Red Hat OpenShift Container Storage 4.6

Scaling storage

Horizontal and vertical scaling options

Red Hat OpenShift Container Storage 4.6 Scaling storage

Horizontal and vertical scaling options

Legal Notice

Abstract

This document explains scaling options for Red Hat OpenShift Container Storage.

Table of Contents

PREFACE

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

- **Scale up storage nodes** - Add storage capacity to the existing Red Hat OpenShift Container Storage worker nodes

- **Scale out storage nodes** - Add new worker nodes containing storage capacity

For scaling your storage in external mode, see Red Hat Ceph Storage documentation .

CHAPTER 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

  - Local storage devices

  - Capacity planning

> **IMPORTANT**
>
> 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.

1.1. SUPPORTED DEPLOYMENTS FOR RED HAT OPENSHIFT CONTAINER STORAGE

- User-provisioned infrastructure:

  - Amazon Web Services (AWS)

  - VMware

  - Bare metal

  - IBM Power Systems

  - IBM Z or LinuxONE

- Installer-provisioned infrastructure:

  - Amazon Web Services (AWS)

  - Microsoft Azure

CHAPTER 2. SCALING UP STORAGE CAPACITY

Depending on the type of your deployment, you can choose one of the following procedures to scale up storage capacity.

- For AWS, VMware, or Azure infrastructures using dynamic or automated provisioning of storage devices, see Section 2.2, "Scaling up storage by adding capacity to your OpenShift Container Storage nodes"

- For bare metal, Amazon EC2 I3, or VMware infrastructures using local storage devices, see Section 2.3, "Scaling up storage by adding capacity to your OpenShift Container Storage nodes using local storage devices"

- For IBM Z or LinuxONE infrastructures using local storage devices, see Section 2.5, "Scaling up storage by adding capacity to your OpenShift Container Storage nodes on IBM Z or LinuxONE infrastructure"

- For IBM Power Systems using local storage devices, see Section 2.4, "Scaling up storage by adding capacity to your OpenShift Container Storage nodes on IBM Power Systems infrastructure using local storage devices"

If you want to scale using a storage class other than the one provisioned during deployment, you must also define an additional storage class before you scale. See Creating a storage class for details.

> **NOTE**
>
> OpenShift Container Storage does not support heterogeneous OSD sizes.

2.1. CREATING A STORAGE CLASS

You can define a new storage class to dynamically provision storage from an existing provider.

> **IMPORTANT**
>
> Using storage classes other than the default for your provider 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 .

Prerequisites

- Administrator access to OpenShift web console.

Procedure

1. Log in to OpenShift Web Console.

2. Click **Storage** → **Storage Classes**.

3. Click **Create Storage Class**

4. Enter the storage class **Name** and **Description**.

5. Select the required **Reclaim Policy** and **Provisioner**.

6. Click **Create** to create the Storage Class.

Verification steps

- Click **Storage → Storage Classes** and verify that you can see the new storage class.

2.2. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSHIFT CONTAINER STORAGE NODES

Use this procedure to add storage capacity and performance to your configured Red Hat OpenShift Container Storage worker nodes on the following infrastructures:

- AWS

- VMware vSphere

- Microsoft Azure

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. Log in to the OpenShift Web Console.

2. Click on **Operators → Installed Operators**.

3. Click **OpenShift Container Storage** Operator.
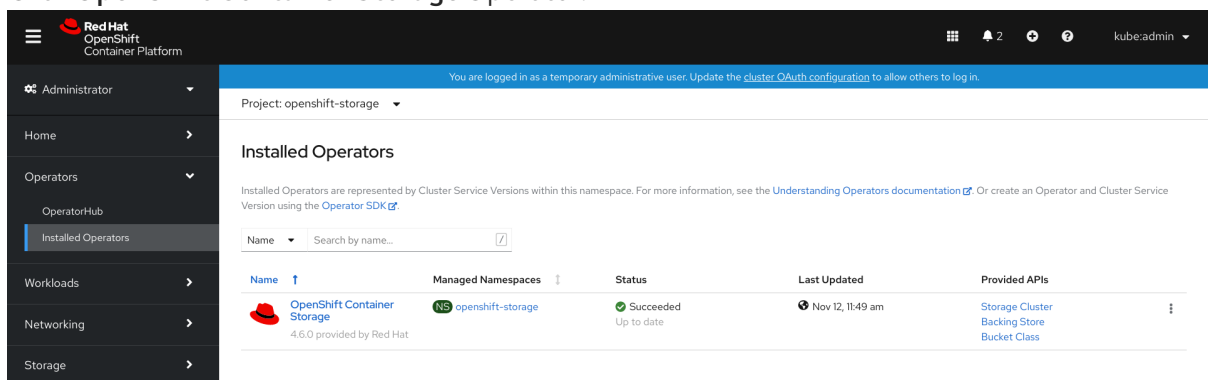
4. Click **Storage Cluster** tab.
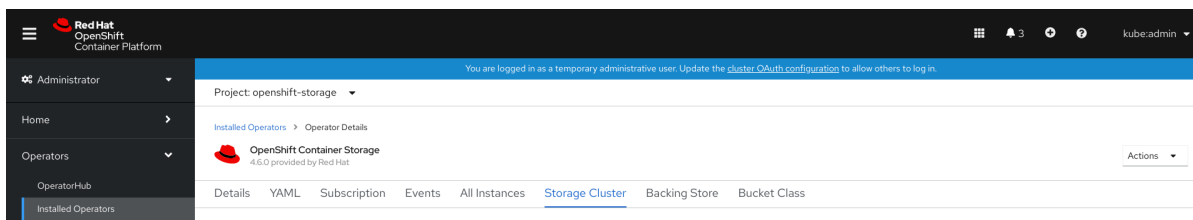
5. The visible list should have only one item. Click (⋮) on the far right to extend the options menu.

