, and manage the lifecycle of containers. Also known as control plane machines.

**Compute node**

Nodes that are responsible for executing workloads for cluster users. Also known as worker nodes.

**Disconnected installation**

In some situations, parts of a data center might not have access to the internet, even through proxy servers. You can still install the OpenShift Container Platform in these environments, but you must download the required software and images and make them available to the disconnected environment.

**The OpenShift Container Platform installation program**

A program that provisions the infrastructure and deploys a cluster.

**Installer-provisioned infrastructure**

The installation program deploys and configures the infrastructure that the cluster runs on.

**Ignition config files**

A file that the Ignition tool uses to configure Red Hat Enterprise Linux CoreOS (RHCOS) during operating system initialization. The installation program generates different Ignition configuration files to initialize bootstrap, control plane, and worker nodes.

**Kubernetes manifests**

Specifications of a Kubernetes API object in a JSON or YAML format. A configuration file can include deployments, config maps, secrets, daemonsets, and so on.

**Kubelet**

A primary node agent that runs on each node in the cluster to ensure that containers are running in a pod.

**Load balancers**

A load balancer serves as the single point of contact for clients. Load balancers for the API distribute incoming traffic across control plane nodes.

**Machine Config Operator**

An Operator that manages and applies configurations and updates of the base operating system and container runtime, including everything between the kernel and kubelet, for the nodes in the cluster.

**Operators**

The preferred method of packaging, deploying, and managing a Kubernetes application in an OpenShift Container Platform cluster. An operator takes human operational knowledge and encodes it into software that is easily packaged and shared with customers.

**User-provisioned infrastructure**

You can install OpenShift Container Platform on infrastructure that you provide. You can use the installation program to generate the assets required to provision the cluster infrastructure, create the cluster infrastructure, and then deploy the cluster to the infrastructure that you provided.

### 1.1.4. Installation process

Except for the Assisted Installer, when you install an OpenShift Container Platform cluster, you must download the installation program from the appropriate [Cluster Type](#) page on the OpenShift Cluster Manager Hybrid Cloud Console. This console manages:

- REST API for accounts.
- Registry tokens, which are the pull secrets that you use to obtain the required components.
- Cluster registration, which associates the cluster identity to your Red Hat account to facilitate the gathering of usage metrics.

In OpenShift Container Platform 4.13, the installation program is a Go binary file that performs a series of file transformations on a set of assets. The way you interact with the installation program differs depending on your installation type. Consider the following installation use cases:

- To deploy a cluster with the Assisted Installer, you must configure the cluster settings by using the [Assisted Installer](#). There is no installation program to download and configure. After you finish setting the cluster configuration, you download a discovery ISO and then boot cluster machines with that image. You can install clusters with the Assisted Installer on Nutanix, vSphere, and bare metal with full integration, and other platforms without integration. If you install on bare metal, you must provide all of the cluster infrastructure and resources, including the networking, load balancing, storage, and individual cluster machines.
- To deploy clusters with the Agent-based Installer, you can download the [Agent-based Installer](#) first. You can then configure the cluster and generate a discovery image. You boot cluster machines with the discovery image, which installs an agent that communicates with the installation program and handles the provisioning for you instead of you interacting with the installation program or setting up a provisioner machine yourself. You must provide all of the cluster infrastructure and resources, including the networking, load balancing, storage, and individual cluster machines. This approach is ideal for disconnected environments.
- For clusters with installer-provisioned infrastructure, you delegate the infrastructure bootstrapping and provisioning to the installation program instead of doing it yourself. The installation program creates all of the networking, machines, and operating systems that are required to support the cluster, except if you install on bare metal. If you install on bare metal, you must provide all of the cluster infrastructure and resources, including the bootstrap machine, networking, load balancing, storage, and individual cluster machines.
- If you provision and manage the infrastructure for your cluster, you must provide all of the cluster infrastructure and resources, including the bootstrap machine, networking, load balancing, storage, and individual cluster machines.

For the installation program, the program uses three sets of files during installation: an installation configuration file that is named **install-config.yaml**, Kubernetes manifests, and Ignition config files for your machine types.



## IMPORTANT

You can modify Kubernetes and the Ignition config files that control the underlying RHCOS operating system during installation. However, no validation is available to confirm the suitability of any modifications that you make to these objects. If you modify these objects, you might render your cluster non-functional. Because of this risk, modifying Kubernetes and Ignition config files is not supported unless you are following documented procedures or are instructed to do so by Red Hat support.

The installation configuration file is transformed into Kubernetes manifests, and then the manifests are wrapped into Ignition config files. The installation program uses these Ignition config files to create the cluster.

The installation configuration files are all pruned when you run the installation program, so be sure to back up all the configuration files that you want to use again.



## IMPORTANT

You cannot modify the parameters that you set during installation, but you can modify many cluster attributes after installation.

### The installation process with the Assisted Installer

Installation with the [Assisted Installer](#) involves creating a cluster configuration interactively by using the web-based user interface or the RESTful API. The Assisted Installer user interface prompts you for required values and provides reasonable default values for the remaining parameters, unless you change them in the user interface or with the API. The Assisted Installer generates a discovery image, which you download and use to boot the cluster machines. The image installs RHCOS and an agent, and the agent handles the provisioning for you. You can install OpenShift Container Platform with the Assisted Installer and full integration on Nutanix, vSphere, and bare metal. Additionally, you can install OpenShift Container Platform with the Assisted Installer on other platforms without integration.

