



Red Hat OpenShift Data Foundation 4.9

Deploying OpenShift Data Foundation using bare metal infrastructure

Instructions on deploying OpenShift Data Foundation using local storage on bare
metal infrastructure

Red Hat OpenShift Data Foundation 4.9 Deploying OpenShift Data Foundation using bare metal infrastructure

Instructions on deploying OpenShift Data Foundation using local storage on bare metal infrastructure

Legal Notice

Copyright © 2022 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 bare metal infrastructure.

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
1.2. ENABLING FILE SYSTEM ACCESS FOR CONTAINERS ON RED HAT ENTERPRISE LINUX BASED NODES	7
1.3. ENABLING KEY VALUE BACKEND PATH AND POLICY IN VAULT	8
CHAPTER 2. DEPLOY OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES	10
2.1. INSTALLING LOCAL STORAGE OPERATOR	10
2.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR	10
2.3. CREATING MULTUS NETWORKS [TECHNOLOGY PREVIEW]	12
2.3.1. Creating network attachment definitions	12
2.4. CREATING OPENSIFT DATA FOUNDATION CLUSTER ON BARE METAL	14
2.5. VERIFYING OPENSIFT DATA FOUNDATION DEPLOYMENT	18
2.5.1. Verifying the state of the pods	18
2.5.2. Verifying the OpenShift Data Foundation cluster is healthy	20
2.5.3. Verifying the Multicloud Object Gateway is healthy	20
2.5.4. Verifying that the OpenShift Data Foundation specific storage classes exist	20
2.5.5. Verifying the Multus networking	21
CHAPTER 3. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY	24
3.1. INSTALLING LOCAL STORAGE OPERATOR	24
3.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR	24
3.3. CREATING STANDALONE MULTICLOUD OBJECT GATEWAY	26
CHAPTER 4. UNINSTALLING OPENSIFT DATA FOUNDATION	29
4.1. UNINSTALLING OPENSIFT DATA FOUNDATION IN INTERNAL MODE	29

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:

- For simple comments on specific passages:
 1. Make sure you are viewing the documentation in the *Multi-page HTML* format. In addition, ensure you see the **Feedback** button in the upper right corner of the document.
 2. Use your mouse cursor to highlight the part of text that you want to comment on.
 3. Click the **Add Feedback** pop-up that appears below the highlighted text.
 4. Follow the displayed instructions.
- For submitting more complex 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 4.9 supports deployment on existing Red Hat OpenShift Container Platform (RHOCP) bare metal clusters in connected or disconnected environments along with out-of-the-box support for proxy environments.

Both internal and external OpenShift Data Foundation clusters are supported on bare metal. 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 local storage devices, you can create internal cluster resources. This approach internally provisions base services. Then, all applications can access additional storage classes.

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

- Enable file system access on Red Hat Enterprise Linux based hosts for worker nodes. See [enable file system access for containers on Red Hat Enterprise Linux based nodes](#) .



NOTE

Skip this step for Red Hat Enterprise Linux CoreOS (RHCOS).

1. Optional: If you want to enable cluster-wide encryption using an external Key Management System (KMS):
 - Ensure that a policy with a token exists and the key value backend path in Vault is enabled. See [enabled the key value backend path and policy in Vault](#) .
 - Ensure that you are using signed certificates on your Vault servers.

After you have addressed the above, follow the below steps in the order given:

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

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 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 to be used by OpenShift Data Foundation.
- 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.

Regional-DR requirements [Developer Preview]

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

- A valid Red Hat OpenShift Data Foundation Advanced entitlement
- A valid Red Hat Advanced Cluster Management for Kubernetes subscription

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

For detailed requirements, see [Regional-DR requirements](#) and [RHACM requirements](#).

Arbiter stretch cluster requirements [Technology Preview]

In this case, a single cluster is stretched across two zones with a third zone as the location for the arbiter. This is a technology preview feature that is currently intended for deployment in the OpenShift Container Platform on-premises.

For detailed requirements and instructions, see [Configuring OpenShift Data Foundation for Metro-DR stretch cluster](#).



NOTE

Flexible scaling and Arbiter both cannot be enabled at the same time as they have conflicting scaling logic. With Flexible scaling, you can add one node at a time to your OpenShift Data Foundation cluster. Whereas in an Arbiter cluster, you need to add at least one node in each of the two data zones.

