



Red Hat OpenShift Data Foundation 4.15

Deploying OpenShift Data Foundation on any platform

Instructions on deploying OpenShift Data Foundation on any platform including virtualized and cloud environments.

Red Hat OpenShift Data Foundation 4.15 Deploying OpenShift Data Foundation on any platform

Instructions on deploying OpenShift Data Foundation on any platform including virtualized and cloud environments.

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

Read this document for instructions about how to install Red Hat OpenShift Data Foundation to use local storage on any platform.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
PREFACE	5
CHAPTER 1. PREPARING TO DEPLOY OPENSIFT DATA FOUNDATION	6
1.1. REQUIREMENTS FOR INSTALLING OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES	6
CHAPTER 2. DEPLOY OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES	8
2.1. INSTALLING LOCAL STORAGE OPERATOR	8
2.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR	8
2.3. CREATING OPENSIFT DATA FOUNDATION CLUSTER ON ANY PLATFORM	10
2.4. VERIFYING OPENSIFT DATA FOUNDATION DEPLOYMENT	15
2.4.1. Verifying the state of the pods	16
2.4.2. Verifying the OpenShift Data Foundation cluster is healthy	18
2.4.3. Verifying the Multicloud Object Gateway is healthy	18
2.4.4. Verifying that the specific storage classes exist	18
2.4.5. Verifying the Multus networking	18
CHAPTER 3. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY	21
3.1. INSTALLING LOCAL STORAGE OPERATOR	21
3.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR	21
3.3. CREATING A STANDALONE MULTICLOUD OBJECT GATEWAY	23
CHAPTER 4. VIEW OPENSIFT DATA FOUNDATION TOPOLOGY	26
CHAPTER 5. UNINSTALLING OPENSIFT DATA FOUNDATION	27
5.1. UNINSTALLING OPENSIFT DATA FOUNDATION IN INTERNAL MODE	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

Red Hat OpenShift Data Foundation supports deployment on any platform that you provision including bare metal, virtualized, and cloud environments. Both internal and external OpenShift Data Foundation clusters are supported on these environments. See [Planning your deployment](#) and [Preparing to deploy OpenShift Data Foundation](#) for more information about deployment requirements.

To deploy OpenShift Data Foundation, follow the appropriate deployment process based on your requirement:

- Internal mode
 - [Deploy using local storage devices](#)
 - [Deploy standalone Multicloud Object Gateway component](#)
- [External mode](#)

CHAPTER 1. PREPARING TO DEPLOY OPENSIFT DATA FOUNDATION

When you deploy OpenShift Data Foundation on OpenShift Container Platform using the local storage devices on any platform, you can create internal cluster resources. This approach internally provisions base services so that all the applications can access additional storage classes.

You can also deploy OpenShift Data Foundation to use an external Red Hat Ceph Storage cluster and IBM FlashSystem. For instructions, see [Deploying OpenShift Data Foundation in external mode](#) . External mode deployment works on clusters that are detected as non-cloud. If your cluster is not detected correctly, open up a bug in [Bugzilla](#).

Before you begin the deployment of Red Hat OpenShift Data Foundation using a local storage, ensure that you meet the resource requirements. See [Requirements for installing OpenShift Data Foundation using local storage devices](#).

- Optional: If you want to enable cluster-wide encryption using an external Key Management System (KMS) follow these steps:
 - Ensure that you have a valid Red Hat OpenShift Data Foundation Advanced subscription. To know how subscriptions for OpenShift Data Foundation work, see [knowledgebase article on OpenShift Data Foundation subscriptions](#).
 - When the Token authentication method is selected for encryption, refer to [Enabling cluster-wide encryption with the Token authentication using KMS](#).
 - When the Kubernetes authentication method is selected for encryption then refer to [Enabling cluster-wide encryption with KMS using the Kubernetes authentication method](#) .
 - Ensure that you are using signed certificates on your vault servers.

After completing the preparatory steps, perform the following procedures:

1. [Install the Local Storage Operator](#) .
2. [Install the Red Hat OpenShift Data Foundation Operator](#) .
3. [Create the OpenShift Data Foundation cluster on any platform](#) .

1.1. REQUIREMENTS FOR INSTALLING OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES

Node requirements

The cluster must consist of at least three OpenShift Container Platform worker or infrastructure nodes with locally attached-storage devices on each of them.