6. Select **Add Capacity** from the options menu.
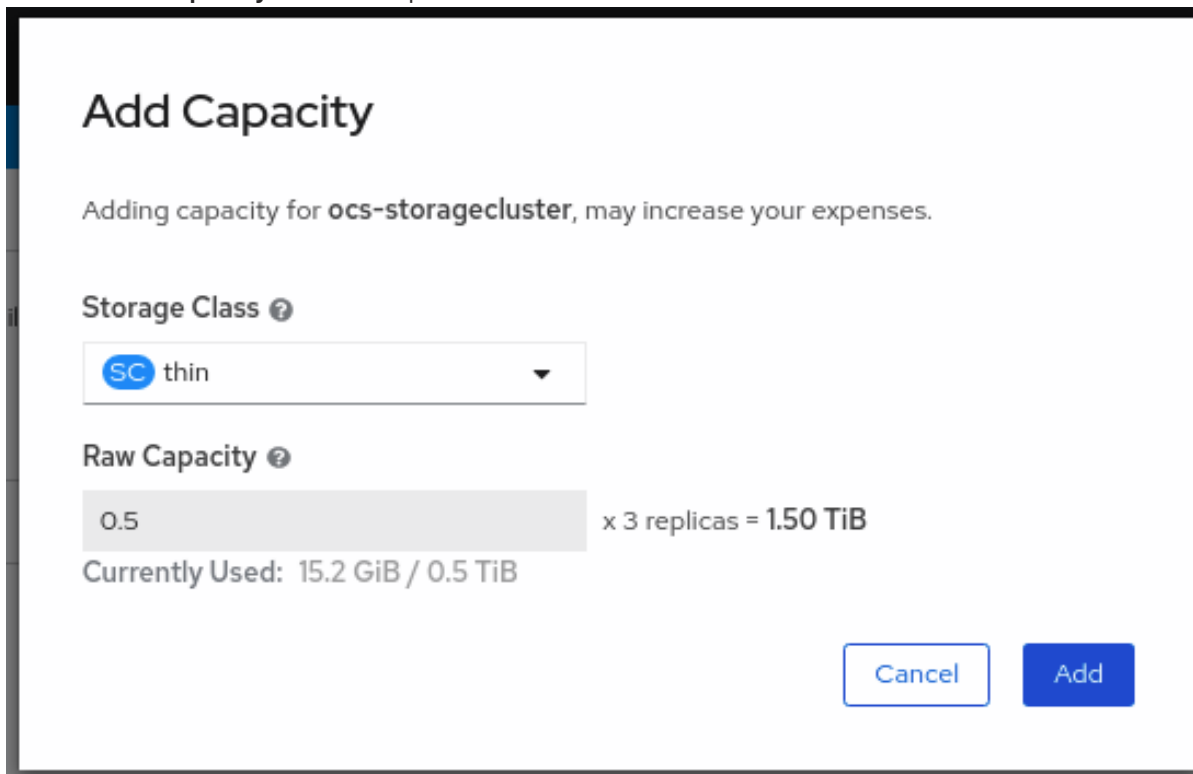
7. Select a storage class.
    Set the storage class to **gp2** on AWS, **thin** on VMware, or **managed_premium** on Microsoft Azure if you are using the default storage class generated during deployment. If you have created other storage classes, select whichever is appropriate.

    IMPORTANT

    Using storage classes other than the default for your provider 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 .

    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.

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

Verification steps

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

Note that the capacity increases based on your selections.

- Verify that the new OSDs and their corresponding new PVCs are created.

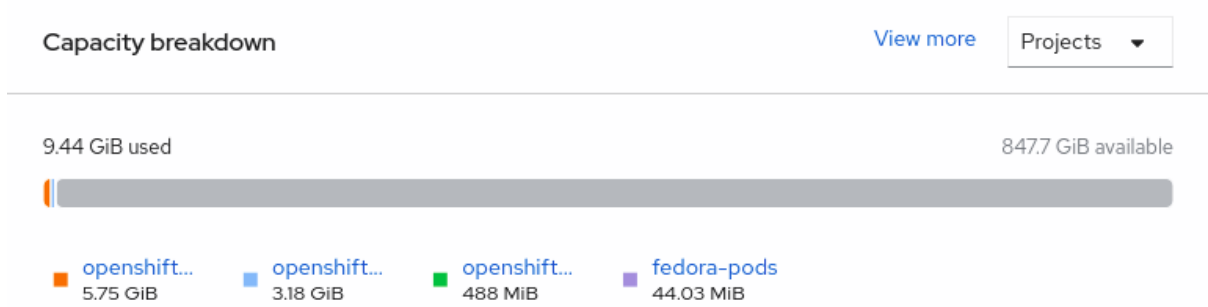 - To view the state of the newly created OSDs:

    a. Click **Workloads → Pods** from the OpenShift Web Console.

    b. Select **openshift-storage** from the **Project** drop-down list.

