



Red Hat OpenShift Data Foundation 4.12

Scaling storage

Instructions for scaling operations in OpenShift Data Foundation

Red Hat OpenShift Data Foundation 4.12 Scaling storage

Instructions for scaling operations in OpenShift Data Foundation

Legal Notice

Copyright © 2024 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	4
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	5
CHAPTER 1. INTRODUCTION TO SCALING STORAGE	6
1.1. SUPPORTED DEPLOYMENTS FOR RED HAT OPENSIFT DATA FOUNDATION	6
CHAPTER 2. REQUIREMENTS FOR SCALING STORAGE	8
CHAPTER 3. SCALING STORAGE CAPACITY OF AWS OPENSIFT DATA FOUNDATION CLUSTER	9
3.1. SCALING UP STORAGE CAPACITY ON A CLUSTER	9
3.2. SCALING OUT STORAGE CAPACITY ON A AWS CLUSTER	11
3.2.1. Adding a node	11
3.2.1.1. Adding a node to an installer-provisioned infrastructure	12
3.2.1.2. Adding a node to an user-provisioned infrastructure	12
3.2.2. Scaling up storage capacity	14
CHAPTER 4. SCALING STORAGE OF BARE METAL OPENSIFT DATA FOUNDATION CLUSTER	15
4.1. SCALING UP A CLUSTER CREATED USING LOCAL STORAGE DEVICES	15
4.2. SCALING OUT STORAGE CAPACITY ON A BARE METAL CLUSTER	17
4.2.1. Adding a node using a local storage device	18
4.2.2. Scaling up storage capacity	20
CHAPTER 5. SCALING STORAGE OF VMWARE OPENSIFT DATA FOUNDATION CLUSTER	21
5.1. SCALING UP STORAGE ON A VMWARE CLUSTER	21
5.2. SCALING UP A CLUSTER CREATED USING LOCAL STORAGE DEVICES	23
5.3. SCALING OUT STORAGE CAPACITY ON A VMWARE CLUSTER	25
5.3.1. Adding a node to an installer-provisioned infrastructure	25
5.3.2. Adding a node to an user-provisioned infrastructure	26
5.3.3. Adding a node using a local storage device	27
5.3.4. Scaling up storage capacity	30
CHAPTER 6. SCALING STORAGE OF MICROSOFT AZURE OPENSIFT DATA FOUNDATION CLUSTER .	31
6.1. SCALING UP STORAGE CAPACITY ON A CLUSTER	31
6.2. SCALING OUT STORAGE CAPACITY ON A MICROSOFT AZURE CLUSTER	33
6.2.1. Adding a node to an installer-provisioned infrastructure	33
6.2.2. Scaling up storage capacity	34
CHAPTER 7. SCALING STORAGE OF RED HAT VIRTUALIZATION OPENSIFT DATA FOUNDATION CLUSTER	35
7.1. SCALING UP STORAGE CAPACITY ON A CLUSTER	35
7.2. SCALING OUT STORAGE CAPACITY ON A RED HAT VIRTUALIZATION CLUSTER	37
7.2.1. Adding a node to an installer-provisioned infrastructure	37
7.2.2. Adding a node using a local storage device	38
7.2.3. Scaling up storage capacity	40
CHAPTER 8. SCALING STORAGE OF IBM Z OR LINUXONE OPENSIFT DATA FOUNDATION CLUSTER	41
8.1. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES ON IBM Z OR LINUXONE INFRASTRUCTURE	41
8.2. SCALING OUT STORAGE CAPACITY ON A IBM Z OR LINUXONE CLUSTER	44
8.2.1. Adding a node using a local storage device	44
8.2.2. Scaling up storage capacity	46
CHAPTER 9. SCALING STORAGE OF IBM POWER OPENSIFT DATA FOUNDATION CLUSTER	47

9.1. SCALING UP STORAGE BY ADDING CAPACITY TO YOUR OPENSIFT DATA FOUNDATION NODES ON IBM POWER INFRASTRUCTURE USING LOCAL STORAGE DEVICES	47
9.2. SCALING OUT STORAGE CAPACITY ON A IBM POWER CLUSTER	50
9.2.1. Adding a node using a local storage device on IBM Power	51
9.2.2. Scaling up storage capacity	53

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**.

