



# Red Hat OpenShift Container Storage 4.6

## Deploying and managing OpenShift Container Storage using IBM Z

How to install and set up your IBM Z environment



# Red Hat OpenShift Container Storage 4.6 Deploying and managing OpenShift Container Storage using IBM Z

---

How to install and set up your IBM Z environment

## 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<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 instructions on installing Red Hat OpenShift Container Storage 4.6 to use local storage on IBM Z infrastructure. While this document refers only to IBM Z, all information in it also applies to LinuxONE. Deploying and managing OpenShift Container Storage on IBM Z infrastructure 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.

## Table of Contents

<b>PREFACE</b> .....	<b>3</b>
<b>CHAPTER 1. DEPLOYING USING LOCAL STORAGE DEVICES</b> .....	<b>4</b>
1.1. REQUIREMENTS FOR INSTALLING OPENSIFT CONTAINER STORAGE USING LOCAL STORAGE DEVICES	4
1.2. INSTALLING RED HAT OPENSIFT CONTAINER STORAGE OPERATOR	4
1.3. INSTALLING LOCAL STORAGE OPERATOR	6
1.4. FINDING AVAILABLE STORAGE DEVICES	7
1.5. CREATING OPENSIFT CONTAINER STORAGE CLUSTER ON IBM Z	8
<b>CHAPTER 2. VERIFYING OPENSIFT CONTAINER STORAGE DEPLOYMENT FOR INTERNAL-ATTACHED DEVICES MODE</b> .....	<b>16</b>
2.1. VERIFYING THE STATE OF THE PODS	16
2.2. VERIFYING THE OPENSIFT CONTAINER STORAGE CLUSTER IS HEALTHY	17
2.3. VERIFYING THE MULTICLOUD OBJECT GATEWAY IS HEALTHY	18
2.4. VERIFYING THAT THE OPENSIFT CONTAINER STORAGE SPECIFIC STORAGE CLASSES EXIST	19
<b>CHAPTER 3. UNINSTALLING OPENSIFT CONTAINER STORAGE IN INTERNAL-ATTACHED DEVICES MODE</b>	<b>20</b>
3.1. UNINSTALLING OPENSIFT CONTAINER STORAGE IN INTERNAL-ATTACHED DEVICES MODE	20
3.1.1. Removing local storage operator configurations	25
3.2. REMOVING MONITORING STACK FROM OPENSIFT CONTAINER STORAGE	26
3.3. REMOVING OPENSIFT CONTAINER PLATFORM REGISTRY FROM OPENSIFT CONTAINER STORAGE	29
3.4. REMOVING THE CLUSTER LOGGING OPERATOR FROM OPENSIFT CONTAINER STORAGE	30
<b>CHAPTER 4. SCALING STORAGE NODES</b> .....	<b>32</b>
4.1. REQUIREMENTS FOR SCALING STORAGE NODES	32
4.2. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT CONTAINER STORAGE NODES ON IBM Z	32
4.3. SCALING OUT STORAGE CAPACITY BY ADDING NEW NODES	35
4.3.1. Adding a node on IBM Z	35
4.3.2. Verifying the addition of a new node	36
4.3.3. Scaling up storage capacity	36
<b>CHAPTER 5. REPLACING STORAGE NODES</b> .....	<b>37</b>
5.1. REPLACING OPERATIONAL NODES ON IBM Z	37
5.2. REPLACING FAILED NODES ON IBM Z	38
<b>CHAPTER 6. REPLACING STORAGE DEVICES</b> .....	<b>41</b>
6.1. REPLACING OPERATIONAL OR FAILED STORAGE DEVICES ON IBM Z	41



---

## PREFACE

Red Hat OpenShift Container Storage 4.6 supports deployment on existing Red Hat OpenShift Container Platform (RHOCP) IBM Z clusters in connected or disconnected environments along with out-of-the-box support for proxy environments.



### NOTE

Only internal Openshift Container Storage clusters are supported on IBM Z. See [Planning your deployment](#) for more information about deployment requirements.

To deploy OpenShift Container Storage, follow the appropriate deployment process for your environment:

- Internal Attached Devices mode
  - [Deploy using local storage devices](#)

# CHAPTER 1. DEPLOYING USING LOCAL STORAGE DEVICES

Deploying OpenShift Container Storage on OpenShift Container Platform using local storage devices provides you with the option to create internal cluster resources. Follow this deployment method to use local storage to back persistent volumes for your OpenShift Container Platform applications.

Use this section to deploy OpenShift Container Storage on IBM Z infrastructure where OpenShift Container Platform is already installed.

To deploy Red Hat OpenShift Container Storage using local storage, follow these steps:

1. Understand the [requirements for installing OpenShift Container Storage using local storage devices](#).
2. [Install the Red Hat OpenShift Container Storage Operator](#) .
3. [Install Local Storage Operator](#) .
4. [Find the available storage devices](#) .
5. [Creating OpenShift Container Storage cluster service on IBM Z](#) .

## 1.1. REQUIREMENTS FOR INSTALLING OPENSIFT CONTAINER STORAGE USING LOCAL STORAGE DEVICES

- You must upgrade to a latest version of OpenShift Container Platform 4.6 before deploying OpenShift Container Storage 4.6. For information, see [Updating OpenShift Container Platform clusters](#) guide.
- The Local Storage Operator version must match the Red Hat OpenShift Container Platform version in order to have the Local Storage Operator fully supported with Red Hat OpenShift Container Storage. The Local Storage Operator does not get upgraded when Red Hat OpenShift Container Platform is upgraded.
- You must have at least three OpenShift Container Platform worker nodes in the cluster with locally attached storage devices on each of them.
  - Each of the three selected nodes must have at least one raw block device available to be used by OpenShift Container Storage.
  - The devices to be used must be empty, that is, there should be no persistent volumes (PVs), volume groups (VGs), or local volumes (LVs) remaining on the disks.
- If you upgraded to OpenShift Container Storage 4.6 from a previous version, ensure that you have followed post-upgrade procedures to create the **LocalVolumeDiscovery** object. See [Post-update configuration changes](#) for details.
- If you upgraded from a previous version of OpenShift Container Storage, create a **LocalVolumeSet** object to enable automatic provisioning of devices as described in [Post-update configuration changes](#).
- For minimum starting node requirements, see [Resource requirements](#) section in Planning guide.

## 1.2. INSTALLING RED HAT OPENSIFT CONTAINER STORAGE OPERATOR



You can install Red Hat OpenShift Container Storage Operator using the Red Hat OpenShift Container Platform Operator Hub. For information about the hardware and software requirements, see [Planning your deployment](#).

## Prerequisites

- You must be logged into the OpenShift Container Platform (RHOC) cluster.
- You must have at least three worker nodes in the RHOC cluster.



## NOTE

- When you need to override the cluster-wide default node selector for OpenShift Container Storage, you can use the following command in command line interface to specify a blank node selector for the **openshift-storage** namespace:

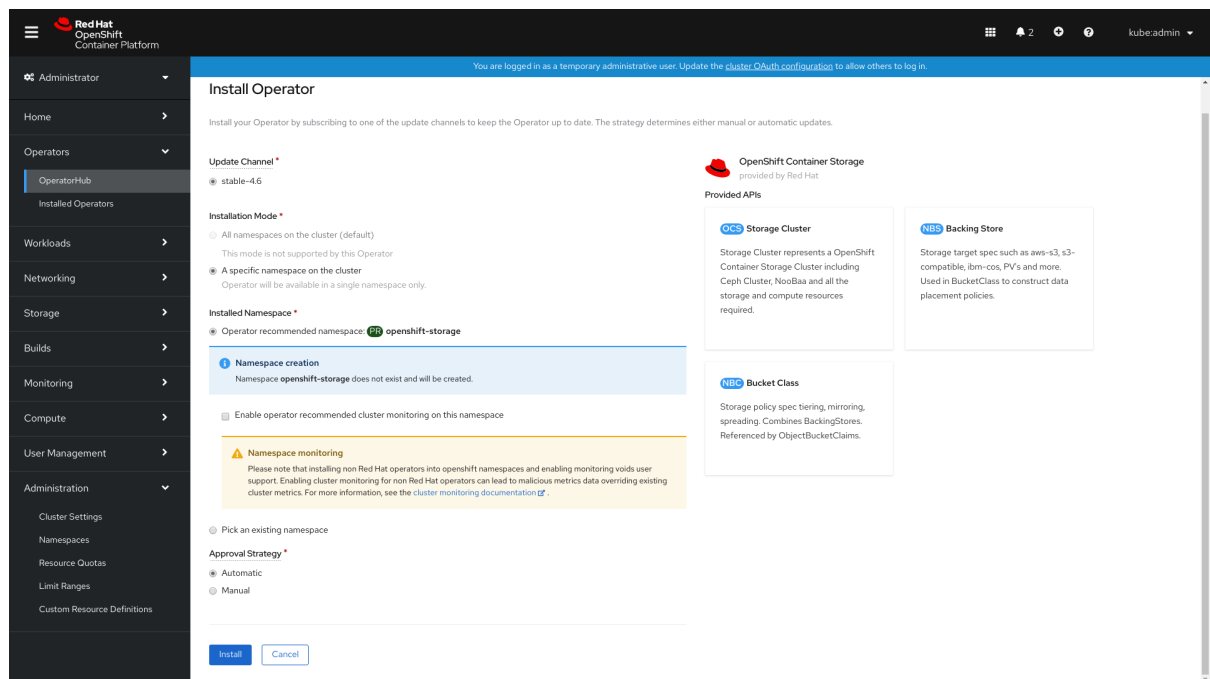
```
$ oc annotate namespace openshift-storage openshift.io/node-selector=
```

- Taint a node as **infra** to ensure only Red Hat OpenShift Container Storage resources are scheduled on that node. This helps you save on subscription costs. For more information, see [How to use dedicated worker nodes for Red Hat OpenShift Container Storage](#) chapter in Managing and Allocating Storage Resources guide.

