



Red Hat Ceph Storage 5

Upgrade Guide

Upgrading a Red Hat Ceph Storage Cluster

Red Hat Ceph Storage 5 Upgrade Guide

Upgrading a Red Hat Ceph Storage Cluster

Legal Notice

Copyright © 2022 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document provides instructions on upgrading a Red Hat Ceph Storage cluster running Red Hat Enterprise Linux on AMD64 and Intel 64 architectures. 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.

Table of Contents

CHAPTER 1. UPGRADING A RED HAT CEPH STORAGE CLUSTER FROM RHCS 4 TO RHCS 5	3
1.1. PREREQUISITES	4
1.2. COMPATIBILITY CONSIDERATIONS BETWEEN RHCS AND PODMAN VERSIONS	5
1.3. PREPARING FOR AN UPGRADE	5
1.4. BACKING UP THE FILES BEFORE THE HOST OS UPGRADE	8
1.5. CONVERTING TO A CONTAINERIZED DEPLOYMENT	9
1.6. UPDATING THE HOST OPERATING SYSTEM	11
1.6.1. Manually upgrading Ceph Monitor nodes and their operating systems	12
1.6.2. Upgrading the OSD nodes	14
1.6.3. Upgrading the Ceph Object Gateway nodes	17
1.6.4. Upgrading the CephFS Metadata Server nodes	18
1.6.5. Manually upgrading the Ceph Dashboard node and its operating system	21
1.6.6. Manually upgrading Ceph Ansible nodes and reconfiguring settings	22
1.7. RESTORING THE BACKUP FILES	23
1.8. BACKING UP THE FILES BEFORE THE RHCS UPGRADE	24
1.9. THE UPGRADE PROCESS	25
1.10. CONVERTING THE STORAGE CLUSTER TO USING CEPHADM	30
1.11. INSTALLING CEPHADM-ANSIBLE ON AN UPGRADED STORAGE CLUSTER	31
 CHAPTER 2. UPGRADING A RED HAT CEPH STORAGE CLUSTER RUNNING RED HAT ENTERPRISE LINUX 8 FROM RHCS 4 TO RHCS 5	 33
2.1. PREREQUISITES	34
2.2. COMPATIBILITY CONSIDERATIONS BETWEEN RHCS AND PODMAN VERSIONS	34
2.3. PREPARING FOR AN UPGRADE	35
2.4. CONVERTING TO A CONTAINERIZED DEPLOYMENT	38
2.5. THE UPGRADE PROCESS	40
2.6. CONVERTING THE STORAGE CLUSTER TO USING CEPHADM	45
2.7. INSTALLING CEPHADM-ANSIBLE ON AN UPGRADED STORAGE CLUSTER	46
 CHAPTER 3. UPGRADE A RED HAT CEPH STORAGE CLUSTER USING CEPHADM	 48
3.1. UPGRADING THE RED HAT CEPH STORAGE CLUSTER	48
3.2. UPGRADING THE RED HAT CEPH STORAGE CLUSTER IN A DISCONNECTED ENVIRONMENT	51
3.3. STAGGERED UPGRADE	54
3.3.1. Staggered upgrade options	55
3.3.2. Performing a staggered upgrade	55
3.3.3. Performing a staggered upgrade from previous releases	58
3.4. MONITORING AND MANAGING UPGRADE OF THE STORAGE CLUSTER	60
3.5. TROUBLESHOOTING UPGRADE ERROR MESSAGES	61

CHAPTER 1. UPGRADING A RED HAT CEPH STORAGE CLUSTER FROM RHCS 4 TO RHCS 5


As a storage administrator, you can upgrade a Red Hat Ceph Storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5. The upgrade process includes the following tasks:

- Upgrade the host OS version on the storage cluster from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8, if your storage cluster is still running Red Hat Enterprise Linux 7.
- Upgrade the host OS version on the Ceph Ansible administration node from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8, if the node is still running Red Hat Enterprise Linux 7.
- Use Ansible playbooks to upgrade a Red Hat Ceph Storage 4 storage cluster to Red Hat Ceph Storage 5.



NOTE

If you are upgrading from Red Hat Ceph Storage 4.3 on Red Hat Enterprise Linux 7.9 to Red Hat Ceph Storage 5.2 on Red Hat Enterprise Linux 9, first upgrade the host OS from Red Hat Enterprise Linux 7.9 to Red Hat Enterprise Linux 8.x, upgrade Red Hat Ceph Storage, and then upgrade to Red Hat Enterprise Linux 9.x.



IMPORTANT

If your Red Hat Ceph Storage 4 cluster is already running Red Hat Enterprise Linux 8, see [Upgrading a Red Hat Ceph Storage running Red Hat Enterprise Linux 8 from RHCS4 to RHCS 5](#).



IMPORTANT

leapp does not support upgrades for encrypted OSDs or OSDs that have encrypted partitions. If your OSDs are encrypted and you are upgrading the host OS, disable **dmccrypt** in **ceph-ansible** before upgrading the OS. For more information about using **leapp**, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) and [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).



IMPORTANT

ceph-ansible is currently not supported with Red Hat Ceph Storage 5. This means that once you have migrated your storage cluster to Red Hat Ceph Storage 5, you must use **cephadm** and **cephadm-ansible** to perform subsequent updates.



IMPORTANT

While upgrading from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, do not set **bluestore_fsck_quick_fix_on_mount** parameter to **true** or do not run the **ceph-bluestore-tool --path PATH_TO_OSD --command quick-fix|repair** commands as it might lead to improperly formatted OMAP keys and cause data corruption.

**WARNING**

Upgrading to Red Hat Ceph Storage 5.2 from Red Hat Ceph Storage 5.0 on Ceph Object Gateway storage clusters (single-site or multi-site) is supported but you must set the **ceph config set mgr mgr/cephadm/no_five_one_rgw true --force** option prior to upgrading your storage cluster.

Upgrading to Red Hat Ceph Storage 5.2 from Red Hat Ceph Storage 5.1 on Ceph Object Gateway storage clusters (single-site or multi-site) is not supported due to a known issue. For more information, see the knowledge base article [Support Restrictions for upgrades for RADOS Gateway \(RGW\) on Red Hat Red Hat Ceph Storage 5.2](#).

**NOTE**

Follow the knowledge base article [How to upgrade from Red Hat Ceph Storage 4.2z4 to 5.0z4](#) with the upgrade procedure if you are planning to upgrade to Red Hat Ceph Storage 5.0z4.

**IMPORTANT**

The option **bluefs_buffered_io** is set to **True** by default for Red Hat Ceph Storage. This option enables BlueFS to perform buffered reads in some cases, and enables the kernel page cache to act as a secondary cache for reads like RocksDB block reads. For example, if the RocksDB block cache is not large enough to hold all blocks during the OMAP iteration, it may be possible to read them from the page cache instead of the disk. This can dramatically improve performance when `osd_memory_target` is too small to hold all entries in the block cache. Currently, enabling **bluefs_buffered_io** and disabling the system level swap prevents performance degradation.

For more information about viewing the current setting for **bluefs_buffered_io**, see the [Viewing the bluefs_buffered_io setting](#) section in the *Red Hat Ceph Storage Administration Guide*.

Red Hat Ceph Storage 5 supports only containerized daemons. It does not support non-containerized storage clusters. If you are upgrading a non-containerized storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, the upgrade process includes the conversion to a containerized deployment.

1.1. PREREQUISITES

- A running Red Hat Ceph Storage 4 cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- Root-level access to all nodes in the storage cluster.
- The Ansible user account for use with the Ansible application.

- Red Hat Ceph Storage tools and Ansible repositories are enabled.



IMPORTANT

You can manually upgrade the Ceph File System (CephFS) Metadata Server (MDS) software on a Red Hat Ceph Storage cluster and the Red Hat Enterprise Linux operating system to a new major release at the same time. The underlying XFS filesystem must be formatted with **ftype=1** or with **d_type** support. Run the command **xfs_info /var** to ensure the **ftype** is set to **1**. If the value of **ftype** is not **1**, attach a new disk or create a volume. On top of this new device, create a new XFS filesystem and mount it on **/var/lib/containers**.

Starting with Red Hat Enterprise Linux 8, **mkfs.xfs** enables **ftype=1** by default.

1.2. COMPATIBILITY CONSIDERATIONS BETWEEN RHCS AND PODMAN VERSIONS

podman and Red Hat Ceph Storage have different end-of-life strategies that might make it challenging to find compatible versions.

If you plan to upgrade from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8 as part of the Ceph upgrade process, make sure that the version of **podman** is compatible with Red Hat Ceph Storage 5.


Red Hat recommends to use the **podman** version shipped with the corresponding Red Hat Enterprise Linux version for Red Hat Ceph Storage 5. See the [Red Hat Ceph Storage: Supported configurations](#) knowledge base article for more details. See the [Contacting Red Hat support for service](#) section in the *Red Hat Ceph Storage Troubleshooting Guide* for additional assistance.

1.3. PREPARING FOR AN UPGRADE

As a storage administrator, you can upgrade your Ceph storage cluster to Red Hat Ceph Storage 5. However, some components of your storage cluster must be running specific software versions before an upgrade can take place. The following list shows the minimum software versions that must be installed on your storage cluster before you can upgrade to Red Hat Ceph Storage 5.

- Red Hat Ceph Storage 4.3 or later.
- Ansible 2.9.
- Ceph-ansible shipped with the latest version of Red Hat Ceph Storage.
- Red Hat Enterprise Linux 8.4 EUS or later.
- FileStore OSDs must be migrated to BlueStore. For more information about converting OSDs from FileStore to BlueStore, refer to [BlueStore](#).

There is no direct upgrade path from Red Hat Ceph Storage versions earlier than Red Hat Ceph Storage 4.3. If you are upgrading from Red Hat Ceph Storage 3, you must first upgrade to Red Hat Ceph Storage 4.3 or later, and then upgrade to Red Hat Ceph Storage 5.




IMPORTANT

You can only upgrade to the latest version of Red Hat Ceph Storage 5. For example, if version 5.1 is available, you cannot upgrade from 4 to 5.0; you must go directly to 5.1.

To upgrade your storage cluster to Red Hat Ceph Storage 5, Red Hat recommends that your cluster be running Red Hat Ceph Storage 4.3 or later. Refer to the Knowledgebase article [What are the Red Hat Ceph Storage Releases?](#). This article contains download links to the most recent versions of the Ceph packages and ceph-ansible.

The upgrade process uses Ansible playbooks to upgrade an Red Hat Ceph Storage 4 storage cluster to Red Hat Ceph Storage 5. If your Red Hat Ceph Storage 4 cluster is a non-containerized cluster, the upgrade process includes a step to transform the cluster into a containerized version. Red Hat Ceph Storage 5 does not run on non-containerized clusters.

If you have a mirroring or multisite configuration, upgrade one cluster at a time. Make sure that each upgraded cluster is running properly before upgrading another cluster.



IMPORTANT

leapp does not support upgrades for encrypted OSDs or OSDs that have encrypted partitions. If your OSDs are encrypted and you are upgrading the host OS, disable **dmcrypt** in **ceph-ansible** before upgrading the OS. For more information about using **leapp**, refer to [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).



IMPORTANT

Perform the first three steps in this procedure *only* if the storage cluster is not already running the latest version of Red Hat Ceph Storage 4. The latest version of Red Hat Ceph Storage 4 should be 4.3 or later.

Prerequisites

- A running Red Hat Ceph Storage 4 cluster.
- Sudo-level access to all nodes in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.
- Red Hat Ceph Storage tools and Ansible repositories are enabled.

Procedure

1. Enable the Ceph and Ansible repositories on the Ansible administration node:

Example

```
[root@admin ceph-ansible]# subscription-manager repos --enable=rhceph-4-tools-for-rhel-8-x86_64-rpms --enable=ansible-2.9-for-rhel-8-x86_64-rpms
```

2. Use the **--extra-vars** option to update the **infrastructure-playbooks/rolling_update.yml** playbook and to change the **health_osd_check_retries** and **health_osd_check_delay** values to **50** and **30**, respectively:

Example

```
[root@admin ceph-ansible]# ansible-playbook -i hosts infrastructure-
playbooks/rolling_update.yml --extra-vars "health_osd_check_retries=50
health_osd_check_delay=30"
```

For each OSD node, these values cause **ceph-ansible** to check the storage cluster health every 30 seconds, up to 50 times. This means that **ceph-ansible** waits up to 25 minutes for each OSD.