CHAPTER 1. INTRODUCTION TO SCALING STORAGE

Red Hat OpenShift Data Foundation is highly scalable storage system. OpenShift Data Foundation allows you to scale by adding the disks in the multiple of three, or three or any number depending upon the deployment type.

- For internal (dynamic provisioning) deployment mode, you can increase the capacity by adding 3 disks at a time.
- For internal-attached (Local Storage Operator based) mode, you can deploy with less than 3 failure domains.

With flexible scale deployment enabled, you can scale up by adding any number of disks. For deployment with 3 failure domains, you will be able to scale up by adding disks in the multiple of 3.

For scaling your storage in external mode, see [Red Hat Ceph Storage documentation](#) .



NOTE

You can use a maximum of nine storage devices per node. The high number of storage devices will lead to a higher recovery time during the loss of a node. This recommendation ensures that nodes stay below the cloud provider dynamic storage device attachment limits, and limits the recovery time after node failure with local storage devices.

While scaling, you must ensure that there are enough CPU and Memory resources as per scaling requirement.

Supported storage classes by default

- **gp2-csi** on AWS
- **thin** on VMware
- **ovirt-csi-sc** on Red Hat Virtualization
- **managed_premium** on Microsoft Azure

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. REQUIREMENTS FOR SCALING STORAGE

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](#)
- [Resource 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](#).

CHAPTER 3. SCALING STORAGE CAPACITY OF AWS OPENSIFT DATA FOUNDATION CLUSTER

To scale the storage capacity of your configured Red Hat OpenShift Data Foundation worker nodes on AWS cluster, you can increase the capacity by adding three disks at a time. Three disks are needed since OpenShift Data Foundation uses a replica count of 3 to maintain the high availability. So the amount of storage consumed is three times the usable space.



NOTE

Usable space may vary when encryption is enabled or replica 2 pools are being used.

3.1. SCALING UP STORAGE CAPACITY ON A CLUSTER

To increase the storage capacity in a dynamically created storage cluster on an user-provisioned infrastructure, you can add storage capacity and performance to your configured Red Hat OpenShift Data Foundation worker nodes.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.
- The disk should be of the same size and type as used during initial deployment.

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**. Choose the storage class which you wish to use to provision new storage devices.
 - 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 -n openshift-storage -o=custom-columns=NODE:.spec.nodeName pod/<OSD-pod-name>
```

<OSD-pod-name>

Is the name of the OSD pod.

For example:

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

Example output:

```
NODE
compute-1
```

- 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.

3.2. SCALING OUT STORAGE CAPACITY ON A AWS CLUSTER

OpenShift Data Foundation is highly scalable. It can be scaled out by adding new nodes with required storage and enough hardware resources in terms of CPU and RAM. Practically there is no limit on the number of nodes which can be added but from the support perspective 2000 nodes is the limit for OpenShift Data Foundation.

Scaling out storage capacity can be broken down into two steps

- Adding new node
- Scaling up the storage capacity



NOTE

OpenShift Data Foundation does not support heterogeneous OSD/Disk sizes.

3.2.1. Adding a node

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs or there are not enough resources to add new OSDs on the existing nodes. It is always recommended to add nodes in the multiple of three, each of them in different failure domains.

While we recommend adding nodes in the multiple of three, you still get the flexibility of adding one node at a time in the flexible scaling deployment. Refer to the Knowledgebase article [Verify if flexible scaling is enabled](#).

**NOTE**


OpenShift Data Foundation does not support heterogeneous disk size and types. The new nodes to be added should have the disk of the same type and size which was used during OpenShift Data Foundation deployment.

3.2.1.1. Adding a node to an installer-provisioned infrastructure**Prerequisites**

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage 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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - `csi-cephfsplugin-*`
 - `csi-rbdplugin-*`

3.2.1.2. Adding a node to an user-provisioned infrastructure**Prerequisites**

- You have administrative privilege to the OpenShift Container Platform Console.

- You have a running OpenShift Data Foundation Storage Cluster.

Procedure

1. Depending on the type of infrastructure, perform the following steps:
 - a. Get a new machine with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new machine.
2. Check for certificate signing requests (CSRs) that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required 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.

Verification steps

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:

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

3.2.2. Scaling up storage capacity

To scale up storage capacity, see [Scaling up storage by adding capacity](#).

CHAPTER 4. SCALING STORAGE OF BARE METAL OPENSIFT DATA FOUNDATION CLUSTER

To scale the storage capacity of your configured Red Hat OpenShift Data Foundation worker nodes on your bare metal cluster, you can increase the capacity by adding three disks at a time. Three disks are needed since OpenShift Data Foundation uses a replica count of 3 to maintain the high availability. So the amount of storage consumed is three times the usable space.



NOTE

Usable space may vary when encryption is enabled or replica 2 pools are being used.

4.1. SCALING UP A CLUSTER CREATED USING LOCAL STORAGE DEVICES

In order to scale up an OpenShift Data Foundation cluster which was created using local storage devices, a new disk needs to be added to the storage node. It is recommended to have the new disks of the same size as used earlier during the deployment as OpenShift Data Foundation does not support heterogeneous disks/OSD's.

For deployments having three failure domains, you can scale up the storage by adding disks in the multiple of three, with the same number of disks coming from nodes in each of the failure domains. For example, if we scale by adding six disks, two disks are taken from nodes in each of the three failure domains. If the number of disks is not in multiples of three, it will only consume the disk to the maximum in the multiple of three while the remaining disks remain unused.

For deployments having less than three failure domains, there is flexibility in adding the number of disks. In this case, you can add any number of disks. In order to check if flexible scaling is enabled or not, refer to the Knowledgebase article [Verify if flexible scaling is enabled](#).



NOTE

Flexible scaling features get enabled at the time of deployment and can not be enabled or disabled later on.


Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.
- Disks to be used for scaling are already attached to the storage node
- **LocalVolumeDiscovery** and **LocalVolumeSet** objects are already created.

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 the **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 -n openshift-storage -o=custom-columns=NODE:.spec.nodeName pod/<OSD-pod-name>
```

<OSD-pod-name>

Is the name of the OSD pod.
For example:

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

Example output:

```
NODE
compute-1
```

- 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.

4.2. SCALING OUT STORAGE CAPACITY ON A BARE METAL CLUSTER

OpenShift Data Foundation is highly scalable. It can be scaled out by adding new nodes with required storage and enough hardware resources in terms of CPU and RAM. There is no limit on the number of nodes which can be added. However, from the technical support perspective, 2000 nodes is the limit for OpenShift Data Foundation.

Scaling out storage capacity can be broken down into two steps

- Adding new node using local storage device
- Scaling up the storage capacity

**NOTE**

OpenShift Data Foundation does not support heterogeneous OSD/Disk sizes.

4.2.1. Adding a node using a local storage device

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs or when there are not enough resources to add new OSDs on the existing nodes.

Add nodes in the multiple of 3, each of them in different failure domains. Though it is recommended to add nodes in multiples of 3 nodes, you have the flexibility to add one node at a time in flexible scaling deployment. See Knowledgebase article [Verify if flexible scaling is enabled](#)

**NOTE**

OpenShift Data Foundation does not support heterogeneous disk size and types. The new nodes to be added should have the disk of the same type and size which was used during initial OpenShift Data Foundation deployment.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.

Procedure

1. Depending on the type of infrastructure, perform the following steps:
 - a. Get a new machine with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new machine.
2. Check for certificate signing requests (CSRs) that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required 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.

