



Red Hat OpenShift Data Foundation 4.10

Scaling storage

Instructions for scaling operations in OpenShift Data Foundation

Red Hat OpenShift Data Foundation 4.10 Scaling storage

Instructions for scaling operations in OpenShift Data Foundation

Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document explains scaling options for Red Hat OpenShift Data Foundation.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
PREFACE	5
CHAPTER 1. REQUIREMENTS FOR SCALING STORAGE NODES	6
1.1. SUPPORTED DEPLOYMENTS FOR RED HAT OPENSIFT DATA FOUNDATION	6
CHAPTER 2. SCALING UP STORAGE CAPACITY	7
2.1. CREATING A STORAGE CLASS	7
2.2. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES	8
2.3. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES USING LOCAL STORAGE DEVICES	10
2.4. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES ON IBM Z OR LINUXONE INFRASTRUCTURE	13
2.5. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES ON IBM POWER INFRASTRUCTURE USING LOCAL STORAGE DEVICES	16
CHAPTER 3. SCALING OUT STORAGE CAPACITY	20
3.1. ADDING A NODE	20
3.1.1. Adding a node on an installer-provisioned infrastructure	20
3.1.2. Adding a node on an user-provisioned infrastructure	21
3.1.3. Adding a node using a local storage device	22
3.1.4. Adding a node using a local storage device on IBM Power	25
3.1.5. Verifying the addition of a new node	26
3.2. ADDING CAPACITY TO A NEWLY ADDED NODE	27
3.2.1. Add capacity with 3 OSDs using the Add Capacity option	27
3.2.2. Add capacity using YAML	27
3.2.2.1. Verifying if flexible scaling is enabled	27
3.2.2.2. Adding capacity using the YAML in multiples of 1 OSD	27

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Do let us know how we can make it better.

To give feedback, create a Bugzilla ticket:

1. Go to the [Bugzilla](#) website.
2. In the **Component** section, choose **documentation**.
3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
4. Click **Submit Bug**.

PREFACE

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

- **Scale up storage nodes** - Add storage capacity to the existing Red Hat OpenShift Data Foundation 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 Data Foundation 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 OPENSIFT DATA FOUNDATION

- User-provisioned infrastructure:
 - Amazon Web Services (AWS)
 - VMware
 - Bare metal
 - IBM Power
 - IBM Z or LinuxONE
- Installer-provisioned infrastructure:
 - Amazon Web Services (AWS)
 - Microsoft Azure
 - Red Hat Virtualization
 - VMware

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, Red Hat Virtualization or Azure infrastructures using dynamic or automated provisioning of storage devices, see [Section 2.2, “Scaling up storage by adding capacity to your OpenShift Data Foundation nodes”](#).
- For bare metal, VMware or Red Hat Virtualization infrastructures using local storage devices, see [Section 2.3, “Scaling up storage by adding capacity to your OpenShift Data Foundation nodes using local storage devices”](#).
- For IBM Z or LinuxONE infrastructures using local storage devices, see [Section 2.4, “Scaling up storage by adding capacity to your OpenShift Data Foundation nodes on IBM Z or LinuxONE infrastructure”](#).
- For IBM Power using local storage devices, see [Section 2.5, “Scaling up storage by adding capacity to your OpenShift Data Foundation nodes on IBM Power 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 Data Foundation 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.

Prerequisites

- Administrator access to OpenShift web console.

Procedure

1. Log in to OpenShift Web Console.
2. Click **Storage** → **StorageClasses**.
3. Click **Create Storage Class**
 - a. Enter the storage class **Name** and **Description**.
 - b. Select the required **Reclaim Policy** and **Provisioner**.
 - c. Click **Create** to create the Storage Class.

Verification steps

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

2.2. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES

You can add storage capacity and performance to your configured Red Hat OpenShift Data Foundation worker nodes on the following infrastructures:

- AWS
- VMware vSphere
- Red Hat Virtualization
- Microsoft Azure

Prerequisites

- A running OpenShift Data Foundation 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 **Operators** → **Installed Operators**.
3. Click **OpenShift Data Foundation** Operator.
4. Click the **Storage Systems** tab.
 - a. Click the **Action Menu** (**⋮**) on the far right of the storage system name to extend the options menu.
 - b. Select **Add Capacity** from the options menu.
 - c. Select the **Storage Class**.
Set the storage class to **gp2** or **gp3** on AWS, **thin** on VMware, **ovirt-csi-sc** on Red Hat Virtualization 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.



