



Red Hat OpenShift Container Storage 4.4

Planning your deployment

Important considerations when deploying RHOCS 4.4

Red Hat OpenShift Container Storage 4.4 Planning your deployment

Important considerations when deploying RHOCS 4.4

Legal Notice

Copyright © 2021 Red Hat, Inc.

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

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

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

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

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

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION TO RED HAT OPENSIFT CONTAINER STORAGE	3
CHAPTER 2. ARCHITECTURE OF OPENSIFT CONTAINER STORAGE	4
2.1. SUPPORTED WORKLOAD TYPES	4
CHAPTER 3. SUPPORTED INFRASTRUCTURE AND PLATFORMS	6
3.1. NODE REQUIREMENTS	6
3.2. PLATFORM REQUIREMENTS	8
CHAPTER 4. SUPPORTED CONFIGURATIONS	9
4.1. USING DYNAMICALLY CREATED STORAGE	9
4.1.1. Storage class requirements	9
4.1.2. Sizing and scaling for dynamic storage	9
4.2. USING LOCAL STORAGE DEVICES	10
4.2.1. Storage sizing and scaling for local devices	10
CHAPTER 5. NEXT STEPS	12

CHAPTER 1. INTRODUCTION TO RED HAT OPENSIFT CONTAINER STORAGE

Red Hat OpenShift Container Storage is software-defined storage that is optimised for container environments. It runs as an operator on Red Hat OpenShift Container Platform to provide highly integrated and simplified persistent storage management for containers.

Red Hat OpenShift Container Storage supports a variety of storage types, including:

- Block storage for databases
- Shared file storage for continuous integration, messaging, and data aggregation
- Object storage for archival, backup, and media storage

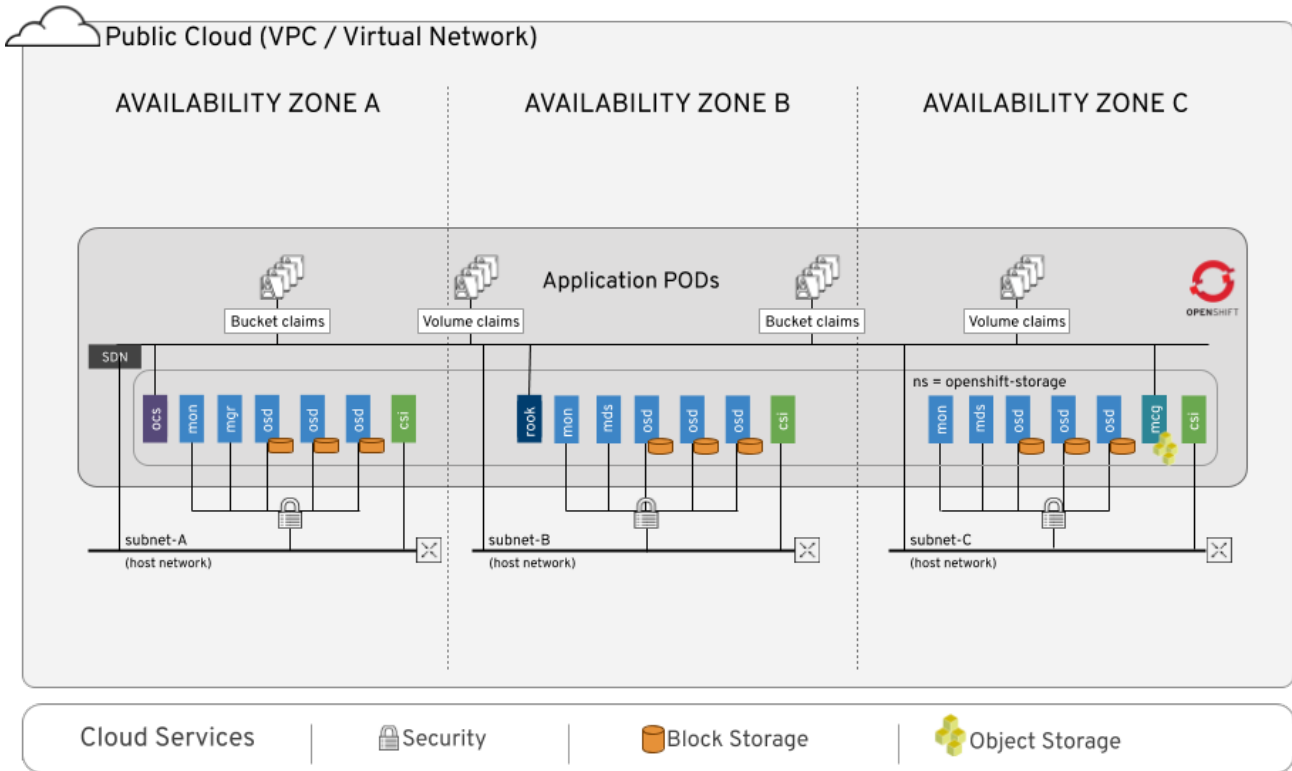
Red Hat OpenShift Container Storage version 4.2 onwards uses:

- Red Hat Ceph Storage to provide the file, block, and object storage that backs persistent volumes
- Rook.io to manage and orchestrate provisioning of persistent volumes and claims.
- NooBaa provides object storage, and its Multicloud Object Gateway enables data federation across multiple cloud environments. NooBaa's enterprise version is named Red Hat Multicloud Object Gateway.

