



Red Hat OpenShift Data Foundation 4.12

Updating OpenShift Data Foundation

Instructions for cluster and storage administrators regarding upgrading

Red Hat OpenShift Data Foundation 4.12 Updating OpenShift Data Foundation

Instructions for cluster and storage administrators regarding upgrading

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document explains how to update previous versions of Red Hat OpenShift Data Foundation.

Table of Contents

| | |
|---|----|
| MAKING OPEN SOURCE MORE INCLUSIVE | 3 |
| PROVIDING FEEDBACK ON RED HAT DOCUMENTATION | 4 |
| CHAPTER 1. OVERVIEW OF THE OPENSIFT DATA FOUNDATION UPDATE PROCESS | 5 |
| CHAPTER 2. OPENSIFT DATA FOUNDATION UPGRADE CHANNELS AND RELEASES | 6 |
| CHAPTER 3. UPDATING RED HAT OPENSIFT DATA FOUNDATION 4.11 TO 4.12 | 7 |
| CHAPTER 4. UPDATING RED HAT OPENSIFT DATA FOUNDATION 4.12.X TO 4.12.Y | 9 |
| CHAPTER 5. CHANGING THE UPDATE APPROVAL STRATEGY | 11 |
| CHAPTER 6. UPDATING THE OPENSIFT DATA FOUNDATION EXTERNAL SECRET | 12 |

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**.

CHAPTER 1. OVERVIEW OF THE OPENSIFT DATA FOUNDATION UPDATE PROCESS

This chapter helps you to upgrade between the minor releases and z-streams for all Red Hat OpenShift Data Foundation deployments (Internal, Internal-Attached and External). The upgrade process remains the same for all deployments.

You can upgrade OpenShift Data Foundation and its components, either between minor releases like 4.11 and 4.12, or between z-stream updates like 4.12.0 and 4.12.1 by enabling automatic updates (if not done so during operator installation) or performing manual updates. When a new z-stream release becomes available, the upgrade process triggers automatically if the update strategy was set to Automatic.

You also need to upgrade the different parts of Red Hat OpenShift Data Foundation in the following order for both internal and external mode deployments:

1. **Update OpenShift Container Platform** according to the [Updating clusters](#) documentation for OpenShift Container Platform.
2. **Update Red Hat OpenShift Data Foundation.**
 - a. **To prepare a disconnected environment for updates**, see [Operators guide to using Operator Lifecycle Manager on restricted networks](#) to be able to update OpenShift Data Foundation as well as Local Storage Operator when in use.
 - b. **For updating between minor releases**, see [Updating Red Hat OpenShift Data Foundation 4.11 to 4.12](#).
 - c. **For updating between z-stream releases**, see [Updating Red Hat OpenShift Data Foundation 4.12.x to 4.12.y](#).
 - d. **For updating external mode deployments**, you must also perform the steps from section [Updating the Red Hat OpenShift Data Foundation external secret](#).
 - e. **If you use local storage, then update the Local Storage operator** See [Checking for Local Storage Operator deployments](#) if you are unsure.

Update considerations

Review the following important considerations before you begin.

- The Red Hat OpenShift Container Platform version is the same as Red Hat OpenShift Data Foundation. See the [Interoperability Matrix](#) for more information about supported combinations of OpenShift Container Platform and Red Hat OpenShift Data Foundation.
- To know whether your cluster was deployed in internal or external mode, refer to the [knowledgebase article](#) on *How to determine if ODF cluster has storage in internal or external mode*.
- The Local Storage Operator is fully supported only when the Local Storage Operator version matches the Red Hat OpenShift Container Platform version.
- The flexible scaling feature is available only in new deployments of OpenShift Data Foundation. For more information, see [Scaling storage guide](#).

CHAPTER 2. OPENSIFT DATA FOUNDATION UPGRADE CHANNELS AND RELEASES

In OpenShift Container Platform 4.1, Red Hat introduced the concept of channels for recommending the appropriate release versions for cluster upgrades. By controlling the pace of upgrades, these upgrade channels allow you to choose an upgrade strategy. As OpenShift Data Foundation gets deployed as an operator in OpenShift Container Platform, it follows the same strategy to control the pace of upgrades by shipping the fixes in multiple channels. Upgrade channels are tied to a minor version of OpenShift Data Foundation.

For example, OpenShift Data Foundation 4.12 upgrade channels recommend upgrades within 4.12. Upgrades to future releases is not recommended. This strategy ensures that administrators can explicitly decide to upgrade to the next minor version of OpenShift Data Foundation.