## Procedure

1. Click **Operators → OperatorHub**.
2. Use Filter by keyword text box or the filter list to search for OpenShift Container Storage from the list of operators.
3. Click **OpenShift Container Storage**.
4. Click **Install** on the OpenShift Container Storage operator page.
5. On the **Install Operator** page, the following required options are selected by default:
  - a. Update Channel as **stable-4.6**
  - b. Installation Mode as **A specific namespace on the cluster**
  - c. Installed Namespace as **Operator recommended namespace openshift-storage**. If Namespace **openshift-storage** does not exist, it will be created during the operator installation.
6. Select **Enable operator recommended cluster monitoring on this namespace** checkbox as this is required for cluster monitoring.
7. Approval Strategy is set to **Automatic** by default.

Figure 1.1. Install Operator page



8. Click **Install**.

### Verification steps

- Verify that **OpenShift Container Storage** Operator shows a green tick indicating successful installation.
- Click **View Installed Operators** in namespace `openshift-storage` link to verify that OpenShift Container Storage Operator shows the **Status** as **Succeeded** on the Installed Operators dashboard.

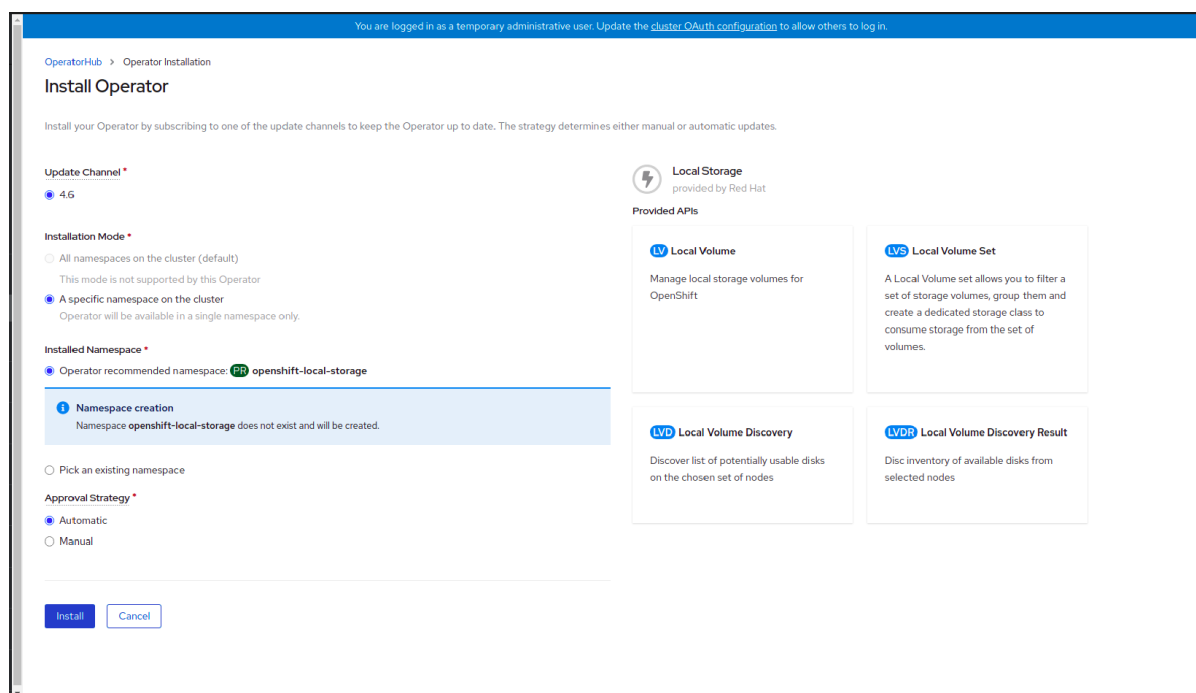
## 1.3. INSTALLING LOCAL STORAGE OPERATOR

Use this procedure to install the Local Storage Operator from the Operator Hub before creating OpenShift Container Storage clusters on local storage devices.

### Procedure

1. Log in to the OpenShift Web Console.
2. Click **Operators** → **OperatorHub**.
3. Search for **Local Storage Operator** from the list of operators and click on it.
4. Click **Install**.

Figure 1.2. Install Operator page



5. Set the following options on the **Install Operator** page:
  - a. Update Channel as **stable-4.6**
  - b. Installation Mode as **A specific namespace on the cluster**
  - c. Installed Namespace as **Operator recommended namespace openshift-local-storage**.
  - d. Approval Strategy as **Automatic**
6. Click **Install**.
7. Verify that the Local Storage Operator shows the Status as **Succeeded**.

## 1.4. FINDING AVAILABLE STORAGE DEVICES

Use this procedure to identify the device names for each of the three or more worker nodes that you have labeled with the OpenShift Container Storage label **cluster.ocs.openshift.io/openshift-storage=** before creating Persistent Volumes (PV) for IBM Z.

### Procedure

1. List and verify the name of the worker nodes with the OpenShift Container Storage label.

```
$ oc get nodes -n openshift-storage
```

Example output:

```
NAME           STATUS    ROLES    AGE    VERSION
bmworker01    Ready    worker   6h45m v1.16.2
bmworker02    Ready    worker   6h45m v1.16.2
bmworker03    Ready    worker   6h45m v1.16.2
```

- Log in to each worker node that is used for OpenShift Container Storage resources and find the unique **by-id** device name for each available raw block device.

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

Example output:

```
$ oc debug node/bmworker01
Starting pod/bmworker01-debug ...
To use host binaries, run `chroot /host`
Pod IP: 10.0.135.71
If you don't see a command prompt, try pressing enter.
sh-4.2# chroot /host
sh-4.4# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
xvda                                202:0   0 120G  0 disk
|-xvda1                             202:1   0  384M  0 part /boot
|-xvda2                             202:2   0  127M  0 part /boot/efi
|-xvda3                             202:3   0    1M  0 part
`-xvda4                             202:4   0 119.5G  0 part
`-coreos-luks-root-nocrypt 253:0   0 119.5G  0 dm  /sysroot
nvme0n1                             259:0   0  931G  0 disk
```

In this example, for **bmworker01**, the available local device is **nvme0n1**.

- Identify the unique ID for each of the devices selected in Step 2.

```
sh-4.4# ls -l /dev/disk/by-id/ | grep nvme0n1
lrwxrwxrwx. 1 root root 13 Mar 17 16:24 nvme-
INTEL_SSDPE2KX010T7_PHLF733402LM1P0GGN -> ../../nvme0n1
```

In the above example, the ID for the local device **nvme0n1**

```
nvme-INTEL_SSDPE2KX010T7_PHLF733402LM1P0GGN
```

- Repeat the above step to identify the device ID for all the other nodes that have the storage devices to be used by OpenShift Container Storage. See this [Knowledge Base article](#) for more details.

## 1.5. CREATING OPENSIFT CONTAINER STORAGE CLUSTER ON IBM Z

Use this procedure to create storage cluster on IBM Z.

### Prerequisites

- Ensure that all the requirements in the [Requirements for installing OpenShift Container Storage using local storage devices](#) section are met.
- You must have three worker nodes with the same storage type and size attached to each node (for example, 2TB NVMe hard drive) to use local storage devices on bare metal.
- Verify your OpenShift Container Platform worker nodes are labeled for OpenShift Container Storage:

```
$ oc get nodes -l cluster.ocs.openshift.io/openshift-storage -o jsonpath='{range .items[*]}
{.metadata.name}{"\n"}'
```

To identify storage devices on each node, refer to [Finding available storage devices](#).

## Procedure

1. Create local persistent volumes (PVs) on the storage nodes using **LocalVolume** custom resource (CR).

Example of **LocalVolume** CR **local-storage-block.yaml** using OpenShift Container Storage label as node selector.

```
apiVersion: local.storage.openshift.io/v1
kind: LocalVolume
metadata:
  name: local-block
  namespace: local-storage
  labels:
    app: ocs-storagecluster
spec:
  nodeSelector:
    nodeSelectorTerms:
      - matchExpressions:
          - key: cluster.ocs.openshift.io/openshift-storage
            operator: In
            values:
              - ""
  storageClassDevices:
    - storageClassName: localblock
      volumeMode: Block
      devicePaths:
        - /dev/disk/by-id/nvme-INTEL_SSDPEKKA128G7_BTPY81260978128A # <-- modify
this line
        - /dev/disk/by-id/nvme-INTEL_SSDPEKKA128G7_BTPY80440W5U128A # <-- modify
this line
        - /dev/disk/by-id/nvme-INTEL_SSDPEKKA128G7_BTPYB85AABDE128A # <-- modify
this line
```

2. Create the **LocalVolume** CR for block PVs.

```
$ oc create -f local-storage-block.yaml
```

3. Check if the pods are created.

Example output:

NAME	READY	STATUS	RESTARTS	AGE
local-block-local-diskmaker-cmfql	1/1	Running	0	31s
local-block-local-diskmaker-g6fzr	1/1	Running	0	31s
local-block-local-diskmaker-jkqxt	1/1	Running	0	31s
local-block-local-provisioner-jgqcc	1/1	Running	0	31s
local-block-local-provisioner-mx49d	1/1	Running	0	31s
local-block-local-provisioner-qbcvp	1/1	Running	0	31s
local-storage-operator-54bc7566c6-ddbrt	1/1	Running	0	12m

