



Red Hat Ceph Storage 8

Dashboard Guide

Monitoring Ceph Cluster with Ceph Dashboard

Red Hat Ceph Storage 8 Dashboard Guide

Monitoring Ceph Cluster with Ceph Dashboard

Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This guide explains how to use the Red Hat Ceph Storage Dashboard for monitoring and management purposes. Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message.

Table of Contents

CHAPTER 1. CEPH DASHBOARD OVERVIEW	6
1.1. CEPH DASHBOARD COMPONENTS	9
1.2. RED HAT CEPH STORAGE DASHBOARD ARCHITECTURE	10
CHAPTER 2. CEPH DASHBOARD INSTALLATION AND ACCESS	11
2.1. NETWORK PORT REQUIREMENTS FOR CEPH DASHBOARD	12
2.2. ACCESSING THE CEPH DASHBOARD	14
2.3. EXPANDING THE CLUSTER ON THE CEPH DASHBOARD	15
2.4. UPGRADING A CLUSTER	19
2.5. TOGGING CEPH DASHBOARD FEATURES	19
2.6. UNDERSTANDING THE LANDING PAGE OF THE CEPH DASHBOARD	21
2.7. CHANGING THE DASHBOARD PASSWORD USING THE CEPH DASHBOARD	24
2.8. CHANGING THE CEPH DASHBOARD PASSWORD USING THE COMMAND LINE INTERFACE	25
2.9. SETTING ADMIN USER PASSWORD FOR GRAFANA	26
2.10. ENABLING RED HAT CEPH STORAGE DASHBOARD MANUALLY	27
2.11. USING SINGLE SIGN-ON WITH THE DASHBOARD	29
2.11.1. Creating an admin account for syncing users to the Ceph dashboard	29
2.11.2. Syncing users to the Ceph dashboard using Red Hat Single Sign-On	31
2.11.3. Enabling Single Sign-On with SAML 2.0 for the Ceph Dashboard	37
2.11.4. Enabling OAuth2 single sign-on (Technology Preview)	39
2.11.5. Disabling Single Sign-On for the Ceph Dashboard	40
CHAPTER 3. MANAGING ROLES ON THE CEPH DASHBOARD	42
3.1. USER ROLES AND PERMISSIONS ON THE CEPH DASHBOARD	42
3.2. CREATING ROLES ON THE CEPH DASHBOARD	44
3.3. EDITING ROLES ON THE CEPH DASHBOARD	46
3.4. CLONING ROLES ON THE CEPH DASHBOARD	47
3.5. DELETING ROLES ON THE CEPH DASHBOARD	48
CHAPTER 4. MANAGING USERS ON THE CEPH DASHBOARD	50
4.1. CREATING USERS ON THE CEPH DASHBOARD	50
4.2. EDITING USERS ON THE CEPH DASHBOARD	51
4.3. DELETING USERS ON THE CEPH DASHBOARD	53
4.4. USER CAPABILITIES	54
4.4.1. Capabilities	54
4.5. ACCESS CAPABILITIES	55
4.6. CREATING USER CAPABILITIES	56
4.7. EDITING USER CAPABILITIES	58
4.8. IMPORTING USER CAPABILITIES	59
4.9. EXPORTING USER CAPABILITIES	60
4.10. DELETING USER CAPABILITIES	61
CHAPTER 5. MANAGING CEPH DAEMONS	63
5.1. DAEMON ACTIONS	63
CHAPTER 6. MONITORING THE CLUSTER ON THE CEPH DASHBOARD	65
6.1. MONITORING HOSTS OF THE CEPH CLUSTER ON THE DASHBOARD	65
6.2. VIEWING AND EDITING THE CONFIGURATION OF THE CEPH CLUSTER ON THE DASHBOARD	67
6.3. VIEWING AND EDITING THE MANAGER MODULES OF THE CEPH CLUSTER ON THE DASHBOARD	68
6.4. MONITORING MONITORS OF THE CEPH CLUSTER ON THE DASHBOARD	69
6.5. MONITORING SERVICES OF THE CEPH CLUSTER ON THE DASHBOARD	71
6.6. MONITORING CEPH OSDS ON THE DASHBOARD	71

6.7. MONITORING HAPROXY ON THE DASHBOARD	73
6.8. VIEWING THE CRUSH MAP OF THE CEPH CLUSTER ON THE DASHBOARD	74
6.9. FILTERING LOGS OF THE CEPH CLUSTER ON THE DASHBOARD	75
6.10. VIEWING CENTRALIZED LOGS OF THE CEPH CLUSTER ON THE DASHBOARD	76
6.11. MONITORING POOLS OF THE CEPH CLUSTER ON THE DASHBOARD	79
6.12. MONITORING CEPH FILE SYSTEMS ON THE DASHBOARD	80
6.13. MONITORING CEPH OBJECT GATEWAY DAEMONS ON THE DASHBOARD	81
6.14. MONITORING BLOCK DEVICE IMAGES ON THE CEPH DASHBOARD	82
CHAPTER 7. MANAGING ALERTS ON THE CEPH DASHBOARD	83
7.1. ENABLING MONITORING STACK	87
7.2. CONFIGURING GRAFANA CERTIFICATE	89
7.3. ADDING ALERTMANAGER WEBHOOKS	91
7.4. VIEWING ALERTS ON THE CEPH DASHBOARD	93
7.5. CREATING A SILENCE ON THE CEPH DASHBOARD	93
7.6. RECREATING A SILENCE ON THE CEPH DASHBOARD	94
7.7. EDITING A SILENCE ON THE CEPH DASHBOARD	95
7.8. EXPIRING A SILENCE ON THE CEPH DASHBOARD	96
CHAPTER 8. MANAGING NFS GANESHA EXPORTS ON THE CEPH DASHBOARD	98
8.1. CONFIGURING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD	98
8.2. CONFIGURING NFS EXPORTS WITH CEPHFS ON THE CEPH DASHBOARD	101
8.3. EDITING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD	104
8.4. DELETING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD	105
CHAPTER 9. MANAGING POOLS ON THE CEPH DASHBOARD	106
9.1. CREATING POOLS ON THE CEPH DASHBOARD	106
9.2. EDITING POOLS ON THE CEPH DASHBOARD	108
9.3. DELETING POOLS ON THE CEPH DASHBOARD	108
CHAPTER 10. MANAGING HOSTS ON THE CEPH DASHBOARD	110
10.1. ENTERING MAINTENANCE MODE	110
10.2. EXITING MAINTENANCE MODE	111
10.3. REMOVING HOSTS USING THE CEPH DASHBOARD	111
CHAPTER 11. MANAGING CEPH OSDS ON THE DASHBOARD	114
11.1. MANAGING THE OSDS ON THE CEPH DASHBOARD	114
11.2. REPLACING THE FAILED OSDS ON THE CEPH DASHBOARD	123
CHAPTER 12. MANAGING CEPH OBJECT GATEWAY USING THE DASHBOARD	127
12.1. MANUALLY ADDING CEPH OBJECT GATEWAY LOGIN CREDENTIALS TO THE DASHBOARD	127
12.2. CREATING THE CEPH OBJECT GATEWAY SERVICES WITH SSL USING THE DASHBOARD	128
12.3. CONFIGURING HIGH AVAILABILITY FOR THE CEPH OBJECT GATEWAY ON THE DASHBOARD	130
12.4. MANAGING CEPH OBJECT GATEWAY USERS ON THE DASHBOARD	132
12.4.1. Creating Ceph object gateway users on the dashboard	133
12.4.2. Adding roles to the Ceph Object Gateway users on the dashboard	134
12.4.3. Creating Ceph object gateway subusers on the dashboard	135
12.4.4. Editing Ceph object gateway users on the dashboard	136
12.4.5. Deleting Ceph Object Gateway users on the dashboard	138
12.5. MANAGING CEPH OBJECT GATEWAY BUCKETS ON THE DASHBOARD	139
12.5.1. Creating Ceph object gateway buckets on the dashboard	139
12.5.2. Editing Ceph object gateway buckets on the dashboard	141
12.5.3. Deleting Ceph Object Gateway buckets on the dashboard	144
12.6. MONITORING MULTI-SITE OBJECT GATEWAY CONFIGURATION ON THE CEPH DASHBOARD	144
12.7. VIEWING CEPH OBJECT GATEWAY PER-USER AND PER-BUCKET PERFORMANCE COUNTERS ON THE	

DASHBOARD	146
12.8. MANAGING CEPH OBJECT GATEWAY BUCKET POLICIES ON THE DASHBOARD	147
12.8.1. Creating and editing Ceph Object Gateway bucket policies on the dashboard	148
12.8.2. Deleting Ceph Object Gateway bucket policies on the dashboard	149
12.9. MANAGING S3 BUCKET LIFECYCLE POLICIES ON THE DASHBOARD	149
12.9.1. Applying and viewing S3 bucket lifecycle policies on the dashboard	150
12.9.2. Deleting S3 bucket lifecycle policies on the dashboard	152
12.10. MANAGEMENT OF BUCKETS OF A MULTI-SITE OBJECT CONFIGURATION ON THE CEPH DASHBOARD	152
12.10.1. Monitoring buckets of a multi-site object	153
12.10.2. Editing buckets of a multi-site Object Gateway configuration on the Ceph Dashboard	153
12.10.3. Deleting buckets of a multi-site Object Gateway configuration on the Ceph Dashboard	154
12.11. CONFIGURING A MULTI-SITE OBJECT GATEWAY ON THE CEPH DASHBOARD	156
CHAPTER 13. MANAGING FILE SYSTEMS USING THE CEPH DASHBOARD	160
13.1. CONFIGURING CEPHFS VOLUMES	160
13.1.1. Creating CephFS volumes	160
13.1.2. Editing CephFS volumes	161
13.1.3. Removing CephFS volumes	161
13.2. CONFIGURING CEPHFS SUBVOLUME GROUPS	161
13.2.1. Creating CephFS subvolume groups	161
13.2.2. Editing CephFS subvolume groups	162
13.2.3. Removing CephFS subvolume groups	162
13.3. CONFIGURING CEPHFS SUBVOLUMES	163
13.3.1. Creating CephFS subvolume	163
13.3.2. Editing CephFS subvolume	164
13.3.3. Removing CephFS subvolume	165
13.4. MANAGING CEPHFS SNAPSHOTS	165
13.4.1. Creating CephFS subvolume snapshots	166
13.4.2. Deleting CephFS subvolume snapshots	167
13.4.3. Cloning CephFS subvolume snapshots	167
13.4.4. Creating CephFS volume snapshots	168
13.4.5. Deleting CephFS volume snapshots	169
13.5. SCHEDULING CEPHFS SNAPSHOTS	169
13.5.1. Creating CephFS snapshot schedule	169
13.5.2. Editing CephFS snapshot schedule	170
13.5.3. Deleting CephFS snapshot schedule	171
13.5.4. Deactivating and activating CephFS snapshot schedule	171
CHAPTER 14. MANAGING BLOCK DEVICES USING THE CEPH DASHBOARD	173
14.1. MANAGING BLOCK DEVICE IMAGES ON THE CEPH DASHBOARD	173
14.1.1. Creating images on the Ceph dashboard	173
14.1.2. Creating namespaces on the Ceph dashboard	174
14.1.3. Editing images on the Ceph dashboard	175
14.1.4. Copying images on the Ceph dashboard	176
14.1.5. Moving images to trash on the Ceph dashboard	178
14.1.6. Purging trash on the Ceph dashboard	178
14.1.7. Restoring images from trash on the Ceph dashboard	179
14.1.8. Deleting images on the Ceph Dashboard	180
14.1.9. Deleting namespaces on the Ceph dashboard.	180
14.1.10. Creating snapshots of images on the Ceph dashboard	181
14.1.11. Renaming snapshots of images on the Ceph dashboard	182
14.1.12. Protecting snapshots of images on the Ceph dashboard	183

14.1.13. Cloning snapshots of images on the Ceph dashboard	183
14.1.14. Copying snapshots of images on the Ceph dashboard	184
14.1.15. Unprotecting snapshots of images on the Ceph dashboard	185
14.1.16. Rolling back snapshots of images on the Ceph dashboard	185
14.1.17. Deleting snapshots of images on the Ceph dashboard	186
14.2. MANAGING MIRRORING FUNCTIONS ON THE CEPH DASHBOARD	187
14.2.1. Mirroring view on the Ceph dashboard	187
14.2.2. Editing mode of pools on the Ceph dashboard	188
14.2.3. Adding peer in mirroring on the Ceph dashboard	189
14.2.4. Editing peer in mirroring on the Ceph dashboard	192
14.2.5. Deleting peer in mirroring on the Ceph dashboard	193
CHAPTER 15. ACTIVATING AND DEACTIVATING TELEMETRY	195
15.1. DEACTIVATING TELEMETRY	196

CHAPTER 1. CEPH DASHBOARD OVERVIEW

As a storage administrator, the Red Hat Ceph Storage Dashboard provides management and monitoring capabilities, allowing you to administer and configure the cluster, as well as visualize information and performance statistics related to it. The dashboard uses a web server hosted by the **ceph-mgr** daemon.

The dashboard is accessible from a web browser and includes many useful management and monitoring features, for example, to configure manager modules and monitor the state of OSDs.

The Ceph dashboard provides the following features:

Multi-user and role management

The dashboard supports multiple user accounts with different permissions and roles. User accounts and roles can be managed using both, the command line and the web user interface. The dashboard supports various methods to enhance password security. Password complexity rules may be configured, requiring users to change their password after the first login or after a configurable time period.

For more information, see [Managing roles on the Ceph Dashboard](#) and [Managing users on the Ceph dashboard](#).

Single Sign-On (SSO)

The dashboard supports authentication with an external identity provider using the SAML 2.0 protocol.

For more information, see [Enabling single sign-on for the Ceph dashboard](#).

Auditing

The dashboard backend can be configured to log all PUT, POST and DELETE API requests in the Ceph manager log.

For more information about using the manager modules with the dashboard, see [Viewing and editing the manager modules of the Ceph cluster on the dashboard](#).

Management features

The Red Hat Ceph Storage Dashboard includes various management features.

Viewing cluster hierarchy

You can view the CRUSH map, for example, to determine which host a specific OSD ID is running on. This is helpful if an issue with an OSD occurs.

For more information, see [Viewing the CRUSH map of the Ceph cluster on the dashboard](#).

Configuring manager modules

You can view and change parameters for Ceph manager modules.

For more information, see [Viewing and editing the manager modules of the Ceph cluster on the dashboard](#).

Embedded Grafana dashboards

Ceph Dashboard Grafana dashboards might be embedded in external applications and web pages to surface information with Prometheus modules gathering the performance metrics.

For more information, see [Ceph Dashboard components](#).

Viewing and filtering logs

You can view event and audit cluster logs and filter them based on priority, keyword, date, or time range.

For more information, see [Filtering logs of the Ceph cluster on the dashboard](#).

Toggling dashboard components

You can enable and disable dashboard components so only the features you need are available.

For more information, see [Toggling Ceph dashboard features](#).

Managing OSD settings

You can set cluster-wide OSD flags using the dashboard. You can also Mark OSDs up, down or out, purge and reweight OSDs, perform scrub operations, modify various scrub-related configuration options, select profiles to adjust the level of backfilling activity. You can set and change the device class of an OSD, display and sort OSDs by device class. You can deploy OSDs on new drives and hosts.

For more information, see [Managing Ceph OSDs on the dashboard](#).

Viewing alerts

The alerts page allows you to see details of current alerts.

For more information, see [Viewing alerts on the Ceph dashboard](#).

Upgrading

You can upgrade the Ceph cluster version using the dashboard.

For more information, see [Upgrading a cluster](#).

Quality of service for images

You can set performance limits on images, for example limiting IOPS or read BPS burst rates.

For more information, see [Managing block device images on the Ceph dashboard](#).

Monitoring features

Monitor different features from within the Red Hat Ceph Storage Dashboard.

Username and password protection

You can access the dashboard only by providing a configurable username and password.

For more information, see [Managing users on the Ceph dashboard](#).

Overall cluster health

Displays performance and capacity metrics. This also displays the overall cluster status, storage utilization, for example, number of objects, raw capacity, usage per pool, a list of pools and their status and usage statistics.

For more information, see [Viewing and editing the configuration of the Ceph cluster on the dashboard](#).

Hosts

Provides a list of all hosts associated with the cluster along with the running services and the installed Ceph version.

For more information, see [Monitoring hosts of the Ceph cluster on the dashboard](#).

Performance counters

Displays detailed statistics for each running service.

For more information, see [Monitoring services of the Ceph cluster on the dashboard](#) .

Monitors

Lists all Monitors, their quorum status and open sessions.

For more information, see [Monitoring monitors of the Ceph cluster on the dashboard](#) .

Configuration editor

Displays all the available configuration options, their descriptions, types, default, and currently set values. These values are editable.

For more information, see [Viewing and editing the configuration of the Ceph cluster on the dashboard](#).

Cluster logs

Displays and filters the latest updates to the cluster's event and audit log files by priority, date, or keyword.

For more information, see [Filtering logs of the Ceph cluster on the dashboard](#) .

Device management

Lists all hosts known by the Orchestrator. Lists all drives attached to a host and their properties.

Displays drive health predictions, SMART data, and blink enclosure LEDs.

For more information, see [Monitoring hosts of the Ceph cluster on the dashboard](#) .

View storage cluster capacity

You can view raw storage capacity of the Red Hat Ceph Storage cluster in the *Capacity* pages of the Ceph dashboard.

For more information, see [Understanding the landing page of the Ceph dashboard](#) .

Pools

Lists and manages all Ceph pools and their details. For example: applications, placement groups, replication size, EC profile, quotas, and CRUSH ruleset.

For more information, see [Understanding the landing page of the Ceph dashboard](#) and [Monitoring pools of the Ceph cluster on the dashboard](#).


OSDs

Lists and manages all OSDs, their status, and usage statistics. **OSDs** also lists detailed information, for example, attributes, OSD map, metadata, and performance counters for read and write operations. **OSDs** also lists all drives that are associated with an OSD.

For more information, see [Monitoring Ceph OSDs on the dashboard](#) .

Images

Lists all Ceph Block Device (RBD) images and their properties such as size, objects, and features. Create, copy, modify and delete RBD images. Create, delete, and rollback snapshots of selected images, protect or unprotect these snapshots against modification. Copy or clone snapshots, flatten cloned images.



NOTE

The performance graph for I/O changes in the *Overall Performance* tab for a specific image shows values only after specifying the pool that includes that image by setting the **rbd_stats_pool** parameter in **Cluster→Manager modules→Prometheus**.

For more information, see [Monitoring block device images on the Ceph dashboard](#).

Block device mirroring

Enables and configures Ceph Block Device (RBD) mirroring to a remote Ceph server. Lists all active sync daemons and their status, pools and RBD images including their synchronization state.

For more information, see [Mirroring view on the Ceph dashboard](#).

Ceph File Systems

Lists all active Ceph File System (CephFS) clients and associated pools, including their usage statistics. Evict active CephFS clients, manage CephFS quotas and snapshots, and browse a CephFS directory structure.

For more information, see [Monitoring Ceph file systems on the dashboard](#).

Object Gateway (RGW)

Lists all active object gateways and their performance counters. Displays and manages, including add, edit, and delete, Ceph Object Gateway users and their details, for example quotas, as well as the users' buckets and their details, for example, owner or quotas.

For more information, see [Monitoring Ceph Object Gateway daemons on the dashboard](#).

NFS

Manages NFS exports of CephFS and Ceph object gateway S3 buckets using the NFS Ganesha.

For more information, see [Managing NFS Ganesha exports on the Ceph dashboard](#).

Security features

The dashboard provides the following security features.

SSL and TLS support

All HTTP communication between the web browser and the dashboard is secured via SSL. A self-signed certificate can be created with a built-in command, but it is also possible to import custom certificates signed and issued by a Certificate Authority (CA).

For more information, see [Ceph Dashboard installation and access](#).

Prerequisites

- System administrator level experience.

1.1. CEPH DASHBOARD COMPONENTS

The functionality of the dashboard is provided by multiple components.

- The Cephadm application for deployment.
- The embedded dashboard **ceph-mgr** module.
- The embedded Prometheus **ceph-mgr** module.
- The Prometheus time-series database.
- The Prometheus node-exporter daemon, running on each host of the storage cluster.


- The Grafana platform to provide monitoring user interface and alerting.

Additional Resources

- For more information, see the [Prometheus website](#).
- For more information, see the [Grafana website](#).

1.2. RED HAT CEPH STORAGE DASHBOARD ARCHITECTURE

The Dashboard architecture depends on the Ceph manager dashboard plugin and other components. See the following diagram to understand how the Ceph manager and dashboard work together.



155_Ceph_0421

CHAPTER 2. CEPH DASHBOARD INSTALLATION AND ACCESS

As a system administrator, you can access the dashboard with the credentials provided on bootstrapping the cluster.

Cephadm installs the dashboard by default. Following is an example of the dashboard URL:

```
URL: https://host01:8443/
User: admin
Password: zbiql951ar
```



NOTE

Update the browser and clear the cookies prior to accessing the dashboard URL.

The following are the Cephadm bootstrap options that are available for the Ceph dashboard configurations:

- `[-initial-dashboard-user INITIAL_DASHBOARD_USER]` - Use this option while bootstrapping to set initial-dashboard-user.
- `[-initial-dashboard-password INITIAL_DASHBOARD_PASSWORD]` - Use this option while bootstrapping to set initial-dashboard-password.
- `[-ssl-dashboard-port SSL_DASHBOARD_PORT]` - Use this option while bootstrapping to set custom dashboard port other than default 8443.
- `[-dashboard-key DASHBOARD_KEY]` - Use this option while bootstrapping to set Custom key for SSL.
- `[-dashboard-crt DASHBOARD_CRT]` - Use this option while bootstrapping to set Custom certificate for SSL.
- `[-skip-dashboard]` - Use this option while bootstrapping to deploy Ceph without dashboard.
- `[-dashboard-password-noupdate]` - Use this option while bootstrapping if you used above two options and don't want to reset password at the first time login.
- `[-allow-fqdn-hostname]` - Use this option while bootstrapping to allow hostname that is fully-qualified.
- `[-skip-prepare-host]` - Use this option while bootstrapping to skip preparing the host.



NOTE

To avoid connectivity issues with dashboard related external URL, use the fully qualified domain names (FQDN) for hostnames, for example, **host01.ceph.redhat.com**.



NOTE

Open the Grafana URL directly in the client internet browser and accept the security exception to see the graphs on the Ceph dashboard. Reload the browser to view the changes.

Example

```
[root@host01 ~]# cephadm bootstrap --mon-ip 127.0.0.1 --registry-json cephadm.txt --initial-
dashboard-user admin --initial-dashboard-password zbiql951ar --dashboard-password-noupdate --
allow-fqdn-hostname
```



NOTE

While bootstrapping the storage cluster using **cephadm**, you can use the **--image** option for either custom container images or local container images.



NOTE

You have to change the password the first time you log into the dashboard with the credentials provided on bootstrapping only if **--dashboard-password-noupdate** option is not used while bootstrapping. You can find the Ceph dashboard credentials in the **var/log/ceph/cephadm.log** file. Search with the "Ceph Dashboard is now available at" string.

This section covers the following tasks:

- [Network port requirements for Ceph dashboard](#) .
- [Accessing the Ceph dashboard](#) .
- [Expanding the cluster on the Ceph dashboard](#) .
- [Upgrading a cluster](#) .
- [Toggling Ceph dashboard features](#) .
- [Understanding the landing page of the Ceph dashboard](#) .
- [Enabling Red Hat Ceph Storage Dashboard manually](#) .
- [Changing the dashboard password using the Ceph dashboard](#) .
- [Changing the Ceph dashboard password using the command line interface](#) .
- [Setting **admin** user password for Grafana](#) .
- [Creating an admin account for syncing users to the Ceph dashboard](#) .
- [Syncing users to the Ceph dashboard using the Red Hat Single Sign-On](#) .
- [Enabling single sign-on for the Ceph dashboard](#) .
- [Disabling single sign-on for the Ceph dashboard](#) .

2.1. NETWORK PORT REQUIREMENTS FOR CEPH DASHBOARD

The Ceph dashboard components use certain TCP network ports which must be accessible. By default, the network ports are automatically opened in **firewalld** during installation of Red Hat Ceph Storage.

Table 2.1. TCP Port Requirements

Port	Use	Originating Host	Destination Host
8443	The dashboard web interface	IP addresses that need access to Ceph Dashboard UI and the host under Grafana server, since the AlertManager service can also initiate connections to the Dashboard for reporting alerts.	The Ceph Manager hosts.
3000	Grafana	IP addresses that need access to Grafana Dashboard UI and all Ceph Manager hosts and Grafana server.	The host or hosts running Grafana server.
2049	NFS-Ganesha	IP addresses that need access to NFS.	The IP addresses that provide NFS services.
9095	Default Prometheus server for basic Prometheus graphs	IP addresses that need access to Prometheus UI and all Ceph Manager hosts and Grafana server or Hosts running Prometheus.	The host or hosts running Prometheus.
9093	Prometheus Alertmanager	IP addresses that need access to Alertmanager Web UI and all Ceph Manager hosts and Grafana server or Hosts running Prometheus.	All Ceph Manager hosts and the host under Grafana server.
9094	Prometheus Alertmanager for configuring a highly available cluster made from multiple instances	All Ceph Manager hosts and the host under Grafana server.	Prometheus Alertmanager High Availability (peer daemon sync), so both src and dst should be hosts running Prometheus Alertmanager.

Port	Use	Originating Host	Destination Host
9100	The Prometheus node-exporter daemon	Hosts running Prometheus that need to view Node Exporter metrics Web UI and All Ceph Manager hosts and Grafana server or Hosts running Prometheus.	All storage cluster hosts, including MONs, OSDs, Grafana server host.
9283	Ceph Manager Prometheus exporter module	Hosts running Prometheus that need access to Ceph Exporter metrics Web UI and Grafana server.	All Ceph Manager hosts.

Additional Resources

- For more information, see the [Red Hat Ceph Storage Installation Guide](#).
- For more information, see [Using and configuring firewalls](#) in [Configuring and managing networking](#).

2.2. ACCESSING THE CEPH DASHBOARD

You can access the Ceph dashboard to administer and monitor your Red Hat Ceph Storage cluster.

Prerequisites

- Successful installation of Red Hat Ceph Storage Dashboard.
- NTP is synchronizing clocks properly.

Procedure

1. Enter the following URL in a web browser:

Syntax

```
https://HOST_NAME:PORT
```

Replace:

- *HOST_NAME* with the fully qualified domain name (FQDN) of the active manager host.
- *PORT* with port **8443**

Example

```
https://host01:8443
```

You can also get the URL of the dashboard by running the following command in the Cephadm shell:


Example

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

This command will show you all endpoints that are currently configured. Look for the **dashboard** key to obtain the URL for accessing the dashboard.

2. On the login page, enter the username **admin** and the default password provided during bootstrapping.
3. You have to change the password the first time you log in to the Red Hat Ceph Storage dashboard.
4. After logging in, the dashboard default landing page is displayed, which provides details, a high-level overview of status, performance, inventory, and capacity metrics of the Red Hat Ceph Storage cluster.

Figure 2.1. Ceph dashboard landing page



5. Click the menu icon (☰) on the dashboard landing page to collapse or display the options in the vertical menu.

Additional Resources

- For more information, see [Changing the dashboard password using the Ceph dashboard](#) in the *Red Hat Ceph Storage Dashboard guide*.

2.3. EXPANDING THE CLUSTER ON THE CEPH DASHBOARD

You can use the dashboard to expand the Red Hat Ceph Storage cluster for adding hosts, adding OSDs, and creating services such as Alertmanager, Cephadm-exporter, CephFS-mirror, Grafana, ingress, MDS, NFS, node-exporter, Prometheus, RBD-mirror, and Ceph Object Gateway.

Once you bootstrap a new storage cluster, the Ceph Monitor and Ceph Manager daemons are created and the cluster is in *HEALTH_WARN* state. After creating all the services for the cluster on the dashboard, the health of the cluster changes from *HEALTH_WARN* to *HEALTH_OK* status.

Prerequisites

- Bootstrapped storage cluster. See [Bootstrapping a new storage cluster](#) section in the *Red Hat Ceph Storage Installation Guide* for more details.
- At least **cluster-manager** role for the user on the Red Hat Ceph Storage Dashboard. See the [User roles and permissions on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

Procedure

1. Copy the admin key from the bootstrapped host to other hosts:

Syntax

```
ssh-copy-id -f -i /etc/ceph/ceph.pub root@HOST_NAME
```

Example

```
[ceph: root@host01 /]# ssh-copy-id -f -i /etc/ceph/ceph.pub root@host02  
[ceph: root@host01 /]# ssh-copy-id -f -i /etc/ceph/ceph.pub root@host03
```

2. Log in to the dashboard with the default credentials provided during bootstrap.
3. Change the password and log in to the dashboard with the new password .
4. On the landing page, click **Expand Cluster**.



