



# Red Hat OpenShift Data Foundation 4.16

## Deploying OpenShift Data Foundation on VMware vSphere

Instructions on deploying OpenShift Data Foundation using VMware vSphere  
infrastructure



# Red Hat OpenShift Data Foundation 4.16 Deploying OpenShift Data Foundation on VMware vSphere

---

Instructions on deploying OpenShift Data Foundation using VMware vSphere infrastructure

## 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> 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<sup>®</sup> 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 VMware vSphere clusters.

## Table of Contents

<b>MAKING OPEN SOURCE MORE INCLUSIVE</b> .....	<b>3</b>
<b>PROVIDING FEEDBACK ON RED HAT DOCUMENTATION</b> .....	<b>4</b>
<b>PREFACE</b> .....	<b>5</b>
<b>CHAPTER 1. PREPARING TO DEPLOY OPENSIFT DATA FOUNDATION</b> .....	<b>6</b>
1.1. REQUIREMENTS FOR INSTALLING OPENSIFT DATA FOUNDATION USING LOCAL STORAGE DEVICES	7
<b>CHAPTER 2. DEPLOY USING DYNAMIC STORAGE DEVICES</b> .....	<b>9</b>
2.1. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR	9
2.2. ENABLING CLUSTER-WIDE ENCRYPTION WITH KMS USING THE TOKEN AUTHENTICATION METHOD	10
2.3. ENABLING CLUSTER-WIDE ENCRYPTION WITH KMS USING THE KUBERNETES AUTHENTICATION METHOD	11
2.4. CREATING AN OPENSIFT DATA FOUNDATION CLUSTER	14
<b>CHAPTER 3. DEPLOY USING LOCAL STORAGE DEVICES</b> .....	<b>19</b>
3.1. INSTALLING LOCAL STORAGE OPERATOR	19
3.2. INSTALLING RED HAT OPENSIFT DATA FOUNDATION OPERATOR	19
3.3. CREATING OPENSIFT DATA FOUNDATION CLUSTER ON VMWARE VSPHERE	21
<b>CHAPTER 4. VERIFYING OPENSIFT DATA FOUNDATION DEPLOYMENT</b> .....	<b>27</b>
4.1. VERIFYING THE STATE OF THE PODS	27
4.2. VERIFYING THE OPENSIFT DATA FOUNDATION CLUSTER IS HEALTHY	29
4.3. VERIFYING THE MULTICLOUD OBJECT GATEWAY IS HEALTHY	29
4.4. VERIFYING THAT THE SPECIFIC STORAGE CLASSES EXIST	29
<b>CHAPTER 5. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY</b> .....	<b>31</b>
5.1. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY USING DYNAMIC STORAGE DEVICES	31
5.1.1. Installing Red Hat OpenShift Data Foundation Operator	31
5.1.2. Creating a standalone Multicloud Object Gateway	32
5.2. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY USING LOCAL STORAGE DEVICES	35
5.2.1. Installing Local Storage Operator	35
5.2.2. Installing Red Hat OpenShift Data Foundation Operator	36
5.2.3. Creating a standalone Multicloud Object Gateway	37
<b>CHAPTER 6. VIEW OPENSIFT DATA FOUNDATION TOPOLOGY</b> .....	<b>42</b>
<b>CHAPTER 7. UNINSTALLING OPENSIFT DATA FOUNDATION</b> .....	<b>43</b>
7.1. UNINSTALLING OPENSIFT DATA FOUNDATION IN INTERNAL MODE	43



## 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 existing Red Hat OpenShift Container Platform (RHOCP) vSphere clusters in connected or disconnected environments along with out-of-the-box support for proxy environments.

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


- Internal mode
  - [Deploy using dynamic storage devices](#)
  - [Deploy using local storage devices](#)
  - [Deploy standalone Multicloud Object Gateway](#)
- External mode
  - [Deploying OpenShift Data Foundation in external mode](#)

# CHAPTER 1. PREPARING TO DEPLOY OPENSIFT DATA FOUNDATION

Deploying OpenShift Data Foundation on OpenShift Container Platform using dynamic or local storage devices provides you with the option to create internal cluster resources. This will result in the internal provisioning of the base services, which helps to make additional storage classes available to applications.