OpenShift Container Platform manages all aspects of the cluster, including the operating system itself. Each machine boots with a configuration that references resources hosted in the cluster that it joins. This configuration allows the cluster to manage itself as updates are applied.

If possible, use the Assisted Installer feature to avoid having to download and configure the Agent-based Installer.

### The installation process with Agent-based infrastructure

Agent-based installation is similar to using the Assisted Installer, except that you must initially download and install the [Agent-based Installer](#). An Agent-based installation is useful when you want the convenience of the Assisted Installer, but you need to install a cluster in a disconnected environment.

If possible, use the Agent-based installation feature to avoid having to create a provisioner machine with a bootstrap VM, and then provision and maintain the cluster infrastructure.

### The installation process with installer-provisioned infrastructure

The default installation type uses installer-provisioned infrastructure. By default, the installation program acts as an installation wizard, prompting you for values that it cannot determine on its own and providing reasonable default values for the remaining parameters. You can also customize the installation process to support advanced infrastructure scenarios. The installation program provisions the underlying infrastructure for the cluster.

You can install either a standard cluster or a customized cluster. With a standard cluster, you provide minimum details that are required to install the cluster. With a customized cluster, you can specify more details about the platform, such as the number of machines that the control plane uses, the type of

virtual machine that the cluster deploys, or the CIDR range for the Kubernetes service network.

If possible, use this feature to avoid having to provision and maintain the cluster infrastructure. In all other environments, you use the installation program to generate the assets that you require to provision your cluster infrastructure.

With installer-provisioned infrastructure clusters, OpenShift Container Platform manages all aspects of the cluster, including the operating system itself. Each machine boots with a configuration that references resources hosted in the cluster that it joins. This configuration allows the cluster to manage itself as updates are applied.

### **The installation process with user-provisioned infrastructure**

You can also install OpenShift Container Platform on infrastructure that you provide. You use the installation program to generate the assets that you require to provision the cluster infrastructure, create the cluster infrastructure, and then deploy the cluster to the infrastructure that you provided.

If you do not use infrastructure that the installation program provisioned, you must manage and maintain the cluster resources yourself. The following list details some of these self-managed resources:


- The underlying infrastructure for the control plane and compute machines that make up the cluster
- Load balancers
- Cluster networking, including the DNS records and required subnets
- Storage for the cluster infrastructure and applications

If your cluster uses user-provisioned infrastructure, you have the option of adding RHEL compute machines to your cluster.

### **Installation process details**

When a cluster is provisioned, each machine in the cluster requires information about the cluster. OpenShift Container Platform uses a temporary bootstrap machine during initial configuration to provide the required information to the permanent control plane. The temporary bootstrap machine boots by using an Ignition config file that describes how to create the cluster. The bootstrap machine creates the control plane machines that make up the control plane. The control plane machines then create the compute machines, which are also known as worker machines. The following figure illustrates this process:

Figure 1.2. Creating the bootstrap, control plane, and compute machines



After the cluster machines initialize, the bootstrap machine is destroyed. All clusters use the bootstrap process to initialize the cluster, but if you provision the infrastructure for your cluster, you must complete many of the steps manually.

## IMPORTANT

- The Ignition config files that the installation program generates contain certificates that expire after 24 hours, which are then renewed at that time. If the cluster is shut down before renewing the certificates and the cluster is later restarted after the 24 hours have elapsed, the cluster automatically recovers the expired certificates. The exception is that you must manually approve the pending **node-bootstrap** certificate signing requests (CSRs) to recover kubelet certificates. See the documentation for *Recovering from expired control plane certificates* for more information.
- Consider using Ignition config files within 12 hours after they are generated, because the 24-hour certificate rotates from 16 to 22 hours after the cluster is installed. By using the Ignition config files within 12 hours, you can avoid installation failure if the certificate update runs during installation.

Bootstrapping a cluster involves the following steps:

1. The bootstrap machine boots and starts hosting the remote resources required for the control plane machines to boot. If you provision the infrastructure, this step requires manual intervention.
2. The bootstrap machine starts a single-node etcd cluster and a temporary Kubernetes control plane.
3. The control plane machines fetch the remote resources from the bootstrap machine and finish booting. If you provision the infrastructure, this step requires manual intervention.

4. The temporary control plane schedules the production control plane to the production control plane machines.
5. The Cluster Version Operator (CVO) comes online and installs the etcd Operator. The etcd Operator scales up etcd on all control plane nodes.
6. The temporary control plane shuts down and passes control to the production control plane.
7. The bootstrap machine injects OpenShift Container Platform components into the production control plane.
8. The installation program shuts down the bootstrap machine. If you provision the infrastructure, this step requires manual intervention.
9. The control plane sets up the compute nodes.
10. The control plane installs additional services in the form of a set of Operators.

The result of this bootstrapping process is a running OpenShift Container Platform cluster. The cluster then downloads and configures remaining components needed for the day-to-day operations, including the creation of compute machines in supported environments.

### 1.1.5. Verifying node state after installation

The OpenShift Container Platform installation completes when the following installation health checks are successful:

- The provisioner can access the OpenShift Container Platform web console.
- All control plane nodes are ready.
- All cluster Operators are available.



#### NOTE

After the installation completes, the specific cluster Operators responsible for the worker nodes continuously attempt to provision all worker nodes. Some time is required before all worker nodes report as **READY**. For installations on bare metal, wait a minimum of 60 minutes before troubleshooting a worker node. For installations on all other platforms, wait a minimum of 40 minutes before troubleshooting a worker node. A **DEGRADED** state for the cluster Operators responsible for the worker nodes depends on the Operators' own resources and not on the state of the nodes.