  - To view the state of the PVCs:

    a. Click **Storage → Persistent Volume Claims** from the OpenShift Web Console.

    b. Select **openshift-storage** from the **Project** drop-down list.

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

  a. Identify the node(s) where the new OSD pod(s) are running.

  ```
  $ oc get -o=custom-columns=NODE:.spec.nodeName pod/<OSD pod name>
  ```

  For example:

  ```
  oc get -o=custom-columns=NODE:.spec.nodeName pod/rook-ceph-osd-0-544db49d7f-qrgqm
  ```

  b. For each of the nodes identified in previous step, do the following:

    i. Create a debug pod and open a chroot environment for the selected host(s).

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

    ii. Run "lsblk" and check for the "crypt" keyword beside the **ocs-deviceset** name(s)

    ```
    $ lsblk
    ```

IMPORTANT

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

2.3. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSHIFT CONTAINER STORAGE NODES USING LOCAL STORAGE DEVICES

Use this procedure to add storage capacity (additional storage devices) to your configured local storage based OpenShift Container Storage worker nodes on bare metal and VMware infrastructures.

IMPORTANT

Scaling up storage on Amazon EC2 I3 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.

NOTE

For Amazon EC2 I3 infrastructure, adding nodes is the only option for adding capacity, as deployment is done using both the available NVMe devices.

Prerequisites

- You must be logged into the OpenShift Container Platform cluster.

- You must have installed local storage operator. Use any of the following procedures applicable to your infrastructure:

    - Installing Local Storage Operator on bare metal

    - Installing Local Storage Operator on vSphere cluster

- If you have 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.

- 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.

- You must have three OpenShift Container Platform worker nodes with the same storage type and size attached to each node (for example, 2TB NVMe drive) as the original OpenShift Container Storage StorageCluster was created with.

Procedure

To add capacity, you can either use a storage class that you provisioned during the deployment or any other storage class that matches the filter.

1. On the OpenShift web console, click on **Operators → Installed Operators**.

2. Click **OpenShift Container Storage** Operator.

3. Click **Storage Cluster** tab.
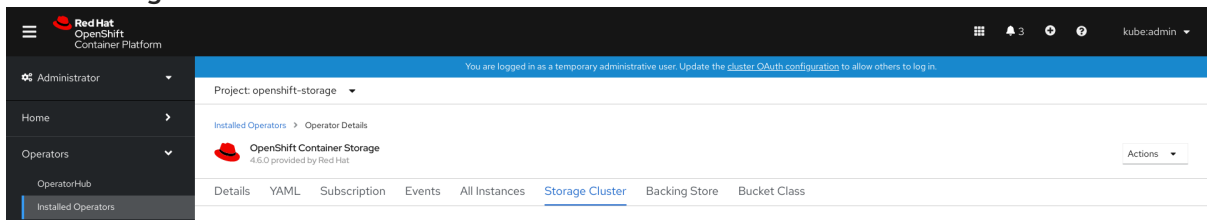
4. The visible list should have only one item. Click (⋮) on the far right to extend the options menu.

5. Select **Add Capacity** from the options menu.

6. Select the **Storage Class** for which you added disks or the new storage class depending on your requirement. Available Capacity displayed is based on the local disks available in storage class.

7. Click **Add**.
   You might need to wait a couple of minutes for the storage cluster to reach **Ready** state.

Verification steps

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

Note that the capacity increases based on your selections.

- Verify that the new OSDs and their corresponding new PVCs are created.

  ○ To view the state of the newly created OSDs:

    a. Click **Workloads → Pods** from the OpenShift Web Console.

    b. Select **openshift-storage** from the **Project** drop-down list.

  ○ To view the state of the PVCs:

    a. Click **Storage → Persistent Volume Claims** from the OpenShift Web Console.

    b. Select **openshift-storage** from the **Project** drop-down list.

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

  a. Identify the node(s) where the new OSD pod(s) are running.

    ```
    $ oc get -o=custom-columns=NODE:.spec.nodeName pod/<OSD pod name>
    ```

    For example:

    ```
    oc get -o=custom-columns=NODE:.spec.nodeName pod/rook-ceph-osd-0-544db49d7f-qrgqm
    ```

  b. For each of the nodes identified in previous step, do the following:

    i. Create a debug pod and open a chroot environment for the selected host(s).

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

    ii. Run "lsblk" and check for the "crypt" keyword beside the **ocs-deviceset** name(s)

    ```
    $ lsblk
    ```

> **IMPORTANT**
>
> OpenShift Container Storage does not support cluster reduction either by reducing OSDs or reducing nodes.