Before you begin the deployment of Red Hat OpenShift Data Foundation using dynamic or local storage, ensure that your resource requirements are met. See the [Resource requirements](#) section in the Planning guide.

1. Verify the rotational flag on your VMDKs before deploying object storage devices (OSDs) on them. For more information, see the knowledgebase article [Override device rotational flag in ODF environment](#).
2. Optional: If you want to enable cluster-wide encryption using the external Key Management System (KMS) HashiCorp Vault, 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 then refer to [Enabling cluster-wide encryption with KMS using the Token authentication method](#).
  - 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.
3. Optional: If you want to enable cluster-wide encryption using the external Key Management System (KMS) Thales CipherTrust Manager, you must first enable the Key Management Interoperability Protocol (KMIP) and use signed certificates on your server. Follow these steps:
  - a. Create a KMIP client if one does not exist. From the user interface, select **KMIP → Client Profile → Add Profile**.
    - i. Add the **CipherTrust** username to the **Common Name** field during profile creation.
  - b. Create a token by navigating to **KMIP → Registration Token → New Registration Token**. Copy the token for the next step.
  - c. To register the client, navigate to **KMIP → Registered Clients → Add Client**. Specify the **Name**. Paste the **Registration Token** from the previous step, then click **Save**.
  - d. Download the Private Key and Client Certificate by clicking **Save Private Key** and **Save Certificate** respectively.
  - e. To create a new KMIP interface, navigate to **Admin Settings → Interfaces → Add Interface**.
    - i. Select **KMIP Key Management Interoperability Protocol** and click **Next**.
    - ii. Select a free **Port**.
    - iii. Select **Network Interface** as **all**.

- iv. Select **Interface Mode** as **TLS**, **verify client cert**, **user name taken from client cert**, **auth request is optional**.
  - v. (Optional) You can enable hard delete to delete both metadata and material when the key is deleted. It is disabled by default.
  - vi. Select the CA to be used, and click **Save**.
- f. To get the server CA certificate, click on the Action menu (  ) on the right of the newly created interface, and click **Download Certificate**.
- g. Optional: If StorageClass encryption is to be enabled during deployment, create a key to act as the Key Encryption Key (KEK):
    - i. Navigate to **Keys** → **Add Key**.
    - ii. Enter **Key Name**.
    - iii. Set the **Algorithm** and **Size** to **AES** and **256** respectively.
    - iv. Enable **Create a key in Pre-Active state** and set the date and time for activation.
    - v. Ensure that **Encrypt** and **Decrypt** are enabled under **Key Usage**.
    - vi. Copy the ID of the newly created Key to be used as the Unique Identifier during deployment.
4. Minimum starting node requirements
 

An OpenShift Data Foundation cluster will be deployed with minimum configuration when the standard deployment resource requirement is not met. See [Resource requirements](#) section in Planning guide.
  5. Disaster recovery requirements [Technology Preview]
 

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 for Kubernetes subscription
 

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

For detailed 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.
  6. For deploying using local storage devices, see [requirements for installing OpenShift Data Foundation using local storage devices](#). These are not applicable for deployment using dynamic storage devices.

## 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.

### Arbiter stretch cluster requirements

In this case, a single cluster is stretched across two zones with a third zone as the location for the arbiter. This solution is currently intended for deployment in the OpenShift Container Platform on-premises and in the same data center. This solution is not recommended for deployments stretching over multiple data centers. Instead, consider Metro-DR as a first option for no data loss DR solution deployed over multiple data centers with low latency networks.

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



#### NOTE

You cannot enable Flexible scaling and Arbiter both 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.

### 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 USING DYNAMIC STORAGE DEVICES

Deploying OpenShift Data Foundation on OpenShift Container Platform using dynamic storage devices provided by VMware vSphere (disk format: thin) provides you with the option to create internal cluster resources. This will result in the internal provisioning of the base services, which helps to make additional storage classes available to applications.



### NOTE

Both internal and external OpenShift Data Foundation clusters are supported on VMware vSphere. 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 an 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 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-4.16**.
  - 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.

