



Red Hat OpenShift Container Storage 4.2

Planning your deployment

Important considerations when deploying RHOCS 4.2

Red Hat OpenShift Container Storage 4.2 Planning your deployment

Important considerations when deploying RHOCS 4.2

Legal Notice

Copyright © 2020 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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 4.2	3
CHAPTER 2. ARCHITECTURE OF OPENSIFT CONTAINER STORAGE	4
2.1. SUPPORTED WORKLOAD TYPES	5
CHAPTER 3. SUPPORTED INFRASTRUCTURE AND PLATFORMS	6
3.1. NODE REQUIREMENTS	6
3.2. STORAGE CLASS REQUIREMENTS	7
3.3. SOFTWARE REQUIREMENTS	7
CHAPTER 4. SUPPORTED CONFIGURATIONS	9
4.1. SIZING AND SCALING RECOMMENDATIONS	9
CHAPTER 5. NEXT STEPS	11

CHAPTER 1. INTRODUCTION TO RED HAT OPENSIFT CONTAINER STORAGE 4.2

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

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.



IMPORTANT

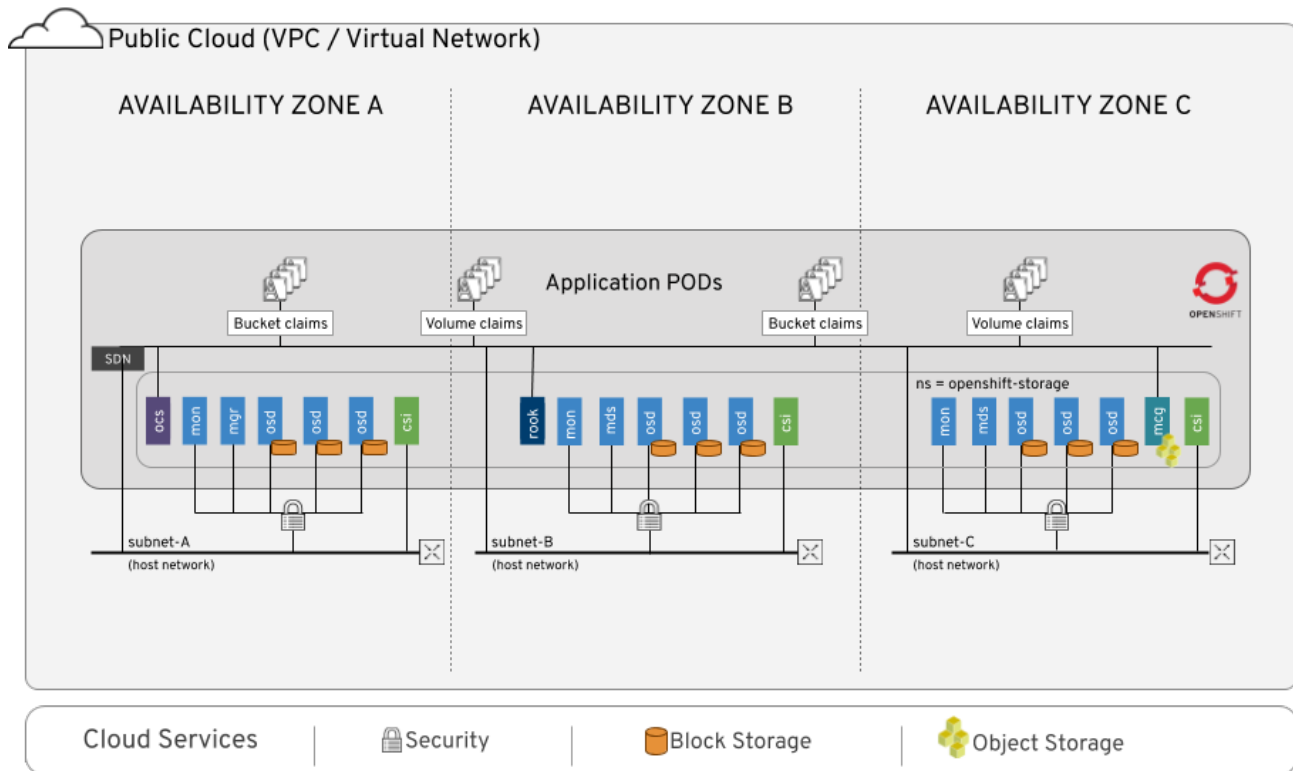
Technology Preview features provide early access to upcoming product innovations, enabling you to test functionality and provide feedback during the development process. However, these features are not fully supported under Red Hat Subscription Level Agreements, may not be functionally complete, and are not intended for production use. As Red Hat considers making future iterations of Technology Preview features generally available, we will attempt to resolve any issues that customers experience when using these features.

