



Red Hat Ceph Storage 2

Installation Guide for Red Hat Enterprise Linux

Installing Red Hat Ceph Storage on Red Hat Enterprise Linux

Red Hat Ceph Storage 2 Installation Guide for Red Hat Enterprise Linux

Installing Red Hat Ceph Storage on Red Hat Enterprise Linux

Legal Notice

Copyright © 2018 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, 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 Software Collections 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 installing Red Hat Ceph Storage on Red Hat Enterprise Linux 7 running on AMD64 and Intel 64 architectures.

Table of Contents

CHAPTER 1. WHAT IS RED HAT CEPH STORAGE?	4
CHAPTER 2. PREREQUISITES	6
2.1. OPERATING SYSTEM	7
2.1.1. Adjusting the PID Count	7
2.2. REGISTERING RED HAT CEPH STORAGE NODES TO CDN AND ATTACHING SUBSCRIPTIONS	8
2.3. ENABLING THE RED HAT CEPH STORAGE REPOSITORIES	9
2.3.1. Content Delivery Network (CDN)	9
2.3.2. Local Repository	10
2.4. CONFIGURING RAID CONTROLLERS	11
2.5. CONFIGURING NETWORK	11
2.6. CONFIGURING FIREWALL	12
2.7. CONFIGURING NETWORK TIME PROTOCOL	14
2.8. CREATING AN ANSIBLE USER WITH SUDO ACCESS	15
2.9. ENABLING PASSWORD-LESS SSH (ANSIBLE DEPLOYMENT ONLY)	15
CHAPTER 3. STORAGE CLUSTER INSTALLATION	17
3.1. INSTALLING RED HAT CEPH STORAGE USING ANSIBLE	17
3.1.1. Installing ceph-ansible	17
3.1.2. Configuring Ceph Global Settings	18
3.1.3. Configuring Monitor Settings	21
3.1.4. Configuring Ceph OSD Settings	21
3.1.5. Overriding Ceph Default Settings	23
3.1.6. Deploying a Ceph Cluster	24
3.1.7. Taking over an Existing Cluster	24
3.1.8. Purging a Ceph Cluster	25
3.2. INSTALLING RED HAT CEPH STORAGE BY USING THE COMMAND-LINE INTERFACE	26
3.2.1. Monitor Bootstrapping	27
3.2.2. OSD Bootstrapping	33
3.2.3. Calamari Server Installation	38
CHAPTER 4. CLIENT INSTALLATION	40
4.1. INSTALLING THE CEPH-CLIENT ROLE	40
4.2. CEPH COMMAND-LINE INTERFACE INSTALLATION	41
4.3. CEPH BLOCK DEVICE INSTALLATION	42
Before you start	42
Installing Ceph Block Devices by Using the Command Line	42
4.4. CEPH OBJECT GATEWAY INSTALLATION	44
4.4.1. Installing Ceph Object Gateway by using Ansible	45
4.4.2. Installing Ceph Object Gateway Manually	46
CHAPTER 5. UPGRADING CEPH STORAGE CLUSTER	49
5.1. UPGRADING FROM RED HAT CEPH STORAGE 1.3 TO 2	49
5.1.1. Upgrading a Ceph Monitor Node	50
5.1.2. Upgrading a Ceph OSD Node	52
5.1.3. Upgrading the Ceph Object Gateway Nodes	55
Before You Start	55
Procedure: Upgrading the Ceph Object Gateway Node	55
See Also	56
5.1.4. Upgrading a Ceph Client Node	56
5.2. UPGRADING BETWEEN MINOR VERSIONS AND APPLYING ASYNCHRONOUS UPDATES	57
5.2.1. Changes Between Ansible 2.1 and 2.2	59

CHAPTER 6. WHAT TO DO NEXT?	61
APPENDIX A. TROUBLESHOOTING	62
A.1. ANSIBLE STOPS INSTALLATION BECAUSE IT DETECTS LESS DEVICES THAN IT EXPECTED	62

CHAPTER 1. WHAT IS RED HAT CEPH STORAGE?

Red Hat Ceph Storage is a scalable, open, software-defined storage platform that combines the most stable version of the Ceph storage system with a Ceph management platform, deployment utilities, and support services.

Red Hat Ceph Storage is designed for cloud infrastructure and web-scale object storage. Red Hat Ceph Storage clusters consist of the following types of nodes:

Red Hat Storage Ansible Administration node

This type of node acts as the traditional Ceph Administration node did for previous versions of Red Hat Ceph Storage. This type of node provides the following functions:

- Centralized storage cluster management
 - [Ansible administration](#)
 - [Ceph Client Command line interface](#)
- The Ceph configuration files and keys
- Optionally, local repositories for installing Ceph on nodes that cannot access the Internet for security reasons



NOTE

In Red Hat Ceph Storage 1.3.x, the Ceph Administration node hosted the Calamari monitoring and administration server, and the `ceph-deploy` utility, which has been deprecated in Red Hat Ceph Storage 2. Use Ceph command-line utility or the Ansible automation application utility instead to install a Red Hat Ceph Storage cluster.

Monitor nodes

Each monitor node runs the monitor daemon (`ceph-mon`), which maintains a master copy of the cluster map. The cluster map includes the cluster topology. A client connecting to the Ceph cluster retrieves the current copy of the cluster map from the monitor which enables the client to read from and write data to the cluster.

Ceph can run with one monitor; however, to ensure high availability in a production cluster, Red Hat recommends to deploy at least three monitor nodes.

OSD nodes

Each Object Storage Device (OSD) node runs the Ceph OSD daemon (`ceph-osd`), which interacts with logical disks attached to the node. Ceph stores data on these OSD nodes.

Ceph can run with very few OSD nodes, which the default is three, but production clusters realize better performance beginning at modest scales, for example 50 OSDs in a storage cluster. Ideally, a Ceph cluster has multiple OSD nodes, allowing isolated failure domains by creating the CRUSH map.

MDS nodes

Each Metadata Server (MDS) node runs the MDS daemon (`ceph-mds`), which manages metadata related to files stored on the Ceph File System (CephFS). The MDS daemon also coordinates access to the shared cluster.

MDS and CephFS are Technology Preview features and as such they are not fully supported yet. For information on MDS installation and configuration, see the [Ceph File System Guide \(Technology Preview\)](#).

Object Gateway node

Ceph Object Gateway node runs the Ceph RADOS Gateway daemon (`ceph-radosgw`), and is an object storage interface built on top of `librados` to provide applications with a RESTful gateway to Ceph Storage Clusters. The Ceph RADOS Gateway supports two interfaces:

S3

Provides object storage functionality with an interface that is compatible with a large subset of the Amazon S3 RESTful API.

Swift

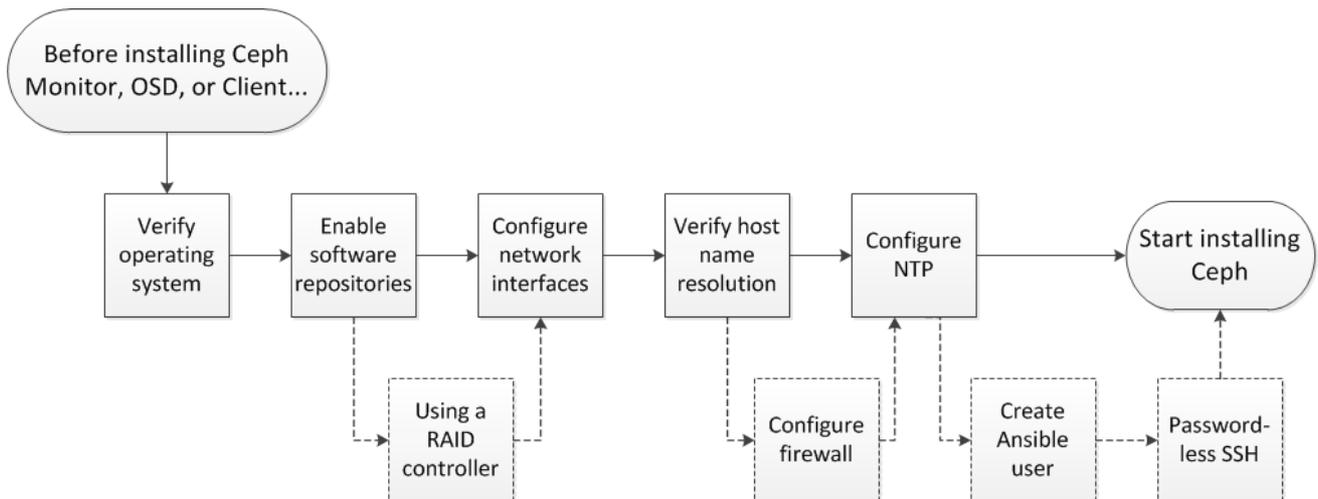
Provides object storage functionality with an interface that is compatible with a large subset of the OpenStack Swift API.

For details on the Ceph architecture, see the [Architecture Guide](#).

For minimum recommended hardware, see the [Hardware Guide](#).

CHAPTER 2. PREREQUISITES

Figure 2.1. Prerequisite Workflow



Before installing Red Hat Ceph Storage, review the following prerequisites first and prepare each Ceph Monitor, OSD, and client nodes accordingly.

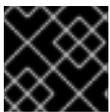
Table 2.1. Prerequisites Checks

Task	Required	Section	Recommendation
Verifying the operating system version	Yes	Section 2.1, “Operating System”	Verify the PID count
Registering Ceph nodes	Yes	Section 2.2, “Registering Red Hat Ceph Storage Nodes to CDN and Attaching Subscriptions”	
Enabling Ceph software repositories	Yes	Section 2.3, “Enabling the Red Hat Ceph Storage Repositories”	Two installation methods: <ul style="list-style-type: none"> Content Delivery Network (CDN) Local Repository (ISO)
Using a RAID controller	No	Section 2.4, “Configuring RAID Controllers”	For OSD nodes only.
Configuring network Interface	Yes	Section 2.5, “Configuring Network”	Using a public network is required. Having a private network for cluster communication is optional, but recommended.

Task	Required	Section	Recommendation
Configuring a firewall	No	Section 2.6, “Configuring Firewall”	
Configuring the Network Time Protocol	Yes	Note	
Creating an Ansible user	No	Section 2.8, “Creating an Ansible User with Sudo Access”	Ansible deployment only. Creating the Ansible user is required on all Ceph nodes.
Enabling password-less SSH	No	Section 2.9, “Enabling Password-less SSH (Ansible Deployment Only)”	Ansible deployment only.

2.1. OPERATING SYSTEM

Red Hat Ceph Storage 2 and later requires Red Hat Enterprise Linux 7 Server with a homogeneous version, for example, Red Hat Enterprise Linux 7.2 running on AMD64 and Intel 64 architectures for all Ceph nodes.



IMPORTANT

Red Hat does not support clusters with heterogeneous operating systems and versions.

[Return to prerequisite checklist](#)

2.1.1. Adjusting the PID Count

Hosts with high numbers of OSDs, that being more than 12, may spawn a lot of threads, especially during recovery and re-balancing events. The kernel defaults to a relatively small maximum number of threads, typically **32768**.

1. Check the current `pid_max` settings:

```
# cat /proc/sys/kernel/pid_max
```

2. As **root**, consider setting `kernel.pid_max` to a higher number of threads. The theoretical maximum is 4,194,303 threads. For example, add the following to the `/etc/sysctl.conf` file to set it to the maximum value:

```
kernel.pid_max = 4194303
```

3. As **root**, to load the changes without a rebooting:

```
# sysctl -p
```

4. As **root**, verify the changes:

```
# sysctl -a | grep kernel.pid_max
```

2.2. REGISTERING RED HAT CEPH STORAGE NODES TO CDN AND ATTACHING SUBSCRIPTIONS

Ceph relies on packages in the Red Hat Enterprise Linux 7 Base content set. Each Ceph node must be able to access the full Red Hat Enterprise Linux 7 Base content.

To do so, register Ceph nodes that can connect to the Internet to the Red Hat Content Delivery Network (CDN) and attach appropriate Ceph subscriptions to the nodes:

Registering Ceph Nodes to CDN

Run all commands in this procedure as **root**.

1. Register a node with the Red Hat Subscription Manager. Run the following command and when prompted, enter your Red Hat Customer Portal credentials:

```
# subscription-manager register
```

2. Pull the latest subscription data from the CDN server:

```
# subscription-manager refresh
```

3. List all available subscriptions and find the appropriate Red Hat Ceph Storage subscription and determine its Pool ID.

```
# subscription-manager list --available
```

4. Attach the subscriptions:

```
# subscription-manager attach --pool=<pool-id>
```

Replace **<pool-id>** with the Pool ID determined in the previous step.

5. Enable the Red Hat Enterprise Linux 7 Server Base repository:

```
# subscription-manager repos --enable=rhel-7-server-rpms
```

6. Enable the Red Hat Enterprise Linux 7 Server Extras repository:

```
# subscription-manager repos --enable=rhel-7-server-extras-rpms
```

7. Update the node:

```
# yum update
```

Once you register the nodes, enable repositories that provide the Red Hat Ceph Storage packages.