Compact mode requirements

OpenShift Data Foundation can be installed on a three-node OpenShift compact bare metal cluster, where all the workloads run on three strong master nodes. There are no worker or storage nodes.

To configure OpenShift Container Platform in compact mode, see [Configuring a three-node cluster](#) and [Delivering a Three-node Architecture for Edge Deployments](#).


Minimum starting node requirements [Technology Preview]

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

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

1.2. ENABLING FILE SYSTEM ACCESS FOR CONTAINERS ON RED HAT ENTERPRISE LINUX BASED NODES

Deploying OpenShift Data Foundation on an OpenShift Container Platform with worker nodes on a Red Hat Enterprise Linux base in a user provisioned infrastructure (UPI) does not automatically provide container access to the underlying Ceph file system.



NOTE

Skip this step for hosts based on Red Hat Enterprise Linux CoreOS (RHCOS).

Procedure

1. Log in to the Red Hat Enterprise Linux based node and open a terminal.
2. For each node in your cluster:

- a. Verify that the node has access to the `rhel-7-server-extras-rpms` repository.

```
# subscription-manager repos --list-enabled | grep rhel-7-server
```

If you do not see both **rhel-7-server-rpms** and **rhel-7-server-extras-rpms** in the output, or if there is no output, run the following commands to enable each repository:

```
# subscription-manager repos --enable=rhel-7-server-rpms
```

```
# subscription-manager repos --enable=rhel-7-server-extras-rpms
```

- b. Install the required packages.

```
# yum install -y policycoreutils container-selinux
```

- c. Persistently enable container use of the Ceph file system in SELinux.

```
# setsebool -P container_use_cephfs on
```

1.3. ENABLING KEY VALUE BACKEND PATH AND POLICY IN VAULT

Prerequisites

- Administrator access to Vault.
- Carefully, choose a unique path name as the backend **path** that follows the naming convention since it cannot be changed later.

Procedure

1. Enable the Key/Value (KV) backend path in Vault.
For Vault KV secret engine API, version 1:

```
$ vault secrets enable -path=odf kv
```

For Vault KV secret engine API, version 2:

```
$ vault secrets enable -path=odf kv-v2
```

2. Create a policy to restrict users to perform a write or delete operation on the secret using the following commands.

```
echo '
path "odf/*" {
  capabilities = ["create", "read", "update", "delete", "list"]
}
path "sys/mounts" {
  capabilities = ["read"]
}' | vault policy write odf -
```

3. Create a token matching the above policy.

```
$ vault token create -policy=odf -format json
```

CHAPTER 2. DEPLOY OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES

You can deploy OpenShift Data Foundation on bare metal infrastructure 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 Local Storage Operator](#).
2. [Install the Red Hat OpenShift Data Foundation Operator](#).
3. [Create OpenShift Data Foundation cluster on bare metal](#).

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 either **4.9** or **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 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 in the command line interface 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 [How to use dedicated worker nodes for Red Hat OpenShift Data Foundation](#) chapter 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-4.9**.
 - 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**.

**NOTE**

We recommend using all default settings. Changing it may result in unexpected behavior. Alter only if you are aware of its result.

Verification steps

- Verify that the **OpenShift Data Foundation** Operator shows a green tick indicating successful installation.
- 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 **Operators** and verify if **OpenShift Data Foundation** is available.

**IMPORTANT**

In case the console plugin option was not automatically enabled after you installed the OpenShift Data Foundation Operator, you need to enable it.

For more information on how to enable the console plugin, see [Enabling the Red Hat OpenShift Data Foundation console plugin](#).

2.3. CREATING MULTUS NETWORKS [TECHNOLOGY PREVIEW]

OpenShift Container Platform uses the Multus CNI plug-in to allow chaining of CNI plug-ins. During cluster installation, you can configure your default pod network. The default network handles all ordinary network traffic for the cluster. You can define an additional network based on the available CNI plug-ins and attach one or more of these networks to your pods. To attach additional network interfaces to a pod, you must create configurations that define how the interfaces are attached. You specify each interface by using a NetworkAttachmentDefinition custom resource (CR). A CNI configuration inside each of the NetworkAttachmentDefinition defines how that interface is created.

OpenShift Data Foundation uses the CNI plug-in called macvlan. Creating a macvlan-based additional network allows pods on a host to communicate with other hosts and pods on those hosts by using a physical network interface. Each pod that is attached to a macvlan-based additional network is provided a unique MAC address.