See [Technology Preview Features Support Scope](#) for more information.

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](#).

OpenShift Container Storage architecture



Red Hat OpenShift Container Storage 4 is comprised 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.

OpenShift Container Storage uses the following operators:

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 (OSDs), 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	Amazon EC2 with a minimum instance type m5.4xlarge
VMware	UPI	vSphere 6.7 Update 2 and higher with vSAN or VMFS datastore. See VMware vSphere infrastructure requirements for details.

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 nodes are required for deployment with at least one Object Storage Daemon (OSD) and one Monitor daemon (MON) on each node.

Table 3.2. Minimum requirements for each starting node

Components	Requirements
CPU	16
Memory	64 GB
Disk	2 TiB storage per disk (1 disk by default; scalable to 3 disks per node) Maintain uniform disk sizes across nodes for storage disks.

Components	Requirements
MON	<p>10 GiB storage per MON on all the storage nodes.</p> <ul style="list-style-type: none"> 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>



NOTE

- In case you plan to run any other workload on a storage node, you must add additional resources (CPU/Memory/Space).
- 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.
 - Red Hat OpenShift Container Storage core-based subscriptions always come in pairs (2 cores).
Example: For a 3 node cluster, a minimum of $3 \times 16 = 48$ units of CPU are required. 48 units of CPU are equivalent to 24 cores which is equivalent to 12 quantity of Red Hat OpenShift Container Storage (2 core) subscriptions.

3.2. STORAGE CLASS REQUIREMENTS

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

These classes are configured on OpenShift Container Platform nodes automatically, but if your 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**.

3.3. SOFTWARE REQUIREMENTS

Red Hat OpenShift Container Storage 4.2 is supported on Red Hat OpenShift Container Platform 4.2.13 and later versions.

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



IMPORTANT

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 4.2 is deployed as a minimal cluster of 3 worker nodes. Spread the nodes across three different availability zones to ensure availability.

Each node in this initial cluster runs one 2 TiB disk, which replicate their contents to each other, providing a usable starting capacity of 2 TiB. The initial cluster will consist of 3 nodes, regardless of the number of nodes selected for storage nodes.

The initial cluster can later be expanded to a maximum of 9 nodes that support up to 9 disks for a 3 node cluster.

The following table shows the supported configurations for OpenShift Container Storage.

Table 4.1. Supported configuration across 3 nodes

	Number of nodes	Disks	Total capacity	Usable storage capacity
Initial configuration	3 Nodes	1 disk of 2 TiB on each node	6 TiB	2 TiB
Possible expansion	Maximum of up to 9 nodes	3 disks of 2TiB on each node, which is a maximum of 27 disks of 2 TiB each	Maximum of 54 TiB	Maximum of 18 TiB



NOTE

The distribution of the disks depends on OpenShift scheduling and available resources.

As of version 4.2 GA, installation is supported only on existing Red Hat OpenShift Container Platform nodes. See [Deploying OpenShift Container Storage](#) for more information.

4.1. SIZING AND SCALING RECOMMENDATIONS

Red Hat OpenShift Container Storage 4.2 supports a minimum of 3 nodes and a maximum of 9 nodes.

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.



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.