NOTE

OpenShift Data Foundation version 4.10 now supports the gp2 and gp3 CSI drivers introduced by AWS. These drivers provide improved storage expansion capabilities (gp2 CSI compared to gp2 in-tree) and a reduced monthly price-point (gp3). In case a high throughput is required, gp3 is recommended to be used when deploying OpenShift Data Foundation.



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 Data Foundation uses a replica count of 3.

- d. Click **Add**.
5. To check the status, navigate to **Storage → Data Foundation** and verify that **Storage System** in the Status card has a green tick.

Verification steps

- Verify the **Raw Capacity** card.
 - a. In the OpenShift Web Console, click **Storage → Data Foundation**.
 - b. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.
 - c. In the **Block and File** tab, check the **Raw Capacity** card. Note that the capacity increases based on your selections.



NOTE

The raw capacity does not take replication into account and shows the full capacity.

- Verify that the new OSDs and their corresponding new Persistent Volume Claims (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.



NOTE

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- 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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- Optional: If cluster-wide encryption is enabled on the cluster, verify that the new OSD devices are encrypted.
 - a. Identify the nodes where the new OSD pods are running.

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

<OSD-pod-name>

Is the name of the OSD pod.

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 the previous step, do the following:
 - i. Create a debug pod and open a chroot environment for the selected hosts.

```
$ oc debug node/<node-name>
```


<node-name>

Is the name of the node.

```
$ chroot /host
```

- ii. Check for the **crypt** keyword beside the **ocs-deviceset** names.

```
$ lsblk
```

**IMPORTANT**

Cluster reduction is supported only with the [Red Hat Support Team's](#) assistance..

2.3. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES USING LOCAL STORAGE DEVICES

You can add storage capacity (additional storage devices) to your configured local storage based OpenShift Data Foundation worker nodes on the following infrastructures:

- Bare metal

- VMware
- Red Hat Virtualization

Prerequisites

- You must be logged into the OpenShift Container Platform cluster.
- You must have installed the 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](#)
 - [Installing Local Storage Operator on Red Hat Virtualization cluster](#)
- 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 Data Foundation 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. In the OpenShift Web Console, click **Operators** → **Installed Operators**.
2. Click **OpenShift Data Foundation Operator**.
3. Click **Storage Systems** tab.
 - a. Click the **Action menu** (**⋮**) next to the visible list to extend the options menu.
 - b. Select **Add Capacity** from the options menu.
 - c. 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.
 - d. Click **Add**.
4. To check the status, navigate to **Storage** → **Data Foundation** and verify that **Storage System** in the **Status** card has a green tick.

Verification steps

- Verify the **Raw Capacity** card.
 - a. In the OpenShift Web Console, click **Storage** → **Data Foundation**.
 - b. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.
 - c. In the **Block and File** tab, check the **Raw Capacity** card.
Note that the capacity increases based on your selections.

**NOTE**

The raw capacity does not take replication into account and shows the full capacity.

- Verify that the new OSDs and their corresponding new Persistent Volume Claims (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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- 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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- Optional: If cluster-wide 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>
```

<OSD-pod-name>

Is the name of the OSD pod.
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 the 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>
```


<node-name>

Is the name of the node.


```
$ chroot /host
```

- ii. Check for the **crypt** keyword beside the **ocs-deviceset** names.

```
$ lsblk
```



IMPORTANT

Cluster reduction is supported only with the [Red Hat Support Team's](#) assistance.

2.4. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES ON IBM Z OR LINUXONE INFRASTRUCTURE

You can add storage capacity and performance to your configured Red Hat OpenShift Data Foundation worker nodes.

Prerequisites

- A running OpenShift Data Foundation 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.

```
$ 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.

```
$ 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 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
zfcplun   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 Data Foundation Operator**.
6. In the top navigation bar, scroll right and click **Storage Systems** tab.
 - a. Click the **Action menu** (**⋮**) next to the visible list to extend the options menu.
 - b. 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 Data Foundation uses a replica count of 3.
 - c. Click **Add**.
7. To check the status, navigate to **Storage → Data Foundation** and verify that **Storage System** in the **Status** card has a green tick.