4. Check if the PVs are created.

```
$ oc get pv
```

Example output:

```
NAME                CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM
STORAGECLASS  REASON  AGE
local-pv-150fdc87  931Gi    RWO           Delete          Available  localblock
2m11s
local-pv-183bfc0a  931Gi    RWO           Delete          Available  localblock
2m15s
local-pv-b2f5cb25  931Gi    RWO           Delete          Available  localblock
2m21s
```

5. Check for the new **localblock StorageClass**.

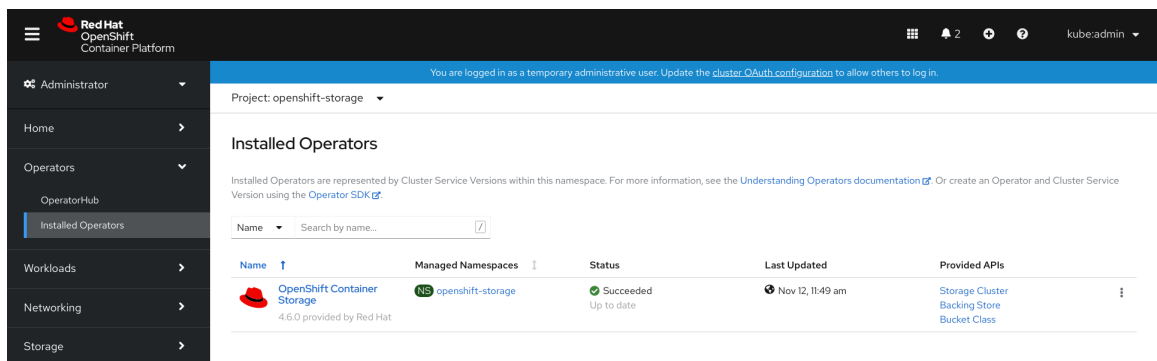
```
$ oc get sc | grep localblock
```

Example output:

```
NAME                PROVISIONER                RECLAIMPOLICY
VOLUMEBINDINGMODE  ALLOWVOLUMEEXPANSION  AGE
localblock          kubernetes.io/no-provisioner  Delete
WaitForFirstConsumer  false                  2m10s
```

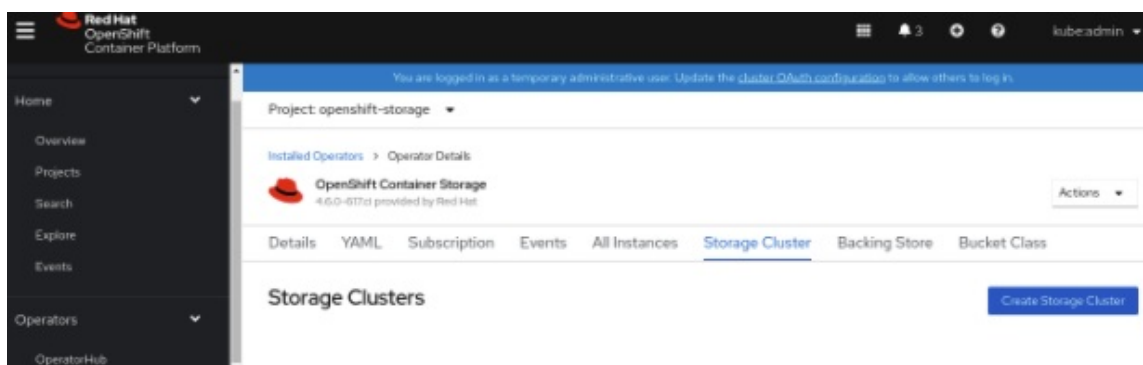
6. Create the OpenShift Container Storage Cluster Service that uses the **localblock** Storage Class.
  - a. Log into the OpenShift Web Console.
  - b. Click **Installed Operators** from the left pane of the OpenShift Web Console to view the installed operators.

Figure 1.3. OpenShift Container Storage Operator page




- c. Click the **OpenShift Container Storage** installed operator.
- d. On the **Operator Details** page, click the **Storage Cluster** link.

Figure 1.4. Storage Cluster tab



e. Click **Create Storage Cluster**.



7. Select **Internal-Attached devices** for the **Select Mode**.
8. Perform the steps to create the storage cluster based on any of the following possible scenarios that you are in.

#### Local Storage Operator is not installed and storage class not created

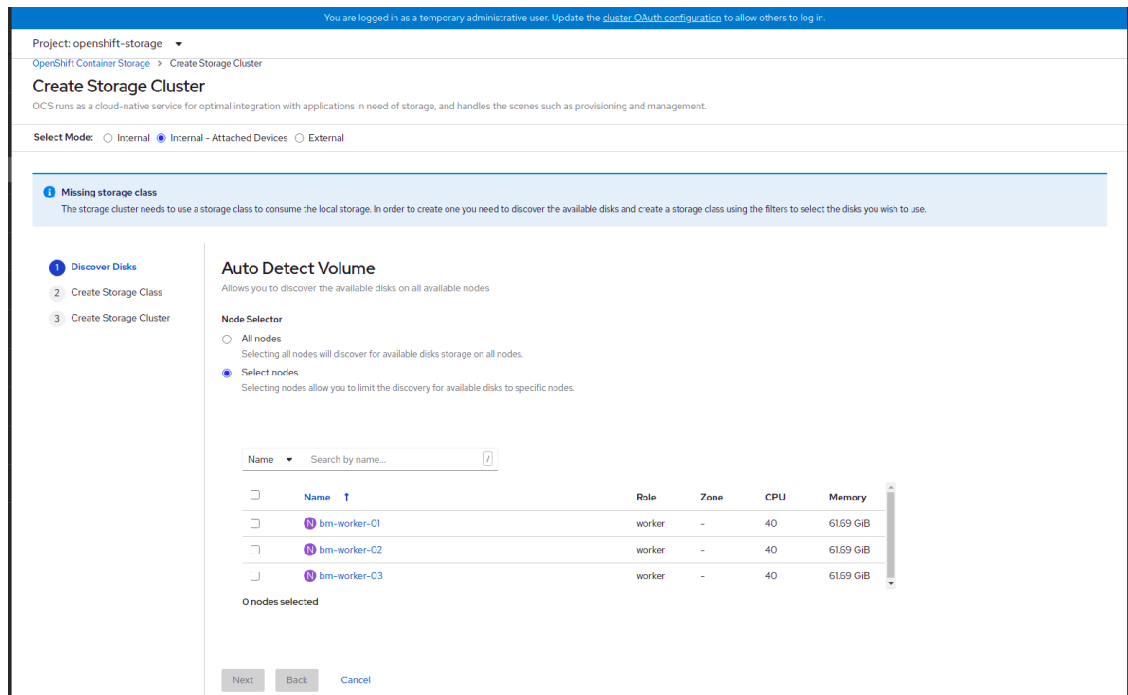
- a. You are prompted to install the Local Storage Operator. Click **Install** and install the operator as described in [Installing Local Storage Operator](#).
- b. Proceed to the storage class creation from the wizard as described in the next steps.

#### Local Storage Operator is installed but the storage class is not created

Perform the following steps in the wizard:

##### Discover Disks

Figure 1.5. Discovery Disks wizard page



a. Choose one of the following:

- **All nodes** to discover disks in all the nodes.
- **Select nodes** to choose a subset of nodes from the nodes listed. To find specific worker nodes in the cluster, you can filter nodes on the basis of Name or Label. Name allows you to search by name of the node and Label allows you to search by selecting the predefined label.

It is recommended that the worker nodes are spread across three different physical nodes, racks or failure domains for high availability.



#### NOTE

Ensure OpenShift Container Storage rack labels are aligned with physical racks in the datacenter to prevent a double node failure at the failure domain level.

If the nodes selected do not match the OpenShift Container Storage cluster requirement of an aggregated 30 CPUs and 72 GiB of RAM, a minimal cluster will be deployed. For minimum starting node requirements, see [Resource requirements](#) section in Planning guide.



#### NOTE

If the nodes to be selected are tainted and not discovered in the wizard, follow the steps provided in the [Red Hat Knowledgebase Solution](#) as a workaround.

b. Click **Next**.

## Create Storage Class



Figure 1.6. Create Storage Class wizard page

Project: opershift-storage

You are logged in as a temporary administrative user. Update the `cluster.OAuth.config` to allow others to log in.

The storage cluster needs to use a storage class to consume the local storage. In order to create one you need to discover the available disks and create a storage class using the filters to select the disks you wish to use.

1 Discover Disks  
2 **Create Storage Class**  
3 Create Storage Cluster

### Local Volume Set

A Local Volume Set allows you to filter a set of storage volumes, group them and create a dedicated storage class to consume storage for them.

Volume Set Name  
localblock

Storage Class Name  
localblock

#### Filter Disks

**Node Selector**

- All nodes (3 nodes)  
Selecting all nodes will use the available disks that match the selected filters on all nodes selected on previous step.
- Select nodes  
Selecting all nodes will use the available disks that match the selected filters only on selected nodes.