Upgrade channels control only release selection and do not impact the version of the cluster that you install; the **odf-operator** decides the version of OpenShift Data Foundation to be installed. By default, it always installs the latest OpenShift Data Foundation release maintaining the compatibility with OpenShift Container Platform. So, on OpenShift Container Platform 4.12, OpenShift Data Foundation 4.12 will be the latest version which can be installed.

OpenShift Data Foundation upgrades are tied to the OpenShift Container Platform upgrade to ensure that compatibility and interoperability are maintained with the OpenShift Container Platform. For OpenShift Data Foundation 4.12, OpenShift Container Platform 4.12 and 4.13 (when generally available) are supported. OpenShift Container Platform 4.13 is supported to maintain forward compatibility of OpenShift Data Foundation with OpenShift Container Platform. Keep the OpenShift Data Foundation version the same as OpenShift Container Platform in order to get the benefit of all the features and enhancements in that release.



IMPORTANT

Due to fundamental Kubernetes design, all OpenShift Container Platform updates between minor versions must be serialized. You must update from OpenShift Container Platform 4.10 to 4.11 and then to 4.12. You cannot update from OpenShift Container Platform 4.10 to 4.12 directly. For more information, see *Preparing to perform an EUS-to-EUS update* of the [Updating clusters](#) guide in OpenShift Container Platform documentation.

OpenShift Data Foundation 4.12 offers the following upgrade channel:

- `stable-4.12`
- `eus-4.y` (only when running an even-numbered 4.y cluster release, like 4.12)

stable-4.12 channel

Once a new version is Generally Available, the stable channel corresponding to the minor version gets updated with the new image which can be used to upgrade. You can use the **stable-4.12** channel to upgrade from OpenShift Data Foundation 4.11 and upgrades within 4.12.

eus-4.y channel

In addition to the stable channel, all even-numbered minor versions of OpenShift Data Foundation offer Extended Update Support (EUS). These EUS versions extend the Full and Maintenance support phases for customers with Standard and Premium Subscriptions to 18 months. The only difference between `stable-4.y` and `eus-4.y` channels is that the EUS channels will include the release only when the next EUS release is available.

CHAPTER 3. UPDATING RED HAT OPENSIFT DATA FOUNDATION 4.11 TO 4.12

This chapter helps you to upgrade between the minor releases for all Red Hat OpenShift Data Foundation deployments (Internal, Internal-Attached and External). The upgrade process remains the same for all deployments. The Only difference is what gets upgraded and what's not.

- For Internal and Internal-attached deployments, upgrading OpenShift Data Foundation upgrades all OpenShift Data Foundation services including the backend Ceph Storage cluster.
- For External mode deployments, upgrading OpenShift Data Foundation only upgrades the OpenShift Data Foundation service while the backend Ceph storage cluster remains untouched and needs to be upgraded separately.

We recommend upgrading RHCS along with OpenShift Data Foundation in order to get new feature support, security fixes, and other bug fixes. Since we do not have a strong dependency on RHCS upgrade, you can upgrade the OpenShift Data Foundation operator first followed by RHCS upgrade or vice-versa. See [solution](#) to know more about Red Hat Ceph Storage releases.



IMPORTANT

Upgrading to 4.12 directly from any version older than 4.11 is unsupported.

Prerequisites

- Ensure that the **OpenShift Container Platform** cluster has been updated to the latest stable release of version 4.12.X, see [Updating Clusters](#).
- Ensure that the **OpenShift Data Foundation** cluster is healthy and data is resilient.
 - Navigate to **Storage → Data Foundation → Storage Systemstab** and then click on the storage system name.
 - Check for the green tick on the status card of both **Overview - Block and File** and **Object** tabs. Green tick indicates that the *storage cluster*, *object service* and *data resiliency* are all healthy.
- Ensure that all **OpenShift Data Foundation** Pods, including the operator pods, are in **Running** state in the **openshift-storage** namespace.
To view the state of the pods, on the OpenShift Web Console, click **Workloads → Pods**. 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.

- Ensure that you have sufficient time to complete the OpenShift Data Foundation update process, as the update time varies depending on the number of OSDs that run in the cluster.

Procedure

1. On the OpenShift Web Console, navigate to **Operators → Installed Operators**.
2. Select **openshift-storage** project.