2.4. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSHIFT CONTAINER STORAGE NODES ON IBM POWER SYSTEMS INFRASTRUCTURE USING LOCAL STORAGE DEVICES

Use this procedure to add storage capacity (additional storage devices) to your configured local storage based OpenShift Container Storage worker nodes on IBM Power Systems infrastructures.

Prerequisites

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

- You must have installed local storage operator. Use the following procedures, see

  - Installing Local Storage Operator on IBM Power Systems

- You must have three OpenShift Container Platform worker nodes with the same storage type and size attached to each node (for example, 0.5TB SSD) as the original OpenShift Container Storage StorageCluster was created with.

Procedure

1. To add storage capacity to OpenShift Container Platform nodes with OpenShift Container Storage installed, you need to

   a. Add a new disk that is, minimum of one device per worker node in your Openshift Container Platform (RHOCP) cluster.

   b. Check if the new disk is added to the node by running lsblk inside node.

      ```
      $ oc debug node/worker-0
      $lsblk
      ```

      Example output:

      ```
      Creating debug namespace/openshift-debug-node-ggrqr ...
      Starting pod/worker-2-debug ...
      To use host binaries, run `chroot /host`
      Pod IP: 192.168.88.23
      If you don't see a command prompt, try pressing enter.
      sh-4.4# chroot /host
      sh-4.4# lsblk
      NAME                    MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
      loop0                     7:0    0  256G  0 loop
      vda                     252:0    0   40G  0 disk
      |-vda1                   252:1    0    4M  0 part
      |-vda2                   252:2    0  384M  0 part /boot
      `-vda4                   252:4    0 39.6G  0 part
        `-coreos-luks-root-nocrypt 253:0    0 39.6G  0 dm   /sysroot
      vdb                     252:16   0  512B  1 disk
      vdc                     252:32   0  256G  0 disk
      vdd                     252:48   0  256G  0 disk
      sh-4.4#
      sh-4.4#
      Removing debug pod ...
      Removing debug namespace/openshift-debug-node-ggrqr ...
      ```

    c.   Newly added disk will automatically gets discovered by LocalVolumeSet.

2. Display the newly created PVs with **storageclass** name used in **localVolumeSet** CR.

```
$ oc get pv | grep localblock | grep Available
```

Example output:

```
local-pv-290020c2   256Gi   RWO     Delete  Available   localblock      2m35s
local-pv-7702952c   256Gi   RWO     Delete  Available   localblock      2m27s
local-pv-a7a567d    256Gi   RWO     Delete  Available   localblock      2m22s
...
```

There are three more available PVs of same size which will be used for new OSDs.

3. Navigate to the OpenShift Web Console.

4. Click on **Operators** on the left navigation bar.

5. Select **Installed Operators**.

6. In the window, click **OpenShift Container Storage** Operator:
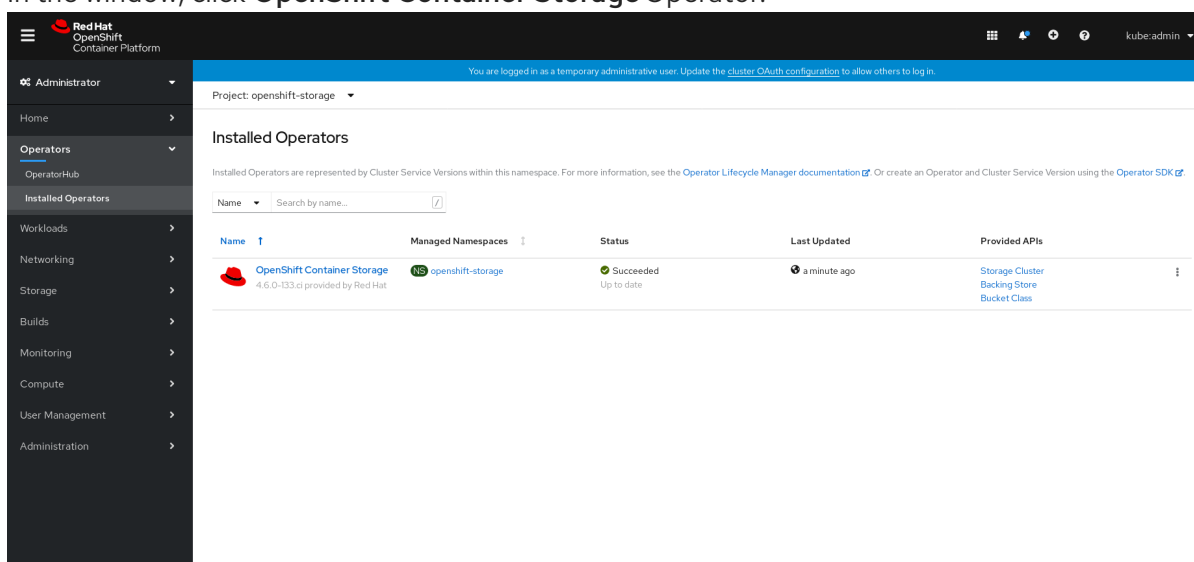
7. In the top navigation bar, scroll right and click **Storage Cluster** tab.
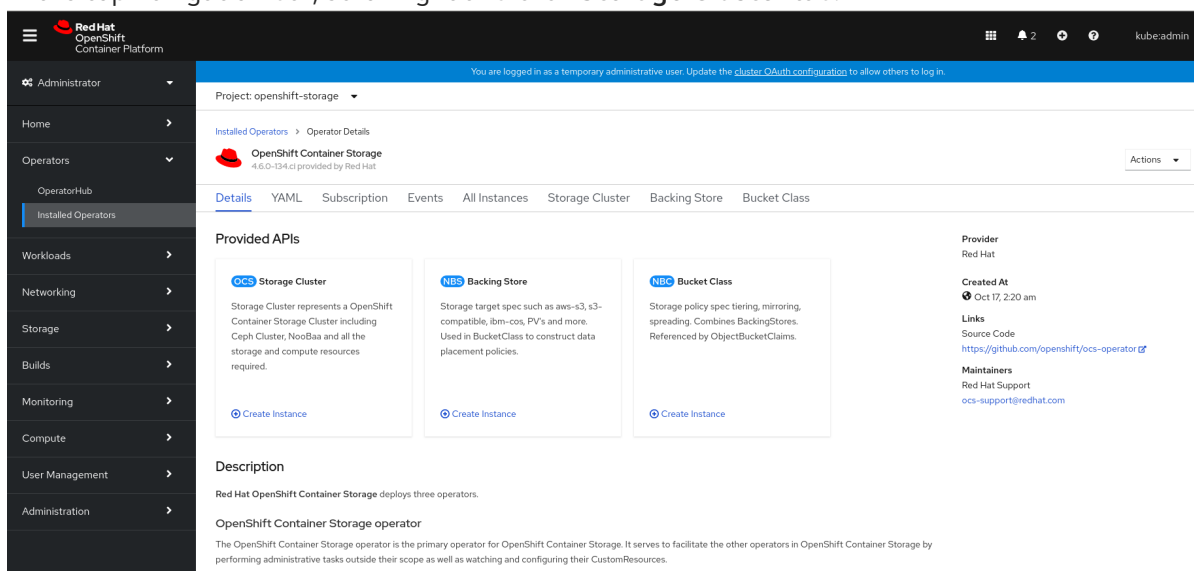
8. The visible list should have only one item. Click (⋮) on the far right to extend the options menu.

