



Red Hat Ceph Storage 6

Dashboard Guide

Monitoring Ceph Cluster with Ceph Dashboard

Red Hat Ceph Storage 6 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	5
1.1. CEPH DASHBOARD COMPONENTS	5
1.2. CEPH DASHBOARD FEATURES	5
1.3. RED HAT CEPH STORAGE DASHBOARD ARCHITECTURE	7
CHAPTER 2. CEPH DASHBOARD INSTALLATION AND ACCESS	9
2.1. NETWORK PORT REQUIREMENTS FOR CEPH DASHBOARD	10
2.2. ACCESSING THE CEPH DASHBOARD	12
2.3. EXPANDING THE CLUSTER ON THE CEPH DASHBOARD	13
2.4. TOGGING CEPH DASHBOARD FEATURES	16
2.5. UNDERSTANDING THE LANDING PAGE OF THE CEPH DASHBOARD	18
2.6. CHANGING THE DASHBOARD PASSWORD USING THE CEPH DASHBOARD	21
2.7. CHANGING THE CEPH DASHBOARD PASSWORD USING THE COMMAND LINE INTERFACE	22
2.8. SETTING ADMIN USER PASSWORD FOR GRAFANA	23
2.9. ENABLING RED HAT CEPH STORAGE DASHBOARD MANUALLY	25
2.10. CREATING AN ADMIN ACCOUNT FOR SYNCING USERS TO THE CEPH DASHBOARD	26
2.11. SYNCING USERS FROM RED HAT SIGN-ON TO THE CEPH DASHBOARD	28
2.12. ENABLING SINGLE SIGN-ON FOR THE CEPH DASHBOARD	34
2.13. DISABLING SINGLE SIGN-ON FOR THE CEPH DASHBOARD	35
CHAPTER 3. MANAGING ROLES ON THE CEPH DASHBOARD	37
3.1. USER ROLES AND PERMISSIONS ON THE CEPH DASHBOARD	37
3.2. CREATING ROLES ON THE CEPH DASHBOARD	39
3.3. EDITING ROLES ON THE CEPH DASHBOARD	42
3.4. CLONING ROLES ON THE CEPH DASHBOARD	43
3.5. DELETING ROLES ON THE CEPH DASHBOARD	44
CHAPTER 4. MANAGING USERS ON THE CEPH DASHBOARD	46
4.1. CREATING USERS ON THE CEPH DASHBOARD	46
4.2. EDITING USERS ON THE CEPH DASHBOARD	47
4.3. DELETING USERS ON THE CEPH DASHBOARD	49
4.4. USER CAPABILITIES	50
4.4.1. Capabilities	50
4.5. ACCESS CAPABILITIES	51
4.6. CREATING USER CAPABILITIES	52
4.7. EDITING USER CAPABILITIES	54
4.8. IMPORTING USER CAPABILITIES	55
4.9. EXPORTING USER CAPABILITIES	56
4.10. DELETING USER CAPABILITIES	57
CHAPTER 5. MANAGING CEPH DAEMONS	59
5.1. DAEMON ACTIONS	59
CHAPTER 6. MONITORING THE CLUSTER ON THE CEPH DASHBOARD	61
6.1. MONITORING HOSTS OF THE CEPH CLUSTER ON THE DASHBOARD	61
6.2. VIEWING AND EDITING THE CONFIGURATION OF THE CEPH CLUSTER ON THE DASHBOARD	63
6.3. VIEWING AND EDITING THE MANAGER MODULES OF THE CEPH CLUSTER ON THE DASHBOARD	64
6.4. MONITORING MONITORS OF THE CEPH CLUSTER ON THE DASHBOARD	65
6.5. MONITORING SERVICES OF THE CEPH CLUSTER ON THE DASHBOARD	67
6.6. MONITORING CEPH OSDS ON THE DASHBOARD	68
6.7. MONITORING HAPROXY ON THE DASHBOARD	70
6.8. VIEWING THE CRUSH MAP OF THE CEPH CLUSTER ON THE DASHBOARD	71

6.9. FILTERING LOGS OF THE CEPH CLUSTER ON THE DASHBOARD	72
6.10. VIEWING CENTRALIZED LOGS OF THE CEPH CLUSTER ON THE DASHBOARD	73
6.11. MONITORING POOLS OF THE CEPH CLUSTER ON THE DASHBOARD	76
6.12. MONITORING CEPH FILE SYSTEMS ON THE DASHBOARD	77
6.13. MONITORING CEPH OBJECT GATEWAY DAEMONS ON THE DASHBOARD	78
6.14. MONITORING BLOCK DEVICE IMAGES ON THE CEPH DASHBOARD.	79
CHAPTER 7. MANAGING ALERTS ON THE CEPH DASHBOARD	81
7.1. ENABLING MONITORING STACK	85
7.2. CONFIGURING GRAFANA CERTIFICATE	87
7.3. ADDING ALERTMANAGER WEBHOOKS	89
7.4. VIEWING ALERTS ON THE CEPH DASHBOARD	91
7.5. CREATING A SILENCE ON THE CEPH DASHBOARD	93
7.6. RE-CREATING 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	97
CHAPTER 8. MANAGING NFS GANESHA EXPORTS ON THE CEPH DASHBOARD	99
8.1. CONFIGURING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD	99
8.2. CONFIGURING NFS EXPORTS WITH CEPHFS ON THE CEPH DASHBOARD	102
8.3. EDITING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD	105
8.4. DELETING NFS GANESHA DAEMONS ON THE CEPH DASHBOARD	107
CHAPTER 9. MANAGING POOLS ON THE CEPH DASHBOARD	109
9.1. CREATING POOLS ON THE CEPH DASHBOARD	109
9.2. EDITING POOLS ON THE CEPH DASHBOARD	111
9.3. DELETING POOLS ON THE CEPH DASHBOARD	112
CHAPTER 10. MANAGING HOSTS ON THE CEPH DASHBOARD	115
10.1. ENTERING MAINTENANCE MODE	115
10.2. EXITING MAINTENANCE MODE	116
10.3. REMOVING HOSTS USING THE CEPH DASHBOARD	117
CHAPTER 11. MANAGING CEPH OSDS ON THE DASHBOARD	120
11.1. MANAGING THE OSDS ON THE CEPH DASHBOARD	120
11.2. REPLACING THE FAILED OSDS ON THE CEPH DASHBOARD	129
CHAPTER 12. MANAGING CEPH OBJECT GATEWAY USING THE DASHBOARD	133
12.1. MANUALLY ADDING CEPH OBJECT GATEWAY LOGIN CREDENTIALS TO THE DASHBOARD	133
12.2. CREATING THE CEPH OBJECT GATEWAY SERVICES WITH SSL USING THE DASHBOARD	134
12.3. CONFIGURING HIGH AVAILABILITY FOR THE CEPH OBJECT GATEWAY ON THE DASHBOARD	136
12.4. MANAGING CEPH OBJECT GATEWAY USERS ON THE DASHBOARD	137
12.4.1. Creating Ceph object gateway users on the dashboard	138
12.4.2. Creating Ceph object gateway subusers on the dashboard	139
12.4.3. Editing Ceph object gateway users on the dashboard	141
12.4.4. Deleting Ceph object gateway users on the dashboard	143
12.5. MANAGING CEPH OBJECT GATEWAY BUCKETS ON THE DASHBOARD	144
12.5.1. Creating Ceph object gateway buckets on the dashboard	144
12.5.2. Editing Ceph object gateway buckets on the dashboard	147
12.5.3. Deleting Ceph object gateway buckets on the dashboard	149
12.6. MONITORING MULTI-SITE OBJECT GATEWAY CONFIGURATION ON THE CEPH DASHBOARD	150
12.7. MANAGING BUCKETS OF A MULTI-SITE OBJECT CONFIGURATION ON THE CEPH DASHBOARD	151
12.7.1. Editing buckets of a multi-site object gateway configuration on the Ceph dashboard	151
12.7.2. Deleting buckets of a multi-site object gateway configuration on the Ceph dashboard	154

CHAPTER 13. MANAGING FILE SYSTEMS USING THE CEPH DASHBOARD	156
13.1. CONFIGURING CEPHFS VOLUMES	156
13.1.1. Creating CephFS volumes	156
13.1.2. Editing CephFS volumes	157
13.1.3. Removing CephFS volumes	157
13.2. CONFIGURING CEPHFS SUBVOLUME GROUPS	157
13.2.1. Creating CephFS subvolume groups	157
13.2.2. Editing CephFS subvolume groups	158
13.2.3. Removing CephFS subvolume groups	158
13.3. CONFIGURING CEPHFS SUBVOLUMES	159
13.3.1. Creating CephFS subvolume	159
13.3.2. Editing CephFS subvolume	160
13.3.3. Removing CephFS subvolume	161
13.4. MANAGING CEPHFS SNAPSHOTS	161
13.4.1. Creating CephFS subvolume snapshots	162
13.4.2. Deleting CephFS subvolume snapshots	163
13.4.3. Cloning CephFS subvolume snapshots	163
13.4.4. Creating CephFS volume snapshots	164
13.4.5. Deleting CephFS volume snapshots	165
13.5. SCHEDULING CEPHFS SNAPSHOTS	165
13.5.1. Creating CephFS snapshot schedule	165
13.5.2. Editing CephFS snapshot schedule	166
13.5.3. Deleting CephFS snapshot schedule	167
13.5.4. Deactivating and activating CephFS snapshot schedule	167
CHAPTER 14. MANAGING BLOCK DEVICES USING THE CEPH DASHBOARD	169
14.1. MANAGING BLOCK DEVICE IMAGES ON THE CEPH DASHBOARD	169
14.1.1. Creating images on the Ceph dashboard	169
14.1.2. Creating namespaces on the Ceph dashboard	170
14.1.3. Editing images on the Ceph dashboard	171
14.1.4. Copying images on the Ceph dashboard	172
14.1.5. Moving images to trash on the Ceph dashboard	173
14.1.6. Purging trash on the Ceph dashboard	174
14.1.7. Restoring images from trash on the Ceph dashboard	175
14.1.8. Deleting images on the Ceph Dashboard	176
14.1.9. Deleting namespaces on the Ceph dashboard.	177
14.1.10. Creating snapshots of images on the Ceph dashboard	178
14.1.11. Renaming snapshots of images on the Ceph dashboard	179
14.1.12. Protecting snapshots of images on the Ceph dashboard	180
14.1.13. Cloning snapshots of images on the Ceph dashboard	181
14.1.14. Copying snapshots of images on the Ceph dashboard	182
14.1.15. Unprotecting snapshots of images on the Ceph dashboard	184
14.1.16. Rolling back snapshots of images on the Ceph dashboard	184
14.1.17. Deleting snapshots of images on the Ceph dashboard	185
14.2. MANAGING MIRRORING FUNCTIONS ON THE CEPH DASHBOARD	186
14.2.1. Mirroring view on the Ceph dashboard	187
14.2.2. Editing mode of pools on the Ceph dashboard	187
14.2.3. Adding peer in mirroring on the Ceph dashboard	188
14.2.4. Editing peer in mirroring on the Ceph dashboard	193
14.2.5. Deleting peer in mirroring on the Ceph dashboard	195
CHAPTER 15. ACTIVATING AND DEACTIVATING TELEMETRY	197
15.1. DEACTIVATING TELEMETRY	198

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.