CHAPTER 2. ARCHITECTURE OF OPENSIFT CONTAINER STORAGE

Red Hat OpenShift Container Storage uses Red Hat OpenShift Container Platform as a base. For information about the architecture and lifecycle of OpenShift Container Platform, see [OpenShift Container Platform architecture](#).

Red Hat OpenShift Container Storage architecture



Red Hat OpenShift Container Storage comprises of 3 main operators, which codify administrative tasks and custom resources so that task and resource characteristics can be more easily automated. Administrators define the desired end state of the cluster, and various operators work to ensure the cluster is either in that state, or approaching that state, with minimal administrator intervention.

The OpenShift Container Storage (OCS) 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.

The Rook-Ceph operator

This operator automates the packaging, deployment, management, upgrading, and scaling of persistent storage provided to container applications, and infrastructure services provided to OpenShift Container Platform. It provides the Ceph cluster resource, which manages the pods that host services such as the Object Storage Daemons, monitors, and the metadata server for the Ceph file system.

The NooBaa operator

This operator provides the Multicloud Object Gateway, an S3 compatible object store service that allows resource access across multiple cloud environments.

2.1. SUPPORTED WORKLOAD TYPES

Red Hat OpenShift Container Storage provides storage appropriate for a number of workload types.

Block storage is suitable for databases and other low-latency transactional workloads. Some examples of supported workloads are Red Hat OpenShift Container Platform logging and monitoring, and PostgreSQL.

Object storage is for video and audio files, compressed data archives, and the data used to train artificial intelligence or machine learning programs. In addition, object storage can be used for any application developed with a cloud-first approach.

File storage is for continuous integration and delivery, web application file storage, and artificial intelligence or machine learning data aggregation. Supported workloads include Red Hat OpenShift Container Platform registry and messaging using JBoss AMQ.

CHAPTER 3. SUPPORTED INFRASTRUCTURE AND PLATFORMS

Red Hat OpenShift Container Storage supports deployment on Red Hat OpenShift Container Platform deployed with Installer Provisioned Infrastructure (Full Stack Automation) and User Provisioned Infrastructure (Pre Existing Infrastructure).

- **Installer Provisioned Infrastructure (IPI)**
With full stack automation, the installer controls all areas of the installation including infrastructure provisioning with an opinionated best practices deployment on Red Hat OpenShift Container Platform.
- **User Provisioned Infrastructure (UPI)**
With pre-existing infrastructure deployments, administrators are responsible for creating and managing their own infrastructure allowing greater customization and operational flexibility.

Red Hat OpenShift Container Storage supports deployment on the following infrastructures:

Table 3.1. Minimum infrastructure requirements

Infrastructure	Deployment type	Requirements
Amazon Web Services	IPI, UPI	<ul style="list-style-type: none"> • Amazon EC2 with <ul style="list-style-type: none"> ◦ a minimum instance type m5.4xlarge ◦ i3en type i3en.2xlarge (Technology preview). See Creating OpenShift Container Storage cluster on Amazon EC2 for more details.
VMware	UPI	vSphere 6.7 Update 2 and higher with vSAN or VMFS datastore. See VMware vSphere infrastructure requirements for details.
Bare Metal	UPI	Storage type must be SSD (NVMe/SATA/SAS)

3.1. NODE REQUIREMENTS

The requirements mentioned in the section apply to pre-existing Red Hat OpenShift Container Platform deployments only and are required for exclusive use by Red Hat OpenShift Container Storage.

A minimum of 3 starter nodes are required for deployment. The starter node minimum requirements support a base set of services, and a single storage device per node. Additional storage devices can be added later, and Kubernetes will schedule OCS pods based on resource availability. In case you plan to run any other workload on a storage node, you must add additional resources (CPU/Memory/Space).

Table 3.2. Minimum requirements for starter nodes

Components	Requirements
CPU	10*
Memory	40 GB*
Disk	<ul style="list-style-type: none"> ● For dynamic storage deployments, <ul style="list-style-type: none"> ○ 1 disk of size 0.5 TiB or 2 TiB or 4 TiB storage ● For Local Storage deployment, <ul style="list-style-type: none"> ○ any disk size of 4 TiB or lesser can be used. ○ disk partitioning is not supported. ○ maintain uniform disk sizes across nodes for storage disks.
MON	<p>10 GiB storage per MON on all the storage nodes.</p> <ul style="list-style-type: none"> ● MON will use 10 GiB storage space under <code>/var/lib/rook</code> for local storage deployments. No PV is necessary for MONs. ● In case a node with MON fails and the MON failover to a new node, 10 GiB space will then be consumed on the new node as well. <p>Note: At a time only 3 nodes require the storage space.</p>

*For AWS EC2, the closest fit instance size provides 16 CPU and 64 GB memory (M5.4xlarge). When the **M5.4xlarge** instance type is used for starter nodes, they will support up to 4 storage devices each. Use of the **i3en.2xlarge** instance type (8 CPUs, 64GB memory) is currently under *technology preview* support.