**IMPORTANT**

Multus support is a Technology Preview feature that is only supported and has been tested on bare metal and VMWare deployments. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information, see [Technology Preview Features Support Scope](#).

2.3.1. Creating network attachment definitions

To utilize Multus, an already working cluster with the correct networking configuration is required, see [Recommended network configuration and requirements for a Multus configuration](#). The newly created

NetworkAttachmentDefinition (NAD) can be selected during the Storage Cluster installation. This is the reason they must be created before the Storage Cluster.

As detailed in the Planning Guide, the Multus networks you create depend on the number of available network interfaces you have for OpenShift Data Foundation traffic. It is possible to separate all of the storage traffic onto one of two interfaces (one interface used for default OpenShift SDN) or to further segregate storage traffic into client storage traffic (public) and storage replication traffic (private or cluster).

The following is an example **NetworkAttachmentDefinition** for all storage traffic, public and cluster, on the same interface. It requires one additional interface on all schedulable nodes (OpenShift default SDN on separate network interface).

```
apiVersion: "k8s.cni.cncf.io/v1"
kind: NetworkAttachmentDefinition
metadata:
  name: ocs-public-cluster
  namespace: openshift-storage
spec:
  config: '{
    "cniVersion": "0.3.1",
    "type": "macvlan",
    "master": "ens2",
    "mode": "bridge",
    "ipam": {
      "type": "whereabouts",
      "range": "192.168.1.0/24"
    }
  }'
```



NOTE

All network interface names must be the same on all the nodes attached to the Multus network (that is, **ens2** for **ocs-public-cluster**).

The following is an example **NetworkAttachmentDefinition** for storage traffic on separate Multus networks, public, for client storage traffic, and cluster, for replication traffic. It requires two additional interfaces on OpenShift nodes hosting OSD pods and one additional interface on all other schedulable nodes (OpenShift default SDN on separate network interface).

```
apiVersion: "k8s.cni.cncf.io/v1"
kind: NetworkAttachmentDefinition
metadata:
  name: ocs-public
  namespace: openshift-storage
spec:
  config: '{
    "cniVersion": "0.3.1",
    "type": "macvlan",
    "master": "ens2",
    "mode": "bridge",
    "ipam": {
      "type": "whereabouts",
```

```

    "range": "192.168.1.0/24"
  }
}'

```

Example **NetworkAttachmentDefinition**:

```

apiVersion: "k8s.cni.cncf.io/v1"
kind: NetworkAttachmentDefinition
metadata:
  name: ocs-cluster
  namespace: openshift-storage
spec:
  config: '{
    "cniVersion": "0.3.1",
    "type": "macvlan",
    "master": "ens3",
    "mode": "bridge",
    "ipam": {
      "type": "whereabouts",
      "range": "192.168.2.0/24"
    }
  }'

```



NOTE

All network interface names must be the same on all the nodes attached to the Multus networks (that is, **ens2** for **ocs-public**, and **ens3** for **ocs-cluster**).

2.4. CREATING OPENSIFT DATA FOUNDATION CLUSTER ON BARE METAL

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 the technology preview feature of multus support, before deployment you must create network attachment definitions (NADs) that 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 the **Create a new StorageClass using the local storage devices** option.
 - b. Expand **Advanced** and select **Full Deployment** for the **Deployment type** option.

c. 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**.

By default, the local volume set name appears for the storage class name. You can change the name.

b. Choose 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. This feature is available only in new deployments of Red Hat OpenShift Data Foundation versions 4.7 and later. Storage clusters upgraded from a previous version to version 4.7 or later do not support flexible scaling. For more information, see *Flexible scaling of OpenShift Data Foundation cluster* in the [New features section of Release Notes](#).
- 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. 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 by default.
Device Type	Select one or more device type 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 PVs that can be created on a node. If this field is left empty, then PVs are created for all the available disks on the matching nodes.
---------------------	--

- 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. 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. Choose 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. Select **Connect to an external key management service** checkbox. This is optional for cluster-wide encryption.
 - i. **Key Management Service Provider** is set to **Vault** by default.
 - ii. Enter Vault **Service Name**, host **Address** of Vault server ('https://<hostname or ip>'), **Port** number and **Token**.
 - iii. Expand **Advanced Settings** to enter additional settings and certificate details based on your Vault configuration:
 - A. Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
 - B. Optional: Enter **TLS Server Name** and **Vault Enterprise Namespace**
 - C. Upload the respective PEM encoded certificate file to provide **CA Certificate**, **Client Certificate** and **Client Private Key**.
 - D. Click **Save**.
 - d. Choose 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**.

