



# Red Hat OpenShift Container Storage 4.6

## Planning your deployment

Important considerations when deploying RHOCS 4.6



## Red Hat OpenShift Container Storage 4.6 Planning your deployment

---

Important considerations when deploying RHOCS 4.6

## Legal Notice

Copyright © 2022 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

Read this document for important considerations when planning your Red Hat OpenShift Container Storage deployment.

## Table of Contents

<b>CHAPTER 1. INTRODUCTION TO RED HAT OPENSIFT CONTAINER STORAGE</b> .....	<b>3</b>
<b>CHAPTER 2. ARCHITECTURE OF OPENSIFT CONTAINER STORAGE</b> .....	<b>4</b>
2.1. ABOUT OPERATORS	4
2.2. STORAGE CLUSTER DEPLOYMENT APPROACHES	5
2.2.1. Internal approach	5
2.2.2. External approach	6
2.3. NODE TYPES	6
<b>CHAPTER 3. INTERNAL STORAGE SERVICES</b> .....	<b>8</b>
<b>CHAPTER 4. EXTERNAL STORAGE SERVICES</b> .....	<b>9</b>
<b>CHAPTER 5. SECURITY CONSIDERATIONS</b> .....	<b>10</b>
5.1. FIPS-140-2	10
5.2. PROXY ENVIRONMENT	10
5.3. DATA ENCRYPTION OPTIONS	10
<b>CHAPTER 6. SUBSCRIPTIONS</b> .....	<b>12</b>
6.1. SUBSCRIPTION OFFERINGS	12
6.2. DISASTER RECOVERY SUBSCRIPTIONS	12
6.3. CORES VERSUS VCPUS AND HYPERTHREADING	12
6.3.1. Cores versus vCPUs and simultaneous multithreading (SMT) for IBM Power Systems	12
6.4. SPLITTING CORES	13
6.4.1. Shared Processor Pools for IBM Power Systems	13
6.5. SUBSCRIPTION REQUIREMENTS	13
<b>CHAPTER 7. INFRASTRUCTURE REQUIREMENTS</b> .....	<b>14</b>
7.1. PLATFORM REQUIREMENTS	14
7.1.1. Amazon EC2	14
7.1.2. Bare Metal	14
7.1.3. VMware vSphere	14
7.1.4. Microsoft Azure	15
7.1.5. Google Cloud [Technology Preview]	15
7.1.6. Red Hat Virtualization Platform [Technology Preview]	15
7.1.7. Red Hat OpenStack Platform [Technology Preview]	15
7.1.8. IBM Power Systems	15
7.1.9. IBM Z and LinuxONE	15
7.2. EXTERNAL MODE REQUIREMENT	15
7.2.1. Red Hat Ceph Storage	15
7.3. RESOURCE REQUIREMENTS	16
7.3.1. Minimum deployment resource requirements [Technology Preview]	16
7.3.2. Compact deployment resource requirements [Technology Preview]	17
7.4. POD PLACEMENT RULES	18
7.5. STORAGE DEVICE REQUIREMENTS	18
7.5.1. Dynamic storage devices	18
7.5.2. Local storage devices	19
7.5.3. Capacity planning	19
<b>CHAPTER 8. DISCONNECTED ENVIRONMENT</b> .....	<b>20</b>
<b>CHAPTER 9. NEXT STEPS</b> .....	<b>21</b>



# CHAPTER 1. INTRODUCTION TO RED HAT OPENSIFT CONTAINER STORAGE

Red Hat OpenShift Container Storage is a highly integrated collection of cloud storage and data services for Red Hat OpenShift Container Platform. It is available as part of the Red Hat OpenShift Container Platform Service Catalog, packaged as an operator to facilitate simple deployment and management.

Red Hat OpenShift Container Storage services are primarily made available to applications by way of storage classes that represent the following components:

- Block storage devices, catering primarily to database workloads. Prime examples include Red Hat OpenShift Container Platform logging and monitoring, and PostgreSQL.
- Shared and distributed file system, catering primarily to software development, messaging, and data aggregation workloads. Examples include Jenkins build sources and artifacts, Wordpress uploaded content, Red Hat OpenShift Container Platform registry, and messaging using JBoss AMQ.
- Multicloud object storage, featuring a lightweight S3 API endpoint that can abstract the storage and retrieval of data from multiple cloud object stores.
- On premises object storage, featuring a robust S3 API endpoint that scales to tens of petabytes and billions of objects, primarily targeting data intensive applications. Examples include the storage and access of row, columnar, and semi-structured data with applications like Spark, Presto, Red Hat AMQ Streams (Kafka), and even machine learning frameworks like TensorFlow and Pytorch.