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. CEPH DASHBOARD FEATURES

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.
- **Single Sign-On (SSO):** The dashboard supports authentication with an external identity provider using the SAML 2.0 protocol.
- **Auditing:** The dashboard backend can be configured to log all PUT, POST and DELETE API requests in the Ceph manager log.


Management features

- **View 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 there is an issue with an OSD.
- **Configure manager modules:** You can view and change parameters for Ceph manager modules.
- **Embedded Grafana Dashboards** Ceph Dashboard Grafana dashboards might be embedded in external applications and web pages to surface information and performance metrics gathered by the Prometheus module.
- **View and filter logs** You can view event and audit cluster logs and filter them based on priority, keyword, date, or time range.
- **Toggle dashboard components:** You can enable and disable dashboard components so only the features you need are available.
- **Manage 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.
- **Viewing Alerts:** The alerts page allows you to see details of current alerts.
- **Quality of Service for images** You can set performance limits on images, for example limiting IOPS or read BPS burst rates.

Monitoring features

- **Username and password protection** You can access the dashboard only by providing a configurable user name and password.
- **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.
- **Hosts:** Provides a list of all hosts associated with the cluster along with the running services and the installed Ceph version.
- **Performance counters:** Displays detailed statistics for each running service.
- **Monitors:** Lists all Monitors, their quorum status and open sessions.
- **Configuration editor:** Displays all the available configuration options, their descriptions, types, default, and currently set values. These values are editable.
- **Cluster logs:** Displays and filters the latest updates to the cluster's event and audit log files by priority, date, or keyword.
- **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.
- **View storage cluster capacity:** You can view raw storage capacity of the Red Hat Ceph Storage cluster in the *Capacity* panels of the Ceph dashboard.

- **Pools:** Lists and manages all Ceph pools and their details. For example: applications, placement groups, replication size, EC profile, quotas, CRUSH ruleset, etc.
- **OSDs:** Lists and manages all OSDs, their status and usage statistics as well as detailed information like attributes, like OSD map, metadata, and performance counters for read and write operations. Lists all drives associated with an OSD.
- **Images:** Lists all 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*.

- **RBD Mirroring:** Enables and configures RBD mirroring to a remote Ceph server. Lists all active sync daemons and their status, pools and RBD images including their synchronization state.
- **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.
- **Object Gateway (RGW):** Lists all active object gateways and their performance counters. Displays and manages, including add, edit, delete, object gateway users and their details, for example quotas, as well as the users' buckets and their details, for example, owner or quotas.
- **NFS:** Manages NFS exports of CephFS and Ceph object gateway S3 buckets using the NFS Ganesha.

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).

Additional Resources

- See [Toggling Ceph dashboard features](#) in the *Red Hat Ceph Storage Dashboard Guide* for more information.

1.3. RED HAT CEPH STORAGE DASHBOARD ARCHITECTURE

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



153_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](#) .
- [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*.

Figure 2.2. Expand cluster



Welcome to Red Hat Ceph Storage Dashboard

Please expand your cluster first

Expand Cluster

Skip

5. Add hosts:

- a. In the *Add Hosts* window, click *+Add*.
- b. Provide the hostname. This is same as the hostname that was provided while copying the key from the bootstrapped host.

**NOTE**

You can use the tool tip in the *Add Hosts* dialog box for more details.

- 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.
- e. Click *Add Host*.
- f. Follow the above steps for all the hosts in the storage cluster.

6. In the *Add Hosts* window, click *Next*.

7. Create OSDs:


- a. In the *Create OSDs* window, 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: Click on the check-box *Encryption* to encrypt the features.
- f. In the *Create OSDs* window, click *Next*.

8. Create services:

- a. In the *Create Services* window, click *+Create*.
- b. In the *Create Service* dialog box,
 - i. Select the type of the service from the drop-down.
 - ii. Provide the service ID, a unique name of the service.
 - iii. Provide the placement by hosts or label.
 - iv. Select the hosts.
 - v. Provide the number of daemons or services that need to be deployed.
- c. Click *Create Service*.

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

10. Review the *Cluster Resources*, *Hosts by Services*, *Host Details*. If you want to edit any parameter, click *Back* and follow the above steps.

Figure 2.3. Review cluster

11. Click *Expand Cluster*.
12. You get a notification that the cluster expansion was successful.
13. The cluster health changes to *HEALTH_OK* status on the dashboard.

Verification

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

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. TOGGLING 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 Filesystem, **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, navigate to *Cluster* drop-down menu.
 - b. Select *Manager Modules*, and then select *Dashboard*.
 - c. In the *Edit Manager module* page, you can enable or disable the dashboard features by checking or unchecking the selection box next to the feature name.
 - d. Once the selections have been made, scroll down and 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.


2.5. 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 navigation bar provides the following options:

- Messages about tasks and notifications.
- Link to the documentation, Ceph Rest API, and details about the Red Hat Ceph Storage Dashboard.
- Link to user management and telemetry configuration.
- Link to change password and sign out of the dashboard.

Figure 2.4. Navigation bar




Apart from that, the individual panel displays specific information about the state of the cluster.

Categories

The landing page organizes panels into the following three categories:

1. **Status**
2. **Capacity**
3. **Performance**

Figure 2.5. Ceph dashboard landing page



Status panel

The status panels display the health of the cluster and host and daemon states.

Cluster Status: Displays the current health status of the Ceph storage cluster.

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

Monitors: Displays the number of Ceph Monitors and the quorum status.

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

Managers: Displays the number and status of the Manager Daemons.

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 Filesystems (CephFS).

Capacity panel

The capacity panel displays storage usage metrics.

Raw Capacity: Displays the utilization and availability of the raw storage capacity of the cluster.

Objects: Displays the total number of objects in the pools and a graph dividing objects into states of *Healthy*, *Misplaced*, *Degraded*, or *Unfound*.

PG Status: Displays the total number of Placement Groups and a graph dividing PGs into states of *Clean*, *Working*, *Warning*, or *Unknown*. To simplify display of PG states *Working* and *Warning* actually each encompass multiple states.

The *Working* state includes PGs with any of these 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 these states:

- backfill_toofull
- backfill_unfound
- down
- incomplete
- inconsistent
- recovery_toofull
- recovery_unfound
- remapped
- snaptrim_error
- stale
- undersized

Pools: Displays the number of storage pools in the Ceph cluster.

PGs per OSD: Displays the number of placement groups per OSD.

Performance panel

The performance panel display information related to data transfer speeds.

Client Read/Write: Displays total input/output operations per second, reads per second, and writes per second.

Client Throughput: Displays total client throughput, read throughput, and write throughput.

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.

Scrubbing: Displays whether Ceph is scrubbing data to verify its integrity.

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.

2.6. 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:

`https://HOST_NAME:8443`

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

Figure 2.6. User management



3. To change the password of **admin**, click it's row.
4. From the *Edit* drop-down menu, select *Edit*.
5. In the *Edit User* window, enter the new password, and change the other parameters, and then Click *Edit User*.

Figure 2.7. Edit user management

User management » Users » Edit

The screenshot shows the 'Edit User' form with the following fields and values:

- Username:** admin
- Password:** Password... (with an eye icon to toggle visibility)
- Confirm password:** Confirm password... (with an eye icon to toggle visibility)
- Password expiration date:** Password expiration date... (with a close icon)
- Full name:** Full name...
- Email:** Email...
- Roles:** administrator (with a plus icon to add more roles)

At the bottom right, there are two buttons: 'Cancel' and 'Edit User'.

You will be logged out and redirected to the log-in screen. A notification appears confirming the password change.

2.7. 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.

2.8. 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.

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

Example

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

4. Optional: Enable the **dashboard** Ceph Manager module:

Example

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

5. 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
```

6. 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/
```

2.9. 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*.

2.10. 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.

- Red hat Single Sign-On installed from a ZIP file. See the [Installing Red Hat Single Sign-On from a zip file](#) for additional information.

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. 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.
5. Optional: For Red Hat Enterprise Linux 8, users might get Certificate Authority (CA) issues. Import the custom certificates from CA and move them into the keystore with the exact java version.

Example

```
[root@host01 ~]# keytool -import -noprompt -trustcacerts -alias ca -file ../ca.cer -keystore
/etc/java/java-1.8.0-openjdk/java-1.8.0-openjdk-1.8.0.272.b10-
3.el8_3.x86_64/lib/security/cacert
```

6. To 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. SYNCING USERS FROM RED HAT SIGN-ON TO THE CEPH DASHBOARD

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.
- Ceph Dashboard is installed.
- Admin level access to the dashboard.
- Users 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*:


The screenshot shows the 'RED HAT SINGLE SIGN-ON' 'Add realm' window. It features a 'Select realm' dropdown on the left. The main area is titled 'Add realm' and contains the following fields and controls:

- Import:** A button with a file icon labeled 'Select file'.
- Name:** A text input field containing 'Ceph_LDAP' with a red asterisk indicating it is required.
- Enabled:** A toggle switch currently set to 'ON'.
- Buttons:** 'Create' and 'Cancel' buttons at the bottom right.

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
Client Protocol	saml	saml
Include AuthnState ment	ON	ON

Name of the parameter	Syntax	Example
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

Name of the parameter	Value
UUID LDAP attribute	nsuniqueid
User Object Classes	inetOrgPerson, 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*:

Trust Email

Use Truststore SPI Only for Idaps

Connection Pooling ☒

Connection Timeout

Read Timeout

Pagination ☒

Kerberos Integration

Allow Kerberos authentication

Use Kerberos For Password Authentication

Sync Settings

Batch Size

Periodic Full Sync


Periodic Changed Users Sync

Cache Settings

Cache Policy

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 Name]

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*.

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.