**NOTE**

For nodes that cannot access the Internet during the installation, provide the Base content by other means. Either use the Red Hat Satellite server in your environment or mount a local Red Hat Enterprise Linux 7 Server ISO image and point the Ceph cluster nodes to it. For additional details, contact the Red Hat Support.

For more information on registering Ceph nodes with the Red Hat Satellite server, see the [How to Register Ceph with Satellite 6](#) and [How to Register Ceph with Satellite 5](#) articles on the [Customer Portal](#).

[Return to prerequisite checklist](#)

2.3. ENABLING THE RED HAT CEPH STORAGE REPOSITORIES

Before you can install Red Hat Ceph Storage, you must choose an installation method. Red Hat Ceph Storage supports two installation methods:

- **Content Delivery Network (CDN)**
For Ceph Storage clusters with Ceph nodes that can connect directly to the Internet, use Red Hat Subscription Manager to enable the required Ceph repositories on each node.
- **Local Repository**
For Ceph Storage clusters where security measures preclude nodes from accessing the Internet, install Red Hat Ceph Storage 2 from a single software build delivered as an ISO image, which will allow you to install local repositories.

**IMPORTANT**

Some Ceph package dependencies require versions that differ from the package versions included in the Extra Packages for Enterprise Linux (EPEL) repository. Disable the EPEL repository to ensure that only the Red Hat Ceph Storage packages are installed.

As **root**, disable the EPEL repository:

```
# yum-config-manager --disable epel
```

This command disables the `epel.repo` file in `/etc/yum.repos.d/`.

2.3.1. Content Delivery Network (CDN)

CDN Installations for...

- **Ansible administration node**

As **root**, enable the Red Hat Ceph Storage 2 Tools repository:

```
# subscription-manager repos --enable=rhel-7-server-rhceph-2-tools-rpms
```

- **Monitor Nodes**

As **root**, enable the Red Hat Ceph Storage 2 Monitor repository:

```
# subscription-manager repos --enable=rhel-7-server-rhceph-2-mon-rpms
```

- **OSD Nodes**

As **root**, enable the Red Hat Ceph Storage 2 OSD repository:

```
# subscription-manager repos --enable=rhel-7-server-rhceph-2-osd-rpms
```

- **RADOS Gateway and Client Nodes**

As **root**, enable the Red Hat Ceph Storage 2 Tools repository:

```
# subscription-manager repos --enable=rhel-7-server-rhceph-2-tools-rpms
```

[Return to prerequisite checklist](#)

2.3.2. Local Repository

For ISO Installations...

- **Download the Red Hat Ceph Storage ISO**

1. Log in to the [Red Hat Customer Portal](#).
2. Click **Downloads** to visit the **Software & Download** center.
3. In the Red Hat Ceph Storage area, click **Download Software** to download the latest version of the software.
4. Copy the ISO image to the node.
5. As **root**, mount the copied ISO image to the `/mnt/rhcs2/` directory:

```
# mkdir -p /mnt/rhcs2
# mount -o loop /<path_to_iso>/rhceph-2.0-rhel-7-x86_64.iso /mnt/rhcs2
```



NOTE

For ISO installations using Ansible to install Red Hat Ceph Storage 2, mounting the ISO and creating a local repository is not required.

- **Create a Local Repository**

1. Copy the ISO image to the node.
2. Follow the steps in this Knowledgebase [solution](#).

**NOTE**

If you are completely disconnected from the Internet, then you must use ISO images to receive any updates.

[Return to prerequisite checklist](#)

2.4. CONFIGURING RAID CONTROLLERS

If a RAID controller with 1-2 GB of cache is installed on a host, enabling write-back caches might result in increased small I/O write throughput. To prevent this problem, the cache must be non-volatile.

Modern RAID controllers usually have super capacitors that provide enough power to drain volatile memory to non-volatile NAND memory during a power loss event. It is important to understand how a particular controller and firmware behave after power is restored.

Some RAID controllers require manual intervention. Hard drives typically advertise to the operating system whether their disk caches should be enabled or disabled by default. However, certain RAID controllers or some firmware do not provide such information, so verify that disk level caches are disabled to avoid file system corruption.

Create a single RAID 0 volume with write-back for each OSD data drive with write-back cache enabled.

If Serial Attached SCSI (SAS) or SATA connected Solid-state Drive (SSD) disks are also present on the controller, investigate whether your controller and firmware support **passthrough** mode.

Passthrough mode helps avoid caching logic, and generally results in much lower latency for fast media.

[Return to prerequisite checklist](#)

2.5. CONFIGURING NETWORK

All Ceph clusters require a public network. You must have a network interface card configured to a public network where Ceph clients can reach Ceph monitors and Ceph OSD nodes.

You might have a network interface card for a cluster network so that Ceph can conduct heart-beating, peering, replication, and recovery on a network separate from the public network.

**IMPORTANT**

Red Hat does not recommend using a single network interface card for both a public and private network.

Configure the network interfaces and ensure to make the changes persistent so that the settings are identical on reboot. Configure the following settings:

- The **BOOTPROTO** parameter is usually set to **none** for static IP addresses.
- The **ONBOOT** parameter must be set to **yes**. If it is set to **no**, Ceph might fail to peer on reboot.
- If you intend to use IPv6 addressing, the IPv6 parameters, for example **IPV6INIT** must be set to **yes** except for the **IPV6_FAILURE_FATAL** parameter. Also, edit the Ceph configuration file to instruct Ceph to use IPv6. Otherwise, Ceph will use IPv4.

Navigate to the `/etc/sysconfig/network-scripts/` directory and ensure that the `ifcfg-<iface>` settings for the public and cluster interfaces are properly configured.

For details on configuring network interface scripts for Red Hat Enterprise Linux 7, see the [Configuring a Network Interface Using ifcfg Files](#) chapter in the Networking Guide for Red Hat Enterprise Linux 7.

For additional information on network configuration see the [Network Configuration Reference](#) chapter in the [Configuration Guide](#) for Red Hat Ceph Storage 2.

[Return to prerequisite checklist](#)

2.6. CONFIGURING FIREWALL

Red Hat Ceph Storage 2 uses the `firewalld` service, which you must configure to suit your environment.

Monitor nodes use port `6789` for communication within the Ceph cluster. The monitor where the `calamari-lite` is running uses port `8002` for access to the Calamari REST-based API.

On each Ceph OSD node, the OSD daemon uses several ports in the range `6800-7300`:

- One for communicating with clients and monitors over the public network
- One for sending data to other OSDs over a cluster network, if available; otherwise, over the public network
- One for exchanging heartbeat packets over a cluster network, if available; otherwise, over the public network

Ceph object gateway nodes use port `7480` by default. However, you can change the default port, for example to port `80`. To use the SSL/TLS service, open port `443`.

For more information about public and cluster network, see [Network](#).

Configuring Access

1. On all Ceph nodes, as `root`, start the `firewalld` service, enable it to run on boot, and ensure that it is running:

```
# systemctl enable firewalld
# systemctl start firewalld
# systemctl status firewalld
```

2. As `root`, on all Ceph Monitor nodes, open port `6789` on the public network:

```
# firewall-cmd --zone=public --add-port=6789/tcp
# firewall-cmd --zone=public --add-port=6789/tcp --permanent
```

To limit access based on the source address, run the following commands:

```
# firewall-cmd --zone=public --add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" port protocol="tcp" \
port="6789" accept"
# firewall-cmd --zone=public --add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" port protocol="tcp" \
```

```
port="6789" accept" --permanent
```

3. If `calamari-lite` is running on the Ceph Monitor node, as `root`, open port **8002** on the public network:

```
# firewall-cmd --zone=public --add-port=8002/tcp
# firewall-cmd --zone=public --add-port=8002/tcp --permanent
```

To limit access based on the source address, run the following commands:

```
# firewall-cmd --zone=public --add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" port protocol="tcp" \
port="8002" accept"
# firewall-cmd --zone=public --add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" port protocol="tcp" \
port="8002" accept" --permanent
```

4. As `root`, on all OSD nodes, open ports **6800-7300**:

```
# firewall-cmd --zone=public --add-port=6800-7300/tcp
# firewall-cmd --zone=public --add-port=6800-7300/tcp --permanent
```

If you have a separate cluster network, repeat the commands with the appropriate zone.

5. As `root`, on all object gateway nodes, open the relevant port or ports on the public network.

- a. To open the default port **7480**:

```
# firewall-cmd --zone=public --add-port=7480/tcp
# firewall-cmd --zone=public --add-port=7480/tcp --permanent
```

To limit access based on the source address, run the following commands:

```
# firewall-cmd --zone=public \
--add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" \
port protocol="tcp" port="7480" accept"
# firewall-cmd --zone=public \
--add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" \
port protocol="tcp" port="7480" accept" --permanent
```

- b. Optionally, as `root`, if you changed the default Ceph object gateway port, for example to port **80**, open this port:

```
# firewall-cmd --zone=public --add-port=80/tcp
# firewall-cmd --zone=public --add-port=80/tcp --permanent
```

To limit access based on the source address, run the following commands:

```
# firewall-cmd --zone=public \
--add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" \
```

```
port protocol="tcp" port="80" accept"
# firewall-cmd --zone=public \
--add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" \
port protocol="tcp" port="80" accept" --permanent
```

- c. Optionally, as **root**, to use SSL/TLS, open port **443**:

```
# firewall-cmd --zone=public --add-port=443/tcp
# firewall-cmd --zone=public --add-port=443/tcp --permanent
```

To limit access based on the source address, run the following commands:

```
# firewall-cmd --zone=public \
--add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" \
port protocol="tcp" port="443" accept"
# firewall-cmd --zone=public \
--add-rich-rule="rule family="ipv4" \
source address="<IP-address>/<prefix>" \
port protocol="tcp" port="443" accept" --permanent
```

For additional details on `firewalld`, see the [Using Firewalls](#) chapter in the [Security Guide](#) for Red Hat Enterprise Linux 7.

[Return to prerequisite checklist](#)

2.7. CONFIGURING NETWORK TIME PROTOCOL



NOTE

If using Ansible to deploy a Red Hat Ceph Storage cluster, then the installation, configuration, and enabling NTP is done automatically during the deployment.

You must configure the Network Time Protocol (NTP) on all Ceph Monitor and OSD nodes. Ensure that Ceph nodes are NTP peers. NTP helps preempt issues that arise from clock drift.

1. As **root**, install the `ntp` package:

```
# yum install ntp
```

2. As **root**, enable the NTP service to be persistent across a reboot:

```
# systemctl enable ntpd
```

3. As **root**, start the NTP service and ensure it is running:

```
# systemctl start ntpd
# systemctl status ntpd
```

4. Ensure that NTP is synchronizing Ceph monitor node clocks properly:

```
$ ntpq -p
```

For additional details on NTP for Red Hat Enterprise Linux 7, see the [Configuring NTP Using ntpd](#) chapter in the [System Administrator's Guide](#) for Red Hat Enterprise Linux 7.

[Return to prerequisite checklist](#)

2.8. CREATING AN ANSIBLE USER WITH SUDO ACCESS

Ansible must login to Ceph nodes as a user that has passwordless `root` privileges, because Ansible needs to install software and configuration files without prompting for passwords.

Red Hat recommends creating an Ansible user on all Ceph nodes in the cluster.



IMPORTANT

Do not use `ceph` as the user name. The `ceph` user name is reserved for the Ceph daemons.

A uniform user name across the cluster can improve ease of use, but avoid using obvious user names, because intruders typically use them to for brute force attacks. For example, `root`, `admin`, or `<productname>` are not advised.

The following procedure, substituting `<username>` for the user name you define, describes how to create an Ansible user with passwordless `root` privileges on a Ceph node.

1. Use the `ssh` command to log in to a Ceph node:

```
$ ssh <user_name>@<hostname>
```

Replace `<hostname>` with the host name of the Ceph node.

2. Create a new Ansible user and set a new password for this user:

```
# adduser <username>
# passwd <username>
```

3. Ensure that the user you added has the `root` privileges:

```
# cat << EOF >/etc/sudoers.d/<username>
<username> ALL = (root) NOPASSWD:ALL
EOF
```

4. Ensure the correct file permissions:

```
# chmod 0440 /etc/sudoers.d/<username>
```

[Return to prerequisite checklist](#)

2.9. ENABLING PASSWORD-LESS SSH (ANSIBLE DEPLOYMENT ONLY)

Since Ansible will not prompt for a password, you must generate SSH keys on the administration node and distribute the public key to each Ceph node.