- Each of the three selected nodes must have at least one raw block device available. OpenShift Data Foundation uses the one or more available raw block devices.
- The devices you use must be empty, the disks must not include Physical Volumes (PVs), Volume Groups (VGs), or Logical Volumes (LVs) remaining on the disk.

For more information, see the *Resource requirements* section in the [Planning guide](#).

Disaster recovery requirements

Disaster Recovery features supported by Red Hat OpenShift Data Foundation require all of the following prerequisites to successfully implement a disaster recovery solution:

- A valid Red Hat OpenShift Data Foundation Advanced subscription.
- A valid Red Hat Advanced Cluster Management (RHACM) for Kubernetes subscription.

To know in detail how subscriptions for OpenShift Data Foundation work, see [knowledgebase article on OpenShift Data Foundation subscriptions](#).

For detailed disaster recovery solution requirements, see [Configuring OpenShift Data Foundation Disaster Recovery for OpenShift Workloads](#) guide, and *Requirements and recommendations* section of the [Install guide](#) in Red Hat Advanced Cluster Management for Kubernetes documentation.

Minimum starting node requirements

An OpenShift Data Foundation cluster is deployed with a minimum configuration when the resource requirement for a standard deployment is not met.

For more information, see the *Resource requirements* section in the [Planning guide](#).

CHAPTER 2. DEPLOY OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES

You can deploy OpenShift Data Foundation on any platform including virtualized and cloud environments where OpenShift Container Platform is already installed.

Also, it is possible to deploy only the Multicloud Object Gateway (MCG) component with OpenShift Data Foundation. For more information, see [Deploy standalone Multicloud Object Gateway](#).

Perform the following steps to deploy OpenShift Data Foundation:

1. [Install the Local Storage Operator](#).
2. [Install the Red Hat OpenShift Data Foundation Operator](#).
3. [Create an OpenShift Data Foundation cluster on any platform](#).

2.1. INSTALLING LOCAL STORAGE OPERATOR

Install the Local Storage Operator from the Operator Hub before creating Red Hat OpenShift Data Foundation clusters on local storage devices.

Procedure

1. Log in to the OpenShift Web Console.
2. Click **Operators → OperatorHub**.
3. Type **local storage** in the **Filter by keyword** box to find the **Local Storage Operator** from the list of operators, and click on it.
4. Set the following options on the **Install Operator** page:
 - a. Update channel as **stable**.
 - b. Installation mode as **A specific namespace on the cluster**
 - c. Installed Namespace as **Operator recommended namespace openshift-local-storage**.
 - d. Update approval as **Automatic**.
5. Click **Install**.

Verification steps

- Verify that the Local Storage Operator shows a green tick indicating successful installation.

2.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR

You can install Red Hat OpenShift Data Foundation Operator using the Red Hat OpenShift Container Platform Operator Hub.

Prerequisites

- Access to an OpenShift Container Platform cluster using an account with **cluster-admin** and operator installation permissions.
- You must have at least three worker or infrastructure nodes in the Red Hat OpenShift Container Platform cluster.
- For additional resource requirements, see the [Planning your deployment](#) guide.



IMPORTANT

- When you need to override the cluster-wide default node selector for OpenShift Data Foundation, you can use the following command to specify a blank node selector for the **openshift-storage** namespace (create **openshift-storage** namespace in this case):


```
$ oc annotate namespace openshift-storage openshift.io/node-selector=
```
- Taint a node as **infra** to ensure only Red Hat OpenShift Data Foundation resources are scheduled on that node. This helps you save on subscription costs. For more information, see the *How to use dedicated worker nodes for Red Hat OpenShift Data Foundation* section in the [Managing and Allocating Storage Resources](#) guide.

Procedure

1. Log in to the OpenShift Web Console.
2. Click **Operators** → **OperatorHub**.
3. Scroll or type **OpenShift Data Foundation** into the **Filter by keyword** box to find the **OpenShift Data Foundation** Operator.
4. Click **Install**.
5. Set the following options on the **Install Operator** page:
 - a. Update Channel as **stable-4.15**.
 - b. Installation Mode as **A specific namespace on the cluster**
 - c. Installed Namespace as **Operator recommended namespace openshift-storage**. If Namespace **openshift-storage** does not exist, it is created during the operator installation.
 - d. Select Approval Strategy as **Automatic** or **Manual**.
If you select **Automatic** updates, then the Operator Lifecycle Manager (OLM) automatically upgrades the running instance of your Operator without any intervention.