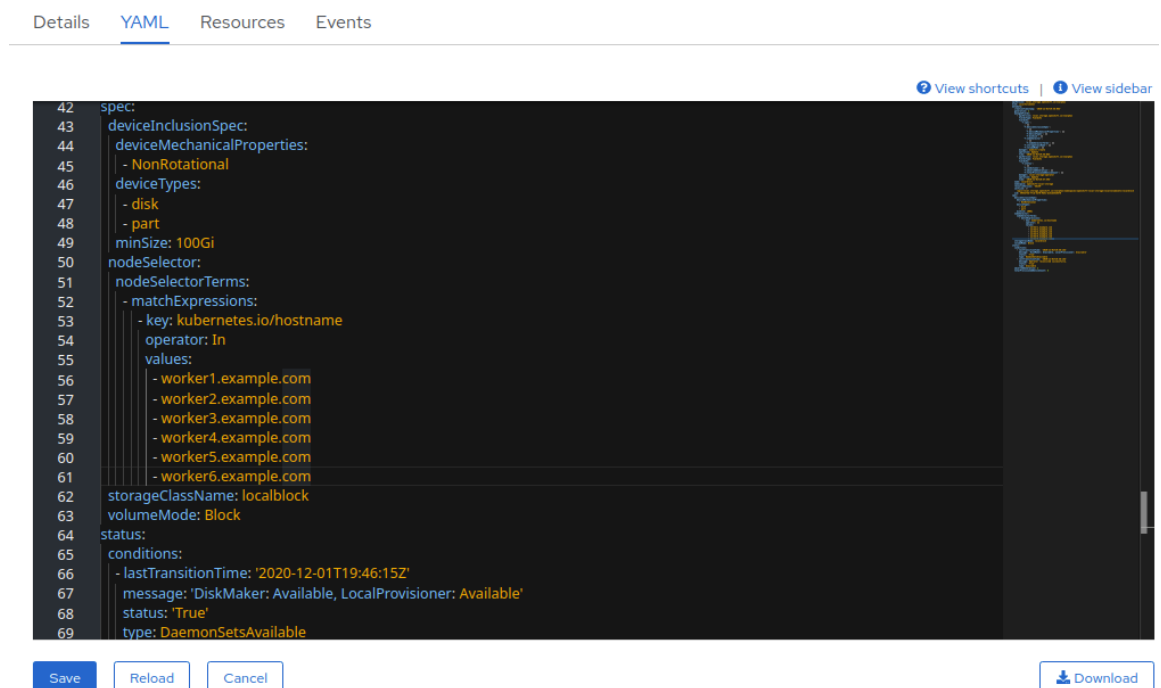
```
$ oc label node <new_node_name> cluster.ocs.openshift.io/openshift-storage=""
```

<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 4.1. YAML showing the addition of new hostnames



- 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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:

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

4.2.2. Scaling up storage capacity

To scale up storage capacity, see [Scaling up storage by adding capacity](#) .

CHAPTER 5. SCALING STORAGE OF VMWARE OPENSIFT DATA FOUNDATION CLUSTER

5.1. SCALING UP STORAGE ON A VMWARE CLUSTER

To increase the storage capacity in a dynamically created storage cluster on a VMware user-provisioned infrastructure, you can add storage capacity and performance to your configured Red Hat OpenShift Data Foundation worker nodes.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.
- The disk should be of the same size and type as used during initial deployment.

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**. Choose the storage class which you wish to use to provision new storage devices.
 - 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 -n openshift-storage -o=custom-columns=NODE:.spec.nodeName pod/<OSD-pod-name>
```

<OSD-pod-name>

Is the name of the OSD pod.

For example:

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

Example output:

```
NODE
compute-1
```

- 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.

5.2. SCALING UP A CLUSTER CREATED USING LOCAL STORAGE DEVICES

In order to scale up an OpenShift Data Foundation cluster which was created using local storage devices, a new disk needs to be added to the storage node. It is recommended to have the new disks of the same size as used earlier during the deployment as OpenShift Data Foundation does not support heterogeneous disks/OSD's.

For deployments having three failure domains, you can scale up the storage by adding disks in the multiple of three, with the same number of disks coming from nodes in each of the failure domains. For example, if we scale by adding six disks, two disks are taken from nodes in each of the three failure domains. If the number of disks is not in multiples of three, it will only consume the disk to the maximum in the multiple of three while the remaining disks remain unused.

For deployments having less than three failure domains, there is flexibility in adding the number of disks. In this case, you can add any number of disks. In order to check if flexible scaling is enabled or not, refer to the Knowledgebase article [Verify if flexible scaling is enabled](#).



NOTE


Flexible scaling features get enabled at the time of deployment and can not be enabled or disabled later on.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.
- Disks to be used for scaling are already attached to the storage node
- **LocalVolumeDiscovery** and **LocalVolumeSet** objects are already created.

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 the **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 -n openshift-storage -o=custom-columns=NODE:.spec.nodeName pod/<OSD-  
pod-name>
```

<OSD-pod-name>

Is the name of the OSD pod.

