Red Hat OpenShift Data Foundation 4.12

Deploying OpenShift Data Foundation using Microsoft Azure

Instructions on deploying OpenShift Data Foundation using Microsoft Azure

Last Updated: 2024-03-06

Red Hat OpenShift Data Foundation 4.12 Deploying OpenShift Data Foundation using Microsoft Azure

Instructions on deploying OpenShift Data Foundation using Microsoft Azure

Legal Notice

Abstract

Read this document for instructions about how to install and manage Red Hat OpenShift Data Foundation using Red Hat OpenShift Container Platform on Microsoft Azure.

Table of Contents

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) Azure clusters.

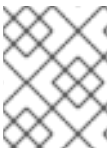

NOTE

Only internal OpenShift Data Foundation clusters are supported on Microsoft Azure. See Planning your deployment 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 appropriate deployment process based on your requirement:

- Deploy OpenShift Data Foundation on Microsoft Azure
- Deploy standalone Multicloud Object Gateway component

CHAPTER 1. PREPARING TO DEPLOY OPENSHIFT DATA FOUNDATION

Deploying OpenShift Data Foundation on OpenShift Container Platform using dynamic 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 OpenShift Data Foundation, follow these steps:

1. Setup a chrony server. See Configuring chrony time service and use knowledgebase solution to create rules allowing all traffic.

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 the Token authentication using KMS.

   - When the Kubernetes authentication method is selected for encryption then refer to Enabling cluster-wide encryption with the Kubernetes authentication using KMS .

   - Ensure that you are using signed certificates on your Vault servers.

   

   NOTE

   If you are using Thales CipherTrust Manager as your KMS, you will enable it during deployment.

3. Minimum starting node requirements
   An OpenShift Data Foundation cluster is deployed with minimum configuration when the standard deployment resource requirement is not met. See Resource requirements section in Planning guide.

4. 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.

CHAPTER 2. DEPLOYING OPENSHIFT DATA FOUNDATION ON MICROSOFT AZURE AND AZURE RED HAT OPENSHIFT

You can deploy OpenShift Data Foundation on OpenShift Container Platform using dynamic storage devices provided by Microsoft Azure installer-provisioned infrastructure (IPI) (type: **managed-csi**) that enables you to create internal cluster resources. This results in 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 Microsoft Azure. See Planning your deployment for more information about deployment requirements.

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 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 nodes in the Red Hat OpenShift Container Platform cluster. Each node should include one disk and requires 3 disks (PVs). However, one PV remains eventually unused by default. This is an expected behavior.
- For additional resource requirements, see the Planning your deployment guide.