If you select **Manual** updates, then the OLM creates an update request. As a cluster administrator, you must then manually approve that update request to update the Operator to a newer version.
 - e. Ensure that the **Enable** option is selected for the **Console plugin**.
 - f. Click **Install**.

Verification steps

- After the operator is successfully installed, a pop-up with a message, **Web console update is available** appears on the user interface. Click **Refresh web console** from this pop-up for the console changes to reflect.
- In the Web Console:
 - Navigate to Installed Operators and verify that the **OpenShift Data Foundation Operator** shows a green tick indicating successful installation.
 - Navigate to **Storage** and verify if the **Data Foundation** dashboard is available.

2.3. CREATING OPENSIFT DATA FOUNDATION CLUSTER ON ANY PLATFORM

Prerequisites

- Ensure that all the requirements in the [Requirements for installing OpenShift Data Foundation using local storage devices](#) section are met.
- If you want to use multus networking, you must create network attachment definitions (NADs) before deployment which is later attached to the cluster. For more information, see [Multi network plug-in \(Multus\) support](#) and [Creating network attachment definitions](#).

Procedure

1. In the OpenShift Web Console, click **Operators → Installed Operators** to view all the installed operators.
Ensure that the **Project** selected is **openshift-storage**.
2. Click on the **OpenShift Data Foundation** operator, and then click **Create StorageSystem**.
3. In the **Backing storage** page, perform the following:
 - a. Select **Full Deployment** for the **Deployment type** option.
 - b. Select the **Create a new StorageClass using the local storage devices** option.
 - c. Optional: Select **Use external PostgreSQL** checkbox to use an external PostgreSQL **[Technology preview]**.
This provides high availability solution for Multicloud Object Gateway where the PostgreSQL pod is a single point of failure.
 - i. Provide the following connection details:
 - **Username**
 - **Password**
 - **Server name and Port**
 - **Database name**
 - ii. Select **Enable TLS/SSL** checkbox to enable encryption for the Postgres server.
 - d. Click **Next**.



IMPORTANT

You are prompted to install the Local Storage Operator if it is not already installed. Click **Install**, and follow the procedure as described in [Installing Local Storage Operator](#).

4. In the **Create local volume set** page, provide the following information:
 - a. Enter a name for the **LocalVolumeSet** and the **StorageClass**.
The local volume set name appears as the default value for the storage class name. You can change the name.
 - b. Select one of the following:
 - **Disks on all nodes**
Uses the available disks that match the selected filters on all the nodes.
 - **Disks on selected nodes**
Uses the available disks that match the selected filters only on the selected nodes.



IMPORTANT

- The flexible scaling feature is enabled only when the storage cluster that you created with three or more nodes are spread across fewer than the minimum requirement of three availability zones.
For information about flexible scaling, see [knowledgebase article on Scaling OpenShift Data Foundation cluster using YAML when flexible scaling is enabled](#).
- Flexible scaling features get enabled at the time of deployment and can not be enabled or disabled later on.
- If the nodes selected do not match the OpenShift Data Foundation cluster requirement of an aggregated 30 CPUs and 72 GiB of RAM, a minimal cluster is deployed if at least 24 CPUs and 72 GiB of RAM is available.
For minimum starting node requirements, see the [Resource requirements](#) section in the *Planning* guide.

- c. From the available list of **Disk Type**, select **SSD/NVMe**.
- d. Expand the **Advanced** section and set the following options:

Volume Mode	Block is selected as the default value.
Device Type	Select one or more device types from the dropdown list.
Disk Size	Set a minimum size of 100GB for the device and maximum available size of the device that needs to be included.


Maximum Disks Limit	This indicates the maximum number of Persistent Volumes (PVs) that you can create on a node. If this field is left empty, then PVs are created for all the available disks on the matching nodes.
---------------------	---

- e. Click **Next**.
A pop-up to confirm the creation of LocalVolumeSet is displayed.
 - f. Click **Yes** to continue.
5. In the **Capacity and nodes** page, configure the following:
 - a. **Available raw capacity** is populated with the capacity value based on all the attached disks associated with the storage class. This takes some time to show up. The **Selected nodes** list shows the nodes based on the storage class.
 - b. In the **Configure performance** section, select one of the following performance profiles:
 - **Lean**
Use this in a resource constrained environment with minimum resources that are lower than the recommended. This profile minimizes resource consumption by allocating fewer CPUs and less memory.
 - **Balanced (default)**
Use this when recommended resources are available. This profile provides a balance between resource consumption and performance for diverse workloads.
 - **Performance**
Use this in an environment with sufficient resources to get the best performance. This profile is tailored for high performance by allocating ample memory and CPUs to ensure optimal execution of demanding workloads.