## NOTE

OpenShift Container Storage 4.6 on IBM Power Systems supports only block and file storage and not the object storage.

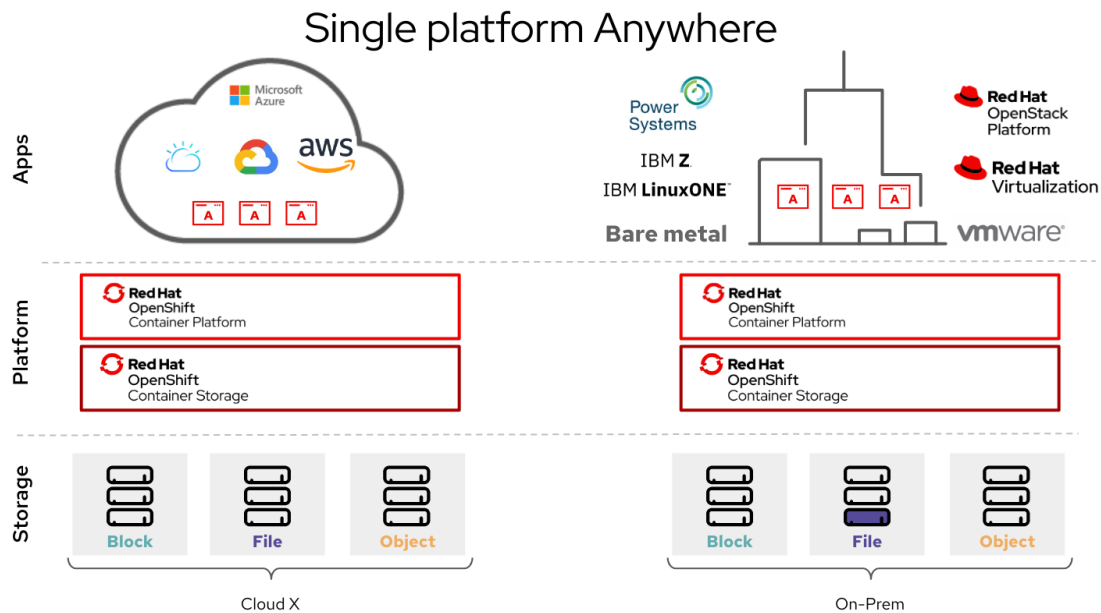
Red Hat OpenShift Container Storage version 4.x integrates a collection of software projects, including:

- Ceph, providing block storage, a shared and distributed file system, and on-premises object storage
- Ceph CSI, to manage provisioning and lifecycle of persistent volumes and claims
- NooBaa, providing a Multicloud Object Gateway
- OpenShift Container Storage, Rook-Ceph, and NooBaa operators to initialize and manage OpenShift Container Storage services.

## CHAPTER 2. ARCHITECTURE OF OPENSIFT CONTAINER STORAGE

Red Hat OpenShift Container Storage provides services for, and can run internally from Red Hat OpenShift Container Platform.

### Red Hat OpenShift Container Storage architecture



Red Hat OpenShift Container Storage supports deployment into Red Hat OpenShift Container Platform clusters deployed on Installer Provisioned Infrastructure or User Provisioned Infrastructure. For details about these two approaches, see [OpenShift Container Platform - Installation process](#). To know more about interoperability of components for the Red Hat OpenShift Container Storage and Red Hat OpenShift Container Platform, see the [interoperability matrix](#).



#### NOTE

For IBM Power Systems refer [OpenShift Container Platform - Installation process](#).

For information about the architecture and lifecycle of OpenShift Container Platform, see [OpenShift Container Platform architecture](#).

## 2.1. ABOUT OPERATORS

Red Hat OpenShift Container Storage comprises three main operators, which codify administrative tasks and custom resources so that task and resource characteristics can be easily automated. Administrators define the desired end state of the cluster, and the OpenShift Container Storage operators ensure the cluster is either in that state, or approaching that state, with minimal administrator intervention.