For example:

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

Example output:

```
NODE  
compute-1
```

- 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.

5.3. SCALING OUT STORAGE CAPACITY ON A VMWARE CLUSTER

5.3.1. Adding a node to an installer-provisioned infrastructure

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage 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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - **csi-cephfsplugin-***
 - **csi-rbdplugin-***

5.3.2. Adding a node to an user-provisioned infrastructure

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.

Procedure

1. Depending on the type of infrastructure, perform the following steps:
 - a. Get a new machine with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new machine.

2. Check for certificate signing requests (CSRs) that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required 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.

Verification steps

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:

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

5.3.3. Adding a node using a local storage device

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs or when there are not enough resources to add new OSDs on the existing nodes.

Add nodes in the multiple of 3, each of them in different failure domains. Though it is recommended to add nodes in multiples of 3 nodes, you have the flexibility to add one node at a time in flexible scaling deployment. See Knowledgebase article [Verify if flexible scaling is enabled](#)



NOTE

OpenShift Data Foundation does not support heterogeneous disk size and types. The new nodes to be added should have the disk of the same type and size which was used during initial OpenShift Data Foundation deployment.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.

Procedure

1. Depending on the type of infrastructure, perform the following steps:
 - a. Get a new machine with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new machine.
2. Check for certificate signing requests (CSRs) that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required 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=""
```

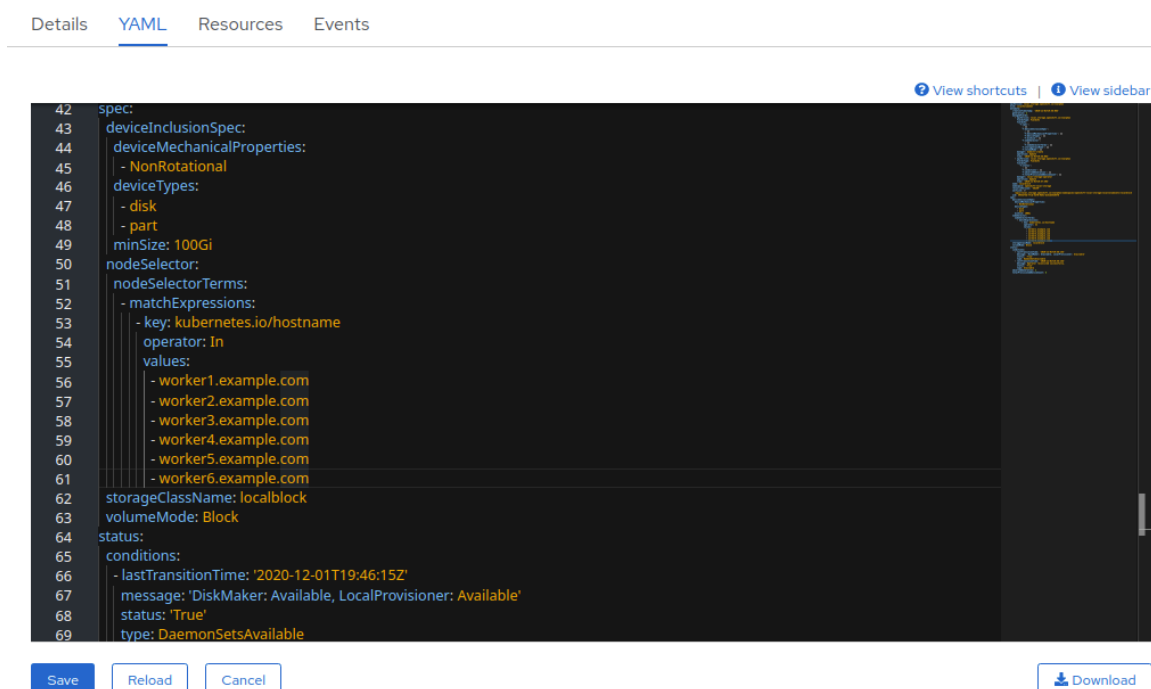
<new_node_name>

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 5.1. 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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:

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

5.3.4. Scaling up storage capacity

To scale up storage capacity:

- For dynamic storage devices, see [Scaling up storage by adding capacity](#) .
- For local storage devices, see [Scaling up a cluster created using local storage devices](#)

CHAPTER 6. SCALING STORAGE OF MICROSOFT AZURE OPENSIFT DATA FOUNDATION CLUSTER

To scale the storage capacity of your configured Red Hat OpenShift Data Foundation worker nodes on Microsoft Azure cluster, you can increase the capacity by adding three disks at a time. Three disks are needed since OpenShift Data Foundation uses a replica count of 3 to maintain the high availability. So the amount of storage consumed is three times the usable space.



NOTE

Usable space may vary when encryption is enabled or replica 2 pools are being used.

6.1. SCALING UP STORAGE CAPACITY ON A CLUSTER

To increase the storage capacity in a dynamically created storage cluster on an user-provisioned infrastructure, you can add storage capacity and performance to your configured Red Hat OpenShift Data Foundation worker nodes.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.
- The disk should be of the same size and type as used during initial deployment.

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**. Choose the storage class which you wish to use to provision new storage devices.
 - 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 -n openshift-storage -o=custom-columns=NODE:.spec.nodeName pod/<OSD-  
pod-name>
```