Verification steps

- Verify the **Raw Capacity** card.
 - a. In the OpenShift Web Console, click **Storage → Data Foundation**.
 - b. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.
 - c. In the **Block and File** tab, check the **Raw Capacity** card.

Note that the capacity increases based on your selections.

**NOTE**

The raw capacity does not take replication into account and shows the full capacity.

- Verify that the new OSDs and their corresponding new Persistent Volume Claims (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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- 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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- Optional: If cluster-wide 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>
```

<OSD-pod-name>

Is the name of the OSD pod.
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 the 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>
```


<node-name>

Is the name of the node.

```
$ chroot /host
```

- ii. Check for the **crypt** keyword beside the **ocs-deviceSet** names.

```
$ lsblk
```

**IMPORTANT**

Cluster reduction is supported only with the [Red Hat Support Team's](#) assistance..

2.5. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES ON IBM POWER INFRASTRUCTURE USING LOCAL STORAGE DEVICES

You can add storage capacity (additional storage devices) to your configured local storage based OpenShift Data Foundation worker nodes on IBM Power infrastructures.

Prerequisites

- You must be logged into OpenShift Container Platform cluster.
- You must have installed the local storage operator. Use the following procedure:
 - [Installing Local Storage Operator on IBM Power](#)
- 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 Data Foundation StorageCluster was created with.

Procedure

1. To add storage capacity to OpenShift Container Platform nodes with OpenShift Data Foundation installed, you need to
 - a. Find the available devices that you want to add, that is, a minimum of one device per worker node. You can follow the procedure for finding available storage devices in the respective deployment guide.

**NOTE**

Make sure you perform this process for all the existing nodes (minimum of 3) for which you want to add storage.

- b. Add the additional disks to the **LocalVolume** custom resource (CR).

```
$ oc edit -n openshift-local-storage localvolume localblock
```

Example output:

```
spec:
  logLevel: Normal
  managementState: Managed
  nodeSelector:
    nodeSelectorTerms:
      - matchExpressions:
          - key: kubernetes.io/hostname
            operator: In
            values:
              - worker-0
              - worker-1
```

```

- worker-2
storageClassDevices:
- devicePaths:
  - /dev/sda
  - /dev/sdx # newly added device
storageClassName: localblock
volumeMode: Block

```

Make sure to save the changes after editing the CR.

Example output:

```
localvolume.local.storage.openshift.io/localblock edited
```

You can see in this CR that new devices are added.

- **sdx**

2. Display the newly created Persistent Volumes (PVs) with the **storageclass** name used in the **localVolume** CR.

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

Example output:

```

local-pv-a04ffd8    500Gi    RWO    Delete    Available    localblock    24s
local-pv-a0ca996b  500Gi    RWO    Delete    Available    localblock    23s
local-pv-c171754a  500Gi    RWO    Delete    Available    localblock    23s


```

3. Navigate to the OpenShift Web Console.
4. Click **Operators** on the left navigation bar.
5. Select **Installed Operators**.
6. In the window, click **OpenShift Data Foundation** Operator.
7. In the top navigation bar, scroll right and click **Storage System** tab.
 - a. Click the **Action menu** (**:**) next to the visible list to extend the options menu.
 - b. Select **Add Capacity** from the options menu.
From this dialog box, set the **Storage Class** name to the name used in the **localVolume** CR.
Available Capacity displayed is based on the local disks available in storage class.
 - c. Click **Add**.
8. To check the status, navigate to **Storage → Data Foundation** and verify that **Storage System** in the **Status** card has a green tick.

Verification steps

- Verify the available Capacity.
 - In the OpenShift Web Console, click **Storage → Data Foundation**.

- Click the **Storage Systems** tab and then click on **ocs-storagecluster-storagesystem**.
- Navigate to **Overview** → **Block and File** tab, then check the **Raw Capacity** card. Note that the capacity increases based on your selections.

**NOTE**

The raw capacity does not take replication into account and shows the full capacity.

- Verify that the new OSDs and their corresponding new Persistent Volume Claims (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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- 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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- Optional: If cluster-wide 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>
```

<OSD-pod-name>

Is the name of the OSD pod.

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 the 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>
```

<node-name>

Is the name of the node.

```
$ chroot /host
```

- ii. Check for the **crypt** keyword beside the **ocs-deviceset** names.

```
$ lsblk
```



IMPORTANT

Cluster reduction is supported only with the [Red Hat Support Team's](#) assistance.