### OpenShift Container Storage operator

A meta-operator that codifies and enforces the recommendations and requirements of a supported Red Hat OpenShift Container Storage deployment by drawing on other operators in specific, tested ways. This operator provides the storage cluster resource that wraps resources provided by the Rook-Ceph and NooBaa operators.



## Rook-Ceph operator

This operator automates the packaging, deployment, management, upgrading, and scaling of persistent storage and file, block, and object services. It creates block and file storage classes for all environments, and creates an object storage class and services object bucket claims made against it in on-premises environments.

Additionally, for internal mode clusters, it provides the Ceph cluster resource, which manages the deployments and services representing the following:

- Object storage daemons (OSDs)
- Monitors (MONs)
- Manager (MGR)
- Metadata servers (MDS)
- Object gateways (RGW) on-premises only

## NooBaa operator

This operator automates the packaging, deployment, management, upgrading, and scaling of the Multicloud Object Gateway object service. It creates an object storage class and services object bucket claims made against it.

Additionally, it provides the NooBaa cluster resource, which manages the deployments and services for NooBaa core, database, and endpoint.

## 2.2. STORAGE CLUSTER DEPLOYMENT APPROACHES

Flexibility is a core tenet of Red Hat OpenShift Container Storage, as evidenced by its growing list of operating modalities. This section provides you with information that will help you to select the most appropriate approach for your environments. OpenShift Container Storage can be deployed either entirely within OpenShift Container Platform (Internal approach) or to make available the services from a cluster running outside of OpenShift Container Platform (External approach).

### 2.2.1. Internal approach

Deployment of Red Hat OpenShift Container Storage entirely within Red Hat OpenShift Container Platform has all the benefits of operator based deployment and management. Internal-attached device approach in the graphical user interface can be used to deploy Red Hat OpenShift Container Storage in internal mode using the local storage operator and local storage devices.

There are two different deployment modalities available when Red Hat OpenShift Container Storage is running entirely within Red Hat OpenShift Container Platform:

- Simple
- Optimized

#### Simple deployment

Red Hat OpenShift Container Storage services run co-resident with applications, managed by operators in Red Hat OpenShift Container Platform.

A simple deployment is best for situations where

- Storage requirements are not clear
- OpenShift Container Storage services will run co-resident with applications
- Creating a node instance of a specific size is difficult (bare metal)

In order for Red Hat OpenShift Container Storage to run co-resident with applications, they must have local storage devices, or portable storage devices attached to them dynamically, like EBS volumes on EC2, or vSphere Virtual Volumes on VMware, or SAN volumes dynamically provisioned by PowerVC.

### Optimized deployment

OpenShift Container Storage services run on dedicated infrastructure nodes managed by Red Hat OpenShift Container Platform.

An optimized approach is best for situations when:

- Storage requirements are clear
- OpenShift Container Storage services run on dedicated infrastructure nodes
- Creating a node instance of a specific size is easy (Cloud, Virtualized environment, etc.)

### 2.2.2. External approach

Red Hat OpenShift Container Storage exposes the Red Hat Ceph Storage services running outside of the OpenShift Container Platform cluster as storage classes.

The external approach is best used when:

- Storage requirements are significant (600+ storage devices)
- Multiple OpenShift Container Platform clusters need to consume storage services from a common external cluster.
- Another team (SRE, Storage, etc.) needs to manage the external cluster providing storage services. Possibly pre-existing.



#### NOTE

External approach is not applicable for IBM Power Systems, IBM Z and LinuxONE architecture.

## 2.3. NODE TYPES

Nodes run the container runtime, as well as services, to ensure that containers are running, and maintain network communication and separation between pods. In OpenShift Container Storage, there are three types of nodes.