## 2.2. ENABLING CLUSTER-WIDE ENCRYPTION WITH KMS USING THE TOKEN AUTHENTICATION METHOD

You can enable the key value backend path and policy in the vault for token authentication.

### Prerequisites

- Administrator access to the vault.
- A valid Red Hat OpenShift Data Foundation Advanced subscription. For more information, see the [knowledgebase article on OpenShift Data Foundation subscriptions](#).

- Carefully, select a unique path name as the backend **path** that follows the naming convention since you cannot change it later.

### Procedure

1. Enable the Key/Value (KV) backend path in the 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 the users to perform a write or delete operation on the secret:

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

3. Create a token that matches the above policy:

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

## 2.3. ENABLING CLUSTER-WIDE ENCRYPTION WITH KMS USING THE KUBERNETES AUTHENTICATION METHOD

You can enable the Kubernetes authentication method for cluster-wide encryption using the Key Management System (KMS).

### Prerequisites

- Administrator access to Vault.
- A valid Red Hat OpenShift Data Foundation Advanced subscription. For more information, see the [knowledgebase article on OpenShift Data Foundation subscriptions](#).
- The OpenShift Data Foundation operator must be installed from the Operator Hub.
- Select a unique path name as the backend **path** that follows the naming convention carefully. You cannot change this path name later.

### Procedure

1. Create a service account:

```
$ oc -n openshift-storage create serviceaccount <serviceaccount_name>
```

where, **<serviceaccount\_name>** specifies the name of the service account.

For example:

```
$ oc -n openshift-storage create serviceaccount odf-vault-auth
```

2. Create **clusterrolebindings** and **clusterroles**:

```
$ oc -n openshift-storage create clusterrolebinding vault-tokenreview-binding --
clusterrole=system:auth-delegator --serviceaccount=openshift-
storage:<serviceaccount_name>
```

For example:

```
$ oc -n openshift-storage create clusterrolebinding vault-tokenreview-binding --
clusterrole=system:auth-delegator --serviceaccount=openshift-storage:odf-vault-auth
```

3. Create a secret for the **serviceaccount** token and CA certificate.

```
$ cat <<EOF | oc create -f -
apiVersion: v1
kind: Secret
metadata:
  name: odf-vault-auth-token
  namespace: openshift-storage
  annotations:
    kubernetes.io/service-account.name: <serviceaccount_name>
type: kubernetes.io/service-account-token
data: {}
EOF
```

where, **<serviceaccount\_name>** is the service account created in the earlier step.

4. Get the token and the CA certificate from the secret.

```
$ SA_JWT_TOKEN=$(oc -n openshift-storage get secret odf-vault-auth-token -o jsonpath="
{.data['token']}" | base64 --decode; echo)
$ SA_CA_CERT=$(oc -n openshift-storage get secret odf-vault-auth-token -o jsonpath="
{.data['ca.crt']}" | base64 --decode; echo)
```

5. Retrieve the OCP cluster endpoint.

```
$ OCP_HOST=$(oc config view --minify --flatten -o jsonpath="{.clusters[0].cluster.server}")
```

6. Fetch the service account issuer:

```
$ oc proxy &
$ proxy_pid=$!
$ issuer="$( curl --silent http://127.0.0.1:8001/.well-known/openid-configuration | jq -r
.issuer)"
$ kill $proxy_pid
```

7. Use the information collected in the previous step to setup the Kubernetes authentication method in Vault:



```
$ vault auth enable kubernetes
```

```
$ vault write auth/kubernetes/config \
  token_reviewer_jwt="$SA_JWT_TOKEN" \
  kubernetes_host="$OCP_HOST" \
  kubernetes_ca_cert="$SA_CA_CERT" \
  issuer="$issuer"
```



### IMPORTANT

To configure the Kubernetes authentication method in Vault when the issuer is empty:

```
$ vault write auth/kubernetes/config \
  token_reviewer_jwt="$SA_JWT_TOKEN" \
  kubernetes_host="$OCP_HOST" \
  kubernetes_ca_cert="$SA_CA_CERT"
```

8. 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
```

9. Create a policy to restrict the users to perform a **write** or **delete** operation on the secret:

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

10. Generate the roles:

```
$ vault write auth/kubernetes/role/odf-rook-ceph-op \
  bound_service_account_names=rook-ceph-system,rook-ceph-osd,noobaa \
  bound_service_account_namespaces=openshift-storage \
  policies=odf \
  ttl=1440h
```

The role **odf-rook-ceph-op** is later used while you configure the KMS connection details during the creation of the storage system.