1. Generate the SSH keys, but do not use `sudo` or the `root` user. Instead, use the Ansible user you created in [Creating an Ansible User with Sudo Access](#). Leave the passphrase empty:

```
$ ssh-keygen

Generating public/private key pair.
Enter file in which to save the key (/ceph-admin/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /ceph-admin/.ssh/id_rsa.
Your public key has been saved in /ceph-admin/.ssh/id_rsa.pub.
```

2. Copy the key to each Ceph Node, replacing `<username>` with the user name you created in [Creating an Ansible User with Sudo Access](#) and `<hostname>` with a host name of a Ceph node:

```
$ ssh-copy-id <username>@<hostname>
```

3. Modify or create (using a utility such as `vi`) the `~/.ssh/config` file of the Ansible administration node so that Ansible can log in to Ceph nodes as the user you created without requiring you to specify the `-u <username>` option each time you execute the `ansible-playbook` command. Replace `<username>` with the name of the user you created and `<hostname>` with a host name of a Ceph node:

```
Host node1
  Hostname <hostname>
  User <username>
Host node2
  Hostname <hostname>
  User <username>
Host node3
  Hostname <hostname>
  User <username>
```

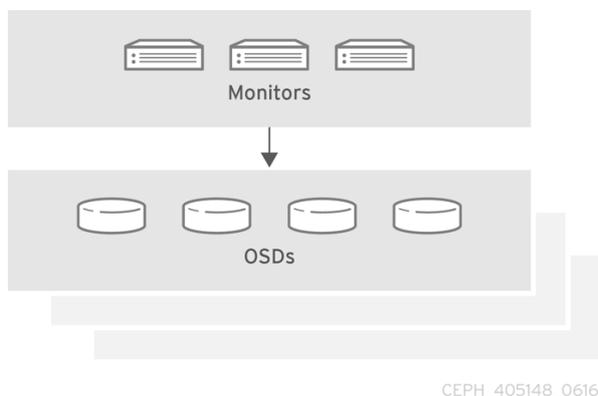
After editing the `~/.ssh/config` file on the Ansible administration node, ensure the permissions are correct:

```
$ chmod 600 ~/.ssh/config
```

[Return to prerequisite checklist](#)

CHAPTER 3. STORAGE CLUSTER INSTALLATION

Production Ceph storage clusters start with a minimum of three monitor hosts and three OSD nodes containing multiple OSDs.



You can install a Red Hat Ceph Storage cluster by using:

- [Ansible automation application](#)
- [Command-line interface](#)

3.1. INSTALLING RED HAT CEPH STORAGE USING ANSIBLE

You can use the Ansible automation application to deploy Red Hat Ceph Storage. Execute the procedures in [Figure 2.1, “Prerequisite Workflow”](#) first.

To add more Monitors or OSDs to an existing storage cluster, see the Red Hat Ceph Storage Administration Guide for details:

- [Adding a Monitor](#)
- [Adding an OSD](#)

3.1.1. Installing ceph-ansible

1. Install the `ceph-ansible` package:

```
# yum install ceph-ansible
```

2. As `root`, add the Ceph hosts to the `/etc/ansible/hosts` file. Remember to comment out example hosts.

If the Ceph hosts have sequential naming, consider using a range:

- a. Add Monitor nodes under the `[mons]` section:

```
[mons]
<monitor-host-name>
<monitor-host-name>
<monitor-host-name>
```

- b. Add OSD nodes under the `[osds]` section:

```
[osds]
<osd-host-name[1:10]>
```

Optionally, use the **devices** parameter to specify devices that the OSD nodes will use. Use a comma-separated list to list multiple devices.

```
[osds]
<ceph-host-name> devices="[ '<device_1>', '<device_2>' ]"
```

For example:

```
[osds]
ceph-osd-01 devices="[ '/dev/sdb', '/dev/sdc' ]"
ceph-osd-02 devices="[ '/dev/sdb', '/dev/sdc', '/dev/sdd' ]"
```

When specifying no devices, then set the **osd_auto_discovery** option to **true** in the `osds.yml` file. See [Section 3.1.4, “Configuring Ceph OSD Settings”](#) for more details.

Using the **devices** parameter is useful when OSDs use devices with different names or when one of the devices failed on one of the OSDs. See [Section A.1, “Ansible Stops Installation Because It Detects Less Devices Than It Expected”](#) for more details.

3. As the Ansible user, ensure that Ansible can reach the Ceph hosts:

```
$ ansible all -m ping
```



NOTE

See [Section 2.8, “Creating an Ansible User with Sudo Access”](#) for more details on creating an Ansible user.

3.1.2. Configuring Ceph Global Settings

1. Create a directory under the home directory so Ansible can write the keys:

```
# cd ~
# mkdir ceph-ansible-keys
```

2. As **root**, create a symbolic link to the Ansible **group_vars** directory in the `/etc/ansible/` directory:

```
# ln -s /usr/share/ceph-ansible/group_vars /etc/ansible/group_vars
```

3. As **root**, create an `all.yml` file from the `all.yml.sample` file and open it for editing:

```
# cd /etc/ansible/group_vars
# cp all.yml.sample all.yml
# vim all.yml
```

4. Uncomment the **fetch_directory** setting under the **GENERAL** section. Then, point it to the directory you created in step 1:

```
fetch_directory: ~/ceph-ansible-keys
```

5. Uncomment the `ceph_repository_type` setting and set it to either `cdn` or `iso`:

```
ceph_repository_type: cdn
```

6. Select the installation method. There are two approaches:

- a. If Ceph hosts have connectivity to the Red Hat Content Delivery Network (CDN), uncomment and set the following:

```
ceph_origin: repository
ceph_repository: rhcs
ceph_repository_type: cdn
ceph_rhcs_version: 2
```

- b. If Ceph nodes cannot connect to the Red Hat Content Delivery Network (CDN), uncomment the `ceph_repository_type` setting and set it to `iso`. This approach is most frequently used in high security environments.

```
ceph_repository_type: iso
```

Then, uncomment the `ceph_rhcs_iso_path` setting and specify the path to the ISO image:

```
ceph_rhcs_iso_path: <path>
```

Example

```
ceph_rhcs_iso_path: /path/to/ISO_file.iso
```

7. Set the `generate_fsid` setting to `false`:

```
generate_fsid: false
```



NOTE

With `generate_fsid` set to `false`, then you must specify a value for the File System Identifier (FSID) manually. For example, using the command-line utility, `uuidgen`, you can generate a Universally Unique Identifier (UUID). Once you generate a UUID, then uncomment the `fsid` setting and specify the generated UUID:

```
fsid: <generated_uuid>
```

With `generate_fsid` set to `true`, then the UUID will be automatically generated. This removes the need to specify the UUID in the `fsid` setting.

8. To enable authentication, uncomment the `cephx` setting under the Ceph Configuration section. Red Hat recommends running Ceph with authentication enabled:

```
cephx: true
```

9. Uncomment the `monitor_interface` setting and specify the network interface:

```
monitor_interface:
```

Example

```
monitor_interface: eth0
```



NOTE

The `monitor_interface` setting will use the IPv4 address. To use an IPv6 address, use the `monitor_address` setting instead.

10. If not using IPv6, then skip this step. Uncomment and set the `ip_version` option:

```
ip_version: ipv6
```

11. Set journal size:

```
journal_size: <size_in_MB>
```

If not filled, the default journal size will be 5 GB. See [Journal Settings](#) for additional details.

12. Set the public network. Optionally, set the cluster network, too:

```
public_network: <public_network>  
cluster_network: <cluster_network>
```

See [Section 2.5, “Configuring Network”](#) and [Network Configuration Reference](#) for additional details.

13. If not using IPv6, then skip this step. Uncomment and set the `radosgw_civetweb_bind_ip` option:

```
radosgw_civetweb_bind_ip: "[&#123;&#123;  
ansible_default_ipv6.address &#125;&#125;]"
```



IMPORTANT

Currently, there is a rendering bug when displaying content within double curly brackets on the Customer Portal. The Customer Portal team is working diligently to resolve this issue.

The HTML escape codes being displayed in the above step represent the left ({) and right (}) curly brackets respectively. For example, written in long hand, the `radosgw_civetweb_bind_ip` option would be the following:

```
radosgw_civetweb_bind_ip: "[<left_curly_bracket>
<left_curly_bracket> ansible_default_ipv6.address
<right_curly_bracket><right_curly_bracket>]"
```

3.1.3. Configuring Monitor Settings

Ansible will create Ceph Monitors without any additional configuration steps. However, you may override default settings for authentication, and for use with OpenStack. By default, the Calamari API is disabled.

To configure monitors, perform the following:

1. Navigate to the `/etc/ansible/group_vars/` directory:

```
# cd /etc/ansible/group_vars/
```

2. As **root**, create an `mons.yml` file from `mons.yml.sample` file and open it for editing:

```
# cp mons.yml.sample mons.yml
# vim mons.yml
```

3. To enable the Calamari API, uncomment the `calamari` setting and set it to **true**:

```
calamari: true
```

4. To configure other settings, uncomment them and set appropriate values.

3.1.4. Configuring Ceph OSD Settings

To configure OSDs:

1. Navigate to the `/etc/ansible/group_vars/` directory:

```
# cd /etc/ansible/group_vars/
```

2. As **root**, create a new `osds.yml` file from the `osds.yml.sample` file and open it for editing:

```
# cp osds.yml.sample osds.yml
# vim osds.yml
```

3. Uncomment and set settings that are relevant for your use case. See [Table 3.1, “What settings are needed for my use case?”](#) for details.

- Once you are done editing the file, save your changes and close the file.

Table 3.1. What settings are needed for my use case?

I want:	Relevant Options	Comments
to have the Ceph journal and OSD data co-located on the same device and to specify OSD disks on my own.	devices journal_collocation: true	The devices setting accepts a list of devices. Ensure that the specified devices correspond to the storage devices on the OSD nodes.
to have the Ceph journal and OSD data co-located on the same device and ceph-ansible to detect and configure all the available devices.	osd_auto_discovery: true journal_collocation: true	
to use one or more dedicated devices to store the Ceph journal.	devices raw_multi_journal: true raw_journal_devices	The devices and raw_journal_devices settings except a list of devices. Ensure that the devices specified correspond to the storage devices on the OSD nodes.
to use directories instead of disks.	osd_directory: true osd_directories	The osd_directories setting accepts a list of directories. IMPORTANT: Red Hat currently does not support this scenario.
to have the Ceph journal and OSD data co-located on the same device and encrypt OSD data.	devices dmccrypt_journal_collocation: true	The devices setting accepts a list of devices. Ensure that the specified devices correspond to the storage devices on the OSD nodes. Note that all OSDs will be encrypted. For details, see the Encryption chapter in the Red Hat Ceph Storage 2 Architecture Guide.
to use one or more dedicated devices to store the Ceph journal and encrypt OSD data.	devices dmccrypt_dedicated_journal: true raw_journal_devices	The devices and raw_journal_devices settings except a list of devices. Ensure that the devices specified correspond to the storage devices on the OSD nodes. Note that all OSDs will be encrypted. For details, see the Encryption chapter in the Red Hat Ceph Storage 2 Architecture Guide.

I want:	Relevant Options	Comments
to use the BlueStore back end instead of the FileStore back end.	devices bluestore: true	The devices setting accepts a list of devices. For details on BlueStore, see the OSD BlueStore (Technology Preview) chapter in the Administration Guide for Red Hat Ceph Storage.

For additional settings, see the `osds.yml.sample` file located in `/usr/share/ceph-ansible/group_vars/`.



WARNING

Some OSD options will conflict with each other. Avoid enabling these sets of options together:

- `journal_collocation` and `raw_multi_journal`
- `journal_collocation` and `osd_directory`
- `raw_multi_journal` and `osd_directory`

In addition, specifying one of these options is required.

3.1.5. Overriding Ceph Default Settings

Unless otherwise specified in the Ansible configuration files, Ceph uses its default settings.

Because Ansible manages the Ceph configuration file, edit the `/etc/ansible/group_vars/all.yml` file to change the Ceph configuration. Use the `ceph_conf_overrides` setting to override the default Ceph configuration.

Ansible supports the same sections as the Ceph configuration file; `[global]`, `[mon]`, `[osd]`, `[mds]`, `[rgw]`, and so on. You can also override particular instances, such as a particular Ceph Object Gateway instance. For example:

```
#####
# CONFIG OVERRIDE #
#####

ceph_conf_overrides:
  client.rgw.rgw1:
    log_file: /var/log/ceph/ceph-rgw-rgw1.log
```

**NOTE**

Ansible does not include braces when referring to a particular section of the Ceph configuration file. Sections and settings names are terminated with a colon.

**IMPORTANT**

Do not set the cluster network with the `cluster_network` parameter in the **CONFIG OVERRIDE** section because this can cause two conflicting cluster networks being set in the Ceph configuration file.

To set the cluster network, use the `cluster_network` parameter in the **CEPH CONFIGURATION** section. For details, see [Configuring Ceph Global Settings](#).