Table 2.1. Types of nodes

Node Type	Description
-----------	-------------

Node Type	Description
Master	<p>These nodes run processes that expose the Kubernetes API, watch and schedule newly created pods, maintain node health and quantity, and control interaction with underlying cloud providers.</p>
Infrastructure (Infra)	<p>Infra nodes run cluster level infrastructure services such as logging, metrics, registry, and routing. These are optional in OpenShift Container Platform clusters. It is recommended to use infra nodes for OpenShift Container Storage in virtualized and cloud environments.</p> <p>To create Infra nodes, you can provision new nodes labeled as <b>infra</b>. See <a href="#">How to use dedicated worker nodes for Red Hat OpenShift Container Storage?</a></p>
Worker	<p>Worker nodes are also known as application nodes since they run applications.</p> <p>When OpenShift Container Storage is deployed in internal mode, a minimal cluster of 3 worker nodes is required, where the nodes are recommended to be spread across three different racks, or availability zones, to ensure availability. In order for OpenShift Container Storage to run on worker nodes, they must either have local storage devices, or portable storage devices attached to them dynamically.</p> <p>When it is deployed in external mode, it runs on multiple nodes to allow rescheduling by K8S on available nodes in case of a failure.</p> <p>Examples of portable storage devices are EBS volumes on EC2, or vSphere Virtual Volumes on VMware, or SAN volumes dynamically provisioned by PowerVC.</p>



## NOTE

Nodes that run only storage workloads require a subscription for Red Hat OpenShift Container Storage. Nodes that run other workloads in addition to storage workloads require both Red Hat OpenShift Container Storage and Red Hat OpenShift Container Platform subscriptions. See [Chapter 6, Subscriptions](#) for more information.

## CHAPTER 3. INTERNAL STORAGE SERVICES

Red Hat OpenShift Container Storage service is available for consumption internally to the Red Hat OpenShift Container Platform running on the following infrastructure:

- Amazon Web Services
- Bare metal
- VMware vSphere
- Microsoft Azure
- Google Cloud [Technology Preview]
- Red Hat Virtualization 4.4.x or higher (IPI) [Technology Preview]
- Red Hat OpenStack 13 or higher (IPI) [Technology Preview]
- IBM Power Systems
- IBM Z and LinuxONE

Ease of deployment and management are the highlights of running OpenShift Container Storage services internally on OpenShift Container Platform. Creation of an internal cluster resource will result in the internal provisioning of the OpenShift Container Storage base services, and make additional storage classes available to applications.

## CHAPTER 4. EXTERNAL STORAGE SERVICES

Red Hat OpenShift Container Storage can make services from an external Red Hat Ceph Storage cluster available for consumption through OpenShift Container Platform clusters running on the following platforms:

- VMware vSphere
- Bare Metal
- Red Hat OpenStack platform (Technology preview)

The OpenShift Container Storage operators create and manage services to satisfy persistent volume and object bucket claims against external services. External cluster can serve Block, File and Object storage classes for applications running on OpenShift Container Platform. External clusters are not deployed or managed by operators.

## CHAPTER 5. SECURITY CONSIDERATIONS

### 5.1. FIPS-140-2

The Federal Information Processing Standard Publication 140-2 (FIPS-140-2) is a standard defining a set of security requirements for the use of cryptographic modules. This standard is mandated by law for US government agencies and contractors and is also referenced in other international and industry specific standards.

Red Hat OpenShift Container Storage is now using FIPS validated cryptographic modules as delivered by Red Hat Enterprise Linux OS/CoreOS (RHCOS).

The cryptography modules are currently being processed by Cryptographic Module Validation Program (CMVP) and their state can be seen at [Modules in Process List](#). For more up-to-date information, see the [knowledge base article](#).



#### NOTE


FIPS mode must be enabled on the OpenShift Container Platform, prior to installing OpenShift Container Storage. OpenShift Container Platform must run on RHCOS nodes, as OpenShift Container Storage deployment on RHEL 7 is not supported for this feature. FIPS is not supported on OpenShift Container Storage 4.6 on IBM Power Systems.

For more information, see [installing a cluster in FIPS mode](#) and [support for FIPS cryptography](#).

### 5.2. PROXY ENVIRONMENT

A proxy environment is a production environment that denies direct access to the internet and provides an available HTTP or HTTPS proxy instead. Red Hat OpenShift Container Platform is configured to use a proxy by modifying the proxy object for existing clusters or by configuring the proxy settings in the `install-config.yaml` file for new clusters.

Red Hat supports deployment of OpenShift Container Storage versions 4.5 and higher in proxy environments when OpenShift Container Platform has been configured according to [configuring the cluster-wide proxy](#).



#### NOTE

Proxy environment is not supported on OpenShift Container Storage 4.6 on IBM Power Systems.

### 5.3. DATA ENCRYPTION OPTIONS

Encryption lets you encode and obscure your data to make it impossible to understand if it is stolen. Red Hat OpenShift Container Storage 4.6 provides support for at-rest encryption of all disks in the storage cluster, meaning that your data is encrypted when it is written to disk, and decrypted when it is read from the disk.