NOTE

Clicking **Expand Cluster** opens a wizard taking you through the expansion steps. To skip and add hosts and services separately, click **Skip**.

Figure 2.2. Expand cluster



5. Add hosts. This needs to be done for each host in the storage cluster.
 - a. In the **Add Hosts** step, click **Add**.
 - b. Provide the hostname. This is same as the hostname that was provided while copying the key from the bootstrapped host.

**NOTE**

Add multiple hosts by using a comma-separated list of host names, a range expression, or a comma separated range expression.

- c. Optional: Provide the respective IP address of the host.
 - d. Optional: Select the labels for the hosts on which the services are going to be created. Click the pencil icon to select or add new labels.
 - e. Click **Add Host**.
The new host is displayed in the Add Hosts pane.
 - f. Click **Next**.
6. Create OSDs:
 - a. In the Create OSDs step, for Primary devices, Click **Add**.
 - b. In the **Primary Devices** window, filter for the device and select the device.
 - c. Click **Add**.
 - d. Optional: In the **Create OSDs** window, if you have any shared devices such as WAL or DB devices, then add the devices.

e. Optional: In the **Features** section, select **Encryption** to encrypt the features.

f. Click **Next**.

7. Create services:

a. In the Create Services step, click **Create**.

b. In the **Create Service** form:

i. Select a service type.

ii. Provide the service ID. The ID is a unique name for the service. This ID is used in the service name, which is **service_type.service_id**.

...Optional: Select if the service is **Unmanaged**.

+ When **Unmanaged** services is selected, the orchestrator will not start or stop any daemon associated with this service. Placement and all other properties are ignored.

i. Select if the placement is by hosts or label.

ii. Select the hosts.

iii. In the **Count** field, provide the number of daemons or services that need to be deployed.

a. Click **Create Service**.


The new service is displayed in the Create Services pane.

1. In the **Create Service** window, Click **Next**.

2. Review the cluster expansion details.

Review the **Cluster Resources**, **Hosts by Services**, **Host Details**. To edit any parameters, click **Back** and follow the previous steps.

Figure 2.3. Review cluster



3. Click **Expand Cluster**.

The **Cluster expansion displayed** notification is displayed and the cluster status changes to *HEALTH_OK* on the dashboard.

Verification

1. Log in to the **cephadm** shell:

Example

Example

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

2. Run the **ceph -s** command.

Example

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

The health of the cluster is *HEALTH_OK*.

Additional Resources

- See the [User roles and permissions on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.
- See the [Red Hat Ceph Storage Installation Guide](#) for more details.

2.4. UPGRADING A CLUSTER

Upgrade Ceph clusters using the dashboard.

Cluster images are pulled automatically from **registry.redhat.io**. Optionally, use custom images for upgrade.

Procedure

1. View if cluster upgrades are available and upgrade as needed from **Administration > Upgrade** on the dashboard.

**NOTE**

If the dashboard displays the **Not retrieving upgrades** message, check if the registries were added to the container configuration files with the appropriate log in credentials to Podman or docker.

Click **Pause** or **Stop** during the upgrade process, if needed. The upgrade progress is shown in the progress bar along with information messages during the upgrade.

**NOTE**

When stopping the upgrade, the upgrade is first paused and then prompts you to stop the upgrade.

2. Optional. View cluster logs during the upgrade process from the **Cluster logs** section of the **Upgrade** page.
3. Verify that the upgrade is completed successfully by confirming that the cluster status displays OK state.

2.5. TOGGING CEPH DASHBOARD FEATURES

You can customize the Red Hat Ceph Storage dashboard components by enabling or disabling features on demand. All features are enabled by default. When disabling a feature, the web-interface elements become hidden and the associated REST API end-points reject any further requests for that feature. Enabling and disabling dashboard features can be done from the command-line interface or the web interface.

Available features:

- Ceph Block Devices:
 - Image management, **rbd**
 - Mirroring, **mirroring**
- Ceph File System, **cephfs**
- Ceph Object Gateway, **rgw**
- NFS Ganesha gateway, **nfs**



NOTE

By default, the Ceph Manager is collocated with the Ceph Monitor.



NOTE

You can disable multiple features at once.



IMPORTANT

Once a feature is disabled, it can take up to 20 seconds to reflect the change in the web interface.

Prerequisites

- Installation and configuration of the Red Hat Ceph Storage dashboard software.
- User access to the Ceph Manager host or the dashboard web interface.
- Root level access to the Ceph Manager host.

Procedure

- To toggle the dashboard features from the dashboard web interface:
 - a. On the dashboard landing page, go to **Administration→Manager Modules** and select the dashboard module.
 - b. Click **Edit**.
 - c. In the *Edit Manager module* form, you can enable or disable the dashboard features by selecting or clearing the check boxes next to the different feature names.
 - d. After the selections are made, click **Update**.
- To toggle the dashboard features from the command-line interface:

- a. Log in to the Cephadm shell:

Example

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

- b. List the feature status:

Example

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

- c. Disable a feature:

```
[ceph: root@host01 /]# ceph dashboard feature disable rgw
```

This example disables the Ceph Object Gateway feature.

- d. Enable a feature:

```
[ceph: root@host01 /]# ceph dashboard feature enable cephfs
```

This example enables the Ceph Filesystem feature.

2.6. UNDERSTANDING THE LANDING PAGE OF THE CEPH DASHBOARD

The landing page displays an overview of the entire Ceph cluster using navigation bars and individual panels.

The menu bar provides the following options:

Tasks and Notifications

Provides task and notification messages.

Help

Provides links to the product and REST API documentation, details about the Red Hat Ceph Storage Dashboard, and a form to report an issue.

Dashboard Settings

Gives access to user management and telemetry configuration.

User

Use this menu to see log in status, to change a password, and to sign out of the dashboard.

Figure 2.4. Menu bar



The navigation menu can be opened or hidden by clicking the navigation menu icon .


Dashboard

The main dashboard displays specific information about the state of the cluster.

The main dashboard can be accessed at any time by clicking **Dashboard** from the navigation menu.

The dashboard landing page organizes the panes into different categories.

Figure 2.5. Ceph dashboard landing page



Details

Displays specific cluster information and if telemetry is active or inactive.

Inventory

Displays the different parts of the cluster, how many are available, and their status. Link directly from **Inventory** to specific inventory items, where available.

Hosts

Displays the total number of hosts in the Ceph storage cluster.

Monitors

Displays the number of Ceph Monitors and the quorum status.

Managers

Displays the number and status of the Manager Daemons.

OSDs

Displays the total number of OSDs in the Ceph Storage cluster and the number that are *up*, and *in*.

Pools

Displays the number of storage pools in the Ceph cluster.

PGs

Displays the total number of placement groups (PGs). The PG states are divided into *Working* and *Warning* to simplify the display. Each one encompasses multiple states. + The *Working* state includes PGs with any of the following states:

- activating
- backfill_wait

- backfilling
- creating
- deep
- degraded
- forced_backfill
- forced_recovery
- peering
- peered
- recovering
- recovery_wait
- repair
- scrubbing
- snaptrim
- snaptrim_wait + The *Warning* state includes PGs with any of the following states:
- backfill_toofull
- backfill_unfound
- down
- incomplete
- inconsistent
- recovery_toofull
- recovery_unfound
- remapped
- snaptrim_error
- stale
- undersized

Object Gateways

Displays the number of Object Gateways in the Ceph storage cluster.

Metadata Servers

Displays the number and status of metadata servers for Ceph File Systems (CephFS).

Status

Displays the health of the cluster and host and daemon states. The current health status of the Ceph storage cluster is displayed. Danger and warning alerts are displayed directly on the landing page. Click **View alerts** for a full list of alerts.

Capacity

Displays storage usage metrics. This is displayed as a graph of used, warning, and danger. The numbers are in percentages and in GiB.

Cluster Utilization

The **Cluster Utilization** pane displays information related to data transfer speeds. Select the time range for the data output from the list. Select a range between the last 5 minutes to the last 24 hours.

Used Capacity (RAW)

Displays usage in GiB.

IOPS

Displays total I/O read and write operations per second.

OSD Latencies

Displays total applies and commits per millisecond.

Client Throughput

Displays total client read and write throughput in KiB per second.

Recovery Throughput

Displays the rate of cluster healing and balancing operations. For example, the status of any background data that may be moving due to a loss of disk is displayed. The information is displayed in bytes per second.

Additional Resources

- For more information, see [Monitoring the cluster on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide* for more information.

2.7. CHANGING THE DASHBOARD PASSWORD USING THE CEPH DASHBOARD

By default, the password for accessing dashboard is randomly generated by the system while bootstrapping the cluster. You have to change the password the first time you log in to the Red Hat Ceph Storage dashboard. You can change the password for the **admin** user using the dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.

Procedure

1. Log in to the dashboard:

Syntax

```
https://HOST_NAME:8443
```

2. Go to **User→Change password** on the menu bar.

3. Enter the old password, for verification.
4. In the **New password** field enter a new password. Passwords must contain a minimum of 8 characters and cannot be the same as the last one.
5. In the **Confirm password** field, enter the new password again to confirm.
6. Click **Change Password**.
You will be logged out and redirected to the login screen. A notification appears confirming the password is changed.

2.8. CHANGING THE CEPH DASHBOARD PASSWORD USING THE COMMAND LINE INTERFACE

If you have forgotten your Ceph dashboard password, you can change the password using the command line interface.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to the host on which the dashboard is installed.

Procedure

1. Log into the Cephadm shell:

Example

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

2. Create the **dashboard_password.yml** file:

Example

```
[ceph: root@host01 /]# touch dashboard_password.yml
```

3. Edit the file and add the new dashboard password:

Example

```
[ceph: root@host01 /]# vi dashboard_password.yml
```

4. Reset the dashboard password:

Syntax

```
ceph dashboard ac-user-set-password DASHBOARD_USERNAME -i PASSWORD_FILE
```

Example

```
[ceph: root@host01 /]# ceph dashboard ac-user-set-password admin -i
```

```
dashboard_password.yml
{"username": "admin", "password":
"$2b$12$i5RmvN1PolR61Fay0mPgt.GDpcga1QpYsaHUbJfoqaHd1rFFx7XS", "roles":
["administrator"], "name": null, "email": null, "lastUpdate": , "enabled": true,
"pwdExpirationDate": null, "pwdUpdateRequired": false}
```

Verification

- Log in to the dashboard with your new password.

2.9. SETTING ADMIN USER PASSWORD FOR GRAFANA

By default, **cephadm** does not create an admin user for Grafana. With the Ceph Orchestrator, you can create an admin user and set the password.

With these credentials, you can log in to the storage cluster's Grafana URL with the given password for the admin user.

Prerequisites

- A running Red Hat Ceph Storage cluster with the monitoring stack installed.
- Root-level access to the **cephadm** host.
- The **dashboard** module enabled.

Procedure

1. As a root user, create a **grafana.yml** file and provide the following details:

Syntax

```
service_type: grafana
spec:
  initial_admin_password: PASSWORD
```

Example

```
service_type: grafana
spec:
  initial_admin_password: mypassword
```

2. Mount the **grafana.yml** file under a directory in the container:

Example

```
[root@host01 ~]# cephadm shell --mount grafana.yml:/var/lib/ceph/grafana.yml
```



NOTE

Every time you exit the shell, you have to mount the file in the container before deploying the daemon.

- Optional: Check if the **dashboard** Ceph Manager module is enabled:

Example

```
[ceph: root@host01 /]# ceph mgr module ls
```

- Optional: Enable the **dashboard** Ceph Manager module:

Example

```
[ceph: root@host01 /]# ceph mgr module enable dashboard
```

- Apply the specification using the **orch** command:

Syntax

```
ceph orch apply -i FILE_NAME.yml
```

Example

```
[ceph: root@host01 /]# ceph orch apply -i /var/lib/ceph/grafana.yml
```

- Redeploy **grafana** service:

Example

```
[ceph: root@host01 /]# ceph orch redeploy grafana
```

This creates an admin user called **admin** with the given password and the user can log in to the Grafana URL with these credentials.

Verification:

- Log in to Grafana with the credentials:

Syntax

```
https://HOST_NAME:PORT
```

Example

```
https://host01:3000/
```

2.10. ENABLING RED HAT CEPH STORAGE DASHBOARD MANUALLY

If you have installed a Red Hat Ceph Storage cluster by using **--skip-dashboard** option during bootstrap, you can see that the dashboard URL and credentials are not available in the bootstrap output. You can enable the dashboard manually using the command-line interface. Although the monitoring stack components such as Prometheus, Grafana, Alertmanager, and node-exporter are deployed, they are disabled and you have to enable them manually.

Prerequisite

- A running Red Hat Ceph Storage cluster installed with **--skip-dashboard** option during bootstrap.
- Root-level access to the host on which the dashboard needs to be enabled.

Procedure

1. Log into the Cephadm shell:

Example

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

2. Check the Ceph Manager services:

Example

```
[ceph: root@host01 /]# ceph mgr services  
  
{  
  "prometheus": "http://10.8.0.101:9283/"  
}
```

You can see that the Dashboard URL is not configured.

3. Enable the dashboard module:

Example

```
[ceph: root@host01 /]# ceph mgr module enable dashboard
```

4. Create the self-signed certificate for the dashboard access:

Example

```
[ceph: root@host01 /]# ceph dashboard create-self-signed-cert
```



NOTE

You can disable the certificate verification to avoid certification errors.

5. Check the Ceph Manager services:

Example

```
[ceph: root@host01 /]# ceph mgr services  
  
{  
  "dashboard": "https://10.8.0.101:8443/",  
  "prometheus": "http://10.8.0.101:9283/"  
}
```

6. Create the admin user and password to access the Red Hat Ceph Storage dashboard:

Syntax

```
echo -n "PASSWORD" > PASSWORD_FILE
ceph dashboard ac-user-create admin -i PASSWORD_FILE administrator
```

Example

```
[ceph: root@host01 /]# echo -n "p@ssw0rd" > password.txt
[ceph: root@host01 /]# ceph dashboard ac-user-create admin -i password.txt administrator
```

7. Enable the monitoring stack. See the [Enabling monitoring stack](#) section in the *Red Hat Ceph Storage Dashboard Guide* for details.

Additional Resources

- See the [Deploying the monitoring stack using the Ceph Orchestrator](#) section in the *Red Hat Ceph Storage Operations Guide*.

2.11. USING SINGLE SIGN-ON WITH THE DASHBOARD

The Ceph Dashboard supports external authentication of users with the choice of either the Security Assertion Markup Language (SAML) 2.0 protocol or with the OAuth2 Proxy (**oauth2-proxy**). Before using single sign-on (SSO) with the Ceph Dashboard, create the dashboard user accounts and assign any required roles. The Ceph Dashboard completes user authorization and then the existing Identity Provider (*IdP*) completes the authentication process. You can enable single sign-on using the SAML protocol or **oauth2-proxy**.

Red Hat Ceph Storage supports dashboard SSO and Multi-Factor Authentication with RHSSO (Keycloak).

OAuth2 SSO uses the **oauth2-proxy** service to work with the Ceph Management gateway (**mgmt-gateway**), providing unified access and improved user experience.



NOTE

The OAuth2 SSO, **mgmt-gateway**, and **oauth2-proxy** services are Technology Preview.

For more information about the Ceph Management gateway and the OAuth2 Proxy service, see [Using the Ceph Management gateway \(mgmt-gateway\)](#) and [Using the OAuth2 Proxy \(oauth2-proxy\) service](#).

For more information about Red Hat build of Keycloak, see [Red Hat build of Keycloak](#) on the [Red Hat Customer Portal](#).

2.11.1. Creating an admin account for syncing users to the Ceph dashboard

You have to create an admin account to synchronize users to the Ceph dashboard.

After creating the account, use Red Hat Single Sign-on (SSO) to synchronize users to the Ceph dashboard. See the [Syncing users to the Ceph dashboard using Red Hat Single Sign-On](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin level access to the dashboard.
- Users are added to the dashboard.
- Root-level access on all the hosts.
- Java OpenJDK installed. For more information, see the [Installing a JRE on RHEL by using yum](#) section of the *Installing and using OpenJDK 8 for RHEL* guide for OpenJDK on the Red Hat Customer Portal.
- Red hat Single Sign-On installed from a ZIP file. See the [Installing RH-SSO from a ZIP File](#) section of the *Server Installation and Configuration Guide* for Red Hat Single Sign-On on the Red Hat Customer Portal.

Procedure

1. Download the [Red Hat Single Sign-On 7.4.0 Server](#) on the system where Red Hat Ceph Storage is installed.

2. Unzip the folder:

```
[root@host01 ~]# unzip rhssso-7.4.0.zip
```

3. Navigate to the **standalone/configuration** directory and open the **standalone.xml** for editing:

```
[root@host01 ~]# cd standalone/configuration  
[root@host01 configuration]# vi standalone.xml
```

4. From the **bin** directory of the newly created **rhssso-7.4.0** folder, run the **add-user-keycloak** script to add the initial administrator user:

```
[root@host01 bin]# ./add-user-keycloak.sh -u admin
```

5. Replace all instances of **localhost** and two instances of **127.0.0.1** with the IP address of the machine where Red Hat SSO is installed.

6. Start the server. From the **bin** directory of **rh-ssso-7.4** folder, run the **standalone** boot script:

```
[root@host01 bin]# ./standalone.sh
```

7. Create the admin account in `https://IP_ADDRESS:8080/auth` with a username and password:



NOTE

You have to create an admin account only the first time that you log into the console.

8. Log into the admin console with the credentials created.

Additional Resources

- For adding roles for users on the dashboard, see the [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more information.
- For creating users on the dashboard, see the [Creating users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

2.11.2. Syncing users to the Ceph dashboard using Red Hat Single Sign-On

You can use Red Hat Single Sign-on (SSO) with Lightweight Directory Access Protocol (LDAP) integration to synchronize users to the Red Hat Ceph Storage Dashboard.

The users are added to specific realms in which they can access the dashboard through SSO without any additional requirements of a password.

Prerequisites


- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin level access to the dashboard.
- Users are added to the dashboard. See the [Creating users on Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.
- Root-level access on all the hosts.
- Admin account created for syncing users. See the [Creating an admin account for syncing users to the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

Procedure


1. To create a realm, click the *Master* drop-down menu. In this realm, you can provide access to users and applications.
2. In the *Add Realm* window, enter a case-sensitive realm name and set the parameter *Enabled* to ON and click *Create*:

3. In the *Realm Settings* tab, set the following parameters and click *Save*:

- a. Enabled - ON
- b. User-Managed Access - ON
- c. Make a note of the link address of SAML 2.0 Identity Provider Metadata to paste in *Client Settings*.



4. In the *Clients* tab, click *Create*:



5. In the *Add Client* window, set the following parameters and click *Save*:

- a. Client ID - BASE_URL:8443/auth/saml2/metadata

Example

`https://example.ceph.redhat.com:8443/auth/saml2/metadata`

- b. Client Protocol - saml

6. In the *Client* window, under *Settings* tab, set the following parameters:

Table 2.2. Client Settings tab

Name of the parameter	Syntax	Example
Client ID	BASE_URL:8443/auth/saml2/metadata	<code>https://example.ceph.redhat.com:8443/auth/saml2/metadata</code>
Enabled	ON	ON

Name of the parameter	Syntax	Example
Client Protocol	saml	saml
Include AuthnState ment	ON	ON
Sign Documents	ON	ON
Signature Algorithm	RSA_SHA1	RSA_SHA1
SAML Signature Key Name	KEY_ID	KEY_ID
Valid Redirect URLs	BASE_URL:8443/*	https://example.ceph.redhat.com:8443/*
Base URL	BASE_URL:8443	https://example.ceph.redhat.com:8443/
Master SAML Processing URL	https://localhost:8080/auth/realms/ <i>REALM_NAME</i> /protocol/saml/descriptor	https://localhost:8080/auth/realms/Ceph_LDAP/protocol/saml/descriptor

**NOTE**

Paste the link of SAML 2.0 Identity Provider Metadata from *Realm Settings* tab.

Under Fine Grain SAML Endpoint Configuration, set the following parameters and click Save:

Table 2.3. Fine Grain SAML configuration

Name of the parameter	Syntax	Example
Assertion Consumer Service POST Binding URL	BASE_URL:8443/#!/dashboard	https://example.ceph.redhat.com:8443/#!/dashboard
Assertion Consumer Service Redirect Binding URL	BASE_URL:8443/#!/dashboard	https://example.ceph.redhat.com:8443/#!/dashboard

Name of the parameter	Syntax	Example
Logout Service Redirect Binding URL	BASE_URL:8443/	https://example.ceph.redhat.com:8443/

7. In the *Clients* window, *Mappers* tab, set the following parameters and click *Save*:

Table 2.4. Client Mappers tab

Name of the parameter	Value
Protocol	saml
Name	username
Mapper Property	User Property
Property	username
SAML Attribute name	username

8. In the *Clients Scope* tab, select *role_list*:
- In *Mappers* tab, select *role list*, set the *Single Role Attribute* to ON.
9. Select *User_Federation* tab:
- In *User Federation* window, select *ldap* from the drop-down menu:
 - In *User_Federation* window, *Settings* tab, set the following parameters and click *Save*:

Table 2.5. User Federation Settings tab

Name of the parameter	Value
Console Display Name	rh-ldap
Import Users	ON
Edit_Mode	READ_ONLY
Username LDAP attribute	username
RDN LDAP attribute	username
UUID LDAP attribute	nsuniqueid
User Object Classes	inetOrgPerson

Name of the parameter	Value
organizationalPerson	rhatPerson
Connection URL	Example: ldap://ldap.corp.redhat.com Click <i>Test Connection</i> . You will get a notification that the LDAP connection is successful.
Users DN	ou=users, dc=example, dc=com
Bind Type	simple

Click *Test authentication*. You will get a notification that the LDAP authentication is successful.


c. In *Mappers* tab, select *first name* row and edit the following parameter and Click *Save*:

- LDAP Attribute - givenName

d. In *User_Federation* tab, *Settings* tab, Click *Synchronize all users*:

You will get a notification that the sync of users is finished successfully.

10. In the *Users* tab, search for the user added to the dashboard and click the Search icon:



RED HAT SINGLE SIGN-ON

Ceph_LDAP

Configure

- Realm Settings
- Clients
- Client Scopes
- Roles
- Identity Providers
- User Federation
- Authentication

Manage

- Groups
- Users**
- Sessions


Users

Lookup

View all users

ID	Username
0edc54ea-8a2b-4d1d-815a-e894e97...	

11. To view the user, click the specific row. You should see the federation link as the name provided for the *User Federation*.



RED HAT SINGLE SIGN-ON

Ceph_LDAP

Configure

- Realm Settings
- Clients
- Client Scopes
- Roles
- Identity Providers
- User Federation
- Authentication

Manage

- Groups
- Users**
- Sessions

Users > [User ID]

Details Attributes Credentials Role Mappings Groups Consents

ID [Redacted]

Created At 9/11/20 5:32:37 PM

Username [Redacted]

Email [Redacted]

First Name [Redacted]

Last Name [Redacted]

User Enabled ☒ ON

Federation Link [rh-ldap](#)



IMPORTANT

Do not add users manually as the users will not be synchronized by LDAP. If added manually, delete the user by clicking *Delete*.



NOTE

If Red Hat SSO is currently being used within your work environment, be sure to first enable SSO. For more information, see the [Enabling Single Sign-On with SAML 2.0 for the Ceph Dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

Verification

- Users added to the realm and the dashboard can access the Ceph dashboard with their mail address and password.

Example

`https://example.ceph.redhat.com:8443`

Additional Resources

- For adding roles for users on the dashboard, see the [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

2.11.3. Enabling Single Sign-On with SAML 2.0 for the Ceph Dashboard

The Ceph Dashboard supports external authentication of users with the Security Assertion Markup Language (SAML) 2.0 protocol. Before using single sign-on (SSO) with the Ceph dashboard, create the dashboard user accounts and assign the desired roles. The Ceph Dashboard performs authorization of the users and the authentication process is performed by an existing Identity Provider (IdP). You can enable single sign-on using the SAML protocol.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Installation of the Ceph Dashboard.
- Root-level access to The Ceph Manager hosts.

Procedure

1. To configure SSO on Ceph Dashboard, run the following command:

Syntax

```
cephadm shell CEPH_MGR_HOST ceph dashboard sso setup saml2
CEPH_DASHBOARD_BASE_URL IDP_METADATA IDP_USERNAME_ATTRIBUTE
IDP_ENTITY_ID SP_X_509_CERT SP_PRIVATE_KEY
```

Example

```
[root@host01 ~]# cephadm shell host01 ceph dashboard sso setup saml2
https://dashboard_hostname.ceph.redhat.com:8443 idp-metadata.xml username
https://10.70.59.125:8080/auth/realms/realm_name /home/certificate.txt /home/private-key.txt
```

Replace

- `CEPH_MGR_HOST` with Ceph **mgr** host. For example, **host01**
- `CEPH_DASHBOARD_BASE_URL` with the base URL where Ceph Dashboard is accessible.
- `IDP_METADATA` with the URL to remote or local path or content of the IdP metadata XML. The supported URL types are http, https, and file.
- **Optional:** `IDP_USERNAME_ATTRIBUTE` with the attribute used to get the username from the authentication response. Defaults to `uid`.
- **Optional:** `IDP_ENTITY_ID` with the IdP entity ID when more than one entity ID exists on the IdP metadata.
- **Optional:** `SP_X_509_CERT` with the file path of the certificate used by Ceph Dashboard for signing and encryption.
- **Optional:** `SP_PRIVATE_KEY` with the file path of the private key used by Ceph Dashboard for signing and encryption.

2. Verify the current SAML 2.0 configuration:

Syntax

```
cephadm shell CEPH_MGR_HOST ceph dashboard sso show saml2
```

Example

```
[root@host01 ~]# cephadm shell host01 ceph dashboard sso show saml2
```

3. To enable SSO, run the following command:

Syntax

```
cephadm shell CEPH_MGR_HOST ceph dashboard sso enable saml2  
SSO is "enabled" with "SAML2" protocol.
```

Example

```
[root@host01 ~]# cephadm shell host01 ceph dashboard sso enable saml2
```

4. Open your dashboard URL.

Example

```
https://dashboard_hostname.ceph.redhat.com:8443
```

5. On the SSO page, enter the login credentials. SSO redirects to the dashboard web interface.

Additional Resources

- To disable single sign-on, see [Disabling Single Sign-on for the Ceph Dashboard](#) in the *Red Hat Ceph Storage Dashboard Guide*.

2.11.4. Enabling OAuth2 single sign-on (Technology Preview)

Enable OAuth2 single sign-on (SSO) for the Ceph Dashboard. OAuth2 SSO uses the `oauth2-proxy` service.



IMPORTANT

Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend using them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. See the support scope for Red Hat Technology Preview features for more details.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Installation of the Ceph Dashboard.
- Root-level access to the Ceph Manager hosts.
- An admin account with Red Hat Single-Sign on 7.6.0. For more information, see [Creating an admin account with Red Hat Single Sign-On 7.6.0](#).
- Enable the Ceph Management gateway (**mgmt-gateway**) service. For more information, see [Enabling the Ceph Management gateway](#).
- Enable the OAuth2 Proxy service (**oauth2-proxy**). For more information, see [Enabling the OAuth2 Proxy service](#).

Procedure

1. Enable Ceph Dashboard OAuth2 SSO access.

Syntax

```
ceph dashboard sso enable oauth2
```

Example

```
[ceph: root@host01 /]# ceph dashboard sso enable oauth2
SSO is "enabled" with "oauth2" protocol.
```

2. Set the valid redirect URL.

Syntax

```
https://HOST_NAME|IP_ADDRESS/oauth2/callback
```

**NOTE**

This URL must be the same redirect URL as configured in the OAuth2 Proxy service.

3. Configure a valid user role.

**NOTE**

For the Administrator role, configure the IDP user with an administrator or read-only access.

4. Open your dashboard URL.

Example

```
https://dashboard_hostname.ceph.redhat.com:8443
```

5. On the SSO page, enter the login credentials. The SSO redirects to the dashboard web interface.

Verification

Check the SSO status at any time with the **cephadm shell ceph dashboard sso status** command.