9. Select **Add Capacity** from the options menu.

Add Capacity

Adding capacity for **ocs-storagecluster**, may increase your expenses.

Storage Class ⓘ

SC localblock ▼

Available capacity: 2.73 TiB / 3 replicas

Cancel   Add

From this dialog box, set the **Storage Class** name to the name used in the **localVolumeset** CR. Available Capacity displayed is based on the local disks available in storage class.

10. Once you are done with your setting, click **Add**. You might need to wait a couple of minutes for the storage cluster to reach **Ready** state.

11. Verify that the new OSDs and their corresponding new PVCs are created.

```
$ oc get -n openshift-storage pods -l app=rook-ceph-osd
```

Example output:

```
NAME                          READY   STATUS   RESTARTS   AGE
rook-ceph-osd-0-6f8655ff7b-gj226   1/1     Running   0          1h
rook-ceph-osd-1-6c66d77f65-cfgfq   1/1     Running   0          1h
rook-ceph-osd-2-69f6b4c597-mtsdv   1/1     Running   0          1h
rook-ceph-osd-3-c784bdbd4-w4cmj    1/1     Running   0          5m
rook-ceph-osd-4-6d99845f5b-k7f8n   1/1     Running   0          5m
rook-ceph-osd-5-fdd9897c9-r9mgb    1/1     Running   0          5m
```

In the above example, osd-3, osd-4, and osd-5 are the newly added pods to the OpenShift Container Storage cluster.

```
$ oc get pvc -n openshift-storage |grep localblock
```

Example output:

```
ocs-deviceset-localblock-0-data-0-sfsgf   Bound   local-pv-8137c873   256Gi   RWO
localblock  1h
ocs-deviceset-localblock-0-data-1-qhs9m   Bound   local-pv-290020c2   256Gi   RWO
localblock  10m
ocs-deviceset-localblock-1-data-0-499r2   Bound   local-pv-ec7f2b80   256Gi   RWO
localblock  1h
```

> ocs-deviceset-localblock-1-data-1-p9rth   Bound   local-pv-a7a567d   256Gi   RWO
> localblock 10m
> ocs-deviceset-localblock-2-data-0-8pzjr   Bound   local-pv-1e31f771   256Gi   RWO
> localblock 1h
> ocs-deviceset-localblock-2-data-1-7zwwn   Bound   local-pv-7702952c   256Gi   RWO
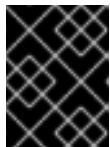> localblock 10m

In the above example, we see three new PVCs are created.