NOTE

You have the option to configure the performance profile even after the deployment using the **Configure performance** option from the options menu of the **StorageSystems** tab.



IMPORTANT

Before selecting a resource profile, make sure to check the current availability of resources within the cluster. Opting for a higher resource profile in a cluster with insufficient resources might lead to installation failures.

For more information about resource requirements, see [Resource requirement for performance profiles](#).

- c. Optional: Select the **Taint nodes** checkbox to dedicate the selected nodes for OpenShift Data Foundation.
 - d. Click **Next**.
6. Optional: In the **Security and network** page, configure the following based on your requirement:

- a. To enable encryption, select **Enable data encryption for block and file storage**
- b. Select one or both of the following **Encryption level**:
 - **Cluster-wide encryption**
Encrypts the entire cluster (block and file).
 - **StorageClass encryption**
Creates encrypted persistent volume (block only) using encryption enabled storage class.
- c. Optional: Select the **Connect to an external key management service** checkbox. This is optional for cluster-wide encryption.
 - i. From the **Key Management Service Provider** drop-down list, either select **Vault** or **Thales CipherTrust Manager (using KMIP)**. If you selected **Vault**, go to the next step. If you selected **Thales CipherTrust Manager (using KMIP)**, go to step iii.
 - ii. Select an **Authentication Method**.

Using Token authentication method

- Enter a unique **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Token**.
- Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:
 - Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
 - Optional: Enter **TLS Server Name** and **Vault Enterprise Namespace**
 - Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key**.
 - Click **Save** and skip to step iv.

Using Kubernetes authentication method

- Enter a unique Vault **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Role** name.
- Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:
 - Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
 - Optional: Enter **TLS Server Name** and **Authentication Path** if applicable.
 - Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key**.
 - Click **Save** and skip to step iv.

iii. To use **Thales CipherTrust Manager (using KMIP)** as the KMS provider, follow the

- iii. To use **Thales CipherTrust Manager (using KMIP)** as the KMIP provider, follow the steps below:
 - A. Enter a unique **Connection Name** for the Key Management service within the project.
 - B. In the **Address** and **Port** sections, enter the IP of Thales CipherTrust Manager and the port where the KMIP interface is enabled. For example:
 - **Address:** 123.34.3.2
 - **Port:** 5696
 - C. Upload the **Client Certificate**, **CA certificate**, and **Client Private Key**.
 - D. If StorageClass encryption is enabled, enter the Unique Identifier to be used for encryption and decryption generated above.
 - E. The **TLS Server** field is optional and used when there is no DNS entry for the KMIP endpoint. For example, **kmip_all_<port>.ciphertrustmanager.local**.

iv. Select a **Network**.

d. Select one of the following:

- **Default (SDN)**
If you are using a single network.
- **Custom (Multus)**
If you are using multiple network interfaces.
 - i. Select a **Public Network Interface** from the dropdown.
 - ii. Select a **Cluster Network Interface** from the dropdown.



NOTE

If you are using only one additional network interface, select the single **NetworkAttachmentDefinition**, that is, **ocs-public-cluster** for the Public Network Interface and leave the Cluster Network Interface blank.

e. Click **Next**.

7. In the **Review and create** page, review the configuration details.

To modify any configuration settings, click **Back** to go back to the previous configuration page.

8. Click **Create StorageSystem**.



NOTE

When your deployment has five or more nodes, racks, or rooms, and when there are five or more number of failure domains present in the deployment, you can configure Ceph monitor counts based on the number of racks or zones. An alert is displayed in the notification panel or Alert Center of the OpenShift Web Console to indicate the option to increase the number of Ceph monitor counts. You can use the **Configure** option in the alert to configure the Ceph monitor counts. For more information, see [Resolving low Ceph monitor count alert](#).