<OSD-pod-name>

Is the name of the OSD pod.

For example:

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

Example output:

```
NODE
compute-1
```

- 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.

6.2. SCALING OUT STORAGE CAPACITY ON A MICROSOFT AZURE CLUSTER

OpenShift Data Foundation is highly scalable. It can be scaled out by adding new nodes with required storage and enough hardware resources in terms of CPU and RAM. Practically there is no limit on the number of nodes which can be added but from the support perspective 2000 nodes is the limit for OpenShift Data Foundation.

Scaling out storage capacity can be broken down into two steps

- Adding new node
- Scaling up the storage capacity



NOTE

OpenShift Data Foundation does not support heterogeneous OSD/Disk sizes.

6.2.1. Adding a node to an installer-provisioned infrastructure

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage 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

1. Execute the following command in the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - **csi-cephfsplugin-***
 - **csi-rbdplugin-***

6.2.2. Scaling up storage capacity

To scale up storage capacity, see [Scaling up storage by adding capacity](#).

CHAPTER 7. SCALING STORAGE OF RED HAT VIRTUALIZATION OPENSIFT DATA FOUNDATION CLUSTER

To scale the storage capacity of your configured Red Hat OpenShift Data Foundation worker nodes on Red Hat Virtualization cluster, you can increase the capacity by adding three disks at a time. Three disks are needed since OpenShift Data Foundation uses a replica count of 3 to maintain the high availability. So the amount of storage consumed is three times the usable space.



NOTE

Usable space may vary when encryption is enabled or replica 2 pools are being used.

7.1. SCALING UP STORAGE CAPACITY ON A CLUSTER

To increase the storage capacity in a dynamically created storage cluster on an user-provisioned infrastructure, you can add storage capacity and performance to your configured Red Hat OpenShift Data Foundation worker nodes.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.
- The disk should be of the same size and type as used during initial deployment.

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**. Choose the storage class which you wish to use to provision new storage devices.
 - 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 -n openshift-storage -o=custom-columns=NODE:.spec.nodeName pod/<OSD-  
pod-name>
```

<OSD-pod-name>

Is the name of the OSD pod.

For example:

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

Example output:


```
NODE
compute-1
```

- 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.

7.2. SCALING OUT STORAGE CAPACITY ON A RED HAT VIRTUALIZATION CLUSTER

OpenShift Data Foundation is highly scalable. It can be scaled out by adding new nodes with required storage and enough hardware resources in terms of CPU and RAM. Practically there is no limit on the number of nodes which can be added but from the support perspective 2000 nodes is the limit for OpenShift Data Foundation.

Scaling out storage capacity can be broken down into two steps

- Adding new node
- Scaling up the storage capacity



NOTE

OpenShift Data Foundation does not support heterogeneous OSD/Disk sizes.

7.2.1. Adding a node to an installer-provisioned infrastructure

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage 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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - **csi-cephfsplugin-***
 - **csi-rbdplugin-***

7.2.2. Adding a node using a local storage device

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs or when there are not enough resources to add new OSDs on the existing nodes.

Add nodes in the multiple of 3, each of them in different failure domains. Though it is recommended to add nodes in multiples of 3 nodes, you have the flexibility to add one node at a time in flexible scaling deployment. See Knowledgebase article [Verify if flexible scaling is enabled](#)



NOTE

OpenShift Data Foundation does not support heterogeneous disk size and types. The new nodes to be added should have the disk of the same type and size which was used during initial OpenShift Data Foundation deployment.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.

Procedure

Procedure

1. Depending on the type of infrastructure, perform the following steps:
 - a. Get a new machine with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new machine.
2. Check for certificate signing requests (CSRs) that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required 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=""
```