**Disk Type**  
All

**Advanced**

**Disk Mode**  
Block

**Disk Size**  
Min: 0 Max: [ ] TiB

**Max Disk Limit**  
Disk limit will set the maximum number of PVs to create on a node. If the field is empty, will create PVs for all available disks on the matching nodes.  
All

Next Back Cancel

Selected Capacity  
3 Nodes 3 Disks  
2.7 TiB  
Out of 2.7 TiB

- Enter the Volume set name.
- Enter the Storage class name. By default, the Volume set name appears for the storage class name.  
You might need to wait for a couple of seconds for the nodes to get populated.
- To select the nodes from the discovered nodes, you can choose one of the following:
  - All nodes** to select all the nodes for which you discovered the available disks.
  - Select nodes** to choose a subset of the nodes for which you discovered the available disks.
- (Optional) You can view the selected capacity of the disks on the selected nodes using the *donut* chart. Also, you can click on the **Nodes** and **Disks** links on the chart which brings up the list of nodes and disks. You can view more details about the selected nodes and disks.

Selected Nodes				
Name	Role	Zone	CPU	Memory
bm-worker-01	worker	-	40	61.69 GiB
bm-worker-02	worker	-	40	61.69 GiB
bm-worker-03	worker	-	40	61.69 GiB

Close

## Selected Disks

Name	Node	Type	Model	Capacity
/dev/sdb	worker-0.bmcluster.com	disk	2107900	150 GiB
/dev/sdb	worker-1.bmcluster.com	disk	2107900	150 GiB
/dev/sdb	worker-2.bmcluster.com	disk	2107900	150 GiB

Close

- e. Select the **Disk type**.
- f. In the **Advanced** section, you can set the following options:

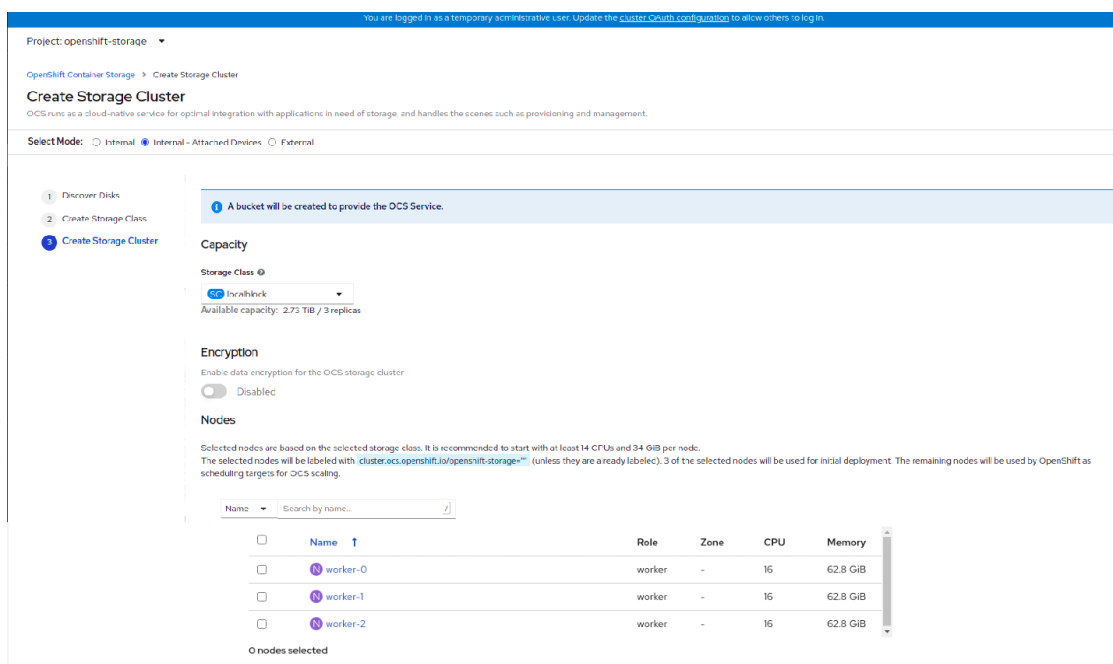
Disk Mode	Block is selected by default.
Disk Size	Minimum and maximum size of the device that needs to be included in <b>Max</b> and <b>Min</b> respectively.
Max Disk Limit	This indicates the maximum number of persistent volumes (PVs) that can be created on a node. If this field is left empty, then PVs are created for all the available disks on the matching nodes.

- g. Click **Next**.
- h. Click **Yes** in the message alert to confirm the creation of the storage class. After the local volume set and storage class are created, it is not possible to go back to the step.

Local Storage Operator is installed and the storage class is created

Create Storage Cluster

Figure 1.7. Create Storage Cluster wizard page



- Select the required storage class.  
You might need to wait a couple of minutes for the storage nodes corresponding to the selected storage class to get populated.
- (Optional) In the Encryption section, set the toggle to Enabled to enable data encryption on the cluster.
- The nodes corresponding to the storage class are displayed based on the storage class that you selected from the drop down list.
- Click **Create**.  
The **Create** button is enabled only when three nodes are selected. A new storage cluster of three volumes will be created with one volume per worker node. The default configuration uses a replication factor of 3.

To expand the capacity of the initial cluster, see [Scaling Storage](#) guide.

## Verification steps

See [Verifying your OpenShift Container Storage installation](#) .

## CHAPTER 2. VERIFYING OPENSIFT CONTAINER STORAGE DEPLOYMENT FOR INTERNAL-ATTACHED DEVICES MODE

Use this section to verify that OpenShift Container Storage is deployed correctly.

### 2.1. VERIFYING THE STATE OF THE PODS

To determine if OpenShift Container storage is deployed successfully, you can verify that the pods are in **Running** state.

#### Procedure

1. Click **Workloads** → **Pods** from the left pane of the OpenShift Web Console.
2. Select **openshift-storage** from the **Project** drop down list.  
For more information on the expected number of pods for each component and how it varies depending on the number of nodes, see [Table 2.1, "Pods corresponding to OpenShift Container storage cluster"](#).
3. Verify that the following pods are in running and completed state by clicking on the **Running** and the **Completed** tabs:

**Table 2.1. Pods corresponding to OpenShift Container storage cluster**


Component	Corresponding pods
OpenShift Container Storage Operator	<b>ocs-operator-*</b> (1 pod on any worker node)
Rook-ceph Operator	<b>rook-ceph-operator-*</b> (1 pod on any worker node)
Multicloud Object Gateway	<ul style="list-style-type: none"> <li>● <b>noobaa-operator-*</b> (1 pod on any worker node)</li> <li>● <b>noobaa-core-*</b> (1 pod on any storage node)</li> <li>● <b>nooba-db-*</b> (1 pod on any storage node)</li> <li>● <b>noobaa-endpoint-*</b> (1 pod on any storage node)</li> </ul>
MON	<b>rook-ceph-mon-*</b> (3 pods distributed across storage nodes)
MGR	<b>rook-ceph-mgr-*</b> (1 pod on any storage node)

Component	Corresponding pods
MDS	<b>rook-ceph-mds-ocs-storagecluster-cephfilesystem-*</b> (2 pods distributed across storage nodes)
RGW	<b>rook-ceph-rgw-ocs-storagecluster-cephobjectstore-*</b> (2 pods distributed across storage nodes)
CSI	<ul style="list-style-type: none"> <li>● <b>cephfs</b> <ul style="list-style-type: none"> <li>○ <b>csi-cephfsplugin-*</b> (1 pod on each worker node)</li> <li>○ <b>csi-cephfsplugin-provisioner-*</b> (2 pods distributed across storage nodes)</li> </ul> </li> <li>● <b>rbd</b> <ul style="list-style-type: none"> <li>○ <b>csi-rbdplugin-*</b> (1 pod on each worker node)</li> <li>○ <b>csi-rbdplugin-provisioner-*</b> (2 pods distributed across storage nodes)</li> </ul> </li> </ul>
rook-ceph-drain-canary	<b>rook-ceph-drain-canary-*</b> (1 pod on each storage node)
rook-ceph-crashcollector	<b>rook-ceph-crashcollector-*</b> (1 pod on each storage node)
OSD	<ul style="list-style-type: none"> <li>● <b>rook-ceph-osd-*</b> (1 pod for each device)</li> <li>● <b>rook-ceph-osd-prepare-ocs-deviceset-*</b> (1 pod for each device)</li> </ul>

## 2.2. VERIFYING THE OPENSIFT CONTAINER STORAGE CLUSTER IS HEALTHY

- Click **Home** → **Overview** from the left pane of the OpenShift Web Console and click **Persistent Storage** tab.
- In the **Status card**, verify that *OCS Cluster* and *Data Resiliency* has a green tick mark as shown in the following image:

Figure 2.1. Health status card in Persistent Storage Overview Dashboard



- In the **Details card**, verify that the cluster information is displayed as follows:

**Service Name**

OpenShift Container Storage

**Cluster Name**

ocs-storagecluster

**Provider**

None

**Mode**

Internal

**Version**


ocs-operator-4.6.0

For more information on the health of OpenShift Container Storage cluster using the persistent storage dashboard, see [Monitoring OpenShift Container Storage](#).

## 2.3. VERIFYING THE MULTICLOUD OBJECT GATEWAY IS HEALTHY

- Click **Home** → **Overview** from the left pane of the OpenShift Web Console and click the **Object Service** tab.
- In the **Status card**, verify that both *Object Service* and *Data Resiliency* are in **Ready** state (green tick).

Figure 2.2. Health status card in Object Service Overview Dashboard



- In the **Details card**, verify that the MCG information is displayed as follows:

**Service Name**

OpenShift Container Storage

**System Name**

Multicloud Object Gateway

**Provider**

None

**Version**

ocs-operator-4.6.0

For more information on the health of the OpenShift Container Storage cluster using the object service dashboard, see [Monitoring OpenShift Container Storage](#).