Adjust the **health_osd_check_retries** option value up or down, based on the used storage capacity of the storage cluster. For example, if you are using 218 TB out of 436 TB, or 50% of the storage capacity, then set the **health_osd_check_retries** option to **50**.

/etc/ansible/hosts is the default location for the Ansible inventory file.

3. If the storage cluster you want to upgrade contains Ceph Block Device images that use the **exclusive-lock** feature, ensure that all Ceph Block Device users have permissions to create a denylist for clients:

Syntax

```
ceph auth caps client.ID mon 'allow r, allow command "osd blacklist"' osd
'EXISTING_OSD_USER_CAPS'
```

4. If the storage cluster was originally installed using Cockpit, create a symbolic link in the **/usr/share/ceph-ansible** directory to the inventory file where Cockpit created it, at **/usr/share/ansible-runner-service/inventory/hosts**:

- a. Change to the **/usr/share/ceph-ansible** directory:

```
# cd /usr/share/ceph-ansible
```

- b. Create the symbolic link:

```
# ln -s /usr/share/ansible-runner-service/inventory/hosts hosts
```

5. To upgrade the cluster using **ceph-ansible**, create the symbolic link in the **etc/ansible/hosts** directory to the **hosts** inventory file:

```
# ln -s /etc/ansible/hosts hosts
```

6. If the storage cluster was originally installed using Cockpit, copy the Cockpit-generated SSH keys to the Ansible user's **~/.ssh** directory:

- a. Copy the keys:

Syntax

```
cp /usr/share/ansible-runner-service/env/ssh_key.pub
```

```
/home/ANSIBLE_USERNAME/.ssh/id_rsa.pub  
cp /usr/share/ansible-runner-service/env/ssh_key  
/home/ANSIBLE_USERNAME/.ssh/id_rsa
```

Replace *ANSIBLE_USERNAME* with the user name for Ansible. The usual default user name is **admin**.

Example

```
# cp /usr/share/ansible-runner-service/env/ssh_key.pub /home/admin/.ssh/id_rsa.pub  
# cp /usr/share/ansible-runner-service/env/ssh_key /home/admin/.ssh/id_rsa
```

- b. Set the appropriate owner, group, and permissions on the key files:

Syntax

```
# chown ANSIBLE_USERNAME:ANSIBLE_USERNAME  
/home/ANSIBLE_USERNAME/.ssh/id_rsa.pub  
# chown ANSIBLE_USERNAME:ANSIBLE_USERNAME  
/home/ANSIBLE_USERNAME/.ssh/id_rsa  
# chmod 644 /home/ANSIBLE_USERNAME/.ssh/id_rsa.pub  
# chmod 600 /home/ANSIBLE_USERNAME/.ssh/id_rsa
```

Replace *ANSIBLE_USERNAME* with the username for Ansible. The usual default user name is **admin**.

Example

```
# chown admin:admin /home/admin/.ssh/id_rsa.pub  
# chown admin:admin /home/admin/.ssh/id_rsa  
# chmod 644 /home/admin/.ssh/id_rsa.pub  
# chmod 600 /home/admin/.ssh/id_rsa
```

Additional Resources

- [What are the Red Hat Ceph Storage Releases?](#)
- For more information about converting from FileStore to BlueStore, refer to [BlueStore](#).

1.4. BACKING UP THE FILES BEFORE THE HOST OS UPGRADE



NOTE

Perform the procedure in this section only if you are upgrading the host OS. If you are not upgrading the host OS, skip this section.

Before you can perform the upgrade procedure, you must make backup copies of the files that you customized for your storage cluster, including keyring files and the **yaml** files for your configuration.

Prerequisites

- A running Red Hat Ceph Storage 4 cluster.

- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.
- Red Hat Ceph Storage Tools and Ansible repositories are enabled.

Procedure

1. Make a backup copy of the **ceph.client.admin.keyring** file.
2. Make backup copies of the **ceph.conf** files from each node.
3. Make backup copies of the **/etc/ganesha/** folder on each node.
4. If the storage cluster has RBD mirroring defined, then make backup copies of the **/etc/ceph** folder and the **group_vars/rbdmirrors.yml** file.

1.5. CONVERTING TO A CONTAINERIZED DEPLOYMENT

This procedure is required for non-containerized clusters. If your storage cluster is a non-containerized cluster, this procedure transforms the cluster into a containerized version.

Red Hat Ceph Storage 5 supports container-based deployments only. A cluster needs to be containerized before upgrading to RHCS 5.x.

If your Red Hat Ceph Storage 4 storage cluster is already containerized, skip this section.



IMPORTANT

This procedure stops and restarts a daemon. If the playbook stops executing during this procedure, be sure to analyze the state of the cluster before restarting.

Prerequisites


- A running Red Hat Ceph Storage non-containerized 4 cluster.
- Root-level access to all nodes in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.

Procedure

1. If you are running a multisite setup, set **rgw_multisite: false** in **all.yml**.
2. Ensure the **group_vars/all.yml** has the following default values for the configuration parameters:

```
ceph_docker_image_tag: "latest"
ceph_docker_registry: "registry.redhat.io"
```

```
ceph_docker_image: rhceph/rhceph-4-rhel8
containerized_deployment: true
```

**NOTE**

These values differ if you use a local registry and a custom image name.

3. Optional: For two-way RBD mirroring configured using the command-line interface in a bare-metal storage cluster, the cluster does not migrate RBD mirroring. For such a configuration, follow the below steps before migrating the non-containerized storage cluster to a containerized deployment:

- a. Create a user on the Ceph client node:

Syntax

```
ceph auth get client.PRIMARY_CLUSTER_NAME -o
/etc/ceph/ceph.PRIMARY_CLUSTER_NAME.keyring
```

Example

```
[root@rbd-client-site-a ~]# ceph auth get client.rbd-mirror.site-a -o
/etc/ceph/ceph.client.rbd-mirror.site-a.keyring
```

- b. Change the username in the **auth** file in **/etc/ceph** directory:

Example

```
[client.rbd-mirror.rbd-client-site-a]
key = AQCbKbVg+E7POBAA7COSZCodvOrg2LWIFc9+3g==
caps mds = "allow *"
caps mgr = "allow *"
caps mon = "allow *"
caps osd = "allow *"
```

- c. Import the **auth** file to add relevant permissions:

Syntax

```
ceph auth import -i PATH_TO_KEYRING
```

Example

```
[root@rbd-client-site-a ~]# ceph auth import -i /etc/ceph/ceph.client.rbd-mirror.rbd-client-
site-a.keyring
```

- d. Check the service name of the RBD mirror node:

Example

```
[root@rbd-client-site-a ~]# systemctl list-units --all
```

```
systemctl stop ceph-rbd-mirror@rbd-client-site-a.service
systemctl disable ceph-rbd-mirror@rbd-client-site-a.service
systemctl reset-failed ceph-rbd-mirror@rbd-client-site-a.service
systemctl start ceph-rbd-mirror@rbd-mirror.rbd-client-site-a.service
systemctl enable ceph-rbd-mirror@rbd-mirror.rbd-client-site-a.service
systemctl status ceph-rbd-mirror@rbd-mirror.rbd-client-site-a.service
```

- e. Add the rbd-mirror node to the **/etc/ansible/hosts** file:

Example

```
[rbdmirrors]
ceph.client.rbd-mirror.rbd-client-site-a
```

4. If you are using daemons that are not containerized, convert them to containerized format:

Syntax

```
ansible-playbook -vvvv -i INVENTORY_FILE infrastructure-playbooks/switch-from-non-
containerized-to-containerized-ceph-daemons.yml
```

The **-vvvv** option collects verbose logs of the conversion process.

Example

```
[ansible@admin ceph-ansible]$ ansible-playbook -vvvv -i hosts infrastructure-
playbooks/switch-from-non-containerized-to-containerized-ceph-daemons.yml
```

5. Once the playbook completes successfully, edit the value of **rgw_multisite: true** in the **all.yml** file and ensure the value of **containerized_deployment** is **true**.

1.6. UPDATING THE HOST OPERATING SYSTEM

Red Hat Ceph Storage 5.2 is supported on Red Hat Enterprise Linux 8.4 EUS, 8.5, 8.6, and 9.0.

This procedure enables you to install Red Hat Ceph Storage 5 and Red Hat Enterprise Linux 8 on the nodes in the storage cluster. If you are already running Red Hat Enterprise Linux 8 on your storage cluster, skip this procedure.

You must manually upgrade all other nodes in the cluster to run the most recent versions of Red Hat Enterprise Linux and Red Hat Ceph Storage.

Prerequisites

- A running Red Hat Ceph Storage 4 storage cluster.
- Sudo-level access to all nodes in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.

- Red Hat Ceph Storage tools and Ansible repositories are enabled.

Procedure

1. Use the **docker-to-podman** playbook to convert docker to podman:

Example


```
[ansible@admin ceph-ansible]$ ansible-playbook -vvvv -i hosts infrastructure-playbooks/  
docker-to-podman.yml
```

Additional Resources

- [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) .
- [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#) .

1.6.1. Manually upgrading Ceph Monitor nodes and their operating systems

As a system administrator, you can manually upgrade the Ceph Monitor software on a Red Hat Ceph Storage cluster node and the Red Hat Enterprise Linux operating system to a new major release at the same time.




IMPORTANT

Perform the procedure on only one Monitor node at a time. To prevent cluster access issues, ensure that the current upgraded Monitor node has returned to normal operation *before* proceeding to the next node.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The nodes are running Red Hat Enterprise Linux 7.9.
- The nodes are using Red Hat Ceph Storage version 4.3 or later.
- Access to the installation source is available for Red Hat Enterprise Linux 8.4 EUS or later.



IMPORTANT

If you are upgrading from Red Hat Ceph Storage 4.3 on Red Hat Enterprise Linux 7.9 to Red Hat Ceph Storage 5.2 on Red Hat Enterprise Linux 9, first upgrade the host OS from Red Hat Enterprise Linux 7.9 to Red Hat Enterprise Linux 8.x, upgrade Red Hat Ceph Storage, and then upgrade to Red Hat Enterprise Linux 9.x.

Procedure

1. Stop the monitor service:

Syntax

```
systemctl stop ceph-mon@MONITOR_ID
```

Replace *MONITOR_ID* with the Monitor node's ID number.

2. If using Red Hat Ceph Storage 4, disable the Red Hat Ceph Storage 4 repositories.

- a. Disable the tools repository:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-4-tools-rpms
```

- b. Disable the mon repository:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-4-mon-rpms
```

3. Install the **leapp** utility.

- For Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).

4. Run through the **leapp** preupgrade checks. See [Assessing upgradability from the command line](#).

5. Set **PermitRootLogin yes** in */etc/ssh/sshd_config*.

6. Restart the OpenSSH SSH daemon:

```
# systemctl restart sshd.service
```

7. Remove the iSCSI module from the Linux kernel:

```
# modprobe -r iscsi
```

8. Reboot the node.

9. Enable the repositories for Red Hat Ceph Storage 5:

- a. Enable the tools repository:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

10. Restore the **ceph-client-admin.keyring** and **ceph.conf** files from a Monitor node which has not been upgraded yet or from a node that has already had those files restored.
11. Verify that the monitor and manager services came back up and that the monitor is in quorum.

Syntax

```
ceph -s
```

On the *mon:* line under *services*, ensure that the node is listed as in *quorum* and not as *out of quorum*.

Example

```
# ceph -s
mon: 3 daemons, quorum node0,node1,node2 (age 2h)
mgr: node0(active, since 2h), standbys: node1, node2
```

12. Repeat the above steps on all Monitor nodes until they have all been upgraded.

Additional Resources

- See [Updating the host operating system](#) for more information.
- See [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) for more information.
- See [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#) for more information.


1.6.2. Upgrading the OSD nodes

As a system administrator, you can manually upgrade the Ceph OSD software on a Red Hat Ceph Storage cluster node and the Red Hat Enterprise Linux operating system to a new major release at the same time.



IMPORTANT

Perform this procedure for each OSD node in the Ceph cluster, but typically only for one OSD node at a time. A maximum of one failure domain's worth of OSD nodes may be performed in parallel. For example, if per-rack replication is in use, one entire rack's OSD nodes can be upgraded in parallel. To prevent data access issues, ensure that the OSDs of the current OSD node have returned to normal operation and that all of the cluster PGs are in the **active+clean** state **before** proceeding to the next OSD.