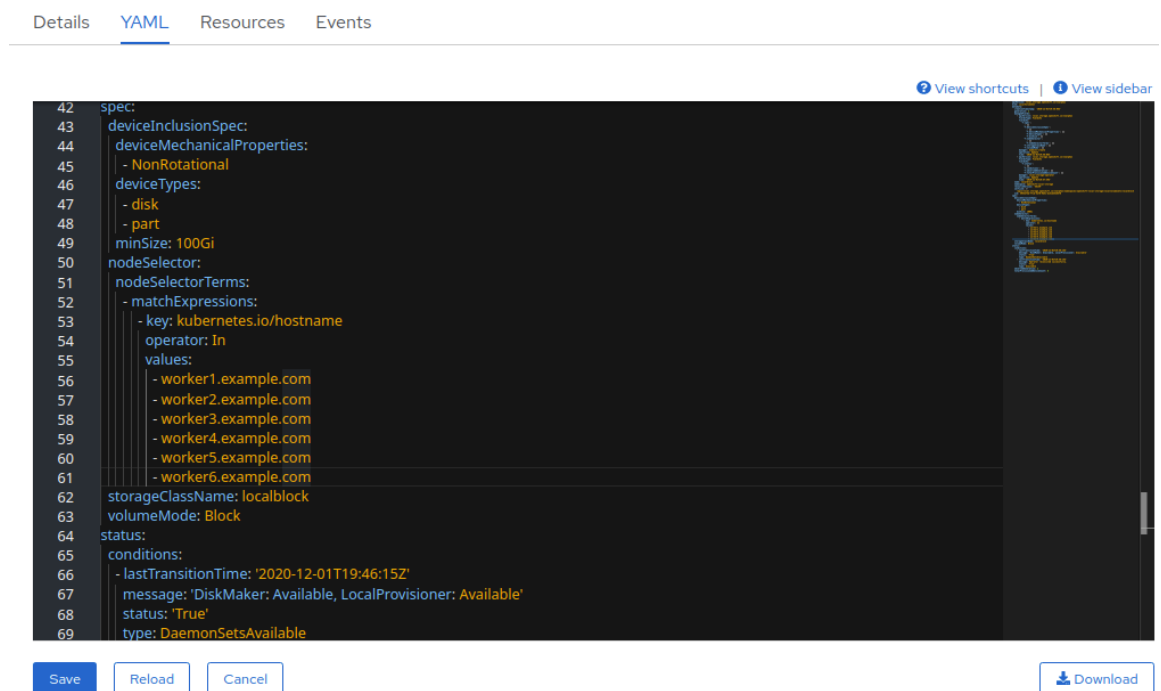
<new_node_name>

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 7.1. 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

1. Execute the following command in the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - **csi-cephfsplugin-***
 - **csi-rbdplugin-***

7.2.3. Scaling up storage capacity

To scale up storage capacity:

- For dynamic storage devices, see [Scaling up storage by adding capacity](#).
- For local storage devices, see [Scaling up a cluster created using local storage devices](#)

CHAPTER 8. SCALING STORAGE OF IBM Z OR LINUXONE OPENSIFT DATA FOUNDATION CLUSTER

8.1. 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.



NOTE

Flexible scaling features get enabled at the time of deployment and can not be enabled or disabled later on.

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
zfcplun   0.0.8204                          yes yes
zfcplun   0.0.8204:0x102107630b1b5060:0x4001402900000000 yes no  sda sg0
zfcplun   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 **zfcplun** 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 nodes where the new OSD pods are running.

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

<OSD-pod-name>

Is the name of the OSD pod.
For example:

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

Example output:

```
NODE
compute-1
```

- 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.