## 2.4. VERIFYING THAT THE OPENSIFT CONTAINER STORAGE SPECIFIC STORAGE CLASSES EXIST

To verify the storage classes exists in the cluster:

- Click **Storage** → **Storage Classes** from the left pane of the OpenShift Web Console.
- Verify that the following storage classes are created with the OpenShift Container Storage cluster creation:
  - **ocs-storagecluster-ceph-rbd**
  - **ocs-storagecluster-cephfs**
  - **openshift-storage.noobaa.io**
  - **ocs-storagecluster-ceph-rgw**

## CHAPTER 3. UNINSTALLING OPENSIFT CONTAINER STORAGE IN INTERNAL-ATTACHED DEVICES MODE

### 3.1. UNINSTALLING OPENSIFT CONTAINER STORAGE IN INTERNAL-ATTACHED DEVICES MODE

Use the steps in this section to uninstall OpenShift Container Storage.

#### Uninstall Annotations

Annotations on the Storage Cluster are used to change the behavior of the uninstall process. To define the uninstall behavior, the following two annotations have been introduced in the storage cluster:

- **uninstall.ocs.openshift.io/cleanup-policy: delete**
- **uninstall.ocs.openshift.io/mode: graceful**

The below table provides information on the different values that can be used with these annotations:

**Table 3.1. uninstall.ocs.openshift.io uninstall annotations descriptions**

Annotation	Value	Default	Behavior
cleanup-policy	delete	Yes	Rook cleans up the physical drives and the <b>DataDirHostPath</b>
cleanup-policy	retain	No	Rook does <b>not</b> clean up the physical drives and the <b>DataDirHostPath</b>
mode	graceful	Yes	Rook and NooBaa <b>pauses</b> the uninstall process until the PVCs and the OBCs are removed by the administrator/user
mode	forced	No	Rook and NooBaa proceeds with uninstall even if PVCs/OBCs provisioned using Rook and NooBaa exist respectively.

You can change the cleanup policy or the uninstall mode by editing the value of the annotation by using the following commands:

```
$ oc annotate storagecluster ocs-storagecluster uninstall.ocs.openshift.io/cleanup-policy="retain" --
overwrite
storagecluster.ocs.openshift.io/ocs-storagecluster annotated
```



```
$ oc annotate storagecluster ocs-storagecluster uninstall.ocs.openshift.io/mode="forced" --overwrite
storagecluster.ocs.openshift.io/ocs-storagecluster annotated
```

## Prerequisites

- Ensure that the OpenShift Container Storage cluster is in a healthy state. The uninstall process can fail when some of the pods are not terminated successfully due to insufficient resources or nodes. In case the cluster is in an unhealthy state, contact Red Hat Customer Support before uninstalling OpenShift Container Storage.
- Ensure that applications are not consuming persistent volume claims (PVCs) or object bucket claims (OBCs) using the storage classes provided by OpenShift Container Storage.
- If any custom resources (such as custom storage classes, cephblockpools) were created by the admin, they must be deleted by the admin after removing the resources which consumed them.

## Procedure

1. Delete the volume snapshots that are using OpenShift Container Storage.

- a. List the volume snapshots from all the namespaces.

```
$ oc get volumesnapshot --all-namespaces
```

- b. From the output of the previous command, identify and delete the volume snapshots that are using OpenShift Container Storage.

```
$ oc delete volumesnapshot <VOLUME-SNAPSHOT-NAME> -n <NAMESPACE>
```

2. Delete PVCs and OBCs that are using OpenShift Container Storage.

In the default uninstall mode (graceful), the uninstaller waits till all the PVCs and OBCs that use OpenShift Container Storage are deleted.

If you wish to delete the Storage Cluster without deleting the PVCs beforehand, you may set the uninstall mode annotation to "forced" and skip this step. Doing so will result in orphan PVCs and OBCs in the system.

- a. Delete OpenShift Container Platform monitoring stack PVCs using OpenShift Container Storage.  
See [Section 3.2, "Removing monitoring stack from OpenShift Container Storage"](#)
- b. Delete OpenShift Container Platform Registry PVCs using OpenShift Container Storage.  
See [Section 3.3, "Removing OpenShift Container Platform registry from OpenShift Container Storage"](#)
- c. Delete OpenShift Container Platform logging PVCs using OpenShift Container Storage.  
See [Section 3.4, "Removing the cluster logging operator from OpenShift Container Storage"](#)
- d. Delete other PVCs and OBCs provisioned using OpenShift Container Storage.
  - Given below is a sample script to identify the PVCs and OBCs provisioned using OpenShift Container Storage. The script ignores the PVCs that are used internally by OpenShift Container Storage.

```
#!/bin/bash
```

```

RBD_PROVISIONER="openshift-storage.rbd.csi.ceph.com"
CEPHFS_PROVISIONER="openshift-storage.cephfs.csi.ceph.com"
NOOBAA_PROVISIONER="openshift-storage.noobaa.io/obc"
RGW_PROVISIONER="openshift-storage.ceph.rook.io/bucket"

NOOBAA_DB_PVC="noobaa-db"
NOOBAA_BACKINGSTORE_PVC="noobaa-default-backing-store-noobaa-pvc"

# Find all the OCS StorageClasses
OCS_STORAGECLASSES=$(oc get storageclasses | grep -e
"$RBD_PROVISIONER" -e "$CEPHFS_PROVISIONER" -e
"$NOOBAA_PROVISIONER" -e "$RGW_PROVISIONER" | awk '{print $1}')

# List PVCs in each of the StorageClasses
for SC in $OCS_STORAGECLASSES
do
    echo
    "=====
=="
    echo "$SC StorageClass PVCs and OBCs"
    echo
    "=====
=="
    oc get pvc --all-namespaces --no-headers 2>/dev/null | grep $SC | grep -v -e
"$NOOBAA_DB_PVC" -e "$NOOBAA_BACKINGSTORE_PVC"
    oc get obc --all-namespaces --no-headers 2>/dev/null | grep $SC
    echo
done

```

**NOTE**

Omit **RGW\_PROVISIONER** for cloud platforms.

- Delete the OBCs.

```
$ oc delete obc <obc name> -n <project name>
```

- Delete the PVCs.

```
$ oc delete pvc <pvc name> -n <project-name>
```

**NOTE**

Ensure that you have removed any custom backing stores, bucket classes, etc., created in the cluster.

3. Delete the Storage Cluster object and wait for the removal of the associated resources.

```
$ oc delete -n openshift-storage storagecluster --all --wait=true
```

4. Check for cleanup pods if the **uninstall.ocs.openshift.io/cleanup-policy** was set to **delete** (default) and ensure that their status is **Completed**.

```
$ oc get pods -n openshift-storage | grep -i cleanup
NAME                READY STATUS  RESTARTS AGE
cluster-cleanup-job-<xx> 0/1   Completed 0      8m35s
cluster-cleanup-job-<yy> 0/1   Completed 0      8m35s
cluster-cleanup-job-<zz> 0/1   Completed 0      8m35s
```

5. Confirm that the directory **/var/lib/rook** is now empty. This directory will be empty only if the **uninstall.ocs.openshift.io/cleanup-policy** annotation was set to **delete**(default).

```
$ for i in $(oc get node -l cluster.ocs.openshift.io/openshift-storage= -o jsonpath='{.items[*].metadata.name }'); do oc debug node/${i} -- chroot /host ls -l /var/lib/rook; done
```

6. If encryption was enabled at the time of install, remove **dm-crypt** managed **device-mapper** mapping from OSD devices on all the OpenShift Container Storage nodes.
  - a. Create a **debug** pod and **chroot** to the host on the storage node.

```
$ oc debug node <node name>
$ chroot /host
```

- b. Get Device names and make note of the OpenShift Container Storage devices.

```
$ dmsetup ls
ocs-deviceset-0-data-0-57snx-block-dmccrypt (253:1)
```

- c. Remove the mapped device.

```
$ cryptsetup luksClose --debug --verbose ocs-deviceset-0-data-0-57snx-block-dmccrypt
```

If the above command gets stuck due to insufficient privileges, run the following commands:

- Press **CTRL+Z** to exit the above command.
- Find PID of the **cryptsetup** process which was stuck.

```
$ ps
```

Example output:

```
PID  TTY  TIME  CMD
778825  ?    00:00:00 cryptsetup
```

Take a note of the **PID** number to kill. In this example, **PID** is **778825**.

- Terminate the process using **kill** command.

```
$ kill -9 <PID>
```

- Verify that the device name is removed.

```
$ dmsetup ls
```

7. Delete the namespace and wait till the deletion is complete. You will need to switch to another project if **openshift-storage** is the active project.

For example:

```
$ oc project default
$ oc delete project openshift-storage --wait=true --timeout=5m
```

The project is deleted if the following command returns a **NotFound** error.

```
$ oc get project openshift-storage
```



#### NOTE

While uninstalling OpenShift Container Storage, if namespace is not deleted completely and remains in **Terminating** state, perform the steps in [Troubleshooting and deleting remaining resources during Uninstall](#) to identify objects that are blocking the namespace from being terminated.