IMPORTANT

If you are upgrading from Red Hat Ceph Storage 4.3 on Red Hat Enterprise Linux 7.9 to Red Hat Ceph Storage 5.2 on Red Hat Enterprise Linux 9, first upgrade the host OS from Red Hat Enterprise Linux 7.9 to Red Hat Enterprise Linux 8.x, upgrade Red Hat Ceph Storage, and then upgrade to Red Hat Enterprise Linux 9.x.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The nodes are running Red Hat Enterprise Linux 7.9.
- The nodes are using Red Hat Ceph Storage version 4.3 or later.
- Access to the installation source for Red Hat Enterprise Linux 8.4 EUS or later.
- FileStore OSDs must be migrated to BlueStore.

Procedure

1. If you have FileStore OSDs that have not been migrated to BlueStore, run the **filestore-to-bluestore** playbook. For more information about converting OSDs from FileStore to BlueStore, refer to [BlueStore](#).
2. Set the OSD **noout** flag to prevent OSDs from getting marked down during the migration:

Syntax

```
ceph osd set noout
```

3. Set the OSD **nobackfill**, **norecover**, **norrebalance**, **noscrub** and **nodeep-scrub** flags to avoid unnecessary load on the cluster and to avoid any data reshuffling when the node goes down for migration:

Syntax

```
ceph osd set nobackfill
ceph osd set norecover
ceph osd set norebalance
ceph osd set noscrub
ceph osd set nodeep-scrub
```

4. Gracefully shut down all the OSD processes on the node:

Syntax

```
systemctl stop ceph-osd.target
```

5. If using Red Hat Ceph Storage 4, disable the Red Hat Ceph Storage 4 repositories.
 - a. Disable the tools repository:

Syntax

```
subscription-manager repos --disable=rhel-7-server-rhceph-4-tools-rpms
```

- b. Disable the osd repository:

Syntax

```
subscription-manager repos --disable=rhel-7-server-rhceph-4-osd-rpms
```

6. Install the **leapp** utility. See [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
7. Run through the **leapp** preupgrade checks. See [Assessing upgradability from the command line](#).
8. Set **PermitRootLogin yes** in **/etc/ssh/sshd_config**.
9. Restart the OpenSSH SSH daemon:

Syntax

```
systemctl restart sshd.service
```

10. Remove the iSCSI module from the Linux kernel:

Syntax

```
modprobe -r iscsi
```

11. Perform the upgrade by following [Performing the upgrade from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) and [Performing the upgrade from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#)

- a. Enable the tools repository:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

12. Restore the **ceph.conf** file.
13. Unset the **noout**, **nobackfill**, **norecover**, **norebalance**, **noscrub** and **nodeep-scrub** flags:

Syntax

```
ceph osd unset noout
ceph osd unset nobackfill
ceph osd unset norecover
ceph osd unset norebalance
ceph osd unset noscrub
ceph osd unset nodeep-scrub
```

14. Verify that the OSDs are **up** and **in**, and that they are in the **active+clean** state.

Syntax

```
ceph -s
```

On the *osd:* line under *services*;, ensure that all OSDs are **up** and **in**:

Example

```
# ceph -s
osd: 3 osds: 3 up (since 8s), 3 in (since 3M)
```


15. Repeat this procedure on all OSD nodes until they have all been upgraded.

Additional Resources

- Refer to [BlueStore](#) for more information about converting OSDs from FileStore to BlueStore.
- For more information about the **leapp** utility, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For more information about converting docker to podman, see [Updating the host operating system](#).

1.6.3. Upgrading the Ceph Object Gateway nodes

As a system administrator, you can manually upgrade the Ceph Object Gateway (RGW) software on a Red Hat Ceph Storage cluster node and the Red Hat Enterprise Linux operating system to a new major release at the same time.



IMPORTANT

Perform this procedure for each RGW node in the Ceph cluster, but only for one RGW node at a time. To prevent client access issues, ensure that the current upgraded RGW has returned to normal operation before proceeding to upgrade the next node.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The nodes are running Red Hat Enterprise Linux 7.9.
- The nodes are using Red Hat Ceph Storage version 4.3 or later.
- Access to the installation source for Red Hat Enterprise Linux 8.4 EUS, 8.5, 8.6, and 9.0.

Procedure

1. Stop the Ceph Object Gateway service:

Syntax

```
# systemctl stop ceph-radosgw.target
```

2. Disable the Red Hat Ceph Storage 4 tools repository:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-4-tools-rpms
```

3. Install the **leapp** utility.

- For Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).

4. Run through the **leapp** preupgrade checks:

- For Red Hat Enterprise Linux 8, see [Assessing upgradability from the command line](#).
- For Red Hat Enterprise Linux 9, see [Assessing upgradability from the command line](#).

5. Set **PermitRootLogin yes** in **/etc/ssh/sshd_config**.

6. Restart the OpenSSH SSH daemon:

```
# systemctl restart sshd.service
```

7. Remove the iSCSI module from the Linux kernel:

```
# modprobe -r iscsi
```

8. Perform the upgrade by following [Performing the upgrade from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).

9. Enable the tools repository:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

10. Restore the **ceph-client-admin.keyring** and **ceph.conf** files.

11. Verify that the daemon is active:

Syntax

```
ceph -s
```

View the *rgw:* line under *services:* to make sure that the RGW daemon is active.

Example

```
rgw: 1 daemon active (node4.rgw0)
```

12. Repeat the above steps on all Ceph Object Gateway nodes until they have all been upgraded.

Additional Resources

- See [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) for more information about the **leapp** utility.
- See [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#) for more information about the **leapp** utility.

1.6.4. Upgrading the CephFS Metadata Server nodes

As a storage administrator, you can manually upgrade the Ceph File System (CephFS) Metadata Server (MDS) software on a Red Hat Ceph Storage cluster and the Red Hat Enterprise Linux operating system to a new major release at the same time.



IMPORTANT

Before you upgrade the storage cluster, reduce the number of active MDS ranks to one per file system. This eliminates any possible version conflicts between multiple MDS. In addition, take all standby nodes offline before upgrading.

This is because the MDS cluster does not possess built-in versioning or file system flags. Without these features, multiple MDS might communicate using different versions of the MDS software, and could cause assertions or other faults to occur.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The nodes are running Red Hat Enterprise Linux 7.9.
- The nodes are using Red Hat Ceph Storage version 4.3 or later.
- Access to the installation source for Red Hat Enterprise Linux 8.4 EUS, 8.5, 8.6, and 9.0.
- Root-level access to all nodes in the storage cluster.

Procedure

1. Reduce the number of active MDS ranks to 1:

Syntax

```
ceph fs set FILE_SYSTEM_NAME max_mds 1
```

Example

```
[root@mds ~]# ceph fs set fs1 max_mds 1
```

2. Wait for the cluster to stop all of the MDS ranks. When all of the MDS have stopped, only rank 0 should be active. The rest should be in standby mode. Check the status of the file system:

```
[root@mds ~]# ceph status
```

3. Use **systemctl** to take all standby MDS offline:

```
[root@mds ~]# systemctl stop ceph-mds.target
```

4. Confirm that only one MDS is online, and that it has rank 0 for the file system:

```
[ceph: root@host01 /]# ceph status
```

5. Disable the Red Hat Ceph Storage 4 tools repository:

```
[root@mds ~]# subscription-manager repos --disable=rhel-7-server-rhceph-4-tools-rpms
```

6. Install the **leapp** utility.

- For Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat](#)

[Enterprise Linux 8.](#)

- For Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9.](#)

7. Run through the **leapp** preupgrade checks:

- For Red Hat Enterprise Linux 8, see [Assessing upgradability from the command line](#) .
- For Red Hat Enterprise Linux 9, see [Assessing upgradability from the command line](#) .

8. Edit **/etc/ssh/sshd_config** and set **PermitRootLogin** to **yes**.

9. Restart the OpenSSH SSH daemon:

```
[root@mds ~]# systemctl restart sshd.service
```

10. Remove the iSCSI module from the Linux kernel:

```
[root@mds ~]# modprobe -r iscsi
```

11. Perform the upgrade:

- For Red Hat Enterprise Linux 8, see [Performing the upgrade from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8.](#)
- For Red Hat Enterprise Linux 9, see [Performing the upgrade from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9.](#)

12. Enable the tools repository:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

13. Restore the **ceph-client-admin.keyring** and **ceph.conf** files.

14. Verify that the daemon is active:

```
[root@mds ~]# ceph -s
```

15. Follow the same processes for the standby daemons.

16. When you have finished restarting all of the MDS in standby, restore the previous value of **max_mds** for your cluster:

Syntax

```
ceph fs set FILE_SYSTEM_NAME max_mds ORIGINAL_VALUE
```

Example

```
[root@mds ~]# ceph fs set fs1 max_mds 5
```

Additional Resources

- See [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) for more information about the **leapp** utility.

1.6.5. Manually upgrading the Ceph Dashboard node and its operating system

As a system administrator, you can manually upgrade the Ceph Dashboard software on a Red Hat Ceph Storage cluster node and the Red Hat Enterprise Linux operating system to a new major release at the same time.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The node is running Red Hat Enterprise Linux 7.9.
- The node is running Red Hat Ceph Storage version 4.3 or later.
- Access is available to the installation source for Red Hat Enterprise Linux 8.4 EUS, 8.5, 8.6, or 9.0.

Procedure

1. Disable the Red Hat Ceph Storage 4 tools repository:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-4-tools-rpms
```

2. Install the **leapp** utility.

- For Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).

3. Run through the **leapp** preupgrade checks:

- For Red Hat Enterprise Linux 8, see [Assessing upgradability from the command line](#).
- For Red Hat Enterprise Linux 9, see [Assessing upgradability from the command line](#).

4. Set **PermitRootLogin yes** in **/etc/ssh/sshd_config**.

5. Restart the OpenSSH SSH daemon:

```
# systemctl restart sshd.service
```

6. Remove the iSCSI module from the Linux kernel:

```
# modprobe -r iscsi
```

7. Perform the upgrade:

- For Red Hat Enterprise Linux 8, see [Performing the upgrade from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For Red Hat Enterprise Linux 9, see [Performing the upgrade from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).

8. Enable the tools repository for Red Hat Ceph Storage 5:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9


```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

Additional Resources

- See [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) for more information.

1.6.6. Manually upgrading Ceph Ansible nodes and reconfiguring settings

Manually upgrade the Ceph Ansible software on a Red Hat Ceph Storage cluster node and the Red Hat Enterprise Linux operating system to a new major release at the same time.



IMPORTANT

Before upgrading the host OS on the Ceph Ansible nodes, back up the **group_vars** and **hosts** files. Use the created backups before reconfiguring the Ceph Ansible nodes.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The node is running Red Hat Enterprise Linux 7.9.
- The node is running Red Hat Ceph Storage version 4.2z2 or later.
- Access is available to the installation source for Red Hat Enterprise Linux 8.4 EUS or Red Hat Enterprise Linux 8.5.

Procedure

1. Disable the tools repository for Red Hat Ceph Storage 4 for Red Hat Enterprise Linux 8:

```
[root@ansible ~]# subscription-manager repos --disable=rhceph-4-tools-for-rhel-8-x86_64-rpms
[root@ansible ~]# subscription-manager repos --disable=ansible-2.9-for-rhel-8-x86_64-rpms
```

2. Install the **leapp** utility.

- For Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).

3. Run through the **leapp** preupgrade checks:

- For Red Hat Enterprise Linux 8, see [Assessing upgradability from the command line](#).
- For Red Hat Enterprise Linux 9, see [Assessing upgradability from the command line](#).

4. Edit **/etc/ssh/sshd_config** and set **PermitRootLogin** to **yes**.

5. Restart the OpenSSH SSH daemon:

```
[root@mds ~]# systemctl restart sshd.service
```

6. Remove the iSCSI module from the Linux kernel:

```
[root@mds ~]# modprobe -r iscsi
```

7. Perform the upgrade:

- For Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).

8. Enable the tools repository for Red Hat Ceph Storage 5:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

9. Restore the **ceph-client-admin.keyring** and **ceph.conf** files.

Additional Resources

- See [Updating the host operating system](#) for more information.
- See [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#) for more information.

1.7. RESTORING THE BACKUP FILES

After you have completed the host OS upgrade on each node in your storage cluster, restore all the files that you backed up earlier to each node so that your upgraded node uses your preserved settings.

Repeat this process on each host in your storage cluster after the OS upgrade process for that host is complete.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to all nodes in the storage cluster.

Procedure

1. Restore the files that you backed up before the host OS upgrade to the host.
2. Restore the **/etc/ceph** folders and their contents to all of the hosts, including the **ceph.client.admin.keyring** and **ceph.conf** files.
3. Restore the **/etc/ganesha/** folder to each node.
4. Check to make sure that the ownership for each of the backed-up files has not changed after the operating system upgrade. The file owner should be **ceph**. If the file owner has been changed to **root**, use the following command on each file to change the ownership back to **ceph**:

Example

```
[root@admin]# chown ceph: ceph.client.rbd-mirror.node1.keyring
```

5. If you upgraded from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8 and the storage cluster had RBD mirroring defined, restore the **/etc/ceph** folder from the backup copy.
6. Restore the **group_vars/rbdmirrors.yml** file that you backed up earlier.

1.8. BACKING UP THE FILES BEFORE THE RHCS UPGRADE

Before you run the **rolling_update.yml** playbook to upgrade Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, make backup copies of all the **yml** files.

Prerequisites

- A Red Hat Ceph Storage 4 cluster running RHCS 4.3 or later.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.
- Red Hat Ceph Storage tools and Ansible repositories are enabled.

Procedure

- Make backup copies of all the **yml** files.

Example

```
[root@admin ceph-ansible]# cp group_vars/all.yml group_vars/all_old.yml
[root@admin ceph-ansible]# cp group_vars/osds.yml group_vars/osds_old.yml
[root@admin ceph-ansible]# cp group_vars/mdss.yml group_vars/mdss_old.yml
[root@admin ceph-ansible]# cp group_vars/rgws.yml group_vars/rgws_old.yml
[root@admin ceph-ansible]# cp group_vars/clients.yml group_vars/clients_old.yml
```

1.9. THE UPGRADE PROCESS

As a storage administrator, you use Ansible playbooks to upgrade an Red Hat Ceph Storage 4 storage cluster to Red Hat Ceph Storage 5. The **rolling_update.yml** Ansible playbook performs upgrades for deployments of Red Hat Ceph Storage. The **ceph-ansible** upgrades the Ceph nodes in the following order:

- Ceph Monitor
- Ceph Manager
- Ceph OSD nodes
- MDS nodes
- Ceph Object Gateway (RGW) nodes
- Ceph RBD-mirror node
- Ceph NFS nodes
- Ceph iSCSI gateway node
- Ceph client nodes
- Ceph-crash daemons
- Node-exporter on all nodes
- Ceph Dashboard



NOTE

Red Hat Ceph Storage 5 supports only containerized deployments.

ceph-ansible is currently not supported with Red Hat Ceph Storage 5. This means that once you have migrated your storage cluster to Red Hat Ceph Storage 5, you must use **cephadm** to perform subsequent updates.



NOTE

Red Hat Ceph Storage 5 also includes a health check function that returns a `DAEMON_OLD_VERSION` warning if it detects that any of the daemons in the storage cluster are running multiple versions of Red Hat Ceph Storage. The warning is triggered when the daemons continue to run multiple versions of Red Hat Ceph Storage beyond the time value set in the `mon_warn_older_version_delay` option. By default, the `mon_warn_older_version_delay` option is set to one week. This setting allows most upgrades to proceed without falsely seeing the warning. If the upgrade process is paused for an extended time period, you can mute the health warning:


```
ceph health mute DAEMON_OLD_VERSION --sticky
```

After the upgrade has finished, unmute the health warning:

```
ceph health unmute DAEMON_OLD_VERSION
```

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to all hosts in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The latest versions of Ansible and **ceph-ansible** available with Red Hat Ceph Storage 5.
- The **ansible** user account for use with the Ansible application.
- The nodes of the storage cluster is upgraded to Red Hat Enterprise Linux 8.4 EUS or later.



IMPORTANT

The Ansible inventory file must be present in the **ceph-ansible** directory.

Procedure

1. Enable the Ceph and Ansible repositories on the Ansible administration node:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms --enable=ansible-2.9-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

2. On the Ansible administration node, ensure that the latest versions of the **ansible** and **ceph-ansible** packages are installed.

Syntax

```
dnf update ansible ceph-ansible
```

3. Navigate to the **/usr/share/ceph-ansible/** directory:

Example

```
[root@admin ~]# cd /usr/share/ceph-ansible
```

4. If upgrading from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, make copies of the **group_vars/osds.yml.sample** and **group_vars/clients.yml.sample** files, and rename them to **group_vars/osds.yml**, and **group_vars/clients.yml** respectively.

Example

```
[root@admin ceph-ansible]# cp group_vars/osds.yml.sample group_vars/osds.yml
[root@admin ceph-ansible]# cp group_vars/mdss.yml.sample group_vars/mdss.yml
[root@admin ceph-ansible]# cp group_vars/rgws.yml.sample group_vars/rgws.yml
[root@admin ceph-ansible]# cp group_vars/clients.yml.sample group_vars/clients.yml
```

5. If upgrading from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, edit the **group_vars/all.yml** file to add Red Hat Ceph Storage 5 details.
6. Once you have done the above two steps, copy the settings from the old **yaml** files to the new **yaml** files. Do not change the values of **ceph_rhcs_version**, **ceph_docker_image**, and **grafana_container_image** as the values for these configuration parameters are for Red Hat Ceph Storage 5. This ensures that all the settings related to your cluster are present in the current **yaml** file.

Example

```
fetch_directory: ~/ceph-ansible-keys
monitor_interface: eth0
public_network: 192.168.0.0/24
ceph_docker_registry_auth: true
ceph_docker_registry_username: SERVICE_ACCOUNT_USER_NAME
ceph_docker_registry_password: TOKEN
dashboard_admin_user:
dashboard_admin_password:
grafana_admin_user:
grafana_admin_password:
radosgw_interface: eth0
ceph_docker_image: "rhceph/rhceph-5-rhel8"
ceph_docker_image_tag: "latest"
ceph_docker_registry: "registry.redhat.io"
node_exporter_container_image: registry.redhat.io/openshift4/ose-prometheus-node-exporter:v4.6
grafana_container_image: registry.redhat.io/rhceph/rhceph-5-dashboard-rhel8:5
prometheus_container_image: registry.redhat.io/openshift4/ose-prometheus:v4.6
alertmanager_container_image: registry.redhat.io/openshift4/ose-prometheus-alertmanager:v4.6
```

**NOTE**

Ensure the Red Hat Ceph Storage 5 container images are set to the default values.

7. Edit the **group_vars/osds.yml** file. Add and set the following options:

Syntax

```
nb_retry_wait_osd_up: 50
delay_wait_osd_up: 30
```

8. Open the **group_vars/all.yml** file and verify the following values are present from the old **all.yml** file.

- a. The **fetch_directory** option is set with the same value from the old **all.yml** file:

Syntax

```
fetch_directory: FULL_DIRECTORY_PATH
```

Replace *FULL_DIRECTORY_PATH* with a writable location, such as the Ansible user's home directory.

- b. If the cluster you want to upgrade contains any Ceph Object Gateway nodes, add the **radosgw_interface** option:

```
radosgw_interface: INTERFACE
```

Replace *INTERFACE* with the interface to which the Ceph Object Gateway nodes listen.

- c. If your current setup has SSL certificates configured, edit the following:

Syntax

```
radosgw_frontend_ssl_certificate: /etc/pki/ca-trust/extracted/CERTIFICATE_NAME
radosgw_frontend_port: 443
```

- d. Uncomment the **upgrade_ceph_packages** option and set it to **True**:

Syntax

```
upgrade_ceph_packages: True
```

- e. If the storage cluster has more than one rgw instance per node, then uncomment the **radosgw_num_instances** setting and set it to the number of instances per node in the cluster:

Syntax

```
radosgw_num_instances : NUMBER_OF_INSTANCES_PER_NODE
```

Example

```
radosgw_num_instances : 2
```

- f. If the storage cluster has RGW multisite defined, check the multisite settings in **all.yml** to make sure that they contain the same values as they did in the old **all.yml** file.
9. Log in as **ansible-user** on the Ansible administration node.
10. Execute the **rolling_update.yml** playbook to convert the storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5:

Syntax

```
ansible-playbook -vvvv infrastructure-playbooks/rolling_update.yml -i INVENTORY_FILE
```

The `-vvvv` option collects verbose logs of the upgrade process.

Example

```
[ansible@admin ceph-ansible]$ ansible-playbook -vvvv infrastructure-  
playbooks/rolling_update.yml -i hosts
```



IMPORTANT

Using the **--limit** Ansible option with the **rolling_update.yml** playbook is not supported.

11. Review the Ansible playbook log output to verify the status of the upgrade.

Verification

1. List all running containers:

Example

```
[root@mon ~]# podman ps
```

2. Check the health status of the cluster. Replace *MONITOR_ID* with the name of the Ceph Monitor container found in the previous step:

Syntax

```
podman exec ceph-mon-MONITOR_ID ceph -s
```

Example

```
[root@mon ~]# podman exec ceph-mon-mon01 ceph -s
```

3. Verify the Ceph cluster daemon versions to confirm the upgrade of all daemons. Replace *MONITOR_ID* with the name of the Ceph Monitor container found in the previous step:

Syntax

```
■
```

```
podman exec ceph-mon-MONITOR_ID ceph --cluster ceph versions
```


Example

```
[root@mon ~]# podman exec ceph-mon-mon01 ceph --cluster ceph versions
```

1.10. CONVERTING THE STORAGE CLUSTER TO USING CEPHADM

After you have upgraded the storage cluster to Red Hat Ceph Storage 5, run the **cephadm-adopt** playbook to convert the storage cluster daemons to run **cephadm**.

The **cephadm-adopt** playbook adopts the Ceph services, installs all **cephadm** dependencies, enables the **cephadm** Orchestrator backend, generates and configures the **ssh** key on all hosts, and adds the hosts to the Orchestrator configuration.



NOTE

After you run the **cephadm-adopt** playbook, remove the **ceph-ansible** package. The cluster daemons no longer work with **ceph-ansible**. You must use **cephadm** to manage the cluster daemons.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to all nodes in the storage cluster.

Procedure

1. Log in to the **ceph-ansible** node and change directory to **/usr/share/ceph-ansible**.
2. Run the **cephadm-adopt** playbook:

Syntax

```
ansible-playbook infrastructure-playbooks/cephadm-adopt.yml -i INVENTORY_FILE
```

Example

```
[ansible@admin ceph-ansible]$ ansible-playbook infrastructure-playbooks/cephadm-adopt.yml -i hosts
```

3. Run the following command to enable applications to run on the NFS-Ganesha pool. *POOL_NAME* is **nfs-ganesha**, and *APPLICATION_NAME* is the name of the application you want to enable, such as **cephfs**, **rbd**, or **rgw**.

Syntax

```
ceph osd pool application enable POOL_NAME APPLICATION_NAME
```

Example

```
[root@node0 ~]# ceph osd pool application enable nfs-ganesha rgw
```

IMPORTANT

The **cephadm-adopt** playbook does not bring up rbd-mirroring after migrating the storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5.

To work around this issue, add the peers manually:

Syntax

```
rbd mirror pool peer add POOL_NAME CLIENT_NAME@CLUSTER_NAME
```

Example

```
[ceph: root@node0 /]# rbd --cluster site-a mirror pool peer add image-pool  
client.rbd-mirror-peer@site-b
```

Additional Resources

- For more information about using **leapp** to upgrade Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For more information about using **leapp** to upgrade Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).
- For more information about converting from FileStore to BlueStore, refer to [BlueStore](#).
- For more information about storage peers, see [Viewing information about peers](#).

1.11. INSTALLING CEPHADM-ANSIBLE ON AN UPGRADED STORAGE CLUSTER

cephadm-ansible is a collection of Ansible playbooks to simplify workflows that are not covered by **cephadm**. After installation, the playbooks are located in **/usr/share/cephadm-ansible/**.



NOTE

Before adding new nodes or new clients to your upgraded storage cluster, run the **cephadm-preflight.yml** playbook.

Prerequisites

- Root-level access to the Ansible administration node.
- A valid Red Hat subscription with the appropriate entitlements.
- An active Red Hat Network (RHN) or service account to access the Red Hat Registry.

Procedure

- Install the **cephadm-ansible** package:

Syntax