3.1.6. Deploying a Ceph Cluster

1. Navigate to the Ansible configuration directory:

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

2. As **root**, create a `site.yml` file from the `site.yml.sample` file:

```
# cp site.yml.sample site.yml
```

3. As the Ansible user, run the Ansible playbook from within the directory where the playbook exists:

```
$ ansible-playbook site.yml [-u <user_name>]
```

Once the playbook runs, it creates a running Ceph cluster.

**NOTE**

During the deployment of a Red Hat Ceph Storage cluster with Ansible, the installation, configuration, and enabling NTP is done automatically on each node in the storage cluster.

4. As **root**, on the Ceph Monitor nodes, create a Calamari user:

Syntax

```
# calamari-ctl add_user --password <password> --email  
<email_address> <user_name>
```

Example

```
# calamari-ctl add_user --password abc123 --email user1@example.com  
user1
```

3.1.7. Taking over an Existing Cluster

You can configure Ansible to use a cluster deployed without Ansible. For example, if you upgraded Red Hat Ceph Storage 1.3 clusters to version 2 manually, configure them to use Ansible by following this procedure:

1. After manually upgrading from version 1.3 to version 2, install and configure Ansible on the administration node. This is the node where the master Ceph configuration file is maintained. See [Section 3.2.1, "Installing Ceph Ansible"](#) for details.
2. Ensure that the Ansible administration node has passwordless `ssh` access to all Ceph nodes in the cluster. See [Section 2.9, "Enabling Password-less SSH \(Ansible Deployment Only\)"](#) for more details.
3. As `root`, create a symbolic link to the Ansible `group_vars` directory in the `/etc/ansible/` directory:

```
# ln -s /usr/share/ceph-ansible/group_vars /etc/ansible/group_vars
```

4. As `root`, create an `all.yml` file from the `all.yml.sample` file and open it for editing:

```
# cd /etc/ansible/group_vars
# cp all.yml.sample all.yml
# vim all.yml
```

5. Set the `generate_fsid` setting to `false` in `group_vars/all.yml`.
6. Get the current cluster `fsid` by executing `ceph fsid`.
7. Set the retrieved `fsid` in `group_vars/all.yml`.
8. Modify the Ansible inventory in `/etc/ansible/hosts` to include Ceph hosts. Add monitors under a `[mons]` section, OSDs under an `[osds]` section and gateways under an `[rgws]` section to identify their roles to Ansible.
9. Make sure `ceph_conf_overrides` is updated with the original `ceph.conf` options used for `[global]`, `[osd]`, `[mon]`, and `[client]` sections in the `all.yml` file. Options like `osd_journal`, `public_network` and `cluster_network` should not be added in `ceph_conf_overrides` because they are already part of `all.yml`. Only the options that are not part of `all.yml` and are in the original `ceph.conf` should be added to `ceph_conf_overrides`.
10. From the `/usr/share/ceph-ansible/` directory run the playbook.

```
# cd /usr/share/ceph-ansible/
# cp infrastructure-playbooks/take-over-existing-cluster.yml .
$ ansible-playbook take-over-existing-cluster.yml -u <username>
```

3.1.8. Purging a Ceph Cluster

If you deployed a Ceph cluster using Ansible and you want to purge the cluster, then use the `purge-cluster.yml` Ansible playbook located in the `infrastructure-playbooks` directory.



IMPORTANT

Purging a Ceph cluster will lose data stored on the cluster's OSDs.

Before purging the Ceph cluster...

Check the `osd_auto_discovery` option in the `osds.yml` file. Having this option set to `true` will cause the purge to fail. To prevent the failure, do the following steps before running the purge:

1. Declare the OSD devices in the `osds.yml` file. See [Section 3.1.4, “Configuring Ceph OSD Settings”](#) for more details.
2. Comment out the `osd_auto_discovery` option in the `osds.yml` file.

To purge the Ceph cluster...

1. As **root**, navigate to the `/usr/share/ceph-ansible/` directory:

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

2. As **root**, copy the `purge-cluster.yml` Ansible playbook to the current directory:

```
# cp infrastructure-playbooks/purge-cluster.yml .
```

3. Run the `purge-cluster.yml` Ansible playbook:

```
$ ansible-playbook purge-cluster.yml
```

3.2. INSTALLING RED HAT CEPH STORAGE BY USING THE COMMAND-LINE INTERFACE

All Ceph clusters require at least one monitor, and at least as many OSDs as copies of an object stored on the cluster. Red Hat recommends using three monitors for production environments and a minimum of three Object Storage Devices (OSD).

Bootstrapping the initial monitor is the first step in deploying a Ceph storage cluster. Ceph monitor deployment also sets important criteria for the entire cluster, such as:

- The number of replicas for pools
- The number of placement groups per OSD
- The heartbeat intervals
- Any authentication requirement

Most of these values are set by default, so it is useful to know about them when setting up the cluster for production.

Installing a Ceph storage cluster by using the command line interface involves these steps:

- Bootstrapping the initial [Monitor node](#)
- Adding an Object Storage Device ([OSD node](#))



IMPORTANT

Red Hat does not support or test upgrading manually deployed clusters. Currently, the only supported way to upgrade to a minor version of Red Hat Ceph Storage 2 is to use the Ansible automation application as described in [Important](#). Therefore, Red Hat recommends to use Ansible to deploy a new cluster with Red Hat Ceph Storage 2. See [Section 3.1, “Installing Red Hat Ceph Storage using Ansible”](#) for details.

You can use command-line utilities, such as Yum, to upgrade manually deployed clusters, but Red Hat does not support or test this.

3.2.1. Monitor Bootstrapping

Bootstrapping a Monitor and by extension a Ceph storage cluster, requires the following data:

Unique Identifier

The File System Identifier (`fsid`) is a unique identifier for the cluster. The `fsid` was originally used when the Ceph storage cluster was principally used for the Ceph file system. Ceph now supports native interfaces, block devices, and object storage gateway interfaces too, so `fsid` is a bit of a misnomer.

Cluster Name

Ceph clusters have a cluster name, which is a simple string without spaces. The default cluster name is `ceph`, but you can specify a different cluster name. Overriding the default cluster name is especially useful when you work with multiple clusters.

When you run multiple clusters in a multi-site architecture, the cluster name for example, `us-west`, `us-east` identifies the cluster for the current command-line session.



NOTE

To identify the cluster name on the command-line interface, specify the Ceph configuration file with the cluster name, for example, `ceph.conf`, `us-west.conf`, `us-east.conf`, and so on.

Example:

```
# ceph --cluster us-west.conf ...
```

Monitor Name

Each Monitor instance within a cluster has a unique name. In common practice, the Ceph Monitor name is the node name. Red Hat recommend one Ceph Monitor per node, and no co-locating the Ceph OSD daemons with the Ceph Monitor daemon. To retrieve the short node name, use the `hostname -s` command.

Monitor Map

Bootstrapping the initial Monitor requires you to generate a Monitor map. The Monitor map requires:

- The File System Identifier (`fsid`)
- The cluster name, or the default cluster name of `ceph` is used
- At least one host name and its IP address.

Monitor Keyring

Monitors communicate with each other by using a secret key. You must generate a keyring with a Monitor secret key and provide it when bootstrapping the initial Monitor.

Administrator Keyring

To use the `ceph` command-line interface utilities, create the `client.admin` user and generate its keyring. Also, you must add the `client.admin` user to the Monitor keyring.

The foregoing requirements do not imply the creation of a Ceph configuration file. However, as a best practice, Red Hat recommends creating a Ceph configuration file and populating it with the `fsid`, the `mon initial members` and the `mon host` settings at a minimum.

You can get and set all of the Monitor settings at runtime as well. However, the Ceph configuration file might contain only those settings which overrides the default values. When you add settings to a Ceph configuration file, these settings override the default settings. Maintaining those settings in a Ceph configuration file makes it easier to maintain the cluster.

To bootstrap the initial Monitor, perform the following steps:

1. Enable the Red Hat Ceph Storage 2 [Monitor repository](#). For ISO-based installations, see the [ISO installation section](#).

2. On your initial Monitor node, install the `ceph-mon` package as `root`:

```
# yum install ceph-mon
```

3. As `root`, create a Ceph configuration file in the `/etc/ceph/` directory. By default, Ceph uses `ceph.conf`, where `ceph` reflects the cluster name:

Syntax

```
# touch /etc/ceph/<cluster_name>.conf
```

Example

```
# touch /etc/ceph/ceph.conf
```

4. As `root`, generate the unique identifier for your cluster and add the unique identifier to the `[global]` section of the Ceph configuration file:

Syntax

```
# echo "[global]" > /etc/ceph/<cluster_name>.conf
# echo "fsid = `uuidgen`" >> /etc/ceph/<cluster_name>.conf
```

Example

```
# echo "[global]" > /etc/ceph/ceph.conf
# echo "fsid = `uuidgen`" >> /etc/ceph/ceph.conf
```

5. View the current Ceph configuration file:

```
$ cat /etc/ceph/ceph.conf
[global]
fsid = a7f64266-0894-4f1e-a635-d0aeaca0e993
```

- As **root**, add the initial Monitor to the Ceph configuration file:

Syntax

```
# echo "mon initial members = <monitor_host_name>[,
<monitor_host_name>]" >> /etc/ceph/<cluster_name>.conf
```

Example

```
# echo "mon initial members = node1" >> /etc/ceph/ceph.conf
```

- As **root**, add the IP address of the initial Monitor to the Ceph configuration file:

Syntax

```
# echo "mon host = <ip-address>[,<ip-address>]" >>
/etc/ceph/<cluster_name>.conf
```

Example

```
# echo "mon host = 192.168.0.120" >> /etc/ceph/ceph.conf
```



NOTE

To use IPv6 addresses, you must set the `ms bind ipv6` option to `true`. See the Red Hat Ceph Storage [Configuration Guide](#) for more details.

- As **root**, create the keyring for the cluster and generate the Monitor secret key:

Syntax

```
# ceph-authtool --create-keyring /tmp/<cluster_name>.mon.keyring --
gen-key -n mon. --cap mon '<capabilities>'
```

Example

```
# ceph-authtool --create-keyring /tmp/ceph.mon.keyring --gen-key -n
mon. --cap mon 'allow *'
creating /tmp/ceph.mon.keyring
```

- As **root**, generate an administrator keyring, generate a `<cluster_name>.client.admin.keyring` user and add the user to the keyring:

Syntax

```
# ceph-authtool --create-keyring
```

```
/etc/ceph/<cluster_name>.client.admin.keyring --gen-key -n
client.admin --set-uid=0 --cap mon '<capabilities>' --cap osd
'<capabilities>' --cap mds '<capabilities>'
```

Example

```
# ceph-authtool --create-keyring /etc/ceph/ceph.client.admin.keyring
--gen-key -n client.admin --set-uid=0 --cap mon 'allow *' --cap osd
'allow *' --cap mds 'allow'
creating /etc/ceph/ceph.client.admin.keyring
```

10. As **root**, add the `<cluster_name>.client.admin.keyring` key to the `<cluster_name>.mon.keyring`:

Syntax

```
# ceph-authtool /tmp/<cluster_name>.mon.keyring --import-keyring
/etc/ceph/<cluster_name>.client.admin.keyring
```

Example

```
# ceph-authtool /tmp/ceph.mon.keyring --import-keyring
/etc/ceph/ceph.client.admin.keyring
importing contents of /etc/ceph/ceph.client.admin.keyring into
/tmp/ceph.mon.keyring
```

11. Generate the Monitor map. Specify using the node name, IP address and the `fsid`, of the initial Monitor and save it as `/tmp/monmap`:

Syntax

```
$ monmaptool --create --add <monitor_host_name> <ip-address> --fsid
<uuid> /tmp/monmap
```

Example

```
$ monmaptool --create --add node1 192.168.0.120 --fsid a7f64266-
0894-4f1e-a635-d0aeaca0e993 /tmp/monmap
monmaptool: monmap file /tmp/monmap
monmaptool: set fsid to a7f64266-0894-4f1e-a635-d0aeaca0e993
monmaptool: writing epoch 0 to /tmp/monmap (1 monitors)
```

12. As **root** on the initial Monitor node, create a default data directory:

Syntax

```
# mkdir /var/lib/ceph/mon/<cluster_name>-<monitor_host_name>
```

Example

```
# mkdir /var/lib/ceph/mon/ceph-node1
```

13. As **root**, populate the initial Monitor daemon with the Monitor map and keyring:

Syntax

```
# ceph-mon [--cluster <cluster_name>] --mkfs -i <monitor_host_name>
--monmap /tmp/monmap --keyring /tmp/<cluster_name>.mon.keyring
```

Example

```
# ceph-mon --mkfs -i node1 --monmap /tmp/monmap --keyring
/tmp/ceph.mon.keyring
ceph-mon: set fsid to a7f64266-0894-4f1e-a635-d0aeaca0e993
ceph-mon: created monfs at /var/lib/ceph/mon/ceph-node1 for
mon.node1
```

14. View the current Ceph configuration file:

```
# cat /etc/ceph/ceph.conf
[global]
fsid = a7f64266-0894-4f1e-a635-d0aeaca0e993
mon_initial_members = node1
mon_host = 192.168.0.120
```

