
Red Hat Ceph Storage Dashboard

As a storage administrator, the Red Hat Ceph Storage Dashboard provides management and monitoring capabilities. Use the dashboard to administer and configure the cluster, and visualize information and performance statistics that are related to it. The dashboard uses a web server that is hosted by the `ceph-mgr` daemon.

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

Note: To use the Red Hat Ceph Storage dashboard, have system administrator level experience.

The 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 by using both the command line and the web user interface. The dashboard supports various methods to enhance password security. Password complexity rules can be configured, requiring users to change their password after the first login or after a configurable time period.

For more information, see [“Managing roles” on page 29](#) and [“Managing users” on page 34](#).

Single Sign-On (SSO)

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

For more information, see [“Enabling single sign-on with SAML 2.0” on page 24](#).

Auditing

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

For more information about using the manager modules with the dashboard, see [“Viewing and editing the manager modules of the Ceph cluster” on page 0](#).

High availability (Technology Preview)

Both Ceph management and Ceph monitoring uses high availability. High availability provides the users with transparent manager failover of the `mgmt-gateway`. High availability redirects to the active manager and to the active instance of Prometheus, Alertmanager or Grafana, when several instances are available.

Important: To use high availability with the Ceph Dashboard, ensure that there are no port conflicts that might disrupt service availability. High-availability configurations and clustering for the `mgmt-gateway` service itself are currently not supported. Services must bind to the appropriate ports based on the applications being proxied.

Security features

The dashboard provides the following security features.

SSL and TLS support

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

For more information, see [“Bootstrap storage cluster along with Dashboard” on page 0](#).

Username and password protection

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

For more information, see [“Managing users” on page 34](#).

Configuring and managing the dashboard interface

Configure and manage the Ceph Dashboard.

Configuring manager modules

You can view and change parameters for Ceph manager modules.

For more information, see [“Viewing and editing the manager modules of the Ceph cluster” on page 0](#).

Embedded Grafana dashboards

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

For more information, see [“Components” on page 4](#).

Viewing and filtering logs

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

For more information, see [“Filtering logs of the Ceph cluster” on page 0](#).

Toggling dashboard components

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

For more information, see [“Toggling Ceph dashboard features” on page 11](#).

Quality of service for images

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

For more information, see [“Managing block device images” on page 105](#).

Ceph management with the dashboard

Manage and monitor Red Hat Ceph Storage with the Ceph Dashboard.

Upgrading

You can upgrade the Ceph cluster version by using the dashboard.

For more information, see [“Upgrading a cluster” on page 18](#).

Overall cluster health

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

For more information, see [“Viewing and editing the configuration of the Ceph cluster” on page 0](#).

Configuration editor

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

For more information, see [“Viewing and editing the configuration of the Ceph cluster” on page 0](#).

Managing OSD settings

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

For more information, see [“Managing Ceph OSDs from the Ceph Dashboard” on page 47](#).

Pools

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

For more information, see [“Understanding the landing page of the Ceph dashboard” on page 6](#) and [“Monitoring pools of the Ceph cluster” on page 0](#).

OSDs

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

For more information, see [“Monitoring Ceph OSDs” on page 0](#).

Images

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

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

For more information, see [“Monitoring block device images” on page 0](#).

Block device mirroring

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

For more information, see [“Mirroring view” on page 113](#).

Ceph File Systems

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

For more information, see [“Monitoring Ceph File Systems” on page 0](#).

Object Gateway (RGW)

Lists all active object gateways and their performance counters. This window displays and manages, including add, edit, and delete, Ceph Object Gateway users and their details. This includes quotas and users’ buckets and their details, for example, owner or quotas.

For more information, see [“Monitoring Ceph Object Gateway daemons” on page 0](#).

Monitoring with the dashboard

Monitor the following components with the dashboard.

Viewing alerts

Use the alerts page to see details of current alerts.

For more information, see [“Viewing alerts” on page 63](#).

Viewing cluster hierarchy

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

For more information, see [“Viewing the CRUSH map of the Ceph cluster” on page 0](#).