8. Unlabel the storage nodes.

```
$ oc label nodes --all cluster.ocs.openshift.io/openshift-storage-
$ oc label nodes --all topology.rook.io/rack-
```

9. Remove the OpenShift Container Storage taint if the nodes were tainted.

```
$ oc adm taint nodes --all node.ocs.openshift.io/storage-
```

10. Confirm all PVs provisioned using OpenShift Container Storage are deleted. If there is any PV left in the **Released** state, delete it.

```
$ oc get pv
$ oc delete pv <pv name>
```

11. Delete the Multicloud Object Gateway storageclass.

```
$ oc delete storageclass openshift-storage.noobaa.io --wait=true --timeout=5m
```

12. Remove **CustomResourceDefinitions**.

```
$ oc delete crd backingstores.noobaa.io bucketclasses.noobaa.io
cephblockpools.ceph.rook.io cephclusters.ceph.rook.io cephfilesystems.ceph.rook.io
cephnfses.ceph.rook.io cephobjectstores.ceph.rook.io cephobjectstoreusers.ceph.rook.io
noobaas.noobaa.io ocsinitializations.ocs.openshift.io
storageclusterinitializations.ocs.openshift.io storageclusters.ocs.openshift.io
cephclients.ceph.rook.io cephobjectrealms.ceph.rook.io cephobjectzonegroups.ceph.rook.io
cephobjectzones.ceph.rook.io cephrbdmirrors.ceph.rook.io --wait=true --timeout=5m
```

13. To ensure that OpenShift Container Storage is uninstalled completely, on the OpenShift Container Platform Web Console,

- a. Click **Home → Overview** to access the dashboard.

- b. Verify that the Persistent Storage and Object Service tabs no longer appear next to the **Cluster** tab.

### 3.1.1. Removing local storage operator configurations

Use the instructions in this section only if you have deployed OpenShift Container Storage using local storage devices.



#### NOTE

For OpenShift Container Storage deployments only using **localvolume** resources, go directly to step 8.

#### Procedure

1. Identify the **LocalVolumeSet** and the corresponding **StorageClassName** being used by OpenShift Container Storage.
2. Set the variable `SC` to the **StorageClass** providing the **LocalVolumeSet**.

```
$ export SC="<StorageClassName>"
```

3. Delete the **LocalVolumeSet**.

```
$ oc delete localvolumesets.local.storage.openshift.io <name-of-volumeset> -n openshift-local-storage
```

4. Delete the local storage PVs for the given **StorageClassName**.

```
$ oc get pv | grep $SC | awk '{print $1}' | xargs oc delete pv
```

5. Delete the **StorageClassName**.

```
$ oc delete sc $SC
```

6. Delete the symlinks created by the **LocalVolumeSet**.

```
[[ ! -z $SC ]] && for i in $(oc get node -l cluster.ocs.openshift.io/openshift-storage= -o jsonpath='{.items[*].metadata.name }'); do oc debug node/${i} -- chroot /host rm -rfv /mnt/local-storage/${SC}/; done
```

7. Delete **LocalVolumeDiscovery**.

```
$ oc delete localvolumediscovery.local.storage.openshift.io/auto-discover-devices -n openshift-local-storage
```

8. Removing **LocalVolume** resources (if any).

Use the following steps to remove the **LocalVolume** resources that were used to provision PVs in the current or previous OpenShift Container Storage version. Also, ensure that these resources are not being used by other tenants on the cluster.

For each of the local volumes, do the following:

- a. Identify the **LocalVolume** and the corresponding **StorageClassName** being used by OpenShift Container Storage.
- b. Set the variable LV to the name of the LocalVolume and variable SC to the name of the StorageClass  
For example:

```
$ LV=local-block
$ SC=localblock
```

- c. Delete the local volume resource.

```
$ oc delete localvolume -n local-storage --wait=true $LV
```

- d. Delete the remaining PVs and StorageClasses if they exist.

```
$ oc delete pv -l storage.openshift.com/local-volume-owner-name=${LV} --wait --
timeout=5m
$ oc delete storageclass $SC --wait --timeout=5m
```

- e. Clean up the artifacts from the storage nodes for that resource.

```
$ [[ ! -z $SC ]] && for i in $(oc get node -l cluster.ocs.openshift.io/openshift-storage= -o
jsonpath='{.items[*].metadata.name}'); do oc debug node/${i} -- chroot /host rm -rfv
/mnt/local-storage/${SC}/; done
```

Example output:

```
Starting pod/node-xxx-debug ...
To use host binaries, run `chroot /host`
removed '/mnt/local-storage/localblock/nvme2n1'
removed directory '/mnt/local-storage/localblock'

Removing debug pod ...
Starting pod/node-yyy-debug ...
To use host binaries, run `chroot /host`
removed '/mnt/local-storage/localblock/nvme2n1'
removed directory '/mnt/local-storage/localblock'

Removing debug pod ...
Starting pod/node-zzz-debug ...
To use host binaries, run `chroot /host`
removed '/mnt/local-storage/localblock/nvme2n1'
removed directory '/mnt/local-storage/localblock'

Removing debug pod ...
```

## 3.2. REMOVING MONITORING STACK FROM OPENSIFT CONTAINER STORAGE

Use this section to clean up the monitoring stack from OpenShift Container Storage.

The PVCs that are created as a part of configuring the monitoring stack are in the **openshift-monitoring** namespace.

## Prerequisites

- PVCs are configured to use OpenShift Container Platform monitoring stack. For information, see [configuring monitoring stack](#).

## Procedure

1. List the pods and PVCs that are currently running in the **openshift-monitoring** namespace.

```
$ oc get pod,pvc -n openshift-monitoring
```

NAME	READY	STATUS	RESTARTS	AGE
pod/alertmanager-main-0	3/3	Running	0	8d
pod/alertmanager-main-1	3/3	Running	0	8d
pod/alertmanager-main-2	3/3	Running	0	8d
pod/cluster-monitoring-operator-84457656d-pkrxm	1/1	Running	0	8d
pod/grafana-79ccf6689f-2ll28	2/2	Running	0	8d
pod/kube-state-metrics-7d86fb966-rvd9w	3/3	Running	0	8d
pod/node-exporter-25894	2/2	Running	0	8d
pod/node-exporter-4dsd7	2/2	Running	0	8d
pod/node-exporter-6p4zc	2/2	Running	0	8d
pod/node-exporter-jbjvg	2/2	Running	0	8d
pod/node-exporter-jj4t5	2/2	Running	0	6d18h
pod/node-exporter-k856s	2/2	Running	0	6d18h
pod/node-exporter-rf8gn	2/2	Running	0	8d
pod/node-exporter-rmb5m	2/2	Running	0	6d18h
pod/node-exporter-zj7kx	2/2	Running	0	8d
pod/openshift-state-metrics-59dbd4f654-4clng	3/3	Running	0	8d
pod/prometheus-adapter-5df5865596-k8dzd	1/1	Running	0	7d23h
pod/prometheus-adapter-5df5865596-n2gj9	1/1	Running	0	7d23h
pod/prometheus-k8s-0	6/6	Running	1	8d
pod/prometheus-k8s-1	6/6	Running	1	8d
pod/prometheus-operator-55cfb858c9-c4zd9	1/1	Running	0	6d21h
pod/telemeter-client-78fc8fc97d-2rgfp	3/3	Running	0	8d

NAME	STATUS	VOLUME CAPACITY	ACCESS MODES	STORAGECLASS	AGE
persistentvolumeclaim/my-alertmanager-claim-alertmanager-main-0	Bound	pvc-0d519c4f-15a5-11ea-baa0-026d231574aa	40Gi	RWO	ocs-storagecluster-ceph-rbd 8d
persistentvolumeclaim/my-alertmanager-claim-alertmanager-main-1	Bound	pvc-0d5a9825-15a5-11ea-baa0-026d231574aa	40Gi	RWO	ocs-storagecluster-ceph-rbd 8d
persistentvolumeclaim/my-alertmanager-claim-alertmanager-main-2	Bound	pvc-0d6413dc-15a5-11ea-baa0-026d231574aa	40Gi	RWO	ocs-storagecluster-ceph-rbd 8d
persistentvolumeclaim/my-prometheus-claim-prometheus-k8s-0	Bound	pvc-0b7c19b0-			

```

15a5-11ea-baa0-026d231574aa 40Gi RWO ocs-storagecluster-ceph-rbd 8d
persistentvolumeclaim/my-prometheus-claim-prometheus-k8s-1 Bound pvc-0b8aed3f-
15a5-11ea-baa0-026d231574aa 40Gi RWO ocs-storagecluster-ceph-rbd 8d

```

2. Edit the monitoring **configmap**.

```
$ oc -n openshift-monitoring edit configmap cluster-monitoring-config
```

3. Remove any **config** sections that reference the OpenShift Container Storage storage classes as shown in the following example and save it.

#### Before editing

```

.
.
.
apiVersion: v1
data:
  config.yaml: |
    alertmanagerMain:
      volumeClaimTemplate:
        metadata:
          name: my-alertmanager-claim
        spec:
          resources:
            requests:
              storage: 40Gi
          storageClassName: ocs-storagecluster-ceph-rbd
    prometheusK8s:
      volumeClaimTemplate:
        metadata:
          name: my-prometheus-claim
        spec:
          resources:
            requests:
              storage: 40Gi
          storageClassName: ocs-storagecluster-ceph-rbd
kind: ConfigMap
metadata:
  creationTimestamp: "2019-12-02T07:47:29Z"
  name: cluster-monitoring-config
  namespace: openshift-monitoring
  resourceVersion: "22110"
  selfLink: /api/v1/namespaces/openshift-monitoring/configmaps/cluster-monitoring-config
  uid: fd6d988b-14d7-11ea-84ff-066035b9efa8
.
.
.

```

#### After editing



```

.
.
.
apiVersion: v1
data:
  config.yaml: |
kind: ConfigMap
metadata:
  creationTimestamp: "2019-11-21T13:07:05Z"
  name: cluster-monitoring-config
  namespace: openshift-monitoring
  resourceVersion: "404352"
  selfLink: /api/v1/namespaces/openshift-monitoring/configmaps/cluster-monitoring-config
  uid: d12c796a-0c5f-11ea-9832-063cd735b81c
.
.
.