For more details on the various Ceph configuration settings, see the [Red Hat Ceph Storage Configuration Guide](#). The following example of a Ceph configuration file lists some of the most common configuration settings:

Example

```
[global]
fsid = <cluster-id>
mon initial members = <monitor_host_name>[, <monitor_host_name>]
mon host = <ip-address>[, <ip-address>]
public network = <network>[, <network>]
cluster network = <network>[, <network>]
auth cluster required = cephx
auth service required = cephx
auth client required = cephx
osd journal size = <n>
filestore xattr use omap = true
osd pool default size = <n> # Write an object n times.
osd pool default min size = <n> # Allow writing n copy in a degraded
state.
osd pool default pg num = <n>
osd pool default pgp num = <n>
osd crush chooseleaf type = <n>
```

15. As **root**, create the **done** file:

Syntax

```
# touch /var/lib/ceph/mon/<cluster_name>-<monitor_host_name>/done
```

Example

```
# touch /var/lib/ceph/mon/ceph-node1/done
```

16. As **root**, update the owner and group permissions on the newly created directory and files:

Syntax

```
# chown -R <owner>:<group> <path_to_directory>
```

Example

```
# chown -R ceph:ceph /var/lib/ceph/mon
# chown -R ceph:ceph /var/log/ceph
# chown -R ceph:ceph /var/run/ceph
# chown ceph:ceph /etc/ceph/ceph.client.admin.keyring
# chown ceph:ceph /etc/ceph/ceph.conf
# chown ceph:ceph /etc/ceph/rbdmap
```

**NOTE**

If the Ceph Monitor node is co-located with an OpenStack Controller node, then the Glance and Cinder keyring files must be owned by **glance** and **cinder** respectively. For example:

```
# ls -l /etc/ceph/
...
-rw-----. 1 glance glance      64 <date>
ceph.client.glance.keyring
-rw-----. 1 cinder cinder      64 <date>
ceph.client.cinder.keyring
...
```

17. For storage clusters with custom names, as **root**, add the the following line:

Syntax

```
# echo "CLUSTER=<custom_cluster_name>" >> /etc/sysconfig/ceph
```

Example

```
# echo "CLUSTER=test123" >> /etc/sysconfig/ceph
```

18. As **root**, start and enable the **ceph-mon** process on the initial Monitor node:

Syntax

```
# systemctl enable ceph-mon.target
# systemctl enable ceph-mon@<monitor_host_name>
# systemctl start ceph-mon@<monitor_host_name>
```

Example

```
# systemctl enable ceph-mon.target
# systemctl enable ceph-mon@node1
# systemctl start ceph-mon@node1
```

19. Verify that Ceph created the default pools:

```
$ ceph osd lspools
0 rbd,
```

20. Verify that the Monitor is running. The status output will look similar to the following example. The Monitor is up and running, but the cluster health will be in a **HEALTH_ERR** state. This error is indicating that placement groups are stuck and inactive. Once OSDs are added to the cluster and active, the placement group health errors will disappear.

Example

```
$ ceph -s
cluster a7f64266-0894-4f1e-a635-d0aeaca0e993
health HEALTH_ERR 192 pgs stuck inactive; 192 pgs stuck unclean; no
osds
monmap e1: 1 mons at {node1=192.168.0.120:6789/0}, election epoch 1,
quorum 0 node1
osdmap e1: 0 osds: 0 up, 0 in
pgmap v2: 192 pgs, 3 pools, 0 bytes data, 0 objects
0 kB used, 0 kB / 0 kB avail
192 creating
```

To add more Red Hat Ceph Storage Monitors to the storage cluster, see the [Red Hat Ceph Storage Administration Guide](#)

3.2.2. OSD Bootstrapping

Once you have your initial monitor running, you can start adding the Object Storage Devices (OSDs). Your cluster cannot reach an **active + clean** state until you have enough OSDs to handle the number of copies of an object.

The default number of copies for an object is three. You will need three OSD nodes at minimum. However, if you only want two copies of an object, therefore only adding two OSD nodes, then update the **osd pool default size** and **osd pool default min size** settings in the Ceph configuration file.

For more details, see the [OSD Configuration Reference](#) section in the Red Hat Ceph Storage Configuration Guide.

After bootstrapping the initial monitor, the cluster has a default CRUSH map. However, the CRUSH map does not have any Ceph OSD daemons mapped to a Ceph node.

To add an OSD to the cluster and updating the default CRUSH map, execute the following on each OSD node:

1. Enable the Red Hat Ceph Storage 2 [OSD repository](#). For ISO-based installations, see the [ISO installation section](#).

2. As **root**, install the **ceph-osd** package on the Ceph OSD node:

```
# yum install ceph-osd
```

3. Copy the Ceph configuration file and administration keyring file from the initial Monitor node to the OSD node:

Syntax

```
# scp <user_name>@<monitor_host_name>:<path_on_remote_system>  
<path_to_local_file>
```

Example

```
# scp root@node1:/etc/ceph/ceph.conf /etc/ceph  
# scp root@node1:/etc/ceph/ceph.client.admin.keyring /etc/ceph
```

4. Generate the Universally Unique Identifier (UUID) for the OSD:

```
$ uuidgen  
b367c360-b364-4b1d-8fc6-09408a9cda7a
```

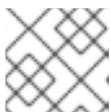
5. As **root**, create the OSD instance:

Syntax

```
# ceph osd create <uuid> [<osd_id>]
```

Example

```
# ceph osd create b367c360-b364-4b1d-8fc6-09408a9cda7a  
0
```



NOTE

This command outputs the OSD number identifier needed for subsequent steps.

6. As **root**, create the default directory for the new OSD:

Syntax

```
# mkdir /var/lib/ceph/osd/<cluster_name>-<osd_id>
```

Example

```
# mkdir /var/lib/ceph/osd/ceph-0
```

7. As **root**, prepare the drive for use as an OSD, and mount it to the directory you just created. Create a partition for the Ceph data and journal. The journal and the data partitions can be located on the same disk. This example is using a 15 GB disk:

Syntax

```
# parted <path_to_disk> mklabel gpt
# parted <path_to_disk> mkpart primary 1 10000
# mkfs -t <fstype> <path_to_partition>
# mount -o noatime <path_to_partition>
/var/lib/ceph/osd/<cluster_name>-<osd_id>
# echo "<path_to_partition> /var/lib/ceph/osd/<cluster_name>-
<osd_id> xfs defaults,noatime 1 2" >> /etc/fstab
```

Example

```
# parted /dev/sdb mklabel gpt
# parted /dev/sdb mkpart primary 1 10000
# parted /dev/sdb mkpart primary 10001 15000
# mkfs -t xfs /dev/sdb1
# mount -o noatime /dev/sdb1 /var/lib/ceph/osd/ceph-0
# echo "/dev/sdb1 /var/lib/ceph/osd/ceph-0 xfs defaults,noatime 1
2" >> /etc/fstab
```

- As root, initialize the OSD data directory:

Syntax

```
# ceph-osd -i <osd_id> --mkfs --mkkey --osd-uuid <uuid>
```

Example

```
# ceph-osd -i 0 --mkfs --mkkey --osd-uuid b367c360-b364-4b1d-8fc6-
09408a9cda7a
... auth: error reading file: /var/lib/ceph/osd/ceph-0/keyring:
can't open /var/lib/ceph/osd/ceph-0/keyring: (2) No such file or
directory
... created new key in keyring /var/lib/ceph/osd/ceph-0/keyring
```



NOTE

The directory must be empty before you run **ceph-osd** with the **--mkkey** option. If you have a custom cluster name, the **ceph-osd** utility requires the **--cluster** option.

- As root, register the OSD authentication key. If your cluster name differs from **ceph**, insert your cluster name instead:

Syntax

```
# ceph auth add osd.<osd_id> osd 'allow *' mon 'allow profile osd' -
i /var/lib/ceph/osd/<cluster_name>-<osd_id>/keyring
```

Example

```
# ceph auth add osd.0 osd 'allow *' mon 'allow profile osd' -i
```

```
/var/lib/ceph/osd/ceph-0/keyring
added key for osd.0
```

10. As **root**, add the OSD node to the CRUSH map:

Syntax

```
# ceph [--cluster <cluster_name>] osd crush add-bucket <host_name>
host
```

Example

```
# ceph osd crush add-bucket node2 host
```

11. As **root**, place the OSD node under the **default** CRUSH tree:

Syntax

```
# ceph [--cluster <cluster_name>] osd crush move <host_name>
root=default
```

Example

```
# ceph osd crush move node2 root=default
```

12. As **root**, add the OSD disk to the CRUSH map

Syntax

```
# ceph [--cluster <cluster_name>] osd crush add osd.<osd_id>
<weight> [<bucket_type>=<bucket-name> ...]
```

Example

```
# ceph osd crush add osd.0 1.0 host=node2
add item id 0 name 'osd.0' weight 1 at location {host=node2} to
crush map
```



NOTE

You can also decompile the CRUSH map, and add the OSD to the device list. Add the OSD node as a bucket, then add the device as an item in the OSD node, assign the OSD a weight, recompile the CRUSH map and set the CRUSH map. For more details, see the Red Hat Ceph Storage [Storage Strategies Guide](#) for more details.

13. As **root**, update the owner and group permissions on the newly created directory and files:

Syntax

```
# chown -R <owner>:<group> <path_to_directory>
```

Example

```
# chown -R ceph:ceph /var/lib/ceph/osd
# chown -R ceph:ceph /var/log/ceph
# chown -R ceph:ceph /var/run/ceph
# chown -R ceph:ceph /etc/ceph
```

14. For storage clusters with custom names, as **root**, add the following line to the `/etc/sysconfig/ceph` file:

Syntax

```
# echo "CLUSTER=<custom_cluster_name>" >> /etc/sysconfig/ceph
```

Example

```
# echo "CLUSTER=test123" >> /etc/sysconfig/ceph
```

15. The OSD node is in your Ceph storage cluster configuration. However, the OSD daemon is **down** and **in**. The new OSD must be **up** before it can begin receiving data. As **root**, enable and start the OSD process:

Syntax

```
# systemctl enable ceph-osd.target
# systemctl enable ceph-osd@<osd_id>
# systemctl start ceph-osd@<osd_id>
```

Example

```
# systemctl enable ceph-osd.target
# systemctl enable ceph-osd@0
# systemctl start ceph-osd@0
```

Once you start the OSD daemon, it is **up** and **in**.

Now you have the monitors and some OSDs up and running. You can watch the placement groups peer by executing the following command:

```
$ ceph -w
```

To view the OSD tree, execute the following command:

```
$ ceph osd tree
```

Example

```
ID  WEIGHT  TYPE NAME          UP/DOWN  REWEIGHT  PRIMARY-AFFINITY
-1   2       root default
-2   2       host node2
```

```

0          1          osd.0          up          1          1
-3         1          host node3
1          1          osd.1          up          1          1

```

To expand the storage capacity by adding new OSDs to the storage cluster, see the Red Hat Ceph Storage [Administration Guide](#) for more details.

3.2.3. Calamari Server Installation

The Calamari server provides a RESTful API for monitoring Ceph storage clusters.

To install `calamari-server`, perform the following steps on all Monitor nodes.

1. As **root**, enable the Red Hat Ceph Storage 2 [Monitor repository](#)
2. As **root**, install `calamari-server`:

```
# yum install calamari-server
```

3. As **root**, initialize the `calamari-server`:

Syntax

```
# calamari-ctl clear --yes-i-am-sure
# calamari-ctl initialize --admin-username <uid> --admin-password
<pwd> --admin-email <email>
```

Example

```
# calamari-ctl clear --yes-i-am-sure
# calamari-ctl initialize --admin-username admin --admin-password
admin --admin-email cephadm@example.com
```



IMPORTANT

The `calamari-ctl clear --yes-i-am-sure` command is only necessary for removing the database of old Calamari server installations. Running this command on a new Calamari server results in an error.



NOTE

During initialization, the `calamari-server` will generate a self-signed certificate and a private key and place them in the `/etc/calamari/ssl/certs/` and `/etc/calamari/ssl/private` directories respectively. Use HTTPS when making requests. Otherwise, user names and passwords are transmitted in clear text.

The `calamari-ctl initialize` process generates a private key and a self-signed certificate, which means there is no need to purchase a certificate from a Certificate Authority (CA).