2.12. ENABLING SINGLE SIGN-ON 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 StorageDashboard Guide*.

2.13. DISABLING SINGLE SIGN-ON FOR THE CEPH DASHBOARD

You can disable single sign-on for Ceph Dashboard using the SAML 2.0 protocol.

Prerequisites

- 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 for the Ceph Dashboard](#) in the *Red Hat Ceph Storage Dashboard 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.

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

☐ All	☐ Read	☐ Create	☐ Update	☐ Delete
☐ cephfs	☐	☐	☐	☐
☐ config-opt	☐	☐	☐	☐
☐ dashboard-settings	☐	☐	☐	☐
☐ grafana	☐	☐	☐	☐
☐ hosts	☐	☐	☐	☐
☐ log	☐	☐	☐	☐
☐ manager	☐	☐	☐	☐
☐ monitor	☐	☐	☐	☐
☐ nfs-ganesha	☐	☐	☐	☐
☐ osd	☐	☐	☐	☐
☐ pool	☐	☐	☐	☐
☐ prometheus	☐	☐	☐	☐
☐ rbd-image	☐	☐	☐	☐
☐ rbd-mirroring	☐	☐	☐	☐
☐ rgw	☐	☐	☐	☐
☐ user	☐	☐	☐	☐

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	
+ Create ▾	
Name 	Description 
> administrator	Administrator
> block-manager	Block Manager
> cephfs-manager	CephFS Manager
> cluster-manager	Cluster Manager
> ganesha-manager	NFS Ganesha Manager
> pool-manager	Pool Manager
> read-only	Read-Only
> rgw-manager	RGW Manager

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 *

NFS-RGW ✓

Description

With this role, you can manage NFS Ganesha and RGW operations ✓

Permissions

☐ All	☐ Read	☐ Create	☐ Update	☐ Delete
☐ cephfs	☐	☐	☐	☐
☐ config-opt	☐	☐	☐	☐
☐ dashboard-settings	☐	☐	☐	☐
☒ grafana	☒	☒	☒	☒
☐ hosts	☐	☐	☐	☐
☐ log	☐	☐	☐	☐
☐ manager	☐	☐	☐	☐
☐ monitor	☐	☐	☐	☐
☒ nfs-ganesha	☒	☒	☒	☒
☐ osd	☐	☐	☐	☐
☐ pool	☐	☐	☐	☐
☐ prometheus	☐	☐	☐	☐
☐ rbd-image	☐	☐	☐	☐
☐ rbd-mirroring	☐	☐	☐	☐
☒ rgw	☒	☒	☒	☒
☐ user	☐	☐	☐	☐

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

ganesha-manager_clone

Description

allows full permissions for the nfs-ganesha scope

Permissions

☐ All
☐ Read
☐ Create
☐ Update
☐ Delete

☐ cephfs	☐	☐	☐	☐
☐ config-opt	☐	☐	☐	☐
☐ dashboard-settings	☐	☐	☐	☐
☐ grafana	☒	☐	☐	☐
☐ hosts	☐	☐	☐	☐
☐ log	☐	☐	☐	☐
☐ manager	☐	☐	☐	☐
☐ monitor	☐	☐	☐	☐
☒ nfs-ganesha	☒	☒	☒	☒
☐ osd	☐	☐	☐	☐
☐ pool	☐	☐	☐	☐
☐ prometheus	☐	☐	☐	☐
☐ rbd-image	☐	☐	☐	☐
☐ rbd-mirroring	☐	☐	☐	☐
☒ rgw	☒	☒	☒	☒
☐ user	☐	☐	☐	☐

Cancel

Edit Role

5. 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*.

Clone Role

×

New name *

ganesha-manager_clone

Cancel

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 click *User management*.



3. On *Roles* tab, click the role you want to delete.
4. Select *Delete* from the *Edit* drop-down menu.
5. In the *Delete Role* dialog box, Click the *Yes, I am sure* box and then 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

 A screenshot of the 'Create User' form. The form has the following fields and options:

- Username ***: A text input field containing 'dashboard_user' with a green checkmark on the right.
- Password ?**: A password input field with a green checkmark and an eye icon on the right.
- Confirm password**: A password input field with a green checkmark and an eye icon on the right.
- Password expiration date ?**: A date input field with a green checkmark and a close icon on the right.
- Full name**: A text input field containing 'Dashboard user'.
- Email**: A text input field containing 'dashboarduser@example.com'.
- Roles**: A section with a blue pencil icon and the text 'There are no roles.'
- Enabled**: A checked checkbox.
- User must change password at next logon**: A checked checkbox.

 At the bottom right, there are two buttons: 'Cancel' and '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

dashboard_user

Password ⓘ

👁

Confirm password

👁

Password expiration date ⓘ

✕

Full name

Dashboard user

Email

dashboarduser@example.com

Roles

✎ read-only ✕

☒ Enabled

☒ User must change password at next logon

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* dialog box, Click the *Yes, I am sure* box and then Click *Delete User* to save the settings.

Delete User



Are you sure that you want to delete **dashboard_user**?

☒ Yes, I am sure.

Cancel

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. Log in to the Dashboard.
2. Under the Cluster drop-down menu, select *Users*.
3. Click the *+ Create*.
4. In the *Create User* window, provide the following details:
 - a. User entity - Provide this as **TYPE.ID**.
 - b. Entity - This can be **mon**, **mgr**, **osd**, or **mds**.
 - c. Entity capabilities - Provide details of the different capabilities that you can 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.



5. Click *Create User*.

6. You get a notification 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. Log in to the Dashboard.
2. Under the Cluster drop-down menu, select *Users*.
3. Select the user whose roles you want to edit.
4. Click *Edit*.
5. In the *Edit User* window, edit the **entity** and **entity capabilities**.



NOTE

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

The screenshot shows the 'Edit User' form in the Red Hat Ceph Storage dashboard. The form is titled 'Edit User' and is located under the 'Cluster > Users' path. It features a 'User entity' field with the value 'client.test11'. Below this is a 'Capabilities' section with four rows, each containing an 'Entity' field and an 'Entity Capabilities' field. The entities are 'mds', 'mgr', 'mon', and 'osd'. The capabilities are 'allow *', 'allow *', 'allow *', and 'allow r' respectively. Each row has a '+' button to add more capabilities and a trash icon to remove the row. At the bottom of the form are 'Cancel' and 'Edit User' buttons.