Cluster logs

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

For more information, see [“Filtering logs of the Ceph cluster” on page 0](#).

Centralized logs

Provides a centralized location for viewing logs in the Red Hat Ceph Storage cluster for more efficient monitoring.

For more information, see [“Viewing centralized logs of the Ceph cluster” on page 0](#).

Hosts

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

For more information, see [“Monitoring hosts of the Ceph cluster” on page 0](#).

Performance counters

Displays detailed statistics for each running service.

For more information, see [“Monitoring services of the Ceph cluster” on page 0](#).

Monitors

Lists all Monitors, their quorum status, and open sessions.

For more information, see [“Monitoring monitors of the Ceph cluster” on page 0](#).

Device management

Lists all hosts that are known by the Orchestrator. Lists all drives that are attached to a host and their properties. Displays drive health predictions, SMART data, and blink enclosure LEDs.
For more information, see [“Monitoring hosts of the Ceph cluster” on page 0](#).

View storage cluster capacity

You can view the raw storage capacity of the Red Hat Ceph Storage cluster in the Capacity pages of the Ceph dashboard.
For more information, see [“Understanding the landing page of the Ceph dashboard” on page 6](#).

Components

Understand the different components of the dashboard.

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.

Related information

[Prometheus website](#)

[Grafana website](#)

Architecture

The Dashboard architecture depends on the Ceph manager dashboard plugin and other components.

[“Figure: Ceph Dashboard architecture” on page 4](#) illustrates how the Ceph manager and dashboard work together.

Figure: Ceph Dashboard architecture

Accessing the Ceph dashboard after bootstrap

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

PREREQUISITE

Be sure that the following is in place before accessing the Ceph dashboard:

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

1. Log into the dashboard through a web browser.
This can be done in one of the following ways:

- Enter the following URL:

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

Replace:

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

For example,

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

- From the Cephadm shell, run the **ceph mgr services** command.
For example,


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

This command displays all endpoints that are currently configured. Look for the dashboard key to obtain the URL for accessing the dashboard.

2. From the login page, enter the username admin and the default password provided during bootstrapping.

Note: During initial login to the Red Hat Ceph Storage dashboard, you must change the password. For more information, see [“Changing the dashboard password” on page 13](#).

3. 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: Red Hat Ceph Storage dashboard landing page



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

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.

Menu bar

The menu bar provides the following options:

Tasks and Notifications

Provides task and notification messages.

Help

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

Dashboard Settings

Gives access to user management and telemetry configuration.

User

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




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

Dashboard

The main dashboard displays specific information about the state of the cluster. The main dashboard can be accessed at any time by clicking **Dashboard** from the navigation menu. The dashboard landing page organizes the panes into different categories.

Figure: Ceph dashboard landing page



Details

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

Inventory

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

Hosts

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

Monitors

Displays the number of Ceph Monitors and the quorum status.

Managers

Displays the number and status of the Manager Daemons.

OSDs

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

Pools

Displays the number of storage pools in the Ceph cluster.

PGs

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

- activating
- backfill_wait
- backfilling
- creating
- deep
- degraded
- forced_backfill
- forced_recovery

- peering
- peered
- recovering
- recovery_wait
- repair
- scrubbing
- snaptrim
- snaptrim_wait

The **Warning** state includes PGs with any of the following states:

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

Object Gateways

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

Metadata Servers

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

Status

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

Capacity

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

Cluster Utilization

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

Used Capacity (RAW)

Displays usage in GiB.

IOPS

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

OSD Latencies

Displays total applies and commits per millisecond.

Client Throughput

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

Recovery Throughput

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

Related information

[Monitor the cluster](#)

Configuring the Dashboard

Configure the Red Hat Ceph Storage to provide a secure, web-based interface for monitoring cluster health, managing Ceph services, and performing administrative tasks.

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.

Prerequisites

- 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 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 [Enabling monitoring stack](#).

Setting message of the day (MOTD)

Sometimes, Ceph Dashboard users need to know about the latest news, updates, and information on Red Hat Ceph Storage.