```
$ vault write auth/kubernetes/role/odf-rook-ceph-osd \
  bound_service_account_names=rook-ceph-osd \
  bound_service_account_namespaces=openshift-storage \
```

```
policies=odf \
ttl=1440h
```

## 2.4. 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](#).
- For VMs on VMware, ensure the **disk.EnableUUID** option is set to **TRUE**. You need to have vCenter account privileges to configure the VMs. For more information, see [Required vCenter account privileges](#). To set the **disk.EnableUUID** option, use the **Advanced** option of the **VM Options** in the **Customize hardware** tab. For more information, see [Installing on vSphere](#).
- Optional: If you want to use thick-provisioned storage for flexibility, you must create a storage class with **zeroedthick** or **eagerzeroedthick** disk format. For information, see [VMware vSphere object definition](#).

### 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 **Full Deployment** for the **Deployment type** option.
  - b. Select the **Use an existing StorageClass** option.
  - c. Select the **Storage Class**.  
By default, it is set to **thin**. If you have created a storage class with **zeroedthick** or **eagerzeroedthick** disk format for thick-provisioned storage, then that storage class is listed in addition to the default, **thin** storage class.
  - d. 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.

- e. Click **Next**.
4. In the **Capacity and nodes** page, provide the necessary information:
- 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).

- In the **Select Nodes** section, select at least three available nodes.
- 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](#).

- Optional: Select the **Taint nodes** checkbox to dedicate the selected nodes for OpenShift Data Foundation.  
Spread the worker nodes across three different physical nodes, racks, or failure domains for high availability.  
  
Use vCenter anti-affinity to align OpenShift Data Foundation rack labels with physical nodes and racks in the data center to avoid scheduling two worker nodes on the same physical chassis.

If the nodes selected do not match the OpenShift Data Foundation cluster requirement of the 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.

Select the **Taint nodes** checkbox to make selected nodes dedicated for OpenShift Data Foundation.

- e. 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**
      - i. Select 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.
      - ii. 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, select one of the following providers and provide the necessary details:
          - **Vault**
            - I. 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**.
              - **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**.
- **Thales CipherTrust Manager (using KMIP)**
    - 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**.
- b. To enable in-transit encryption, select **In-transit encryption**.
    - i. Select a **Network**.
    - ii. Click **Next**.
6. In the **Data Protection** page, if you are configuring Regional-DR solution for Openshift Data Foundation then select the **Prepare cluster for disaster recovery (Regional-DR only)** checkbox, else click **Next**.
  7. In the **Review and create** page, review the configuration details. To modify any configuration settings, click **Back**.
  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** → **ocs-storagecluster-storagesystem** → **Resources**.
  - b. Verify that **Status** of **StorageCluster** is **Ready** and has a green tick mark next to it.
    1. To verify that all components for OpenShift Data Foundation are successfully installed, see [Verifying your OpenShift Data Foundation deployment](#).

## Additional resources

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

## CHAPTER 3. DEPLOY USING LOCAL STORAGE DEVICES

Deploying OpenShift Data Foundation on OpenShift Container Platform using local storage devices provides you with the option to create internal cluster resources. This will result in the internal provisioning of the base services, which helps to make additional storage classes available to applications.

Use this section to deploy OpenShift Data Foundation on VMware where OpenShift Container Platform is already installed.

Also, ensure that you have addressed the requirements in [Preparing to deploy OpenShift Data Foundation](#) chapter before proceeding with the next steps.

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

### 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-4.16**.
  - 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 OPENSIFT DATA FOUNDATION CLUSTER ON VMWARE VSPHERE

VMware vSphere supports the following three types of local storage:

- Virtual machine disk (VMDK)
- Raw device mapping (RDM)
- VMDirectPath I/O

### Prerequisites

- Ensure that all the requirements in the [Requirements for installing OpenShift Data Foundation using local storage devices](#) section are met.
- You must have a minimum of three worker nodes with the same storage type and size attached to each node to use local storage devices on VMware.
- For VMs on VMware vSphere, ensure the **disk.EnableUUID** option is set to **TRUE**. You need to have vCenter account privileges to configure the VMs. For more information, see [Required vCenter account privileges](#). To set the **disk.EnableUUID** option, use the **Advanced** option of the **VM Options** in the **Customize hardware** tab. For more information, see [Installing on vSphere](#).

### 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**.



#### NOTE

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. Select one of the following:
    - **Disks on all nodes** to use the available disks that match the selected filters on all nodes.
    - **Disks on selected nodes** to use the available disks that match the selected filters only on selected nodes.



#### IMPORTANT

- The flexible scaling feature is enabled only when the storage cluster that you created with 3 or more nodes is spread across fewer than the minimum requirement of 3 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.  
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 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 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. 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 of the following **Encryption level**:
    - **Cluster-wide encryption** to encrypt the entire cluster (block and file).
    - **StorageClass encryption** to create 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 the **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 the **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 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. Click **Next**.
- 7. In the **Data Protection** page, if you are configuring Regional-DR solution for Openshift Data Foundation then select the **Prepare cluster for disaster recovery (Regional-DR only)** checkbox, else click **Next**.
- 8. 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.
- 9. 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** → **ocs-storagecluster-storage-system** → **Resources**.

- b. Verify that the **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**.
    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, the flexible scaling feature is enabled.

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

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

#### Additional resources

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

## CHAPTER 4. VERIFYING OPENSIFT DATA FOUNDATION DEPLOYMENT

Use this section to verify that OpenShift Data Foundation is deployed correctly.

### 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 the following table:

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

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

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



## 4.2. VERIFYING THE OPENSIFT 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](#).

## 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](#).



### IMPORTANT

The Multicloud Object Gateway only has a single copy of the database (NooBaa DB). This means if NooBaa DB PVC gets corrupted and we are unable to recover it, can result in total data loss of applicative data residing on the Multicloud Object Gateway. Because of this, Red Hat recommends taking a backup of NooBaa DB PVC regularly. If NooBaa DB fails and cannot be recovered, then you can revert to the latest backed-up version. For instructions on backing up your NooBaa DB, follow the steps in [this knowledgebase article](#).

## 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:
  - **ocs-storagecluster-ceph-rbd**

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

## CHAPTER 5. 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. You can deploy the Multicloud Object Gateway component either using dynamic storage devices or using the local storage devices.

### 5.1. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY USING DYNAMIC STORAGE DEVICES

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

#### 5.1.1. Installing Red Hat OpenShift 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-4.16**.
  - 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.

### 5.1.2. 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**.
  - b. Select the **Use an existing StorageClass** option.
  - c. Click **Next**.
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, **knip\_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
-----------	--------------------

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

## 5.2. DEPLOY STANDALONE MULTICLOUD OBJECT GATEWAY USING LOCAL STORAGE DEVICES

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

### 5.2.1. 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.

## 5.2.2. Installing Red Hat OpenShift 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-4.16**.
  - 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.

### 5.2.3. 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**.

- b. Select the **Create a new StorageClass using the local storage devices** option.
- c. Click **Next**.

**NOTE**

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

Volume Mode	Filesystem is selected by default. Always ensure that the Filesystem is selected for <b>Volume Mode</b> .
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 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: 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**.
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 → 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"> <li>● <b>ocs-operator-*</b> (1 pod on any storage node)</li> <li>● <b>ocs-metrics-exporter-*</b> (1 pod on any storage node)</li> <li>● <b>odf-operator-controller-manager-*</b> (1 pod on any storage node)</li> <li>● <b>odf-console-*</b> (1 pod on any storage node)</li> <li>● <b>csi-addons-controller-manager-*</b> (1 pod on any storage node)</li> </ul>

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

## CHAPTER 6. 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 7. UNINSTALLING OPENSIFT DATA FOUNDATION

### 7.1. UNINSTALLING OPENSIFT DATA FOUNDATION IN INTERNAL MODE

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