```

In this example, **alertmanagerMain** and **prometheusK8s** monitoring components are using the OpenShift Container Storage PVCs.

4. Delete relevant PVCs. Make sure you delete all the PVCs that are consuming the storage classes.

```
$ oc delete -n openshift-monitoring pvc <pvc-name> --wait=true --timeout=5m
```

### 3.3. REMOVING OPENSIFT CONTAINER PLATFORM REGISTRY FROM OPENSIFT CONTAINER STORAGE

Use this section to clean up OpenShift Container Platform registry from OpenShift Container Storage. If you want to configure an alternative storage, see [image registry](#)

The PVCs that are created as a part of configuring OpenShift Container Platform registry are in the **openshift-image-registry** namespace.

#### Prerequisites

- The image registry should have been configured to use an OpenShift Container Storage PVC.

#### Procedure

1. Edit the **configs.imageregistry.operator.openshift.io** object and remove the content in the **storage** section.

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

Before editing

```

.
.
.
storage:
  pvc:
    claim: registry-cephfs-rwx-pvc
.
.
.

```

After editing

```

.
.
.
storage:
  emptyDir: {}
.
.
.

```

In this example, the PVC is called **registry-cephfs-rwx-pvc**, which is now safe to delete.

2. Delete the PVC.

```
$ oc delete pvc <pvc-name> -n openshift-image-registry --wait=true --timeout=5m
```

### 3.4. REMOVING THE CLUSTER LOGGING OPERATOR FROM OPENSIFT CONTAINER STORAGE

Use this section to clean up the cluster logging operator from OpenShift Container Storage.

The PVCs that are created as a part of configuring cluster logging operator are in the **openshift-logging** namespace.

#### Prerequisites

- The cluster logging instance should have been configured to use OpenShift Container Storage PVCs.

#### Procedure

1. Remove the **ClusterLogging** instance in the namespace.

```
$ oc delete clusterlogging instance -n openshift-logging --wait=true --timeout=5m
```

The PVCs in the **openshift-logging** namespace are now safe to delete.

2. Delete PVCs.

```
█ $ oc delete pvc <pvc-name> -n openshift-logging --wait=true --timeout=5m
```

## CHAPTER 4. SCALING STORAGE NODES

To scale the storage capacity of OpenShift Container Storage, you can do either of the following:

- **Scale up storage nodes** - Add storage capacity to the existing OpenShift Container Storage worker nodes
- **Scale out storage nodes** - Add new worker nodes containing storage capacity

### 4.1. REQUIREMENTS FOR SCALING STORAGE NODES

Before you proceed to scale the storage nodes, refer to the following sections to understand the node requirements for your specific Red Hat OpenShift Container Storage instance:

- [Platform requirements](#)
- Storage device requirements
  - [Dynamic storage devices](#)
  - [Capacity planning](#)



#### WARNING

Always ensure that you have plenty of storage capacity.

If storage ever fills completely, it is not possible to add capacity or 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.

### 4.2. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT CONTAINER STORAGE NODES ON IBM Z

Use this procedure to add storage capacity and performance to your configured Red Hat OpenShift Container Storage worker nodes.

#### Prerequisites

- A running OpenShift Container Storage Platform.
- Administrative privileges on the OpenShift Web Console.
- To scale using a storage class other than the one provisioned during deployment, first define an additional storage class. See [Creating a storage class](#) for details.

## Procedure

1. Add additional hardware resources with zFCP disks

- a. List all the disks with the following command.

```
$ lszdev
```

Example output:

```
TYPE      ID                                ON PERS NAMES
zfcplib  0.0.8204                            yes yes
zfcplib  0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no  sda sg0
zfcplib  0.0.8204:0x500407630c0b50a4:0x3002b03000000000 yes yes  sdb sg1
qeth     0.0.bdd0:0.0.bdd1:0.0.bdd2          yes no  encbdd0
generic-ccw 0.0.0009                            yes no
```

A SCSI disk is represented as a **zfcplib** with the structure **<device-id>:<wwpn>:<lun-id>** in the ID section. The first disk is used for the operating system. The device id for the new disk can be the same.

- b. Append a new SCSI disk with the following command.

```
$ chzdev -e 0.0.8204:0x400506630b1b50a4:0x3001301a00000000
```




### NOTE

The device ID for the new disk must be the same as the disk to be replaced. The new disk is identified with its WWPN and LUN ID.


- c. List all the FCP devices to verify the new disk is configured.

```
$ lsdev zfcplib
TYPE      ID                                ON PERS NAMES
zfcplib  0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no  sda sg0
zfcplib  0.0.8204:0x500507630b1b50a4:0x4001302a00000000 yes yes  sdb sg1
zfcplib  0.0.8204:0x400506630b1b50a4:0x3001301a00000000 yes yes  sdc sg2
```

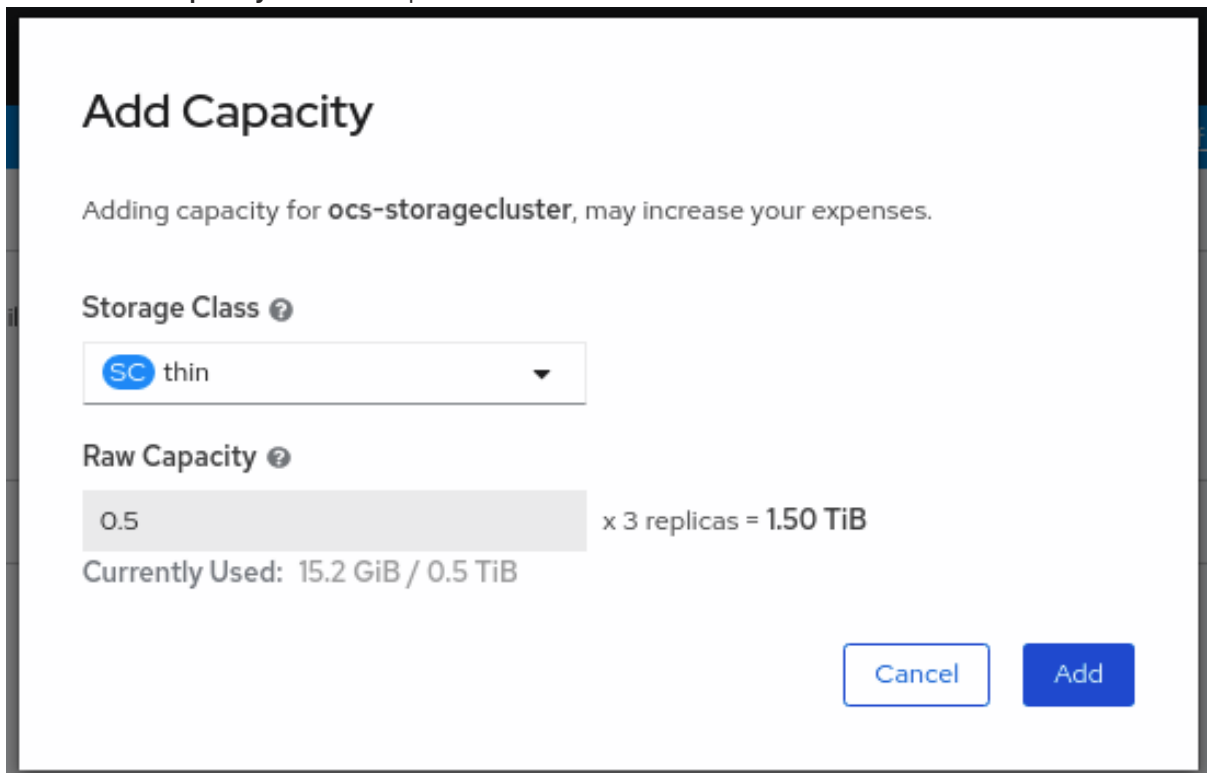
2. Navigate to the OpenShift Web Console.
3. Click **Operators** on the left navigation bar.
4. Select **Installed Operators**.
5. In the window, click **OpenShift Container Storage Operator**:



6. In the top navigation bar, scroll right and click **Storage Cluster** tab.



7. Click ( : ) next to the visible list to extend the options menu.
8. Select **Add Capacity** from the options menu.

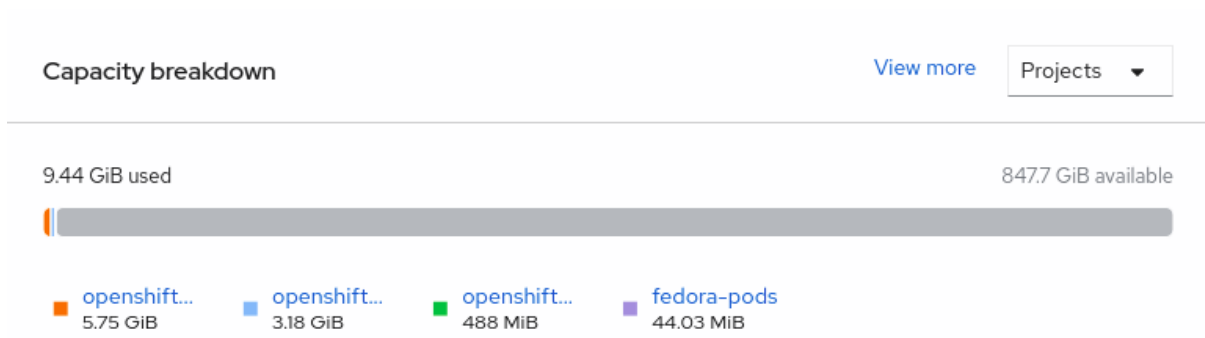


The **Raw Capacity** field shows the size set during storage class creation. The total amount of storage consumed is three times this amount, because OpenShift Container Storage uses a replica count of 3.