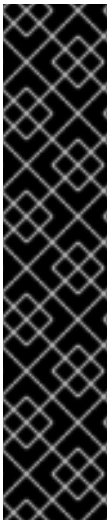

IMPORTANT

- 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-4.12**.
   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 **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_CRT=$(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="$SA_CA_CRT" \
        issuer="$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_CRT"
```

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 OPENSHIFT 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 using the Operator Hub .

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 **managed-csi**.
   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.

c. Optional: Select the **Taint nodes** checkbox to dedicate the selected nodes for OpenShift Data Foundation.
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 requirements 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.

d. 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. 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.

c. Optional: Select the **Connect to an external key management service** checkbox. This is optional for cluster-wide encryption.

i. From the **Key Management Service Provider** drop-down list, either select **Vault** or **Thales CipherTrust Manager (using KMIP)**. If you selected **Vault**, go to the next step. If you selected **Thales CipherTrust Manager (using KMIP)**, go to step iii.

ii. Select an **Authentication Method**.

Using Token authentication method

- Enter a unique **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Token**.

- Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:

  - Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.

  - Optional: Enter **TLS Server Name** and **Vault Enterprise Namespace**.

  - Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key**.

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

Using Kubernetes authentication method

- Enter a unique Vault **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Role** name.

- Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:
  - Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
  - Optional: Enter **TLS Server Name** and **Authentication Path** if applicable.
  - Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key** .
  - Click **Save** and skip to step iv.

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

v. 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-storagesystem → Resources**.
  b. Verify that **Status** of **StorageCluster** is **Ready** and has a green tick mark next to it.
- 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. DEPLOYING OPENSHIFT DATA FOUNDATION ON AZURE RED HAT OPENSHIFT

The Azure Red Hat OpenShift service enables you to deploy fully managed OpenShift clusters. Red Hat OpenShift Data Foundation can be deployed on Azure Red Hat OpenShift service.

IMPORTANT

OpenShift Data Foundation on Azure Red Hat OpenShift is not a managed service offering. Red Hat OpenShift Data Foundation subscriptions are required to have the installation supported by the Red Hat support team. Open support cases by choosing the product as **Red Hat OpenShift Data Foundation** with the Red Hat support team (and not Microsoft) if you need any assistance for Red Hat OpenShift Data Foundation on Azure Red Hat OpenShift.

To install OpenShift Data Foundation on Azure Red Hat OpenShift, follow sections:

1. Getting a Red Hat pull secret for new deployment of Azure Red Hat OpenShift .
2. Preparing a Red Hat pull secret for existing Azure Red Hat OpenShift clusters .
3. Adding the pull secret to the cluster .
4. Validating your Red Hat pull secret is working .
5. Install the Red Hat OpenShift Data Foundation Operator .
6. Create the OpenShift Data Foundation Cluster Service .

3.1. GETTING A RED HAT PULL SECRET FOR NEW DEPLOYMENT OF AZURE RED HAT OPENSHIFT

A Red Hat pull secret enables the cluster to access Red Hat container registries along with additional content.

Prerequisites

- A Red Hat portal account.
- OpenShift Data Foundation subscription.

Procedure

To get a Red Hat pull secret for a new deployment of Azure Red Hat OpenShift, follow the steps in the section Get a Red Hat pull secret in the official Microsoft Azure documentation.

Note that while creating the Azure Red Hat OpenShift cluster , you may need larger worker nodes, controlled by *--worker-vm-size* or more worker nodes, controlled by *--worker-count*. The recommended *worker-vm-size* is **Standard_D16s_v3**. You can also use dedicated worker nodes, for more information, see How to use dedicated worker nodes for Red Hat OpenShift Data Foundation in the *Managing and allocating storage resources* guide.

3.2. PREPARING A RED HAT PULL SECRET FOR EXISTING AZURE RED HAT OPENSHIFT CLUSTERS

When you create an Azure Red Hat OpenShift cluster without adding a Red Hat pull secret, a pull secret is still created on the cluster automatically. However, this pull secret is not fully populated.

Use this section to update the automatically created pull secret with the additional values from the Red Hat pull secret.

Prerequisites

- Existing Azure Red Hat OpenShift cluster without a Red Hat pull secret.

Procedure

To prepare a Red Hat pull secret for existing an existing Azure Red Hat OpenShift clusters, follow the steps in the section Prepare your pull secret in the official Mircosoft Azure documentation.

3.3. ADDING THE PULL SECRET TO THE CLUSTER

Prerequisites

- A Red Hat pull secret.

Procedure

- Run the following command to update your pull secret.

NOTE

Running this command causes the cluster nodes to restart one by one as they are updated.

```
oc set data secret/pull-secret -n openshift-config --from-file=.dockerconfigjson=./pull-secret.json
```

After the secret is set, you can enable the Red Hat Certified Operators.

3.3.1. Modifying the configuration files to enable Red Hat operators

To modify the configuration files to enable Red Hat operators, follow the steps in the section Modify the configuration files in the official Microsoft Azure documentation.

3.4. VALIDATING YOUR RED HAT PULL SECRET IS WORKING

After you add the pull secret and modify the configuration files, the cluster can take several minutes to get updated.

To check if the cluster has been updated, run the following command to show the **Certified Operators** and **Red Hat Operators** sources available:

```
$ oc get catalogsource -A
```

```
NAMESPACE               NAME                DISPLAY
openshift-marketplace   redhat-operators    Red Hat Operators