Verification steps

- To verify the final Status of the installed storage cluster:
 - a. In the OpenShift Web Console, navigate to **Installed Operators → OpenShift Data Foundation → Storage System**
 - b. Click **ocs-storagecluster-storagesystem → Resources**.
 - c. Verify that the **Status** of the **StorageCluster** is **Ready** and has a green tick mark next to it.
- To verify if the flexible scaling is enabled on your storage cluster, perform the following steps (for arbiter mode, flexible scaling is disabled):
 1. In the OpenShift Web Console, navigate to **Installed Operators → OpenShift Data Foundation → Storage System**
 2. Click **ocs-storagecluster-storagesystem → Resources → ocs-storagecluster**.
 3. In the YAML tab, search for the keys **flexibleScaling** in the **spec** section and **failureDomain** in the **status** section. If **flexible scaling** is true and **failureDomain** is set to host, flexible scaling feature is enabled:

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

- To verify that all the components for OpenShift Data Foundation are successfully installed, see [Verifying your OpenShift Data Foundation installation](#).
- To verify the multi networking (Multus), see [Verifying the Multus networking](#).

Additional resources

- To expand the capacity of the initial cluster, see the [Scaling Storage](#) guide and follow the instructions in the "Scaling storage of bare metal OpenShift Data Foundation cluster" section.

2.4. VERIFYING OPENSIFT DATA FOUNDATION DEPLOYMENT

To verify that OpenShift Data Foundation is deployed correctly:

1. [Verify the state of the pods](#).
2. [Verify that the OpenShift Data Foundation cluster is healthy](#).

3. [Verify that the Multicloud Object Gateway is healthy](#) .
4. [Verify that the OpenShift Data Foundation specific storage classes exist](#) .
5. [Verify the Multus networking](#).

2.4.1. Verifying the state of the pods

Procedure

1. Click **Workloads** → **Pods** from the OpenShift Web Console.
2. 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.

For more information on the expected number of pods for each component and how it varies depending on the number of nodes, see [Table 2.1, “Pods corresponding to OpenShift Data Foundation cluster”](#).

3. Set filter for Running and Completed pods to verify that the following pods are in **Running** and **Completed** state:

Table 2.1. Pods corresponding to OpenShift Data Foundation cluster

Component	Corresponding pods
OpenShift Data Foundation Operator	<ul style="list-style-type: none"> ● ocs-operator-* (1 pod on any storage node) ● ocs-metrics-exporter-* (1 pod on any storage node) ● odf-operator-controller-manager-* (1 pod on any storage node) ● odf-console-* (1 pod on any storage node) ● csi-addons-controller-manager-* (1 pod on any storage node)
Rook-ceph Operator	<p>rook-ceph-operator-*</p> <p>(1 pod on any storage node)</p>

Component	Corresponding pods
Multicloud Object Gateway	<ul style="list-style-type: none"> ● noobaa-operator-* (1 pod on any storage node) ● noobaa-core-* (1 pod on any storage node) ● noobaa-db-pg-* (1 pod on any storage node) ● noobaa-endpoint-* (1 pod on any storage node)
MON	<p>rook-ceph-mon-*</p> <p>(3 pods distributed across storage nodes)</p>
MGR	<p>rook-ceph-mgr-*</p> <p>(1 pod on any storage node)</p>
MDS	<p>rook-ceph-mds-ocs-storagecluster-cephfilesystem-*</p> <p>(2 pods distributed across storage nodes)</p>
RGW	<p>rook-ceph-rgw-ocs-storagecluster-cephobjectstore-* (1 pod on any storage node)</p>
CSI	<ul style="list-style-type: none"> ● cephfs <ul style="list-style-type: none"> ○ csi-cephfsplugin-* (1 pod on each storage node) ○ csi-cephfsplugin-provisioner-* (2 pods distributed across storage nodes) ● rbd <ul style="list-style-type: none"> ○ csi-rbdplugin-* (1 pod on each storage node) ○ csi-rbdplugin-provisioner-* (2 pods distributed across storage nodes)
rook-ceph-crashcollector	<p>rook-ceph-crashcollector-*</p> <p>(1 pod on each storage node)</p>

Component	Corresponding pods
OSD	<ul style="list-style-type: none"> ● rook-ceph-osd-* (1 pod for each device) ● rook-ceph-osd-prepare-ocs-deviceset-* (1 pod for each device)

2.4.2. Verifying the OpenShift Data Foundation cluster is healthy

Procedure

1. In the OpenShift Web Console, click **Storage → Data Foundation**.
2. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.
3. In the **Status** card of the **Block and File** tab, verify that the *Storage Cluster* has a green tick.
4. In the **Details** card, verify that the cluster information is displayed.