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 Data Foundation 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 or Red Hat Virtualization installer-provisioned infrastructures, see [Adding a node on an installer-provisioned infrastructure](#).
- For AWS or VMware user-provisioned infrastructure, see [Adding a node on an user-provisioned infrastructure](#).
- For bare metal, IBM Z or LinuxONE, VMware, or Red Hat Virtualization infrastructures, see [Adding a node using a local storage device](#).
- For IBM Power, see [Adding a node using a local storage device on IBM Power](#).

3.1.1. Adding a node on an installer-provisioned infrastructure

You can add a node on the following installer provisioned infrastructures:

- AWS
- Azure
- Red Hat Virtualization
- VMware

Prerequisites

- You must be logged into OpenShift Container Platform cluster.

Procedure

1. Navigate to **Compute → Machine Sets**.
2. On the machine set where you want to add nodes, select **Edit Machine Count**

- a. Add the amount of nodes, and click **Save**.
 - b. Click **Compute** → **Nodes** and confirm if the new node is in **Ready** state.
3. Apply the OpenShift Data Foundation label to the new node.
 - a. For the new node, click **Action menu (⋮)** → **Edit Labels**.
 - b. Add `cluster.ocs.openshift.io/openshift-storage`, and click **Save**.



NOTE

It is recommended to add 3 nodes, one 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

You can add a node on an AWS or VMware user-provisioned infrastructure.

Prerequisites

- You must be logged into OpenShift Container Platform cluster.

Procedure

1. Depending on whether you are adding a node on an AWS 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 virtual machine (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 Data Foundation that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required OpenShift Data Foundation CSRs for the new node.

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

<Certificate_Name>

Is the name of the CSR.

4. Click **Compute** → **Nodes**, confirm if the new node is in **Ready** state.
5. Apply the OpenShift Data Foundation 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

- Apply the OpenShift Data Foundation label to the new node.

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

<new_node_name>

Is the name of the new node.

**NOTE**

It is recommended to add 3 nodes, one 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

You can add a node on the following:

- Bare metal
- IBM Z or LinuxONE
- VMware
- Red Hat Virtualization

Prerequisites

- You must be logged into the OpenShift Container Platform cluster.
- 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 Data Foundation StorageCluster was created with.

Procedure

1. Depending on the type of infrastructure, perform the following steps:

- For VMware:
 - a. Create a new virtual machine (VM) on vSphere with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new VM.
 - For Red Hat Virtualization:
 - a. Create a new VM on Red Hat Virtualization 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 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 Data Foundation that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required OpenShift Data Foundation CSRs for the new node.

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

<Certificate_Name>

Is the name of the CSR.

4. Click **Compute** → **Nodes**, confirm if the new node is in **Ready** state.
5. Apply the OpenShift Data Foundation 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

- Apply the OpenShift Data Foundation label to the new node.