To verify access to the HTTPS API through a web browser, go to the following URL. Click through the untrusted certificate warnings, because the auto-generated certificate is self-signed:

-

```
https://<calamari_hostname>:8002/api/v2/cluster
```

To use a key and certificate from a CA:

1. Purchase a certificate from a CA. During the process, you will generate a private key and a certificate for CA. Or you can also use the self-signed certificate generated by Calamari.
2. Save the private key associated to the certificate to a path, preferably under `/etc/calamari/ssl/private/`.
3. Save the certificate to a path, preferably under `/etc/calamari/ssl/certs/`.
4. Open the `/etc/calamari/calamari.conf` file.
5. Under the `[calamari_web]` section, modify `ssl_cert` and `ssl_key` to point to the respective certificate and key path, for example:

```
[calamari_web]
...
ssl_cert = /etc/calamari/ssl/certs/calamari-lite-bundled.crt
ssl_key = /etc/calamari/ssl/private/calamari-lite.key
```

6. As **root**, re-initialize Calamari:

```
# calamari-ctl initialize
```

CHAPTER 4. CLIENT INSTALLATION

Red Hat Ceph Storage supports the following types of Ceph clients:

Ceph CLI

The Ceph command-line interface (CLI) enables administrators to execute Ceph administrative commands. See [Section 4.2, “Ceph Command-line Interface Installation”](#) for information on installing the Ceph CLI.

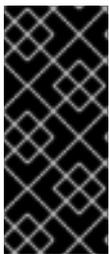
Block Device

The Ceph Block Device is a thin-provisioned, resizable block device. See [Section 4.3, “Ceph Block Device Installation”](#) for information on installing Ceph Block Devices.

Object Gateway

The Ceph Object Gateway provides its own user management and Swift- and S3-compliant APIs. See [Section 4.4, “Ceph Object Gateway Installation”](#) for information on installing Ceph Object Gateways.

In addition, the `ceph-ansible` utility provides the `ceph-client` role that copies the Ceph configuration file and the administration keyring to nodes. See [Section 4.1, “Installing the ceph-client role”](#) for details.



IMPORTANT

To use Ceph clients, you must have a Ceph cluster storage running, preferably in the `active + clean` state.

In addition, before installing the Ceph clients, ensure to perform the tasks listed in the [Figure 2.1, “Prerequisite Workflow”](#) section.

4.1. INSTALLING THE CEPH-CLIENT ROLE

The `ceph-client` role copies the Ceph configuration file and administration keyring to a node. In addition, you can use this role to create custom pools and clients.

Perform the following tasks on the Ansible administration node, see [Installing ceph-ansible](#) for details.

1. Add a new section `[clients]` to the `/etc/ansible/hosts` file:

```
[clients]
<client-hostname>
```

Replace `<client-hostname>` with the host name of the node where you want to install the `ceph-client` role.

2. Navigate to the `/etc/ansible/group_vars/` directory:

```
$ cd /etc/ansible/group_vars
```

3. Create a new copy of the `clients.yml.sample` file named `clients.yml`:

```
# cp clients.yml.sample clients.yml
```

4. Optionally, instruct `ceph-client` to create pools and clients.

a. Update `clients.yml`.

- Uncomment the `user_config` setting and set it to `true`.
- Uncomment the `pools` and `keys` sections and update them as required. You can define custom pools and client names altogether with the `cephx` capabilities.

b. Add the `osd_pool_default_pg_num` setting to the `ceph_conf_overrides` section in the `all.yml` file:

```
ceph_conf_overrides:
  global:
    osd_pool_default_pg_num: <number>
```

Replace `<number>` with the default number of placement groups.

5. Run the Ansible playbook:

```
$ ansible-playbook site.yml
```

4.2. CEPH COMMAND-LINE INTERFACE INSTALLATION

The Ceph command-line interface (CLI) is provided by the `ceph-common` package and includes the following utilities:

- `ceph`
- `ceph-authtool`
- `ceph-dencoder`
- `rados`

To install the Ceph CLI:

1. On the client node, enable the [Tools repository](#).
2. On the client node, install the `ceph-common` package:

```
# yum install ceph-common
```

3. From the initial monitor node, copy the Ceph configuration file, in this case `ceph.conf`, and the administration keyring to the client node:

Syntax

```
# scp /etc/ceph/<cluster_name>.conf
  <user_name>@<client_host_name>:/etc/ceph/
# scp /etc/ceph/<cluster_name>.client.admin.keyring
  <user_name>@<client_host_name>:/etc/ceph/
```

Example

```
# scp /etc/ceph/ceph.conf root@node1:/etc/ceph/
```

```
# scp /etc/ceph/ceph.client.admin.keyring root@node1:/etc/ceph/
```

Replace `<client_host_name>` with the host name of the client node.

4.3. CEPH BLOCK DEVICE INSTALLATION

The following procedure shows how to install and mount a thin-provisioned, resizable Ceph Block Device.



IMPORTANT

Ceph Block Devices must be deployed on separate nodes from the Ceph Monitor and OSD nodes. Running kernel clients and kernel server daemons on the same node can lead to kernel deadlocks.

Before you start

- Ensure to perform the tasks listed in the [Section 4.2, “Ceph Command-line Interface Installation”](#) section.
- If you use Ceph Block Devices as a back end for virtual machines (VMs) that use QEMU, increase the default file descriptor. See the [Ceph - VM hangs when transferring large amounts of data to RBD disk](#) Knowledgebase article for details.

Installing Ceph Block Devices by Using the Command Line

1. Create a Ceph Block Device user named `client.rbd` with full permissions to files on OSD nodes (`osd 'allow rwx'`) and output the result to a keyring file:

```
ceph auth get-or-create client.rbd mon 'allow r' osd 'allow rwx
pool=<pool_name>' \
-o /etc/ceph/rbd.keyring
```

Replace `<pool_name>` with the name of the pool that you want to allow `client.rbd` to have access to, for example `rbd`:

```
# ceph auth get-or-create \
client.rbd mon 'allow r' osd 'allow rwx pool=rbd' \
-o /etc/ceph/rbd.keyring
```

See the [User Management](#) section in the Red Hat Ceph Storage Administration Guide for more information about creating users.

2. Create a block device image:

```
rbd create <image_name> --size <image_size> --pool <pool_name> \
--name client.rbd --keyring /etc/ceph/rbd.keyring
```

Specify `<image_name>`, `<image_size>`, and `<pool_name>`, for example:

```
$ rbd create image1 --size 4096 --pool rbd \
--name client.rbd --keyring /etc/ceph/rbd.keyring
```



WARNING

The default Ceph configuration includes the following Ceph Block Device features:

- **layering**
- **exclusive-lock**
- **object-map**
- **deep-flatten**
- **fast-diff**

If you use the kernel RBD (**krbd**) client, you will not be able to map the block device image because the current kernel version included in Red Hat Enterprise Linux 7.3 does not support **object-map**, **deep-flatten**, and **fast-diff**.

To work around this problem, disable the unsupported features. Use one of the following options to do so:

- Disable the unsupported features dynamically:

```
rbd feature disable <image_name> <feature_name>
```

For example:

```
# rbd feature disable image1 object-map deep-
flatten fast-diff
```

- Use the **--image-feature layering** option with the **rbd create** command to enable only **layering** on newly created block device images.
- Disable the features by default in the Ceph configuration file:

```
rbd_default_features = 1
```

This is a known issue, for details see the [Release Notes Red Hat Ceph Storage 2.2](#).

All these features work for users that use the user-space RBD client to access the block device images.

3. Map the newly created image to the block device:

```
rbd map <image_name> --pool <pool_name>\
--name client.rbd --keyring /etc/ceph/rbd.keyring
```

For example:

```
# rbd map image1 --pool rbd --name client.rbd \
--keyring /etc/ceph/rbd.keyring
```



IMPORTANT

Kernel block devices currently only support the legacy straw bucket algorithm in the CRUSH map. If you have set the CRUSH tunables to optimal, you must set them to legacy or an earlier major release, otherwise, you will not be able to map the image.

Alternatively, replace `straw2` with `straw` in the CRUSH map. For details, see the [Editing a CRUSH Map](#) chapter in the Storage Strategies Guide for Red Hat Ceph Storage 2.

4. Use the block device by creating a file system:

```
mkfs.ext4 -m5 /dev/rbd/<pool_name>/<image_name>
```

Specify the pool name and the image name, for example:

```
# mkfs.ext4 -m5 /dev/rbd/rbd/image1
```

This can take a few moments.

5. Mount the newly created file system:

```
mkdir <mount_directory>
mount /dev/rbd/<pool_name>/<image_name> <mount_directory>
```

For example:

```
# mkdir /mnt/ceph-block-device
# mount /dev/rbd/rbd/image1 /mnt/ceph-block-device
```

For additional details, see the Red Hat Ceph Storage [Block Device Guide](#).

4.4. CEPH OBJECT GATEWAY INSTALLATION

The Ceph object gateway, also known as the RADOS gateway, is an object storage interface built on top of the `librados` API to provide applications with a RESTful gateway to Ceph storage clusters.

For more information about the Ceph object gateway, see the Object Gateway Guide for [Red Hat Enterprise Linux](#).

There are two ways to install the Ceph object gateway:

- Using the Ansible automation application, see [Section 4.4.1, “Installing Ceph Object Gateway by using Ansible”](#) for details
- Using the command-line interface, see [Section 4.3.2, “Installing Ceph Object Gateway Manually”](#) for details

4.4.1. Installing Ceph Object Gateway by using Ansible

Perform the following tasks on the Ansible administration node, see [Install Ceph Ansible](#) for details.

1. As root, create the `rgws` file from the sample file:

```
# cd /etc/ansible/group_vars
# cp rgws.yml.sample rgws.yml
```

2. To copy the administrator key to the Ceph Object Gateway node, uncomment the `copy_admin_key` setting in the `/etc/ansible/group_vars/rgws.yml` file:

```
copy_admin_key: true
```

3. The `rgws.yml` file may specify a different default port than the default port `7480`. For example:

```
ceph_rgw_civetweb_port: 80
```

4. Generally, to change default settings, uncomment the settings in the `rgw.yml` file, and make changes accordingly. To make additional changes to settings that aren't in the `rgw.yml` file, use `ceph_conf_overrides:` in the `all.yml` file. For example, set the `rgw_dns_name:` with the host of the DNS server and ensure the cluster's DNS server to configure it for wild cards to enable S3 subdomains.

```
ceph_conf_overrides:
  client.rgw.rgw1:
    rgw_dns_name: <host_name>
    rgw_override_bucket_index_max_shards: 16
    rgw_bucket_default_quota_max_objects: 1638400
```

For advanced configuration details, refer to the [Ceph Object Gateway for Production](#) guide. Advanced topics include:

- [Configuring Ansible Groups](#)
 - [Developing Storage Strategies](#). See [Creating the Root Pool](#), [Creating System Pools](#), and [Creating Data Placement Strategies](#) for additional details on how create and configure the pools. See [Bucket Sharding](#) for configuration details on bucket sharding.
5. Add the `radosgw_interface` setting to the `group_vars/all.yml` file, and specify the network interface for the Ceph Object Gateway. For example:

```
radosgw_interface: eth0
```

6. Add gateway hosts to the `/etc/ansible/hosts` file under the `[rgws]` section to identify their roles to Ansible. If the hosts have sequential naming, you can use a range. For example:

```
[rgws]
<rgw_host_name_1>
<rgw_host_name_2>
<rgw_host_name[3..10]>
```

7. Navigate to the Ansible configuration directory, `/etc/ansible/`:

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

8. Run the Ansible playbook:

```
$ ansible-playbook site.yml
```



NOTE

Ansible ensures that each Ceph Object Gateway is running.

For a single site configuration, add Ceph Object Gateways to the Ansible configuration.

For multi-site deployments, you should have an Ansible configuration for each zone. That is, Ansible will create a Ceph storage cluster and gateway instances for that zone.

After installation for a multi-site cluster is complete, proceed to the [Multi-site](#) chapter in the [Object Gateway Guide](#) for Red Hat Enterprise Linux for details on configuring a cluster for multi-site.

4.4.2. Installing Ceph Object Gateway Manually

1. Enable the Red Hat Ceph Storage 2 [Tools repository](#). For ISO-based installations, see the [ISO installation section](#).
2. On the Object Gateway node, install the `ceph-radosgw` package:

```
# yum install ceph-radosgw
```

3. On the initial Monitor node, do the following steps.

- a. Update the Ceph configuration file as follows:

```
[client.rgw.<obj_gw_hostname>]
host = <obj_gw_hostname>
rgw frontends = "civetweb port=80"
rgw dns name = <obj_gw_hostname>.example.com
```

Where `<obj_gw_hostname>` is a short host name of the gateway node. To view the short host name, use the `hostname -s` command.