3. Click the OpenShift Data Foundation operator name.
4. Click the **Subscription** tab and click the link under **Update Channel**.
5. Select the **Stable-4.12** update channel and **Save** it.
6. If the **Upgrade status** shows **requires approval**, click on **requires approval**.
 - a. On the Install Plan Details page, click **Preview Install Plan**.
 - b. Review the install plan and click **Approve**.
Wait for the **Status** to change from **Unknown** to **Created**.
7. Navigate to **Operators → Installed Operators**.
8. Select the **openshift-storage** project.
Wait for the OpenShift Data Foundation Operator **Status** to change to **Up to date**.

Verification steps

- Check the **Version** below the OpenShift Data Foundation name and check the operator status.
 - Navigate to **Operators → Installed Operators** and select the **openshift-storage** project.
 - When the upgrade completes, the version updates to a new version number for OpenShift Data Foundation and status changes to **Succeeded** with a green tick.
- Verify that the **OpenShift Data Foundation** cluster is healthy and data is resilient.
 - Navigate to **Storage → Data Foundation → Storage Systems** tab and then click on the storage system name.
 - Check for the green tick on the status card of **Overview- Block and File** and **Object** tabs.
Green tick indicates that the storage cluster, object service and data resiliency is healthy.
- If verification steps fail, [contact Red Hat Support](#).



IMPORTANT

After updating external mode deployments, you must also update the external secret. For instructions, see [Updating the OpenShift Data Foundation external secret](#).

Additional Resources

If you face any issues while updating OpenShift Data Foundation, see the *Commonly required logs for troubleshooting* section in the [Troubleshooting guide](#).

CHAPTER 4. UPDATING RED HAT OPENSIFT DATA FOUNDATION 4.12.X TO 4.12.Y

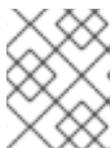
This chapter helps you to upgrade between the z-stream release for all Red Hat OpenShift Data Foundation deployments (Internal, Internal-Attached and External). The upgrade process remains the same for all deployments. The Only difference is what gets upgraded and what's not.

- For Internal and Internal-attached deployments, upgrading OpenShift Data Foundation upgrades all OpenShift Data Foundation services including the backend Ceph Storage cluster.
- For External mode deployments, upgrading OpenShift Data Foundation only upgrades the OpenShift Data Foundation service while the backend Ceph storage cluster remains untouched and needs to be upgraded separately.
Hence, we recommend upgrading RHCS along with OpenShift Data Foundation in order to get new feature support, security fixes, and other bug fixes. Since we do not have a strong dependency on RHCS upgrade, you can upgrade the OpenShift Data Foundation operator first followed by RHCS upgrade or vice-versa. See [solution](#) to know more about Red Hat Ceph Storage releases.

When a new z-stream release becomes available, the upgrade process triggers automatically if the update strategy was set to **Automatic**. If the update strategy is set to **Manual** then use the following procedure.

Prerequisites

- Ensure that the **OpenShift Container Platform** cluster has been updated to the latest stable release of version 4.12.X, see [Updating Clusters](#).
- Ensure that the **OpenShift Data Foundation** cluster is healthy and data is resilient.
 - Navigate to **Storage → Data Foundation → Storage Systemstab** and then click on the storage system name.
 - Check for the green tick on the status card of **Overview - Block and File** and **Object** tabs. Green tick indicates that the storage cluster, object service and data resiliency is healthy.
- Ensure that all **OpenShift Data Foundation** Pods, including the operator pods, are in **Running** state in the **openshift-storage** namespace.
To view the state of the pods, on the OpenShift Web Console, click **Workloads → Pods** 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.

- Ensure that you have sufficient time to complete the OpenShift Data Foundation update process, as the update time varies depending on the number of OSDs that run in the cluster.

Procedure

1. On the OpenShift Web Console, navigate to **Operators → Installed Operators**.
2. Select **openshift-storage** project.

**NOTE**

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

3. Click the **OpenShift Data Foundation** operator name.
4. Click the **Subscription** tab.
5. If the **Upgrade Status** shows **require approval**, click on **requires approval** link.
6. On the **InstallPlan Details** page, click **Preview Install Plan**.
7. Review the install plan and click **Approve**.
8. Wait for the Status to change from **Unknown** to **Created**.