6. Click *Edit User*.

7. You get a notification that the user was edited successfully.

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 = AQD9S29kmjgJFxAkvhFar6Af3AWKDY2DsULRg==
caps mds = "allow *"
caps mgr = "allow *"
caps mon = "allow *"
caps osd = "allow r"
```

2. Log in to the Dashboard.
3. Under the Cluster drop-down menu, select *Users*.
4. Select the user whose roles you want to import.
5. From the *Edit* drop-down menu, select *Import*.
6. In the *Import User* window, click *Choose file*, select the appropriate file.
7. Click *Import User*

Cluster » Users

8. You get a notification that the keys are imported successfully.

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. Log in to the Dashboard.
2. Under the Cluster drop-down menu, select *Users*.

3. Select the user whose roles you want to export.
4. From the *Edit* drop-down menu, select *Export*.
5. In the *Ceph user export data* dialog box, click the *Copy to Clipboard* icon.



6. You get a notification that the keys are copied successfully.
7. On your local system, create a keyring file and paste the keys:

Example

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

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

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. Log in to the Dashboard.

2. Under the Cluster drop-down menu, select *Users*.
3. Select the user you want to delete.
4. From the *Edit* drop-down menu, select *Delete*.
5. In the *Delete entity* window, select *Yes, I am sure*.



6. Click *Delete entity*.
7. You get a notification that the user was 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

- 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. Log in to the dashboard.
2. From the *Cluster* drop-down menu, select *Services*.
3. View the details of the service with the daemon to perform the action on by clicking the *Expand/Collapse* icon on its row.
4. In *Details*, select the drop down next to the desired daemon to perform *Start*, *Stop*, *Restart*, or *Redeploy*.

Figure 5.1. Managing daemons



From the Hosts page:

1. Log in to the dashboard.
2. From the *Cluster* drop-down menu, select *Hosts*.
3. From the *Hosts List*, select the host with the daemon to perform the action on.

4. In the *Daemon* tab of the host, click the daemon.
5. Use the drop down at the top to perform *Start*, *Stop*, *Restart*, or *Redeploy*.

Figure 5.2. Managing daemons

The screenshot shows the Red Hat Ceph Storage 6 Dashboard. The left sidebar contains navigation links: Dashboard, Cluster, Hosts, Physical Disks, Monitors, Services, OSDs, Configuration, CRUSH map, Manager Modules, Logs, Monitoring, Pools, and Block. The main content area is titled 'Cluster > Hosts' and has tabs for 'Hosts List' and 'Overall Performance'. The 'Hosts List' tab is active, showing a table of hosts. The first host is 'magna08lceph.redhat.com' with labels 'mgr', 'mon', and 'admin'. Below the host list, the 'Daemons' tab is selected, showing a table of daemons. The table has columns: Daemon name, Version, Status, Last Refreshed, CPU Usage, Memory Usage, and Daemon Events. Two daemons are listed: 'alertmanager' and 'crash.magna08l'. The 'alertmanager' daemon is running, has a version of '10.2.8-68.el8cp', a status of 'running', last refreshed 'A few seconds ago', CPU usage of 2%, and memory usage of 31.0 MB. The 'crash.magna08l' daemon is also running, has a version of '10.2.8-68.el8cp', a status of 'running', last refreshed 'A few seconds ago', CPU usage of 4%, and memory usage of 7 MB. The dashboard includes a sidebar with navigation links and a top navigation bar with tabs for Hosts List and Overall Performance.

Hostname	Service Instances	Labels	Status	Model	CPU	Cores	Total Memory	Raw Capacity	HDDs	Flash	NICs
magna08lceph.redhat.com	mgr, mon	_admin, mgr, mon	To be filled by O.E.M. (SYS-F627R3-RTB*)		1	6	31.2 GiB	3.6 TiB	4	0	2

Daemon name	Version	Status	Last Refreshed	CPU Usage	Memory Usage	Daemon Events
alertmanager	10.2.8-68.el8cp	running	A few seconds ago	2%	31.0 MB	5 hours ago - daemon.alertmanager.magna08l Reconfigured alertmanager.magna08l on host 'magna08lceph.redhat.com'
crash.magna08l	10.2.8-68.el8cp	running	A few seconds ago	4%	7 MB	5 hours ago - daemon.crash.magna08l Reconfigured crash.magna08l on host

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:

- **Devices** – This tab has details such as device ID, state of health, device name, and the daemons on the hosts.
- **Inventory** – This tab shows 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.
- **Daemons** – This tab shows all services that have been deployed on the selected host, which container they are running in and their current status. It has details such as hostname, daemon type, daemon ID, container ID, container image name, container image ID, version status and last refreshed time.

- **Performance details** - This tab has details such as OSDs deployed, CPU utilization, RAM usage, network load, network drop rate, and OSD disk performance statistics.
- **Device health** - For SMART-enabled devices, you can get the individual health status and SMART data only on the OSD deployed hosts.

Prerequisites

- 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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Hosts*.
3. To view the details of a specific host, click the *Expand/Collapse* icon on it's row.
4. You can view the details such as *Devices*, *Inventory*, *Daemons*, *Performance Details*, and *Device Health* by clicking the respective tabs.

Figure 6.1. Monitoring hosts of the Ceph cluster

Cluster » Hosts

Hosts List Overall Performance

Edit

Hostname	Services
> ceph-adm1	mgr.ceph-adm1.ubzhck,
▼ ceph-adm2	mgr.ceph-adm2.zecld,

Devices Inventory Daemons Performance Details Device health

Device ID	State of Health
QEMU_QEMU_HARDDISK_33ad2f18-e635-4288-bcb0-611c0a5fcc9b	Unknown
QEMU_QEMU_HARDDISK_5f378165-1e92-497b-84bf-429360e661ea	Unknown

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

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

Procedure

1. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Configuration*.
3. Optional: You can search for the configuration using the *Search* box:
4. Optional: You can filter for a specific configuration using 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
5. To view the details of the configuration, click the *Expand/Collapse* icon on it's row.

Figure 6.2. Configuration options

Cluster » Configuration



Name	Description
> client_cache_size	soft maximum number of directory entries in client cache
> cluster_addr	cluster-facing address to bind to
▼ container_image	container image (used by cephadm orchestrator)
Name	container_image
Description	container image (used by ce
Long description	

6. To edit a configuration, click its row and click *Edit*.
 - a. In the edit dialog window, edit the required parameters and Click *Update*.
7. You get a notification 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

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

Viewing the manager modules

1. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Manager Modules*.

3. To view the details of a specific manager module, click the *Expand/Collapse* icon on it's row.

Figure 6.3. Manager modules

Cluster » Manager Modules

Edit

Name

>	alerts
▼	balancer
active	true
begin_time	0000
begin_weekday	0
crush_compat_max_iterations	25
crush_compat_metrics	pgs,objects,bytes
crush_compat_step	0.5

Enabling a manager module

1. Select the row.
2. From the *Edit* drop-down menu, select *Enable*.

Disabling a manager module

1. Select the row.
2. From the *Edit* drop-down menu, select *Disable*.

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*.
3. You get a notification 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* tab.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Monitors*.
3. The *Monitors* overview page displays information about the overall monitor status as well as tables of *In Quorum* and *Not in quorum* Monitor hosts.

Cluster > Monitors

Status		In Quorum		
Cluster ID	62a081d6-88aa-7eb-a367-001a4a000672			
monmap modified	A day ago			
monmap epoch	4			
quorum con	4540138297036906239			
quorum mon	kraken,luminous,mimic,osdmep-prune,nautilus,octopus,pacific,elector-pinging			
required con	244995874737026820			
required mon	kraken,luminous,mimic,osdmep-prune,nautilus,octopus,pacific,elector-pinging			

Name	Rank	Public Address	Open Sessions
ceph-adm2	0	10.74.249.163:6789/0	0
ceph-adm3	1	10.74.254.129:6789/0	0

2 total

4. To see the number of open sessions, hover the cursor over the blue dotted trail.
 5. To see performance counters for any monitor, click its hostname.
- View the performance counter of the monitor:

Cluster > Monitors > Performance Counters

mon.ceph-adm2

Name	Description	Value
.cache_bytes	current memory available for caches.	1020054731
.heap_bytes	aggregate bytes in use by the heap	952180736
.mapped_bytes	total bytes mapped by the process	942153728
.target_bytes	target process memory usage in bytes	2147483648
.unmapped_bytes	unmapped bytes that the kernel has yet to reclaim	10027008
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	2

98 total

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

- 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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Services*.
3. To view the details of a specific service, click the *Expand/Collapse* icon on it's row.

Figure 6.4. Monitoring services of the Ceph cluster

Cluster » Services

Delete

Service  Placement 

 alertmanager count:1;label:monitoring

Hostname 	Daemon type 	Daemon ID 	Container ID 
ceph-sangadi-1624428547341-node2-mon-mgr-osd-grafana	alertmanager	ceph-sangadi-1624428547341-node2-mon-mgr-osd-grafana	1d6fdbdf98f7

1 total

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.

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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *OSDs*.
3. To view the details of a specific OSD, click the *Expand/Collapse* icon on it's row.

Figure 6.5. 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
☒ v 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_PHKE91360037375A	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. Launch the Ceph dashboard and log in with your credentials.

Example


```
https://dashboard_url:8443
```

4. From the *Cluster* drop-down menu, select *Object Gateway*.
5. Select *Daemons*.
6. Select the *Overall Performance* tab.

Verification

- Verify the Ceph Object Gateway HAProxy metrics:

Figure 6.6. 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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *CRUSH Map*.
3. To view the details of the specific OSD, click it's row.

Figure 6.7. 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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Logs*.

Figure 6.8. Cluster logs

Cluster » Logs

Priority: All ▼ Keyword: Q ✕ Date: YYYY-MM-DD

Cluster Logs Audit Logs

```

5/27/21 7:31:00 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm12...
5/27/21 7:31:00 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm11...
5/27/21 7:31:00 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm10...
5/27/21 7:30:00 PM [INF] overall HEALTH_OK
5/27/21 7:29:47 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm12...
5/27/21 7:29:47 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm11...
5/27/21 7:29:47 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm10...
5/27/21 7:28:40 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm12...
5/27/21 7:28:40 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm11...
5/27/21 7:28:40 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm10...
5/27/21 7:27:34 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm12...
5/27/21 7:27:34 PM [INF] Applying service osd.dashboard-admin-1621424607412 on host ceph-adm11...

```

- To filter by priority, click the **Priority** drop-down menu and select either *Debug*, *Info*, *Warning*, *Error*, or *All*.
 - To filter by keyword, enter text into the **Keyword** field.
 - To filter by date, click the **Date** field and either use the date picker to select a date from the menu, or enter a date in the form of *YYYY-MM-DD*.
 - To filter by time, enter a range in the **Time range** fields using the *HH:MM - HH:MM* format. Hours must be entered using numbers **0** to **23**.
 - To combine filters, set two or more filters.
- Click the *Download* icon or *Copy to Clipboard* icon to download the logs.

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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, click *Services*.
3. In the *Services* window, click *+ Create* from the drop-down menu.
4. In the *Create Service* window, from the *Type* field, choose **loki**, fill in the remaining details, and click *Create Service*.
5. Repeat the previous step to create the **Promtail** service as well by choosing **promtail** in the *Type* field. On successful creation, you can see **loki** and **promtail** services running in the list of services.

Figure 6.9. Creating Loki and Promtail services

Create Service

×

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.

6. To enable logging to files, from the *Cluster* drop-down menu, click *Configuration*.
7. Search for **log_to_file** in the search bar and click *Edit*.
8. In the *Edit log_to_file* window, set *global* to **true**.

Figure 6.10. 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

9. Repeat steps 6 to 8 to configure *global* to **true** for the **mon_cluster_log_to_file** file as well, by searching for **mon_cluster_log_to_file** in the search bar.




NOTE

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

10. Navigate to *Logs* under *Cluster* and click *Daemon Logs* tab to view the centralized logs. Use *Log browser* to select files and click *Show logs* to view the logs from that file.

Figure 6.11. View centralized logs



NOTE

If you do not see the logs, you need to sign in to Grafana and reload the page.

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. Log in to the Dashboard.
2. On the navigation bar, select *Pools*.
3. View the pools list 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.
4. To view more information about a pool, click the *Expand/Collapse* icon on it's row.

Figure 6.12. Monitoring pools



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. There are four main tabs in *File Systems*:

- *Details* - View the metadata servers (MDS) and their rank plus any standby daemons, pools and their usage, and performance counters.
- *Clients* - View list of clients that have mounted the file systems.
- *Directories* - View list of directories.
- *Performance* - View the performance of the file systems.

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. Log in to the dashboard.
2. On the navigation bar, click *Filesystems*.

- 3. To view more information about the file system, click the *Expand/Collapse* icon on it's row.

Figure 6.13. Monitoring Ceph File Systems



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. Log in to the dashboard.
2. On the navigation bar, click *Object Gateway*.

- To view more information about the Ceph object gateway daemon, click the *Expand/Collapse* icon on it's row.

Figure 6.14. Monitoring Ceph object gateway daemons

Selected Object Gateway: rgw.def2.new.magna123.byabeh (us)

Object Gateway » Daemons

Daemons List Overall Performance Sync Performance

ID 	Hostname 
▼ rgw.def.magna125.ndsaiv	magna125

Details Performance Counters Performance Details

arch	x86_64
ceph_release	pacific
ceph_version	ceph version 16.2.0-46.el8cp (66a64d4C
ceph_version_short	16.2.0-46.el8cp
container_hostname	magna125
container_image	registry-proxy.engineering.redhat.com/rh
cpu	Intel(R) Xeon(R) CPU E5-2620 v2 @ 2.10
distro	rhel

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. Log in to the dashboard.
2. On the navigation bar, click *Block*.
3. To view more information about the images, click the *Expand/Collapse* icon on it's row.

Figure 6.15. Monitoring Block device images

Block » Images

Images Namespaces Trash Overall Performance

Edit

Name	Pool	Namespace
test_image	pool_test_1	

Details Snapshots Configuration Performance

Name	test_image
Pool	pool_test_1
Data Pool	-
Created	6/3/21 3:23:42 PM
Size	10 GiB
Objects	2.6 k
Object size	4 MiB
Features	deep-flatten exclusive-lock
Provisioned	0 B

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:

- 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
```

- 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 'http://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.



NOTE

SSL is not supported in Red Hat Ceph Storage 6 cluster.

Prerequisite

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

Procedure

1. Log in to the Dashboard.
2. Customize the *alerts* module on the dashboard to get an email alert for the storage cluster:
 - a. On the navigation menu, click *Cluster*.
 - b. Select *Manager modules*.
 - c. Select *alerts* module.
 - d. In the *Edit* drop-down menu, select *Edit*.
 - e. In the *Edit Manager module* window, update the required parameters and click *Update*.