Example

```
[root@host01 ~]# cephadm shell ceph dashboard sso status
SSO is "enabled" with "oauth2" protocol.
```

2.11.5. Disabling Single Sign-On for the Ceph Dashboard

You can disable SAML 2.0 and OAuth2 SSO for the Ceph Dashboard at any time.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Installation of the Ceph Dashboard.
- Root-level access to The Ceph Manager hosts.
- Single sign-on enabled for Ceph Dashboard

Procedure

1. To view status of SSO, run the following command:

Syntax

```
cephadm shell CEPH_MGR_HOST ceph dashboard sso status
```

Example

```
[root@host01 ~]# cephadm shell host01 ceph dashboard sso status
SSO is "enabled" with "SAML2" protocol.
```

2. To disable SSO, run the following command:

Syntax

```
cephadm shell CEPH_MGR_HOST ceph dashboard sso disable
SSO is "disabled".
```

Example

```
[root@host01 ~]# cephadm shell host01 ceph dashboard sso disable
```

Additional Resources

- To enable single sign-on, see [Enabling Single Sign-On with SAML 2.0 for the Ceph Dashboard](#) in the *Red Hat Ceph StorageDashboard Guide*.

CHAPTER 3. MANAGING ROLES ON THE CEPH DASHBOARD

As a storage administrator, you can create, edit, clone, and delete roles on the dashboard.

By default, there are eight system roles. You can create custom roles and give permissions to those roles. These roles can be assigned to users based on the requirements.

This section covers the following administrative tasks:

- [User roles and permissions on the Ceph dashboard](#).
- [Creating roles on the Ceph dashboard](#).
- [Editing roles on the Ceph dashboard](#).
- [Cloning roles on the Ceph dashboard](#).
- [Deleting roles on the Ceph dashboard](#).

3.1. USER ROLES AND PERMISSIONS ON THE CEPH DASHBOARD

User accounts are associated with a set of roles that define the specific dashboard functionality which can be accessed. View user roles and permissions by going to **Dashboard settings→User management**.

The Red Hat Ceph Storage dashboard functionality or modules are grouped within a security scope. Security scopes are predefined and static. The current available **security scopes** on the Red Hat Ceph Storage dashboard are:

- **cephfs**: Includes all features related to CephFS management.
- **config-opt**: Includes all features related to management of Ceph configuration options.
- **dashboard-settings**: Allows to edit the dashboard settings.
- **grafana**: Include all features related to Grafana proxy.
- **hosts**: Includes all features related to the Hosts menu entry.
- **log**: Includes all features related to Ceph logs management.
- **manager**: Includes all features related to Ceph manager management.
- **monitor**: Includes all features related to Ceph monitor management.
- **nfs-ganesha**: Includes all features related to NFS-Ganesha management.
- **osd**: Includes all features related to OSD management.
- **pool**: Includes all features related to pool management.
- **prometheus**: Include all features related to Prometheus alert management.
- **rbd-image**: Includes all features related to RBD image management.
- **rbd-mirroring**: Includes all features related to RBD mirroring management.
- **rgw**: Includes all features related to Ceph object gateway (RGW) management.

A role specifies a set of mappings between a security scope and a set of permissions. There are four types of **permissions**:

- **Read**
- **Create**
- **Update**
- **Delete**

User management / Users

Users Roles						
Search Refresh Export Create +						
Username	↑	Name	Email	Roles	Enabled	Password expires
admin		-	-	administrator	✓	
Items per page: 10 1-1 of 1 item 1 of 1 page						

The list of **system roles** are:

- **administrator**: Allows full permissions for all security scopes.
- **block-manager**: Allows full permissions for RBD-image and RBD-mirroring scopes.
- **cephfs-manager**: Allows full permissions for the Ceph file system scope.
- **cluster-manager**: Allows full permissions for the hosts, OSDs, monitor, manager, and config-opt scopes.
- **ganesha-manager**: Allows full permissions for the NFS-Ganesha scope.
- **pool-manager**: Allows full permissions for the pool scope.
- **read-only**: Allows read permission for all security scopes except the dashboard settings and config-opt scopes.
- **rgw-manager**: Allows full permissions for the Ceph object gateway scope.

User management / Roles

Users Roles			
Search Refresh Export Create +			
Name	↑	Description	System Role
administrator		allows full permissions for all security scopes	✓
block-manager		allows full permissions for rbd-image, rbd-mirroring, and iscsi scopes	✓
cephfs-manager		allows full permissions for the cephfs scope	✓
cluster-manager		allows full permissions for the hosts, osd, mon, mgr, and config-opt scopes	✓
ganesha-manager		allows full permissions for the nfs-ganesha scope	✓
pool-manager		allows full permissions for the pool scope	✓
read-only		allows read permission for all security scope except dashboard settings and config-opt	✓
rgw-manager		allows full permissions for the rgw scope	✓
Items per page: 10 1-8 of 8 items 1 of 1 page			

For example, you need to provide **rgw-manager** access to the users for all Ceph object gateway operations.

Additional Resources

- For creating users on the Ceph dashboard, see [Creating users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For creating roles on the Ceph dashboard, see [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.

3.2. CREATING ROLES ON THE CEPH DASHBOARD


You can create custom roles on the dashboard and these roles can be assigned to users based on their roles.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.

Procedure

1. Log in to the Dashboard.
2. Click the *Dashboard Settings* icon and then click *User management*.



3. On *Roles* tab, click *Create*.
4. In the *Create Role* window, set the *Name*, *Description*, and select the *Permissions* for this role, and then click the *Create Role* button.

User management / [Roles](#) / Create

Create Role

Name (required)

NFS-RGW

Description

With this role you can manage NFS Ganesha and Object Gateway operations

Permissions

☐ All	☐ Read	☐ Create	☐ Update	☐ Delete
☐ cephfs	☐	☐	☐	☐
☐ config-opt	☐	☐	☐	☐
☐ dashboard-settings	☐	☐	☐	☐
☒ grafana	☒	☒	☒	☒
☐ hosts	☐	☐	☐	☐
☐ iscsi	☐	☐	☐	☐
☐ log	☐	☐	☐	☐
☐ manager	☐	☐	☐	☐
☐ monitor	☐	☐	☐	☐
☒ nfs-ganesha	☒	☒	☒	☒
☐ nvme-of	☐	☐	☐	☐
☐ osd	☐	☐	☐	☐
☐ pool	☐	☐	☐	☐
☐ prometheus	☐	☐	☐	☐
☐ rbd-image	☐	☐	☐	☐
☐ rbd-mirroring	☐	☐	☐	☐
☒ rgw	☒	☒	☒	☒
☐ user	☐	☐	☐	☐
Items per page: ▼	1-18 of 18 items			1 ▼ of 1 page ◀ ▶

Cancel

Create Role

In this example, the user assigned with **ganesha-manager** and **rgw-manager** roles can manage all NFS-Ganesha gateway and Ceph object gateway operations.

- You get a notification that the role was created successfully.
- Click on the *Expand/Collapse* icon of the row to view the details and permissions given to the roles.

Additional Resources

- See the [User roles and permissions on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.
- See the [Creating users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

3.3. EDITING ROLES ON THE CEPH DASHBOARD


The dashboard allows you to edit roles on the dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.
- A role is created on the dashboard.

Procedure

1. Log in to the Dashboard.
2. Click the *Dashboard Settings* icon and then click *User management*.



3. On *Roles* tab, click the role you want to edit.
4. In the *Edit Role* window, edit the parameters, and then click *Edit Role*.

User management / [Roles](#) / Edit

Edit Role

Name (required)

NFS-RGW

Description

With this role you can manage NFS Ganesha and Object Gateway operations

Permissions

☐ All	☐ Read	☐ Create	☐ Update	☐ Delete
☐ cephfs	☐	☐	☐	☐
☐ config-opt	☐	☐	☐	☐
☐ dashboard-settings	☐	☐	☐	☐
☐ grafana	☒	☐	☐	☐
☐ hosts	☐	☐	☐	☐
☐ iscsi	☐	☐	☐	☐
☐ log	☐	☐	☐	☐
☐ manager	☐	☐	☐	☐
☐ monitor	☐	☐	☐	☐
☒ nfs-ganesha	☒	☒	☒	☒
☐ nvme-of	☐	☐	☐	☐
☐ osd	☐	☐	☐	☐
☐ pool	☐	☐	☐	☐
☐ prometheus	☐	☐	☐	☐
☐ rbd-image	☐	☐	☐	☐
☐ rbd-mirroring	☐	☐	☐	☐
☒ rgw	☒	☒	☒	☒
☐ user	☐	☐	☐	☐
Items per page: ▼	1-18 of 18 items			1 ▼ of 1 page ◀ ▶

Cancel

Edit Role

- You get a notification that the role was updated successfully.

Additional Resources

- See the [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

3.4. CLONING ROLES ON THE CEPH DASHBOARD

When you want to assign additional permissions to existing roles, you can clone the system roles and edit it on the Red Hat Ceph Storage Dashboard.


Prerequisites

- A running Red Hat Ceph Storage cluster.


- Dashboard is installed.
- Admin-level access to the dashboard.
- Roles are created on the dashboard.

Procedure

1. Log in to the Dashboard.
2. Click the *Dashboard Settings* icon and then click *User management*.



3. On *Roles* tab, click the role you want to clone.
4. Select *Clone* from the *Edit* drop-down menu.
5. In the *Clone Role* dialog box, enter the details for the role, and then click *Clone Role*.



6. Once you clone the role, you can customize the permissions as per the requirements.

Additional Resources

- See the [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

3.5. DELETING ROLES ON THE CEPH DASHBOARD

You can delete the custom roles that you have created on the Red Hat Ceph Storage dashboard.



NOTE


You cannot delete the system roles of the Ceph Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.
- A custom role is created on the dashboard.

Procedure

1. Log in to the Dashboard.
2. Click the **Dashboard Settings** icon and then select **User management**.



3. On the **Roles** tab, click the role you want to delete and select **Delete** from the action drop-down.
4. In the **Delete Role** notification, select **Yes, I am sure** and click **Delete Role**.

Additional Resources

- See the [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

CHAPTER 4. MANAGING USERS ON THE CEPH DASHBOARD

As a storage administrator, you can create, edit, and delete users with specific roles on the Red Hat Ceph Storage dashboard. Role-based access control is given to each user based on their roles and the requirements.

You can also create, edit, import, export, and delete Ceph client authentication keys on the dashboard. Once you create the authentication keys, you can rotate keys using command-line interface (CLI). Key rotation meets the current industry and security compliance requirements.

This section covers the following administrative tasks:

- [Creating users on the Ceph dashboard](#).
- [Editing users on the Ceph dashboard](#).
- [Deleting users on the Ceph dashboard](#).
- [User capabilities](#)
- [Access capabilities](#)
- [Creating user capabilities](#)
- [Editing user capabilities](#)
- [Importing user capabilities](#)
- [Exporting user capabilities](#)
- [Deleting user capabilities](#)

4.1. CREATING USERS ON THE CEPH DASHBOARD

You can create users on the Red Hat Ceph Storage dashboard with adequate roles and permissions based on their roles. For example, if you want the user to manage Ceph object gateway operations, then you can give **rgw-manager** role to the user.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.




NOTE

The Red Hat Ceph Storage Dashboard does not support any email verification when changing a users password. This behavior is intentional, because the Dashboard supports Single Sign-On (SSO) and this feature can be delegated to the SSO provider.

Procedure

1. Log in to the Dashboard.

- Click the *Dashboard Settings* icon and then click *User management*.



- On *Users* tab, click *Create*.
- In the *Create User* window, set the *Username* and other parameters including the roles, and then click *Create User*.

User management / Users / Create

Create User

Username *

johnd

✓

Password ⓘ

.....

✓

👁

Confirm password

.....

✓

👁

Password expiration date ⓘ

Password expiration date...

✓

✕

Full name

John Doe

✓

Email

johnd@example.com

✓

Roles

✎ There are no roles.

☒ Enabled

☒ User must change password at next logon

Cancel

Create User

- You get a notification that the user was created successfully.

Additional Resources

- See the [Creating roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.
- See the [User roles and permissions on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

4.2. EDITING USERS ON THE CEPH DASHBOARD


You can edit the users on the Red Hat Ceph Storage dashboard. You can modify the user's password and roles based on the requirements.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.
- User created on the dashboard.

Procedure

1. Log in to the Dashboard.
2. Click the *Dashboard Settings* icon and then click *User management*.



3. To edit the user, click the row.
4. On *Users* tab, select *Edit* from the *Edit* drop-down menu.
5. In the *Edit User* window, edit parameters like password and roles, and then click *Edit User*.

User management / Users / Edit

Edit User

Username

johnd

Password ⓘ

Password...

👁

Confirm password

Confirm password...

👁

Password expiration date ⓘ

Password expiration date...

✕

Full name

John Doe

Email

johnd@example.com

Roles

✎

read-only ✕

☒ Enabled

☒ User must change password at next login

Cancel

Edit User 🖱

**NOTE**

If you want to disable any user's access to the Ceph dashboard, you can uncheck *Enabled* option in the *Edit User* window.

6. You get a notification that the user was created successfully.

Additional Resources

- See the [Creating users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

4.3. DELETING USERS ON THE CEPH DASHBOARD


You can delete users on the Ceph dashboard. Some users might be removed from the system. The access to such users can be deleted from the Ceph dashboard.

Prerequisites


- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.
- User created on the dashboard.

Procedure

1. Log in to the Dashboard.
2. Click the *Dashboard Settings* icon and then click *User management*.



3. On *Users* tab, click the user you want to delete.
4. select *Delete* from the *Edit* drop-down menu.
5. In the **Delete User** notification, select **Yes, I am sure** and click **Delete User**.



Additional Resources

- See the [Creating users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

4.4. USER CAPABILITIES

Ceph stores data RADOS objects within pools irrespective of the Ceph client used. Ceph users must have access to a given pool to read and write data, and must have executable permissions to use Ceph administrative's commands. Creating users allows you to control their access to your Red Hat Ceph Storage cluster, its pools, and the data within the pools.

Ceph has a concept of **type** of user which is always **client**. You need to define the user with the **TYPE.ID** where ID is the user ID, for example, **client.admin**. This user typing is because the Cephx protocol is used not only by clients but also non-clients, such as Ceph Monitors, OSDs, and Metadata Servers. Distinguishing the user type helps to distinguish between client users and other users. This distinction streamlines access control, user monitoring, and traceability.

4.4.1. Capabilities

Ceph uses capabilities (caps) to describe the permissions granted to an authenticated user to exercise the functionality of the monitors, OSDs, and metadata servers. The capabilities restrict access to data within a pool, a namespace within a pool, or a set of pools based on their applications tags. A Ceph administrative user specifies the capabilities of a user when creating or updating the user.

You can set the capabilities to monitors, managers, OSDs, and metadata servers.

- The Ceph Monitor capabilities include **r**, **w**, and **x** access settings. These can be applied in aggregate from pre-defined profiles with **profile NAME**.
- The OSD capabilities include **r**, **w**, **x**, **class-read**, and **class-write** access settings. These can be applied in aggregate from pre-defined profiles with **profile NAME**.
- The Ceph Manager capabilities include **r**, **w**, and **x** access settings. These can be applied in aggregate from pre-defined profiles with **profile NAME**.
- For administrators, the metadata server (MDS) capabilities include **allow ***.

**NOTE**

The Ceph Object Gateway daemon (**radosgw**) is a client of the Red Hat Ceph Storage cluster and is not represented as a Ceph storage cluster daemon type.

Additional Resources

- See [Access capabilities](#) for more details.

4.5. ACCESS CAPABILITIES

This section describes the different access or entity capabilities that can be given to a Ceph user or a Ceph client such as Block Device, Object Storage, File System, and native API.

Additionally, you can describe the capability profiles while assigning roles to clients.

allow, Description

Precedes access settings for a daemon. Implies **rw** for MDS only

r, Description

Gives the user *read* access. Required with monitors to retrieve the CRUSH map.

w, Description

Gives the user *write* access to objects.

x, Description

Gives the user the capability to call class methods, that is, both *read* and *write*, and to conduct **auth** operations on monitors.

class-read, Description

Gives the user the capability to call class read methods. Subset of **x**.

class-write, Description

Gives the user the capability to call class write methods. Subset of **x**.

***, all, Description**

Gives the user *read*, *write*, and *execute* permissions for a particular daemon or a pool, as well as the ability to execute admin commands.

The following entries describe valid capability profile:

profile osd, Description

This is applicable to Ceph Monitor only. Gives a user permissions to connect as an OSD to other OSDs or monitors. Conferred on OSDs to enable OSDs to handle replication heartbeat traffic and status reporting.

profile mds, Description

This is applicable to Ceph Monitor only. Gives a user permissions to connect as an MDS to other MDSs or monitors.

profile bootstrap-osd, Description

This is applicable to Ceph Monitor only. Gives a user permissions to bootstrap an OSD. Conferred on deployment tools, such as **ceph-volume** and **cephadm**, so that they have permissions to add keys when bootstrapping an OSD.

profile bootstrap-mds, Description

This is applicable to Ceph Monitor only. Gives a user permissions to bootstrap a metadata server. Conferred on deployment tools, such as **cephadm**, so that they have permissions to add keys when bootstrapping a metadata server.

profile bootstrap-rbd, Description

This is applicable to Ceph Monitor only. Gives a user permissions to bootstrap an RBD user. Conferred on deployment tools, such as **cephadm**, so that they have permissions to add keys when bootstrapping an RBD user.

profile bootstrap-rbd-mirror, Description

This is applicable to Ceph Monitor only. Gives a user permissions to bootstrap an **rbd-mirror** daemon user. Conferred on deployment tools, such as **cephadm**, so that they have permissions to add keys when bootstrapping an **rbd-mirror** daemon.

profile rbd, Description

This is applicable to Ceph Monitor, Ceph Manager, and Ceph OSDs. Gives a user permissions to manipulate RBD images. When used as a Monitor cap, it provides the user with the minimal privileges required by an RBD client application; such privileges include the ability to blocklist other client users. When used as an OSD cap, it provides an RBD client application with read-write access to the specified pool. The Manager cap supports optional **pool** and **namespace** keyword arguments.

profile rbd-mirror, Description

This is applicable to Ceph Monitor only. Gives a user permissions to manipulate RBD images and retrieve RBD mirroring config-key secrets. It provides the minimal privileges required for the user to manipulate the **rbd-mirror** daemon.

profile rbd-read-only, Description

This is applicable to Ceph Monitor and Ceph OSDs. Gives a user read-only permissions to RBD images. The Manager cap supports optional **pool** and **namespace** keyword arguments.

profile simple-rados-client, Description

This is applicable to Ceph Monitor only. Gives a user read-only permissions for monitor, OSD, and PG data. Intended for use by direct librados client applications.

profile simple-rados-client-with-blocklist, Description

This is applicable to Ceph Monitor only. Gives a user read-only permissions for monitor, OSD, and PG data. Intended for use by direct librados client applications. Also includes permissions to add blocklist entries to build high-availability (HA) applications.

profile fs-client, Description

This is applicable to Ceph Monitor only. Gives a user read-only permissions for monitor, OSD, PG, and MDS data. Intended for CephFS clients.

profile role-definer, Description

This is applicable to Ceph Monitor and Auth. Gives user **all** permissions for the auth subsystem, read-only access to monitors, and nothing else. Useful for automation tools. **WARNING:** Do not assign this unless you really, know what you are doing, as the security ramifications are substantial and pervasive.

profile crash, Description

This is applicable to Ceph Monitor and Ceph Manager. Gives a user read-only access to monitors. Used in conjunction with the manager crash module to upload daemon **crash** dumps into monitor storage for later analysis.

Additional Resources

- See [User capabilities_](#) for more details.

4.6. CREATING USER CAPABILITIES

Create role-based access users with different capabilities on the Ceph dashboard.

For details on different user capabilities, see [User capabilities](#) and [Access capabilities](#)

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.

Procedure

1. From the dashboard navigation, go to **Administration→Ceph Users**.
2. Click **Create**.
3. In the **Create User** form, provide the following details:
 - a. **User entity:** Enter as **TYPE.ID**.
 - b. **Entity:** This can be **mon**, **mgr**, **osd**, or **mds**.
 - c. **Entity Capabilities:** Enter the capabilities that you can to provide to the user. For example, 'allow *' and **profile crash** are some of the capabilities that can be assigned to the client.



NOTE

You can add more entities to the user, based on the requirement.

Create User

User entity *

client.test1 ✓

Capabilities

Entity *	Entity Capabilities *		
mon ✓	allow * ✓	+	🗑️
mds ✓	allow * ✓	+	🗑️
osd ✓	allow r ✓	+	🗑️
mgr ✓	allow * ✓	+	🗑️

Cancel

Create User

4. Click **Create User**.

A notification displays that the user is created successfully.

4.7. EDITING USER CAPABILITIES

Edit the roles of users or clients on the dashboard.

For details on different user capabilities, see [User capabilities](#) and [Access capabilities](#)


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.

Procedure

1. From the dashboard navigation, go to **Administration→Ceph Users**.
2. Select the user whose roles you want to edit.

3. Click **Edit**.
4. In the **Edit User** form, edit the **Entity** and **Entity Capabilities**, as needed.

**NOTE**

You can add more entities to the user based on the requirement.

5. Click **Edit User**.
A notification displays that the user is successfully edited.

4.8. IMPORTING USER CAPABILITIES

Import the roles of users or clients from the the local host to the client on the dashboard.

For details on different user capabilities, see [User capabilities](#) and [Access capabilities](#)

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.

Procedure

1. Create a keyring file on the local host:

Example

```
[localhost:~]$ cat import.keyring

[client.test11]
key = AQD9S29kmjgJFxAakvhFar6Af3AWKDY2DsULRg==
caps mds = "allow *"
caps mgr = "allow *"
caps mon = "allow *"
caps osd = "allow r"
```

2. From the dashboard navigation, go to **Administration→Ceph Users**.
3. Select the user whose roles you want to export.
4. Select **Edit→Import**.
5. In the **Import User** form, click **Choose File**.
6. Browse to the file on your local host and select.
7. Click **Import User**.

Cluster » Users

Import User

User file import ⓘ

Choose File

test.keyring

Cancel

Import User

A notification displays that the keys are successfully imported.

4.9. EXPORTING USER CAPABILITIES

Export the roles of the users or clients from the dashboard to a the local host.


For details on different user capabilities, see [User capabilities](#) and [Access capabilities](#)

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.

Procedure

1. From the dashboard navigation, go to **Administration→Ceph Users**.
2. Select the user whose roles you want to export.
3. Select **Export** from the action drop-down.
4. From the **Ceph user export data** dialog, click **Copy to Clipboard**.



A notification displays that the keys are successfully copied.

- On your local system, create a keyring file and paste the keys:

Example

```
[localhost:~]$ cat exported.keyring

[client.test11]
key = AQD9S29kmjgJFxAkvhFar6Af3AWKDY2DsULRg==
caps mds = "allow *"
caps mgr = "allow *"
caps mon = "allow *"
caps osd = "allow r"
```

- Click **Close**.

4.10. DELETING USER CAPABILITIES

Delete the roles of users or clients on the dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level access to the dashboard.

Procedure


1. From the dashboard navigation, go to **Administration→Ceph Users**.
2. Select the user that you want to delete and select **Delete** from the action drop-down.
3. In the **Delete user** dialog, select **Yes, I am sure..**
4. Click **Delete user**.
A notification displays that the user is deleted successfully.

CHAPTER 5. MANAGING CEPH DAEMONS

As a storage administrator, you can manage Ceph daemons on the Red Hat Ceph Storage dashboard.

5.1. DAEMON ACTIONS

The Red Hat Ceph Storage dashboard allows you to start, stop, restart, and redeploy daemons.



NOTE

These actions are supported on all daemons except monitor and manager daemons.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- At least one daemon is configured in the storage cluster.

Procedure

You can manage daemons two ways.

From the Services page:

1. From the dashboard navigation, go to **Administration→Services**.
2. Expand the service with the daemon that will have the action run on.



NOTE

The row can be collapsed at any time.

3. On the **Daemons** tab, select the row with the daemon.



NOTE

The **Daemons** table can be searched and filtered.

4. Select the action that needs to be run on the daemon. The options are **Start**, **Stop**, **Restart**, and **Redeploy**.

Figure 5.1. Managing daemons from Services

Administration / Services

Search

Service	Placement	Running	Last Refreshed
alertmanager	count:1	1 / 1	3 minutes ago
ceph-exporter	*	3 / 3	3 minutes ago

DaemonsService Events

Search

HostnameAny

Hostname	Daemon name	Version	Status	Last Refreshed	CPU Usage	Memory Usage	Daemon Events
ceph-vpap-i80-9bx9bs-node1-installer	ceph-exporter.ceph-vpap-i80-9bx9bs-node1-installer	19.2.0-44.el9cp	running	3 minutes ago	2%	13.2 MiB	
ceph-vpap-i80-9bx9bs-node2	ceph-exporter.ceph-vpap-i80-9bx9bs-node2	19.2.0-44.el9cp	running	3 minutes ago	0%	5.1 MiB	
ceph-vpap-i80-9bx9bs-node3	ceph-exporter.ceph-vpap-i80-9bx9bs-node3	19.2.0-44.el9cp	running	3 minutes ago	0%	5.1 MiB	

Items per page: 101-3 of 3 items

crash* 3 / 3 | 3 minutes ago | |

StartStopRestartRedeploy

From the Hosts page:

1. From the dashboard navigation, go to **Cluster→Hosts**.

2. On the **Hosts List** tab, expand the host row and select the host with the daemon to perform the action on.

3. On the **Daemons** tab of the host, select the row with the daemon.

NOTE

The **Daemons** table can be searched and filtered.

4. Select the action that needs to be run on the daemon. The options are **Start**, **Stop**, **Restart**, and **Redeploy**.

Figure 5.2. Managing daemons from Hosts

Cluster / Hosts

Hosts ListOverall Performance

Search

Expand ClusterAdd

Hostname	Service Instances	Labels	Status	Model	CPUs	Cores	Total Memory	Raw Capacity	HDDs	Flash	NICs
ceph-vpap-i80-9bx9bs-node1-installer (10.0.66.213)	mon: 1mgr: 1ceph-exporter: 1crash: 1node-exporter: 1grafana: 1prometheus: 1alertmanager: 1osd: 6	_admin, prometheus, osd, mgr, crash, alertmanager, installer, mon, grafana, node-exporter	available	Virtual Machine (OpenStack Compute)	4	1	7.5 GiB	150 GiB	7	0	1

DevicesPhysical DisksDaemonsPerformance DetailsDevice health

Search

Daemon nameAny

Daemon name	Version	Status	Last Refreshed	CPU Usage	Memory Usage	Daemon Events
alertmanager.ceph-vpap-i80-9bx9bs-node1-installer	0.26.0	running	2 minutes ago	0%	26.3 MiB	
ceph-exporter.ceph-vpap-i80-9bx9bs-node1-installer	19.2.0-44.el9cp	running	2 minutes ago	2%	13.2 MiB	
crash.ceph-vpap-i80-9bx9bs-node1-installer	19.2.0-44.el9cp	running	2 minutes ago	0%	6.7 MiB	
grafana.ceph-vpap-i80-9bx9bs-node1-installer	10.4.8	running	2 minutes ago	1%	83.3 MiB	
mgr.ceph-vpap-i80-9bx9bs-node1-installer.suhhdm	19.2.0-44.el9cp	running	2 minutes ago	5%	633.3 MiB	

StartStopRestartRedeploy

64

CHAPTER 6. MONITORING THE CLUSTER ON THE CEPH DASHBOARD

As a storage administrator, you can use Red Hat Ceph Storage Dashboard to monitor specific aspects of the cluster based on types of hosts, services, data access methods, and more.

This section covers the following administrative tasks:

- [Monitoring hosts of the Ceph cluster on the dashboard](#) .
- [Viewing and editing the configuration of the Ceph cluster on the dashboard](#) .
- [Viewing and editing the manager modules of the Ceph cluster on the dashboard](#) .
- [Monitoring monitors of the Ceph cluster on the dashboard](#) .
- [Monitoring services of the Ceph cluster on the dashboard](#) .
- [Monitoring Ceph OSDs on the dashboard](#) .
- [Monitoring HAProxy on the dashboard](#) .
- [Viewing the CRUSH map of the Ceph cluster on the dashboard](#) .
- [Filtering logs of the Ceph cluster on the dashboard](#) .
- [Viewing centralized logs of the Ceph cluster on the dashboard](#) .
- [Monitoring pools of the Ceph cluster on the dashboard](#).
- [Monitoring Ceph file systems on the dashboard](#).
- [Monitoring Ceph Object Gateway daemons on the dashboard](#).
- [Monitoring block device images on the Ceph dashboard](#).

6.1. MONITORING HOSTS OF THE CEPH CLUSTER ON THE DASHBOARD

You can monitor the hosts of the cluster on the Red Hat Ceph Storage Dashboard.

The following are the different tabs on the hosts page. Each tab contains a table with the relevant information. The tables are searchable and customizable by column and row.

To change the order of the columns, select the column name and drag to place within the table.

To select which columns are displaying, click the **toggle columns** button and select or clear column names.

Enter the number of rows to be displayed in the row selector field.