- b. Copy the updated configuration file to the new Object Gateway node and all other nodes in the Ceph storage cluster:

Syntax

```
# scp /etc/ceph/<cluster_name>.conf
<user_name>@<target_host_name>:/etc/ceph
```

Example

```
# scp /etc/ceph/ceph.conf root@node1:/etc/ceph/
```

- c. Copy the `<cluster_name>.client.admin.keyring` file to the new Object Gateway node:

Syntax

```
# scp /etc/ceph/<cluster_name>.client.admin.keyring
<user_name>@<target_host_name>:/etc/ceph/
```

Example

```
# scp /etc/ceph/ceph.client.admin.keyring root@node1:/etc/ceph/
```

4. On the Object Gateway node, create the data directory:

Syntax

```
# mkdir -p /var/lib/ceph/radosgw/<cluster_name>-rgw.`hostname` -s`
```

Example

```
# mkdir -p /var/lib/ceph/radosgw/ceph-rgw.`hostname` -s`
```

5. On the Object Gateway node, add a user and keyring to bootstrap the object gateway:

Syntax

```
# ceph auth get-or-create client.rgw.`hostname` -s` osd 'allow rwx'
mon 'allow rw' -o /var/lib/ceph/radosgw/<cluster_name>-rgw.`hostname`
-s`/keyring
```

Example

```
# ceph auth get-or-create client.rgw.`hostname` -s` osd 'allow rwx'
mon 'allow rw' -o /var/lib/ceph/radosgw/ceph-rgw.`hostname` -
s`/keyring
```



IMPORTANT

When you provide capabilities to the gateway key you must provide the read capability. However, providing the Monitor write capability is optional; if you provide it, the Ceph Object Gateway will be able to create pools automatically.

In such a case, ensure to specify a reasonable number of placement groups in a pool. Otherwise, the gateway uses the default number, which might not be suitable for your needs. See [Ceph Placement Groups \(PGs\) per Pool Calculator](#) for details.

6. On the Object Gateway node, create the `done` file:

Syntax

```
# touch /var/lib/ceph/radosgw/<cluster_name>-rgw.`hostname` -s`/done
```

Example

```
# touch /var/lib/ceph/radosgw/ceph-rgw.`hostname` -s`/done
```

7. On the Object Gateway node, change the owner and group permissions:

```
# chown -R ceph:ceph /var/lib/ceph/radosgw
# chown -R ceph:ceph /var/log/ceph
# chown -R ceph:ceph /var/run/ceph
# chown -R ceph:ceph /etc/ceph
```

8. For storage clusters with custom names, as **root**, add the the following line:

Syntax

```
# echo "CLUSTER=<custom_cluster_name>" >> /etc/sysconfig/ceph
```

Example

```
# echo "CLUSTER=test123" >> /etc/sysconfig/ceph
```

9. On the Object Gateway node, open TCP port 80:

```
# firewall-cmd --zone=public --add-port=80/tcp
# firewall-cmd --zone=public --add-port=80/tcp --permanent
```

10. On the Object Gateway node, start and enable the **ceph-radosgw** process:

Syntax

```
# systemctl enable ceph-radosgw.target
# systemctl enable ceph-radosgw@rgw.<rgw_hostname>
# systemctl start ceph-radosgw@rgw.<rgw_hostname>
```

Example

```
# systemctl enable ceph-radosgw.target
# systemctl enable ceph-radosgw@rgw.node1
# systemctl start ceph-radosgw@rgw.node1
```

Once installed, the Ceph Object Gateway automatically creates pools if the write capability is set on the Monitor. See the [Pools](#) chapter in the [Storage Strategies Guide](#) for information on creating pools manually.

CHAPTER 5. UPGRADING CEPH STORAGE CLUSTER

There are two main upgrading paths:

- from Red Hat Ceph Storage 1.3 to 2 ([Important](#))
- between minor versions of Red Hat Ceph Storage 2 or between asynchronous updates ([Important](#))

5.1. UPGRADING FROM RED HAT CEPH STORAGE 1.3 TO 2



IMPORTANT

Please contact [Red Hat support](#) prior to upgrading, if you have a large Ceph Object Gateway storage cluster with millions of objects present in buckets.

For more details refer to the Red Hat Ceph Storage 2.5 [Release Notes](#), under the *Slow OSD startup after upgrading to Red Hat Ceph Storage 2.5* heading.

You can upgrade the Ceph Storage Cluster in a rolling fashion and while the cluster is running. Upgrade each node in the cluster sequentially, only proceeding to the next node after the previous node is done.

Red Hat recommends upgrading the Ceph components in the following order:

- Monitor nodes
- OSD nodes
- Ceph Object Gateway nodes
- All other Ceph client nodes

Two methods are available to upgrade a Red Hat Ceph Storage 1.3.2 to 2.0:

- Using Red Hat's Content Delivery Network (CDN)
- Using a Red Hat provided ISO image file

After upgrading the storage cluster you might have a health warning regarding the CRUSH map using legacy tunables. See the Red Hat Ceph Storage [Strategies Guide](#) for more information.

Example

```
$ ceph -s
  cluster 848135d7-cdb9-4084-8df2-fb5e41ae60bd
  health HEALTH_WARN
        crush map has legacy tunables (require bobtail, min is
firefly)
  monmap e1: 1 mons at {ceph1=192.168.0.121:6789/0}
        election epoch 2, quorum 0 ceph1
  osdmap e83: 2 osds: 2 up, 2 in
  pgmap v1864: 64 pgs, 1 pools, 38192 kB data, 17 objects
        10376 MB used, 10083 MB / 20460 MB avail
        64 active+clean
```

**IMPORTANT**

Red Hat recommends all Ceph clients to be running the same version as the Ceph storage cluster.

5.1.1. Upgrading a Ceph Monitor Node

Red Hat recommends a minimum of three Monitors for a production storage cluster. There must be an odd number of Monitors. While you are upgrading one Monitor, the storage cluster will still have quorum.

For Red Hat Ceph Storage 1.3.2 Monitor nodes running on Red Hat Enterprise Linux 7, perform the following steps on each Monitor node in the storage cluster. Sequentially upgrading one Monitor node at a time.

1. As **root**, disable any Red Hat Ceph Storage 1.3.x repositories:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-1.3-mon-
rpms --disable=rhel-7-server-rhceph-1.3-installer-rpms --
disable=rhel-7-server-rhceph-1.3-calamari-rpms
```

**NOTE**

If an ISO-based installation was performed for Red Hat Ceph Storage 1.3.x, then skip this first step.

2. Enable the Red Hat Ceph Storage 2 [Monitor repository](#). For ISO-based installations, see the [ISO installation section](#).
3. As **root**, stop the Monitor process:

Syntax

```
# service ceph stop <daemon_type>.<monitor_host_name>
```

Example

```
# service ceph stop mon.node1
```

4. As **root**, update the **ceph-mon** package:

```
# yum update ceph-mon
```

5. As **root**, update the owner and group permissions:

Syntax

```
# chown -R <owner>:<group> <path_to_directory>
```

Example

```
# chown -R ceph:ceph /var/lib/ceph/mon
```

```
# chown -R ceph:ceph /var/log/ceph
# chown -R ceph:ceph /var/run/ceph
# chown ceph:ceph /etc/ceph/ceph.client.admin.keyring
# chown ceph:ceph /etc/ceph/ceph.conf
# chown ceph:ceph /etc/ceph/rbdmap
```



NOTE

If the Ceph Monitor node is co-located with an OpenStack Controller node, then the Glance and Cinder keyring files must be owned by **glance** and **cinder** respectively. For example:

```
# ls -l /etc/ceph/
...
-rw-----. 1 glance glance      64 <date>
ceph.client.glance.keyring
-rw-----. 1 cinder cinder      64 <date>
ceph.client.cinder.keyring
...
```

6. If SELinux is set to enforcing or permissive mode, then set a relabelling of the SELinux context on files for the next reboot:

```
# touch /.autorelabel
```



WARNING

Relabeling will take a long time to complete, because SELinux must traverse every file system and fix any mislabeled files. To exclude directories from being relabelled, add the directory to the `/etc/selinux/fixfiles_exclude_dirs` file before rebooting.



NOTE

In environments with large number of objects per placement group (PG), the directory enumeration speed will decrease, causing a negative impact to performance. This is caused by the addition of xattr queries which verifies the SELinux context. Setting the context at mount time removes the xattr queries for context and helps overall disk performance, especially on slower disks.

Add the following line to the `[osd]` section in the `/etc/ceph/ceph.conf` file:

```
osd_mount_options_xfs=rw,noatime,inode64,context="system_u:object_r:ceph_var_lib_t:s0"
```

1. As **root**, replay device events from the kernel:

```
# udevadm trigger
```

2. As **root**, enable the **ceph-mon** process:

```
# systemctl enable ceph-mon.target
# systemctl enable ceph-mon@<monitor_host_name>
```

3. As **root**, reboot the Monitor node:

```
# shutdown -r now
```

4. Once the Monitor node is up, check the health of the Ceph storage cluster before moving to the next Monitor node:

```
# ceph -s
```

To add more Red Hat Ceph Storage Monitors to the storage cluster, see the [Red Hat Ceph Storage Administration Guide](#)

5.1.2. Upgrading a Ceph OSD Node

Red Hat recommends having a minimum of three OSD nodes in the Ceph storage cluster. For Red Hat Ceph Storage 1.3.2 OSD nodes running on Red Hat Enterprise Linux 7, perform the following steps on each OSD node in the storage cluster. Sequentially upgrading one OSD node at a time.

During the upgrade of an OSD node, some placement groups will become degraded, because the OSD might be down or restarting. You will need to tell the storage cluster not to mark an OSD out, because you do not want to trigger a recovery. The default behavior is to mark an OSD out of the CRUSH map after five minutes.

On a Monitor node, set **noout** and **norebalance** flags for the OSDs:

```
# ceph osd set noout
# ceph osd set norebalance
```

Perform the following steps on each OSD node in the storage cluster. Sequentially upgrading one OSD node at a time. If an ISO-based installation was performed for Red Hat Ceph Storage 1.3, then skip this first step.

1. As **root**, disable the Red Hat Ceph Storage 1.3 repositories:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-1.3-osd-rpms --disable=rhel-7-server-rhceph-1.3-installer-rpms --disable=rhel-7-server-rhceph-1.3-calamari-rpms
```

2. Enable the Red Hat Ceph Storage 2 [OSD repository](#). For ISO-based installations, see the [ISO installation section](#).

3. As **root**, stop any running OSD process:

Syntax

```
# service ceph stop <daemon_type>.<osd_id>
```

-

Example

```
# service ceph stop osd.0
```

4. As **root**, update the **ceph-osd** package:

```
# yum update ceph-osd
```

5. As **root**, update the owner and group permissions on the newly created directory and files:

Syntax

```
# chown -R <owner>:<group> <path_to_directory>
```

Example

```
# chown -R ceph:ceph /var/lib/ceph/osd
# chown -R ceph:ceph /var/log/ceph
# chown -R ceph:ceph /var/run/ceph
# chown -R ceph:ceph /etc/ceph
```

**NOTE**

Running the following **find** command might speed up the process of changing ownership by running the **chown** command in parallel on a Ceph storage cluster with a large number of disks:

```
# find /var/lib/ceph/osd -maxdepth 1 -mindepth 1 -print |
xargs -P12 -n1 chown -R ceph:ceph
```

6. If SELinux is set to enforcing or permissive mode, then set a relabelling of the SELinux context on files for the next reboot:

```
# touch /.autorelabel
```

**WARNING**

Relabeling will take a long time to complete, because SELinux must traverse every file system and fix any mislabeled files. To exclude directories from being relabelled, add the directory to the `/etc/selinux/fixfiles_exclude_dirs` file before rebooting.

**NOTE**

In environments with large number of objects per placement group (PG), the directory enumeration speed will decrease, causing a negative impact to performance. This is caused by the addition of xattr queries which verifies the SELinux context. Setting the context at mount time removes the xattr queries for context and helps overall disk performance, especially on slower disks.

Add the following line to the `[osd]` section in the `/etc/ceph/ceph.conf` file:

```
osd_mount_options_xfs=rw,noatime,inode64,context="system_u:object_r:ceph_var_lib_t:s0"
```

1. As **root**, replay device events from the kernel:

```
# udevadm trigger
```

2. As **root**, enable the **ceph-osd** process:

```
# systemctl enable ceph-osd.target
# systemctl enable ceph-osd@<osd_id>
```

3. As **root**, reboot the OSD node:

```
# shutdown -r now
```

4. Move to the next OSD node.

**NOTE**

While the **noout** and **norebalance** flags are set, the storage cluster will have the **HEALTH_WARN** status:

```
$ ceph health
HEALTH_WARN noout,norebalance flag(s) set
```

Once you are done upgrading the Ceph storage cluster, unset the previously set OSD flags and verify the storage cluster status.

On a Monitor node, and after all OSD nodes have been upgraded, unset the **noout** and **norebalance** flags:

```
# ceph osd unset noout
# ceph osd unset norebalance
```

In addition, set the **require_jewel_osds** flag. This flag ensures that no more OSDs with Red Hat Ceph Storage 1.3 can be added to the storage cluster. If you do not set this flag, the storage status will be **HEALTH_WARN**.

```
# ceph osd set require_jewel_osds
```

To expand the storage capacity by adding new OSDs to the storage cluster, see the Red Hat Ceph Storage [Administration Guide](#) for more details.

5.1.3. Upgrading the Ceph Object Gateway Nodes

This section describes steps to upgrade a Ceph Object Gateway node to a later version.