After your installation completes, you can continue to monitor the condition of the nodes in your cluster.

#### Prerequisites

- The installation program resolves successfully in the terminal.

#### Procedure

1. Show the status of all worker nodes:

```
$ oc get nodes
```

### Example output

NAME	STATUS	ROLES	AGE	VERSION
example-compute1.example.com	Ready	worker	13m	v1.21.6+bb8d50a
example-compute2.example.com	Ready	worker	13m	v1.21.6+bb8d50a
example-compute4.example.com	Ready	worker	14m	v1.21.6+bb8d50a
example-control1.example.com	Ready	master	52m	v1.21.6+bb8d50a
example-control2.example.com	Ready	master	55m	v1.21.6+bb8d50a
example-control3.example.com	Ready	master	55m	v1.21.6+bb8d50a

2. Show the phase of all worker machine nodes:

```
$ oc get machines -A
```

### Example output

NAMESPACE	NAME	PHASE	TYPE	REGION	ZONE	AGE
openshift-machine-api	example-zbbt6-master-0	Running			95m	
openshift-machine-api	example-zbbt6-master-1	Running			95m	
openshift-machine-api	example-zbbt6-master-2	Running			95m	
openshift-machine-api	example-zbbt6-worker-0-25bhp	Running			49m	
openshift-machine-api	example-zbbt6-worker-0-8b4c2	Running			49m	
openshift-machine-api	example-zbbt6-worker-0-jkbqt	Running			49m	
openshift-machine-api	example-zbbt6-worker-0-ql5b	Running			49m	

### Additional resources

- [Getting the BareMetalHost resource](#)
- [Following the progress of the installation](#)
- [Validating an installation](#)
- [Agent-based Installer](#)
- [Assisted Installer for OpenShift Container Platform](#)

### Installation scope

The scope of the OpenShift Container Platform installation program is intentionally narrow. It is designed for simplicity and ensured success. You can complete many more configuration tasks after installation completes.

### Additional resources

- See [Available cluster customizations](#) for details about OpenShift Container Platform configuration resources.

## 1.1.6. OpenShift Local overview

OpenShift Local supports rapid application development to get started building OpenShift Container Platform clusters. OpenShift Local is designed to run on a local computer to simplify setup and testing, and to emulate the cloud development environment locally with all of the tools needed to develop container-based applications.

Regardless of the programming language you use, OpenShift Local hosts your application and brings a minimal, preconfigured Red Hat OpenShift Container Platform cluster to your local PC without the need for a server-based infrastructure.

On a hosted environment, OpenShift Local can create microservices, convert them into images, and run them in Kubernetes-hosted containers directly on your laptop or desktop running Linux, macOS, or Windows 10 or later.


For more information about OpenShift Local, see [Red Hat OpenShift Local Overview](#).

## 1.2. SUPPORTED PLATFORMS FOR OPENSIFT CONTAINER PLATFORM CLUSTERS

In OpenShift Container Platform 4.13, you can install a cluster that uses installer-provisioned infrastructure on the following platforms:

- Alibaba Cloud
- Amazon Web Services (AWS)
- Bare metal
- Google Cloud Platform (GCP)
- IBM Cloud® VPC
- Microsoft Azure
- Microsoft Azure Stack Hub
- Nutanix
- Red Hat OpenStack Platform (RHOSP)
  - The latest OpenShift Container Platform release supports both the latest RHOSP long-life release and intermediate release. For complete RHOSP release compatibility, see the [OpenShift Container Platform on RHOSP support matrix](#).
- VMware Cloud (VMC) on AWS
- VMware vSphere

For these clusters, all machines, including the computer that you run the installation process on, must have direct internet access to pull images for platform containers and provide telemetry data to Red Hat.



### IMPORTANT

After installation, the following changes are not supported:

- Mixing cloud provider platforms.
- Mixing cloud provider components. For example, using a persistent storage framework from a another platform on the platform where you installed the cluster.



In OpenShift Container Platform 4.13, you can install a cluster that uses user-provisioned infrastructure on the following platforms:

- AWS
- Azure
- Azure Stack Hub
- Bare metal
- GCP
- IBM Power
- IBM Z or IBM® LinuxONE
- RHOSP
  - The latest OpenShift Container Platform release supports both the latest RHOSP long-life release and intermediate release. For complete RHOSP release compatibility, see the [OpenShift Container Platform on RHOSP support matrix](#).
- VMware Cloud on AWS
- VMware vSphere

Depending on the supported cases for the platform, you can perform installations on user-provisioned infrastructure, so that you can run machines with full internet access, place your cluster behind a proxy, or perform a disconnected installation.

In a disconnected installation, you can download the images that are required to install a cluster, place them in a mirror registry, and use that data to install your cluster. While you require internet access to pull images for platform containers, with a disconnected installation on vSphere or bare metal infrastructure, your cluster machines do not require direct internet access.

The [OpenShift Container Platform 4.x Tested Integrations](#) page contains details about integration testing for different platforms.

### Additional resources

- See [Supported installation methods for different platforms](#) for more information about the types of installations that are available for each supported platform.
- See [Selecting a cluster installation method and preparing it for users](#) for information about choosing an installation method and preparing the required resources.

## CHAPTER 2. SELECTING A CLUSTER INSTALLATION METHOD AND PREPARING IT FOR USERS

Before you install OpenShift Container Platform, decide what kind of installation process to follow and verify that you have all of the required resources to prepare the cluster for users.