Devices

This tab has a table that details the device ID, state of the device health, life expectancy, device name, prediction creation date, and the daemons on the hosts.

Physical Disks

This tab has a table that details all disks attached to a selected host, as well as their type, size and others. It has details such as device path, type of device, available, vendor, model, size, and the OSDs deployed. To identify which disk is where on the physical device, select the device and click **Identify**. Select the duration of how long the LED should blink for to find the selected disk.

Daemons

This tab has a table that details all services that have been deployed on the selected host, which container they are running in, and their current status. The table has details such as daemon name, daemon version, status, when the daemon was last refreshed, CPU usage, memory usage (in MiB), and daemon events. Daemon actions can be run from this tab. For more details, see [Daemon actions](#).

Performance Details

This tab has details such as OSDs deployed, CPU utilization, RAM usage, network load, network drop rate, and OSD disk performance statistics. View performance information through the embedded Grafana Dashboard.

Device health

For SMART-enabled devices, you can get the individual health status and SMART data only on the OSD deployed hosts.


Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Hosts are added to the storage cluster.
- All the services, monitor, manager, and OSD daemons are deployed on the storage cluster.

Procedure

1. From the dashboard navigation, go to **Cluster→Hosts**.
2. On the **Hosts List** tab, expand the host row and select the host with the daemon to perform the action on.
3. On the **Daemons** tab of the host, select the row with the daemon.



NOTE

The **Daemons** table can be searched and filtered.

4. Select the action that needs to be run on the daemon. The options are **Start**, **Stop**, **Restart**, and **Redeploy**.

Figure 6.1. Monitoring hosts of the Ceph cluster

The screenshot shows the Ceph dashboard interface. At the top, there's a navigation bar with 'Cluster / Hosts' and tabs for 'Hosts List' and 'Overall Performance'. Below this is a search bar and buttons for 'Add' and 'Expand Cluster'. The main table lists hosts with various metrics. The bottom section, 'Daemons', shows a list of running services like 'alertmanager', 'ceph-exporter', 'crash', and 'grafana'. A context menu is visible over the 'alertmanager' row, offering actions like 'Start', 'Stop', 'Restart', and 'Redeploy'.

Additional Resources

- See the [Ceph performance counters](#) in the *Red Hat Ceph Storage Administration Guide* for more details.

6.2. VIEWING AND EDITING THE CONFIGURATION OF THE CEPH CLUSTER ON THE DASHBOARD

You can view various configuration options of the Ceph cluster on the dashboard. You can edit only some configuration options.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- All the services are deployed on the storage cluster.

Procedure

1. From the dashboard navigation, go to **Administration→Configuration**.
2. To view the details of the configuration, expand the row contents.

Figure 6.2. Configuration options

Administration / Configuration

Search Level basic Edit +

Level: basic X Clear filters

Name	Description	Current value	Default	Editable
client_cache_size	soft maximum number of directory entries in client cache		16384	✓
Name client_cache_size Description soft maximum number of directory entries in client cache Long description Current values Default 16384 Daemon default Type size Min Max Flags Services mds_client Source Level basic Can be updated at runtime (editable) Yes Tags Enum values See also				
cluster_addr	cluster-facing address to bind to		-	
container_image	container image (used by cephadm orchestrator)	global: registry-proxy.engineering.redhat.com/rh-osbs/rhceph@sha256:29f33c984392ee19c48c955de74d8ac5dee7ee74613f01499f55db745cd6fbb	docker.io/ceph/daemon-base:latest-master-devel	

- Optional: Use the search field to find a configuration.
- Optional: You can filter for a specific configuration. Use the following filters:
 - Level** - Basic, advanced, or dev
 - Service** - Any, mon, mgr, osd, mds, common, mds_client, rgw, and similar filters.
 - Source** - Any, mon, and similar filters
 - Modified** - yes or no
- To edit a configuration, select the configuration row and click **Edit**.
 - Use the Edit form to edit the required parameters, and click **Update**.
A notification displays that the configuration was updated successfully.

Additional Resources

- See the [Ceph Network Configuration](#) chapter in the *Red Hat Ceph Storage Configuration Guide* for more details.

6.3. VIEWING AND EDITING THE MANAGER MODULES OF THE CEPH CLUSTER ON THE DASHBOARD

Manager modules are used to manage module-specific configuration settings. For example, you can enable alerts for the health of the cluster.

You can view, enable or disable, and edit the manager modules of a cluster on the Red Hat Ceph Storage dashboard.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.

Viewing the manager modules

1. From the dashboard navigation, go to **Administration→Manager Modules**.
2. To view the details of a specific manager module, expand the row contents.

Figure 6.3. Manager modules

Administrator / Manager Modules

Search

Name	↑ Enabled	Always-On
alerts		
balancer	✓	✓
Key	↑	Value
active		true
begin_time		0000
begin_weekday		0
crush_compat_max_iterations		25
crush_compat_metrics		pgs.objects.bytes
crush_compat_step		0.5
end_time		2359
end_weekday		0
log_level		
log_to_cluster		false
cephadm	✓	

Enabling a manager module

Select the row and click **Enable** from the action drop-down.

Disabling a manager module

Select the row and click **Disable** from the action drop-down.

Editing a manager module

1. Select the row:



NOTE

Not all modules have configurable parameters. If a module is not configurable, the **Edit** button is disabled.

2. Edit the required parameters and click **Update**.
A notification displays that the module was updated successfully.

6.4. MONITORING MONITORS OF THE CEPH CLUSTER ON THE DASHBOARD

You can monitor the performance of the Ceph monitors on the landing page of the Red Hat Ceph Storage dashboard. You can also view the details such as status, quorum, number of open session, and performance counters of the monitors in the **Monitors** panel.

Prerequisites


Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Monitors are deployed in the storage cluster.

Procedure

1. From the dashboard navigation, go to **Cluster→Monitors**.
The **Monitors** panel displays information about the overall monitor status and monitor hosts that are in and out of quorum.

2. To see the number of open sessions, in the **In Quorum** table, hover the cursor over the **Open Sessions**.
In Quorum




Name	Rank	Public Address	Open Sessions
ceph-vpap-80-nfb57h-node1-installer	0	10.0.64.136:6789/0	
ceph-vpap-80-nfb57h-node2	2	10.0.65.146:6789/0	16
ceph-vpap-80-nfb57h-node3	1	10.0.66.231:6789/0	

3. To see performance counters for any monitor, click the **Name** in the **In Quorum** and **Not In Quorum** tables.

Figure 6.4. Viewing monitor Performance Counters

Cluster / Monitors / Performance Counters

mon.ceph-vpap-80-nfb57h-node1-installer



Name	Description	Value
mon.election_call	Elections started	0
mon.election_lose	Elections lost	0
mon.election_win	Elections won	0
mon.num_elections	Elections participated in	0
mon.num_sessions	Open sessions	16
mon.session_add	Created sessions	0
mon.session_rm	Removed sessions	0
mon.session_trim	Trimmed sessions	0
paxos.accept_timeout	Accept timeouts	0
paxos.begin	Started and handled begins	11999109715713614

Additional Resources

- See the [Ceph monitors](#) section in the *Red Hat Ceph Storage Operations guide*.

- See the [Ceph performance counters](#) in the *Red Hat Ceph Storage Administration Guide* for more details.

6.5. MONITORING SERVICES OF THE CEPH CLUSTER ON THE DASHBOARD

You can monitor the services of the cluster on the Red Hat Ceph Storage Dashboard. You can view the details such as hostname, daemon type, daemon ID, container ID, container image name, container image ID, version status and last refreshed time.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Hosts are added to the storage cluster.
- All the services are deployed on the storage cluster.

Procedure

1. From the dashboard navigation, go to **Administration→Services**.
2. Expand the service for more details.

Figure 6.5. Monitoring services of the Ceph cluster

Administration / Services

Search ↻ 📄 Create +

Service	↑ Placement	Running	Last Refreshed
^ alertmanager	count:1	1 / 1	7 minutes ago

Daemons Service Events

Search ▽ Hostname... ▽ Any ↻ 📄

Hostname	Daemon name	↑ Version	Status	Last Refreshed	CPU Usage	Memory Usage	Daemon Events
ceph-vpap-80-nfb57h-node1-installer	alertmanager.ceph-vpap-80-nfb57h-node1-installer	0.26.0	running	7 minutes ago	0%	27 MiB	⋮

Items per page: 10 ▾ 1-1 of 1 item

1 ▾ of 1 page ⏪ ⏩

▽ ceph-exporter * 3 / 3 7 minutes ago ⋮

Additional Resources

- See the [Ceph Orchestrators](#) in the *Red Hat Ceph Storage Operations Guide* for more details.

6.6. MONITORING CEPH OSDS ON THE DASHBOARD

You can monitor the status of the Ceph OSDs on the landing page of the Red Hat Ceph Storage Dashboard. You can also view the details such as host, status, device class, number of placement groups (PGs), size flags, usage, and read or write operations time in the *OSDs* tab.

The following are the different tabs on the OSDs page:

- **Devices** - This tab has details such as Device ID, state of health, life expectancy, device name, and the daemons on the hosts.
- **Attributes (OSD map)** - This tab shows the cluster address, details of heartbeat, OSD state, and the other OSD attributes.
- **Metadata** - This tab shows the details of the OSD object store, the devices, the operating system, and the kernel details.
- **Device health** - For SMART-enabled devices, you can get the individual health status and SMART data.
- **Performance counter** - This tab gives details of the bytes written on the devices.
- **Performance Details** - This tab has details such as OSDs deployed, CPU utilization, RAM usage, network load, network drop rate, and OSD disk performance statistics. View performance information through the embedded Grafana Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Hosts are added to the storage cluster.
- All the services including OSDs are deployed on the storage cluster.

Procedure

1. From the dashboard navigation, go to **Cluster→OSDs**.
2. To view the details of a specific OSD, from the **OSDs List** tab, expand an OSD row.

Figure 6.6. Monitoring OSDs of the Ceph cluster

Cluster » OSDs

OSDs List Overall Performance

+ Create Cluster-wide configuration ▾

ID	Host	Status	Device class	PGs
0	depressa003	in up	ssd	176
1	depressa003	in up	ssd	184
2	depressa003	in up	ssd	184
3	depressa003	in up	ssd	8
4	depressa003	in up	ssd	4

Devices Attributes (OSD map) Metadata Device health Performance counter Performance Details

Device ID State of Health Life Expectancy

INTEL_SSDPE21K375GA_PHKE91360037375AC Unknown

You can view additional details such as **Devices**, **Attributes (OSD map)**, **Metadata**, **Device Health**, **Performance counter**, and **Performance Details**, by clicking on the respective tabs.

Additional Resources

- See the [Ceph Orchestrators](#) in the *Red Hat Ceph Storage Operations Guide* for more details.

6.7. MONITORING HAProxy ON THE DASHBOARD

The Ceph Object Gateway allows you to assign many instances of the object gateway to a single zone, so that you can scale out as load increases. Since each object gateway instance has its own IP address, you can use HAProxy to balance the load across Ceph Object Gateway servers.

You can monitor the following HAProxy metrics on the dashboard:

- Total responses by HTTP code.
- Total requests/responses.
- Total number of connections.
- Current total number of incoming / outgoing bytes.

You can also get the Grafana details by running the **ceph dashboard get-grafana-api-url** command.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Admin level access on the storage dashboard.
- An existing Ceph Object Gateway service, without SSL. If you want SSL service, the certificate should be configured on the ingress service, not the Ceph Object Gateway service.
- Ingress service deployed using the Ceph Orchestrator.
- Monitoring stack components are created on the dashboard.

Procedure

1. Log in to the Grafana URL and select the **RGW_Overview** panel:

Syntax

```
https://DASHBOARD_URL:3000
```

Example

```
https://dashboard_url:3000
```

2. Verify the HAProxy metrics on the Grafana URL.
3. From the Ceph dashboard navigation, go to **Object→Gateways**.
4. From the **Overall Performance** tab, verify the Ceph Object Gateway HAProxy metrics.

Figure 6.7. HAProxy metrics



Additional Resources

- See the [Configuring high availability for the Ceph Object Gateway](#) in the *Red Hat Ceph Storage Object Gateway Guide* for more details.

6.8. VIEWING THE CRUSH MAP OF THE CEPH CLUSTER ON THE DASHBOARD

You can view the The CRUSH map that contains a list of OSDs and related information on the Red Hat Ceph Storage dashboard. Together, the CRUSH map and CRUSH algorithm determine how and where data is stored. The dashboard allows you to view different aspects of the CRUSH map, including OSD hosts, OSD daemons, ID numbers, device class, and more.

The CRUSH map allows you to determine which host a specific OSD ID is running on. This is helpful if there is an issue with an OSD.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- OSD daemons deployed on the storage cluster.

Procedure

1. From the dashboard navigation, go to **Cluster→CRUSH map**.
2. To view the details of the specific OSD, click it's row.

Figure 6.8. CRUSH Map detail view



Additional Resources

- For more information about the CRUSH map, see [CRUSH admin overview](#) in the *Red Hat Ceph Storage Storage strategies guide*.

6.9. FILTERING LOGS OF THE CEPH CLUSTER ON THE DASHBOARD

You can view and filter logs of the Red Hat Ceph Storage cluster on the dashboard based on several criteria. The criteria includes *Priority*, *Keyword*, *Date*, and *Time range*.

You can download the logs to the system or copy the logs to the clipboard as well for further analysis.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- The Dashboard is installed.
- Log entries have been generated since the Ceph Monitor was last started.



NOTE

The Dashboard logging feature only displays the thirty latest high level events. The events are stored in memory by the Ceph Monitor. The entries disappear after restarting the Monitor. If you need to review detailed or older logs, refer to the file based logs.

Procedure

1. From the dashboard navigation, go to **Observability**→**Logs**.
2. From the **Cluster Logs** tab, view cluster logs.

Figure 6.9. Cluster logs

Observability » Logs

Priority: All Keyword: Date: YYYY-MM-DD Time range: 00 : 00 - 23 : 59

Cluster Logs Audit Logs Daemon Logs

```

11/4/24 10:27 AN [Obs] pgmap v1160: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 1.5 KiB/s rd, 1 op/s
11/4/24 10:27 AN [Obs] pgmap v1160: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 2.2 KiB/s rd, 2 op/s
11/4/24 10:27 AN [Obs] pgmap v1164: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 1.5 KiB/s rd, 1 op/s
11/4/24 10:27 AN [Obs] pgmap v1163: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 1.7 KiB/s rd, 1 op/s
11/4/24 10:27 AN [Obs] pgmap v1162: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 2.2 KiB/s rd, 2 op/s
11/4/24 10:27 AN [Obs] pgmap v1161: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 1.7 KiB/s rd, 1 op/s
11/4/24 10:27 AN [Obs] pgmap v1160: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 2.5 KiB/s rd, 2 op/s
11/4/24 10:27 AN [Obs] pgmap v1159: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 1.7 KiB/s rd, 1 op/s
11/4/24 10:27 AN [Obs] pgmap v1158: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 2.0 KiB/s rd, 2 op/s
11/4/24 10:27 AN [Obs] pgmap v1157: 977 pgs: 977 active+clean; 1.4 GiB data, 8.4 GiB used, 261 GiB / 270 GiB avail; 2.2 KiB/s rd, 2 op/s
  
```

- a. Use the **Priority** filter to filter by *Debug, Info, Warning, Error, or All*.
 - b. Use the **Keyword** field to enter text to search by keyword.
 - c. Use the **Date** picker to filter by a specific date.
 - d. Use the **Time range** fields to enter a range, using the *HH:MM - HH:MM* format. Hours must be entered using numbers **0** to **23**.
 - e. To combine filters, set two or more filters.
3. To save the logs, use the **Download** or **Copy to Clipboard** buttons.

Additional Resources

- See the [Configuring Logging](#) chapter in the *Red Hat Ceph Storage Troubleshooting Guide* for more information.
- See the [Understanding Ceph Logs](#) section in the *Red Hat Ceph Storage Troubleshooting Guide* for more information.

6.10. VIEWING CENTRALIZED LOGS OF THE CEPH CLUSTER ON THE DASHBOARD

Ceph Dashboard allows you to view logs from all the clients in a centralized space in the Red Hat Ceph Storage cluster for efficient monitoring. This is achieved through using Loki, a log aggregation system designed to store and query logs, and Promtail, an agent that ships the contents of local logs to a private Grafana Loki instance.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Grafana is configured and logged into on the cluster.

Procedure

1. From the dashboard navigation, go to **Administration→Services**.
2. From **Services**, click **Create**.
3. In the **Create Service** form, from the **Type** list, select **loki**. Fill in the remaining details, and click **Create Service**.
4. Repeat the previous step to create the **Promtail** service. Select **promtail** from the **Type** list. The **loki** and **promtail** services are displayed in the **Services** table, after being created successfully.

Figure 6.10. Creating Loki and Promtail services



Create Service X

Type * -- Select a service type --

Id

☐ Unmanaged

Placement Hosts

Hosts ✎ There are no hosts.

Count ?

Cancel Create Service



NOTE

By default, Promtail service is deployed on all the running hosts.

5. Enable logging to files.
 - a. Go to **Administration→Configuration**.
 - b. Select **log_to_file** and click **Edit**.

- c. In the **Edit log_to_file** form, set the **global** value to **true**.

Figure 6.11. Configuring log files

Edit log_to_file

Name	log_to_file
Description	send log lines to a file
Default	true

Values

global	true ▼
mon	-- Default -- ▼
mgr	-- Default -- ▼
osd	true ▼
mds	-- Default -- ▼
client	-- Default -- ▼

Cancel
Update

- d. Click **Update**.
The **Updated config option log_to_file** notification displays and you are returned to the **Configuration** table.
- e. Repeat these steps for **mon_cluster_log_to_file**, setting the **global** value to **true**.



NOTE

Both **log_to_file** and **mon_cluster_log_to_file** files need to be configured.

6. **Optional:** To view the Ceph Object Gateway 'ops_log', **rgw_enable_ops_log** must be set to **true** by using the following command:


```
$ ceph config set client.rgw rgw_enable_ops_log true
```

To do it from the dashboard, follow the below steps:

- Go to **Administration → Configuration**.
- Change level from 'basic' to 'Dev'.
- Search for **rgw_enable_ops_log** and edit the value to **true**.

- d. Next, under the **Daemon Logs** tab, locate the logs file in the **filename** field and run the query to view the ops log.
7. To view the centralized logs, go to **Observability→Logs** and switch to the **Daemon Logs** tab. Use **Log browser** to select files and click **Show logs** to view the logs from that file.

Figure 6.12. View centralized logs



6.11. MONITORING POOLS OF THE CEPH CLUSTER ON THE DASHBOARD

You can view the details, performance details, configuration, and overall performance of the pools in a cluster on the Red Hat Ceph Storage Dashboard.

A pool plays a critical role in how the Ceph storage cluster distributes and stores data. If you have deployed a cluster without creating a pool, Ceph uses the default pools for storing data.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Pools are created

Procedure

1. From the dashboard navigation, go to **Cluster→Pools**.
2. View the **Pools List** tab, which gives the details of Data protection and the application for which the pool is enabled. Hover the mouse over **Usage**, **Read bytes**, and **Write bytes** for the required details.
3. Expand the pool row for detailed information about a specific pool.

Figure 6.13. Monitoring pools



4. For general information, go to the **Overall Performance** tab.

Additional Resources

- For more information about pools, see [Ceph pools](#) in the *Red Hat Ceph Storage Architecture guide*.
- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

6.12. MONITORING CEPH FILE SYSTEMS ON THE DASHBOARD

You can use the Red Hat Ceph Storage Dashboard to monitor Ceph File Systems (CephFS) and related components.

For each File System listed, the following tabs are available:

Details

View the metadata servers (MDS) and their rank plus any standby daemons, pools and their usage, and performance counters.

Directories

View list of directories, their quotas and snapshots. Select a directory to set and unset maximum file and size quotas and to create and delete snapshots for the specific directory.

Subvolumes

Create, edit, and view subvolume information. These can be filtered by subvolume groups.

Subvolume groups

Create, edit, and view subvolume group information.

Snapshots

Create, clone, and view snapshot information. These can be filtered by subvolume groups and subvolumes.

Snapshot schedules

Enable, create, edit, and delete snapshot schedules.

Clients

View and evict Ceph File System client information.

Performance Details

View the performance of the file systems through the embedded Grafana Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- MDS service is deployed on at least one of the hosts.
- Ceph File System is installed.

Procedure

1. From the dashboard navigation, go to **File→File Systems**.
2. To view more information about an individual file system, expand the file system row.

Additional Resources

- For more information, see the [File System Guide](#).

6.13. MONITORING CEPH OBJECT GATEWAY DAEMONS ON THE DASHBOARD

You can use the Red Hat Ceph Storage Dashboard to monitor Ceph object gateway daemons. You can view the details, performance counters, and performance details of the Ceph object gateway daemons.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- At least one Ceph object gateway daemon configured in the storage cluster.

Procedure

1. From the dashboard navigation, go to **Object→Gateways**.
2. View information about individual gateways, from the **Gateways List** tab.
3. To view more information about an individual gateway, expand the gateway row.
4. If you have configured multiple Ceph Object Gateway daemons, click on **Sync Performance** tab and view the multi-site performance counters.

Additional Resources

- For more information, see the [Red Hat Ceph Storage Ceph object gateway Guide](#).

6.14. MONITORING BLOCK DEVICE IMAGES ON THE CEPH DASHBOARD

You can use the Red Hat Ceph Storage Dashboard to monitor and manage Block Device images. You can view the details, snapshots, configuration details, and performance details of the images.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. Expand the image row to see detailed information.

Figure 6.14. Monitoring Block Device images



Additional Resources

- See the [Creating images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

CHAPTER 7. MANAGING ALERTS ON THE CEPH DASHBOARD

As a storage administrator, you can see the details of alerts and create silences for them on the Red Hat Ceph Storage dashboard. This includes the following pre-defined alerts:

- CephadmDaemonFailed
- CephadmPaused
- CephadmUpgradeFailed
- CephDaemonCrash
- CephDeviceFailurePredicted
- CephDeviceFailurePredictionTooHigh
- CephDeviceFailureRelocationIncomplete
- CephFilesystemDamaged
- CephFilesystemDegraded
- CephFilesystemFailureNoStandby
- CephFilesystemInsufficientStandby
- CephFilesystemMDSRanksLow
- CephFilesystemOffline
- CephFilesystemReadOnly
- CephHealthError
- CephHealthWarning
- CephMgrModuleCrash
- CephMgrPrometheusModuleInactive
- CephMonClockSkew
- CephMonDiskspaceCritical
- CephMonDiskspaceLow
- CephMonDown
- CephMonDownQuorumAtRisk
- CephNodeDiskspaceWarning
- CephNodeInconsistentMTU
- CephNodeNetworkPacketDrops
- CephNodeNetworkPacketErrors

- CephNodeRootFilesystemFull
- CephObjectMissing
- CephOSDBackfillFull
- CephOSDDown
- CephOSDDownHigh
- CephOSDFlapping
- CephOSDFull
- CephOSDHostDown
- CephOSDInternalDiskSizeMismatch
- CephOSDNearFull
- CephOSDReadErrors
- CephOSDTimeoutsClusterNetwork
- CephOSDTimeoutsPublicNetwork
- CephOSDTooManyRepairs
- CephPGBackfillAtRisk
- CephPGImbalance
- CephPGNotDeepScrubbed
- CephPGNotScrubbed
- CephPGRecoveryAtRisk
- CephPGsDamaged
- CephPGsHighPerOSD
- CephPGsInactive
- CephPGsUnclean
- CephPGUnavailableBlockingIO
- CephPoolBackfillFull
- CephPoolFull
- CephPoolGrowthWarning
- CephPoolNearFull
- CephSlowOps

- PrometheusJobMissing

Figure 7.1. Pre-defined alerts

Cluster » [Monitoring](#) » Alerts

Active Alerts Alerts Silences

	Name 	Severity 
>	CephadmDaemonFailed	critical
>	CephadmPaused	warning
>	CephadmUpgradeFailed	critical
>	CephDaemonCrash	critical
>	CephDeviceFailurePredicted	warning
>	CephDeviceFailurePredictionTooHigh	critical
>	CephDeviceFailureRelocationIncomplete	warning
>	CephFilesystemDamaged	critical
>	CephFilesystemDegraded	critical
>	CephFilesystemFailureNoStandby	critical
>	CephFilesystemInsufficientStandby	warning
>	CephFilesystemMDSRanksLow	warning
>	CephFilesystemOffline	critical
>	CephFilesystemReadOnly	critical
>	CephHealthError	critical
>	CephHealthWarning	warning
>	CephMgrModuleCrash	critical
>	CephMgrPrometheusModuleInactive	critical
>	CephMonClockSkew	warning

	CephMonDown	warning
>	CephMonDiskspaceCritical	critical
>	CephMonDiskspaceLow	warning
>	CephMonDown	warning
>	CephMonDownQuorumAtRisk	critical
>	CephNodeDiskspaceWarning	warning
>	CephNodeInconsistentMTU	warning
>	CephNodeNetworkPacketDrops	warning
>	CephNodeNetworkPacketErrors	warning
>	CephNodeRootFilesystemFull	critical
>	CephObjectMissing	critical
>	CephOSDBackfillFull	warning
>	CephOSDDown	warning
>	CephOSDDownHigh	critical
>	CephOSDFlapping	warning
>	CephOSDFull	critical
>	CephOSDHostDown	warning
>	CephOSDInternalDiskSizeMismatch	warning
>	CephOSDNearFull	warning
>	CephOSDReadErrors	warning
>	CephOSDTimeoutsClusterNetwork	warning
>	CephOSDTimeoutsPublicNetwork	warning
>	CephOSDTooManyRepairs	warning
>	CephPGBackfillAtRisk	critical
>	CephPGImbalance	warning
>	CephPGNotDeepScrubbed	warning

>	CephPGNotScrubbed	warning
>	CephPGRecoveryAtRisk	critical
>	CephPGsDamaged	critical
>	CephPGsHighPerOSD	warning
>	CephPGsInactive	critical
>	CephPGsUnclean	warning
>	CephPGUnavailableBlockingIO	critical
>	CephPoolBackfillFull	warning
>	CephPoolFull	critical
>	CephPoolGrowthWarning	warning
>	CephPoolNearFull	warning
>	CephSlowOps	warning
>	PrometheusJobMissing	critical

You can also monitor alerts using simple network management protocol (SNMP) traps.

7.1. ENABLING MONITORING STACK

You can manually enable the monitoring stack of the Red Hat Ceph Storage cluster, such as Prometheus, Alertmanager, and Grafana, using the command-line interface.

You can use the Prometheus and Alertmanager API to manage alerts and silences.

Prerequisite

- A running Red Hat Ceph Storage cluster.
- root-level access to all the hosts.

Procedure

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

Example

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

2. Set the APIs for the monitoring stack:

- a. Specify the host and port of the Alertmanager server:

Syntax

```
ceph dashboard set-alertmanager-api-host ALERTMANAGER_API_HOST:PORT
```

Example

```
[ceph: root@host01 /]# ceph dashboard set-alertmanager-api-host http://10.0.0.101:9093
Option ALERTMANAGER_API_HOST updated
```

- b. To see the configured alerts, configure the URL to the Prometheus API. Using this API, the Ceph Dashboard UI verifies that a new silence matches a corresponding alert.

Syntax

```
ceph dashboard set-prometheus-api-host PROMETHEUS_API_HOST:PORT
```

Example

```
[ceph: root@host01 /]# ceph dashboard set-prometheus-api-host http://10.0.0.101:9095
Option PROMETHEUS_API_HOST updated
```

After setting up the hosts, refresh your browser's dashboard window.

- Specify the host and port of the Grafana server:

Syntax

```
ceph dashboard set-grafana-api-url GRAFANA_API_URL:PORT
```

Example

```
[ceph: root@host01 /]# ceph dashboard set-grafana-api-url https://10.0.0.101:3000
Option GRAFANA_API_URL updated
```

3. Get the Prometheus, Alertmanager, and Grafana API host details:

Example

```
[ceph: root@host01 /]# ceph dashboard get-alertmanager-api-host
http://10.0.0.101:9093
[ceph: root@host01 /]# ceph dashboard get-prometheus-api-host
http://10.0.0.101:9095
[ceph: root@host01 /]# ceph dashboard get-grafana-api-url
http://10.0.0.101:3000
```

4. Optional: If you are using a self-signed certificate in your Prometheus, Alertmanager, or Grafana setup, disable the certificate verification in the dashboard. This avoids refused connections caused by certificates signed by an unknown Certificate Authority (CA) or that do not match the hostname.