OpenShift Container Storage 4.6 uses Linux Unified Key System (LUKS) version 2 based encryption with a key size of 512 bits and the **aes-xts-plain64** cipher. Each device has a different encryption key, which is stored as a Kubernetes secret.

You can enable or disable encryption for your whole cluster during cluster deployment. It is disabled by default. Working with encrypted data incurs only a very small penalty to performance.

Data encryption is only supported for new clusters deployed using OpenShift Container Storage 4.6. It is not supported on existing clusters that are upgraded to version 4.6.

## CHAPTER 6. SUBSCRIPTIONS

### 6.1. SUBSCRIPTION OFFERINGS

Red Hat OpenShift Container Storage subscription is based on “core-pairs,” similar to Red Hat OpenShift Container Platform. The Red Hat OpenShift Container Storage 2-core subscription is based on the number of logical cores on the CPUs in the system where OpenShift Container Platform runs.

As with OpenShift Container Platform:

- OpenShift Container Storage subscriptions are stackable to cover larger hosts.
- Cores can be distributed across as many virtual machines (VMs) as needed. For example, ten 2-core subscriptions will provide 20 cores and in case of IBM Power Systems a 2-core subscription at SMT level of 8 will provide 2 cores or 16 vCPUs that can be used across any number of VMs.
- OpenShift Container Storage subscriptions are available with Premium or Standard support.

### 6.2. DISASTER RECOVERY SUBSCRIPTIONS

Red Hat OpenShift Container Storage does not offer disaster recovery (DR), cold backup, or other subscription types. Any system with OpenShift Container Storage installed, powered-on or powered-off, running workload or not, requires an active subscription.

### 6.3. CORES VERSUS VCPUS AND HYPERTHREADING

Making a determination about whether or not a particular system consumes one or more cores is currently dependent on whether or not that system has hyperthreading available. Hyperthreading is only a feature of Intel CPUs. Visit the Red Hat Customer Portal to determine whether a particular system supports hyperthreading.

For systems where hyperthreading is enabled and where one hyperthread equates to one visible system core, the [calculation of cores](#) is a ratio of 2 cores to 4 vCPUs. Therefore, a 2-core subscription covers 4 vCPUs in a hyperthreaded system. A large virtual machine (VM) might have 8 vCPUs, equating to 4 subscription cores. As subscriptions come in 2-core units, you will need two 2-core subscriptions to cover these 4 cores or 8 vCPUs.

Where hyperthreading is not enabled, and where each visible system core correlates directly to an underlying physical core, the calculation of cores is a ratio of 2 cores to 2 vCPUs.

#### 6.3.1. Cores versus vCPUs and simultaneous multithreading (SMT) for IBM Power Systems

Making a determination about whether or not a particular system consumes one or more cores is currently dependent on the level of simultaneous multithreading configured (SMT). IBM Power systems provide simultaneous multithreading levels of 1, 2, 4 or 8.

For systems where SMT is configured the calculation of cores depends on the SMT level. Therefore, a 2-core subscription 2 vCPU on SMT level of 1, 4 vCPUs on SMT level of 2, 8 vCPUs on SMT level of 4 and 16 vCPUs on SMT level of 8. A large virtual machine (VM) might have 16 vCPUs, which at a SMT level 8 will be equivalent of 2 subscription cores. As subscriptions come in 2-core units, you will need 1 2-core subscription to cover these 2 cores or 16 vCPUs.



## 6.4. SPLITTING CORES

Systems that require an odd number of cores need to consume a full 2-core subscription. For example, a system that is calculated to require only 1 core will end up consuming a full 2-core subscription once it is registered and subscribed.

When a single virtual machine (VM) with 2 vCPUs uses hyperthreading resulting in 1 calculated vCPU, a full 2-core subscription is required; a single 2-core subscription may not be split across two VMs with 2 vCPUs using hyperthreading. See section [Cores versus vCPUs and hyperthreading](#) for more information.

It is recommended that virtual instances be sized so that they require an even number of cores.

### 6.4.1. Shared Processor Pools for IBM Power Systems

IBM Power Systems have a notion of shared processor pools. The processors in a shared processor pool can be shared across the nodes in the cluster. The aggregate compute capacity required for a OpenShift Container Storage should be a multiple of core-pairs.

