



Red Hat OpenShift Data Foundation 4.9

Deploying OpenShift Data Foundation using Amazon Web Services

Instructions for deploying OpenShift Data Foundation using Amazon Web Services
for cloud storage

Red Hat OpenShift Data Foundation 4.9 Deploying OpenShift Data Foundation using Amazon Web Services

Instructions for deploying OpenShift Data Foundation using Amazon Web Services for cloud storage

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 using Red Hat OpenShift Container Platform on Amazon Web Services.

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. ENABLING FILE SYSTEM ACCESS FOR CONTAINERS ON RED HAT ENTERPRISE LINUX BASED NODES | 6 |
| 1.2. ENABLING KEY VALUE BACKEND PATH AND POLICY IN VAULT | 7 |
| CHAPTER 2. DEPLOY OPENSIFT DATA FOUNDATION USING DYNAMIC STORAGE DEVICES | 9 |
| 2.1. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR | 9 |
| 2.2. CREATING AN OPENSIFT DATA FOUNDATION CLUSTER | 11 |
| 2.3. VERIFYING OPENSIFT DATA FOUNDATION DEPLOYMENT | 13 |
| 2.3.1. Verifying the state of the pods | 13 |
| 2.3.2. Verifying the OpenShift Data Foundation cluster is healthy | 15 |
| 2.3.3. Verifying the Multicloud Object Gateway is healthy | 15 |
| 2.3.4. Verifying that the OpenShift Data Foundation specific storage classes exist | 15 |
| CHAPTER 3. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY | 17 |
| 3.1. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR | 17 |
| 3.2. CREATING STANDALONE MULTICLOUD OBJECT GATEWAY | 18 |
| CHAPTER 4. UNINSTALLING OPENSIFT DATA FOUNDATION | 21 |
| 4.1. UNINSTALLING OPENSIFT DATA FOUNDATION IN INTERNAL MODE | 21 |

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) AWS clusters in connected or disconnected environments along with out-of-the-box support for proxy environments.



NOTE

Only internal OpenShift Data Foundation clusters are supported on AWS. See [Planning your deployment](#) and [Preparing to deploy OpenShift Data Foundation](#) for more information about deployment requirements.

To deploy OpenShift Data Foundation, start with the requirements in [Preparing to deploy OpenShift Data Foundation](#) chapter and then follow the deployment process for your environment based on your requirement:


- [Deploy using dynamic storage devices](#)
- [Deploy standalone Multicloud Object Gateway component](#)

CHAPTER 1. PREPARING TO DEPLOY OPENSIFT DATA FOUNDATION

Deploying OpenShift Data Foundation on OpenShift Container Platform using dynamic storage devices provides you with the option to create internal cluster resources.

Before you begin the deployment of Red Hat OpenShift Data Foundation, follow these steps:

1. For Red Hat Enterprise Linux based hosts for worker nodes, [enable file system access for containers on Red Hat Enterprise Linux based nodes](#).



NOTE


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

2. 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.
3. 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. See [Resource requirements](#) section in the *Planning guide*.
4. 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](#).

1.1. 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
```

1.2. 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 DYNAMIC STORAGE DEVICES

You can deploy OpenShift Data Foundation on OpenShift Container Platform using dynamic storage devices provided by Amazon Web Services (AWS) EBS (type, **gp2**) that provides you with the option to create internal cluster resources. This results in the internal provisioning of the base services, which helps to make additional storage classes available to applications.

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](#).



NOTE

Only internal OpenShift Data Foundation clusters are supported on AWS. See [Planning your deployment](#) for more information about deployment requirements.

Also, ensure that you have addressed the requirements in [Preparing to deploy OpenShift Data Foundation](#) chapter before proceeding with the below steps for deploying using dynamic storage devices:


1. [Install the Red Hat OpenShift Data Foundation Operator](#).
2. [Create the OpenShift Data Foundation Cluster](#).

2.1. 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](#).

2.2. CREATING AN OPENSIFT DATA FOUNDATION CLUSTER

Create an OpenShift Data Foundation cluster after you install the OpenShift Data Foundation operator.

Prerequisites

- The OpenShift Data Foundation operator must be installed from the Operator Hub. For more information, see [Installing OpenShift Data Foundation Operator](#).

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, select the following:
 - a. Select the **Use an existing StorageClass** option.
 - b. Select the **Storage Class**.
By default, it is set to **gp2**.
 - c. Expand **Advanced** and select **Full Deployment** for the **Deployment type** option.
 - d. Click **Next**.
4. In the **Capacity and nodes** page, provide the necessary information:
 - a. Select a value for **Requested Capacity** from the dropdown list. It is set to **2 TiB** by default.



NOTE

Once you select the initial storage capacity, cluster expansion is performed only using the selected usable capacity (three times of raw storage).

- b. In the **Select Nodes** section, select at least three available nodes.
For cloud platforms with multiple availability zones, ensure that the Nodes are spread across different Locations/availability zones.

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. Click **Next**.
5. Optional: In the **Security and network** page, configure the following based on your requirements:
 - a. To enable encryption, select **Enable data encryption for block and file storage**
 - b. Choose either one or both the encryption levels:
 - **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 the **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 the **CA Certificate**, **Client Certificate** and **Client Private Key**.
 - D. Click **Save**.
 - d. Click **Next**.
6. In the **Review and create** page, review the configuration details.
To modify any configuration settings, click **Back**.
7. 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-storage-system** → **Resources**.
 - b. Verify that **Status** of **StorageCluster** is **Ready** and has a green tick mark next to it.

1. To verify that all the components for OpenShift Data Foundation are successfully installed, see [Verifying OpenShift Data Foundation deployment](#).

Additional resources

To enable Overprovision Control alerts, refer to [Alerts](#) in Monitoring guide.

2.3. 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](#).

2.3.1. 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) |
| 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) |

2.3.2. 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](#).

2.3.3. 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](#).

2.3.4. 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

CHAPTER 3. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY

Deploying only the Multicloud Object Gateway component with 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 Red Hat OpenShift Data Foundation Operator
- Creating standalone Multicloud Object Gateway

3.1. 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](#).

3.2. 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) |
| Rook-ceph Operator | <p>rook-ceph-operator-*</p> <p>(1 pod on any worker node)</p> |
| 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](#).