NOTE

Do not select the **smtp_ssl** parameter.

Figure 7.2. Edit Manager module for alerts

Cluster » Manager Modules » Alerts » Edit

Edit Manager module

interval ?

5

log_level

log_to_cluster

☒

log_to_cluster_level

info

log_to_file

☒

smtp_destination ?

test@test.com

✓

smtp_from_name ?

Ceph_Dashboard

✓

smtp_host ?

smtp.corp.redhat.com

✓

smtp_password ?

✓

smtp_port ?

25

✓

smtp_sender ?

ceph_test@redhat.com

✓

smtp_ssl ?

☐

smtp_user ?

Test_user

✓

Cancel

Update

- 3. On the navigation menu, click *Cluster*.
- 4. Select *Monitoring* from the drop-down menu.
- 5. To view details of the alert, click the *Expand/Collapse* icon on it's row.

Figure 7.3. Viewing alerts

Cluster » Monitoring » Active Alerts

Active Alerts

Alerts

Silences

+ Create Silence

Name	Job	Severity	State
▼ CephHealthWarning	ceph	warning	active

alername

CephHealthWarning

description

Ceph has been in HEALTH_WARN for more than 15 minutes. Please check "ceph health detail"

endsAt

8/31/22 2:03:10 PM

fingerprint

5a89c41212847278

- 6. 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. Log in to the Dashboard.
2. On the navigation menu, click *Cluster*.
3. Select *Monitoring* from the drop-down menu.
4. To create silence for an alert, select it's row.
5. Click *+Create Silence*.
6. In the *Create Silence* window, Add the details for the *Duration* and click *Create Silence*.

Figure 7.4. Create Silence

Cluster » Monitoring » Silences » Create

Create Silence

Creator *

sunil

Comment *

Start time * ⓘ

2021-06-14 17:07

Duration *

2h

End time *

2021-06-14 19:07

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

Create Silence

7. You get a notification that the silence was created successfully.

7.6. RE-CREATING A SILENCE ON THE CEPH DASHBOARD

You can re-create 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. Log in to the Dashboard.

94

2. On the navigation menu, click *Cluster*.
3. Select *Monitoring* from the drop-down menu.
4. Click the *Silences* tab.
5. To recreate an expired silence, click it's row.
6. Click the *Recreate* button.
7. In the *Recreate Silence* window, add the details and click *Recreate Silence*.

Figure 7.5. Recreate silence

Cluster » Monitoring » Silences » Recreate

Recreate 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

Recreate Silence

8. You get a notification that the silence was recreated successfully.

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 7.6. Edit silence

Cluster » Monitoring » Silences » Edit

Edit Silence ?

Creator *

admin

Comment *

test

Start time * ?

2021-06-14 17:07

Duration *

2m

✓

End time *

2021-06-14 17:09

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. 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 expire a silence, click it's row.
6. In the *Edit* drop-down menu, select *Expire*.
7. In the *Expire Silence* dialog box, select *Yes, I am sure* , and then click *Expire Silence*.

Figure 7.7. Expire Silence



Expire Silence

Are you sure that you want to expire **ffdb900f-2fff-4959-810c-cdd09c7b9f57**?

☒ **Yes, I am sure.**

Cancel Expire Silence

8. You get a notification that the silence was expired successfully.

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.



NOTE

Red Hat Ceph Storage 5 supports only NFSv4 protocol.

Prerequisites

- 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. On the dashboard landing page, click *NFS*.
6. Select *Create*.
7. In the *Create NFS export* window, set the following parameters and click *Create NFS export*:
 - a. Cluster - Name of the cluster.
 - b. Daemons - You can select all daemons.
 - c. Storage Backend - You can select Object Gateway.
 - d. Object Gateway User - Select the user created. In this example, it is test_user.
 - e. Path - Any directory.
 - f. NFS Protocol - NFSv4 is selected by default.
 - g. Pseudo - root path
 - h. Access Type - The supported access types are RO, RW, and NONE.
 - i. Squash
 - j. Transport Protocol

k. Clients

NFS » Create

Create NFS export

Cluster *

rgw-nfs

Daemons

+ Add all daemons

Storage Backend *

-- Select the storage backend --

NFS Protocol *

☐ NFSv3
 ☒ NFSv4

Pseudo * ?

Access Type *

RW

Allows all operations

Squash *

--Select what kind of user id squashing is performed --

Transport Protocol *

☒ UDP
 ☒ TCP

Clients

Any client can access

+ Add clients

Cancel

Create NFS export

8. Verify the NFS daemon is configured:

Example

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

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

Example

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

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

11. 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 under buckets in the Ceph Object Gateway.

Object Gateway » Buckets

+ Create ▾		
Name		Owner
☐ ▾ nfsexport		nfsuserexport

Additional Resources

- For more information on deploying Ceph object gateway, see the [. Deploying the Ceph object gateway using the command line interface](#) section in the *Red Hat Ceph Storage Operations guide*.
- 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.



NOTE

Red Hat Ceph Storage 5 supports only NFSv4 protocol.

Prerequisites

- 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

1. 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. In the *Create NFS export* window, set the following parameters and click *Create NFS export*:
 - a. Cluster - Name of the cluster.

- b. Daemons – You can select all daemons.
- c. Storage Backend – You can select CephFS.
- d. CephFS User ID – Select the service where the NFS cluster is created.
- e. CephFS Name – Provide a user name.
- f. CephFs Path – Any directory.
- g. NFS Protocol – NFSv4 is selected by default.
- h. Pseudo – root path
- i. Access Type – The supported access types are RO, RW, and NONE.
- j. Squash – Select the squash type.
- k. Transport Protocol – Select either the UDP or TCP protocol.
- l. Clients

Figure 8.1. CephFS NFS export window

NFS » Create

Create NFS export

Cluster *

test

✓ ↕

Daemons

test.mgmt-0

+ Remove all daemons

Storage Backend *

CephFS

✓ ↕

CephFS User ID *

nfs.test.mgmt-0

✓ ↕

CephFS Name *

cephfs

✓ ↕

Security Label

☐ Enable security label

CephFS Path *

/cephfs

✓

New directory will be created

NFS Protocol *

☐ NFSv3

☒ NFSv4

Pseudo * ⓘ

/nfsfsexport

✓

Access Type *

RW

✓ ↕

Allows all operations

Squash *

no_root_squash

✓ ↕

Transport Protocol *

☒ UDP

☒ TCP

Clients

Any client can access

+ Add clients

Cancel

Create NFS export

10. 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 **ganesha-manager** level of access on the Ceph dashboard.
- NFS Ganesha daemon configured on the dashboard.

Procedure

1. On the dashboard, click *NFS*.
2. Click the row that needs to be edited.
3. From the *Edit* drop-down menu, click *Edit*.
4. In the *Edit NFS export* window, edit the required parameters and click *Edit NFS export*.

NFS » Edit

Edit NFS export

Cluster * ?

nfs-rgw-service

Storage Backend *

CephFS

Volume *

cephfs

Security Label

☐ Enable security label

CephFS Path * ?

/

NFS Protocol *

☒ NFSv4

Pseudo * ?

/cephobject

Access Type *

RO

Allows only operations that do not modify the server

Squash ?

no_root_squash

Transport Protocol *

☒ UDP

☒ TCP

Clients

Any client can access

+ Add clients

Cancel

Edit NFS export

Verification

- You will get a notification that the NFS ganeshda is 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. On the dashboard, click *NFS*.
2. Click the row that needs to be delete.
3. From the *Edit* drop-down menu, click *Delete*.
4. In the *Delete NFS export* dialog box, check *Yes, I am sure* and click *Delete NFS export*.



Delete NFS export ×

Are you sure that you want to delete foo:1?

☒ Yes, I am sure.

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. Log in to the dashboard.
2. On the navigation menu, click *Pools*.
3. Click *Create*.
4. In the *Create Pool* window, set the following parameters:

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

- Set the name of the pool and select the pool type.
- Select either replicated or Erasure Coded (EC) pool type.
- Set the Placement Group (PG) number.
- 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.

5. Click *Create Pool*.
6. You get a notification 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. Log in to the dashboard.
2. On the navigation menu, click *Pools*.
3. To edit the pool, click its row.
4. Select *Edit* in the *Edit* drop-down.
5. In the *Edit Pool* window, edit the required parameters and click *Edit Pool*:

Figure 9.2. Editing pools

[Pools](#) » Edit

Edit 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 

0 B

Max objects 

0

RBD Configuration

Quality of Service 

Cancel

Edit Pool

6. You get a notification that the pool was created 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.

112

Procedure

1. Log in to the dashboard.
2. On the navigation bar, in *Cluster* drop-down menu, click *Configuration*.
3. In the *Level* drop-down menu, select **Advanced**:
4. Search for **mon_allow_pool_delete**, click Edit
5. Set all the values to **true**:

Figure 9.3. Configuration to delete pools

Cluster » Configuration » Edit

Edit mon_allow_pool_delete

Name

mon_allow_pool_delete

Description

allow pool deletions

Default

false

✓

Services

mon

Values


global	true	✓	⬆
mon	true	✓	⬆
mgr	true	✓	⬆
osd	true	✓	⬆
mds	true	✓	⬆
client	true	✓	⬆

Cancel

Update

6. On the navigation bar, click *Pools*:
7. To delete the pool, click on its row:
8. From *Edit* drop-down menu, select *Delete*.
9. In the *Delete Pool* window, Click the *Yes, I am sure* box and then Click *Delete Pool* to save the settings:

Figure 9.4. Delete pools



A confirmation dialog box titled "Delete Pool" with a close button (X) in the top right corner. The dialog asks, "Are you sure that you want to delete **test_pool**?" Below the question is a checked checkbox with the text "Yes, I am sure." At the bottom right are two buttons: "Cancel" and "Delete Pool".

Delete Pool X

Are you sure that you want to delete **test_pool**?

☒ Yes, I am sure.

Cancel Delete Pool

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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Hosts*.
3. Select a host from the list.
4. From the *Edit* drop-down menu, click *Enter Maintenance*.

Figure 10.1. Entering maintenance mode



NOTE

When a host enters maintenance, all daemons are stopped. You can check the status of the daemons under the *Daemons* tab of a host.

Verification

- 1. You get a notification that the host is successfully moved to maintenance and a *maintenance* label appears in the **Status** column.