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 → ocs-storagecluster-storagesystem → Resources**.
 - b. Verify that **Status** of **StorageCluster** is **Ready** and has a green tick mark next to it.
- To verify if 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 → ocs-storagecluster-storagesystem → Resources → ocs-storagecluster**.
 2. In the YAML tab, search for the keys **flexibleScaling** in **spec** section and **failureDomain** in **status** section. If **flexible scaling** is true and **failureDomain** is set to host, flexible scaling feature is enabled.

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

- 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

Additional resources

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

2.5. 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. [Verifying the Multus networking.](#)

2.5.1. 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 about 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. Click the **Running** and **Completed** tabs 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 worker node) • ocs-metrics-exporter-* (1 pod on any worker node) • odf-operator-controller-manager-* (1 pod on any worker node) • odf-console-* (1 pod on any worker node)

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

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)

2.5.2. Verifying the OpenShift Data Foundation cluster is healthy

Procedure

1. In the OpenShift Web Console, click **Storage** → **OpenShift 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 *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](#).

2.5.3. Verifying the Multicloud Object Gateway is healthy

Procedure

1. In the OpenShift Web Console, click **Storage** → **OpenShift 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](#).

2.5.4. Verifying that the OpenShift Data Foundation 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:
 - **ocs-storagecluster-ceph-rbd**
 - **ocs-storagecluster-cephfs**

- **openshift-storage.noobaa.io**
- **ocs-storagecluster-ceph-rgw**

2.5.5. 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 either **4.9** or **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 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 in the command line interface 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 [How to use dedicated worker nodes for Red Hat OpenShift Data Foundation](#) chapter 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-4.9**.
 - 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**.



NOTE

We recommend using all default settings. Changing it may result in unexpected behavior. Alter only if you are aware of its result.

Verification steps

- Verify that the **OpenShift Data Foundation** Operator shows a green tick indicating successful installation.
- 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 **Operators** and verify if **OpenShift Data Foundation** is available.



IMPORTANT

In case the console plugin option was not automatically enabled after you installed the OpenShift Data Foundation Operator, you need to enable it.

For more information on how to enable the console plugin, see [Enabling the Red Hat OpenShift Data Foundation console plugin](#).

3.3. CREATING STANDALONE MULTICLOUD OBJECT GATEWAY

Use this section to create only the Multicloud Object Gateway component with OpenShift Data Foundation.

Prerequisites

- Ensure that OpenShift Data Foundation Operator is installed.
- (For deploying using local storage devices only) Ensure that Local Storage Operator is installed.
- Ensure that you have a storage class and is set as the default.

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, expand **Advanced**.
4. Select **Multicloud Object Gateway** for **Deployment type**.
5. Click **Next**.
6. Optional: In the **Security** page, select **Connect to an external key management service**
 - a. **Key Management Service Provider** is set to **Vault** by default.
 - b. Enter Vault **Service Name**, host Address of Vault server ('https:// <hostname or ip>'), **Port number**, and **Token**.
 - c. Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:

- i. Enter the Key Value secret path in the **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
 - ii. Optional: Enter **TLS Server Name** and **Vault Enterprise Namespace**
 - iii. Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate**, and **Client Private Key**.
 - iv. Click **Save**.
- d. Click **Next**.
7. In the **Review and create** page, review the configuration details:
To modify any configuration settings, click **Back**.
 8. Click **Create StorageSystem**.

Verification steps

Verifying that the OpenShift Data Foundation cluster is healthy

1. In the OpenShift Web Console, click **Storage** → **OpenShift 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.

Verify 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 worker node) ● ocs-metrics-exporter-* (1 pod on any worker node) ● odf-operator-controller-manager-* (1 pod on any worker node) ● odf-console-* (1 pod on any worker node)

Component	Corresponding pods
Rook-ceph Operator	rook-ceph-operator-* (1 pod on any worker node)
Multicloud Object Gateway	<ul style="list-style-type: none">● noobaa-operator-* (1 pod on any worker node)● noobaa-core-* (1 pod on any worker node)● noobaa-db-pg-* (1 pod on any worker node)● noobaa-endpoint-* (1 pod on any worker node)

CHAPTER 4. UNINSTALLING OPENSIFT DATA FOUNDATION

4.1. 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](#).