 TYPE   PUBLISHER   AGE
 grpc   Red Hat     11s
```

If you do not see the Red Hat Operators, wait a few minutes and try again.

To ensure that your pull secret has been updated and is working correctly, open **Operator Hub** and check for any Red Hat verified Operator. For example, check if the OpenShift Data Foundation Operator is available, and see if you have permissions to install it.

3.5. 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 nodes in the Red Hat OpenShift Container Platform cluster. Each node should include one disk and requires 3 disks (PVs). However, one PV remains eventually unused by default. This is an expected behavior.
- For additional resource requirements, see the Planning your deployment guide.

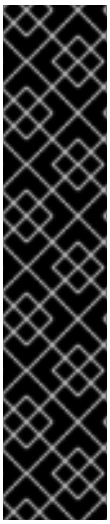

IMPORTANT

- 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-4.12**.

    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 **Data Foundation** dashboard is available.

3.6. CREATING AN OPENSHIFT 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 using the Operator Hub .

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 **managed-csi**.

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.

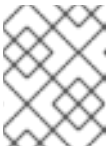

NOTE

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.

c. Optional: Select the **Taint nodes** checkbox to dedicate the selected nodes for OpenShift Data Foundation.
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 requirements 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.

d. 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. 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.

c. Optional: Select the **Connect to an external key management service** checkbox. This is optional for cluster-wide encryption.

i. From the **Key Management Service Provider** drop-down list, either select **Vault** or **Thales CipherTrust Manager (using KMIP)**. If you selected **Vault**, go to the next step. If you selected **Thales CipherTrust Manager (using KMIP)**, go to step iii.

ii. Select an **Authentication Method**.

Using Token authentication method

- Enter a unique **Connection Name**, host **Address** of the Vault server ('https://<hostname or ip>'), **Port** number and **Token**.

- Expand **Advanced Settings** to enter additional settings and certificate details based on your **Vault** configuration:

- Enter the Key Value secret path in **Backend Path** that is dedicated and unique to OpenShift Data Foundation.
- Optional: Enter **TLS Server Name** and **Vault Enterprise Namespace**.
- Upload the respective PEM encoded certificate file to provide the **CA Certificate**, **Client Certificate** and **Client Private Key**.
- Click **Save** and skip to step iv.

Using Kubernetes authentication method

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

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

v. 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-storagesystem → Resources**.
    b. Verify that **Status** of **StorageCluster** is **Ready** and has a green tick mark next to it.
- 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 4. VERIFYING OPENSHIFT 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.

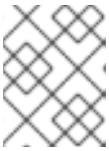

NOTE

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 Table 4.1, "Pods corresponding to OpenShift Data Foundation cluster".

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

Table 4.1. Pods corresponding to OpenShift Data Foundation cluster

Component	Corresponding pods
OpenShift Data Foundation Operator	• **ocs-operator-*** (1 pod on any storage node) • **ocs-metrics-exporter-*** (1 pod on any storage node) • **odf-operator-controller-manager-*** (1 pod on any storage node) • **odf-console-*** (1 pod on any storage node) • **csi-addons-controller-manager-*** (1 pod on any storage node)
Rook-ceph Operator	**rook-ceph-operator-*** (1 pod on any storage node)

Component	Corresponding pods
Multicloud Object Gateway	• **noobaa-operator-*** (1 pod on any storage node) • **noobaa-core-*** (1 pod on any storage node) • **noobaa-db-pg-*** (1 pod on any storage node) • **noobaa-endpoint-*** (1 pod on any storage node)
MON	**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	• **cephfs**
	  ◦ **csi-cephfsplugin-*** (1 pod on each storage node)
	  ◦ **csi-cephfsplugin-provisioner-*** (2 pods distributed across storage nodes)
	• **rbd**
	  ◦ **csi-rbdplugin-*** (1 pod on each storage node)
	  ◦ **csi-rbdplugin-provisioner-*** (2 pods distributed across storage nodes)

rook-ceph-crashcollector	**rook-ceph-crashcollector-*** (1 pod on each storage node)
OSD	• **rook-ceph-osd-*** (1 pod for each device) • **rook-ceph-osd-prepare-ocs-deviceset-*** (1 pod for each device)

4.2. VERIFYING THE OPENSHIFT DATA FOUNDATION CLUSTER IS HEALTHY

Procedure

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

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**

CHAPTER 5. 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

5.1. 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 nodes in the Red Hat OpenShift Container Platform cluster. Each node should include one disk and requires 3 disks (PVs). However, one PV remains eventually unused by default. This is an expected behavior.
- 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-4.12**.

    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 **Data Foundation** dashboard is available.

5.2. 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, **kmip_all_<port>.ciphertrustmanager.local**.

d. Select a **Network**.

e. Click **Next**.

5. In the **Review and create** page, review the configuration details:
To modify any configuration settings, click **Back**.

6. Click **Create StorageSystem**.

Verification steps

Verifying that the OpenShift Data Foundation cluster is healthy

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

2. In the **Status** card of the **Overview** tab, click **Storage System** and then click the storage system link from the pop up that appears.

   a. In the **Status card** of the **Object** tab, verify that both *Object Service* and *Data Resiliency* have a green tick.

   b. In the **Details** card, verify that the MCG information is displayed.

Verifying the state of the pods

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

2. Select **openshift-storage** from the **Project** drop-down list and verify that the following pods are in **Running** state.

NOTE

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

Component	Corresponding pods

Component	Corresponding pods
OpenShift Data Foundation Operator	• **ocs-operator-*** (1 pod on any storage node) • **ocs-metrics-exporter-*** (1 pod on any storage node) • **odf-operator-controller-manager-*** (1 pod on any storage node) • **odf-console-*** (1 pod on any storage node) • **csi-addons-controller-manager-*** (1 pod on any storage node)
Rook-ceph Operator	**rook-ceph-operator-*** (1 pod on any storage node)
Multicloud Object Gateway	• **noobaa-operator-*** (1 pod on any storage node) • **noobaa-core-*** (1 pod on any storage node) • **noobaa-db-pg-*** (1 pod on any storage node) • **noobaa-endpoint-*** (1 pod on any storage node)

CHAPTER 6. UNINSTALLING OPENSHIFT DATA FOUNDATION

6.1. UNINSTALLING OPENSHIFT DATA FOUNDATION IN INTERNAL MODE

To uninstall OpenShift Data Foundation in Internal mode, refer to the knowledge base article on Uninstalling OpenShift Data Foundation.