Verification steps

1. Navigate to **Overview → Persistent Storage** tab, then check the **Capacity breakdown** card.
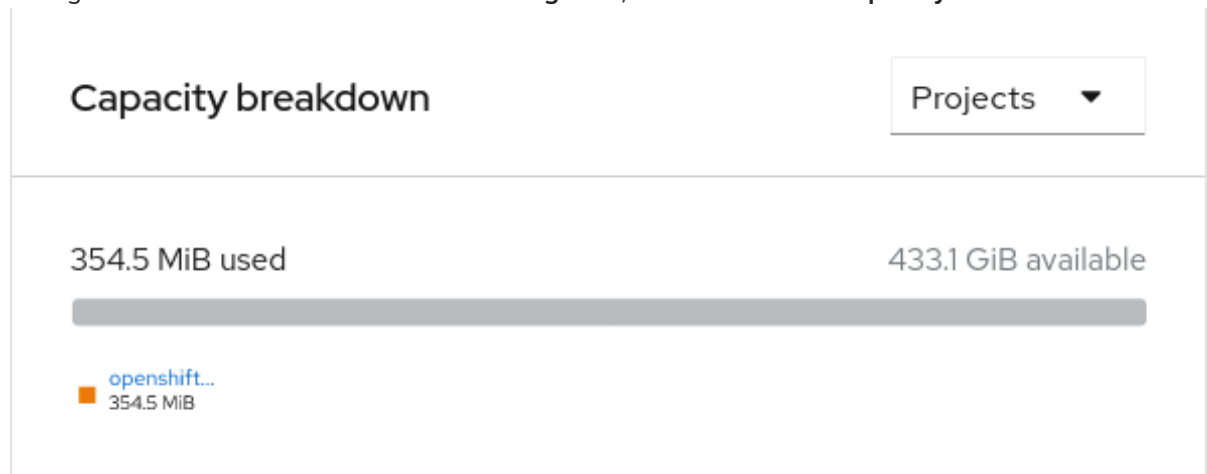
Note that the capacity increases based on your selections.

> **IMPORTANT**
>
> OpenShift Container Storage does not support cluster reduction either by
> reducing OSDs or reducing nodes.

2.5. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSHIFT CONTAINER STORAGE NODES ON IBM Z OR LINUXONE INFRASTRUCTURE

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
zfcp-host   0.0.8204                                yes yes
zfcp-lun    0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no   sda sg0
zfcp-lun    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 **zfcp-lun** 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.

```
$ lszdev zfcp-lun
TYPE        ID                                      ON  PERS  NAMES
zfcp-lun    0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no   sda sg0
zfcp-lun    0.0.8204:0x500507630b1b50a4:0x4001302a00000000 yes yes  sdb sg1
zfcp-lun    0.0.8204:0x400506630b1b50a4:0x3001301a00000000 yes yes  sdc sg2
```
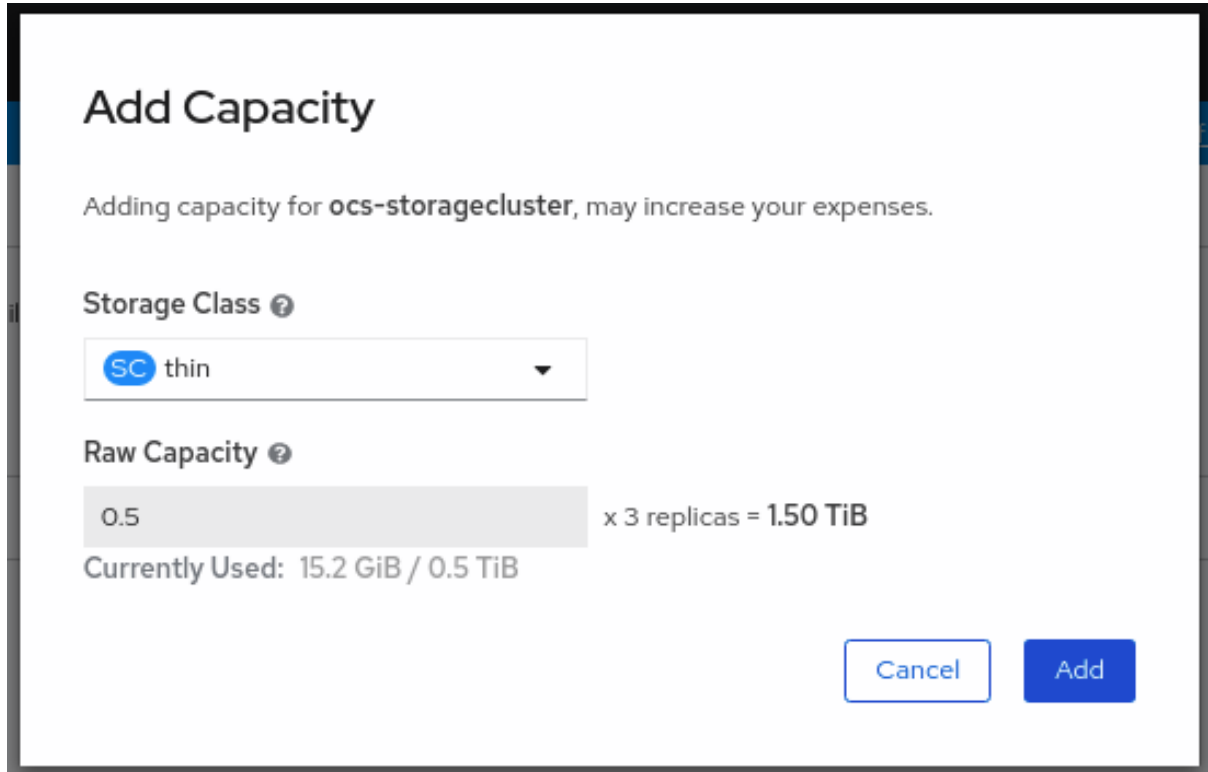
2. Navigate to the OpenShift Web Console.