NOTE

If the maintenance mode fails, you get a notification 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. Log in to the Dashboard.
- 2. From the *Cluster* drop-down menu, select *Hosts*.
- 3. From the *Hosts* List, select the host in maintenance.




NOTE

You can identify the host in maintenance by checking for the *maintenance* label in the **Status** column.

- 4. From the *Edit* drop-down menu, click *Exit Maintenance*.

Figure 10.2. Exiting maintenance mode



After exiting the maintenance mode, you need to create the required services on the host by default-crash and the node-exporter gets deployed.

Verification

1. You get a notification that the host has been successfully moved out of maintenance and the maintenance label is removed from the Status column.

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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Hosts*.
3. From the *Hosts List*, select the host you want to remove.
4. From the *Edit* drop-down menu, click *Start Drain*.

Figure 10.3. 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* option from the *Edit* drop-down menu.
5. Check if all the daemons are removed from the host.
 - a. Click the *Expand/Collapse* icon on it's row
 - b. Select *Daemons*. No daemons should be listed.

Figure 10.4. Checking the status of host daemons




IMPORTANT

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


6. Remove the host.

- a. From the *Edit* drop-down menu, click *Remove*.

Figure 10.5. Removing the host



- b. In the *Remove Host* dialog box, check **Yes, I am sure.** and click **Remove Host**.



Verification

1. You get a notification after the successful removal of the host from the *Hosts* List.

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

1. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *OSDs*.

Creating an OSD

1. To create the OSD, click *Create*.

Figure 11.1. Add device for OSDs


Primary devices x

⚠ 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. You can check for available devices in *Physical Disks* under the *Cluster* drop-down menu.

- In the *Create OSDs* window, from *Deployment Options*, select one of the below 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 NVMEs are used to deploy OSDs.
 - From the *Advanced Mode*, you can 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, NVMEs or SSDs.
 - **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, NVMEs or SSDs).
 - If you want to encrypt your data for security purposes, under *Features*, select *encryption*.
 - Click the *Preview* button and in the OSD Creation Preview dialog box, Click *Create*.
 - In the *OSD Creation Preview* dialog box, Click *Create*.
- You get a notification that the OSD was created successfully.
 - The OSD status changes from *in and down* to *in and up*.

Editing an OSD

1. To edit an OSD, select the row.
 - a. From *Edit* drop-down menu, select *Edit*.
 - b. Edit the device class.
 - c. Click *Edit OSD*.

Figure 11.2. Edit an OSD



Dialog box titled "Edit OSD: 9" with a close button (X) in the top right corner.

Below the title bar, there is a horizontal line.

Underneath the line, the label "Device class" is displayed, followed by a red asterisk (*).

To the right of the label is a text input field containing the value "hdd".

Below the input field is another horizontal line.


At the bottom right of the dialog, there are two buttons: "Cancel" and "Edit OSD".

- d. You get a notification that the OSD was updated successfully.

Marking the Flags of OSDs

1. To mark the flag of the OSD, select the row.
 - a. From *Edit* drop-down menu, select *Flags*.
 - b. Mark the Flags with *No Up*, *No Down*, *No In*, or *No Out*.
 - c. Click *Update*.

Figure 11.3. Marking Flags of an OSD



Individual OSD Flags X

☐ **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

- d. You get a notification that the flags of the OSD was updated successfully.

Scrubbing the OSDs

1. To scrub the OSD, select the row.
 - a. From *Edit* drop-down menu, select *Scrub*.
 - b. In the *OSDs Scrub* dialog box, click *Update*.

Figure 11.4. Scrubbing an OSD



OSDs Scrub X

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


Cancel Update

- c. You get a notification that the scrubbing of the OSD was initiated successfully.

Deep-scrubbing the OSDs

1. To deep-scrub the OSD, select the row.
 - a. From *Edit* drop-down menu, select *Deep scrub*.
 - b. In the *OSDs Deep Scrub* dialog box, click *Update*.

Figure 11.5. Deep-scrubbing an OSD



- c. You get a notification that the deep scrubbing of the OSD was initiated successfully.

Reweighting the OSDs

1. To reweight the OSD, select the row.
 - a. From *Edit* drop-down menu, select *Reweight*.
 - b. In the *Reweight OSD* dialog box, enter a value between zero and one.
 - c. Click *Reweight*.


Figure 11.6. Reweighting an OSD



Marking OSDs Out

1. To mark the OSD out, select the row.
 - a. From *Edit* drop-down menu, select *Mark Out*.
 - b. In the *Mark OSD out* dialog box, click *Mark Out*.

Figure 11.7. Marking OSDs out




- c. The status of the OSD will change to *out*.

Marking OSDs In

1. To mark the OSD in, select the OSD row that is in *out* status.
 - a. From *Edit* drop-down menu, select *Mark In*.
 - b. In the *Mark OSD in* dialog box, click *Mark In*.

Figure 11.8. Marking OSDs in




- c. The status of the OSD will change to *in*.

Marking OSDs Down

1. To mark the OSD down, select the row.
 - a. From *Edit* drop-down menu, select *Mark Down*.
 - b. In the *Mark OSD down* dialog box, 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

- c. The status of the OSD will change to *down*.

Marking OSDs Lost

1. To mark the OSD lost, select the OSD in *out* and *down* status.
 - a. From *Edit* drop-down menu, select *Mark Lost*.
 - b. In the *Mark OSD Lost* dialog box, check *Yes, I am sure* option, 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 the OSD, select the OSD in *down* status.
 - a. From *Edit* drop-down menu, select *Purge*.
 - b. In the *Purge OSDs* dialog box, check *Yes, I am sure* option, and click *Purge OSD*.

Figure 11.11. Purging OSDs



- c. All the flags are reset and the OSD is back in *in* and *up* status.

Destroying OSDs

1. To destroy the OSD, select the OSD in *down* status.
 - a. From *Edit* drop-down menu, select *Destroy*.
 - b. In the *Destroy OSDs* dialog box, check *Yes, I am sure* option, 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

- c. The status of the OSD changes to *destroyed*.

Deleting OSDs

1. To delete the OSD, select the OSD in *down* status.
 - a. From *Edit* drop-down menu, select *Delete*.
 - b. In the *Destroy OSDs* dialog box, check *Yes, I am sure* option, 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 page showing failed OSDs.



In this example, you can see that one of the OSDs is down on the landing page of the dashboard.

Apart from this, on the physical drive, you can view the LED lights blinking if one of the OSDs is down.

2. Click *OSDs*.
3. Select the **out** and **down** OSD:
 - a. From the *Edit* drop-down menu, select *Flags* and select *No Up* and click *Update*.
 - b. From the *Edit* drop-down menu, select *Delete*.
 - c. In the *Delete OSD* dialog box, select the *Preserve OSD ID(s) for replacement* and *Yes, I am sure* check boxes.
 - d. Click *Delete OSD*.
 - e. Wait till the status of the OSD changes to **out** and **destroyed** status.
4. Optional: If you want to change the *No Up* Flag for the entire cluster, in the *Cluster-wide configuration* drop-down menu, select *Flags*.
 - a. In *Cluster-wide OSDs Flags* dialog box, select *No Up* and click *Update*.
5. 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
```

6. From the *Create* drop-down menu, select *Create*.
7. In the *Create OSDs* window, click *+Add* for Primary devices.
 - a. In the *Primary devices* dialog box, from the *Hostname* drop-down list, select any one filter. From *Any* drop-down list, select the respective option.



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 from *Any* list select **hdd**. Select *Vendor* and from *Any* list, select **ATA**

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

- b. Click *Add*.
 - c. In the *Create OSDs window*, click the *Preview* button.
 - d. In the *OSD Creation Preview* dialog box, Click *Create*.
 - e. You will get a notification that the OSD is created. The OSD will be in **out** and **down** status.
8. Select the newly created OSD that has *out* and *down* status.
 - a. In the *Edit* drop-down menu, select *Mark-in*.
 - b. In the *Mark OSD in* window, select *Mark in*.
 - c. In the *Edit* drop-down menu, select *Flags*.
 - d. Uncheck *No Up* and click *Update*.
 9. 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* dialog box, uncheck *No Up* 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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Services*.
3. Click *+Create*.
4. In the *Create Service* window, select **rgw** service.
5. Select **SSL** and upload the *Certificate* in **.pem** format.

Figure 12.1. Creating Ceph Object Gateway service

Cluster » Services » Create

Create Service

Type *

rgw ✓

Id *

rgw.ssl

☐ Unmanaged

Placement

Hosts ✓

Hosts

[ceph-ssl-rhcs5-8dcxv6-node2](#) ✕

Count ?

Port

443 ✓

☒ SSL

Certificate ?