## 6.5. SUBSCRIPTION REQUIREMENTS

OpenShift Container Storage components can run on either OpenShift Container Platform worker or infrastructure nodes, for which either Red Hat CoreOS (RHCOS) or Red Hat Enterprise Linux (RHEL) 7 can be used as the host operating system. When worker nodes are used for OpenShift Container Storage components, those nodes are required to have subscriptions for both OpenShift Container Platform and OpenShift Container Storage. When infrastructure nodes are used, those nodes are only required to have OpenShift Container Storage subscriptions. Labels are used to indicate whether a node should be considered a worker or infrastructure node, see [How to use dedicated worker nodes for Red Hat OpenShift Container Storage](#) in the [Managing and Allocating Storage Resources](#) guide.

## CHAPTER 7. INFRASTRUCTURE REQUIREMENTS

### 7.1. PLATFORM REQUIREMENTS

Red Hat OpenShift Container Storage can be combined with an OpenShift Container Platform release that is one minor release behind or ahead of the OpenShift Container Storage version.

OpenShift Container Storage 4.6 can run on:

- OpenShift Container Platform 4.5 (one version behind)
- OpenShift Container Platform 4.6 (same version)
- OpenShift Container Platform 4.7 (one version ahead)



#### NOTE

For OpenShift Container Storage 4.6 on IBM Power Systems, IBM Z and LinuxONE infrastructure, only OpenShift Container Platform 4.6 or later is supported.

For external cluster subscription requirements, see this [Red Hat Knowledgebase article](#).

For a complete list of supported platform versions, see the [Red Hat OpenShift Container Storage and Red Hat OpenShift Container Platform interoperability matrix](#).

#### 7.1.1. Amazon EC2

Supports internal Red Hat OpenShift Container Storage clusters only.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing either

- EBS storage via the aws-efs provisioner
- Instance storage via the Local Storage Operator

#### 7.1.2. Bare Metal

Supports internal clusters and consuming external clusters.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing local SSD (NVMe/SATA/SAS, SAN) via the Local Storage Operator.

#### 7.1.3. VMware vSphere

Supports internal clusters and consuming external clusters.

Recommended version: vSphere 6.7, Update 2

See [VMware vSphere infrastructure requirements](#) for details.

Additionally, an Internal cluster must both meet [storage device requirements](#) and have a storage class providing either

- vSAN or VMFS datastore via the vsphere-volume provisioner
- VMDK, RDM, or DirectPath storage devices via the Local Storage Operator.

#### 7.1.4. Microsoft Azure

Supports internal Red Hat Openshift Container Storage clusters only.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing

- Azure disk via the azure-disk provisioner

#### 7.1.5. Google Cloud [Technology Preview]

Supports internal Red Hat Openshift Container Storage clusters only.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing

- GCE Persistent Disk via the gce-pd provisioner

#### 7.1.6. Red Hat Virtualization Platform [Technology Preview]

Supports internal Red Hat Openshift Container Storage clusters only.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing local SSD (NVMe/SATA/SAS, SAN) via the Local Storage Operator.

#### 7.1.7. Red Hat OpenStack Platform [Technology Preview]

Supports internal Red Hat Openshift Container Storage clusters and consuming external clusters.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing

- standard disk via the Cinder provisioner

#### 7.1.8. IBM Power Systems

Supports internal Red Hat Openshift Container Storage clusters only.

An Internal cluster must both meet [storage device requirements](#) and have a storage class providing local SSD (NVMe/SATA/SAS, SAN) via the Local Storage Operator.

#### 7.1.9. IBM Z and LinuxONE

Supports internal Red Hat Openshift Container Storage clusters only.


An Internal cluster must both meet [storage device requirements](#) and have a storage class providing local SSD (NVMe/SATA/SAS, SAN) via the Local Storage Operator.

## 7.2. EXTERNAL MODE REQUIREMENT

### 7.2.1. Red Hat Ceph Storage

Red Hat Ceph Storage (RHCS) version 4.2z1 or later is required. For more information on versions supported, see this [knowledge base article on Red Hat Ceph Storage releases and corresponding Ceph package versions](#).

For instructions regarding how to install a RHCS 4 cluster, see [Installation guide](#).



#### NOTE

External mode is not applicable for IBM Power Systems architecture.