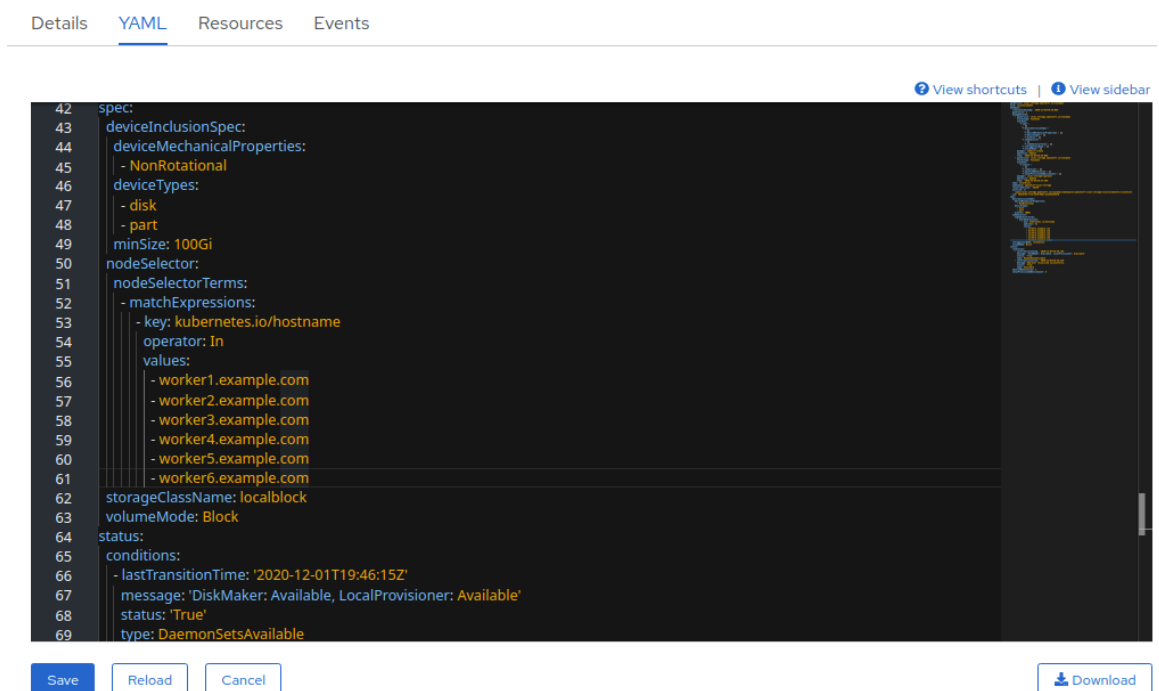
```
$ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
```

<new_node_name>

Is the name of the new node.

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 **Local Storage**.
8. Click the **Local Volume Discovery** tab.
 - a. Beside the **LocalVolumeDiscovery**, click Action menu (⋮) → **Edit Local Volume Discovery**.
 - b. In the YAML, add the hostname of the new node in the **values** field under the node selector.
 - c. Click **Save**.
9. Click the **Local Volume Sets** tab.
 - a. Beside the **LocalVolumeSet**, click Action menu (⋮) → **Edit Local Volume Set**
 - b. 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



```

42 spec:
43   deviceInclusionSpec:
44     deviceMechanicalProperties:
45       - NonRotational
46     deviceTypes:
47       - disk
48       - part
49     minSize: 100Gi
50   nodeSelector:
51     nodeSelectorTerms:
52       - matchExpressions:
53         - key: kubernetes.io/hostname
54           operator: In
55           values:
56             - worker1.example.com
57             - worker2.example.com
58             - worker3.example.com
59             - worker4.example.com
60             - worker5.example.com
61             - worker6.example.com
62   storageClassName: localblock
63   volumeMode: Block
64   status:
65     conditions:
66       - lastTransitionTime: '2020-12-01T19:46:15Z'
67         message: 'DiskMaker: Available, LocalProvisioner: Available'
68         status: 'True'
69         type: DaemonSetsAvailable

```

Save Reload Cancel Download

- c. Click **Save**.



NOTE

It is recommended to add 3 nodes, one 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. Adding a node using a local storage device on IBM Power

Prerequisites

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

Procedure

1. Get a new IBM Power machine with the required infrastructure. See [Platform requirements](#).
2. Create a new OpenShift Container Platform node using the new IBM Power machine.
 - a. Check for certificate signing requests (CSRs) related to OpenShift Data Foundation that are in **Pending** state.

```
$ oc get csr
```

- b. Approve all the required OpenShift Data Foundation CSRs for the new node.

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

<Certificate_Name>

Is the name of the CSR.

- c. Click **Compute** → **Nodes**, confirm if the new node is in **Ready** state.
- d. Apply the OpenShift Data Foundation label to the new node using any one of the following:

From User interface

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

From Command line interface

- Apply the OpenShift Data Foundation label to the new node.

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

<new_node_name>

Is the name of the new node.

- e. 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.

- f. Click **Local Storage**.
- g. Click the **Local Volume** tab.
 - i. Beside the **LocalVolume**, click Action menu (⋮) → **Edit Local Volume**.
 - ii. In the YAML, add the hostname of the new node in the **values** field under the **node selector**.