```
-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQEA+Cf4190agD6x67HhdCy4Asqw89Zz9ZuGbH50/71tIMQpJJU0
gu90bNtIoC0zabJ7n1jujueYgIp0qGnhRSvsGJ1EkgN81NLQ9rqAVaGpadjrNLcM
bpgqJCZj0vzzmtFBCtenpb51/EccMFcAydGtGeLP33SaWiZ4Rne56GBInk6SATI/
JSKweGD1y5GiAwipBR4C74HiAW9q6hCOuSdp/2WQxwT3T1j2sj1qxkHdtInUtw0m
-----
```

Choose File rgw.crt

Cancel

Create Service

6. Click *Create Service*.
7. 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. Log in to the Dashboard.
2. From the *Cluster* drop-down menu, select *Services*.
3. Click *Create*.
4. In the *Create Service* window, select **ingress** service.
5. Select backend service and edit the required parameters.

Figure 12.2. Creating **ingress** service

Create Service

X

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

6. Click *Create Service*.

7. You get a notification 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. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Users* and then Click *Create*.
4. In the *Create User* window, set the following parameters:
 - a. Set the user name, full name, and edit the maximum number of buckets if required.
 - b. Optional: Set an email address or suspended status.
 - c. Optional: Set a custom access key and secret key by unchecking *Auto-generate key*.
 - d. Optional: Set a user quota.
 - e. Check *Enabled* under *User quota*.
 - f. Uncheck *Unlimited size* or *Unlimited objects*.
 - g. Enter the required values for *Max. size* or *Max. objects*.
 - h. Optional: Set a bucket quota.
 - i. Check *Enabled* under *Bucket quota*.
 - j. Uncheck *Unlimited size* or *Unlimited objects*:

- k. Enter the required values for *Max. size* or *Max. objects*:
5. Click *Create User*.

Figure 12.3. Create Ceph object gateway user

The screenshot shows the 'Create User' form with the following fields and values:

- User ID ***: rgw-test-11 (with a green checkmark)
- Show Tenant**: ☐
- Full name ***: rgw-test (with a green checkmark)
- Email address**: rgw-test@mail.com (with a green checkmark)
- Max. buckets**: Custom (dropdown menu)
- 1000**: (input field)
- Suspended**: ☐
- S3 key**:
 - Auto-generate key**: ☒
- User quota**:
 - Enabled**: ☐
- Bucket quota**:
 - Enabled**: ☐

At the bottom right, there are two buttons: 'Cancel' and 'Create User' (highlighted in blue).

6. You get a notification 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. 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. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Users*.
4. Select the user by clicking its row.
5. From *Edit* drop-down menu, select *Edit*.
6. In the *Edit User* window, click *+Create Subuser*.
7. In the *Create Subuser* dialog box, enter the user name and select the appropriate permissions.
8. Check the *Auto-generate secret* box and then click *Create Subuser*.

Figure 12.4. 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 clicking *Auto-generate-secret* checkbox, the secret key for object gateway is generated automatically.

9. In the *Edit User* window, click the *Edit user* button
10. You get a notification that the user was updated successfully.

12.4.3. 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. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Users*.
4. To edit the user capabilities, click its row.
5. From the *Edit* drop-down menu, select *Edit*.
6. In the *Edit User* window, edit the required parameters.
7. Click *Edit User*.

Figure 12.5. 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
----	--

8. You get a notification 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.

12.4.4. 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. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Users*.
4. To delete the user, click its row.
5. From the *Edit* drop-down menu, select *Delete*.
6. In the *Edit User* window, edit the required parameters.
7. In the *Delete user* dialog window, Click the *Yes, I am sure* box and then Click *Delete User* to save the settings:

Figure 12.6. 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. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Buckets* and then click *Create*.
4. In the *Create Bucket* window, enter a value for *Name* and select a user that is not suspended. Select a placement target.

Figure 12.7. Create Ceph object gateway bucket

Create Bucket

Name *

Owner *

Placement target *

Locking

☐ Enabled ?

Security

☐ Encryption ?



NOTE

A bucket's placement target is selected on creation and can not be modified.

5. Optional: Enable *Locking* for the objects in the bucket. Locking can only be enabled while creating a bucket. Once locking is enabled, you also have to choose the lock mode, *Compliance* or *Governance* and the lock retention period in either days or years, not both.
6. Optional: Enable *Security* to encrypt the objects in the bucket. To enable encryption on a bucket, you need to set the configuration values for SSE-S3.
 - a. To set the configuration values, hover the cursor over the question mark and click *Click here*.

- b. In the *Update RGW Encryption Configurations* window, select **SSE-S3** as the *Encryption Type*, provide the required details, and click *Submit*.

Figure 12.8. Encrypt objects in the bucket

Update RGW Encryption Configurations ×

Encryption Type: ☒ SSE-S3 Encryption ☐ SSE-KMS Encryption

Key management service provider *

Authentication Method * ✓

Secret Engine * ✓

Secret Path

Namespace

Vault Address *

Role *

CA Certificate ?

Client Certificate ?

Client Private Key ?

**NOTE**

When using **SSE-S3** encryption type, Ceph manages the encryption keys that are stored in the vault by the user.

7. Click *Create bucket*.
8. You get a notification 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 12.9. Edit Ceph object gateway bucketSelected 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. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*.
3. Click *Buckets*.
4. To delete the bucket, click it's row.
5. From the *Edit* drop-down select *Delete*.
6. In the *Delete Bucket* dialog box, Click the *Yes, I am sure* box and then Click *Delete bucket* to save the settings:

Figure 12.10. 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

- 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. On the Dashboard landing page of the secondary site, in the vertical menu bar, click *Object Gateway* drop-down list.
2. Select *Buckets*.
3. You can see those object gateway buckets on the secondary landing page that were created for the object gateway users on the primary site.

Figure 12.11. Multisite 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. MANAGING 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.

12.7.1. 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 multiste 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. On the Dashboard landing page of the secondary site, in the vertical menu bar, click *Object Gateway* drop-down list.
2. Select *Buckets*.
3. You can see those object gateway buckets on the secondary landing page that were created for the object gateway users on the primary site.

Figure 12.12. Monitoring 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

4. Click the row of the bucket that you want to edit.
5. From the *Edit* drop-down menu, select *Edit*.
6. In the *Edit Bucket* window, edit the required parameters and click *Edit Bucket*.

Figure 12.13. 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

Verification

- You will get a notification that the bucket is updated successfully.

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 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*.

12.7.2. 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 multiste 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. On the Dashboard landing page of the primary site, in the vertical menu bar, click *Object Gateway* drop-down list.
2. Select *Buckets*.
3. You can see those object gateway buckets of the secondary site here.
4. Click the row of the bucket that you want to delete.
5. From the *Edit* drop-down menu, select *Delete*.
6. In the *Delete Bucket* dialog box, select *Yes, I am sure* checkbox, and click *Delete Bucket*.

Verification

- The selected row of the bucket is deleted successfully.

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 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*.

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. Log in to the dashboard.
2. On the navigation menu, click the *Block* drop-down menu.
3. Select *Images*.
4. Click *Create*.
5. In the *Create RBD* window, enter the parameters.
6. Optional: Click *Advanced* and set the parameters.
7. Click *Create RBD*.
8. Create Block device image.

Figure 14.1. Create Block device image

The screenshot shows the 'Create RBD' form in the Red Hat Ceph Storage dashboard. The left sidebar contains navigation links: Dashboard, Cluster, Pools, Block (selected), Images, Mirroring, NFS, File Systems, and Object Gateway. The main content area shows the 'Create RBD' form with the following fields and options:

- Name ***: A text input field with a placeholder 'Name...'.
- Pool ***: A dropdown menu with 'testing' selected.
- Use a dedicated data pool**: An unchecked checkbox with a help icon.
- Size ***: A text input field with a placeholder 'e.g., 10GiB'.
- Features**: A list of checkboxes, all of which are checked:
 - Deep flatten
 - Layering
 - Exclusive lock
 - Object map (requires exclusive-lock)
 - Fast diff (interlocked with object-map)
 - Mirroring (unchecked with a help icon)

At the bottom right of the form, there is a 'Cancel' button and a 'Create RBD' button. An 'Advanced...' link is also visible near the bottom right.

9. You get a notification 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. Log in to the Dashboard.
2. On the navigation menu, click the *Block* drop-down menu.

3. Select *Images*.
4. To create the namespace, in the *Namespaces* tab, click *Create*.
5. In the *Create Namespace* window, select the pool and enter a name for the namespace.
6. Click *Create*.

Figure 14.2. Create namespace

The screenshot shows a 'Create Namespace' window. It has a title bar with a close button (X). Below the title bar, there are two input fields. The first is labeled 'Pool' with a red asterisk, and it contains a dropdown menu with 'testbench' selected. The second is labeled 'Name' with a red asterisk, and it contains a text input field with the placeholder text 'Namespace name...'. At the bottom right of the window, there are two buttons: 'Cancel' and 'Create'.

7. You get a notification that the namespace was created successfully.

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. Log in to the Dashboard.
2. On the navigation bar, click the *Block* drop-down menu.

3. Select *Images*.
4. To edit the image, click its row.
5. In the *Edit* drop-down menu, select *Edit*.
6. In the *Edit RBD* window, edit the required parameters and click *Edit RBD*.

Figure 14.3. Edit Block device image

Block » Images » Edit

Edit RBD

Name *

test_image ✓

Pool

pool_test_1

☐ Use a dedicated data pool [?](#)

Size *

10 GiB

Features

☒ Deep flatten
☒ Layering
☒ Exclusive lock
☒ Object map (requires exclusive-lock)
☐ Journaling (requires exclusive-lock)
☒ Fast diff (interlocked with object-map)

[Advanced...](#)

Cancel

Edit RBD

7. You get a notification that the image was updated 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.

14.1.4. Copying images on the Ceph dashboard

You can copy 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. Log in to the Dashboard.
2. On the navigation bar, click the *Block* drop-down menu.
3. Select *Images*.
4. To copy the image, click its row.
5. In the *Edit* drop-down menu, select *Copy*.
6. In the *Copy RBD* window, set the required parameters and click *Copy RBD*.

Figure 14.4. Copy Block device image

Block » Images » Copy

Copy RBD

Copy from

rbd_primary_pool/image_test

Name *

Name...

Pool *

rbd_primary_pool

☐ Use a dedicated data pool

Size *

1 GiB

Features

☒ Deep flatten
☒ Layering
☒ Exclusive lock
☒ Object map (requires exclusive-lock)
☐ Journaling (requires exclusive-lock)
☒ Fast diff (interlocked with object-map)

[Advanced...](#)

Cancel

Copy RBD

7. You get a notification that the image was copied 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.

14.1.5. 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images* from the drop-down menu.
4. To move the image to trash, click its row.
5. Select *Move to Trash* in the *Edit* drop-down.
6. In the *Moving an image to trash* window, edit the date till which the image needs protection, and then click *Move*.