For more information on the health of the OpenShift Data Foundation cluster using the **Block and File** dashboard, see [Monitoring OpenShift Data Foundation](#).

2.4.3. Verifying the Multicloud Object Gateway is healthy

Procedure

1. In the OpenShift Web Console, click **Storage → Data Foundation**.
2. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.
 - a. In the **Status** card of the **Object** tab, verify that both *Object Service* and *Data Resiliency* have a green tick.
 - b. In the **Details** card, verify that the MCG information is displayed.

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

2.4.4. Verifying that the specific storage classes exist

Procedure

1. Click **Storage → Storage Classes** from the left pane of the OpenShift Web Console.
2. Verify that the following storage classes are created with the OpenShift Data Foundation cluster creation:

2.4.5. Verifying the Multus networking

To determine if Multus is working in your cluster, verify the Multus networking.

Procedure

Based on your Network configuration choices, the OpenShift Data Foundation operator will do one of the following:

- If only a single NetworkAttachmentDefinition (for example, **ocs-public-cluster**) was selected for the Public Network Interface, then the traffic between the application pods and the OpenShift Data Foundation cluster will happen on this network. Additionally the cluster will be self configured to also use this network for the replication and rebalancing traffic between OSDs.
- If both NetworkAttachmentDefinitions (for example, **ocs-public** and **ocs-cluster**) were selected for the Public Network Interface and the Cluster Network Interface respectively during the Storage Cluster installation, then client storage traffic will be on the public network and cluster network for the replication and rebalancing traffic between OSDs.

To verify the network configuration is correct, complete the following:

In the OpenShift console, navigate to **Installed Operators** → **OpenShift Data Foundation** → **Storage System** → **ocs-storagecluster-storagesystem** → **Resources** → **ocs-storagecluster**.

In the YAML tab, search for **network** in the **spec** section and ensure the configuration is correct for your network interface choices. This example is for separating the client storage traffic from the storage replication traffic.

Sample output:

```
[..]
spec:
  [..]
  network:
    ipFamily: IPv4
    provider: multus
    selectors:
      cluster: openshift-storage/ocs-cluster
      public: openshift-storage/ocs-public
  [..]
```

To verify the network configuration is correct using the command line interface, run the following commands:

```
$ oc get storagecluster ocs-storagecluster \
-n openshift-storage \
-o=jsonpath='{.spec.network}'
```

Sample output:

```
{"ipFamily":"IPv4","provider":"multus","selectors":{"cluster":"openshift-storage/ocs-cluster","public":"openshift-storage/ocs-public"}}
```

Confirm the OSD pods are using correct network

In the **openshift-storage** namespace use one of the OSD pods to verify the pod has connectivity to the correct networks. This example is for separating the client storage traffic from the storage replication traffic.

**NOTE**

Only the OSD pods will connect to both Multus public and cluster networks if both are created. All other OCS pods will connect to the Multus public network.

```
$ oc get -n openshift-storage $(oc get pods -n openshift-storage -o name -l app=rook-ceph-osd | grep 'osd-0') -o=jsonpath='{.metadata.annotations.k8s\.v1\.cni\.cncf\.io/network-status}'
```

Sample output:

```
[{
  "name": "openshift-sdn",
  "interface": "eth0",
  "ips": [
    "10.129.2.30"
  ],
  "default": true,
  "dns": {}
},{
  "name": "openshift-storage/ocs-cluster",
  "interface": "net1",
  "ips": [
    "192.168.2.1"
  ],
  "mac": "e2:04:c6:81:52:f1",
  "dns": {}
},{
  "name": "openshift-storage/ocs-public",
  "interface": "net2",
  "ips": [
    "192.168.1.1"
  ],
  "mac": "ee:a0:b6:a4:07:94",
  "dns": {}
}]
```

To confirm the OSD pods are using correct network using the command line interface, run the following command (requires the jq utility):

```
$ oc get -n openshift-storage $(oc get pods -n openshift-storage -o name -l app=rook-ceph-osd | grep 'osd-0') -o=jsonpath='{.metadata.annotations.k8s\.v1\.cni\.cncf\.io/network-status}' | jq -r '.[].name'
```

Sample output:

```
openshift-sdn
openshift-storage/ocs-cluster
openshift-storage/ocs-public
```


CHAPTER 3. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY

Deploying only the Multicloud Object Gateway component with the OpenShift Data Foundation provides the flexibility in deployment and helps to reduce the resource consumption. Use this section to deploy only the standalone Multicloud Object Gateway component, which involves the following steps:

- Installing the Local Storage Operator.
- Installing Red Hat OpenShift Data Foundation Operator
- Creating standalone Multicloud Object Gateway

3.1. INSTALLING LOCAL STORAGE OPERATOR

Install the Local Storage Operator from the Operator Hub before creating Red Hat OpenShift Data Foundation clusters on local storage devices.