Before You Start

- Red Hat recommends putting a Ceph Object Gateway behind a load balancer, such as [HAProxy](#). If you use a load balancer, remove the Ceph Object Gateway from the load balancer once no requests are being served.
- If you use a custom name for the region pool, specified in the `rgw_region_root_pool` parameter, add the `rgw_zonegroup_root_pool` parameter to the `[global]` section of the Ceph configuration file. Set the value of `rgw_zonegroup_root_pool` to be the same as `rgw_region_root_pool`, for example:

```
[global]
rgw_zonegroup_root_pool = .us.rgw.root
```

Procedure: Upgrading the Ceph Object Gateway Node

1. If you used online repositories to install Red Hat Ceph Storage, disable the 1.3 repositories.

```
# subscription-manager repos --disable=rhel-7-server-rhceph-1.3-
tools-rpms --disable=rhel-7-server-rhceph-1.3-installer-rpms --
disable=rhel-7-server-rhceph-1.3-calamari-rpms
```

2. Enable the Red Hat Ceph Storage 2 [Tools repository](#). For ISO-based installations, see the [ISO Installation](#) section.
3. Stop the Ceph Object Gateway process (`ceph-radosgw`):

```
# service ceph-radosgw stop
```

4. Update the `ceph-radosgw` package:

```
# yum update ceph-radosgw
```

5. Change the owner and group permissions on the newly created `/var/lib/ceph/radosgw/` and `/var/log/ceph/` directories and their content to `ceph`.

```
# chown -R ceph:ceph /var/lib/ceph/radosgw
# chown -R ceph:ceph /var/log/ceph
```

6. If SELinux is set to run in enforcing or permissive mode, instruct it to relabel SELinux context on the next boot.

```
# touch /.autorelabel
```



IMPORTANT

Relabeling takes a long time to complete, because SELinux must traverse every file system and fix any mislabeled files. To exclude directories from being relabeled, add them to the `/etc/selinux/fixfiles_exclude_dirs` file before rebooting.



NOTE

In environments with large number of objects per placement group (PG), the directory enumeration speed will decrease, causing a negative impact to performance. This is caused by the addition of xattr queries which verifies the SELinux context. Setting the context at mount time removes the xattr queries for context and helps overall disk performance, especially on slower disks.

Add the following line to the `[osd]` section in the `/etc/ceph/ceph.conf` file:

```
osd_mount_options_xfs=rw,noatime,inode64,context="system_u:object_r:ceph_var_lib_t:s0"
```

1. Enable the `ceph-radosgw` process.

```
# systemctl enable ceph-radosgw.target
# systemctl enable ceph-radosgw@rgw.<hostname>
```

Replace `<hostname>` with the name of the Ceph Object Gateway host, for example `gateway-node`.

```
# systemctl enable ceph-radosgw.target
# systemctl enable ceph-radosgw@rgw.gateway-node
```

2. Reboot the Ceph Object Gateway node.

```
# shutdown -r now
```

3. If you use a load balancer, add the Ceph Object Gateway node back to the load balancer.
4. Repeat these steps on a next Ceph Object Gateway node.

See Also

- The [Ceph Object Gateway Guide for Red Hat Enterprise Linux](#)

5.1.4. Upgrading a Ceph Client Node

Ceph clients can be the RADOS Gateway, RADOS block devices, the Ceph command-line interface (CLI), Nova compute nodes, `qemu-kvm`, or any custom application using the Ceph client-side libraries. Red Hat recommends all Ceph clients to be running the same version as the Ceph storage cluster.

**IMPORTANT**

Red Hat recommends stopping all IO running against a Ceph client node while the packages are being upgraded. Not stopping all IO might cause unexpected errors to occur.

1. As **root**, disable any Red Hat Ceph Storage 1.3 repositories:

```
# subscription-manager repos --disable=rhel-7-server-rhceph-1.3-
tools-rpms --disable=rhel-7-server-rhceph-1.3-installer-rpms --
disable=rhel-7-server-rhceph-1.3-calamari-rpms
```

**NOTE**

If an ISO-based installation was performed for Red Hat Ceph Storage 1.3.x clients, then skip this first step.

2. On the client node, enable the [Tools repository](#).
3. On the client node, update the **ceph-common** package:

```
# yum update ceph-common
```

Any application depending on the Ceph client-side libraries will have to be restarted after upgrading the Ceph client package.

**NOTE**

For Nova compute nodes with running **qemu-kvm** instances or if using a dedicated **qemu-kvm** client, then stopping and starting the **qemu-kvm** instance processes is required. A simple restart will not work here.

5.2. UPGRADING BETWEEN MINOR VERSIONS AND APPLYING ASYNCHRONOUS UPDATES

**IMPORTANT**

Please contact [Red Hat support](#) prior to upgrading, if you have a large Ceph Object Gateway storage cluster with millions of objects present in buckets.

For more details refer to the Red Hat Ceph Storage 2.5 [Release Notes](#), under the *Slow OSD startup after upgrading to Red Hat Ceph Storage 2.5* heading.

Use the Ansible `rolling_update.yml` playbook located in the `infrastructure-playbooks` directory from the administration node to upgrade between two minor versions of Red Hat Ceph Storage 2 or to apply asynchronous updates.

Currently, this is the only supported way to upgrade to a minor version. If you use a cluster that was not deployed by using Ansible, see [Section 3.1.7, “Taking over an Existing Cluster”](#) for details on configuring Ansible to use a cluster that was deployed without it.

Ansible upgrades the Ceph nodes in the following order:

- Monitor nodes
- OSD nodes
- MDS nodes
- Ceph Object Gateway nodes
- All other Ceph client nodes



NOTE

Upgrading encrypted OSD nodes is the same as upgrading OSD nodes that are not encrypted.

Before you Start

- On the Ansible Administration node, enable the Red Hat Ceph Storage 2 Tools repository:

```
# subscription-manager repos --enable=rhel-7-server-rhceph-2-tools-rpms
```

- On the Ansible Administration node, ensure the latest version of `ceph-ansible` is installed:

```
# yum update ceph-ansible
```

- In the `rolling_update.yml` playbook, change the `health_osd_check_retries` and `health_osd_check_delay` values to **40** and **30** respectively. For each OSD node, Ansible will wait up to 20 minutes. Also, Ansible will check the cluster health every 30 seconds, waiting before continuing the upgrade process. Set the following values:

```
health_osd_check_retries: 40
health_osd_check_delay: 30
```

- If the Ceph nodes are not connected to the Red Hat Content Delivery Network (CDN) and you used an ISO image to install Red Hat Ceph Storage, update the local repository with the latest version of Red Hat Ceph Storage. See [Section 2.3, “Enabling the Red Hat Ceph Storage Repositories”](#) for details.
- If you upgrade from Red Hat Ceph Storage 2.1 to 2.2, review the [Section 5.2.1, “Changes Between Ansible 2.1 and 2.2”](#) section first. Ansible 2.2 uses slightly different file names and setting.

Procedure: Updating the Ceph Storage Cluster by using Ansible

1. On the Ansible administration node, edit the `/etc/ansible/hosts` file with custom `osd_scenarios` if your cluster has any.
2. On the Ansible administration node, navigate to the `/usr/share/ceph-ansible/` directory:

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

3. In the `group_vars/all.yml` file, uncomment the `upgrade_ceph_packages` option and set it to **True**:

```
upgrade_ceph_packages: True
```

4. In the `group_vars/all.yml` file, set `generate_fsid` to `false`.
5. Get the current cluster `fsid` by executing `ceph fsid`. Set the retrieved `fsid` in `group_vars/all.yml`.
6. If the cluster you want to upgrade contains any Ceph Object Gateway nodes, add the `radosgw_interface` parameter to the `group_vars/all.yml` file.

```
radosgw_interface: <interface>
```

Replace:

- `<interface>` with the interface that the Ceph Object Gateway nodes listen to

7. Run the `rolling_update.yml` playbook:

```
# cp infrastructure-playbooks/rolling_update.yml .
$ ansible-playbook rolling_update.yml
```

When upgrading from version 2.4 to 2.5, run the playbook using the following command:

```
$ ansible-playbook rolling_update.yml -e jewel_minor_update=true
```

8. From the RBD mirroring daemon node, upgrade `rbd-mirror` manually:

```
# yum upgrade rbd-mirror
```

Restart the daemon:

```
# systemctl restart ceph-rbd-mirror@<client-id>
```



IMPORTANT

The `rolling_update.yml` playbook includes the `serial` variable that adjusts the number of nodes to be updated simultaneously. Red Hat strongly recommends to use the default value (1), which ensures that hosts will be upgraded one by one.

5.2.1. Changes Between Ansible 2.1 and 2.2

Red Hat Ceph Storage 2.2 includes Ansible 2.2 that introduces the following changes:

- Files in the `group_vars` directory have the `.yaml` extension. Before updating to 2.2, you must rename them. To do so:
Navigate to the Ansible directory:

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

Change the names of the files in `group_vars`:

```
# mv groups_vars/all groups_vars/all.yaml
```

```
# mv groups_vars/mons groups_vars/mons.yml
# mv groups_vars/osds groups_vars/osds.yml
# mv groups_vars/mdss groups_vars/mdss.yml
# mv groups_vars/rgws groups_vars/rgws.yml
```

- Ansible 2.2 uses different variable names and handles this change automatically when updating to version 2.2. See [Table 5.1, “Differences in Variable Names Between Ansible 2.1 and 2.2”](#) table for details.

Table 5.1. Differences in Variable Names Between Ansible 2.1 and 2.2

Ansible 2.1 variable name	Ansible 2.2 variable name
ceph_stable_rh_storage	ceph_rhcs
ceph_stable_rh_storage_version	ceph_rhcs_version
ceph_stable_rh_storage_cdn_install	ceph_rhcs_cdn_install
ceph_stable_rh_storage_iso_install	ceph_rhcs_iso_install
ceph_stable_rh_storage_iso_path	ceph_rhcs_iso_path
ceph_stable_rh_storage_mount_path	ceph_rhcs_mount_path
ceph_stable_rh_storage_repository_path	ceph_rhcs_repository_path

CHAPTER 6. WHAT TO DO NEXT?

This is only the beginning of what Red Hat Ceph Storage can do to help you meet the challenging storage demands of the modern data center. Here are links to more information on a variety of topics:

- Benchmarking performance and accessing performance counters, see the [Red Hat Ceph Storage Administration Guide](#).
- Creating and managing snapshots, see the [Red Hat Ceph Storage Block Device Guide](#) .
- Expanding the Red Hat Ceph Storage cluster, see the [Red Hat Ceph Storage Administration Guide](#).
- Mirroring RADOS Block Devices, see the [Red Hat Ceph Storage Block Device Guide](#) .
- Process management, enabling debug logging, and related topics, see the [Red Hat Ceph Storage Administration Guide](#).
- Tunable parameters, see the [Red Hat Ceph Storage Configuration Guide](#) .
- Using Ceph as the back end storage for OpenStack, see the [Red Hat OpenStack Platform Storage Guide](#).

APPENDIX A. TROUBLESHOOTING

A.1. ANSIBLE STOPS INSTALLATION BECAUSE IT DETECTS LESS DEVICES THAN IT EXPECTED

The Ansible automation application stops the installation process and returns the following error:

```
- name: fix partitions gpt header or labels of the osd disks (autodiscover
disks)
  shell: "sgdisk --zap-all --clear --mbrtogpt -- '/dev/{{ item.0.item.key
}}' || sgdisk --zap-all --clear --mbrtogpt -- '/dev/{{ item.0.item.key
}}'"
  with_together:
    - "{{ osd_partition_status_results.results }}"
    - "{{ ansible_devices }}"
  changed_when: false
  when:
    - ansible_devices is defined
    - item.0.item.value.removable == "0"
    - item.0.item.value.partitions|count == 0
    - item.0.rc != 0
```

What this means:

When the `osd_auto_discovery` parameter is set to `true` in the `/etc/ansible/group_vars/osds.yml` file, Ansible automatically detects and configures all the available devices. During this process, Ansible expects that all OSDs use the same devices. The devices get their names in the same order in which Ansible detects them. If one of the devices fails on one of the OSDs, Ansible fails to detect the failed device and stops the whole installation process.

Example situation:

1. Three OSD nodes (`host1`, `host2`, `host3`) use the `/dev/sdb`, `/dev/sdc`, and `dev/sdd` disks.
2. On `host2`, the `/dev/sdc` disk fails and is removed.
3. Upon the next reboot, Ansible fails to detect the removed `/dev/sdc` disk and expects that only two disks will be used for `host2`, `/dev/sdb` and `/dev/sdc` (formerly `/dev/sdd`).
4. Ansible stops the installation process and returns the above error message.

To fix the problem:

In the `/etc/ansible/hosts` file, specify the devices used by the OSD node with the failed disk (`host2` in the Example situation above):

```
[osds]
host1
host2 devices="[ '/dev/sdb', '/dev/sdc' ]"
host3
```

See [Installing Ceph Ansible](#) for details.