### 2.1. SELECTING A CLUSTER INSTALLATION TYPE

Before you install an OpenShift Container Platform cluster, you need to select the best installation instructions to follow. Think about your answers to the following questions to select the best option.

#### 2.1.1. Do you want to install and manage an OpenShift Container Platform cluster yourself?

If you want to install and manage OpenShift Container Platform yourself, you can install it on the following platforms:

- Alibaba Cloud
- Amazon Web Services (AWS) on 64-bit x86 instances
- Amazon Web Services (AWS) on 64-bit ARM instances
- Microsoft Azure on 64-bit x86 instances
- Microsoft Azure on 64-bit ARM instances
- Microsoft Azure Stack Hub
- Google Cloud
- Red Hat OpenStack Platform (RHOSP)
- Red Hat Virtualization (RHV)
- IBM Cloud VPC
- IBM Z or IBM® LinuxONE
- IBM Z or IBM® LinuxONE for Red Hat Enterprise Linux (RHEL) KVM
- IBM Power
- IBM Power Virtual Server
- Nutanix
- VMware vSphere
- VMware Cloud (VMC) on AWS
- Bare metal or other platform agnostic infrastructure

You can deploy an OpenShift Container Platform 4 cluster to both on-premise hardware and to cloud hosting services, but all of the machines in a cluster must be in the same data center or cloud hosting service.

If you want to use OpenShift Container Platform but do not want to manage the cluster yourself, you have several managed service options. If you want a cluster that is fully managed by Red Hat, you can use [OpenShift Dedicated](#) or [OpenShift Online](#). You can also use OpenShift as a managed service on Azure, AWS, IBM Cloud VPC, or Google Cloud. For more information about managed services, see the [OpenShift Products](#) page. If you install an OpenShift Container Platform cluster with a cloud virtual machine as a virtual bare metal, the corresponding cloud-based storage is not supported.

### 2.1.2. Have you used OpenShift Container Platform 3 and want to use OpenShift Container Platform 4?

If you used OpenShift Container Platform 3 and want to try OpenShift Container Platform 4, you need to understand how different OpenShift Container Platform 4 is. OpenShift Container Platform 4 weaves the Operators that package, deploy, and manage Kubernetes applications and the operating system that the platform runs on, Red Hat Enterprise Linux CoreOS (RHCOS), together seamlessly. Instead of deploying machines and configuring their operating systems so that you can install OpenShift Container Platform on them, the RHCOS operating system is an integral part of the OpenShift Container Platform cluster. Deploying the operating system for the cluster machines is part of the installation process for OpenShift Container Platform. See [Differences between OpenShift Container Platform 3 and 4](#).

Because you need to provision machines as part of the OpenShift Container Platform cluster installation process, you cannot upgrade an OpenShift Container Platform 3 cluster to OpenShift Container Platform 4. Instead, you must create a new OpenShift Container Platform 4 cluster and migrate your OpenShift Container Platform 3 workloads to them. For more information about migrating, see [Migrating from OpenShift Container Platform 3 to 4 overview](#). Because you must migrate to OpenShift Container Platform 4, you can use any type of production cluster installation process to create your new cluster.

### 2.1.3. Do you want to use existing components in your cluster?

Because the operating system is integral to OpenShift Container Platform, it is easier to let the installation program for OpenShift Container Platform stand up all of the infrastructure. These are called *installer provisioned infrastructure* installations. In this type of installation, you can provide some existing infrastructure to the cluster, but the installation program deploys all of the machines that your cluster initially needs.

You can deploy an installer-provisioned infrastructure cluster without specifying any customizations to the cluster or its underlying machines to [Alibaba Cloud](#), [AWS](#), [Azure](#), [Azure Stack Hub](#), [Google Cloud](#), [Nutanix](#), or [VMC on AWS](#). These installation methods are the fastest way to deploy a production-capable OpenShift Container Platform cluster.

If you need to perform basic configuration for your installer-provisioned infrastructure cluster, such as the instance type for the cluster machines, you can customize an installation for [Alibaba Cloud](#), [AWS](#), [Azure](#), [Google Cloud](#), [Nutanix](#), or [VMC on AWS](#).

For installer-provisioned infrastructure installations, you can use an existing [VPC in AWS](#), [vNet in Azure](#), or [VPC in Google Cloud](#). You can also reuse part of your networking infrastructure so that your cluster in [AWS](#), [Azure](#), [Google Cloud](#), or [VMC on AWS](#) can coexist with existing IP address allocations in your environment and integrate with existing MTU and VXLAN configurations. If you have existing accounts and credentials on these clouds, you can re-use them, but you might need to modify the accounts to have the required permissions to install OpenShift Container Platform clusters on them.

You can use the installer-provisioned infrastructure method to create appropriate machine instances on your hardware for [RHOSP](#), [RHOSP with Kuryr](#), [RHV](#), [vSphere](#), and [bare metal](#). Additionally, for [vSphere](#), [VMC on AWS](#), you can also customize additional network parameters during installation.

For some installer-provisioned infrastructure installations, for example on the VMware vSphere and bare metal platforms, the external traffic that reaches the ingress virtual IP (VIP) is not balanced between the default **IngressController** replicas. For vSphere and bare metal installer-provisioned infrastructure installations where exceeding the baseline **IngressController** router performance is expected, you must configure an external load balancer. Configuring an external load balancer achieves the performance of multiple **IngressController** replicas. For more information about the baseline **IngressController** performance, see [Baseline Ingress Controller \(router\) performance](#). For more information about configuring an external load balancer, see [Configuring an external load balancer](#).