Figure 3.2. YAML showing the addition of new hostnames

```

1  apiVersion: local.storage.openshift.io/v1
2  kind: LocalVolume
3  metadata:
4    creationTimestamp: '2021-07-09T08:12:56Z'
5  > finalizers: ...
7  generation: 2
8  > managedFields: ...
47  name: localblock
48  namespace: openshift-local-storage
49  resourceVersion: '1988620'
50  uid: 89d78619-a76b-4b97-af9d-6760c57ab531
51  spec:
52    logLevel: Normal
53    managementState: Managed
54    nodeSelector:
55      nodeSelectorTerms:
56      - matchExpressions:
57        - key: kubernetes.io/hostname
58          operator: In
59          values:
60            - worker-0
61            - worker-1
62            - worker-2
63            - worker-3
64    storageClassDevices:

```

- iii. Click **Save**.

**NOTE**

It is recommended to add 3 nodes, one 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](#).

3.1.5. 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. ADDING CAPACITY TO A NEWLY ADDED NODE

To add capacity to a newly added node, either use the **Add Capacity** option to expand the storage cluster with 3 OSDs or use the new flexible scaling feature that allows you to expand the storage cluster by any number of OSDs if it is enabled.

3.2.1. Add capacity with 3 OSDs using the Add Capacity option

To add capacity by 3 OSDs for dynamic and local storage using the **Add Capacity** option in the user interface, see [Scaling up storage by adding capacity](#). The **Add Capacity** option is available for storage clusters with or without the flexible scaling feature enabled.

3.2.2. Add capacity using YAML

With flexible scaling enabled, you can add capacity by 1 or more OSDs at a time using the YAML instead of the default set of 3 OSDs. However, you need to make sure that you add disks in a way that the cluster remains balanced.

Flexible scaling is supported only for the internal-attached mode of storage cluster creation.

To enable flexible scaling, create a cluster with 3 nodes, and fewer than 3 availability zones. The OpenShift Web Console detects the 3 nodes spread across fewer than 3 availability zones and enables flexible scaling.



IMPORTANT

You can not enable or disable the flexible scaling feature after creating the storage cluster.

3.2.2.1. Verifying if flexible scaling is enabled

Procedure

- In the Web Console, click **Home** → **Search**.
 - a. Enter **StorageCluster** in the search field.
 - b. Click **ocs-storagecluster**.
 - c. In the YAML tab, search for the keys **flexibleScaling** in **spec** section and **failureDomain** in **status** section.
If **flexible scaling** is true and **failureDomain** is set to host, flexible scaling feature is enabled.

```
spec:
  flexibleScaling: true
  [...]
status:
  failureDomain: host
```

3.2.2.2. Adding capacity using the YAML in multiples of 1 OSD


To add OSDs to your storage cluster flexibly through the YAML, perform the following steps:

Prerequisites

- Administrator access to the OpenShift Container Platform web console.
- A storage cluster with flexible scaling enabled.
- Additional disks available for adding capacity.

Procedure

1. In the OpenShift Web Console, click **Home** → **Search**.
 - a. Search for **ocs-storagecluster** in the search field and click on **ocs-storagecluster** from the search result.
 - b. Click the Action menu (**⋮**) next to the storage cluster you want to scale up.
 - c. Click **Edit Storage Cluster**. You are redirected to the YAML.
2. In the YAML, search for the key **count**. This count parameter scales up the capacity.
3. Increase the count by the number of OSDs you want to add to your cluster.



IMPORTANT

Ensure the **count** parameter in the YAML is incremented depending on the number of available disks and also make sure that you add disks in a way that the cluster remains balanced.

4. Click **Save**.



NOTE

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

Verification steps

- Verify the **Raw Capacity** card.
 - a. In the OpenShift Web Console, click **Storage** → **Data Foundation**.
 - b. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.
 - c. In the **Block and File** tab, check the **Raw Capacity** card.
Note that the capacity increases based on your selections.



NOTE

The raw capacity does not take replication into account and shows the full capacity.

- Verify that the new OSDs and their corresponding new Persistent Volume Claims (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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- 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.

**NOTE**

If the **Show default projects** option is disabled, use the toggle button to list all the default projects.

- Optional: If cluster-wide 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>
```

<OSD-pod-name>

Is the name of the OSD pod.
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 the 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>
```


<node name>

Is the name of the node.

```
$ chroot /host
```

- ii. Check for the **crypt** keyword beside the **ocs-device** name(s).

```
$ lsblk
```



IMPORTANT

Cluster reduction is supported only with the [Red Hat Support Team's](#) assistance.