8.2. SCALING OUT STORAGE CAPACITY ON A IBM Z OR LINUXONE CLUSTER

8.2.1. Adding a node using a local storage device

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs or when there are not enough resources to add new OSDs on the existing nodes.

Add nodes in the multiple of 3, each of them in different failure domains. Though it is recommended to add nodes in multiples of 3 nodes, you have the flexibility to add one node at a time in flexible scaling deployment. See Knowledgebase article [Verify if flexible scaling is enabled](#)



NOTE

OpenShift Data Foundation does not support heterogeneous disk size and types. The new nodes to be added should have the disk of the same type and size which was used during initial OpenShift Data Foundation deployment.

Prerequisites

- You have administrative privilege to the OpenShift Container Platform Console.
- You have a running OpenShift Data Foundation Storage Cluster.

Procedure

1. Depending on the type of infrastructure, perform the following steps:
 - a. Get a new machine with the required infrastructure. See [Platform requirements](#).
 - b. Create a new OpenShift Container Platform worker node using the new machine.
2. Check for certificate signing requests (CSRs) that are in **Pending** state.

```
$ oc get csr
```

3. Approve all the required 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 8.1. YAML showing the addition of new hostnames

Details **YAML** Resources Events

View shortcuts | View sidebar

```

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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - **csi-cephfsplugin-***
 - **csi-rbdplugin-***

8.2.2. Scaling up storage capacity

To scale up storage capacity, see [Scaling up storage by adding capacity](#).

CHAPTER 9. SCALING STORAGE OF IBM POWER OPENSIFT DATA FOUNDATION CLUSTER

To scale the storage capacity of your configured Red Hat OpenShift Data Foundation worker nodes on IBM Power cluster, you can increase the capacity by adding three disks at a time. Three disks are needed since OpenShift Data Foundation uses a replica count of 3 to maintain the high availability. So the amount of storage consumed is three times the usable space.



NOTE

Usable space may vary when encryption is enabled or replica 2 pools are being used.

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

In order to scale up an OpenShift Data Foundation cluster which was created using local storage devices, a new disk needs to be added to the storage node. It is recommended to have the new disks of the same size as used earlier during the deployment as OpenShift Data Foundation does not support heterogeneous disks/OSD's.

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



NOTE

Flexible scaling features get enabled at the time of deployment and can not be enabled or disabled later on.

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 nodes where the new OSD pods are running.

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

<OSD-pod-name>

Is the name of the OSD pod.
For example:

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

Example output:

```
NODE
compute-1
```

- 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.

9.2. SCALING OUT STORAGE CAPACITY ON A IBM POWER CLUSTER

OpenShift Data Foundation is highly scalable. It can be scaled out by adding new nodes with required storage and enough hardware resources in terms of CPU and RAM. Practically there is no limit on the number of nodes which can be added but from the support perspective 2000 nodes is the limit for OpenShift Data Foundation.

Scaling out storage capacity can be broken down into two steps:

- Adding new node
- Scaling up the storage capacity



NOTE

OpenShift Data Foundation does not support heterogeneous OSD/Disk sizes.

9.2.1. Adding a node using a local storage device on IBM Power

You can add nodes to increase the storage capacity when existing worker nodes are already running at their maximum supported OSDs or when there are not enough resources to add new OSDs on the existing nodes.

Add nodes in the multiple of 3, each of them in different failure domains. Though it is recommended to add nodes in multiples of 3 nodes, you have the flexibility to add one node at a time in flexible scaling deployment. See Knowledgebase article [Verify if flexible scaling is enabled](#)



NOTE

OpenShift Data Foundation does not support heterogeneous disk size and types. The new nodes to be added should have the disk of the same type and size which was used during initial OpenShift Data Foundation deployment.

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) that are in **Pending** state.

```
$ oc get csr
```

- b. Approve all the required 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 9.1. 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

1. Execute the following command the terminal 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. On the OpenShift web console, click **Workloads** → **Pods**, confirm that at least the following pods on the new node are in **Running** state:
 - **csi-cephfsplugin-***
 - **csi-rbdplugin-***

9.2.2. Scaling up storage capacity

To scale up storage capacity, see [Scaling up storage by adding capacity](#) .