- For Prometheus:

Example

```
[ceph: root@host01 /]# ceph dashboard set-prometheus-api-ssl-verify False
```

- For Alertmanager:

Example

```
[ceph: root@host01 /]# ceph dashboard set-alertmanager-api-ssl-verify False
```

- For Grafana:

Example

```
[ceph: root@host01 /]# ceph dashboard set-grafana-api-ssl-verify False
```

5. Get the details of the self-signed certificate verification setting for Prometheus, Alertmanager, and Grafana:

Example

```
[ceph: root@host01 /]# ceph dashboard get-prometheus-api-ssl-verify
[ceph: root@host01 /]# ceph dashboard get-alertmanager-api-ssl-verify
[ceph: root@host01 /]# ceph dashboard get-grafana-api-ssl-verify
```

6. Optional: If the dashboard does not reflect the changes, you have to disable and then enable the dashboard:

Example

```
[ceph: root@host01 /]# ceph mgr module disable dashboard
[ceph: root@host01 /]# ceph mgr module enable dashboard
```

Additional Resources

- See the [Bootstrap command options](#) section in the *Red Hat Ceph Storage Installation Guide*.
- See the [Red Hat Ceph Storage installation](#) chapter in the *Red Hat Ceph Storage Installation Guide*.
- See the [Deploying the monitoring stack using the Ceph Orchestrator](#) section in the *Red Hat Ceph Storage Operations Guide*.

7.2. CONFIGURING GRAFANA CERTIFICATE

The **cephadm** deploys Grafana using the certificate defined in the ceph key/value store. If a certificate is not specified, **cephadm** generates a self-signed certificate during the deployment of the Grafana service.

You can configure a custom certificate with the **ceph config-key set** command.

Prerequisite

- A running Red Hat Ceph Storage cluster.

Procedure

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

Example

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

2. Configure the custom certificate for Grafana:

Example

```
[ceph: root@host01 /]# ceph config-key set mgr/cephadm/grafana_key -i $PWD/key.pem
[ceph: root@host01 /]# ceph config-key set mgr/cephadm/grafana_cert -i
$PWD/certificate.pem
```

3. If Grafana is already deployed, then run **reconfig** to update the configuration:

Example

```
[ceph: root@host01 /]# ceph orch reconfig grafana
```

4. Every time a new certificate is added, follow the below steps:

- a. Make a new directory

Example

```
[root@host01 ~]# mkdir /root/internalca
[root@host01 ~]# cd /root/internalca
```

- b. Generate the key:

Example

```
[root@host01 internalca]# openssl ecparam -genkey -name secp384r1 -out $(date
+ %F).key
```

- c. View the key:

Example

```
[root@host01 internalca]# openssl ec -text -in $(date + %F).key | less
```

- d. Make a request:

Example

```
[root@host01 internalca]# umask 077; openssl req -config openssl-san.cnf -new -sha256
-key $(date +%F).key -out $(date +%F).csr
```

- e. Review the request prior to sending it for signature:

Example

```
[root@host01 internalca]# openssl req -text -in $(date +%F).csr | less
```

- f. As the CA sign:

Example

```
[root@host01 internalca]# openssl ca -extensions v3_req -in $(date +%F).csr -out $(date
+%F).crt -extfile openssl-san.cnf
```

- g. Check the signed certificate:

Example

```
[root@host01 internalca]# openssl x509 -text -in $(date +%F).crt -noout | less
```

Additional Resources

- See the [Using shared system certificates](#) for more details.

7.3. ADDING ALERTMANAGER WEBHOOKS

You can add new webhooks to an existing Alertmanager configuration to receive real-time alerts about the health of the storage cluster. You have to enable incoming webhooks to allow asynchronous messages into third-party applications.

For example, if an OSD is down in a Red Hat Ceph Storage cluster, you can configure the Alertmanager to send notification on Google chat.

Prerequisite

- A running Red Hat Ceph Storage cluster with monitoring stack components enabled.
- Incoming webhooks configured on the receiving third-party application.

Procedure

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

Example

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

2. Configure the Alertmanager to use the webhook for notification:

Syntax

```

service_type: alertmanager
spec:
  user_data:
    default_webhook_urls:
      - "_URLS_"

```

The **default_webhook_urls** is a list of additional URLs that are added to the default receivers' **webhook_configs** configuration.

Example

```

service_type: alertmanager
spec:
  user_data:
    webhook_configs:
      - url: 'http:127.0.0.10:8080'

```

3. Update Alertmanager configuration:

Example

```
[ceph: root@host01 /]# ceph orch reconfig alertmanager
```

Verification

- An example notification from Alertmanager to Gchat:

Example

```

using: https://chat.googleapis.com/v1/spaces/(xx- space identifier -xx)/messages
posting: {'status': 'resolved', 'labels': {'alertname': 'PrometheusTargetMissing', 'instance':
'postgres-exporter.host03.chest
response: 200
response: {
  "name": "spaces/(xx- space identifier -xx)/messages/3PYDBOslofE.3PYDBOslofE",
  "sender": {
    "name": "users/114022495153014004089",
    "displayName": "monitoring",
    "avatarUrl": "",
    "email": "",
    "domainId": "",
    "type": "BOT",
    "isAnonymous": false,
    "caaEnabled": false
  },
  "text": "Prometheus target missing (instance postgres-exporter.cluster.local:9187)\n\nA
Prometheus target has disappeared. An e
  "cards": [],
  "annotations": [],
  "thread": {
    "name": "spaces/(xx- space identifier -xx)/threads/3PYDBOslofE"
  },
  "space": {
    "name": "spaces/(xx- space identifier -xx)",

```

```

"type": "ROOM",
"singleUserBotDm": false,
"threaded": false,
"displayName": "_privmon",
"legacyGroupChat": false
},
"fallbackText": "",
"argumentText": "Prometheus target missing (instance postgres-exporter.cluster.local:9187)\n\nA Prometheus target has disappea
"attachment": [],
"createTime": "2022-06-06T06:17:33.805375Z",
"lastUpdateTime": "2022-06-06T06:17:33.805375Z"

```

7.4. VIEWING ALERTS ON THE CEPH DASHBOARD


After an alert has fired, you can view it on the Red Hat Ceph Storage Dashboard. You can edit the **Manager module** settings to trigger a mail when an alert is fired.

Prerequisite

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A running simple mail transfer protocol (SMTP) configured.
- An alert emitted.

Procedure

1. From the dashboard navigation, go to **Observability→Alerts**.
2. View active Prometheus alerts from the **Active Alerts** tab.
3. View all alerts from the **Alerts** tab.
To view alert details, expand the alert row.
4. To view the source of an alert, click on its row, and then click **Source**.



Additional resources

- See the [Using the Ceph Manager alerts module](#) for more details to configure SMTP.

7.5. CREATING A SILENCE ON THE CEPH DASHBOARD

You can create a silence for an alert for a specified amount of time on the Red Hat Ceph Storage Dashboard.

Prerequisite

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- An alert fired.

Procedure

1. From the dashboard navigation, go to **Observability→Alerts**.
2. On the **Silences** tab, click **Create**.
3. In the **Create Silence** form, fill in the required fields.
 - a. Use the **Add matcher** to add silence requirements.

Figure 7.2. Creating a silence

Observability » Alerts » Silences » Create Silence

Create Silence

Creator * admin ✓

Comment * creating a silence ✓

Start time * 2024-05-13 08:36 ✓

Duration * 2h ✓

End time * 2024-05-13 10:36 ✓

Matchers *

alertname = CephadmPaused

+ Add matcher

Matches 1 rule with no active alerts.

Cancel Create Silence

4. Click **Create Silence**.
A notification displays that the silence was created successfully and the **Alerts Silenced** updates in the **Silences** table.

7.6. RECREATING A SILENCE ON THE CEPH DASHBOARD

You can recreate a silence from an expired silence on the Red Hat Ceph Storage Dashboard.

Prerequisite

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- An alert fired.
- A silence created for the alert.

Procedure

1. From the dashboard navigation, go to **Observability→Alerts**.
2. On the **Silences** tab, select the row with the alert that you want to recreate, and click **Recreate** from the action drop-down.
3. Edit any needed details, and click **Recreate Silence** button.
A notification displays indicating that the silence was edited successfully and the status of the silence is now **active**.

7.7. EDITING A SILENCE ON THE CEPH DASHBOARD

You can edit an active silence, for example, to extend the time it is active on the Red Hat Ceph Storage Dashboard. If the silence has expired, you can either recreate a silence or create a new silence for the alert.

Prerequisite

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- An alert fired.
- A silence created for the alert.

Procedure

1. Log in to the Dashboard.
2. On the navigation menu, click *Cluster*.
3. Select *Monitoring* from the drop-down menu.
4. Click the *Silences* tab.
5. To edit the silence, click its row.
6. In the *Edit* drop-down menu, select *Edit*.
7. In the *Edit Silence* window, update the details and click *Edit Silence*.

Figure 7.3. Edit silence

Cluster » Monitoring » Silences » Edit

Edit Silence ?

Creator *

Comment *

Start time * ?

Duration *

End time *

Matchers *

	alertname	>_	OSD down		☐	
	oid	>_	1.3.6.1.4.1.50495.15.1. ...		☐	
	severity	>_	warning		☐	
	type	>_	ceph_default		☐	

+ Add matcher

Your matcher seems to match no currently defined rule or active alert.

Cancel

Edit Silence

- You get a notification that the silence was updated successfully.

7.8. EXPIRING A SILENCE ON THE CEPH DASHBOARD

You can expire a silence so any matched alerts will not be suppressed on the Red Hat Ceph Storage Dashboard.

Prerequisite

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- An alert fired.
- A silence created for the alert.

Procedure

1. From the dashboard navigation, go to **Observability→Alerts**.
2. On the **Silences** tab, select the row with the alert that you want to expire, and click **Expire** from the action drop-down.
3. In the **Expire Silence** notification, select **Yes, I am sure** and click **Expire Silence**.
A notification displays indicating that the silence was expired successfully and the Status of the alert is **expired**, in the Silences table.

Additional Resources

- For more information, see the [Red Hat Ceph Storage Troubleshooting Guide](#).

CHAPTER 8. MANAGING NFS GANESHA EXPORTS ON THE CEPH DASHBOARD

As a storage administrator, you can manage the NFS Ganesha exports that use Ceph Object Gateway as the backstore on the Red Hat Ceph Storage dashboard. You can deploy and configure, edit and delete the NFS ganesha daemons on the dashboard.

The dashboard manages NFS-Ganesha configuration files stored in RADOS objects on the Ceph Cluster. NFS-Ganesha must store part of their configuration in the Ceph cluster.

8.1. CONFIGURING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD

You can configure NFS Ganesha on the dashboard after configuring the Ceph Object Gateway and enabling a dedicated pool for NFS-Ganesha using the command line interface.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Ceph Object gateway login credentials are added to the dashboard.
- A dedicated pool enabled and tagged with custom tag of **nfs**.
- At least **ganesha-manager** level of access on the Ceph dashboard.

Procedure

1. Log into the Cephadm shell:

Example

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

2. Create the RADOS pool, namespace, and enable **rgw**:

Syntax

```
ceph osd pool create POOL_NAME _  
ceph osd pool application enable POOL_NAME freeform/rgw/rbd/cephfs/nfs
```

Example

```
[ceph: root@host01 /]# ceph osd pool create nfs-ganesha  
[ceph: root@host01 /]# ceph osd pool application enable nfs-ganesha rgw
```

3. Deploy NFS-Ganesha gateway using placement specification in the command line interface:

Syntax

```
ceph orch apply nfs SERVICE_ID --placement="NUMBER_OF_DAEMONS HOST_NAME_1  
HOST_NAME_2 HOST_NAME_3"
```

Example

```
[ceph: root@host01 /]# ceph orch apply nfs foo --placement="2 host01 host02"
```

This deploys an NFS-Ganesha cluster **nfsganesha** with one daemon on **host01** and **host02**.

4. Update **ganesha-clusters-rados-pool-namespace** parameter with the namespace and the service_ID:

Syntax

```
ceph dashboard set-ganesha-clusters-rados-pool-namespace POOL_NAME/SERVICE_ID
```

Example

```
[ceph: root@host01 /]# ceph dashboard set-ganesha-clusters-rados-pool-namespace nfs-  
ganesha/foo
```

5. From the dashboard navigation, go to **Object→NFS**.
6. Click **Create**.
7. Complete the **Create NFS export** form and click **Create NFS export** to save and continue.

Object / [NFS](#) / Create

Create NFS export

Cluster (required)

test1

This is the ID of an NFS Service

Type

Bucket

Bucket (required)

bucket1

NFS Protocol (required)

☒ NFSv3☒ NFSv4

Pseudo (required)

/2

The position this export occupies in the Pseudo FS. It must be unique.
By using different Pseudo options, the same Path may be exported multiple times.

Access Type (required)

RO

Allows only operations that do not modify the server

Squash (required)

no_root_squash

Allows the root user on the NFS client to retain full root privileges on the NFS server, which may pose security risks.

Transport Protocol (required)

☒ UDP☒ TCP

Clients

Any client can access

Add clients +

Cancel

Create NFS export

- Verify the NFS daemon is configured:

Example

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

- As a root user, check if the NFS-service is active and running:

Example

```
[root@host01 ~]# systemctl list-units | grep nfs
```

- Mount the NFS export and perform a few I/O operations.

- Once the NFS service is up and running, in the NFS-RGW container, comment out the **dir_chunk=0** parameter in **etc/ganesha/ganesha.conf** file. Restart the NFS-Ganesha service. This allows proper listing at the NFS mount.

Verification

- You can view the NFS daemon by going to **File→NFS**.

Object / NFS

User	Bucket	Pseudo	Cluster	Storage Backend	Access Type
dashboard	bucket1	/2	test1	Object Gateway	RO

Items per page: 10 1-1 of 1 item 1 of 1 page

Additional Resources

- For more information on adding object gateway login credentials to the dashboard, see the [Manually adding object gateway login credentials to the dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on creating object gateway users on the dashboard, see the [Creating Ceph Object Gateway users on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on creating object gateway buckets on the dashboard, see the [Creating Ceph Object Gateway buckets on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on system roles, see the [Managing roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

8.2. CONFIGURING NFS EXPORTS WITH CEPHFS ON THE CEPH DASHBOARD

You can create, edit, and delete NFS exports on the Ceph dashboard after configuring the Ceph File System (CephFS) using the command-line interface. You can export the CephFS namespaces over the NFS Protocol.

You need to create an NFS cluster which creates a common recovery pool for all the NFS Ganesha daemons, new user based on the *CLUSTER_ID*, and a common NFS Ganesha config RADOS objects.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Root-level access to the bootstrapped host.
- At least **ganesha-manager** level of access on the Ceph dashboard.

Procedure

- Log in to the **cephadm** shell:

Example

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

2. Create the CephFS storage in the backend:

Syntax

```
ceph fs volume create CEPH_FILE_SYSTEM
```

Example

```
[ceph: root@host01 /]# ceph fs volume create cephfs
```

3. Enable the Ceph Manager NFS module:

Example

```
[ceph: root@host01 /]# ceph mgr module enable nfs
```

4. Create an NFS Ganesha cluster:

Syntax

```
ceph nfs cluster create NFS_CLUSTER_NAME "HOST_NAME_PLACEMENT_LIST"
```

Example

```
[ceph: root@host01 /]# ceph nfs cluster create nfs-cephfs host02  
NFS Cluster Created Successfully
```

5. Get the dashboard URL:

Example

```
[ceph: root@host01 /]# ceph mgr services  
{  
  "dashboard": "https://10.00.00.11:8443/",  
  "prometheus": "http://10.00.00.11:9283/"  
}
```

6. Log in to the Ceph dashboard with your credentials.
7. On the dashboard landing page, click *NFS*.
8. Click *Create*.
9. Complete the *Create NFS export* form and click *Create NFS export* to save and continue.

Figure 8.1. CephFS NFS export form

File / NFS / Create

Create NFS export

Cluster (required)

test1

This is the ID of an NFS Service

Volume (required)

cephfs

☐ Enable security label

Subvolume Group

_nogroup

Subvolume

subvolume1

Path (required)

/volumes/_nogroup/subvolume1/73a909c5-48ba-4cab-890d-5b5b5b5c5a2c

A path in a CephFS file system.

NFS Protocol (required)

☒ NFSv3

☒ NFSv4

Pseudo (required)

/1

The position this export occupies in the Pseudo FS. It must be unique.
By using different Pseudo options, the same Path may be exported multiple times.

Access Type (required)

RW

Allows all operations

Squash (required)

no_root_squash

Allows the root user on the NFS client to retain full root privileges on the NFS server, which may pose security risks.

Transport Protocol (required)

☒ UDP

☒ TCP

Clients

Any client can access

Add clients +

Cancel Create NFS export

- As a root user on the client host, create a directory and mount the NFS export:

Syntax

```
mkdir -p /mnt/nfs/
mount -t nfs -o port=2049 HOSTNAME:EXPORT_NAME _MOUNT_DIRECTORY_
```

Example

```
[root@ client ~]# mkdir -p /mnt/nfs/
[root@ client ~]# mount -t nfs -o port=2049 host02:/export1 /mnt/nfs/
```

Verification

- Verify if the NFS daemon is configured:

Example

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

Additional Resources

- See [Creating the NFS-Ganesha cluster using the Ceph Orchestrator](#) section in the *Red Hat Ceph Storage Operations Guide* for more details.

8.3. EDITING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD

You can edit the NFS Ganesha daemons on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- At least **ganesh-manager** level of access on the Ceph dashboard.
- NFS Ganesha daemon configured on the dashboard.

Procedure

1. From the dashboard navigation, go to **Object → NFS**.
2. Select the row that needs to be edited. and click **Edit**.
3. In the **Edit NFS export** window, edit the required parameters.
4. Complete by clicking **Edit NFS export**
A notification displays that the NFS object was updated successfully.

Additional Resources

- For more information on configuring NFS Ganesha, see [Configuring NFS Ganesha daemons on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on adding object gateway login credentials to the dashboard, see the [Manually adding Ceph Object Gateway login credentials to the dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on creating object gateway users on the dashboard, see the [Creating object gateway users on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on creating object gateway buckets on the dashboard, see the [Creating Ceph Object Gateway buckets on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on system roles, see the [Managing roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

8.4. DELETING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD

The Ceph dashboard allows you to delete the NFS Ganesha daemons.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- At least **ganesha-manager** level of access on the Ceph dashboard.
- NFS Ganesha daemon configured on the dashboard.

Procedure

1. From the dashboard navigation, go to **File→NFS**.
2. Select the row that needs to be deleted and click **Delete** from the action drop-down.
3. In the **Delete NFS export** notification, select **Yes, I am sure** and click **Delete NFS export**

Verification

- The selected row is deleted successfully.

Additional Resources

- For more information on configuring NFS Ganesha, see [Configuring NFS Ganesha daemons on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on adding object gateway login credentials to the dashboard, see the [Manually adding Ceph Object Gateway login credentials to the dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on creating object gateway users on the dashboard, see the [Creating object gateway users on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on creating object gateway buckets on the dashboard, see the [Creating Ceph Object Gateway buckets on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on system roles, see the [Managing roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

CHAPTER 9. MANAGING POOLS ON THE CEPH DASHBOARD

As a storage administrator, you can create, edit, and delete pools on the Red Hat Ceph Storage dashboard.

This section covers the following administrative tasks:

- [Creating pools on the Ceph dashboard](#).
- [Editing pools on the Ceph dashboard](#).
- [Deleting pools on the Ceph dashboard](#).

9.1. CREATING POOLS ON THE CEPH DASHBOARD

When you deploy a storage cluster without creating a pool, Ceph uses the default pools for storing data. You can create pools to logically partition your storage objects on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.

Procedure

1. From the dashboard navigation, go to **Cluster→Pools**.
2. Click **Create**.
3. Fill out the **Create Pool** form.

Figure 9.1. Creating pools

Pools » Create

Create Pool

Name *

test_pool

Pool type *

replicated ✓ ↕

PG Autoscale

on ▼

Replicated size *

3

Applications

rbd ✕

CRUSH

Crush ruleset

replicated_rule ⓘ + 🗑️

Compression

Mode

none ▼

Quotas

Max bytes ⓘ

e.g., 10GiB

Max objects ⓘ

0

RBD Configuration

Quality of Service +

Cancel

Create Pool

**NOTE**

The form changes based off of selection. Not all fields are mandatory.

- Set the name of the pool and select the pool type.
- Select the **Pool type**, either **replicated** or **erasure**. Erasure is referred to as Erasure Coded (EC).
- Optional: Select if the **PG Autoscale** is **on**, **off**, or **warn**.
- Optional: If using a replicated pool type, set the replicated size.
- Optional: If using an EC pool type configure the following additional settings.
- Optional: To see the settings for the currently selected EC profile, click the question mark.
- Optional: Add a new EC profile by clicking the plus symbol.
- Optional: Click the pencil symbol to select an application for the pool.
- Optional: Set the CRUSH rule, if applicable.

- Optional: If compression is required, select *passive*, *aggressive*, or *force*.
 - Optional: Set the Quotas.
 - Optional: Set the Quality of Service configuration.
4. To save the changes and complete creating the pool, click **Create Pool**.
A notification displays that the pool was created successfully.

Additional Resources

- For more information, see [Ceph pools](#) section in the *Red Hat Ceph Storage Architecture Guide* for more details.

9.2. EDITING POOLS ON THE CEPH DASHBOARD

You can edit the pools on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool is created.

Procedure

1. From the dashboard navigation, go to **Cluster→Pools**.
2. To edit the pool, select the pool row and click **Edit**.
3. In the **Edit Pool** form, edit the required parameters.
4. Save changes, by clicking **Edit Pool**.
A notification displays that the pool was updated successfully.

Additional Resources

- See the [Ceph pools](#) in the *Red Hat Ceph Storage Architecture Guide* for more information.
- See the [Pool values](#) in the *Red Hat Ceph Storage Storage Strategies Guide* for more information on Compression Modes.

9.3. DELETING POOLS ON THE CEPH DASHBOARD

You can delete the pools on the Red Hat Ceph Storage Dashboard. Ensure that value of **mon_allow_pool_delete** is set to **True** in Manager modules.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.

- A pool is created.

Procedure

1. From the dashboard navigation, go to **Administration→Configuration**.
2. From the **Configuration** table, select **mon_allow_pool_delete**, and click **Edit**.



NOTE

If needed, clear filters and search for the configuration.

3. From the **Edit mon_allow_pool_delete** form, in **Values**, set all values to **true**.
4. Click **Update**.
A notification displays that the configuration was updated successfully.
5. Go to **Cluster→Pools**.
6. Select the pool to be deleted, and click **Delete** from the action drop-down.
7. In the **Delete Pool** dialog, select **Yes, I am sure.** and complete, by clicking **Delete Pool**.
A notification displays that the pool was deleted successfully.

Additional Resources

- See the [Ceph pools](#) in the *Red Hat Ceph Storage Architecture Guide* for more information.
- See the [Pool values](#) in the *Red Hat Ceph Storage Storage Strategies Guide* for more information on Compression Modes.

CHAPTER 10. MANAGING HOSTS ON THE CEPH DASHBOARD

As a storage administrator, you can enable or disable maintenance mode for a host in the Red Hat Ceph Storage Dashboard. The maintenance mode ensures that shutting down the host, to perform maintenance activities, does not harm the cluster.

You can also remove hosts using *Start Drain* and *Remove* options in the Red Hat Ceph Storage Dashboard.

This section covers the following administrative tasks:

- [Entering maintenance mode.](#)
- [Exiting maintenance mode.](#)
- [Removing hosts using the Ceph Dashboard.](#)

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Hosts, Ceph Monitors and Ceph Manager Daemons are added to the storage cluster.

10.1. ENTERING MAINTENANCE MODE

You can enter a host into the maintenance mode before shutting it down on the Red Hat Ceph Storage Dashboard. If the maintenance mode gets enabled successfully, the host is taken offline without any errors for the maintenance activity to be performed. If the maintenance mode fails, it indicates the reasons for failure and the actions you need to take before taking the host down.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- All other prerequisite checks are performed internally by Ceph and any probable errors are taken care of internally by Ceph.

Procedure

1. From the dashboard navigation, go to **Cluster→Hosts**.
2. Select the host to enter maintenance mode, and click **Enter Maintenance** from the action drop-down.



NOTE

If the host contains Ceph Object Gateway (RGW) daemons, a warning displays that removing RGW daemons can cause clients to lose connectivity. Click **Continue** to enter maintenance.

**NOTE**

When a host enters maintenance mode, all daemons are stopped. Check the status of the daemons of a host by expanding the host view and switching to the **Daemons** tab.

A notification displays that the host was moved to maintenance successfully.

Verification

1. The **maintenance** label displays in the **Status** column of the Host table.

**NOTE**

If the maintenance mode fails, a notification displays, indicating the reasons for failure.

10.2. EXITING MAINTENANCE MODE

To restart a host, you can move it out of maintenance mode on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- All other prerequisite checks are performed internally by Ceph and any probable errors are taken care of internally by Ceph.

Procedure

1. From the dashboard navigation, go to **Cluster→Hosts**.
2. Select the host currently in maintenance mode, and click **Exit Maintenance** from the action drop-down.

**NOTE**

Identify which host is maintenance by checking for the **maintenance** label in the **Status** column of the Host table.

A notification displays that the host was moved out of maintenance successfully.

3. Create the required services on the host. By default, **crash** and **node-exporter** get deployed.

Verification

1. The **maintenance** label is removed from the **Status** column of the Host table.

10.3. REMOVING HOSTS USING THE CEPH DASHBOARD

To remove a host from a Ceph cluster, you can use *Start Drain* and *Remove* options in Red Hat Ceph Storage Dashboard.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- All other prerequisite checks are performed internally by Ceph and any probable errors are taken care of internally by Ceph.

Procedure

1. From the dashboard navigation, go to **Cluster→Hosts**.
2. Select the host that is to be removed, and click **Start Drain** from the action drop-down.

Figure 10.1. Selecting Start Drain option



This option drains all the daemons from the host.



NOTE

The **`_no_schedule`** label is automatically applied to the host, which blocks the deployment of daemons on this host.


- a. Optional: To stop the draining of daemons from the host, click **Stop Drain** from the action drop-down.
3. Check that all the daemons are removed from the host.
 - a. Expand the host row.
 - b. Go to the **Daemons** tab.
No daemons should be listed.



IMPORTANT

A host can be safely removed from the cluster after all the daemons are removed from it.

4. Select the host that is to be removed, and click **Remove** from the action drop-down.
 - a. In the **Remove Host** notification, select **Yes, I am sure** and click **Remove Host**.

A confirmation dialog box titled "Remove Host" with a close button (X) in the top right corner. The main text asks, "Are you sure that you want to remove **ceph-vpap-80-nfb57h-node1-installer**?" Below this, there is a checkbox labeled "Yes, I am sure." which is currently checked. At the bottom, there are two buttons: "Cancel" on the left and "Remove Host" on the right. The "Remove Host" button is highlighted in red, and a mouse cursor is pointing at it.

Remove Host

Are you sure that you want to remove **ceph-vpap-80-nfb57h-node1-installer**?

☒ Yes, I am sure.

Cancel Remove Host

A notification displays that the host is removed successfully.

CHAPTER 11. MANAGING CEPH OSDS ON THE DASHBOARD

As a storage administrator, you can monitor and manage OSDs on the Red Hat Ceph Storage Dashboard.

Some of the capabilities of the Red Hat Ceph Storage Dashboard are:

- List OSDs, their status, statistics, information such as attributes, metadata, device health, performance counters and performance details.
- Mark OSDs down, in, out, lost, purge, reweight, scrub, deep-scrub, destroy, delete, and select profiles to adjust backfilling activity.
- List all drives associated with an OSD.
- Set and change the device class of an OSD.
- Deploy OSDs on new drives and hosts.

Prerequisites

- A running Red Hat Ceph Storage cluster
- **cluster-manager** level of access on the Red Hat Ceph Storage dashboard

11.1. MANAGING THE OSDS ON THE CEPH DASHBOARD

You can carry out the following actions on a Ceph OSD on the Red Hat Ceph Storage Dashboard:

- Create a new OSD.
- Edit the device class of the OSD.
- Mark the Flags as *No Up*, *No Down*, *No In*, or *No Out*.
- Scrub and deep-scrub the OSDs.
- Reweight the OSDs.
- Mark the OSDs *Out*, *In*, *Down*, or *Lost*.
- Purge the OSDs.
- Destroy the OSDs.
- Delete the OSDs.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Hosts, Monitors, and Manager Daemons are added to the storage cluster.

Procedure

From the dashboard navigation, go to **Cluster→OSDs**.

Creating an OSD

1. To create the OSD, from the **OSDs List** table, click **Create**.

Figure 11.1. Add device for OSDs

Primary devices ×

⚠ At least one of these filters must be applied in order to proceed: Type Vendor Model Size

Hostname	Device path	Type	Vendor	Model	Size
ceph-adm2	/dev/sdc	HDD	QEMU	QEMU HARDDISK	30 GiB

1 total

Cancel Add



NOTE


Ensure you have an available host and a few available devices. Check for available devices in **Cluster→Physical Disks** and filter for *Available*.