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.
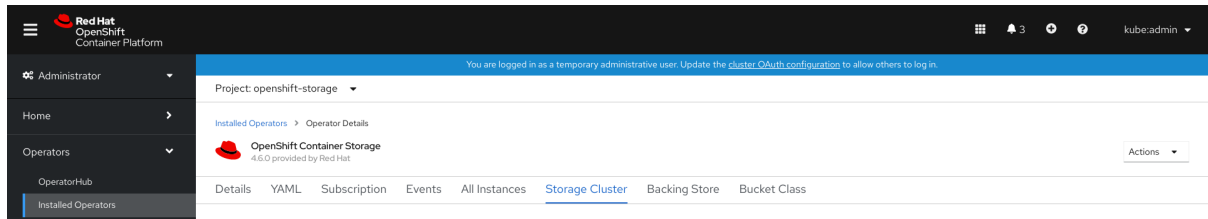
7. Click (⋮) next to the visible list to extend the options menu.

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.
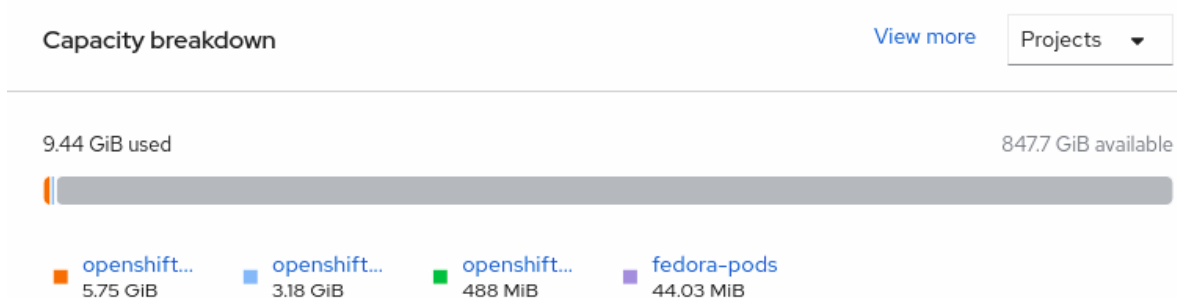
2. Note that the capacity increases based on your selections.

IMPORTANT

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

CHAPTER 3. SCALING OUT STORAGE CAPACITY

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

- Add a new node

- Verify that the new node is added successfully

- Scale up the storage capacity
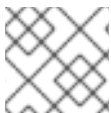
> **NOTE**
>
> OpenShift Container Storage does not support heterogeneous OSD sizes.

3.1. ADDING A NODE

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs, which is increment of 3 OSDs of the capacity selected during initial configuration.

Depending on the type of your deployment, you can choose one of the following procedures to add a storage node:

- For AWS or Azure installer-provisioned infrastructures, see Adding a node on an AWS installer-provisioned infrastructure

- For AWS or VMware user-provisioned infrastructure, see Adding a node on an AWS or a VMware user-provisioned infrastructure

- For bare metal, IBM Power Systems, IBM Z or LinuxONE, Amazon EC2 I3, or VMware infrastructures, see Adding a node using a local storage device

3.1.1. Adding a node on an installer-provisioned infrastructure

Use this proecdure to add a node on an AWS or Azure installer provisioned infrastructure.

Prerequisites

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

Procedure

1. Navigate to **Compute → Machine Sets**.

2. On the machine set where you want to add nodes, select **Edit Machine Count**

3. Add the amount of nodes, and click **Save**.

4. Click **Compute → Nodes** and confirm if the new node is in **Ready** state.

5. Apply the OpenShift Container Storage label to the new node.

   a. For the new node, **Action menu (⋮) → Edit Labels**.

   b. 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 .

3.1.2. Adding a node on an user-provisioned infrastructure

Use this procedure to add a node on an AWS or VMware user-provisioned infrastructure.

Prerequisites

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

Procedure

1. Depending on whether you are adding a node on an AWS user provisioned infrastructure or a VMware user-provisioned infrastructure, perform the following steps:

   - For AWS

     a. Create a new AWS machine instance with the required infrastructure. See Platform requirements.

     b. Create a new OpenShift Container Platform node using the new AWS machine instance.

   - For VMware:

     a. Create a new VM on vSphere with the required infrastructure. See Platform requirements.

     b. Create a new OpenShift Container Platform worker node using the new VM.