Verification steps

- Check the **Version** below the OpenShift Data Foundation name and check the operator status.
 - Navigate to **Operators → Installed Operators** and select the **openshift-storage** project.
 - When the upgrade completes, the version updates to a new version number for OpenShift Data Foundation and status changes to **Succeeded** with a green tick.
- Verify that the OpenShift Data Foundation cluster is healthy and data is resilient.
 - Navigate to **Storage → Data Foundation → Storage Systemstab** and then click on the storage system name.
 - Check for the green tick on the status card of **Overview - Block and File** and **Object** tabs. Green tick indicates that the storage cluster, object service and data resiliency is healthy
- If verification steps fail, [contact Red Hat Support](#).

CHAPTER 5. CHANGING THE UPDATE APPROVAL STRATEGY

To ensure that the storage system gets updated automatically when a new update is available in the same channel, we recommend keeping the update approval strategy to **Automatic**. Changing the update approval strategy to **Manual** will need manual approval for each upgrade.

Procedure

1. Navigate to **Operators → Installed Operators**.
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.

3. Click on **OpenShift Data Foundation** operator name
4. Go to the **Subscription** tab.
5. Click on the **pencil** icon for changing the **Update approval**.
6. Select the update approval strategy and click **Save**.

Verification steps

- Verify that the Update approval shows the newly selected approval strategy below it.

CHAPTER 6. UPDATING THE OPENSIFT DATA FOUNDATION EXTERNAL SECRET

Update the OpenShift Data Foundation external secret after updating to the latest version of OpenShift Data Foundation.



NOTE

Updating the external secret is not required for batch updates. For example, when updating from OpenShift Data Foundation 4.12.x to 4.12.y.

Prerequisites

- Update the OpenShift Container Platform cluster to the latest stable release of 4.12.z, see [Updating Clusters](#).
- Ensure that the OpenShift Data Foundation cluster is healthy and the data is resilient. Navigate to **Storage** → **Data Foundation** → **Storage Systems** tab and then click on the storage system name.
 - On the **Overview - Block and File** tab, check the Status card and confirm that the *Storage cluster* has a green tick indicating it is healthy.
 - Click the **Object** tab and confirm *Object Service* and *Data resiliency* has a green tick indicating it is healthy. The RADOS Object Gateway is only listed in case RADOS Object Gateway endpoint details are included while deploying OpenShift Data Foundation in external mode.
- Red Hat Ceph Storage must have a Ceph dashboard installed and configured.

Procedure

1. Download the OpenShift Data Foundation version of the **ceph-external-cluster-details-exporter.py** python script.

```
# oc get csv $(oc get csv -n openshift-storage | grep ocs-operator | awk '{print $1}') -n
openshift-storage -o
jsonpath='{.metadata.annotations.external\.features\.ocs\.openshift\.io/export-script}' | base64
--decode > ceph-external-cluster-details-exporter.py
```

2. Update permission caps on the external Red Hat Ceph Storage cluster by running **ceph-external-cluster-details-exporter.py** on any client node in the external Red Hat Ceph Storage cluster. You may need to ask your Red Hat Ceph Storage administrator to do this.

```
# python3 ceph-external-cluster-details-exporter.py --upgrade
```

The updated permissions for the user are set as:

```
client.csi-cephfs-node
key: AQCZYz0piYgu/IRAAipji4C8+Lfymu9vOrox3zQ==
caps: [mds] allow rw
caps: [mgr] allow rw
caps: [mon] allow r, allow command 'osd blacklist'
caps: [osd] allow rw tag cephfs =
```

```

client.csi-cephfs-provisioner
key: AQCZY0piDUMSIxAARuGUyhLXFO9u4zQeRG65pQ==
caps: [mgr] allow rw
caps: [mon] allow r, allow command 'osd blocklist'
caps: [osd] allow rw tag cephfs metadata=*
client.csi-rbd-node
key: AQCZY0pi88IKHhAAvzRN4fD90nkb082ldrTaHA==
caps: [mon] profile rbd, allow command 'osd blocklist'
caps: [osd] profile rbd
client.csi-rbd-provisioner
key: AQCZY0pi6W8IIBAAgRJfrAW7kZfucNdqJqS9dQ==
caps: [mgr] allow rw
caps: [mon] profile rbd, allow command 'osd blocklist'
caps: [osd] profile rbd

```

3. Run the previously downloaded python script and save the JSON output that gets generated, from the external Red Hat Ceph Storage cluster.
 - a. Run the previously downloaded python script:

```

# python3 ceph-external-cluster-details-exporter.py --rbd-data-pool-name <rbd block pool name> --monitoring-endpoint <ceph mgr prometheus exporter endpoint> --monitoring-endpoint-port <ceph mgr prometheus exporter port> --rgw-endpoint <rgw endpoint> --run-as-user <ocs_client_name> [optional arguments]

```

--rbd-data-pool-name

Is a mandatory parameter used for providing block storage in OpenShift Data Foundation.