- In the **Create OSDs** form, in the **Deployment Options** section, select one of the following options:
 - **Cost/Capacity-optimized:** The cluster gets deployed with all available HDDs.
 - **Throughput-optimized:** Slower devices are used to store data and faster devices are used to store journals/WALs.
 - **IOPS-optimized:** All the available NVMe devices are used to deploy OSDs.
- In the **Advanced Mode** section, add primary, WAL, and DB devices by clicking **Add**.
 - **Primary devices:** Primary storage devices contain all OSD data.
 - **WAL devices:** Write-Ahead-Log devices are used for BlueStore's internal journal and are used only if the WAL device is faster than the primary device. For example, NVMe or SSD devices.
 - **DB devices:** DB devices are used to store BlueStore's internal metadata and are used only if the DB device is faster than the primary device. For example, NVMe or SSD devices.
- To encrypt your data, for security purposes, from the **Features** section of the form, select **Encryption**.
- Click **Preview**.
- In the **OSD Creation Preview** dialog review the OSD and click **Create**.
A notification displays that the OSD was created successfully and the OSD status changes from *in and down* to *in and up*.

Editing an OSD

1. To edit an OSD, select the row and click **Edit**.
 - a. From the **Edit OSD** form, edit the device class.
 - b. Click **Edit OSD**.

Figure 11.2. Edit an OSD



The screenshot shows a modal window titled "Edit OSD: 9" with a close button (X) in the top right corner. Below the title bar, there is a form with a label "Device class" followed by a red asterisk (*). To the right of the label is a text input field containing the text "hdd". At the bottom right of the form, there are two buttons: a "Cancel" button and an "Edit OSD" button.

A notification displays that the OSD was updated successfully.

Marking the OSD flags

1. To mark the flag of the OSD, select the row and click **Flags** from the action drop-down.
2. In the **Individual OSD Flags** form, select the OSD flags needed.
3. Click **Update**.

Figure 11.3. Marking OSD flags

Individual OSD Flags ×

☐ **No Up**
OSDs are not allowed to start

☐ **No Down**
OSD failure reports are being ignored, such that the monitors will not mark OSDs down

☐ **No In**
OSDs that were previously marked out will not be marked back in when they start

☐ **No Out**
OSDs will not automatically be marked out after the configured interval

Restore previous selection
Cancel
Update

A notification displays that the OSD flags updated successfully.

Scrubbing an OSD

1. To scrub an OSD, select the row and click **Scrub** from the action drop-down.
2. In the **OSDs Scrub** notification, click **Update**.

Figure 11.4. Scrubbing an OSD

OSDs Scrub ×

You are about to apply a scrub to the OSD(s): 9.


Cancel
Update

A notification displays that the scrubbing of the OSD was initiated successfully.

Deep-scrubbing the OSDs

1. To deep-scrub the OSD, select the row and click **Deep Scrub** from the action drop-down.
2. In the **OSDs Deep Scrub** notification, click **Update**.

Figure 11.5. Deep-scrubbing an OSD




The dialog has a title bar with the text "OSDs Deep Scrub" and a close button (X) on the right. Below the title bar is a message: "You are about to apply a deep scrub to the OSD(s): 9." At the bottom right of the dialog are two buttons: "Cancel" and "Update".

A notification displays that the deep scrubbing of the OSD was initiated successfully.

Reweighting the OSDs

1. To reweight the OSD, select the row and click **Reweight** from the action drop-down.
2. In the **Reweight OSD** form enter a value between 0 and 1.
3. Click **Reweight**.


Figure 11.6. Reweighting an OSD



The dialog has a title bar with the text "Reweight OSD: 9" and a close button (X) on the right. Below the title bar is a form with a label "Weight" and a text input field containing the value "0.8". To the right of the input field are two small circular buttons with up and down arrows, and a green checkmark. At the bottom right of the dialog are two buttons: "Cancel" and "Reweight".

Marking OSDs out


1. To mark an OSD as *out*, select the row and click **Mark Out** from the action drop-down.
2. In the **Mark OSD out** notification, click **Mark Out**.

Figure 11.7. Marking OSDs *out*

The OSD status changes to *out*.

Marking OSDs *in*


1. To mark an OSD as *in*, select the OSD row that is in *out* status and click **Mark In** from the action drop-down.
2. In the **Mark OSD in** notification, click **Mark In**.

Figure 11.8. Marking OSDs *in*

The OSD status changes to *in*.

Marking OSDs *down*

1. To mark an OSD *down*, select the row and click **Mark Down** from the action drop-down.
2. In the **Mark OSD down** notification, click **Mark Down**.

Figure 11.9. Marking OSDs *down*


Mark OSD down


OSD(s) 9 will be marked **down** if you proceed.

Cancel Mark down

The OSD status changes to *down*.

Marking OSDs *lost*

1. To mark an OSD *lost*, select the OSD in *out* and *down* status and click **Mark Lost** from the action drop-down.
2. In the **Mark OSD Lost** notification, select **Yes, I am sure** and click **Mark Lost**.

Figure 11.10. Marking OSDs *lost*


Mark OSD lost

OSD 9 will be marked **lost** if you proceed.

Are you sure that you want to mark the selected OSD lost?

☒ Yes, I am sure.

Cancel Mark OSD lost

Purging OSDs

1. To purge an OSD, select the OSD in *down* status and click **Purge** from the action drop-down.
2. In the **Purge OSDs** notification, select **Yes, I am sure** and click **Purge OSD**.

Figure 11.11. Purging OSDs

Purge OSD



Warning
The OSD is not safe to be purged!

OSD 9 will be **purged** if you proceed.

Are you sure that you want to purge the selected OSD?

☒ **Yes, I am sure.**

CancelPurge OSD

All the flags are reset and the OSD is back in *in* and *up* status.

Destroying OSDs

1. To destroy an OSD, select the OSD in *down* status and click **Destroy** from the action drop-down.
2. In the **Destroy OSDs** notification, select **Yes, I am sure** and click **Destroy OSD**.

Figure 11.12. Destroying OSDs

Destroy OSD



Warning
The OSD is not safe to be destroyed!

OSD 9 will be **destroyed** if you proceed.

Are you sure that you want to destroy the selected OSD?

☒ **Yes, I am sure.**

Cancel

Destroy OSD

The OSD status changes to *destroyed*.

Deleting OSDs

1. To delete an OSD, select the OSD and click **Delete** from the action drop-down.
2. In the **Delete OSDs** notification, select **Yes, I am sure** and click **Delete OSD**.



NOTE

You can preserve the OSD_ID when you have to replace the failed OSD.

Figure 11.13. Deleting OSDs

Delete OSD
✕

OSD 8 will be **deleted** if you proceed.

Are you sure that you want to delete the selected OSD?

☒ Preserve OSD ID(s) for replacement.

☒ Yes, I am sure.

Cancel
Delete OSD

11.2. REPLACING THE FAILED OSDS ON THE CEPH DASHBOARD

You can replace the failed OSDs in a Red Hat Ceph Storage cluster with the **cluster-manager** level of access on the dashboard. One of the highlights of this feature on the dashboard is that the OSD IDs can be preserved while replacing the failed OSDs.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- At least **cluster-manager** level of access to the Ceph Dashboard.
- At least one of the OSDs is **down**


Procedure

1. On the dashboard, you can identify the failed OSDs in the following ways:

- Dashboard AlertManager pop-up notifications.
- Dashboard landing page showing *HEALTH_WARN* status.
- Dashboard landing page showing failed OSDs.
- Dashboard OSD list showing failed OSDs.

In the following example, you can see that one of the OSDs is down and one is out on the landing page of the dashboard.

Figure 11.14. OSD status on the Ceph Dashboard landing page



You can also view the LED blinking lights on the physical drive if one of the OSDs is down.

2. From **Cluster→OSDs**, on the **OSDs List** table, select the **out** and **down** OSD.
 - a. Click **Flags** from the action drop-down, select **No Up** in the **Individual OSD Flags** form, and click **Update**.
 - b. Click **Delete** from the action drop-down. In the **Delete OSD** notification, select **Preserve OSD ID(s) for replacement** and **Yes, I am sure** and click **Delete OSD**.
 - c. Wait until the status of the OSD changes to *out* and *destroyed*.
3. Optional: To change the *No Up* Flag for the entire cluster, from the **Cluster-wide configuration** menu, select **Flags**.
 - a. In **Cluster-wide OSDs Flags** form, select **No Up** and click **Update**.
4. Optional: If the OSDs are down due to a hard disk failure, replace the physical drive:
 - If the drive is hot-swappable, replace the failed drive with a new one.
 - If the drive is not hot-swappable and the host contains multiple OSDs, you might have to shut down the whole host and replace the physical drive. Consider preventing the cluster from backfilling. See the [Stopping and Starting Rebalancing](#) chapter in the *Red Hat Ceph Storage Troubleshooting Guide* for details.

- When the drive appears under the **/dev/** directory, make a note of the drive path.
- If you want to add the OSD manually, find the OSD drive and format the disk.
- If the new disk has data, zap the disk:

Syntax

```
ceph orch device zap HOST_NAME PATH --force
```

Example

```
ceph orch device zap ceph-adm2 /dev/sdc --force
```

5. From the Ceph Dashboard **OSDs List**, click **Create**.
6. In the **Create OSDs** form **Advanced Mode** section, add a primary device.
 - a. In the **Primary devices** dialog, select a **Hostname** filter.
 - b. Select a device type from the list.




NOTE

You have to select the Hostname first and then at least one filter to add the devices.

For example, from *Hostname* list, select **Type** and then **hdd**.

- c. Select **Vendor** and from device list, select **ATA**.

Figure 11.15. Using the Primary devices filter



- d. Click **Add**.
- e. In the **Create OSDs** form, click **Preview**.
- f. In the **OSD Creation Preview** dialog, click **Create**.
A notification displays that the OSD is created successfully and the OSD changes to be in the **out** and **down** status.

7. Select the newly created OSD that has *out* and *down* status.
 - a. Click **Mark In** from the action drop-down.
 - b. In the **Mark OSD in** notification, click **Mark In**.
The OSD status changes to *in*.
 - c. Click **Flags** from the action drop-down.
 - d. Clear the **No Up** selection and click **Update**.
8. Optional: If you have changed the *No Up* flag before for cluster-wide configuration, in the **Cluster-wide configuration** menu, select **Flags**.
 - a. In **Cluster-wide OSDs Flags** form, clear the **No Up** selection and click **Update**.

Verification

1. Verify that the OSD that was destroyed is created on the device and the OSD ID is preserved.

Devices	Attributes (OSD map)	Metadata	Device health	Performance counter	Performance Details
Device ID State of Health Life Expectancy Device Name Daemons					
QEMU_QEMU_HARDDISK_83d13ce8 f3d8-4a63-9f00-9650807b1248 Unknown sdc osd.4					
1 total					

Additional Resources

- For more information on Down OSDs, see the [Down OSDs](#) section in the *Red Hat Ceph Storage Troubleshooting Guide*.
- For additional assistance see the [Red Hat Support for service](#) section in the *Red Hat Ceph Storage Troubleshooting Guide*.
- For more information on system roles, see the [Managing roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

CHAPTER 12. MANAGING CEPH OBJECT GATEWAY USING THE DASHBOARD

As a storage administrator, the Ceph Object Gateway functions of the dashboard allow you to manage and monitor the Ceph Object Gateway.

You can also create the Ceph Object Gateway services with Secure Sockets Layer (SSL) using the dashboard.

For example, monitoring functions allow you to view details about a gateway daemon such as its zone name, or performance graphs of GET and PUT rates. Management functions allow you to view, create, and edit both users and buckets.

Ceph Object Gateway functions are divided between user functions and bucket functions.

12.1. MANUALLY ADDING CEPH OBJECT GATEWAY LOGIN CREDENTIALS TO THE DASHBOARD

The Red Hat Ceph Storage Dashboard can manage the Ceph Object Gateway, also known as the RADOS Gateway, or RGW. When Ceph Object Gateway is deployed with **cephadm**, the Ceph Object Gateway credentials used by the dashboard is automatically configured. You can also manually force the Ceph object gateway credentials to the Ceph dashboard using the command-line interface.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Ceph Object Gateway is installed.

Procedure

1. Log into the Cephadm shell:

Example

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

2. Set up the credentials manually:

Example

```
[ceph: root@host01 /]# ceph dashboard set-rgw-credentials
```

This creates a Ceph Object Gateway user with UID **dashboard** for each realm in the system.

3. Optional: If you have configured a custom **admin** resource in your Ceph Object Gateway admin API, you have to also set the the admin resource:

Syntax

```
ceph dashboard set-rgw-api-admin-resource RGW_API_ADMIN_RESOURCE
```

Example

```
[ceph: root@host01 /]# ceph dashboard set-rgw-api-admin-resource admin
Option RGW_API_ADMIN_RESOURCE updated
```

- Optional: If you are using HTTPS with a self-signed certificate, disable certificate verification in the dashboard to avoid refused connections.
Refused connections can happen when the certificate is signed by an unknown Certificate Authority, or if the host name used does not match the host name in the certificate.

Syntax

```
ceph dashboard set-rgw-api-ssl-verify false
```

Example

```
[ceph: root@host01 /]# ceph dashboard set-rgw-api-ssl-verify False
Option RGW_API_SSL_VERIFY updated
```

- Optional: If the Object Gateway takes too long to process requests and the dashboard runs into timeouts, you can set the timeout value:

Syntax

```
ceph dashboard set-rest-requests-timeout _TIME_IN_SECONDS_
```

The default value of 45 seconds.

Example

```
[ceph: root@host01 /]# ceph dashboard set-rest-requests-timeout 240
```

12.2. CREATING THE CEPH OBJECT GATEWAY SERVICES WITH SSL USING THE DASHBOARD

After installing a Red Hat Ceph Storage cluster, you can create the Ceph Object Gateway service with SSL using two methods:

- Using the command-line interface.
- Using the dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- SSL key from Certificate Authority (CA).

**NOTE**

Obtain the SSL certificate from a CA that matches the hostname of the gateway host. Red Hat recommends obtaining a certificate from a CA that has subject alternate name fields and a wildcard for use with S3-style subdomains.

Procedure

1. From the dashboard navigation, go to **Administration→Services**.
2. Click **Create**.
3. Fill in the **Create Service** form.
 - a. Select **rgw** from the **Type** service list.
 - b. Enter the ID that is used in **service_id**.
 - c. Select **SSL**.
 - d. Click **Choose File** and upload the SSL certificate **.pem** format.

Figure 12.1. Creating Ceph Object Gateway service

The screenshot shows the 'Create Service' dialog box with the following configuration:

- Type:** rgw
- Id:** rgw.ssl
- Realm:** (empty)
- Zone Group:** default
- Zone:** default
- Unmanaged:** ☐
- Placement:** Hosts
- Hosts:** ceph-vpap-doc-s3lpr0-node1-installer
- Count:** (empty)
- Port:** 443
- SSL:** ☒
- Certificate:** -----BEGIN RSA PRIVATE KEY-----
MIICXAIBAAKBgQCSws3GV/yiiuwf519LHMrdgC25Z1zbEWqh3
XhAmxYPFgr64x1Y01JL+jumespJz/+swjyzdpbIX7awp3912k

At the bottom right, there are two buttons: 'Cancel' and 'Create Service'.

e. Click **Create Service**.

4. Check the Ceph Object Gateway service is up and running.

Additional Resources

- See the [Configuring SSL for Beast](#) section in the *Red Hat Ceph Storage Object Gateway Guide*.

12.3. CONFIGURING HIGH AVAILABILITY FOR THE CEPH OBJECT GATEWAY ON THE DASHBOARD

The **ingress** service provides a highly available endpoint for the Ceph Object Gateway. You can create and configure the **ingress** service using the Ceph Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A minimum of two Ceph Object Gateway daemons running on different hosts.
- Dashboard is installed.
- A running **rgw** service.

Procedure

1. From the dashboard navigation, go to **Administration→Services**.
2. Click **Create**.
3. In the **Create Service** form, select **ingress** service.
4. Select backend service and edit the required parameters.

Figure 12.2. Creating **ingress** service

Create Service
×

Type *

ingress

Backend Service *

rgw.rgw

Id *

rgw.rgw

☐ Unmanaged

Placement

Hosts

Hosts

There are no hosts.

Count ?

2

Virtual IP ? *

10.240.3.68

Frontend Port ? *

2049

Monitor Port ? *

9000

CIDR Networks ?

☐ SSL

Cancel

Create Service

- Click **Create Service**.

A notification displays that the **ingress** service was created successfully.

Additional Resources

- See [High availability for the Ceph Object Gateway](#) for more information about the **ingress** service.

12.4. MANAGING CEPH OBJECT GATEWAY USERS ON THE DASHBOARD

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and manage Ceph Object Gateway users.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.

12.4.1. Creating Ceph object gateway users on the dashboard

You can create Ceph object gateway users on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.

Procedure

1. From the dashboard navigation, go to **Object→Users**.
2. On the **Users** tab, click **Create**.
3. **Create User** form, set the following parameters:
 - a. Enter the **User ID** and **Full name**.
 - b. If required, edit the maximum number of buckets.
 - c. Optional: Fill in an **Email address**
 - d. Optional: Select if the user is **Suspended** or a **System user**.
 - e. Optional: In the **S3 key** section, set a custom access key and secret key by clearing the **Auto-generate key** selection.
 - f. Optional: In the **User quota** section, select if the user quota is **Enabled, Unlimited size**, or has **Unlimited objects**. If there is a limited size enter the maximum size. If there are limited objects, enter the maximum objects.
 - g. Optional: In the **Bucket quota** section, select if the bucket quota is **Enabled, Unlimited size**, or has **Unlimited objects**. If there is a limited size enter the maximum size. If there are limited objects, enter the maximum objects.
4. Click **Create User**.

Figure 12.3. Create Ceph object gateway user

Create User

User ID *

rgw-test-11

✓

☐ Show Tenant

Full name *

rgw-test

✓

Email address

rgw-test@mail.com

✓

Max. buckets

Custom

1000

☐ Suspended

S3 key

☒ Auto-generate key

User quota

☐ Enabled

Bucket quota

☐ Enabled

Cancel

Create User

A notification displays that the user was created successfully.

Additional Resources

- See the [Manually adding Ceph object gateway login credentials to the dashboard](#) section in the Red Hat Ceph Storage Dashboard guide for more information.
- See the [Red Hat Ceph Storage Object Gateway Guide](#) for more information.

12.4.2. Adding roles to the Ceph Object Gateway users on the dashboard

You can add a role to a specific Ceph object gateway user on the Red Hat Ceph Storage dashboard.

Prerequisites

- Ceph Object Gateway is installed.
- Ceph Object gateway login credentials are added to the dashboard.

- Ceph Object gateway user is created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Roles*.
4. Select the user by clicking the relevant row.
5. From *Edit* drop-down menu, select *Create Role*.
6. In the *Create Role* window, configure *Role name*, *Path*, and *Assume Role Policy Document*.

Figure 12.4. Create Ceph object gateway subuser

Create Role

Role name *

s3access

Path *

/

Assume Role Policy Document * ?

```
{ "Version": "2012-10-17", "Statement": [ { "Effect": "Allow", "Principal": { "AWS": [ "arn:aws:iam:::user/kvm" ] }, "Action": [ "sts:AssumeRole", "s3:PutObject" ] } ] }
```

Cancel Create Role

7. Click *Create Role*.

12.4.3. Creating Ceph object gateway subusers on the dashboard

A subuser is associated with a user of the S3 interface. You can create a sub user for a specific Ceph object gateway user on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- Object gateway user is created.

Procedure

1. From the dashboard navigation, go to **Object→Users**.
2. On the **Users** tab, select a user and click **Edit**.
3. In the **Edit User** form, click **Create Subuser**.
4. In the **Create Subuser** dialog, enter the username and select the appropriate permissions.
5. Select the **Auto-generate secret** box and then click **Create Subuser**.

Figure 12.5. Create Ceph object gateway subuser

Create Subuser

×

Username

rgw-test-11

Subuser *

rgw-test-sub-user ✓

Permission *

read, write ✓ ▴ ▾

Swift key

☒ Auto-generate secret

Cancel

Create Subuser



NOTE

By selecting **Auto-generate-secret**, the secret key for Object Gateway is generated automatically.

6. In the **Edit User** form, click **Edit user**.
A notification displays that the user was updated successfully.

12.4.4. Editing Ceph object gateway users on the dashboard

You can edit Ceph object gateway users on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- A Ceph object gateway user is created.

Procedure

1. From the dashboard navigation, go to **Object→Users**.
2. On the **Users** tab, select the user row and click **Edit**.
3. In the **Edit User** form, edit the required parameters and click **Edit User**.

Figure 12.6. Edit Ceph object gateway user

Selected Object Gateway: rgw.def2.new.magna123.byabeh (us)

Object Gateway » Users » Edit

Edit User

User ID	rgw-test-11	✓
	☒ Show Tenant	
Tenant		✓
Full name	rgw-test-primary	✓
Email address	rgw-test@mail.com	
Max. buckets	Custom 1000	
	☐ Suspended	

Subusers

There are no subusers.

[+ Create Subuser](#)

Keys

S3	rgw-test-11 eye x
----	--

A notification displays that the user was updated successfully.

Additional Resources

- See the [Manually adding Ceph object gateway login credentials to the dashboard](#) section in the Red Hat Ceph Storage Dashboard guide for more information.
- See the [Red Hat Ceph Storage Object Gateway Guide](#) for more information.

12.4.5. Deleting Ceph Object Gateway users on the dashboard

You can delete Ceph object gateway users on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- A Ceph object gateway user is created.

Procedure

1. From the dashboard navigation, go to **Object→Users**.
2. Select the **Username** to delete, and click **Delete** from the action drop-down.
3. In the **Delete user** notification, select **Yes, I am sure** and click **Delete User**.
The user is removed from the **Users** table.

Figure 12.7. Delete Ceph object gateway user

Delete user

×

Are you sure that you want to delete **rgw-test-11**?

☒ Yes, I am sure.

CancelDelete user

Additional Resources

- See the [Manually adding Ceph object gateway login credentials to the dashboard](#) section in the Red Hat Ceph Storage Dashboard guide for more information.
- See the [Red Hat Ceph Storage Object Gateway Guide](#) for more information.

12.5. MANAGING CEPH OBJECT GATEWAY BUCKETS ON THE DASHBOARD

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and manage Ceph Object Gateway buckets.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- At least one Ceph Object Gateway user is created.
- Object gateway login credentials are added to the dashboard.

12.5.1. Creating Ceph object gateway buckets on the dashboard


You can create Ceph object gateway buckets on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- Object gateway user is created and not suspended.

Procedure

1. From the dashboard navigation, go to **Object→Buckets**.
2. Click **Create**.
The **Create Bucket** form displays.
3. Enter a **Name** for the bucket.
4. Select an **Owner**. The owner is a user that is not suspended.
5. Select a **Placement target**

**IMPORTANT**

A bucket's placement target cannot be changed after creation.

Figure 12.8. Create Ceph object gateway bucket

Create Bucket

Name *

Name...

Owner *

-- Select a user --

Placement target *

default-placement (pool: delhi.rgw.buckets.data)

Locking

☐ Enabled ?


Security

☐ Encryption ?

Cancel

Create Bucket

6. Optional: In the **Locking** section, select **Enabled** to enable locking for the bucket objects.

**IMPORTANT**

Locking can only be enabled while creating a bucket and cannot be changed after creation.

- a. Select the **Mode**, either **Compliance** or **Governance**.
- b. In the **Days** field, select the default retention period that is applied to new objects placed in this bucket.
7. Optional: In the **Security** section, select **Security** to encrypt objects in the bucket.
 - a. Set the configuration values for SSE-S3. Click the **Encryption** information icon and then **Click here**.

**NOTE**

When using **SSE-S3** encryption type, Ceph manages the encryption keys that are stored in the vault by the user.

- i. In the **Update RGW Encryption Configurations** dialog, ensure that **SSE-S3** is selected as the **Encryption Type**.
- ii. Fill the other required information.

- iii. Click **Submit**.

Figure 12.9. Encrypt objects in the bucket

Update RGW Encryption Configurations

Encryption Type:
☒ SSE-S3 Encryption
☐ SSE-KMS Encryption

Key management service provider *

vault

Authentication Method *

agent

Secret Engine *

transit

Secret Path

/v1/transit

Namespace

Vault Address *

Role *

CA Certificate ?

Choose file

Client Certificate ?

Choose file

Client Private Key ?

Choose file

Cancel

Submit

8. Click **Create bucket**.

A notification displays that the bucket was created successfully.

12.5.2. Editing Ceph object gateway buckets on the dashboard

You can edit Ceph object gateway buckets on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- Object gateway user is created and not suspended.
- A Ceph Object Gateway bucket created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Buckets*.
4. To edit the bucket, click it's row.
5. From the *Edit* drop-down select *Edit*.
6. In the *Edit bucket* window, edit the *Owner* by selecting the user from the dropdown.

Figure 12.10. Edit Ceph object gateway bucket

Selected Object Gateway: rgw.def2.new.magna123.byabeh (us)

Object Gateway » Buckets » Edit

Edit Bucket

Id	2026be62-1lad-4fbd-a82e-86f809aa1c90.105829.3
Name	rgw-test
Owner *	rgw-test-1l ▾
Placement target	default-placement

Versioning

☐ Enabled ?

Multi-Factor Authentication

☐ Delete enabled ?

Locking

☐ Enabled ?

Cancel Edit Bucket

- a. Optional: Enable *Versioning* if you want to enable versioning state for all the objects in an existing bucket.
 - To enable versioning, you must be the owner of the bucket.
 - If *Locking* is enabled during bucket creation, you cannot disable the versioning.
 - All objects added to the bucket will receive a unique version ID.
 - If the versioning state has not been set on a bucket, then the bucket will not have a versioning state.
- b. Optional: Check *Delete enabled* for *Multi-Factor Authentication*. Multi-Factor Authentication(MFA) ensures that users need to use a one-time password(OTP) when removing objects on certain buckets. Enter a value for *Token Serial Number* and *Token PIN*.

**NOTE**

The buckets must be configured with versioning and MFA enabled which can be done through the S3 API.

7. Click *Edit Bucket*.

8. You get a notification that the bucket was updated successfully.

12.5.3. Deleting Ceph Object Gateway buckets on the dashboard

You can delete Ceph object gateway buckets on the Red Hat Ceph Storage once the credentials are set-up using the CLI.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object Gateway login credentials are added to the dashboard.
- Object Gateway user is created and not suspended.
- A Ceph Object Gateway bucket created.

Procedure

1. From the dashboard navigation, go to **Object→Buckets**.
2. Select the bucket to be deleted, and click **Delete** from the action drop-down.
3. In the **Delete Bucket** notification, select **Yes, I am sure** and click **Delete bucket**.

Figure 12.11. Delete Ceph Object Gateway bucket



12.6. MONITORING MULTI-SITE OBJECT GATEWAY CONFIGURATION ON THE CEPH DASHBOARD

The Red Hat Ceph Storage dashboard supports monitoring the users and buckets of one zone in another zone in a multi-site object gateway configuration. For example, if the users and buckets are created in a zone in the primary site, you can monitor those users and buckets in the secondary zone in the secondary site.

Prerequisites

Prerequisites

- At least one running Red Hat Ceph Storage cluster deployed on both the sites.
- Dashboard is installed.
- The multi-site object gateway is configured on the primary and secondary sites.
- Object gateway login credentials of the primary and secondary sites are added to the dashboard.
- Object gateway users are created on the primary site.
- Object gateway buckets are created on the primary site.

Procedure

1. From the dashboard navigation of the secondary site, go to **Object→Buckets**.
2. View the Object Gateway buckets on the secondary landing page that were created for the Object Gateway users on the primary site.

Figure 12.12. Multi-site Object Gateway monitoring

Object Gateway » Buckets

[+ Create](#)

Name	Owner
☐ test	test-user-2

Name	test
ID	19969065-6089-49c5-89e0-6c3537b356bc.45549.4
Owner	test-user-2
Index type	Normal
Placement rule	default-placement
Marker	19969065-6089-49c5-89e0-6c3537b356bc.45549.4
Maximum marker	0#1#2#3#4#5#6#7#8#9#10#
Version	0#1#1#2#1#3#1#4#1#5#1#6#1#7#1#8#1#9#1#10#1
Master version	0#0#1#0#2#0#3#0#4#0#5#0#6#0#7#0#8#0#9#0#10#0

Additional Resources

- For more information on configuring multi-site, see the [Multi-site configuration and administration](#) section of the *Red Hat Ceph Storage Object Gateway* guide.
- For more information on adding object gateway login credentials to the dashboard, see the [Manually adding Ceph Object Gateway login credentials to the dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on creating object gateway users on the dashboard, see the [Creating Ceph Object Gateway users on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.