```
dnf install cephadm-ansible
```

Additional Resources

- [Running the preflight playbook](#)
- [Adding hosts](#)
- [Adding Monitor service](#)
- [Adding Manager service](#)
- [Adding OSDs](#)
- For more information about configuring clients and services, see [Red Hat Ceph Storage Operations Guide](#).
- For more information about the **cephadm-ansible** playbooks, see [The cephadm-ansible playbooks](#).

CHAPTER 2. UPGRADING A RED HAT CEPH STORAGE CLUSTER RUNNING RED HAT ENTERPRISE LINUX 8 FROM RHCS 4 TO RHCS 5


As a storage administrator, you can upgrade a Red Hat Ceph Storage cluster running Red Hat Enterprise Linux 8 from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5. The upgrade process includes the following tasks:

- Use Ansible playbooks to upgrade a Red Hat Ceph Storage 4 storage cluster to Red Hat Ceph Storage 5.



IMPORTANT

ceph-ansible is currently not supported with Red Hat Ceph Storage 5. This means that once you have migrated your storage cluster to Red Hat Ceph Storage 5, you must use **cephadm** and **cephadm-ansible** to perform subsequent updates.



IMPORTANT

While upgrading from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, do not set **bluestore_fsck_quick_fix_on_mount** parameter to **true** or do not run the **ceph-bluestore-tool --path *PATH_TO_OSD* --command quick-fix|repair** commands as it might lead to improperly formatted OMAP keys and cause data corruption.



WARNING


Upgrading to Red Hat Ceph Storage 5.2 from Red Hat Ceph Storage 5.0 on Ceph Object Gateway storage clusters (single-site or multi-site) is supported but you must set the **ceph config set mgr mgr/cephadm/no_five_one_rgw true --force** option prior to upgrading your storage cluster.

Upgrading to Red Hat Ceph Storage 5.2 from Red Hat Ceph Storage 5.1 on Ceph Object Gateway storage clusters (single-site or multi-site) is not supported due to a known issue. For more information, see the knowledge base article [Support Restrictions for upgrades for RADOS Gateway \(RGW\) on Red Hat Red Hat Ceph Storage 5.2](#).



NOTE

Follow the knowledge base article [How to upgrade from Red Hat Ceph Storage 4.2z4 to 5.0z4](#) with the upgrade procedure if you are planning to upgrade to Red Hat Ceph Storage 5.0z4.



IMPORTANT


The option **bluefs_buffered_io** is set to **True** by default for Red Hat Ceph Storage. This option enables BlueFS to perform buffered reads in some cases, and enables the kernel page cache to act as a secondary cache for reads like RocksDB block reads. For example, if the RocksDB block cache is not large enough to hold all blocks during the OMAP iteration, it may be possible to read them from the page cache instead of the disk. This can dramatically improve performance when `osd_memory_target` is too small to hold all entries in the block cache. Currently, enabling **bluefs_buffered_io** and disabling the system level swap prevents performance degradation.

For more information about viewing the current setting for **bluefs_buffered_io**, see the [Viewing the **bluefs_buffered_io** setting](#) section in the *Red Hat Ceph Storage Administration Guide*.

Red Hat Ceph Storage 5 supports only containerized daemons. It does not support non-containerized storage clusters. If you are upgrading a non-containerized storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, the upgrade process includes the conversion to a containerized deployment.

2.1. PREREQUISITES

- A Red Hat Ceph Storage 4 cluster running Red Hat Enterprise Linux 8.4 or later.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- Root-level access to all nodes in the storage cluster.
- The Ansible user account for use with the Ansible application.
- Red Hat Ceph Storage tools and Ansible repositories are enabled.



IMPORTANT

You can manually upgrade the Ceph File System (CephFS) Metadata Server (MDS) software on a Red Hat Ceph Storage cluster and the Red Hat Enterprise Linux operating system to a new major release at the same time. The underlying XFS filesystem must be formatted with **ftype=1** or with **d_type** support. Run the command **xfs_info /var** to ensure the **ftype** is set to **1**. If the value of **ftype** is not **1**, attach a new disk or create a volume. On top of this new device, create a new XFS filesystem and mount it on **/var/lib/containers**.

Starting with Red Hat Enterprise Linux 8, **mkfs.xfs** enables **ftype=1** by default.

2.2. COMPATIBILITY CONSIDERATIONS BETWEEN RHCS AND PODMAN VERSIONS

podman and Red Hat Ceph Storage have different end-of-life strategies that might make it challenging to find compatible versions.

If you plan to upgrade from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8 as part of the Ceph upgrade process, make sure that the version of **podman** is compatible with Red Hat Ceph Storage 5.

Red Hat recommends to use the **podman** version shipped with the corresponding Red Hat Enterprise Linux version for Red Hat Ceph Storage 5. See the [Red Hat Ceph Storage: Supported configurations](#) knowledge base article for more details. See the [Contacting Red Hat support for service](#) section in the *Red Hat Ceph Storage Troubleshooting Guide* for additional assistance.

2.3. PREPARING FOR AN UPGRADE

As a storage administrator, you can upgrade your Ceph storage cluster to Red Hat Ceph Storage 5. However, some components of your storage cluster must be running specific software versions before an upgrade can take place. The following list shows the minimum software versions that must be installed on your storage cluster before you can upgrade to Red Hat Ceph Storage 5.

- Red Hat Ceph Storage 4.3 or later.
- Ansible 2.9.
- Ceph-ansible shipped with the latest version of Red Hat Ceph Storage.
- Red Hat Enterprise Linux 8.4 EUS or later.
- FileStore OSDs must be migrated to BlueStore. For more information about converting OSDs from FileStore to BlueStore, refer to [BlueStore](#).

There is no direct upgrade path from Red Hat Ceph Storage versions earlier than Red Hat Ceph Storage 4.3. If you are upgrading from Red Hat Ceph Storage 3, you must first upgrade to Red Hat Ceph Storage 4.3 or later, and then upgrade to Red Hat Ceph Storage 5.



IMPORTANT

You can only upgrade to the latest version of Red Hat Ceph Storage 5. For example, if version 5.1 is available, you cannot upgrade from 4 to 5.0; you must go directly to 5.1.

To upgrade your storage cluster to Red Hat Ceph Storage 5, Red Hat recommends that your cluster be running Red Hat Ceph Storage 4.3 or later. Refer to the Knowledgebase article [What are the Red Hat Ceph Storage Releases?](#). This article contains download links to the most recent versions of the Ceph packages and ceph-ansible.


The upgrade process uses Ansible playbooks to upgrade an Red Hat Ceph Storage 4 storage cluster to Red Hat Ceph Storage 5. If your Red Hat Ceph Storage 4 cluster is a non-containerized cluster, the upgrade process includes a step to transform the cluster into a containerized version. Red Hat Ceph Storage 5 does not run on non-containerized clusters.

If you have a mirroring or multisite configuration, upgrade one cluster at a time. Make sure that each upgraded cluster is running properly before upgrading another cluster.



IMPORTANT

leapp does not support upgrades for encrypted OSDs or OSDs that have encrypted partitions. If your OSDs are encrypted and you are upgrading the host OS, disable **dmccrypt** in **ceph-ansible** before upgrading the OS. For more information about using **leapp**, refer to [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).



IMPORTANT

Perform the first three steps in this procedure *only* if the storage cluster is not already running the latest version of Red Hat Ceph Storage 4. The latest version of Red Hat Ceph Storage 4 should be 4.3 or later.

Prerequisites

- A running Red Hat Ceph Storage 4 cluster.
- Sudo-level access to all nodes in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.
- Red Hat Ceph Storage tools and Ansible repositories are enabled.

Procedure

1. Enable the Ceph and Ansible repositories on the Ansible administration node:

Example

```
[root@admin ceph-ansible]# subscription-manager repos --enable=rhceph-4-tools-for-rhel-8-x86_64-rpms --enable=ansible-2.9-for-rhel-8-x86_64-rpms
```

2. Use the **--extra-vars** option to update the **infrastructure-playbooks/rolling_update.yml** playbook and to change the **health_osd_check_retries** and **health_osd_check_delay** values to **50** and **30**, respectively:

Example

```
[root@admin ceph-ansible]# ansible-playbook -i hosts infrastructure-playbooks/rolling_update.yml --extra-vars "health_osd_check_retries=50 health_osd_check_delay=30"
```

For each OSD node, these values cause **ceph-ansible** to check the storage cluster health every 30 seconds, up to 50 times. This means that **ceph-ansible** waits up to 25 minutes for each OSD.

Adjust the **health_osd_check_retries** option value up or down, based on the used storage capacity of the storage cluster. For example, if you are using 218 TB out of 436 TB, or 50% of the storage capacity, then set the **health_osd_check_retries** option to **50**.

/etc/ansible/hosts is the default location for the Ansible inventory file.

3. If the storage cluster you want to upgrade contains Ceph Block Device images that use the **exclusive-lock** feature, ensure that all Ceph Block Device users have permissions to create a denylist for clients:

Syntax

```
ceph auth caps client.ID mon 'allow r, allow command "osd blacklist"' osd
'EXISTING_OSD_USER_CAPS'
```

4. If the storage cluster was originally installed using Cockpit, create a symbolic link in the **/usr/share/ceph-ansible** directory to the inventory file where Cockpit created it, at **/usr/share/ansible-runner-service/inventory/hosts**:

- a. Change to the **/usr/share/ceph-ansible** directory:

```
# cd /usr/share/ceph-ansible
```

- b. Create the symbolic link:

```
# ln -s /usr/share/ansible-runner-service/inventory/hosts hosts
```

5. To upgrade the cluster using **ceph-ansible**, create the symbolic link in the **etc/ansible/hosts** directory to the **hosts** inventory file:

```
# ln -s /etc/ansible/hosts hosts
```

6. If the storage cluster was originally installed using Cockpit, copy the Cockpit-generated SSH keys to the Ansible user's **~/.ssh** directory:

- a. Copy the keys:

Syntax

```
cp /usr/share/ansible-runner-service/env/ssh_key.pub
/home/ANSIBLE_USERNAME/.ssh/id_rsa.pub
cp /usr/share/ansible-runner-service/env/ssh_key
/home/ANSIBLE_USERNAME/.ssh/id_rsa
```

Replace **ANSIBLE_USERNAME** with the user name for Ansible. The usual default user name is **admin**.

Example

```
# cp /usr/share/ansible-runner-service/env/ssh_key.pub /home/admin/.ssh/id_rsa.pub
# cp /usr/share/ansible-runner-service/env/ssh_key /home/admin/.ssh/id_rsa
```

- b. Set the appropriate owner, group, and permissions on the key files:

Syntax

```
# chown ANSIBLE_USERNAME:ANSIBLE_USERNAME
/home/ANSIBLE_USERNAME/.ssh/id_rsa.pub
# chown ANSIBLE_USERNAME:ANSIBLE_USERNAME
/home/ANSIBLE_USERNAME/.ssh/id_rsa
# chmod 644 /home/ANSIBLE_USERNAME/.ssh/id_rsa.pub
# chmod 600 /home/ANSIBLE_USERNAME/.ssh/id_rsa
```

Replace **ANSIBLE_USERNAME** with the username for Ansible. The usual default user name is **admin**.

Example

```
# chown admin:admin /home/admin/.ssh/id_rsa.pub
# chown admin:admin /home/admin/.ssh/id_rsa
# chmod 644 /home/admin/.ssh/id_rsa.pub
# chmod 600 /home/admin/.ssh/id_rsa
```

Additional Resources


- [What are the Red Hat Ceph Storage Releases?](#)
- For more information about converting from FileStore to BlueStore, refer to [BlueStore](#).

2.4. CONVERTING TO A CONTAINERIZED DEPLOYMENT

This procedure is required for non-containerized clusters. If your storage cluster is a non-containerized cluster, this procedure transforms the cluster into a containerized version.

Red Hat Ceph Storage 5 supports container-based deployments only. A cluster needs to be containerized before upgrading to RHCS 5.x.

If your Red Hat Ceph Storage 4 storage cluster is already containerized, skip this section.



IMPORTANT

This procedure stops and restarts a daemon. If the playbook stops executing during this procedure, be sure to analyze the state of the cluster before restarting.

Prerequisites

- A running Red Hat Ceph Storage non-containerized 4 cluster.
- Root-level access to all nodes in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The Ansible user account for use with the Ansible application.

Procedure

1. If you are running a multisite setup, set **rgw_multisite: false** in **all.yml**.
2. Ensure the **group_vars/all.yml** has the following default values for the configuration parameters:

```
ceph_docker_image_tag: "latest"
ceph_docker_registry: "registry.redhat.io"
ceph_docker_image: rhceph/rhceph-4-rhel8
containerized_deployment: true
```



NOTE

These values differ if you use a local registry and a custom image name.

3. Optional: For two-way RBD mirroring configured using the command-line interface in a bare-metal storage cluster, the cluster does not migrate RBD mirroring. For such a configuration, follow the below steps before migrating the non-containerized storage cluster to a containerized deployment:

- a. Create a user on the Ceph client node:

Syntax

```
ceph auth get client.PRIMARY_CLUSTER_NAME -o
/etc/ceph/ceph.PRIMARY_CLUSTER_NAME.keyring
```

Example

```
[root@rbd-client-site-a ~]# ceph auth get client.rbd-mirror.site-a -o
/etc/ceph/ceph.client.rbd-mirror.site-a.keyring
```

- b. Change the username in the **auth** file in **/etc/ceph** directory:

Example

```
[client.rbd-mirror.rbd-client-site-a]
key = AQCbKbVg+E7POBAA7COSZCodvOrg2LWIFc9+3g==
caps mds = "allow *"
caps mgr = "allow *"
caps mon = "allow *"
caps osd = "allow *"
```

- c. Import the **auth** file to add relevant permissions:

Syntax

```
ceph auth import -i PATH_TO_KEYRING
```

Example

```
[root@rbd-client-site-a ~]# ceph auth import -i /etc/ceph/ceph.client.rbd-mirror.rbd-client-
site-a.keyring
```

- d. Check the service name of the RBD mirror node:

Example

```
[root@rbd-client-site-a ~]# systemctl list-units --all

systemctl stop ceph-rbd-mirror@rbd-client-site-a.service
systemctl disable ceph-rbd-mirror@rbd-client-site-a.service
systemctl reset-failed ceph-rbd-mirror@rbd-client-site-a.service
```

```
systemctl start ceph-rbd-mirror@rbd-mirror.rbd-client-site-a.service
systemctl enable ceph-rbd-mirror@rbd-mirror.rbd-client-site-a.service
systemctl status ceph-rbd-mirror@rbd-mirror.rbd-client-site-a.service
```

- e. Add the rbd-mirror node to the **/etc/ansible/hosts** file:

Example

```
[rbdmirrors]
ceph.client.rbd-mirror.rbd-client-site-a
```

4. If you are using daemons that are not containerized, convert them to containerized format:

Syntax

```
ansible-playbook -vvvv -i INVENTORY_FILE infrastructure-playbooks/switch-from-non-
containerized-to-containerized-ceph-daemons.yml
```

The **-vvvv** option collects verbose logs of the conversion process.

Example

```
[ansible@admin ceph-ansible]$ ansible-playbook -vvvv -i hosts infrastructure-
playbooks/switch-from-non-containerized-to-containerized-ceph-daemons.yml
```

5. Once the playbook completes successfully, edit the value of **rgw_multisite: true** in the **all.yml** file and ensure the value of **containerized_deployment** is **true**.

2.5. THE UPGRADE PROCESS

As a storage administrator, you use Ansible playbooks to upgrade an Red Hat Ceph Storage 4 storage cluster to Red Hat Ceph Storage 5. The **rolling_update.yml** Ansible playbook performs upgrades for deployments of Red Hat Ceph Storage. The **ceph-ansible** upgrades the Ceph nodes in the following order:

- Ceph Monitor
- Ceph Manager
- Ceph OSD nodes
- MDS nodes
- Ceph Object Gateway (RGW) nodes
- Ceph RBD-mirror node
- Ceph NFS nodes
- Ceph iSCSI gateway node
- Ceph client nodes
- Ceph-crash daemons

- Node-exporter on all nodes
- Ceph Dashboard



NOTE

Red Hat Ceph Storage 5 supports only containerized deployments.

ceph-ansible is currently not supported with Red Hat Ceph Storage 5. This means that once you have migrated your storage cluster to Red Hat Ceph Storage 5, you must use **cephadm** to perform subsequent updates.



NOTE

Red Hat Ceph Storage 5 also includes a health check function that returns a `DAEMON_OLD_VERSION` warning if it detects that any of the daemons in the storage cluster are running multiple versions of Red Hat Ceph Storage. The warning is triggered when the daemons continue to run multiple versions of Red Hat Ceph Storage beyond the time value set in the **mon_warn_older_version_delay** option. By default, the **mon_warn_older_version_delay** option is set to one week. This setting allows most upgrades to proceed without falsely seeing the warning. If the upgrade process is paused for an extended time period, you can mute the health warning:

```
ceph health mute DAEMON_OLD_VERSION --sticky
```

After the upgrade has finished, unmute the health warning:

```
ceph health unmute DAEMON_OLD_VERSION
```

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to all hosts in the storage cluster.
- A valid customer subscription.
- Root-level access to the Ansible administration node.
- The latest versions of Ansible and **ceph-ansible** available with Red Hat Ceph Storage 5.
- The **ansible** user account for use with the Ansible application.
- The nodes of the storage cluster is upgraded to Red Hat Enterprise Linux 8.4 EUS or later.



IMPORTANT

The Ansible inventory file must be present in the **ceph-ansible** directory.

Procedure

1. Enable the Ceph and Ansible repositories on the Ansible administration node:

Red Hat Enterprise Linux 8

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-8-x86_64-rpms --
enable=ansible-2.9-for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
subscription-manager repos --enable=rhceph-5-tools-for-rhel-9-x86_64-rpms
```

- On the Ansible administration node, ensure that the latest versions of the **ansible** and **ceph-ansible** packages are installed.

Syntax

```
dnf update ansible ceph-ansible
```

- Navigate to the **/usr/share/ceph-ansible/** directory:

Example

```
[root@admin ~]# cd /usr/share/ceph-ansible
```

- If upgrading from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, make copies of the **group_vars/osds.yml.sample** and **group_vars/clients.yml.sample** files, and rename them to **group_vars/osds.yml**, and **group_vars/clients.yml** respectively.

Example

```
[root@admin ceph-ansible]# cp group_vars/osds.yml.sample group_vars/osds.yml
[root@admin ceph-ansible]# cp group_vars/mdss.yml.sample group_vars/mdss.yml
[root@admin ceph-ansible]# cp group_vars/rgws.yml.sample group_vars/rgws.yml
[root@admin ceph-ansible]# cp group_vars/clients.yml.sample group_vars/clients.yml
```

- If upgrading from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, edit the **group_vars/all.yml** file to add Red Hat Ceph Storage 5 details.
- Once you have done the above two steps, copy the settings from the old **yaml** files to the new **yaml** files. Do not change the values of **ceph_rhcs_version**, **ceph_docker_image**, and **grafana_container_image** as the values for these configuration parameters are for Red Hat Ceph Storage 5. This ensures that all the settings related to your cluster are present in the current **yaml** file.

Example

```
fetch_directory: ~/ceph-ansible-keys
monitor_interface: eth0
public_network: 192.168.0.0/24
ceph_docker_registry_auth: true
ceph_docker_registry_username: SERVICE_ACCOUNT_USER_NAME
ceph_docker_registry_password: TOKEN
dashboard_admin_user:
dashboard_admin_password:
grafana_admin_user:
grafana_admin_password:
radosgw_interface: eth0
```

```
ceph_docker_image: "rhceph/rhceph-5-rhel8"
ceph_docker_image_tag: "latest"
ceph_docker_registry: "registry.redhat.io"
node_exporter_container_image: registry.redhat.io/openshift4/ose-prometheus-node-exporter:v4.6
grafana_container_image: registry.redhat.io/rhceph/rhceph-5-dashboard-rhel8:5
prometheus_container_image: registry.redhat.io/openshift4/ose-prometheus:v4.6
alertmanager_container_image: registry.redhat.io/openshift4/ose-prometheus-alertmanager:v4.6
```



NOTE

Ensure the Red Hat Ceph Storage 5 container images are set to the default values.

7. Edit the **group_vars/osds.yml** file. Add and set the following options:

Syntax

```
nb_retry_wait_osd_up: 50
delay_wait_osd_up: 30
```

8. Open the **group_vars/all.yml** file and verify the following values are present from the old **all.yml** file.

- a. The **fetch_directory** option is set with the same value from the old **all.yml** file:

Syntax

```
fetch_directory: FULL_DIRECTORY_PATH
```

Replace *FULL_DIRECTORY_PATH* with a writable location, such as the Ansible user's home directory.

- b. If the cluster you want to upgrade contains any Ceph Object Gateway nodes, add the **radosgw_interface** option:

```
radosgw_interface: INTERFACE
```

Replace *INTERFACE* with the interface to which the Ceph Object Gateway nodes listen.

- c. If your current setup has SSL certificates configured, edit the following:

Syntax

```
radosgw_frontend_ssl_certificate: /etc/pki/ca-trust/extracted/CERTIFICATE_NAME
radosgw_frontend_port: 443
```

- d. Uncomment the **upgrade_ceph_packages** option and set it to **True**:

Syntax

```
upgrade_ceph_packages: True
```

- e. If the storage cluster has more than one rgw instance per node, then uncomment the **radosgw_num_instances** setting and set it to the number of instances per node in the cluster:

Syntax

```
radosgw_num_instances : NUMBER_OF_INSTANCES_PER_NODE
```

Example

```
radosgw_num_instances : 2
```

- f. If the storage cluster has RGW multisite defined, check the multisite settings in **all.yml** to make sure that they contain the same values as they did in the old **all.yml** file.
9. Log in as **ansible-user** on the Ansible administration node.
 10. Execute the **rolling_update.yml** playbook to convert the storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5:


Syntax

```
ansible-playbook -vvvv infrastructure-playbooks/rolling_update.yml -i INVENTORY_FILE
```

The **-vvvv** option collects verbose logs of the upgrade process.

Example

```
[ansible@admin ceph-ansible]$ ansible-playbook -vvvv infrastructure-  
playbooks/rolling_update.yml -i hosts
```



IMPORTANT

Using the **--limit** Ansible option with the **rolling_update.yml** playbook is not supported.

11. Review the Ansible playbook log output to verify the status of the upgrade.

Verification

1. List all running containers:

Example

```
[root@mon ~]# podman ps
```

2. Check the health status of the cluster. Replace *MONITOR_ID* with the name of the Ceph Monitor container found in the previous step:

Syntax

```
podman exec ceph-mon-MONITOR_ID ceph -s
```

Example

```
[root@mon ~]# podman exec ceph-mon-mon01 ceph -s
```

3. Verify the Ceph cluster daemon versions to confirm the upgrade of all daemons. Replace *MONITOR_ID* with the name of the Ceph Monitor container found in the previous step:

Syntax

```
podman exec ceph-mon-MONITOR_ID ceph --cluster ceph versions
```

Example

```
[root@mon ~]# podman exec ceph-mon-mon01 ceph --cluster ceph versions
```

2.6. CONVERTING THE STORAGE CLUSTER TO USING **CEPHADM**

After you have upgraded the storage cluster to Red Hat Ceph Storage 5, run the **cephadm-adopt** playbook to convert the storage cluster daemons to run **cephadm**.

The **cephadm-adopt** playbook adopts the Ceph services, installs all **cephadm** dependencies, enables the **cephadm** Orchestrator backend, generates and configures the **ssh** key on all hosts, and adds the hosts to the Orchestrator configuration.



NOTE

After you run the **cephadm-adopt** playbook, remove the **ceph-ansible** package. The cluster daemons no longer work with **ceph-ansible**. You must use **cephadm** to manage the cluster daemons.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to all nodes in the storage cluster.

Procedure

1. Log in to the **ceph-ansible** node and change directory to **/usr/share/ceph-ansible**.
2. Run the **cephadm-adopt** playbook:

Syntax

```
ansible-playbook infrastructure-playbooks/cephadm-adopt.yml -i INVENTORY_FILE
```

Example

```
[ansible@admin ceph-ansible]$ ansible-playbook infrastructure-playbooks/cephadm-adopt.yml -i hosts
```

- Run the following command to enable applications to run on the NFS-Ganesha pool. *POOL_NAME* is **nfs-ganesha**, and *APPLICATION_NAME* is the name of the application you want to enable, such as **cephfs**, **rbd**, or **rgw**.