9. Click **Add** and wait for the cluster state to change to **Ready**.

## Verification steps

1. Navigate to **Overview** → **Persistent Storage** tab, then check the **Capacity breakdown** card.



- Note that the capacity increases based on your selections.



### IMPORTANT

Cluster reduction is not currently supported, regardless of whether reduction would be done by removing nodes or OSDs.

## 4.3. SCALING OUT STORAGE CAPACITY BY ADDING NEW NODES

To scale out storage capacity, you need to perform the following:

- Add a new node to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs, which is the increment of 3 OSDs of the capacity selected during initial configuration.
- Verify that the new node is added successfully
- Scale up the storage capacity after the node is added


### 4.3.1. Adding a node on IBM Z

#### Prerequisites

- You must be logged into OpenShift Container Platform (RHOC) cluster.

#### Procedure

- Navigate to **Compute** → **Machine Sets**.
- On the machine set where you want to add nodes, select **Edit Machine Count**
- Add the amount of nodes, and click **Save**.
- Click **Compute** → **Nodes** and confirm if the new node is in **Ready** state.
- Apply the OpenShift Container Storage label to the new node.
  - For the new node, **Action menu ( ⋮ )** → **Edit Labels**.
  - Add `cluster.ocs.openshift.io/openshift-storage` and click **Save**.



## NOTE

It is recommended to add 3 nodes each in different zones. You must add 3 nodes and perform this procedure for all of them.

### Verification steps

- To verify that the new node is added, see [Verifying the addition of a new node](#).

### 4.3.2. Verifying the addition of a new node

1. Execute the following command and verify that the new node is present in the output:

```
$ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
```

2. Click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
  - **csi-cephfsplugin-\***
  - **csi-rbdplugin-\***

### 4.3.3. Scaling up storage capacity

After you add a new node to OpenShift Container Storage, you must scale up the storage capacity as described in [Scaling up storage by adding capacity](#).



## CHAPTER 5. REPLACING STORAGE NODES

You can choose one of the following procedures to replace storage nodes:

- [Replacing operational nodes on IBM Z infrastructure](#)
- [Replacing failed nodes on IBM Z infrastructure](#)

### 5.1. REPLACING OPERATIONAL NODES ON IBM Z

Use this procedure to replace an operational node on IBM Z.

#### Procedure

1. Log in to OpenShift Web Console.
2. Click **Compute** → **Nodes**.
3. Identify the node that needs to be replaced. Take a note of its **Machine Name**.
4. Mark the node as unschedulable using the following command:

```
$ oc adm cordon <node_name>
```

5. Drain the node using the following command:


```
$ oc adm drain <node_name> --force --delete-local-data --ignore-daemonsets
```



#### IMPORTANT

This activity may take at least 5-10 minutes. Ceph errors generated during this period are temporary and are automatically resolved when the new node is labeled and functional.

6. Click **Compute** → **Machines**. Search for the required machine.
7. Besides the required machine, click the **Action menu ( ⋮ )** → **Delete Machine**.
8. Click **Delete** to confirm the machine deletion. A new machine is automatically created.
9. Wait for the new machine to start and transition into **Running** state.



#### IMPORTANT

This activity may take at least 5-10 minutes.

10. Click **Compute** → **Nodes**, confirm if the new node is in **Ready** state.
11. Apply the OpenShift Container Storage label to the new node using any one of the following:

#### From User interface

- a. For the new node, click **Action Menu ( ⋮ )** → **Edit Labels**

- b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

#### From command line interface

- Execute the following command to apply the OpenShift Container Storage label to the new node:

```
$ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
```

#### Verification steps

1. Execute the following command and verify that the new node is present in the output:

```
$ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= |cut -d' ' -f1
```

2. Click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:

- **csi-cephfsplugin-\***
- **csi-rbdplugin-\***

3. Verify that all other required OpenShift Container Storage pods are in **Running** state.
4. Verify that new OSD pods are running on the replacement node.

```
$ oc get pods -o wide -n openshift-storage| egrep -i new-node-name | egrep osd
```

5. (Optional) If data encryption is enabled on the cluster, verify that the new OSD devices are encrypted.

- a. Create a debug pod and open a **chroot** environment for the host.

```
$ oc debug node new-node-name
$ chroot /host
```

- b. Verify the devices are encrypted.

```
$ dmsetup ls | grep ocs-deviceset
ocs-deviceset-0-data-0-57snx-block-dmccrypt (253:1)
```

```
$ lsblk | grep ocs-deviceset
`-ocs-deviceset-0-data-0-57snx-block-dmccrypt 253:1 0 512G 0 crypt
```

6. If verification steps fail, [contact Red Hat Support](#).

## 5.2. REPLACING FAILED NODES ON IBM Z

Perform this procedure to replace a failed node which is not operational on IBM Z for OpenShift Container Storage.

### Procedure

1. Log in to OpenShift Web Console and click **Compute** → **Nodes**.
2. Identify the faulty node and click on its **Machine Name**.
3. Click **Actions** → **Edit Annotations**, and click **Add More**.
4. Add **machine.openshift.io/exclude-node-draining** and click **Save**.
5. Click **Actions** → **Delete Machine**, and click **Delete**.
6. A new machine is automatically created, wait for new machine to start.



### IMPORTANT

This activity may take at least 5–10 minutes. Ceph errors generated during this period are temporary and are automatically resolved when the new node is labeled and functional.

7. Click **Compute** → **Nodes**, confirm if the new node is in **Ready** state.
8. Apply the OpenShift Container Storage label to the new node using any one of the following:

#### From the web user interface

- a. For the new node, click **Action Menu ( ⋮ )** → **Edit Labels**
- b. Add **cluster.ocs.openshift.io/openshift-storage** and click **Save**.

#### From the command line interface

- Execute the following command to apply the OpenShift Container Storage label to the new node:

```
$ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
```

9. Execute the following command and verify that the new node is present in the output:

```
$ oc get nodes --show-labels | grep cluster.ocs.openshift.io/openshift-storage= | cut -d' ' -f1
```

10. Click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:

- **csi-cephfsplugin-\***
- **csi-rbdplugin-\***

11. Verify that all other required OpenShift Container Storage pods are in **Running** state.
12. Verify that new OSD pods are running on the replacement node.

```
$ oc get pods -o wide -n openshift-storage | egrep -i new-node-name | egrep osd
```

13. (Optional) If data encryption is enabled on the cluster, verify that the new OSD devices are encrypted.

- a. Create a debug pod and open a **chroot** environment for the host.

```
$ oc debug node new-node-name  
$ chroot /host
```

- b. Verify the devices are encrypted.

```
$ dmsetup ls | grep ocs-deviceset  
ocs-deviceset-0-data-0-57snx-block-dmccrypt (253:1)
```

```
$ lsblk | grep ocs-deviceset  
`-ocs-deviceset-0-data-0-57snx-block-dmccrypt 253:1 0 512G 0 crypt
```

14. If verification steps fail, [contact Red Hat Support](#).

## CHAPTER 6. REPLACING STORAGE DEVICES

### 6.1. REPLACING OPERATIONAL OR FAILED STORAGE DEVICES ON IBM Z

You can replace operational or failed storage devices on IBM Z with new SCSI disks.

IBM Z supports SCSI FCP disk logical units (SCSI disks) as persistent storage devices from external disk storage. A SCSI disk can be identified by using its FCP Device number, two target worldwide port names (WWPN1 and WWPN2), and the logical unit number (LUN). For more information, see [https://www.ibm.com/support/knowledgecenter/SSB27U\\_6.4.0/com.ibm.zvm.v640.hcpa5/scsiover.html](https://www.ibm.com/support/knowledgecenter/SSB27U_6.4.0/com.ibm.zvm.v640.hcpa5/scsiover.html)

#### Procedure

1. List all the disks with the following command.

```
$ lszdev
```

Example output:

```
TYPE      ID
zfcplib  0.0.8204                yes yes
zfcplib  0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no  sda sg0
zfcplib  0.0.8204:0x500407630c0b50a4:0x3002b03000000000 yes yes  sdb sg1
qeth     0.0.bdd0:0.0.bdd1:0.0.bdd2                yes no  encbdd0
generic-ccw 0.0.0009                yes no
```

A SCSI disk is represented as a **zfcplib** with the structure **<device-id>:<wwpn>:<lun-id>** in the **ID** section. The first disk is used for the operating system. If one storage device fails, it can be replaced with a new disk.

2. Remove the disk.

Run the following command on the disk, replacing **scsi-id** with the SCSI disk identifier of the disk to be replaced.

```
$ chzdev -d scsi-id
```

For example, the following command removes one disk with the device ID **0.0.8204**, the WWPN **0x500507630a0b50a4**, and the LUN **0x4002403000000000** with the following command:

```
$ chzdev -d 0.0.8204:0x500407630c0b50a4:0x3002b03000000000
```

3. Append a new SCSI disk with the following command:

```
$ chzdev -e 0.0.8204:0x500507630b1b50a4:0x4001302a00000000
```



#### NOTE

The device ID for the new disk must be the same as the disk to be replaced. The new disk is identified with its WWPN and LUN ID.

4. List all the FCP devices to verify the new disk is configured.

```
$ lsdev zfcplun
TYPE      ID                                     ON PERS NAMES
zfcplun   0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no  sda sg0
zfcplun   0.0.8204:0x500507630b1b50a4:0x4001302a00000000 yes yes  sdb sg1
```