- For more information on creating object gateway buckets on the dashboard, see the [Creating Ceph Object Gateway buckets on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.

12.7. VIEWING CEPH OBJECT GATEWAY PER-USER AND PER-BUCKET PERFORMANCE COUNTERS ON THE DASHBOARD

You can view the Ceph Object Gateway performance counters per user per bucket in the Grafana dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Grafana is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- Object gateway user is created and not suspended.
- Configure below parameters to Ceph Object Gateway service:

```
# ceph config set <rgw-service> <param> <value>
"rgw_bucket_counters_cache": "true"
"rgw_user_counters_cache": "true"
```

Procedure

1. Log in to the Grafana URL.


Syntax

```
https://DASHBOARD_URL:3000
```


Example

```
https://dashboard_url:3000
```

2. Go to the 'Dashboard' tab and search for 'RGW S3 Analytics'.
3. To view per-bucket Ceph Object gateway operations, select the 'Bucket' panel:



4. To view user-level Ceph Object gateway operations, select the 'User' panel:



NOTE

The output of per-bucket/per-user **get operation count** command increases by two for each 'get' operation run from client: s3cmd. This is a known issue.

12.8. MANAGING CEPH OBJECT GATEWAY BUCKET POLICIES ON THE DASHBOARD

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and manage Ceph Object Gateway bucket policies.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.

- The Ceph Object Gateway is installed.
- At least one Ceph object gateway user is created.
- Ceph Object Gateway login credentials are added to the dashboard.
- At least one Ceph Object Gateway bucket. For more information about creating a bucket, see [Creating Ceph Object Gateway buckets on the dashboard](#).

12.8.1. Creating and editing Ceph Object Gateway bucket policies on the dashboard

You can create and edit Ceph Object Gateway bucket policies on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- At least one Ceph object gateway user is created.
- Ceph Object Gateway login credentials are added to the dashboard.
- At least one Ceph Object Gateway bucket. For more information about creating a bucket, see [Creating Ceph Object Gateway buckets on the dashboard](#).

Procedure

1. From the dashboard, go to **Object → Buckets**.
2. Create or modify a bucket policy for an existing bucket.



NOTE

To create a bucket policy during bucket creation, click **Create** and fill in the bucket policy information in the **Policies** section of the **Create Bucket** form.

Select the bucket for which the bucket policy will be created or modified, and then click **Edit**.

3. In the **Create Bucket** form, go to **Policies**.
4. Enter or modify the policy in JSON format.
Use the following links from within the form to help create your bucket policy. These links open a new tab in your browser.
 - **Policy generator** is an external tool from AWS to generate a bucket policy. For more information, see [AWS Policy Generator](#).



NOTE

You can use the policy generator with the **S3 Bucket Policy** type as a guideline for building your Ceph Object Gateway bucket policies.

- **Policy examples** takes you to AWS documentation with examples of bucket policies.

5. To save the bucket policy, click **Edit Bucket**.



NOTE

When creating a bucket policy during an initial bucket creation, click **Create Bucket**.

When the bucket policy is saved, the **Updated Object Gateway bucket `bucketname`** notification is displayed.

12.8.2. Deleting Ceph Object Gateway bucket policies on the dashboard

You can delete Ceph Object Gateway bucket policies on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- At least one Ceph object gateway user is created.
- Ceph Object Gateway login credentials are added to the dashboard.
- At least one Ceph Object Gateway bucket. For more information about creating a bucket, see [Creating Ceph Object Gateway buckets on the dashboard](#).

Procedure

1. From the dashboard, go to **Object → Buckets**.
2. Select the bucket for which the bucket policy will be created or modified, and then click **Edit**.
3. In the **Edit Bucket** form, go to **Policies**.
4. Click **Clear**.
5. To complete the bucket policy deletion, click **Edit Bucket**.
When the bucket policy is deleted, the **Updated Object Gateway bucket `bucketname`** notification is displayed.

12.9. MANAGING S3 BUCKET LIFECYCLE POLICIES ON THE DASHBOARD

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and manage S3 bucket lifecycle policies on the dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.

- At least one Ceph object gateway user is created.
- Ceph Object Gateway login credentials are added to the dashboard.
- At least one Ceph Object Gateway bucket. For more information about creating a bucket, see [Creating Ceph Object Gateway buckets on the dashboard](#) .

12.9.1. Applying and viewing S3 bucket lifecycle policies on the dashboard

You can apply and manage S3 bucket lifecycle policies on the Red Hat Ceph Storage dashboard.



PROCEDURE

Bucket lifecycle policiy cannot be applied during the creation of the bucket. They can be applied only after a bucket is created.

1. From the dashboard, go to **Object → Buckets**.
2. Select the bucket for which the lifecycle policy needs to be applied and click **Edit**.
3. In the **Edit Bucket** form, go to **Policies** and apply the lifecycle rule in the **Lifecycle** field in JSON format.
4. To save the bucket lifecycle policy, click **Edit Bucket**.

Figure 12.13. Apply bucket lifecycle policy

Enable
☐

Enables encryption for the objects in the bucket. To enable encryption on a bucket you need to set the configuration values for SSE-S3 or SSE-KMS. To set the configuration values [Click here](#)

Replication

Enable
☐

Enables replication for the objects in the bucket.

Tags

Tagging provides a way to categorize storage

+ Add tag

Policies

Bucket policy

{ }

✕ Clear

[Policy examples](#)
[Policy generator](#)

Lifecycle [?](#)

```

{
  "LifecycleConfiguration": {
    "Rule": {
      "ID": "Delete lc",
      "Prefix": "",
      "Status": "Enabled",
      "Expiration": {
        "Days": "1"
      }
    }
  }
}

```

✕ Clear

[Policy examples](#)

ACL [?](#)

Grantee [?](#)

Owner

▼

Permissions [?](#)

Full control

▼

Cancel

Edit Bucket

- After the bucket lifecycle policy is applied, it can viewed in the bucket listing screen by expanding the relevant bucket entry.

151

Figure 12.14. View bucket lifecycle policy

Object > Buckets

	Name	Owner	Used Capacity	Capacity Limit %	Objects	Object Limit %
☐	> test-ck1	ckulal	0 B	No Limit	0	No Limit
☐	> test-ck2	ckulal	0 B	No Limit	0	No Limit
☐	> testbkt3	test1	0 B	No Limit	0	No Limit
☐	▼ testbkt4	ckulal	0 B	No Limit	0	No Limit

Details Policies

Bucket policy

Lifecycle

JSON XML

```
{
  "LifecycleConfiguration": {
    "Rule": {
      "ID": "Rule 1",
      "Prefix": "",
      "Status": "Enabled",
      "Expiration": {
        "Date": "2022-04-05"
      }
    }
  }
}
```

Replication policy

```
{
  "Role": ""
}
```

ACL

Grantee	Permissions
Bucket Owner	FULL_CONTROL
Everyone	-
Authenticated users group	-

0 selected / 4 total

12.9.2. Deleting S3 bucket lifecycle policies on the dashboard

You can delete S3 bucket lifecycle policies on the Red Hat Ceph Storage dashboard.

Procedure

1. From the dashboard, go to **Object → Buckets**.
2. Select the bucket for which the bucket lifecycle policy needs to be deleted, and click **Edit**.
3. In the **Edit Bucket** form, go to **Policies**.
4. Click **Clear**.
5. To complete the bucket lifecycle policy deletion, click **Edit Bucket**.

12.10. MANAGEMENT OF BUCKETS OF A MULTI-SITE OBJECT CONFIGURATION ON THE CEPH DASHBOARD

As a storage administrator, you can edit buckets of one zone in another zone on the Red Hat Ceph Storage Dashboard. However, you can delete buckets of secondary sites in the primary site. You cannot delete the buckets of master zones of primary sites in other sites. For example, If the buckets are created in a zone in the secondary site, you can edit and delete those buckets in the master zone in the primary site.

Prerequisites

- At least one running Red Hat Ceph Storage cluster deployed on both the sites.
- Dashboard is installed.
- The multi-site object gateway is configured on the primary and secondary sites.

- Object gateway login credentials of the primary and secondary sites are added to the dashboard.
- Object gateway users are created on the primary site.
- Object gateway buckets are created on the primary site.
- At least **rgw-manager** level of access on the Ceph dashboard.

12.10.1. Monitoring buckets of a multi-site object

Monitor the multi-site sync status of a bucket on the dashboard. You can view the source zones and sync status from **Object→Multi-site** on the Ceph Dashboard.

The multi-site sync status is divided into two sections:

Primary Source Zone

Displays the default realm, zonegroup, and the zone the Ceph Object Gateway is connected to.

Source Zones

View both the metadata sync status and data sync information progress. When you click the status, a breakdown of the shard syncing is displayed. The sync status shows the **Last Synced** time stamp with the relative time of the last sync occurrence in relation to the current time. When the sync is complete, this shows as **Up to Date**. When a sync is not caught up the **status** shows as **Syncing**. However, the **Last sync** shows the number of days the sync is not caught up. By clicking **Syncing**, it displays the details about shards which are not synced.

12.10.2. Editing buckets of a multi-site Object Gateway configuration on the Ceph Dashboard

You can edit and update the details of the buckets of one zone in another zone on the Red Hat Ceph Storage Dashboard in a multi-site object gateway configuration. You can edit the owner, versioning, multi-factor authentication and locking features of the buckets with this feature of the dashboard.

Prerequisites

- At least one running Red Hat Ceph Storage cluster deployed on both the sites.
- Dashboard is installed.
- The multi-site object gateway is configured on the primary and secondary sites.
- Object gateway login credentials of the primary and secondary sites are added to the dashboard.
- Object gateway users are created on the primary site.
- Object gateway buckets are created on the primary site.
- At least **rgw-manager** level of access on the Ceph dashboard.

Procedure

1. From the dashboard navigation of the secondary site, go to **Object→Buckets**.
The Object Gateway buckets from the primary site are displayed.

2. Select the bucket that you want to edit, and click **Edit** from the action drop-down.
3. In the **Edit Bucket** form, edit the required parameters, and click **Edit Bucket**.
A notification is displayed that the bucket is updated successfully.

Figure 12.15. Edit buckets in a multi-site

Object Gateway » Buckets » Edit

Edit Bucket

Id

19969065-6089-49c5-89e0-6c3537b356bc.45549.4

Name

test

Owner *

test-user

Placement target

default-placement

Versioning

☐ Enabled ⓘ

Multi-Factor Authentication

☐ Delete enabled ⓘ

Locking

☐ Enabled ⓘ

Cancel

Edit Bucket

Additional Resources

- For more information on adding object gateway login credentials to the dashboard, see the [Manually adding Ceph Object Gateway login credentials to the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on creating object gateway users on the dashboard, see the [Creating Ceph Object Gateway users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on creating object gateway buckets on the dashboard, see the [Creating Ceph Object Gateway buckets on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard guide*.
- For more information on system roles, see the [Managing roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

12.10.3. Deleting buckets of a multi-site Object Gateway configuration on the Ceph Dashboard

You can delete buckets of secondary sites in primary sites on the Red Hat Ceph Storage Dashboard in a multi-site Object Gateway configuration.



IMPORTANT

Red Hat does not recommend to delete buckets of primary site from secondary sites.

Prerequisites

- At least one running Red Hat Ceph Storage cluster deployed on both the sites.
- Dashboard is installed.
- The multi-site object gateway is configured on the primary and secondary sites.
- Object Gateway login credentials of the primary and secondary sites are added to the dashboard.
- Object Gateway users are created on the primary site.
- Object Gateway buckets are created on the primary site.
- At least **rgw-manager** level of access on the Ceph dashboard.

Procedure

1. From the dashboard navigation of the primary site, go to **Object→Buckets**.
2. Select the bucket of the secondary site to be deleted, and click **Delete** from the action drop-down.
3. In the **Delete Bucket** notification, select **Yes, I am sure** and click **Delete bucket**.
The bucket is deleted from the **Buckets** table.

Additional Resources

- For more information on configuring multi-site, see the [Multi-site configuration and administration](#) section of the *Red Hat Ceph Storage Object Gateway* guide.
- For more information on adding object gateway login credentials to the dashboard, see the [Manually adding object gateway login credentials to the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on creating object gateway users on the dashboard, see the [Creating object gateway users on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on creating object gateway buckets on the dashboard, see the [Creating object gateway buckets on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard* guide.
- For more information on system roles, see the [System roles on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide*.

12.11. CONFIGURING A MULTI-SITE OBJECT GATEWAY ON THE CEPH DASHBOARD

You can configure Ceph Object Gateway multi-site on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster deployed on both the sites.
- At least one Ceph Object Gateway service installed at both the sites.

Procedure

1. Enable the Ceph Object Gateway module for import/export on both the the primary and secondary sites.
 - a. From the dashboard navigation of the secondary site, go to **Object→Multi-site**.
 - b. In the **In order to access the import/export feature, the rgw module must be enabled** note, click **Enable**.
2. On the primary site dashboard, create a default realm, zonegroup, and zone.
 - a. Click **Create Realm**.
 - b. In the **Create Realm** form, provide a realm name, and select **Default**.
 - c. Click **Create Realm**.
 - d. Click **Create Zone Group** from the action drop-down.
 - e. In the **Create Zone Group** form, provide a zone group name, the Ceph Object Gateway endpoints, and select **Default**.
 - f. Click **Create Zone Group**.
 - g. Click **Create Zone** from the action drop-down.
 - h. In the **Create Zone** form, provide a **Zone Name**, select **Default**, and provide the Ceph Object Gateway endpoints of the primary site. For the user, provide the access and secret key of the user with system privileges.

NOTE

While creating a zone, Red Hat recommends to give access key and secret key of the dashboard default user, **dashboard**.
 - i. Click **Create Zone**.

A warning is displayed to restart the Ceph Object Gateway service to complete the zone creation.
3. Restart the Ceph Object Gateway service.
 - a. From the dashboard navigation of the secondary site, go to **Administration→Services**.
 - b. Select the Ceph Object Gateway service row and expand the row.

- c. From the **Daemons** tab, select the hostname.
- d. Click **Restart** from the action drop-down.
4. From the dashboard navigation, in **Object→Overview** you get an error that "The Object Gateway Service is not configured". This bug is a known issue. See [BZ#2231072](#).
 - a. As a workaround, set the Ceph Object Gateway credentials on the command-line interface.

Syntax

```
ceph dashboard set-rgw-credentials
RGW credentials configured
```

- b. Go to **Object→Overview** to verify that you are able to access the Ceph Object Gateway on the dashboard.
5. Create a replication user on the primary site. You can use the following two options:
 - Create user using the CLI:

Example

```
[ceph: root@host01 /]# radosgw-admin user create --uid="uid" --display-
name="displayname" --system
```

- Create user from the dashboard and modify the user from the CLI:

Example

```
[ceph: root@host01 /]# radosgw-admin user modify --uid="uid" --system
```

6. From the dashboard navigation, go to **Object→Users**.
7. Expand the user row and from **Keys**, click **Show**.
 - a. Use the **Copy to Clipboard** to copy the access and secret keys. These will be used in a later step.
8. From the primary site dashboard, go to **Object→Multi-site**.
 - a. From the **Topology Viewer**, select the zone and click the **Edit** icon.
 - b. From the **Edit Zone** form, paste the access key in the **S3 access key** field and the secret key in the **S3 secret key** field. Use the keys that were copied previously.
 - c. Click **Edit Zone**.
9. Click **Export**.
 - a. From the **Export Multi-site Realm Token** dialog, copy the token.
10. From the secondary site, go to **Object→Multi-site**.
11. Import the token from the primary zone, by clicking **Import**.

- a. In the **Import Multi-site Token** dialog, in the **Zone** section, paste the token that was copied earlier, and provide a secondary zone name.
 - b. In the **Service** section, select the placement and the port where the new Ceph Object Gateway service is going to be created.
 - c. Click **Import**.
A warning is displayed to restart the Ceph Object Gateway service.
12. Restart the Ceph Object Gateway service.
 - a. From the dashboard navigation of the secondary site, go to **Administration→Services**.
 - b. Select the Ceph Object Gateway service row and expand the row.
 - c. From the **Daemons** tab, select the hostname.
 - d. Click **Restart** from the action drop-down.
Wait until the users are synced to the secondary site.
 13. Verify that the sync is complete using the following commands:

Syntax

```
radosgw-admin sync status
radosgw-admin user list
```

Example

```
[ceph: root@host01 /]# radosgw-admin sync status
[ceph: root@host01 /]# radosgw-admin user list
```


14. In **Object→Overview** you get an error that "The Object Gateway Service is not configured". This bug is a known issue. See [BZ#2231072](#).
 - a. As a workaround, set the Ceph Object Gateway credentials on the command-line interface.

Syntax

```
ceph dashboard set-rgw-credentials
RGW credentials configured
```

- b. Go to **Object→Overview** to verify that you are able to access the Ceph Object Gateway on the dashboard.
15. On the primary site, **Object→Overview**, in the **Multi-Site Sync Status** section, an error is displayed because on the secondary zone you can see that the endpoints and the hostname are not the IP address. This bug is a known issue while configuring multi-site. See [BZ#2242994](#).
 - a. As a workaround, from the secondary site dashboard, go to **Object→Multi-site**.
 - b. Select the secondary zone and click the **Edit** icon.
 - c. Edit the endpoints to reflect the IP address.
 - d. Click **Edit Zone**.

16. On the primary site and secondary site dashboards, from **Object→Overview**, in the **Multi-Site Sync Status** section, the status displays.



Verification

- Create a user on the primary site. You see that the user syncs to the secondary site.

CHAPTER 13. MANAGING FILE SYSTEMS USING THE CEPH DASHBOARD

As a storage administrator, you can create, edit, delete, and manage accesses of file systems on the Red Hat Ceph Storage dashboard.

13.1. CONFIGURING CEPHFS VOLUMES

As a storage administrator, you can configure Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

13.1.1. Creating CephFS volumes

You can create Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

Prerequisites

- A working Red Hat Ceph Storage cluster with MDS deployed.

Procedure

1. From the dashboard navigation, go to File > File Systems.
2. Click Create.
3. In the Create Volume window, set the following parameters:
 - a. Name: Set the name of the volume.
 - b. Placement (Optional): Select the placement of the volume. You can set it as either Hosts or Label.
 - c. Hosts/Label (Optional): If placement is selected as Hosts in the previous option, then select the appropriate host from the list. If placement is selected as Label in the previous option, then enter the label.
[IMPORTANT]: To identify the label that has been created for the hosts while using it as a placement, you can run the following command from the CLI:

```
[ceph: root@ceph-hk-ds-uoayxl-node1-installer /]# ceph orch host ls
```

```
HOST ADDR LABELS STATUS
```

```
host01 10.0.210.182 _admin,installer,mon,mgr
host02 10.0.96.72 mon,mgr
host03 10.0.99.37 mon,mds
host04 10.0.99.244 osd,mds
host05 10.0.98.118 osd,mds
host06 10.0.98.66 osd,nfs,mds
host07 10.0.98.23 nfs,mds
7 hosts in cluster
```

4. Click Create Volume.
 - a. A notification displays that the volume was created successfully.

13.1.2. Editing CephFS volumes

You can edit Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

Procedure

1. From the dashboard navigation, go to File > File Systems.
2. From the listed volumes, select the volume to be edited and click Edit.
3. In the Edit File System window, rename the volume as required and click Edit File System.
A notification displays that the volume was edited successfully.

13.1.3. Removing CephFS volumes

You can remove Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

Procedure

1. From the dashboard navigation, go to File > File Systems.
2. From the listed volumes, select the volume to be removed and click Remove from the action drop-down.
3. In the Remove File System window, select Yes, I am sure and click Remove File System.
A notification displays that the volume was removed successfully.

13.2. CONFIGURING CEPHFS SUBVOLUME GROUPS

As a storage administrator, you can configure Ceph File System (CephFS) subvolume groups on the Red Hat Ceph Storage dashboard.

13.2.1. Creating CephFS subvolume groups

You can create subvolume groups to create subvolume on the dashboard. You can also use subvolume groups to apply policies across a set of subvolumes.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.

2. From the listed volumes, select the row for which you want to create subvolumes and expand the row.
3. From the **Subvolume groups** tab, click 'Create' to create a subvolume group.
4. In the **Create Subvolume group** window, enter the following parameters:
 - a. **Name**: Set the name of the subvolume group.
 - b. **Volume name**: Validate that the correct name of the volume is selected.
 - c. **Size**: Set the size of the subvolume group. If left blank or entered as 0, then, by default, the size will be set as infinite.
 - d. **Pool**: Set the pool of the subvolume group. By default, data_pool_layout of the parent directory is selected.
 - e. **UID**: Set the UID of the subvolume group.
 - f. **GID**: Set the GID of the subvolume group.
 - g. **Mode**: Set the permissions for the directory. By default, the mode is 755 which is rwxr-xr-x.
5. Click **Create Subvolume group**.
A notification displays that the subvolume group was created successfully.

13.2.2. Editing CephFS subvolume groups

You can edit the subvolume groups on the dashboard.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. From the listed volumes, select the row for which you want to edit a subvolumes and expand the row.
3. From the **Subvolume groups** tab, select the row containing the group that you want to edit.
4. Click **Edit** to edit a subvolume group.
5. In the **Edit Subvolume group** window, edit the needed parameters and click **Edit Subvolume group**.
A notification displays that the subvolume group was edited successfully.

13.2.3. Removing CephFS subvolume groups

You can remove Ceph File System (CephFS) subvolume groups on the Red Hat Ceph Storage dashboard.

**WARNING**

Ensure to remove the subvolume within the subvolume group before removing the subvolume group.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. From the listed volumes, select the row for which you want to remove a subvolume and expand the row.
3. Under the **Subvolume groups** tab, select the subvolume group you want to remove, and click **Remove** from the action drop-down.
4. In the **Remove Subvolume group** window, select 'Yes, I am sure', and click **Remove Subvolume group**.
A notification displays that the subvolume group was removed successfully.

13.3. CONFIGURING CEPHFS SUBVOLUMES

As a storage administrator, you can configure Ceph File System (CephFS) subvolumes on the Red Hat Ceph Storage dashboard.

13.3.1. Creating CephFS subvolume

You can create Ceph File System (CephFS) subvolume on the Red Hat Ceph Storage dashboard.

**WARNING**

Ensure to remove the subvolume within the subvolume group before removing the subvolume group.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.

2. From the listed volumes, select the row.
3. Under the Subvolume tab, select the Subvolume group in which you want to create the subvolume.

**NOTE**

If you select the Default option, the subvolumes are not created under any subvolume groups.

4. Click Create to create a subvolume.
5. In the Create Subvolume window, enter the following parameters:
 - a. Name: Set the name of the subvolume group.
 - b. Subvolume name: Set the name of the volume.
 - c. Size: Set the size of the subvolume group. If left blank or entered as 0, then, by default, the size will be set as infinite.
 - d. Pool: Set the pool of the subvolume group. By default, data_pool_layout of the parent directory is selected.
 - e. UID: Set the UID of the subvolume group.
 - f. GID: Set the GID of the subvolume group.
 - g. Mode: Set the permissions for the directory. By default, the mode is 755 which is rwxr-xr-x.
 - h. Isolated Namespace: If you want to create the subvolume in a separate RADOS namespace, select this option.
6. Click **Create Subvolume**.

A notification that the subvolume was created successfully is displayed.

13.3.2. Editing CephFS subvolume

You can edit Ceph File System (CephFS) subvolume on the Red Hat Ceph Storage dashboard.

**NOTE**

You can only edit the size of the subvolumes on the dashboard.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.
- A subvolume created.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. From the listed volumes, select the row.

3. Under the **Subvolume** tab, select the subvolume you want to edit, and click **Edit** within the volume.
4. Click Create to create a subvolume.
5. In the **Edit Subvolume** window, enter the following parameters:
 - a. **Name:** Set the name of the subvolume group.
 - b. **Subvolume name:** Set the name of the volume.
 - c. **Size:** Set the size of the subvolume group. If left blank or entered as 0, then, by default, the size will be set as infinite.
 - d. **Pool:** Set the pool of the subvolume group. By default, data_pool_layout of the parent directory is selected.
 - e. **UID:** Set the UID of the subvolume group.
 - f. **GID:** Set the GID of the subvolume group.
 - g. **Mode:** Set the permissions for the directory. By default, the mode is 755 which is rwxr-xr-x.
6. Click **Edit Subvolume**.
A notification that the subvolume was created successfully is displayed.

13.3.3. Removing CephFS subvolume

You can remove Ceph File System (CephFS) subvolume on the Red Hat Ceph Storage dashboard.

Prerequisites

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.
- A subvolume created.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. From the listed volumes, select the row.
3. Navigate to **Subvolume** tab, select the subvolume you want to remove, and click **Remove**.
4. In the **Remove Subvolume** window, confirm whether you want to remove the selected subvolume and click **Remove Subvolume**.
A notification that the subvolume was removed successfully is displayed.

13.4. MANAGING CEPHFS SNAPSHOTS

As a storage administrator, you can create Ceph File System (CephFS) volume and subvolume snapshots on the Red Hat Ceph Storage dashboard.

The Ceph File System (CephFS) snapshots create an immutable, point-in-time view of a Ceph File System. CephFS snapshots are asynchronous and are kept in a special hidden directory in the CephFS directory named .snap. You can specify snapshot creation for any directory within a Ceph File System.

When specifying a directory, the snapshot also includes all the subdirectories beneath it.

CephFS snapshot feature is enabled by default on new Ceph File Systems, but it must be manually enabled on existing Ceph File Systems.



WARNING

Each Ceph Metadata Server (MDS) cluster allocates the snap identifiers independently. Using snapshots for multiple Ceph File Systems that are sharing a single pool causes snapshot collisions, and results in missing file data.

13.4.1. Creating CephFS subvolume snapshots

As a storage administrator, you can create Ceph File System (CephFS) subvolume snapshot on the Red Hat Ceph Storage dashboard. You can create an immutable, point-in-time view of a Ceph File System (CephFS) by creating a snapshot.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A subvolume group with corresponding subvolumes.

Procedure

1. Log into the dashboard.
2. On the dashboard navigation menu, click **File > File Systems**.
3. Select the CephFS where you want to create the subvolume snapshot.
If there are no file systems available, create a new file system.
4. Go to the **Snapshots** tab. There are three columns – **Groups**, **Subvolumes**, and **Create**.
5. From the **Groups** and **Subvolumes** columns, select the group and the subvolume for which you want to create a snapshot.
6. Click Create in the third column. The Create snapshot form opens.
 - a. **Name:** A default name (date and time of the creation of snapshot) is already added. You can edit this name and add a new name.
 - b. **Volume name:** The volume name is the file system name. It is already added to the form as per your selection.
 - c. **Subvolume group:** The subvolume group name is already added as per your selection. Alternatively, you select a different subvolume group from the dropdown list.
 - d. **Subvolume:** The subvolume name is already added as per your selection. Alternatively, you can select a different subvolume from the dropdown list.

7. Click **Create Snapshot**.

A notification is displayed that the snapshot is created successfully.

Verification

- Navigate to the **Snapshots** tab, select the subvolume and subvolume group for which the snapshot is created.

13.4.2. Deleting CephFS subvolume snapshots

As a storage administrator, you can create Ceph File System (CephFS) subvolume snapshot on the Red Hat Ceph Storage dashboard. You can create an immutable, point-in-time view of a Ceph File System (CephFS) by creating a snapshot.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A subvolume group with corresponding subvolumes.

Procedure

1. Log into the dashboard.
2. On the dashboard navigation menu, click **File > File Systems**.
3. Select the CephFS where you want to delete the subvolume snapshot.
4. Go to the **Snapshots** tab.
5. Select the snapshot you want to delete. You can list the snapshot by selecting the subvolume group and subvolume for which the snapshot is created.
6. Check the box 'Yes, I am sure' to confirm you want to delete the snapshot.
7. Click **Delete Snapshot**.
A notification is displayed that the snapshot is deleted successfully.

13.4.3. Cloning CephFS subvolume snapshots

As a storage administrator, you can clone Ceph File System (CephFS) subvolume snapshot on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A subvolume snapshot.

Procedure

1. Log into the dashboard.
2. On the dashboard navigation menu, click **File > File Systems**.

3. Select the CephFS where you want to clone the subvolume snapshot.
4. Go to the **Snapshots** tab.
5. Select the snapshot you want to clone. You can list the snapshot by selecting the subvolume group and subvolume for which the snapshot is created. Check the box 'Yes, I am sure' to confirm you want to delete the snapshot.
6. Click the arrow next to the Delete button.
7. Click **Clone** and configure the following.
 - a. **Name:** A default name (date and time of the creation of clone) is already added. You can change this name and add a new name.
 - b. **Group name:** The subvolume group name is already added as per your selection. Alternatively, you select a different subvolume group from the dropdown list.
8. Click Create Clone.
A notification is displayed that the clone is created successfully.
9. You can verify that the clone is created by going to the **Snapshots** tab, in the **Subvolume** column. Select the subvolume group for which the clone is created.