If you want to reuse extensive cloud infrastructure, you can complete a *user-provisioned infrastructure* installation. With these installations, you manually deploy the machines that your cluster requires during the installation process. If you perform a user-provisioned infrastructure installation on [AWS](#), [Azure](#), [Azure Stack Hub](#), [Google Cloud](#), or [VMC on AWS](#), you can use the provided templates to help you stand up all of the required components. You can also reuse a shared [VPC on Google Cloud](#). Otherwise, you can use the [provider-agnostic](#) installation method to deploy a cluster into other clouds.

You can also complete a user-provisioned infrastructure installation on your existing hardware. If you use [RHOSP](#), [RHV](#), [IBM Z or IBM® LinuxONE](#), [IBM Z and IBM® LinuxONE with RHEL KVM](#), [IBM Power](#), or [vSphere](#), use the specific installation instructions to deploy your cluster. If you use other supported hardware, follow the [bare metal installation](#) procedure. For some of these platforms, such as [RHOSP](#), [vSphere](#), [VMC on AWS](#), and [bare metal](#), you can also customize additional network parameters during installation.

#### 2.1.4. Do you need extra security for your cluster?

If you use a user-provisioned installation method, you can configure a proxy for your cluster. The instructions are included in each installation procedure.

If you want to prevent your cluster on a public cloud from exposing endpoints externally, you can deploy a private cluster with installer-provisioned infrastructure on [AWS](#), [Azure](#), or [Google Cloud](#).

If you need to install your cluster that has limited access to the internet, such as a disconnected or restricted network cluster, you can [mirror the installation packages](#) and install the cluster from them. Follow detailed instructions for user provisioned infrastructure installations into restricted networks for [AWS](#), [Google Cloud](#), [IBM Z or IBM® LinuxONE](#), [IBM Z or IBM® LinuxONE with RHEL KVM](#), [IBM Power](#), [vSphere](#), [VMC on AWS](#), or [bare metal](#). You can also install a cluster into a restricted network using installer-provisioned infrastructure by following detailed instructions for [AWS](#), [Google Cloud](#), [Nutanix](#), [VMC on AWS](#), [RHOSP](#), [RHV](#), and [vSphere](#).

If you need to deploy your cluster to an [AWS GovCloud region](#), [AWS China region](#), or [Azure government region](#), you can configure those custom regions during an installer-provisioned infrastructure installation.

## 2.2. PREPARING YOUR CLUSTER FOR USERS AFTER INSTALLATION

Some configuration is not required to install the cluster but recommended before your users access the cluster. You can customize the cluster itself by [customizing](#) the Operators that make up your cluster and integrate you cluster with other required systems, such as an identity provider.

For a production cluster, you must configure the following integrations:

- [Persistent storage](#)

- [An identity provider](#)
- [Monitoring core OpenShift Container Platform components](#)

## 2.3. PREPARING YOUR CLUSTER FOR WORKLOADS

Depending on your workload needs, you might need to take extra steps before you begin deploying applications. For example, after you prepare infrastructure to support your application [build strategy](#), you might need to make provisions for [low-latency](#) workloads or to [protect sensitive workloads](#). You can also configure [monitoring](#) for application workloads. If you plan to run [Windows workloads](#), you must enable [hybrid networking with OVN-Kubernetes](#) during the installation process; hybrid networking cannot be enabled after your cluster is installed.

## 2.4. SUPPORTED INSTALLATION METHODS FOR DIFFERENT PLATFORMS

You can perform different types of installations on different platforms.



### NOTE

Not all installation options are supported for all platforms, as shown in the following tables. A checkmark indicates that the option is supported and links to the relevant section.

Table 2.1. Installer-provisioned infrastructure options

	Al ib a b a	A W S (6 4- bi t x8 6)	A W S (6 4- bi t A R M )	A zu re (6 4- bi t A R M )	A zu re (6 4- bi t A R M )	A zu re St ac k H u b	G o o g l e C l o u d	N ut an ix	R H O S P	R H V	B ar e m et al (6 4- bi t x8 6)	B ar e m et al (6 4- bi t A R M )	vS p h er e	V M C	IB M Cl o u d V P C	IB M Z	IB M P o w er	IB M P o w er Vi rt u al S er v er
D ef a ul t	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓			
C us to m	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓	✓			✓

	Al ib a b a	A W S (6 4- bit x8 6)	A W S (6 4- bit A R M )	A zu re (6 4- bit x8 6)	A zu re (6 4- bit A R M )	A zu re St ack H u b	G o o g l e C l o u d	N u t a n ix	R H O S P	R H V	B a r e m e t a l (6 4- bit x8 6)	B a r e m e t a l (6 4- bit A R M )	vS p h e r e	V M C	IB M C l o u d V P C	IB M Z	IB M P o w er	IB M P o w er V i r t u a l S e r v e r
N e t w o r k c u s t o m i z a t i o n	✓	✓	✓	✓	✓	✓	✓		✓				✓	✓	✓			
R e s t r i c t e d n e t w o r k		✓	✓				✓	✓	✓	✓	✓	✓	✓	✓				✓
P r i v a t e c l u s t e r s		✓	✓	✓	✓		✓								✓			✓

	Al ib a b a	A W S (6 4- bit x8 6)	A W S (6 4- bit A R M )	A zu re (6 4- bit x8 6)	A zu re (6 4- bit A R M )	A zu re St ac k H u b	G o o g l e C l o u d	N u t a n ix	R H O S P	R H V	B a r e m et al (6 4- bit x8 6)	B a r e m et al (6 4- bit A R M )	vS p h e r e	V M C	IB M C l o u d V P C	IB M Z	IB M P o w er	IB M P o w er V i r t u a l S e r v e r
E x i s t i n g v i r t u a l p r i v a t e n e t w o r k s		✓	✓	✓	✓		✓								✓			✓
G o v e r n m e n t r e g i o n s		✓		✓														