Procedure

1. Log in to the OpenShift Web Console.
2. Click **Operators** → **OperatorHub**.
3. Type **local storage** in the **Filter by keyword** box to find the **Local Storage Operator** from the list of operators, and click on it.
4. Set the following options on the **Install Operator** page:
 - a. Update channel as **stable**.
 - b. Installation mode as **A specific namespace on the cluster**
 - c. Installed Namespace as **Operator recommended namespace openshift-local-storage**.
 - d. Update approval as **Automatic**.
5. Click **Install**.

Verification steps

- Verify that the Local Storage Operator shows a green tick indicating successful installation.

3.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR

You can install Red Hat OpenShift Data Foundation Operator using the Red Hat OpenShift Container Platform Operator Hub.

Prerequisites

- Access to an OpenShift Container Platform cluster using an account with **cluster-admin** and operator installation permissions.

- You must have at least three worker or infrastructure nodes in the Red Hat OpenShift Container Platform cluster.
- For additional resource requirements, see the [Planning your deployment](#) guide.



IMPORTANT

- When you need to override the cluster-wide default node selector for OpenShift Data Foundation, you can use the following command to specify a blank node selector for the **openshift-storage** namespace (create **openshift-storage** namespace in this case):

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

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

Procedure

1. Log in to the OpenShift Web Console.
2. Click **Operators** → **OperatorHub**.
3. Scroll or type **OpenShift Data Foundation** into the **Filter by keyword** box to find the **OpenShift Data Foundation** Operator.
4. Click **Install**.
5. Set the following options on the **Install Operator** page:
 - a. Update Channel as **stable-4.15**.
 - b. Installation Mode as **A specific namespace on the cluster**
 - c. Installed Namespace as **Operator recommended namespace openshift-storage**. If Namespace **openshift-storage** does not exist, it is created during the operator installation.
 - d. Select Approval Strategy as **Automatic** or **Manual**.
If you select **Automatic** updates, then the Operator Lifecycle Manager (OLM) automatically upgrades the running instance of your Operator without any intervention.

If you select **Manual** updates, then the OLM creates an update request. As a cluster administrator, you must then manually approve that update request to update the Operator to a newer version.
 - e. Ensure that the **Enable** option is selected for the **Console plugin**.
 - f. Click **Install**.

Verification steps

- After the operator is successfully installed, a pop-up with a message, **Web console update is available** appears on the user interface. Click **Refresh web console** from this pop-up for the console changes to reflect.
- In the Web Console:
 - Navigate to Installed Operators and verify that the **OpenShift Data Foundation Operator** shows a green tick indicating successful installation.
 - Navigate to **Storage** and verify if the **Data Foundation** dashboard is available.

3.3. CREATING A STANDALONE MULTICLOUD OBJECT GATEWAY

You can create only the standalone Multicloud Object Gateway component while deploying OpenShift Data Foundation.

Prerequisites

- Ensure that the OpenShift Data Foundation Operator is installed.

Procedure

1. In the OpenShift Web Console, click **Operators** → **Installed Operators** to view all the installed operators.
Ensure that the **Project** selected is **openshift-storage**.
2. Click **OpenShift Data Foundation** operator and then click **Create StorageSystem**.
3. In the **Backing storage** page, select the following:
 - a. Select **Multicloud Object Gateway** for **Deployment type**.
4. Optional: Select the **Connect to an external key management service** checkbox. This is optional for cluster-wide encryption.
 - a. From the **Key Management Service Provider** drop-down list, either select **Vault** or **Thales CipherTrust Manager (using KMIP)**. If you selected **Vault**, go to the next step. If you selected **Thales CipherTrust Manager (using KMIP)**, go to step iii.
 - b. Select an **Authentication Method**.