13.4.4. Creating CephFS volume snapshots

As a storage administrator, you can create Ceph File System (CephFS) volume snapshot on the IBM Storage Ceph dashboard. You can create an immutable, point-in-time view of a Ceph File System (CephFS) volume by creating a snapshot.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A subvolume group with corresponding subvolumes.

Procedure

1. Log into the dashboard.
2. On the dashboard navigation menu, click **File > File Systems**.
3. Select the CephFS where you want to create the subvolume snapshot.
If there are no file systems available, create a new file system.
4. Go to the Directories tab.
5. From the list of the subvolume and volumes, select the volume for which you want to create the snapshot.
6. Click Create in the Snapshot row. The Create Snapshot form opens.
 - a. **Name:** A default name (date and time of the creation of snapshot) is already added. You can edit this name and add a new name.
7. Click **Create Snapshot**.
A notification is displayed that the snapshot is created successfully.

8. You can verify that the snapshot is created from the **Snapshots** row in the **Directories** tab

13.4.5. Deleting CephFS volume snapshots

As a storage administrator, you can delete Ceph File System (CephFS) volume snapshot on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A volume snapshot.

Procedure

1. Log into the dashboard.
2. On the dashboard navigation menu, click **File > File Systems**.
3. Select the CephFS where you want to delete the subvolume snapshot.
4. Go to the **Directories** tab.
5. Select the snapshot you want to delete. You can see list of snapshots in the Snapshot row.
6. Click 'Delete'.
7. Check the box *Yes, I am sure* to confirm you want to delete the snapshot.
8. Click **Delete CephFS Snapshot**
A notification is displayed that the snapshot is deleted successfully.

13.5. SCHEDULING CEPHFS SNAPSHOTS

As a storage administrator, you can schedule Ceph File System (CephFS) snapshots on the Red Hat Ceph Storage dashboard.

Scheduling Ceph File System (CephFS) snapshots ensures consistent reliable backups at regular intervals, reducing the risk of data loss. Scheduling snapshots also provides ease of management by reducing administrative overhead of manually managing backups.

A list of available snapshots for a particular file system, subvolume, or directory is found on the File > File Systems page. Use the snapshot list for creating scheduled backups.

13.5.1. Creating CephFS snapshot schedule

As a storage administrator, you can create Ceph File System (CephFS) snapshot schedules on the Red Hat Ceph Storage dashboard.

Create a policy for automatic creation of Ceph File System (CephFS) snapshot of a volume or a certain directory. You can also define how often you want to schedule it (for example: every hour, every day, every week, etc.)

You can specify the number of snapshots or a period for which you want to keep the snapshots. All the older snapshots are deleted when the number of snapshots exceeds or when the retention period is over

Prerequisites

- A running Red Hat Ceph Storage cluster.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
File system volumes are listed. Select the file system where you want to create the snapshot schedule.
2. Go to **Snapshot schedules** tab.



NOTE

The **Enable** is available only if the **snapshot_scheduler** is disabled on the cluster. .. Optional: Click **Enable** to enable the **snapshot_scheduler** module. .. After enabling the scheduler, wait for the dashboard to reload and navigate back to the Snapshot schedule page. .. Click 'Create'.

- + The **Create Snapshot schedule** form opens.
 - a. Enter the directory name, start date, start time, and schedule.
 - i. **Directory**: You can search by typing in the path of the directory or subvolume and select the directory from the suggested list where you want to create the snapshot schedule.
 - ii. **Start date**: Enter date on which you want the scheduler to start creating the snapshots. By default the current date is selected.
 - iii. **Start time**: Enter the time at which you want the scheduler to start creating the snapshots. By default the current time is added.
 - iv. **Schedule**: Enter the number of snapshots and the frequency at which you want to create the snapshots. Frequency can be hourly, daily, weekly, monthly, yearly or latest snapshots.
- 3. Optional: Click **Add retention policy** if you want to add retention policy to the schedule. You can add multiple retention policies. Enter the number of snapshots and the frequency at which you want to retain the snapshots. Frequency can be hourly, daily, weekly, monthly, yearly or latest snapshots.
- 4. Click **Create snapshot schedule**
A notification is displayed that the snapshot schedule is created successfully.

13.5.2. Editing CephFS snapshot schedule

As a storage administrator, you can edit Ceph File System (CephFS) snapshot schedules on the IBM Storage Ceph dashboard. You can only edit the retention policy of the snapshot schedule. You can add another retention policy or delete the existing policy.

Prerequisites

- A running Red Hat Ceph Storage cluster.

Procedure

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. Select the CephFS where you want to edit snapshot schedules and click **Snapshot schedules**.
3. Select the snapshot schedule that you want to edit.
4. Click 'Edit'. An **Edit Snapshot schedule** dialog box appears.
5. Optional: From the **Edit Snapshot schedule** dialog, add a schedule retention policy, by selecting **Add retention policy**.
Enter the number of snapshots and the frequency at which you want to retain the snapshots. Frequency can be one of the following. .. Weekly or latest snapshots. For example, enter the number of snapshots as 10 and select 'latest snapshots' to retain last 10 snapshots irrespective of the frequency at which they were created. .. Daily .. Hourly .. Monthly .. Yearly
6. Click 'Edit snapshot schedule' to save.
A notification is displayed that the retention policy is created successfully.
7. Click the trash icon to delete the existing retention policy.
A notification is displayed that the retention policy is deleted successfully.

13.5.3. Deleting CephFS snapshot schedule

As a storage administrator, you can delete Ceph File System (CephFS) snapshot schedules on the IBM Storage Ceph dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A snapshot schedule.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. Select the CephFS where you want to edit snapshot schedules.
3. Click 'Snapshot schedules'.
4. Select the snapshot schedule that you want to delete.
5. Click 'Edit' from the dropdown menu.
6. Click 'Delete'.
A Delete Snapshot schedule dialog box appears.
7. Select *Yes, I am sure* to confirm if you want to delete the schedule.
8. Click 'Delete' snapshot schedule.
A notification is displayed that the snapshot schedule is deleted successfully.

13.5.4. Deactivating and activating CephFS snapshot schedule

As a storage administrator, you can deactivate and activate Ceph File System (CephFS) snapshot schedules on the IBM Storage Ceph dashboard. The snapshot schedule is activated by default and runs according to how it is configured. Deactivating excludes the snapshot from scheduling until it is activated again.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- A snapshot schedule.

Procedure

1. From the dashboard navigation, go to **File > File Systems**.
2. Select the CephFS where you want to edit snapshot schedules.
3. Click 'Snapshot schedules'.
4. Select the snapshot schedule that you want to deactivate and click 'Deactivate' from the action drop-down.
5. Select *Yes, I am sure* to confirm if you want to deactivate the schedule.
6. Click 'Deactivate snapshot schedule'.
A notification is displayed that the snapshot schedule is deactivated successfully.
7. You can activate the snapshot schedule by clicking 'Activate' in the action drop-down.

CHAPTER 14. MANAGING BLOCK DEVICES USING THE CEPH DASHBOARD

As a storage administrator, you can manage and monitor block device images on the Red Hat Ceph Storage dashboard. The functionality is divided between generic image functions and mirroring functions. For example, you can create new images, view the state of images mirrored across clusters, and set IOPS limits on an image.

14.1. MANAGING BLOCK DEVICE IMAGES ON THE CEPH DASHBOARD

As a storage administrator, you can create, edit, copy, purge, and delete images using the Red Hat Ceph Storage dashboard.

You can also create, clone, copy, rollback, and delete snapshots of the images using the Ceph dashboard.



NOTE

The Block Device images table is paginated for use with 10000+ image storage clusters to reduce Block Device information retrieval costs.

14.1.1. Creating images on the Ceph dashboard

You can create block device images on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, click **Create**.
3. Complete the **Create Images** form.
4. Optional: Click **Advanced** to set advanced parameters, such as **Striping** and **Quality of Service**.

Figure 14.1. Create a Block Device imageBlock / [Images](#) / Create

Create Image

Name (required)

Name...

Pool (required)

rbd



Mirroring

Allow data to be asynchronously mirrored between two Ceph clusters

You need to set **mirror mode** in the selected pool to enable mirroring.[Set mode](#)

Size (required)

e.g., 10GiB

Advanced

[Cancel](#)[Create Image](#)

5. Click **Create Image**.

A notification displays that the image was created successfully.

Additional Resources

- See the [Red Hat Ceph Storage Block Device Guide](#) for more information on Images.
- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.2. Creating namespaces on the Ceph dashboard

You can create namespaces for the block device images on the Red Hat Ceph Storage dashboard.

Once the namespaces are created, you can give access to the users for those namespaces.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Namespaces** tab, click **Create**.
3. In the **Create Namespace** dialog, select the pool and enter a name for the namespace.
4. Click **Create**.
A notification displays that the namespace was created successfully.

Figure 14.2. Create namespace



Additional Resources

- See the Knowledgebase article [Segregate Block device images within isolated namespaces](#) for more details.

14.1.3. Editing images on the Ceph dashboard

You can edit block device images on the Red Hat Ceph Storage dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. Go to **Edit** from the menu of the image that needs to be edited.
3. In the **Edit Image** form, edit the required parameters and click **Edit Image**.
A notification displays that the image was updated successfully.

Figure 14.3. Edit image

Block / [Images](#) / Edit

Edit Image

Name (required)

test_image1

Pool

test-pool_replicated



Mirroring

Allow data to be asynchronously mirrored between two Ceph clusters

You need to set **mirror mode** in the selected pool to enable mirroring.[Set mode](#)

Use a dedicated data pool

Use a dedicated pool to store the image data. If not selected, the image data will be stored in the same pool as the image metadata.

Size (required)

5 GiB

Advanced



Cancel

Edit Image

Additional Resources

- See the [Red Hat Ceph Storage Block Device Guide](#) for more information on Images.
- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.4. Copying images on the Ceph dashboard

You can copy block device images on the Red Hat Ceph Storage dashboard.

Prerequisites

Before you begin, make sure that you have the following prerequisites in place:

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.

- An image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, click **Copy** from the menu on the image that you want to copy.
3. In the **Copy Image** form, set the required parameters and click **Copy Image**.
A notification displays that the image was copied successfully.

Figure 14.4. Copy block image

Block / [Images](#) / Copy

Copy Image

Copy from

rbd/test_image

Name (required)

test_image1b

Pool (required)

rbd

☒ **Mirroring**
Allow data to be asynchronously mirrored between two Ceph clusters

☒ **Journal** ⓘ ☐ **Snapshot** ⓘ

☐ **Use a dedicated data pool**
Use a dedicated pool to store the image data. If not selected, the image data will be stored in the same pool as the image metadata.

Size (required)

5 GiB

Advanced

Cancel Copy Image

Additional Resources

- See the [Red Hat Ceph Storage Block Device Guide](#) for more information on Images.

- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.5. Moving images to trash on the Ceph dashboard

You can move the block device images to trash before it is deleted on the Red Hat Ceph Storage dashboard.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, select the image to move to trash, and click **Move to Trash** from the image menu.
3. In the **Move an image to trash** dialog, change the **Protection expires at** field and click **Move**. A notification displays that the image was moved to trash successfully.

Figure 14.5. Moving images to trash



14.1.6. Purging trash on the Ceph dashboard

You can purge trash using the Red Hat Ceph Storage dashboard.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.

- A pool with the *rbd* application enabled is created.
- An image is trashed.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Trash** tab, click **Purge Trash**.




IMPORTANT

To be able to restore an image, be sure that you set an expiry date before moving to trash.

3. In the **Purge Trash** dialog, select the pool, and click **Purge Trash**.
A notification displays that the pools in the trash were purged successfully.

Figure 14.6. Purge trash



Additional resources

- See the [Purging the Block Device Snapshots](#) section in the *Red Hat Ceph Storage Block Device Guide* for more details.

14.1.7. Restoring images from trash on the Ceph dashboard

You can restore the images that were trashed and has an expiry date on the Red Hat Ceph Storage Dashboard.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is trashed.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Trash** tab, select the row of the image to restore.
3. Click **Restore** from the protected image menu.

Figure 14.7. Restore images from the trash



4. In the **Restore Image** dialog, review the image name and click **Restore**.
A notification displays that the image was restored successfully.

Additional resources

- See the [Creating images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.

14.1.8. Deleting images on the Ceph Dashboard

You can delete the images from the cluster on the Ceph Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. Select the row to be deleted and click **Delete** from the action drop-down.
3. In the **Delete RBD** notification, select **Yes, I am sure** and click **Delete RBD**.
A notification displays that the image was deleted successfully.

Additional resources

- See the [Moving images to trash on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.

14.1.9. Deleting namespaces on the Ceph dashboard.

You can delete the namespaces of the images on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- A namespace is created in the pool.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Namespaces** tab, select the namespace and click **Delete** from the action drop-down.
3. In the **Delete Namespace** notification, select **Yes, I am sure** and click **Delete Namespace**.
A notification displays that the namespace was deleted successfully.

14.1.10. Creating snapshots of images on the Ceph dashboard

You can take snapshots of the Ceph block device images on the Red Hat Ceph Storage Dashboard.


Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row.
3. In the **Snapshots** tab, click **Create**.
4. In the **Create RBD Snapshot** dialog, enter the snapshot name and click **Create RBD Snapshot**

Figure 14.8. Creating snapshot of images



A notification displays that the snapshot was created successfully.

Additional Resources

- See the [Creating a block device snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information on creating snapshots.
- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the [Creating images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.11. Renaming snapshots of images on the Ceph dashboard

You can rename the snapshots of the Ceph block device images on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row.
3. In the **Snapshots** tab, click **Rename**.
4. In the **Rename RBD Snapshot** dialog, enter the new name and click **Rename RBD Snapshot**.

Additional Resources

- See the [Renaming a block device snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the [Creating images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.12. Protecting snapshots of images on the Ceph dashboard

You can protect the snapshots of the Ceph block device images on the Red Hat Ceph Storage Dashboard.

This is required when you need to clone the snapshots.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row and click on the **Snapshots** tab.
3. Select the snapshot to protect, and click **Protect** from the action drop-down.
The snapshot updates and the **State** changes from **Unprotected** to **Protected**.

Additional Resources

- See the [Protecting a block device snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.

14.1.13. Cloning snapshots of images on the Ceph dashboard

You can clone the snapshots of images on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.

- An image is created.
- A snapshot of the image is created and protected.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row.
3. In the **Snapshots** tab, select the snapshot to clone and click **Clone** from the action drop-down.
4. In the **Clone RBD** form, fill in the required details and click **Clone RBD**.
A notification displays that the snapshots was cloned successfully and the new image displays in the **Images** table.

Additional Resources

- See the [Protecting a Block device Snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the [Protecting snapshots of images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.14. Copying snapshots of images on the Ceph dashboard

You can copy the snapshots of images on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row.
3. In the **Snapshots** tab, select the snapshot to clone and click **Copy** from the action drop-down.
4. In the **Copy RBD** form, fill in the required details and click **Copy RBD**.
A notification displays that the snapshots was cloned successfully and the new image displays in the **Images** table.

Additional Resources

- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the [Creating images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.15. Unprotecting snapshots of images on the Ceph dashboard

You can unprotect the snapshots of the Ceph block device images on the Red Hat Ceph Storage Dashboard.

This is required when you need to delete the snapshots.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created and protected.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row and click on the **Snapshots** tab.
3. Select the protected snapshot, and click **Unprotect** from the action drop-down. The snapshot updates and the **State** changes from **Protected** to **Unprotected***.

Additional Resources

- See the [Unprotecting a block device snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the [Protecting snapshots of images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.16. Rolling back snapshots of images on the Ceph dashboard

You can rollback the snapshots of the Ceph block device images on the Red Hat Ceph Storage Dashboard. Rolling back an image to a snapshot means overwriting the current version of the image with data from a snapshot. The time it takes to execute a rollback increases with the size of the image. It is faster to clone from a snapshot than to rollback an image to a snapshot, and it is the preferred method of returning to a pre-existing state.

Prerequisites


- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.

- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row and click on the **Snapshots** tab.
3. Select the snapshot to rollback, and click **Rollback** from the snapshot menu.
4. In the **RBD snapshot rollback** dialog, click **Rollback**.

Figure 14.9. Rolling back snapshot of images



Additional Resources

- See the [Rolling a block device snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the [Creating pools on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the [Creating images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.1.17. Deleting snapshots of images on the Ceph dashboard

You can delete the snapshots of the Ceph block device images on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created and is unprotected.

Procedure

1. From the dashboard navigation, go to **Block→Images**.
2. From the **Images** tab, expand an image row.
3. In the **Snapshots** tab, select the snapshot to delete and click **Delete** from the action drop-down.
4. In the **Delete RBD Snapshot** dialog, select **Yes, I am sure** and click **Delete RBD Snapshot**. A notification displays that the snapshot was created successfully.

Additional Resources

- See the [Deleting a block device snapshot](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the [Unprotecting snapshots of images on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

14.2. MANAGING MIRRORING FUNCTIONS ON THE CEPH DASHBOARD

As a storage administrator, you can manage and monitor mirroring functions of the Block devices on the Red Hat Ceph Storage Dashboard.

You can add another layer of redundancy to Ceph block devices by mirroring data images between storage clusters. Understanding and using Ceph block device mirroring can provide you protection against data loss, such as a site failure. There are two configurations for mirroring Ceph block devices, one-way mirroring or two-way mirroring, and you can configure mirroring on pools and individual images.

14.2.1. Mirroring view on the Ceph dashboard

You can view the Block device mirroring on the Red Hat Ceph Storage Dashboard.

You can view the daemons, the site details, the pools, and the images that are configured for block device mirroring.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Mirroring is configured.

Procedure

- From the dashboard navigation, go to **Block→Mirroring**.

Figure 14.10. View mirroring of Ceph Block Devices

Block / Mirroring

Site Name
ceph-rbd1

Import Bootstrap Token Create Bootstrap Token

Daemons

Instance	ID	Hostname	Version	Health
No data to display				

Pools

Name	Mode	Leader	# Local	# Remote	Health
rbd	image				Warning
test-pool_replicated	disabled				Disabled
test_pool2	pool				OK

Images

Issues (2) Syncing (0) Ready (0)

Pool	Image	State	Issue
rbd	test_image	Unknown	
rbd	test_image3	Unknown	

Additional Resources

- For more information on mirroring, see [Mirroring Ceph block devices](#) section in the *Red Hat Ceph Storage Block Device Guide*.

14.2.2. Editing mode of pools on the Ceph dashboard

You can edit mode of the overall state of mirroring functions, which includes pools and images on the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- Mirroring is configured.

Procedure

- From the dashboard navigation, go to **Block→Mirroring**.
 - In the **Pools** table, select the pool to edit and click **Edit Mode**.
 - In the **Edit Mode** dialog, select the mode and click **Update**.
A notification displays that the mode was updated successfully and the **Mode** updates in the **Pools** table.

Additional Resources


- See the [Ceph Block Device Mirroring](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.

14.2.3. Adding peer in mirroring on the Ceph dashboard

You can add storage cluster peer for the **rbd-daemon** mirror to discover its peer storage cluster on the Red Hat Ceph Storage Dashboard.

Prerequisites

- Two healthy running Red Hat Ceph Storage clusters.
- Dashboard is installed on both the clusters.
- Pools created with the same name.
- **rbd** application enabled on both the clusters.



NOTE

Ensure that mirroring is enabled for the pool in which images are created.

Procedure

Site A

1. From the dashboard navigation, go to **Block → Mirroring** and click **Create Bootstrap Token**
2. From the *Navigation* menu, click the *Block* drop-down menu, and click *Mirroring*.
3. Click *Create Bootstrap Token* and configure the following in the window:

Figure 14.11. Create bootstrap token

Create Bootstrap Token

To create a bootstrap token which can be imported by a peer site cluster, provide the local site's name, select which pools will have mirroring enabled, and click **Generate**.

Site Name (required)

ceph-rbd2

Pools

☒ rbd

Generate

Token

eyJmc2lkjoINDdKNTQ1MzAtOTUwNS0xMWVmLTlmMzAtZmExNjNlMGIxMGEzliwiY2xpZW50X2lkjoicmJkLW1pcnJvci1wZWVylwiw2V5ljoIQVFEV1dSOW5pcU1EQ3hBQTllQUYwM3RhNjRMejEyOWE0Z2haM2c9PSlsm1vbl9ob3N0ljoIW3YyOjEwLjAuNjYuMjEzOjMzMdAvMCx2MToxMC4wLjY2LjlxMzo2Nzg5LzBdLFT2MjoxMC4wLjY2LjIOMjoxMzAwLzAsdJE6MTAuMC42Ni4yNDI6Njc4OS8wXSxbdjI6MTAuMC42Ny4yMTk6MzMwMC8wLHYxOjEwLjAuNjcuMjE5OjY3ODkvMF0ifQ==

Close

- a. For the provided site name, select the pools to be mirrored.
 - b. For the selected pools, generate a new bootstrap token by clicking **Generate**.
 - c. Click **Copy to Clipboard**.
 - d. Click **Close**.
4. Enable the pool mirror mode.
 - a. Select the pool.
 - b. Click **Edit Mode**.
 - c. In the **Edit pool mirror mode** dialog, select **Image** from the **Mode** list.
 - d. Click **Update**.
A notification displays that the pool was updated successfully.

Site B

1. From the dashboard navigation, go to **Block → Mirroring** and click **Import Bootstrap Token** from the action drop-down.



NOTE

Ensure that mirroring mode is enabled for the specific pool for which you are importing the bootstrap token.

2. In the **Import Bootstrap Token** dialog, select the direction, and paste the token copied earlier, from site A.

Figure 14.12. Import bootstrap token

Import Bootstrap Token

To import a bootstrap token which was created by a peer site cluster, provide the local site's name, select which pools will have mirroring enabled, provide the generated token, and click **Import**.

Site Name (required)

ceph-rbd1

Direction

Bidirectional

Pools

☒ rbd

Token (required)


eyJmc2lkjoiNDdkNTQ1MzAtOTUwNS0xMWVvLTlmMzAtZmExNjNlMGlxMGEzliwiY2xpZW50X2lkjoicmJkLW1pcnJvci1wZWVvY2V5ljoQVFEV1dSOW5pcU1EQ3hBQTllQUYwM3RhNjRMejEyOWE0Z2haM2c9PSlslm1vbl9ob3N0ljoiw3YyOjEwLjAuNjYuMjEzOjMzMdAvMCx2MToxMC4wLjY2LjlxMzo2Nzg5LzBdLFT2MjoxMC4wLjY2LjI0MjoxMzAwLzAsdjE6MTAuMC42Ni4yNDI6Njc4OS8wXSxbdjI6MTAuMC42Ny4yMTk6MzMwMwMC8wLHYxOjEwLjAuNjcuMjE5OjY3ODkvMF0ifQ==

Cancel Submit

3. Click **Submit**.
The peer is added and the images are mirrored in the cluster at site B.
4. On the **Block → Mirroring** page, in the **Pool** table, verify the health of the pool is in the *OK* state.

Site A

1. Create an image with *Mirroring* enabled.
 - a. From the dashboard navigation, go to **Block → Images**.
 - b. On the **Images** tab, click **Create**.
 - c. In the **Create Image** form, fill in the **Name** and **Size**.
 - d. Select **Mirroring**.



NOTE

Select mirroring with either **Journal** or **Snapshot**.

- e. Click **Create Image**.

Figure 14.13. Create mirroring image

Block / [Images](#) / Create

Create Image

Name (required)

test_image2

Pool (required)

rbd



Mirroring

Allow data to be asynchronously mirrored between two Ceph clusters



Journal ⓘ



Snapshot ⓘ



Use a dedicated data pool

Use a dedicated pool to store the image data. If not selected, the image data will be stored in the same pool as the image metadata.

Namespace

-- Select a namespace --

Namespace allows you to logically group RBD images within your Ceph Cluster. Choosing a namespace makes it easier to locate and manage related RBD images efficiently

Size (required)

10 GiB

Advanced

Cancel

Create Image

2. Verify the image is available at both the sites.
 - a. From the **Images** table, verify that the image in site A is set to **primary** and that the image in site B is set to **secondary**.

Additional Resources

- See the [Configuring two-way mirroring using the command-line interface](#) section in the *Red Hat Ceph Storage Block Device Guide* for more information.

14.2.4. Editing peer in mirroring on the Ceph dashboard

You can edit storage cluster peer for the **rbd-daemon** mirror to discover its peer storage cluster in the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rd* application enabled is created.
- An image is created.
- Mirroring is configured.
- A peer is added.

Procedure

- From the dashboard navigation, go to **Block→Mirroring**.
 - a. From the **Pools** table, click **Edit Peer** from the pool menu.
 1. In the **Edit pool mirror peer** dialog, edit the parameters, and click **Submit**.
A notification displays that the peer was updated successfully.

Figure 14.14. Editing peer in mirroring

Edit pool mirror peer ×

Edit the pool mirror peer attributes for pool **rd** and click **Submit**.

Cluster Name (required)
cluster1

CephX ID (required)
cluster-id_123

Monitor Addresses
Comma-delimited addresses...

CephX Key
Base64-encoded key...

Cancel Submit

Additional Resources

- See the [Adding peer in mirroring on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

14.2.5. Deleting peer in mirroring on the Ceph dashboard

You can edit storage cluster peer for the `rd-daemon`` mirror to discover its peer storage cluster in the Red Hat Ceph Storage Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- Mirroring is configured.
- A peer is added.

Procedure

1. From the dashboard navigation, go to **Block→Mirroring**.
2. From the **Pools** table, select the pool to edit and click **Delete Peer** from the action drop-down.
3. In the **Delete mirror peer** dialog, select **Yes, I am sure** and click **Delete mirror peer**.
A notification displays that the peer deleted successfully.

Additional Resources

- See the [Adding peer in mirroring on the Ceph dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

CHAPTER 15. ACTIVATING AND DEACTIVATING TELEMETRY

Activate the telemetry module to help Ceph developers understand how Ceph is used and what problems users might be experiencing. This helps improve the dashboard experience. Activating the telemetry module sends anonymous data about the cluster back to the Ceph developers.

View the telemetry data that is sent to the Ceph developers on the [public telemetry dashboard](#). This allows the community to easily see summary statistics on how many clusters are reporting, their total capacity and OSD count, and version distribution trends.

The telemetry report is broken down into several channels, each with a different type of information. Assuming telemetry has been enabled, you can turn on and off the individual channels. If telemetry is off, the per-channel setting has no effect.

Basic

Provides basic information about the cluster.

Crash

Provides information about daemon crashes.

Device

Provides information about device metrics.


Ident

Provides user-provided identifying information about the cluster.

Perf

Provides various performance metrics of the cluster.

The data reports contain information that help the developers gain a better understanding of the way Ceph is used. The data includes counters and statistics on how the cluster has been deployed, the version of Ceph, the distribution of the hosts, and other parameters.



IMPORTANT

The data reports do not contain any sensitive data like pool names, object names, object contents, hostnames, or device serial numbers.




NOTE

Telemetry can also be managed by using an API. For more information, see the [Telemetry](#) chapter in the *Red Hat Ceph Storage Developer Guide*.


Procedure

1. Activate the telemetry module in one of the following ways:
 - From the banner within the Ceph dashboard.



- Go to **Settings→Telemetry configuration**.

2. Select each channel that telemetry should be enabled on.



NOTE

For detailed information about each channel type, click **More Info** next to the channels.

3. Complete the **Contact Information** for the cluster. Enter the contact, Ceph cluster description, and organization.
4. Optional: Complete the **Advanced Settings** field options.

Interval

Set the interval by hour. The module compiles and sends a new report per this hour interval. The default interval is 24 hours.

Proxy

Use this to configure an HTTP or HTTPs proxy server if the cluster cannot directly connect to the configured telemetry endpoint. Add the server in one of the following formats:

<https://10.0.0.1:8080> or <https://ceph:telemetry@10.0.01:8080>

The default endpoint is **telemetry.ceph.com**.

5. Click **Next**. This displays the **Telemetry report preview** before enabling telemetry.
6. Review the **Report preview**.



NOTE

The report can be downloaded and saved locally or copied to the clipboard.

7. Select **I agree to my telemetry data being submitted under the Community Data License Agreement**.
8. Enable the telemetry module by clicking **Update**.
The following message is displayed, confirming the telemetry activation:

The Telemetry module has been configured and activated successfully

15.1. DEACTIVATING TELEMETRY

To deactivate the telemetry module, go to **Settings→Telemetry configuration** and click **Deactivate**.