--rgw-endpoint

Is optional. Provide this parameter if object storage is to be provisioned through Ceph Rados Gateway for OpenShift Data Foundation. Provide the endpoint in the following format: **<ip_address>:<port>**.

--monitoring-endpoint

Is optional. It accepts comma separated list of IP addresses of active and standby mgrs reachable from the OpenShift Container Platform cluster. If not provided, the value is automatically populated.

--monitoring-endpoint-port

Is optional. It is the port associated with the ceph-mgr Prometheus exporter specified by **--monitoring-endpoint**. If not provided, the value is automatically populated.

--run-as-user

The client name used during OpenShift Data Foundation cluster deployment. Use the default client name **client.healthchecker** if a different client name was not set.



NOTE

Ensure that all the parameters, including the optional arguments, except for monitoring-endpoint and monitoring-endpoint-port, are the same as what was used during the deployment of OpenShift Data Foundation in external mode.

- b. Save the JSON output generated after running the script in the previous step.

Example output:

```
{{"name": "rook-ceph-mon-endpoints", "kind": "ConfigMap", "data": {"data": "xxx.xxx.xxx.xxx:xxxx", "maxMonId": "0", "mapping": "{}"}}, {"name": "rook-ceph-mon", "kind": "Secret", "data": {"admin-secret": "admin-secret", "fsid": "<fs-id>", "mon-secret": "mon-secret"}}, {"name": "rook-ceph-operator-creds", "kind": "Secret", "data": {"userID": "<user-id>", "userKey": "<user-key>"}}, {"name": "rook-csi-rbd-node", "kind": "Secret", "data": {"userID": "csi-rbd-node", "userKey": "<user-key>"}}, {"name": "ceph-rbd", "kind": "StorageClass", "data": {"pool": "<pool>"}}, {"name": "monitoring-endpoint", "kind": "CephCluster", "data": {"MonitoringEndpoint": "xxx.xxx.xxx.xxx", "MonitoringPort": "xxxx"}}, {"name": "rook-ceph-dashboard-link", "kind": "Secret", "data": {"userID": "ceph-dashboard-link", "userKey": "<user-key>"}}, {"name": "rook-csi-rbd-provisioner", "kind": "Secret", "data": {"userID": "csi-rbd-provisioner", "userKey": "<user-key>"}}, {"name": "rook-csi-cephfs-provisioner", "kind": "Secret", "data": {"adminID": "csi-cephfs-provisioner", "adminKey": "<admin-key>"}}, {"name": "rook-csi-cephfs-node", "kind": "Secret", "data": {"adminID": "csi-cephfs-node", "adminKey": "<admin-key>"}}, {"name": "cephfs", "kind": "StorageClass", "data": {"fsName": "cephfs", "pool": "cephfs_data"}}, {"name": "ceph-rgw", "kind": "StorageClass", "data": {"endpoint": "xxx.xxx.xxx.xxx", "poolPrefix": "default"}}, {"name": "rgw-admin-ops-user", "kind": "Secret", "data": {"accessKey": "<access-key>", "secretKey": "<secret-key>"}]}
```

4. Upload the generated JSON file.
 - a. Log in to the OpenShift Web Console.
 - b. Click **Workloads → Secrets**
 - c. Set project to **openshift-storage**.
 - d. Click **rook-ceph-external-cluster-details**.
 - e. Click **Actions (⋮) → Edit Secret**
 - f. Click **Browse** and upload the JSON file.
 - g. Click **Save**.

Verification steps

- To verify that the OpenShift Data Foundation cluster is healthy and data is resilient, navigate to **Storage → Data Foundation → Storage Systems** tab and then click on the storage system name.
 - On the **Overview → Block and File** tab, check the Details card to verify that the RHCS dashboard link is available and also check the Status card to confirm that the *Storage Cluster* has a green tick indicating it is healthy.
 - Click the **Object** tab and confirm *Object Service* and *Data resiliency* has a green tick indicating it is healthy. The RADOS Object Gateway is only listed in case RADOS Object Gateway endpoint details are included while deploying OpenShift Data Foundation in external mode.
- If verification steps fail, [contact Red Hat Support](#).