	Al ib a b a	A W S (6 4- bit x8 6)	A W S (6 4- bit A R M )	A zu re (6 4- bit x8 6)	A zu re (6 4- bit A R M )	A zu re St ac k H u b	G o o g l e C l o u d	N ut an ix	R H O S P	R H V	B ar e m et al (6 4- bit x8 6)	B ar e m et al (6 4- bit A R M )	vS p h er e	V M C	IB M Cl o u d V P C	IB M Z	IB M P o w er	IB M P o w er Vi rt u al S er v er
S e c r e t r e g i o n s		✓																
C h i n a r e g i o n s		✓																

Table 2.2. User-provisioned infrastructure options



	Al ib a b a	A W S (6 4 - bi t x 8 6)	A W S (6 4 - bi t A R M )	A z ur e (6 4 - bi t x 8 6)	A z ur e (6 4 - bi t A R M )	A z ur e St a ck H u b	G o o g l e C l o u d	N u t a n i x	R H O S P	R H V	B a r e m e t a l ( 6 4 - bi t x 8 6 )	B a r e m e t a l ( 6 4 - bi t A R M )	v S p h e r e	V M C	I B M C l o u d V P C	I B M Z	I B M Z w i t h R H E L K V M	I B M P o w e r	Pl a t f o r m a g n o s t i c
C u s t o m		✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓		✓	✓	✓	✓
N e t w o r k c u s t o m i z a t i o n									✓		✓	✓	✓	✓					
R e s t r i c t e d n e t w o r k		✓	✓				✓			✓	✓	✓	✓	✓		✓	✓	✓	

Platform	IBM Power	IBM Z	IBM Cloud VPC	VSphere	Barometal (64-bit ARM)	Barometal (64-bit x86)	RHV	RHOSP	Nutanix	Google Cloud	Azure Stack Hub	Azure (64-bit ARM)	Azure (64-bit x86)	AWS (64-bit ARM)	AWS (64-bit x86)	Alibaba
Shared VPC Hosted outside of cluster project										✓						

## CHAPTER 3. CLUSTER CAPABILITIES

Cluster administrators can use cluster capabilities to enable or disable optional components prior to installation. Cluster administrators can enable cluster capabilities at anytime after installation.



### NOTE

Cluster administrators cannot disable a cluster capability after it is enabled.

### 3.1. SELECTING CLUSTER CAPABILITIES

You can select cluster capabilities by following one of the installation methods that include customizing your cluster, such as "Installing a cluster on AWS with customizations" or "Installing a cluster on GCP with customizations".

During a customized installation, you create an **install-config.yaml** file that contains the configuration parameters for your cluster.



### NOTE

If you customize your cluster by enabling or disabling specific cluster capabilities, you are responsible for manually maintaining your **install-config.yaml** file. New OpenShift Container Platform updates might declare new capability handles for existing components, or introduce new components altogether. Users who customize their **install-config.yaml** file should consider periodically updating their **install-config.yaml** file as OpenShift Container Platform is updated.

You can use the following configuration parameters to select cluster capabilities:

capabilities:

```
baselineCapabilitySet: v4.11 1
additionalEnabledCapabilities: 2
- CSISnapshot
- Console
- Storage
```

- <sup>1</sup> Defines a baseline set of capabilities to install. Valid values are **None**, **vCurrent** and **v4.x**. If you select **None**, all optional capabilities will be disabled. The default value is **vCurrent**, which enables all optional capabilities.



### NOTE

**v4.x** refers to any value up to and including the current cluster version. For example, valid values for a OpenShift Container Platform 4.12 cluster are **v4.11** and **v4.12**.

- <sup>2</sup> Defines a list of capabilities to explicitly enable. These will be enabled in addition to the capabilities specified in **baselineCapabilitySet**.

**NOTE**

In this example, the default capability is set to **v4.11**. The **additionalEnabledCapabilities** field enables additional capabilities over the default **v4.11** capability set.

The following table describes the **baselineCapabilitySet** values.

**Table 3.1. Cluster capabilities **baselineCapabilitySet** values description**

Value	Description
<b>vCurrent</b>	Specify this option when you want to automatically add new, default capabilities that are introduced in new releases.
<b>v4.11</b>	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.11. By specifying <b>v4.11</b> , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.11 are <b>baremetal</b> , <b>marketplace</b> , and <b>openshift-samples</b> .
<b>v4.12</b>	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.12. By specifying <b>v4.12</b> , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.12 are <b>baremetal</b> , <b>marketplace</b> , <b>openshift-samples</b> , <b>Console</b> , <b>Insights</b> , <b>Storage</b> and <b>CSISnapshot</b> .
<b>v4.13</b>	Specify this option when you want to enable the default capabilities for OpenShift Container Platform 4.13. By specifying <b>v4.13</b> , capabilities that are introduced in newer versions of OpenShift Container Platform are not enabled. The default capabilities in OpenShift Container Platform 4.13 are <b>baremetal</b> , <b>marketplace</b> , <b>openshift-samples</b> , <b>Console</b> , <b>Insights</b> , <b>Storage</b> , <b>CSISnapshot</b> and <b>NodeTuning</b> .
<b>None</b>	Specify when the other sets are too large, and you do not need any capabilities or want to fine-tune via <b>additionalEnabledCapabilities</b> .