## 7.3. RESOURCE REQUIREMENTS

OpenShift Container Storage services consist of an initial set of base services, and can be extended with additional device sets. All of these OpenShift Container Storage services pods are scheduled by kubernetes on OpenShift Container Platform nodes according to [resource requirements](#). Expanding the cluster in multiples of three, one node in each failure domain, is an easy way to satisfy [pod placement rules](#).

**Table 7.1. Aggregate available resource requirements**

Deployment Mode	Base services	Additional device Set
Internal	<ul style="list-style-type: none"> <li>● 30 CPU (logical)</li> <li>● 72 GiB memory</li> <li>● 3 storage devices</li> </ul>	<ul style="list-style-type: none"> <li>● 6 CPU (logical)</li> <li>● 15 GiB memory</li> <li>● 3 storage devices</li> </ul>
External	<ul style="list-style-type: none"> <li>● 4 CPU (logical)</li> <li>● 16 GiB memory</li> </ul>	Not applicable

Example: For a 3 node cluster in an internal mode deployment with a single device set, a minimum of  $3 \times 10 = 30$  units of CPU are required.

For more information, see [Chapter 6, Subscriptions](#) and [CPU units](#).

### CPU units

In this section, 1 CPU Unit maps to the Kubernetes concept of 1 CPU unit.

- 1 unit of CPU is equivalent to 1 core for non-hyperthreaded CPUs.
- 2 units of CPU are equivalent to 1 core for hyperthreaded CPUs.
- OpenShift Container Storage core-based subscriptions always come in pairs (2 cores).

For additional guidance with designing your OpenShift Container Storage cluster, see the [OCS Sizing Tool](#).

### 7.3.1. Minimum deployment resource requirements [Technology Preview]

An OpenShift Container Storage cluster will be deployed with minimum configuration when the standard deployment resource requirement is not met.

**Table 7.2. Aggregate resource requirements**

Deployment Mode	Base services
Internal	<ul style="list-style-type: none"> <li>● 24 CPU (logical)</li> <li>● 72 GiB memory</li> <li>● 3 storage devices</li> </ul>

If you want to add additional device sets, we recommend converting your minimum deployment to standard deployment.



### IMPORTANT

Deployment of OpenShift Container Storage with minimum configuration is a Technology Preview feature. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information, see [Technology Preview Features Support Scope](#).


### 7.3.2. Compact deployment resource requirements [Technology Preview]

OpenShift Container Storage can now be installed on a three-node OpenShift compact bare metal cluster, where all the workloads run on three strong master nodes. There are no worker or storage nodes.

**Table 7.3. Aggregate resource requirements for OpenShift Container Storage only**

Deployment Mode	Base services	Additional device Set
Internal	<ul style="list-style-type: none"> <li>● 24 CPU (logical)</li> <li>● 72 GiB memory</li> <li>● 3 storage devices</li> </ul>	<ul style="list-style-type: none"> <li>● 6 CPU (logical)</li> <li>● 15 GiB memory</li> <li>● 3 storage devices</li> </ul>

To configure OpenShift Container Platform on a compact bare metal cluster, see [Configuring a three-node cluster](#) and [Delivering a Three-node Architecture for Edge Deployments](#).



## IMPORTANT

Deployment of OpenShift Container Storage on a three-node compact cluster is a Technology Preview feature. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information, see [Technology Preview Features Support Scope](#).

## 7.4. POD PLACEMENT RULES

Kubernetes is responsible for pod placement based on declarative placement rules. The OpenShift Container Storage base service placement rules for Internal cluster can be summarized as follows:

- Nodes are labeled with the **cluster.ocs.openshift.io/openshift-storage** key
- Nodes are sorted into pseudo failure domains if none exist
- Components requiring high availability are spread across failure domains
- A storage device must be accessible in each failure domain

This leads to the requirement that there be at least three nodes, and that nodes be in three distinct rack or zone failure domains in the case of pre-existing [topology labels](#).

For additional device sets, there must be a storage device, and sufficient resources for the pod consuming it, in each of the three failure domains. Manual placement rules can be used to override default placement rules, but generally this approach is only suitable for bare metal deployments.

## 7.5. STORAGE DEVICE REQUIREMENTS

Use this section to understand the different storage capacity requirements that you can consider when planning internal mode deployments and upgrades. We generally recommend 9 devices or less per node. This recommendation ensures both that nodes stay below cloud provider dynamic storage device attachment limits, and to limit the recovery time after node failures with local storage devices. Expanding the cluster in multiples of three, one node in each failure domain, is an easy way to satisfy [pod placement rules](#).