Syntax

```
ceph osd pool application enable POOL_NAME APPLICATION_NAME
```

Example

```
[root@node0 ~]# ceph osd pool application enable nfs-ganesha rgw
```



IMPORTANT

The **cephadm-adopt** playbook does not bring up rbd-mirroring after migrating the storage cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5.

To work around this issue, add the peers manually:

Syntax

```
rbd mirror pool peer add POOL_NAME CLIENT_NAME@CLUSTER_NAME
```

Example

```
[ceph: root@node0 /]# rbd --cluster site-a mirror pool peer add image-pool  
client.rbd-mirror-peer@site-b
```

Additional Resources

- For more information about using **leapp** to upgrade Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8, see [Upgrading from Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8](#).
- For more information about using **leapp** to upgrade Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9, see [Upgrading from Red Hat Enterprise Linux 8 to Red Hat Enterprise Linux 9](#).
- For more information about converting from FileStore to BlueStore, refer to [BlueStore](#).
- For more information about storage peers, see [Viewing information about peers](#).

2.7. INSTALLING CEPHADM-ANSIBLE ON AN UPGRADED STORAGE CLUSTER

cephadm-ansible is a collection of Ansible playbooks to simplify workflows that are not covered by **cephadm**. After installation, the playbooks are located in **/usr/share/cephadm-ansible/**.



NOTE

Before adding new nodes or new clients to your upgraded storage cluster, run the **cephadm-preflight.yml** playbook.

Prerequisites

- Root-level access to the Ansible administration node.
- A valid Red Hat subscription with the appropriate entitlements.
- An active Red Hat Network (RHN) or service account to access the Red Hat Registry.

Procedure

- Install the **cephadm-ansible** package:

Syntax

```
dnf install cephadm-ansible
```

Additional Resources

- [Running the preflight playbook](#)
- [Adding hosts](#)
- [Adding Monitor service](#)
- [Adding Manager service](#)
- [Adding OSDs](#)
- For more information about configuring clients and services, see [Red Hat Ceph Storage Operations Guide](#).
- For more information about the **cephadm-ansible** playbooks, see [The cephadm-ansible playbooks](#).

CHAPTER 3. UPGRADE A RED HAT CEPH STORAGE CLUSTER USING `CEPHADM`

As a storage administrator, you can use the **cephadm** Orchestrator to upgrade Red Hat Ceph Storage 5.0 and later.

The automated upgrade process follows Ceph best practices. For example:

- The upgrade order starts with Ceph Managers, Ceph Monitors, then other daemons.
- Each daemon is restarted only after Ceph indicates that the cluster will remain available.

The storage cluster health status is likely to switch to **HEALTH_WARNING** during the upgrade. When the upgrade is complete, the health status should switch back to **HEALTH_OK**.



WARNING

Upgrading to Red Hat Ceph Storage 5.2 on Ceph Object Gateway storage clusters (single-site or multi-site) is supported but you must set the **ceph config set mgr mgr/cephadm/no_five_one_rgw true --force** option prior to upgrading your storage cluster.

For more information, see the knowledge base article [Support Restrictions for upgrades for RADOS Gateway \(RGW\) on Red Hat Red Hat Ceph Storage 5.2](#) and [Known issues](#) section in the Red Hat Ceph Storage 5.2 Release Notes.



NOTE

ceph-ansible is currently not supported with Red Hat Ceph Storage 5. This means that once you have migrated your storage cluster to Red Hat Ceph Storage 5, you must use **cephadm** and **cephadm-ansible** to perform subsequent updates.



NOTE

You do not get a message once the upgrade is successful. Run **ceph versions** and **ceph orch ps** commands to verify the new image ID and the version of the storage cluster.

3.1. UPGRADING THE RED HAT CEPH STORAGE CLUSTER

You can use **ceph orch upgrade** command for upgrading a Red Hat Ceph Storage 5.0 cluster.

Prerequisites

- A running Red Hat Ceph Storage cluster 5.
- Red Hat Enterprise Linux 8.4 EUS or later.
- Root-level access to all the nodes.

- Ansible user with sudo and passwordless **ssh** access to all nodes in the storage cluster.
- At least two Ceph Manager nodes in the storage cluster: one active and one standby.

NOTE

Red Hat Ceph Storage 5 also includes a health check function that returns a `DAEMON_OLD_VERSION` warning if it detects that any of the daemons in the storage cluster are running multiple versions of RHCS. The warning is triggered when the daemons continue to run multiple versions of Red Hat Ceph Storage beyond the time value set in the `mon_warn_older_version_delay` option. By default, the `mon_warn_older_version_delay` option is set to 1 week. This setting allows most upgrades to proceed without falsely seeing the warning. If the upgrade process is paused for an extended time period, you can mute the health warning:

```
ceph health mute DAEMON_OLD_VERSION --sticky
```

After the upgrade has finished, unmute the health warning:

```
ceph health unmute DAEMON_OLD_VERSION
```

Procedure

1. Update the **cephadm** and **cephadm-ansible** package:

Example

```
[root@admin ~]# dnf update cephadm
[root@admin ~]# dnf update cephadm-ansible
```

2. Navigate to the `/usr/share/cephadm-ansible/` directory:

Example

```
[root@admin ~]# cd /usr/share/cephadm-ansible
```

3. Run the preflight playbook with the `upgrade_ceph_packages` parameter set to **true** on the bootstrapped host in the storage cluster:

Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars "ceph_origin=rhcs
upgrade_ceph_packages=true"
```

Example

```
[ansible@admin cephdm-ansible]$ ansible-playbook -i /etc/ansible/hosts cephadm-
preflight.yml --extra-vars "ceph_origin=rhcs upgrade_ceph_packages=true"
```

This package upgrades **cephadm** on all the nodes.

4. Log into the **cephadm** shell:

Example

```
[root@host01 ~]# cephadm shell
```

5. Ensure all the hosts are online and that the storage cluster is healthy:

Example

```
[ceph: root@host01 /]# ceph -s
```

6. Set the OSD **noout**, **noscrub**, and **nodeep-scrub** flags to prevent OSDs from getting marked out during upgrade and to avoid unnecessary load on the cluster:

Example

```
[ceph: root@host01 /]# ceph osd set noout
[ceph: root@host01 /]# ceph osd set noscrub
[ceph: root@host01 /]# ceph osd set nodeep-scrub
```

7. Check service versions and the available target containers:

Syntax

```
ceph orch upgrade check IMAGE_NAME
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade check registry.redhat.io/rhceph/rhceph-5-
rhel8:latest
```



NOTE

The image name is applicable for both Red Hat Enterprise Linux 8 and Red Hat Enterprise Linux 9.


8. Upgrade the storage cluster:

Syntax

```
ceph orch upgrade start IMAGE_NAME
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade start registry.redhat.io/rhceph/rhceph-5-
rhel8:latest
```



NOTE

To perform a staggered upgrade, see [Performing a staggered upgrade](#).

While the upgrade is underway, a progress bar appears in the **ceph status** output.

Example

```
[ceph: root@host01 /]# ceph status
[...]
progress:
  Upgrade to 16.2.0-146.el8cp (1s)
  [.....]
```

9. Verify the new *IMAGE_ID* and *VERSION* of the Ceph cluster:

Example

```
[ceph: root@host01 /]# ceph versions
[ceph: root@host01 /]# ceph orch ps
```



NOTE

If you are not using the **cephadm-ansible** playbooks, after upgrading your Ceph cluster, you must upgrade the **ceph-common** package and client libraries on your client nodes.

Example

```
[root@client01 ~] dnf update ceph-common
```

Verify you have the latest version:

Example

```
[root@client01 ~] ceph --version
```

10. When the upgrade is complete, unset the **noout**, **noscrub**, and **nodeep-scrub** flags:

Example

```
[ceph: root@host01 /]# ceph osd unset noout
[ceph: root@host01 /]# ceph osd unset noscrub
[ceph: root@host01 /]# ceph osd unset nodeep-scrub
```

3.2. UPGRADING THE RED HAT CEPH STORAGE CLUSTER IN A DISCONNECTED ENVIRONMENT

You can upgrade the storage cluster in a disconnected environment by using the **--image** tag.

You can use **ceph orch upgrade** command for upgrading a Red Hat Ceph Storage 5 cluster.

Prerequisites

- A running Red Hat Ceph Storage cluster 5.

- Red Hat Enterprise Linux 8.4 EUS or later.
- Root-level access to all the nodes.
- Ansible user with sudo and passwordless **ssh** access to all nodes in the storage cluster.
- At least two Ceph Manager nodes in the storage cluster: one active and one standby.
- Register the nodes to CDN and attach subscriptions.
- Check for the customer container images in a disconnected environment and change the configuration, if required. See the [Changing configurations of custom container images for disconnected installations](#) section in the *Red Hat Ceph Storage Installation Guide* for more details.

By default, the monitoring stack components are deployed based on the primary Ceph image. For disconnected environment of the storage cluster, you have to use the latest available monitoring stack component images.

Table 3.1. Custom image details for monitoring stack

Monitoring stack component	Red Hat Ceph Storage version	Image details
Prometheus	All Red Hat Ceph Storage 5 versions	registry.redhat.io/openshift4/ose-prometheus:v4.6
Grafana	All Red Hat Ceph Storage 5 versions	registry.redhat.io/rhceph/rhceph-5-dashboard-rhel8:latest
Node-exporter	Red Hat Ceph Storage 5.0	registry.redhat.io/openshift4/ose-prometheus-node-exporter:v4.5
	Red Hat Ceph Storage 5.0z1 onwards	registry.redhat.io/openshift4/ose-prometheus-node-exporter:v4.6
AlertManager	Red Hat Ceph Storage 5.0	registry.redhat.io/openshift4/ose-prometheus-alertmanager:v4.5
	Red Hat Ceph Storage 5.0z1 onwards	registry.redhat.io/openshift4/ose-prometheus-alertmanager:v4.6
HAProxy	Red Hat Ceph Storage 5.1	registry.redhat.io/rhceph/rhceph-haproxy-rhel8:latest
Keepalived	Red Hat Ceph Storage 5.1	registry.redhat.io/rhceph/keepalived-rhel8:latest
SNMP Gateway	All Red Hat Ceph Storage 5 versions	registry.redhat.io/rhceph/snmp-notifier-rhel8:latest

Procedure

1. Update the **cephadm** and **cephadm-ansible** package.

Example

```
[root@admin ~] dnf update cephadm
[root@admin ~] dnf update cephadm-ansible
```

2. Run the preflight playbook with the **upgrade_ceph_packages** parameter set to **true** and the **ceph_origin** parameter set to **custom** on the bootstrapped host in the storage cluster:

Syntax

```
ansible-playbook -i INVENTORY_FILE cephadm-preflight.yml --extra-vars
"ceph_origin=custom upgrade_ceph_packages=true"
```

Example

```
[ansible@admin ~]$ ansible-playbook -i /etc/ansible/hosts cephadm-preflight.yml --extra-vars
"ceph_origin=custom upgrade_ceph_packages=true"
```

This package upgrades **cephadm** on all the nodes.

3. Log into the **cephadm** shell:

Example

```
[root@host01 ~]# cephadm shell
```

4. Ensure all the hosts are online and that the storage cluster is healthy:

Example

```
[ceph: root@host01 /]# ceph -s
```

5. Set the OSD **noout**, **noscrub**, and **nodeep-scrub** flags to prevent OSDs from getting marked out during upgrade and to avoid unnecessary load on the cluster:

Example

```
[ceph: root@host01 /]# ceph osd set noout
[ceph: root@host01 /]# ceph osd set noscrub
[ceph: root@host01 /]# ceph osd set nodeep-scrub
```

6. Check service versions and the available target containers:

Syntax

```
ceph orch upgrade check IMAGE_NAME
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade check  
LOCAL_NODE_FQDN:5000/rhceph/rhceph-5-rhel8
```

7. Upgrade the storage cluster:

Syntax

```
ceph orch upgrade start IMAGE_NAME
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade start LOCAL_NODE_FQDN:5000/rhceph/rhceph-5-rhel8
```

While the upgrade is underway, a progress bar appears in the **ceph status** output.

Example

```
[ceph: root@host01 /]# ceph status  
[...]  
progress:  
  Upgrade to 16.2.0-115.el8cp (1s)  
  [.....]
```

8. Verify the new *IMAGE_ID* and *VERSION* of the Ceph cluster:

Example