**Additional resources**

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

## 3.2. OPTIONAL CLUSTER CAPABILITIES IN OPENSIFT CONTAINER PLATFORM 4.13

Currently, cluster Operators provide the features for these optional capabilities. The following summarizes the features provided by each capability and what functionality you lose if it is disabled.

#### Additional resources

- [Cluster Operators reference](#)


### 3.2.1. Bare-metal capability

The Cluster Baremetal Operator provides the features for the **baremetal** capability.

The Cluster Baremetal Operator (CBO) deploys all the components necessary to take a bare-metal server to a fully functioning worker node ready to run OpenShift Container Platform compute nodes. The CBO ensures that the metal3 deployment, which consists of the Bare Metal Operator (BMO) and Ironic containers, runs on one of the control plane nodes within the OpenShift Container Platform cluster. The CBO also listens for OpenShift Container Platform updates to resources that it watches and takes appropriate action.

The bare-metal capability is required for deployments using installer-provisioned infrastructure. Disabling the bare-metal capability can result in unexpected problems with these deployments.

It is recommended that cluster administrators only disable the bare-metal capability during installations with user-provisioned infrastructure that do not have any **BareMetalHost** resources in the cluster.



#### IMPORTANT

If the bare-metal capability is disabled, the cluster cannot provision or manage bare-metal nodes. Only disable the capability if there are no **BareMetalHost** resources in your deployment.


#### Additional resources

- [Deploying installer-provisioned clusters on bare metal](#)
- [Preparing for bare metal cluster installation](#)
- [Bare metal configuration](#)

### 3.2.2. Cluster storage capability

The Cluster Storage Operator provides the features for the **Storage** capability.

The Cluster Storage Operator sets OpenShift Container Platform cluster-wide storage defaults. It ensures a default **storageclass** exists for OpenShift Container Platform clusters. It also installs Container Storage Interface (CSI) drivers which enable your cluster to use various storage backends.



#### IMPORTANT

If the cluster storage capability is disabled, the cluster will not have a default **storageclass** or any CSI drivers. Users with administrator privileges can create a default **storageclass** and manually install CSI drivers if the cluster storage capability is disabled.

#### Notes

- The storage class that the Operator creates can be made non-default by editing its annotation, but this storage class cannot be deleted as long as the Operator runs.

### 3.2.3. Console capability

The Console Operator provides the features for the **Console** capability.

The Console Operator installs and maintains the OpenShift Container Platform web console on a cluster. The Console Operator is installed by default and automatically maintains a console.

#### Additional resources

- [Web console overview](#)

### 3.2.4. CSI snapshot controller capability

The Cluster CSI Snapshot Controller Operator provides the features for the **CSISnapshot** capability.

The Cluster CSI Snapshot Controller Operator installs and maintains the CSI Snapshot Controller. The CSI Snapshot Controller is responsible for watching the **VolumeSnapshot** CRD objects and manages the creation and deletion lifecycle of volume snapshots.

#### Additional resources

- [CSI volume snapshots](#)

### 3.2.5. Insights capability

The Insights Operator provides the features for the **Insights** capability.

The Insights Operator gathers OpenShift Container Platform configuration data and sends it to Red Hat. The data is used to produce proactive insights recommendations about potential issues that a cluster might be exposed to. These insights are communicated to cluster administrators through the Insights advisor service on [console.redhat.com](https://console.redhat.com).

#### Notes

Insights Operator complements OpenShift Container Platform Telemetry.

#### Additional resources

- [Using Insights Operator](#)

### 3.2.6. Marketplace capability

The Marketplace Operator provides the features for the **marketplace** capability.

The Marketplace Operator simplifies the process for bringing off-cluster Operators to your cluster by using a set of default Operator Lifecycle Manager (OLM) catalogs on the cluster. When the Marketplace Operator is installed, it creates the **openshift-marketplace** namespace. OLM ensures catalog sources installed in the **openshift-marketplace** namespace are available for all namespaces on the cluster.

If you disable the **marketplace** capability, the Marketplace Operator does not create the **openshift-marketplace** namespace. Catalog sources can still be configured and managed on the cluster manually,

but OLM depends on the **openshift-marketplace** namespace in order to make catalogs available to all namespaces on the cluster. Users with elevated permissions to create namespaces prefixed with **openshift-**, such as system or cluster administrators, can manually create the **openshift-marketplace** namespace.

If you enable the **marketplace** capability, you can enable and disable individual catalogs by configuring the Marketplace Operator.

#### Additional resources

- [Red Hat-provided Operator catalogs](#)

### 3.2.7. Node Tuning capability

The Node Tuning Operator provides features for the **NodeTuning** capability.

The Node Tuning Operator helps you manage node-level tuning by orchestrating the TuneD daemon and achieves low latency performance by using the Performance Profile controller. The majority of high-performance applications require some level of kernel tuning. The Node Tuning Operator provides a unified management interface to users of node-level sysctls and more flexibility to add custom tuning specified by user needs.

If you disable the NodeTuning capability, some default tuning settings will not be applied to the control-plane nodes. This might limit the scalability and performance of large clusters with over 900 nodes or 900 routes.

#### Additional resources

- [Using the Node Tuning Operator](#)

### 3.2.8. OpenShift samples capability

The Cluster Samples Operator provides the features for the **openshift-samples** capability.

The Cluster Samples Operator manages the sample image streams and templates stored in the **openshift** namespace.