PREREQUISITE

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

As a storage administrator, you can configure a message of the day (MOTD) in the command-line interface (CLI). When the user logs in to the Ceph Dashboard, the configured MOTD is displayed in a banner within the Ceph Dashboard, similar to the Telemetry module. You can configure MOTD based on the severity, such as `info`, `warning`, or `danger`.

You can close an MOTD with a `info` or `warning` severity. The `info` MOTD is not displayed anymore until the local storage cookies are cleared or a new MOTD with a different severity is displayed. A MOTD with a `warning` severity is displayed again in a new session.

1. Configure a MOTD for the dashboard.

Syntax

```
ceph dashboard motd set SEVERITY EXPIRES MESSAGE
```

Replace:

- `SEVERITY` can be `info`, `warning`, or `danger`.
- `EXPIRES` can be for seconds (`s`), minutes (`m`), hours (`h`), days (`d`), weeks (`w`), or never expires (`0`).
- `MESSAGE` can be any custom message that users can view when they log in to the dashboard.

For example,

```
[ceph: root@host01 /]# ceph dashboard motd set danger 2d "Custom login message"
Message of the day has been set.
```

2. **Optional:** Set the MOTD that does not expire.

For example,

```
[ceph: root@host01 /]# ceph dashboard motd set danger 0 "Custom login message"
Message of the day has been set.
```

3. Get the configured MOTD.
For example,

```
[ceph: root@host01 /]# ceph dashboard motd get
Message="Custom login message", severity="danger",
expires="2022-09-08T07:38:52.963882Z"
```

4. **Optional:** Clear the configured MOTD with the `clear` option.
For example,

```
[ceph: root@host01 /]# ceph dashboard motd clear
Message of the day has been cleared.
```

5. Verify the message of the day.

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

Related information

[Activating and deactivating telemetry](#)

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.

PREREQUISITE

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

- 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.

When disabling a feature, the web-interface elements become hidden and the associated REST API end-points reject any further requests for that feature. Enabling and disabling dashboard features can be done from the command-line interface or the web interface.

Available features:

- Ceph Block Devices:
 - Image management, `rbd`
 - Mirroring, `mirroring`
- Ceph File System, `cephfs`
- Ceph Object Gateway, `rgw`

Note:

- By default, the Ceph Manager is colocated with the Ceph Monitor.
- 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.

Toggling the dashboard features from the web interface

1. From the dashboard navigation, go to › **Administration › Manager Modules**.
2. Click **Edit**.
3. In the **Edit Manager module** form, you can enable or disable the dashboard features by selecting or clearing the check boxes next to the different feature names.

Figure: Editing manager modules

Edit Manager module

ACCOUNT_LOCKOUT_AT
TEMPTS

10

ALERTMANAGER_API_H
OST

http://host.containers.internal:9093

ALERTMANAGER_API_S
SL_VERIFY

☒

AUDIT_API_ENABLED

☐

AUDIT_API_LOG_PAYLO
AD

☒

ENABLE_BROWSABLE_
API

☒

FEATURE_TOGGLE_CEP
HFS

☒

FEATURE_TOGGLE_ISC
SI

☒

FEATURE_TOGGLE_MIR
RORING

☒

FEATURE_TOGGLE_NFS

☒

FEATURE_TOGGLE_RBD

☒

4. After the selections are made, click **Update**.

Toggling the dashboard features from the command-line interface

1. Login to cephadm shell.

```
cephadm shell
```

2. List the feature status.

```
ceph dashboard feature status
```

For example,

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

3. Disable a feature.

```
ceph dashboard feature disable FEATURE
```

The following example disables the Ceph Object Gateway feature.

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

4. Enable a feature.

```
ceph dashboard feature enable FEATURE
```

The following example enables the Ceph File System feature.

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

Changing the dashboard password

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*.
3. To change the password of **admin**, click its 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*.