Table 3.3. Additional requirement per storage device

Components	Requirements
CPU	2
Memory	8 GB
Note: Storage devices must be provisioned in sets of three	

CPU units

- In this section, 1 CPU Unit maps to the Kubernetes concept of 1 CPU unit. For more information, see [CPU units](#).
 - 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). Example: For a 3 node cluster, a minimum of $3 \times 10 = 30$ units of CPU are required. 30 units of CPU are equivalent to 15 cores which is equivalent to ~8 quantities of Red Hat OpenShift Container Storage (2 core) subscriptions.

3.2. PLATFORM REQUIREMENTS

Red Hat OpenShift Container Storage is compatible only with the latest Red Hat OpenShift Container Platform versions.

For Red Hat OpenShift Container Storage 4.4, it is recommended to use Red Hat OpenShift Container Platform 4.4 for flexibility in deployment.



IMPORTANT

When upgrading Red Hat OpenShift Container Platform, you must upgrade Local Storage Operator version to match with the Red Hat OpenShift Container Platform version in order to have the Local Storage Operator fully supported with Red Hat OpenShift Container Storage.

For more information, see [Red Hat OpenShift Container Storage and Red Hat OpenShift Container Platform interoperability matrix](#).



WARNING

Red Hat OpenShift Container Platform must not be installed or running with Federal Information Processing Standards (FIPS) mode enabled.

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.

CHAPTER 4. SUPPORTED CONFIGURATIONS

Red Hat OpenShift Container Storage is deployed as a minimal cluster of 3 worker nodes. Spread the nodes across three different availability zones to ensure availability.

4.1. USING DYNAMICALLY CREATED STORAGE

Red Hat OpenShift Container Storage supports different storage disk sizes with a capacity of 0.5 TiB, 2 TiB and 4 TiB when storage is dynamically created.

4.1.1. Storage class requirements

Red Hat OpenShift Container Storage makes use of the Red Hat OpenShift Container Platform default storage class, and expects a certain default storage class depending on your infrastructure provider.

These classes are configured on Red Hat OpenShift Container Platform nodes automatically, but if your Red Hat OpenShift Container Platform node uses a different storage class as the default, you must change the default storage class back to the appropriate storage class for your infrastructure provider.

- On Amazon Web Services, the default storage class must be **gp2**.
- On VMware vSphere, the default storage class must be **thin**.

4.1.2. Sizing and scaling for dynamic storage

The initial cluster of 3 nodes can later be expanded with a recommendation of up to 9 disks in each node. In case of more than 3 worker nodes, the distribution of the disks depends on OpenShift scheduling and available resources.

Expand the cluster in sets of three nodes to ensure that your storage is replicated, and to ensure you can use at least three availability zones.



NOTE

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

The following tables shows the supported configurations for Red Hat OpenShift Container Storage.

Table 4.1. Initial configuration across 3 nodes

Disks	Disks 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 4.2. Example of an expanded configuration with 9 nodes (N)

Disk size (D)	Disks per node (M)	Total capacity (D * M * N)	Usable storage capacity (D*M*N/3)
0.5 TiB	3	13.5 TiB	4.5 TiB
2 TiB	6	108 TiB	36 TiB
4 TiB	9	324 TiB	108 TiB



WARNING

Always ensure that you have plenty of storage capacity.

If storage ever fills completely, it is not possible to add capacity or to delete or migrate content away from the storage to free up space. Completely full storage is very difficult to recover.

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.

As of Red Hat OpenShift Container Storage 4.4, installation is supported only on existing Red Hat OpenShift Container Platform nodes. See [Deploying OpenShift Container Storage](#) for more information.

4.2. USING LOCAL STORAGE DEVICES

Use of local devices requires the installation of the local storage operator using Operator Hub. See [Installing OpenShift Container Storage using local storage devices](#).

4.2.1. Storage sizing and scaling for local devices

For local storage deployment, any disk size of 4 TiB or lesser can be used, and all disks must be in the same size and type.

The initial cluster of 3 nodes can later be expanded with a recommendation of up to 9 disks in each node. In case of more than 3 worker nodes, the distribution of the disks depend on OpenShift scheduling and available resources.

**WARNING**

Always ensure that you have plenty of storage capacity.

If storage ever fills completely, it is not possible to add capacity or to delete or migrate content away from the storage to free up space. Completely full storage is very difficult to recover.

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.

CHAPTER 5. NEXT STEPS

Go to [Deploying OpenShift Container Storage](#) to start deploying your container storage solution.