On initial start up, the Operator creates the default samples configuration resource to initiate the creation of the image streams and templates. The configuration object is a cluster scoped object with the key **cluster** and type **configs.samples**.

The image streams are the Red Hat Enterprise Linux CoreOS (RHCOS)-based OpenShift Container Platform image streams pointing to images on **registry.redhat.io**. Similarly, the templates are those categorized as OpenShift Container Platform templates.

If you disable the samples capability, users cannot access the image streams, samples, and templates it provides. Depending on your deployment, you might want to disable this component if you do not need it.

#### Additional resources

- [Configuring the Cluster Samples Operator](#)

## 3.3. ADDITIONAL RESOURCES

- [Enabling cluster capabilities after installation](#)



## CHAPTER 4. SUPPORT FOR FIPS CRYPTOGRAPHY

You can install an OpenShift Container Platform cluster that uses FIPS Validated / Modules in Process cryptographic libraries on **x86\_64**, **ppc64le**, and **s390x** architectures.

For the Red Hat Enterprise Linux CoreOS (RHCOS) machines in your cluster, this change is applied when the machines are deployed based on the status of an option in the **install-config.yaml** file, which governs the cluster options that a user can change during cluster deployment. With Red Hat Enterprise Linux (RHEL) machines, you must enable FIPS mode when you install the operating system on the machines that you plan to use as worker machines. These configuration methods ensure that your cluster meet the requirements of a FIPS compliance audit: only FIPS Validated / Modules in Process cryptography packages are enabled before the initial system boot.

Because FIPS must be enabled before the operating system that your cluster uses boots for the first time, you cannot enable FIPS after you deploy a cluster.

### 4.1. FIPS VALIDATION IN OPENSIFT CONTAINER PLATFORM

OpenShift Container Platform uses certain FIPS Validated / Modules in Process modules within RHEL and RHCOS for the operating system components that it uses. See [RHEL8 core crypto components](#). For example, when users SSH into OpenShift Container Platform clusters and containers, those connections are properly encrypted.

OpenShift Container Platform components are written in Go and built with Red Hat's go lang compiler. When you enable FIPS mode for your cluster, all OpenShift Container Platform components that require cryptographic signing call RHEL and RHCOS cryptographic libraries.

**Table 4.1. FIPS mode attributes and limitations in OpenShift Container Platform 4.13**

Attributes	Limitations
FIPS support in RHEL 8 and RHCOS operating systems.	The FIPS implementation does not offer a single function that both computes hash functions and validates the keys that are based on that hash. This limitation will continue to be evaluated and improved in future OpenShift Container Platform releases.
FIPS support in CRI-O runtimes.	
FIPS support in OpenShift Container Platform services.	
FIPS Validated / Modules in Process cryptographic module and algorithms that are obtained from RHEL 8 and RHCOS binaries and images.	
Use of FIPS compatible go lang compiler.	TLS FIPS support is not complete but is planned for future OpenShift Container Platform releases.
FIPS support across multiple architectures.	FIPS is currently only supported on OpenShift Container Platform deployments using <b>x86_64</b> , <b>ppc64le</b> , and <b>s390x</b> architectures.

### 4.2. FIPS SUPPORT IN COMPONENTS THAT THE CLUSTER USES

Although the OpenShift Container Platform cluster itself uses FIPS Validated / Modules in Process modules, ensure that the systems that support your OpenShift Container Platform cluster use FIPS Validated / Modules in Process modules for cryptography.

#### 4.2.1. etcd

To ensure that the secrets that are stored in etcd use FIPS Validated / Modules in Process encryption, boot the node in FIPS mode. After you install the cluster in FIPS mode, you can [encrypt the etcd data](#) by using the FIPS-approved **aes cbc** cryptographic algorithm.

#### 4.2.2. Storage

For local storage, use RHEL-provided disk encryption or Container Native Storage that uses RHEL-provided disk encryption. By storing all data in volumes that use RHEL-provided disk encryption and enabling FIPS mode for your cluster, both data at rest and data in motion, or network data, are protected by FIPS Validated / Modules in Process encryption. You can configure your cluster to encrypt the root filesystem of each node, as described in [Customizing nodes](#).


#### 4.2.3. Runtimes

To ensure that containers know that they are running on a host that is using FIPS Validated / Modules in Process cryptography modules, use CRI-O to manage your runtimes. CRI-O supports FIPS mode, in that it configures the containers to know that they are running in FIPS mode.

### 4.3. INSTALLING A CLUSTER IN FIPS MODE

To install a cluster in FIPS mode, follow the instructions to install a customized cluster on your preferred infrastructure. Ensure that you set **fips: true** in the **install-config.yaml** file before you deploy your cluster.

- [Amazon Web Services](#)
- [Alibaba Cloud](#)
- [Microsoft Azure](#)
- [Bare metal](#)
- [Google Cloud](#)
- [IBM Cloud VPC](#)
- [IBM Power](#)
- [IBM Z and IBM® LinuxONE](#)
- [IBM Z and IBM® LinuxONE with RHEL KVM](#)
- [Red Hat OpenStack Platform \(RHOSP\)](#)
- [VMware vSphere](#)



#### NOTE

If you are using Azure File storage, you cannot enable FIPS mode.

To apply **AES CBC** encryption to your etcd data store, follow the [Encrypting etcd data](#) process after you install your cluster.

If you add RHEL nodes to your cluster, ensure that you enable FIPS mode on the machines before their initial boot. See [Adding RHEL compute machines to an OpenShift Container Platform cluster](#) and [Enabling FIPS Mode](#) in the RHEL 8 documentation.