Using Token authentication method

- Enter a unique **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Token**.
- Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:
 - Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
 - Optional: Enter **TLS Server Name** and **Vault Enterprise Namespace**
 - Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key**.

- Click **Save** and skip to step iv.

Using Kubernetes authentication method

- Enter a unique Vault **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Role** name.
 - Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:
 - Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
 - Optional: Enter **TLS Server Name** and **Authentication Path** if applicable.
 - Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key**.
 - Click **Save** and skip to step iv.
- c. To use **Thales CipherTrust Manager (using KMIP)** as the KMS provider, follow the steps below:
- i. Enter a unique **Connection Name** for the Key Management service within the project.
 - ii. In the **Address** and **Port** sections, enter the IP of Thales CipherTrust Manager and the port where the KMIP interface is enabled. For example:
 - **Address:** 123.34.3.2
 - **Port:** 5696
 - iii. Upload the **Client Certificate**, **CA certificate**, and **Client Private Key**.
 - iv. If StorageClass encryption is enabled, enter the Unique Identifier to be used for encryption and decryption generated above.
 - v. The **TLS Server** field is optional and used when there is no DNS entry for the KMIP endpoint. For example, **kmip_all_<port>.ciphertrustmanager.local**.
- d. Select a **Network**.
- e. Click **Next**.
5. In the **Review and create** page, review the configuration details:
To modify any configuration settings, click **Back**.
6. Click **Create StorageSystem**.

Verification steps

Verifying that the OpenShift Data Foundation cluster is healthy

1. In the OpenShift Web Console, click **Storage** → **Data Foundation**.
2. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.

- a. In the **Status card** of the **Object** tab, verify that both *Object Service* and *Data Resiliency* have a green tick.
- b. In the **Details** card, verify that the MCG information is displayed.

Verifying the state of the pods

1. Click **Workloads** → **Pods** from the OpenShift Web Console.
2. Select **openshift-storage** from the **Project** drop-down list and verify that the following pods are in **Running** state.



NOTE

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

Component	Corresponding pods
OpenShift Data Foundation Operator	<ul style="list-style-type: none"> ● ocs-operator-* (1 pod on any storage node) ● ocs-metrics-exporter-* (1 pod on any storage node) ● odf-operator-controller-manager-* (1 pod on any storage node) ● odf-console-* (1 pod on any storage node) ● csi-addons-controller-manager-* (1 pod on any storage node)
Rook-ceph Operator	<p>rook-ceph-operator-*</p> <p>(1 pod on any storage node)</p>
Multicloud Object Gateway	<ul style="list-style-type: none"> ● noobaa-operator-* (1 pod on any storage node) ● noobaa-core-* (1 pod on any storage node) ● noobaa-db-pg-* (1 pod on any storage node) ● noobaa-endpoint-* (1 pod on any storage node)

CHAPTER 4. VIEW OPENSIFT DATA FOUNDATION TOPOLOGY

The topology shows the mapped visualization of the OpenShift Data Foundation storage cluster at various abstraction levels and also lets you to interact with these layers. The view also shows how the various elements compose the Storage cluster altogether.

Procedure

1. On the OpenShift Web Console, navigate to **Storage → Data Foundation → Topology**.
The view shows the storage cluster and the zones inside it. You can see the nodes depicted by circular entities within the zones, which are indicated by dotted lines. The label of each item or resource contains basic information such as status and health or indication for alerts.
2. Choose a node to view node details on the right-hand panel. You can also access resources or deployments within a node by clicking on the search/preview decorator icon.
3. To view deployment details
 - a. Click the preview decorator on a node. A modal window appears above the node that displays all of the deployments associated with that node along with their statuses.
 - b. Click the **Back to main view** button in the model's upper left corner to close and return to the previous view.
 - c. Select a specific deployment to see more information about it. All relevant data is shown in the side panel.
4. Click the **Resources** tab to view the pods information. This tab provides a deeper understanding of the problems and offers granularity that aids in better troubleshooting.
5. Click the pod links to view the pod information page on OpenShift Container Platform. The link opens in a new window.

CHAPTER 5. UNINSTALLING OPENSIFT DATA FOUNDATION

5.1. UNINSTALLING OPENSIFT DATA FOUNDATION IN INTERNAL MODE

To uninstall OpenShift Data Foundation in Internal mode, refer to the [knowledge base article on Uninstalling OpenShift Data Foundation](#).