Figure 14.5. Moving images to trash

Move an image to trash

×

This image contains snapshot(s), which will prevent it from being removed after moved to trash.

To move `rbd_primary_pool/image_test` to trash, click **Move**. Optionally, you can pick an expiration date.

Protection expires at

NOT PROTECTED

Cancel

Move

7. You get a notification that the image was moved to trash successfully.

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. Log in to the Dashboard.
2. On the navigation bar, click *Block*:
3. Select *Images*.
4. In the *Trash* tab, click *Purge Trash*.
5. In the *Purge Trash* window, select the pool, and then click *Purge Trash*.

Figure 14.6. Purge trash



Purge Trash

To purge, select **All** or one pool and click **Purge**.

Pool:

All

Cancel Purge

6. You get a notification that the pools in the trash were purged successfully.

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. Log in to the Dashboard.
2. On the navigation menu, click *Block*:
3. Select *Images*.
4. To restore the image from Trash, in the *Trash* tab, click its row:
5. Select *Restore* in the *Restore* drop-down.
6. In the *Restore Image* window, enter the new name of the image , and then click *Restore*.

Figure 14.7. Restore images from trash



Restore Image

To restore `rbd_primary_pool/image_test@450cb496723a`, type the image's new name and click **Restore**.

New Name

image_test

Cancel Restore

7. You get a notification 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


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. Log in to the Dashboard.
2. On the navigation bar, click *Block*
3. Select *Images*.
4. To delete the image, select the row.
5. From the *Edit* drop-down, select *Delete*.
6. In the *Delete RBD* dialog box, click the *Yes, I am sure* box and then Click *Delete RBD* to save the settings.

Figure 14.8. Deleting images



7. You get a notification 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. Log in to the Dashboard.
2. On the navigation bar, click *Block*
3. Select *Images*.
4. To delete the namespace, in the *Namespaces* tab, click its row.
5. Click *Delete*.
6. In the *Delete Namespace* dialog box, click the *Yes, I am sure* box and then Click *Delete Namespace* to save the settings:

Figure 14.9. Deleting namespaces



7. You get a notification 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To take the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Create* in the *Create* drop-down.
6. In the *Create RBD Snapshot* dialog, enter the name and click *Create RBD Snapshot*:

Figure 14.10. Creating snapshot of images



Create RBD Snapshot ×

Name * ✓

Cancel Create RBD Snapshot

7. You get a notification 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To rename the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Rename* in the the *Rename* drop-down.
6. In the *Rename RBD Snapshot* dialog box, enter the name and click *Rename RBD Snapshot*:

Figure 14.11. Renaming snapshot of images



Rename RBD Snapshot

Name *

Cancel 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To protect the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Protect* in the the *Rename* drop-down.
6. The *State* of the snapshot 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.

4. To protect the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Clone* in the the *Rename* drop-down.
6. In the *Clone RBD* window, edit the parameters and click *Clone RBD*.

Figure 14.12. Cloning snapshot of images

Block » Images » Clone

The screenshot shows the 'Clone RBD' dialog box. It has a title bar 'Clone RBD'. Below it, there are several fields and checkboxes:

- Clone from:** A text field containing 'rbd_primary_pool/image_test@image_snap'.
- Name:** A text field with a red asterisk, containing 'Name...'.
- Pool:** A dropdown menu with a red asterisk, showing 'rbd_primary_pool'.
- Use a dedicated data pool:** An unchecked checkbox.
- Size:** A text field with a red asterisk, containing '1 GiB'.
- Features:** A list of checkboxes:
 - ☒ Deep flatten
 - ☒ Layering
 - ☒ Exclusive lock
 - ☒ Object map (requires exclusive-lock)
 - ☐ Journaling (requires exclusive-lock)
 - ☒ Fast diff (interlocked with object-map)

At the bottom right, there are two buttons: 'Cancel' and 'Clone RBD'. There is also a link 'Advanced...' in the bottom right corner.

7. You get a notification that the snapshot was cloned successfully. You can search for the cloned image in the *Images* tab.

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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To protect the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Copy* in the the *Rename* drop-down menu.
6. In the *Copy RBD* window, enter the parameters and click the *Copy RBD* button:

Figure 14.13. Copying snapshot of images

The screenshot shows the 'Copy RBD' form in the Ceph Dashboard. The left sidebar has a dark theme with the following menu items: Dashboard, Cluster, Pools, Block (selected), Images, Mirroring (1), NFS, File Systems, and Object Gateway. The main content area is titled 'Copy RBD' and has a breadcrumb 'Block > Images > Copy'. The form fields are as follows:

- Copy from:** test_pool/test_image
- Name:** pool_test1 (with a green checkmark)
- Pool:** test_pool (with a green checkmark)
- Namespace:** -- Select a namespace -- (dropdown menu)
- Use a dedicated data pool:** ☐ (with a help icon)
- Size:** 10 GiB
- Features:**
 - ☒ Deep flatten
 - ☒ Layering
 - ☒ Exclusive lock
 - ☒ Object map (requires exclusive-lock)
 - ☒ Fast diff (interlocked with object-map)
 - ☐ Mirroring (with a help icon)

At the bottom right, there are 'Advanced...' and 'Copy RBD' buttons, and a 'Cancel' button.

7. You get a notification that the snapshot was copied successfully. You can search for the copied image in the *Images* tab.

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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To unprotect the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *UnProtect* in the the *Rename* drop-down.
6. The *State* of the snapshot 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To rollback the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Rollback* in the the *Rename* drop-down.
6. In the *RBD snapshot rollback* dialog box, click *Rollback*.

Figure 14.14. 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. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Select *Images*.
4. To take the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab.
5. Select *Delete* in the the *Rename* drop-down:

Figure 14.15. Deleting snapshot of images



6. You get a notification that the snapshot was deleted 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

1. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Click *Mirroring*.

Figure 14.16. View mirroring of Block devices

The screenshot shows the Red Hat Ceph Storage Dashboard in the 'Block > Mirroring' view. At the top, the 'Site Name' is 'f64f341c-655d-11eb-8778-fa163e914bcc'. Below this, there are three main sections: 'Daemons', 'Pools', and 'Images'.

Daemons: A table with columns: Instance, ID, Hostname, Version, and Health. It lists two daemons: '14677' (ceph-ek6-0-pr23qs-node1-installer.rndsey) and '14680' (ceph-ek6-0-pr23qs-node2.jgqxd0). Both are in 'OK' health. A '2 total' summary is shown at the bottom.

Pools: A table with columns: Name and Mode. It lists one pool: 'test_pool' with mode 'pool'. An 'Edit Mode' button is present. A '1 selected / 1 total' summary is shown at the bottom.

Images: A table with columns: Pool, Image, State, and Issue. It lists one image: 'test_pool' with image 'test_image' and state 'Unknown'. A '1 total' summary is shown at the bottom.

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

1. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Click *Mirroring*.
4. In the *Pools* tab, click the peer you want to edit.
5. In the *Edit Mode* drop-down, select *Edit Mode*.
6. In the *Edit pool mirror mode* window, select the mode from the drop-down, and then click *Update*. Pool is updated successfully.

Figure 14.17. Editing mode in mirroring



Edit pool mirror mode

To edit the mirror mode for pool `g_mirror_pool`, select a new mode from the list and click **Update**

Mode

Image

Cancel Update

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. Log in to the dashboard.
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 14.18. 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 *

7d8f3cac-1ee2-11ee-92eb-fa163eb0ad2d ✓

Pools *

☒ testing

At least one pool is required.

Token Generate

eyJmc2lkjoiN2Q4ZjNjYWMtMWVlMiOxMWVlLTkyZWltZmExNjNlYjBhZD
JkliwiY2xpZW50X2lkjoicmJkLWlpcnJvcilwZWVylwia2V5ljoIQVFESytLe ✓



Close

- a. For the provided site name, choose the pool to be mirrored.
 - b. For the selected pool, generate a new bootstrap token by clicking *Generate*.
 - c. Click the *Copy* icon to copy the token to clipboard.
 - d. Click *Close*.
4. Enable pool mirror mode.
 - a. Select the pool.
 - b. Click *Edit Mode*.
 - c. From the *Edit pool mirror mode* window, select *Image* from the drop-down.

- d. Click *Update*.

Site B

1. Log in to the dashboard.
2. From the *Navigation* menu, click the *Block* drop-down menu, and click *Mirroring*.
3. From the *Create Bootstrap token* drop-down, select *Import Bootstrap Token*.



NOTE

Ensure that mirroring mode is enabled for the specific pool for which you are importing the bootstrap token.

4. In the *Import Bootstrap Token* window, choose the direction, and paste the token copied earlier from site A.

Figure 14.19. 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 *

7d8f3cac-1ee2-11ee-92eb-fa163eb0ad2d ✓

Direction

Bidirectional

Pools *

☒ testing

Token *

JA0NJC40S8WXSxbDj10MTAdiMC4yMDCuMijE0MIZiMWMIC8WLFYxOjEwLjAuMjA3LjlxOjY3ODkvMF0sW3YyOjEwLjAuMjA1LjMyOjMzMdAvMCx2MToxMC4wLjlwNS4zMjo2Nzg5LzBdlnO= ✓

Cancel Submit

5. Click *Submit*.
The peer is added and the images are mirrored in the cluster at site B.
6. Verify the health of the pool is in OK state.
 - In the *Navigation* menu, under *Block*, select *Mirroring*. The health of the pool is *OK*.

Site A

1. Create an image with *Mirroring* enabled.
 - a. From the *Navigation* menu, click the *Block* drop-down menu.
 - b. Click *Images*.

- c. Click *Create*.
- d. In the *Create RBD* window, provide the *Name*, *Size* and enable *Mirroring*.



NOTE

You can either choose *Journal* or *Snapshot*.

- e. Click *Create RBD*.

Figure 14.20. Create mirroring image

Red Hat Ceph Storage

Block > Images > Create

Create RBD

Name *

Pool *

☐ Use a dedicated data pool

Size *

Features

- ☒ Deep flatten
- ☒ Layering
- ☒ Exclusive lock
- ☒ Object map (requires exclusive-lock)
- ☒ Fast diff (interlocked with object-map)
- ☒ Mirroring
 - ☐ Journal
 - ☒ Snapshot

Schedule Interval

Advanced...

Cancel Create RBD

2. Verify the image is available at both the sites.

- In the *Navigation* menu, under *Block*, select *Images*. The image in site A is **primary** while the image in site B is **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 *rbd* application enabled is created.
- An image is created.
- Mirroring is configured.
- A peer is added.

Procedure

1. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Click *Mirroring*.
4. In the *Pools* tab, click the peer you want to edit.
5. In the *Edit Mode* drop-down, select *Edit peer*.
6. In the *Edit pool mirror peer* window, edit the parameters, and then click *Submit*:

Figure 14.21. Editing peer in mirroring

Edit pool mirror peer [X]

Edit the pool mirror peer attributes for pool **testing** and click **Submit**.

Cluster Name *

7d8f3cac-1ee2-11ee-92eb-fa163eb0ad2d ✓

CephX ID *

rbd-mirror-peer

Monitor Addresses

[v2:10.0.206.202:3300/0,v1:10.0.206.202:6789/0],[v2:10.0.206.203:3300/0...

CephX Key

AQDK+Kxkl0fTKxAAykGul00tal2crje8mufv6g==

Cancel Submit

7. You get a notification that the peer was updated 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.

14.2.5. Deleting 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 *rbd* application enabled is created.
- An image is created.
- Mirroring is configured.
- A peer is added.

Procedure

1. Log in to the Dashboard.
2. On the navigation menu, click *Block*.
3. Click *Mirroring*.
4. In the *Pools* tab, click the peer you want to delete.
5. In the *Edit Mode* drop-down, select *Delete peer*.
6. In the *Delete mirror peer* dialog window, Click the *Yes, I am sure* box and then Click *Delete mirror peer* to save the settings:

Figure 14.22. Delete peer in mirroring



7. You get a notification that the peer was 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**.