2. Check for certificate signing requests (CSRs) related to OpenShift Container Storage that are in **Pending** state:

   ```
   $ oc get csr
   ```

3. Approve all required OpenShift Container Storage CSRs for the new node:

   ```
   $ oc adm certificate approve <Certificate_Name>
   ```

4. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

5. 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=""
  ```

> **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 .

3.1.3. Adding a node using a local storage device

Use this procedure to add a node on bare metal, IBM Power Systems, IBM Z or LinuxONE, Amazon EC2, and VMware infrastructures.

> **IMPORTANT**
>
> Scaling storage nodes for Amazon EC2 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.

Prerequisites

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

- 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.

- You must have three OpenShift Container Platform worker nodes with the same storage type and size attached to each node (for example, 2TB SSD or 2TB NVMe drive) as the original OpenShift Container Storage StorageCluster was created with.

- If you have 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.

Procedure

1. Depending on whether you are adding a node on bare metal, IBM Power Systems, IBM Z or LinuxONE, Amazon EC2, or VMware infrastructure, perform the following steps:

   - For Amazon EC2

     a. Create a new Amazon EC2 I3 machine instance with the required infrastructure. See Creating a MachineSet in AWS and Platform requirements.

b. Create a new OpenShift Container Platform node using the new Amazon EC2 I3 machine instance.

- For VMware:

  a. Create a new VM on vSphere with the required infrastructure. See Platform requirements.

  b. Create a new OpenShift Container Platform worker node using the new VM.

- For bare metal:

  a. Get a new bare metal machine with the required infrastructure. See Platform requirements.

  b. Create a new OpenShift Container Platform node using the new bare metal machine.

- For IBM Power Systems:

  a. Get a new IBM Power machine with the required infrastructure. See Platform requirements.

  b. Create a new OpenShift Container Platform node using the new IBM Power machine.

- For IBM Z or LinuxONE:

  a. Get a new IBM Z or LinuxONE machine with the required infrastructure. See Platform requirements.

  b. Create a new OpenShift Container Platform node using the new IBM Z or LinuxONE machine.

2. Check for certificate signing requests (CSRs) related to OpenShift Container Storage that are in **Pending** state:

```
$ oc get csr
```

3. Approve all required OpenShift Container Storage CSRs for the new node:

```
$ oc adm certificate approve <Certificate_Name>
```

4. Click **Compute → Nodes**, confirm if the new node is in **Ready** state.

5. 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**.
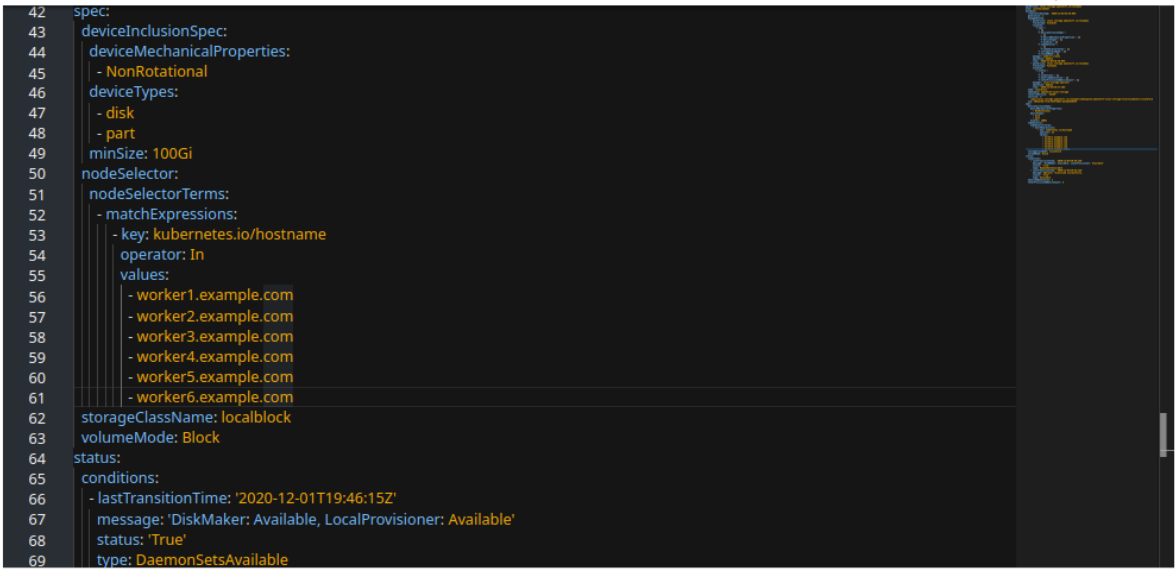
From Command line interface

- 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=""
```

6. Click **Operators → Installed Operators** from the OpenShift Web Console.
   From the **Project** drop-down list, make sure to select the project where the Local Storage
   Operator is installed.

7. Click on **Local Storage**.

8. Click the **Local Volume Discovery** tab

9. Beside the **LocalVolumeDiscovery**, click Action menu **(⋮)** → Edit **Local Volume Discovery**.

10. In the YAML, add the hostname of the new node in the values field under the **node selector**.

11. Click **Save**.

12. Click the **Local Volume Sets** tab.

13. Beside the **LocalVolumeSet**, click Action menu **(⋮)** → **Edit Local Volume Set**.

14. In the YAML, add the hostname of the new node in the **values** field under the **node selector**.

    Figure 3.1. YAML showing the addition of new hostnames

15. 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 .

3.1.4. 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-***

3.2. SCALING UP STORAGE CAPACITY

To scale up storage capacity, see Scaling up storage by adding capacity .