```
[ceph: root@host01 /]# ceph version  
[ceph: root@host01 /]# ceph versions  
[ceph: root@host01 /]# ceph orch ps
```

9. When the upgrade is complete, unset the **noout**, **noscrub**, and **nodeep-scrub** flags:

Example

```
[ceph: root@host01 /]# ceph osd unset noout  
[ceph: root@host01 /]# ceph osd unset noscrub  
[ceph: root@host01 /]# ceph osd unset nodeep-scrub
```

Additional Resources

- See the [Registering Red Hat Ceph Storage nodes to the CDN and attaching subscriptions](#) section in the *Red Hat Ceph Storage Installation Guide*.
- See the [Configuring a private registry for a disconnected installation](#) section in the *Red Hat Ceph Storage Installation Guide*.

3.3. STAGGERED UPGRADE

As a storage administrator, you can upgrade Red Hat Ceph Storage components in phases rather than all at once. Starting with Red Hat Ceph Storage 5.2, the **ceph orch upgrade** command enables you to specify options to limit which daemons are upgraded by a single upgrade command.



NOTE

If you want to upgrade from a version that does not support staggered upgrades, you must first manually upgrade the Ceph Manager (**ceph-mgr**) daemons. For more information on performing a staggered upgrade from previous releases, see [Performing a staggered upgrade from previous releases](#).

3.3.1. Staggered upgrade options

Starting with Red Hat Ceph Storage 5.2, the **ceph orch upgrade** command supports several options to upgrade cluster components in phases. The staggered upgrade options include:

- **--daemon_types**: The **--daemon_types** option takes a comma-separated list of daemon types and will only upgrade daemons of those types. Valid daemon types for this option include **mgr**, **mon**, **crash**, **osd**, **mds**, **rgw**, **rbd-mirror**, **cephfs-mirror**, **iscsi**, and **nfs**.
- **--services**: The **--services** option is mutually exclusive with **--daemon-types**, only takes services of one type at a time, and will only upgrade daemons belonging to those services. For example, you cannot provide an OSD and RGW service simultaneously.
- **--hosts**: You can combine the **--hosts** option with **--daemon_types**, **--services**, or use it on its own. The **--hosts** option parameter follows the same format as the command line options for orchestrator CLI placement specification.
- **--limit**: The **--limit** option takes an integer greater than zero and provides a numerical limit on the number of daemons **cephadm** will upgrade. You can combine the **--limit** option with **--daemon_types**, **--services**, or **--hosts**. For example, if you specify to upgrade daemons of type **osd** on **host01** with a limit set to **3**, **cephadm** will upgrade up to three OSD daemons on host01.

3.3.2. Performing a staggered upgrade

As a storage administrator, you can use the **ceph orch upgrade** options to limit which daemons are upgraded by a single upgrade command.

Cephadm strictly enforces an order for the upgrade of daemons that is still present in staggered upgrade scenarios. The current upgrade order is:

- Ceph Manager nodes
- Ceph Monitor nodes
- Ceph-crash daemons
- Ceph OSD nodes
- Ceph Metadata Server (MDS) nodes
- Ceph Object Gateway (RGW) nodes
- Ceph RBD-mirror node
- CephFS-mirror node

- Ceph iSCSI gateway node
- Ceph NFS nodes



NOTE

If you specify parameters that upgrade daemons out of order, the upgrade command blocks and notes which daemons you need to upgrade before you proceed.

Example

```
[ceph: root@host01 /]# ceph orch upgrade start --image
registry.redhat.io/rhceph/rhceph-5-rhel8:latest --hosts host02
```

Error EINVAL: Cannot start upgrade. Daemons with types earlier in upgrade order than daemons on given host need upgrading.

Please first upgrade mon.ceph-host01

NOTE: Enforced upgrade order is: mgr -> mon -> crash -> osd -> mds -> rgw -> rbd-mirror -> cephfs-mirror -> iscsi -> nfs

Prerequisites

- A cluster running Red Hat Ceph Storage 5.2 or later.
- Root-level access to all the nodes.
- At least two Ceph Manager nodes in the storage cluster: one active and one standby.

Procedure

1. Log into the **cephadm** shell:

Example

```
[root@host01 ~]# cephadm shell
```

2. Ensure all the hosts are online and that the storage cluster is healthy:

Example

```
[ceph: root@host01 /]# ceph -s
```

3. Set the OSD **noout**, **noscrub**, and **nodeep-scrub** flags to prevent OSDs from getting marked out during upgrade and to avoid unnecessary load on the cluster:

Example

```
[ceph: root@host01 /]# ceph osd set noout
[ceph: root@host01 /]# ceph osd set noscrub
[ceph: root@host01 /]# ceph osd set nodeep-scrub
```

4. Check service versions and the available target containers:

Syntax

```
ceph orch upgrade check IMAGE_NAME
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade check registry.redhat.io/rhceph/rhceph-5-rhel8:latest
```

5. Upgrade the storage cluster:

- a. To upgrade specific daemon types on specific hosts:

Syntax

```
ceph orch upgrade start --image IMAGE_NAME --daemon-types  
DAEMON_TYPE1,DAEMON_TYPE2 --hosts HOST1,HOST2
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade start --image registry.redhat.io/rhceph/rhceph-5-rhel8:latest --daemon-types mgr,mon --hosts host02,host03
```

- b. To specify specific services and limit the number of daemons to upgrade:

Syntax

```
ceph orch upgrade start --image IMAGE_NAME --services SERVICE1,SERVICE2 --limit  
LIMIT_NUMBER
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade start --image registry.redhat.io/rhceph/rhceph-5-rhel8:latest --services rgw.example1,rgw1.example2 --limit 2
```



NOTE

In staggered upgrade scenarios, if using a limiting parameter, the monitoring stack daemons, including Prometheus and **node-exporter**, are refreshed after the upgrade of the Ceph Manager daemons. As a result of the limiting parameter, Ceph Manager upgrades take longer to complete. The versions of monitoring stack daemons might not change between Ceph releases, in which case, they are only redeployed.



NOTE

Upgrade commands with limiting parameters validates the options before beginning the upgrade, which can require pulling the new container image. As a result, the **upgrade start** command might take a while to return when you provide limiting parameters.

6. To see which daemons you still need to upgrade, run the **ceph orch upgrade check** or **ceph versions** command:

Example

```
[ceph: root@host01 /]# ceph orch upgrade check --image registry.redhat.io/rhceph/rhceph-5-rhel8:latest
```

7. To complete the staggered upgrade, verify the upgrade of all remaining services:

Syntax

```
ceph orch upgrade start --image IMAGE_NAME
```

Example

```
[ceph: root@host01 /]# ceph orch upgrade start --image registry.redhat.io/rhceph/rhceph-5-rhel8:latest
```

8. When the upgrade is complete, unset the **noout**, **noscrub**, and **nodeep-scrub** flags:

Example

```
[ceph: root@host01 /]# ceph osd unset noout
[ceph: root@host01 /]# ceph osd unset noscrub
[ceph: root@host01 /]# ceph osd unset nodeep-scrub
```

Verification

- Verify the new *IMAGE_ID* and *VERSION* of the Ceph cluster:

Example

```
[ceph: root@host01 /]# ceph versions
[ceph: root@host01 /]# ceph orch ps
```

Additional Resources

- For more information about upgrading a Red Hat Ceph Storage cluster using **cephadm**, see [Upgrade a Red Hat Ceph Storage cluster using cephadm](#).

3.3.3. Performing a staggered upgrade from previous releases

Starting with Red Hat Ceph Storage 5.2, you can perform a staggered upgrade on your storage cluster by providing the necessary arguments. If you want to upgrade from a version that does not support staggered upgrades, you must first manually upgrade the Ceph Manager (**ceph-mgr**) daemons. Once you have upgraded the Ceph Manager daemons, you can pass the limiting parameters to complete the staggered upgrade.



IMPORTANT

Verify you have at least two running Ceph Manager daemons before attempting this procedure.

Prerequisites

- A cluster running Red Hat Ceph Storage 5.0 or later.
- At least two Ceph Manager nodes in the storage cluster: one active and one standby.

Procedure

1. Log into the Cephadm shell:

Example

```
[root@host01 ~]# cephadm shell
```

2. Determine which Ceph Manager is active and which are standby:

Example

```
[ceph: root@host01 /]# ceph -s
cluster:
  id: 266ee7a8-2a05-11eb-b846-5254002d4916
  health: HEALTH_OK

services:
  mon: 2 daemons, quorum host01,host02 (age 92s)
  mgr: host01.ndtpjh(active, since 16h), standbys: host02.pzgrhz
```

3. Manually upgrade each standby Ceph Manager daemon:

Syntax

```
ceph orch daemon redeploy mgr.ceph-HOST.MANAGER_ID --image IMAGE_ID
```

Example

```
[ceph: root@host01 /]# ceph orch daemon redeploy mgr.ceph-host02.pzgrhz --image
registry.redhat.io/rhceph/rhceph-5-rhel8:latest
```

4. Fail over to the upgraded standby Ceph Manager:

Example

```
[ceph: root@host01 /]# ceph mgr fail
```

5. Check that the standby Ceph Manager is now active:

Example

```
[ceph: root@host01 /]# ceph -s
cluster:
  id: 266ee7a8-2a05-11eb-b846-5254002d4916
  health: HEALTH_OK

services:
  mon: 2 daemons, quorum host01,host02 (age 1h)
  mgr: host02.pzgrhz(active, since 25s), standbys: host01.ndtpjh
```

6. Verify that the active Ceph Manager is upgraded to the new version:

Syntax

```
ceph tell mgr.ceph-HOST.MANAGER_ID version
```

Example

```
[ceph: root@host01 /]# ceph tell mgr.host02.pzgrhz version
{
  "version": "16.2.8-12.el8cp",
  "release": "pacific",
  "release_type": "stable"
}
```

7. Repeat steps 2 - 6 to upgrade the remaining Ceph Managers to the new version.
8. Check that all Ceph Managers are upgraded to the new version:

Example

```
[ceph: root@host01 /]# ceph mgr versions
{
  "ceph version 16.2.8-12.el8cp (600e227816517e2da53d85f2fab3cd40a7483372) pacific
  (stable)": 2
}
```

9. Once you upgrade all your Ceph Managers, you can specify the limiting parameters and complete the remainder of the staggered upgrade.

Additional Resources

- For more information about performing a staggered upgrade and staggered upgrade options, see [Performing a staggered upgrade](#).

3.4. MONITORING AND MANAGING UPGRADE OF THE STORAGE CLUSTER

After running the **ceph orch upgrade start** command to upgrade the Red Hat Ceph Storage cluster, you can check the status, pause, resume, or stop the upgrade process. The health of the cluster changes to **HEALTH_WARNING** during an upgrade. If the host of the cluster is offline, the upgrade is paused.

**NOTE**

You have to upgrade one daemon type after the other. If a daemon cannot be upgraded, the upgrade is paused.

Prerequisites

- A running Red Hat Ceph Storage cluster 5.
- Root-level access to all the nodes.
- At least two Ceph Manager nodes in the storage cluster: one active and one standby.
- Upgrade for the storage cluster initiated.

Procedure

1. Determine whether an upgrade is in process and the version to which the cluster is upgrading:

Example

```
[ceph: root@node0 /]# ceph orch upgrade status
```

**NOTE**

You do not get a message once the upgrade is successful. Run **ceph versions** and **ceph orch ps** commands to verify the new image ID and the version of the storage cluster.

2. Optional: Pause the upgrade process:

Example

```
[ceph: root@node0 /]# ceph orch upgrade pause
```

3. Optional: Resume a paused upgrade process:

Example

```
[ceph: root@node0 /]# ceph orch upgrade resume
```

4. Optional: Stop the upgrade process:

Example

```
[ceph: root@node0 /]# ceph orch upgrade stop
```

3.5. TROUBLESHOOTING UPGRADE ERROR MESSAGES

The following table shows some **cephadm** upgrade error messages. If the **cephadm** upgrade fails for any reason, an error message appears in the storage cluster health status.

Error Message	Description
UPGRADE_NO_STANDBY_MGR	Ceph requires both active and standby manager daemons to proceed, but there is currently no standby.
UPGRADE_FAILED_PULL	Ceph was unable to pull the container image for the target version. This can happen if you specify a version or container image that does not exist (e.g., 1.2.3), or if the container registry is not reachable from one or more hosts in the cluster.