### NOTE

You can expand the storage capacity only in the increment of the capacity selected at the time of installation.

### 7.5.1. Dynamic storage devices

Red Hat OpenShift Container Storage permits the selection of either 0.5 TiB, 2 TiB or 4 TiB capacities as the request size for dynamic storage device sizes. The number of dynamic storage devices that can run per node is a function of the node size, underlying provisioner limits and [resource requirements](#).



### NOTE

OpenShift Container Storage 4.6 on IBM Power Systems does not support dynamic storage devices.

## 7.5.2. Local storage devices

For local storage deployment, any disk size of 4 TiB or less can be used, and all disks should be of the same size and type. The number of local storage devices that can run per node is a function of the node size and [resource requirements](#). Expanding the cluster in multiples of three, one node in each failure domain, is an easy way to satisfy [pod placement rules](#).

## 7.5.3. Capacity planning

Always ensure that available storage capacity stays ahead of consumption. Recovery is difficult if available storage capacity is completely exhausted, and requires more intervention than simply adding capacity or deleting or migrating content.

Capacity alerts are issued when cluster storage capacity reaches 75% (near-full) and 85% (full) of total capacity. Always address capacity warnings promptly, and review your storage regularly to ensure that you do not run out of storage space. If you do run out of storage space completely, contact [Red Hat Customer Support](#).

The following tables show example node configurations for Red Hat OpenShift Container Storage with dynamic storage devices.

**Table 7.4. Example initial configurations with 3 nodes**

Storage Device size	Storage Devices per node	Total capacity	Usable storage capacity
0.5 TiB	1	1.5 TiB	0.5 TiB
2 TiB	1	6 TiB	2 TiB
4 TiB	1	12 TiB	4 TiB


**Table 7.5. Example of expanded configurations with 30 nodes (N)**

Storage Device size (D)	Storage Devices per node (M)	Total capacity (D * M * N)	Usable storage capacity (D*M*N/3)
0.5 TiB	3	45 TiB	15 TiB
2 TiB	6	360 TiB	120 TiB
4 TiB	9	1080 TiB	360 TiB

## CHAPTER 8. DISCONNECTED ENVIRONMENT

Disconnected environment is a network restricted environment where Operator Lifecycle Manager (OLM) cannot access the default Operator Hub and image registries, which require Internet connectivity.

Red Hat supports deployment of OpenShift Container Storage in disconnected environments where OpenShift Container Platform is installed in restricted networks.



### NOTE

When you install OpenShift Container Storage in a restricted network environment, apply a custom Network Time Protocol (NTP) configuration to the nodes, because by default, internet connectivity is assumed in OpenShift Container Platform and chronyd is configured to use **\*.rhel.pool.ntp.org** servers. See [Red Hat Knowledgebase article](#) and [Configuring chrony time service](#) for more details.

For more information, see [Operators guide to using Operator Lifecycle Manager on restricted networks](#) .



### NOTE

OpenShift Container Storage 4.6 on IBM Power Systems does not support Disconnected environment.



---

## CHAPTER 9. NEXT STEPS

To start deploying your OpenShift Container Storage, you can use the internal mode within OpenShift Container Platform or use external mode to make available services from a cluster running outside of OpenShift Container Platform.

Depending on your requirement, go to the respective deployment guides.

### Internal mode

- [Deploying OpenShift Container Storage using Amazon web services](#)
- [Deploying OpenShift Container Storage using Bare Metal](#)
- [Deploying OpenShift Container Storage using VMWare vSphere](#)
- [Deploying OpenShift Container Storage using Microsoft Azure and Azure Red Hat OpenShift](#)
- [Deploying OpenShift Container Storage using Google Cloud](#) [Technology Preview]
- [Deploying OpenShift Container Storage using Red Hat OpenStack Platform](#) [Technology Preview]
- [Deploying OpenShift Container Storage using Red Hat Virtualization Platform](#) [Technology Preview]
- [Deploying OpenShift Container Storage on IBM Power Systems](#)
- [Deploying OpenShift Container Storage on IBM Z](#)

### External mode

- [Deploying OpenShift Container Storage in external mode](#)