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

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.GDpcga1QpYsaHUbJfoqaHd1rfFFx7XS", "roles":
["administrator"], "name": null, "email": null, "lastUpdate": , "enabled": true,
"pwdExpirationDate": null, "pwdUpdateRequired": false}
```

Verification

- Log in to the dashboard with your new password.

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.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster with the monitoring stack installed.
 - Root-level access to the `cephadm` host.
 - The dashboard module enabled.
1. As a root user, create a `grafana.yml` file. Provide the following details:

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

2. Mount the `grafana.yml` file under a directory in the container. For example,

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

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

- Optional: Verify and enable that the dashboard Ceph Manager module is enabled.

- Check that the dashboard module is enabled.

```
ceph mgr module ls
```

- Enable the dashboard module.

```
ceph mgr enable dashboard
```

- Apply the specification file.

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

For example,

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

- Redeploy the grafana service.

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

```
ceph orch redeploy grafana
```

AFTER COMPLETING THE TASK

Log into Grafana with the proper credentials.

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

For example,

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

Configuring Prometheus data retention

You can control how long Prometheus stores monitoring data and how much disk space it consumes by configuring retention settings in the Prometheus service specification. Prometheus is the component responsible for storing metrics; therefore, retention policies must be applied to the Prometheus spec.

PREREQUISITE

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

- Prometheus is deployed and managed by Cephadm.
- You have cluster administrator privileges.
- Prometheus is running.

```
ceph orch ls | grep prometheus
```

- Create or edit the prometheus.yml specification file.

```
service_type: prometheus
service_name: prometheus
placement:
  count: 1
spec:
  retention_time: "30d"      # Retain data for 30 days
  retention_size: "50GB"    # Limit storage to 50 GB
```

- Review the parameter descriptions:

retention_time

Defines how long Prometheus retains monitoring data. Supported units: y, w, d, h, m, s. The default value is 15 days.

retention_size

Defines the maximum size of Prometheus time-series database (TSDB). Supported units: B, KB, MB, GB, TB, PB, EB. The default value is 0 (disabled)

3. Apply the updated specification file.

```
ceph orch apply -i prometheus.yaml
```

4. Redeploy Prometheus to activate changes.

```
ceph orch redeploy prometheus
```

5. Verify that the retention configuration is active:

```
ceph orch ps | grep prometheus
```

Example

The following example configures Prometheus to retain monitoring data for one year or up to 100 GB.

```
service_type: prometheus
service_name: prometheus
placement:
  count: 1
spec:
  retention_time: "1y"
  retention_size: "100GB"
```

Prometheus automatically deletes older data when either the time-based or size-based threshold is reached.

Prometheus now enforces the configured retention policy. If Prometheus was already deployed before applying the updated specification, redeployment ensures that the new settings take effect.

Expanding the cluster

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, 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. For more information, see [Bootstrapping a new storage cluster](#).
- At least `cluster-manager` role for the user on the Red Hat Ceph Storage Dashboard. For more information, see [“User roles and permissions” on page 30](#).

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*.
 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.
 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*.

Upgrading a cluster

Upgrade Ceph clusters by using the dashboard.

PREREQUISITE

Verify that your upgrade version path and operating system is supported before starting the upgrade process. For more information, see [Compatibility matrix](#) and [Supported upgrade and crossgrade paths](#). Before starting to upgrade Red Hat Ceph Storage, see [Installation and upgrade requirements](#).

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

Important: These items cannot be done through the dashboard and must be completed manually, through the command-line interface, before continuing to upgrade the cluster from the dashboard.

Note: For detailed information, see [Upgrading cluster in a disconnected environment](#) and complete steps [1](#) through [9](#).

1. The latest `cephadm`.

```
dnf update cephadm
```

2. The latest `cephadm-ansible`.

```
dnf update cephadm-ansible
```

3. The latest `cephadm pre-flight` playbook.

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

4. Run the following Ceph commands to avoid alerts and rebalancing of the data during the cluster upgrade:

```
ceph health mute DAEMON_OLD_VERSION --sticky  
ceph osd set noout  
ceph osd set noscrub  
ceph osd set nodeep-scrub
```

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

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

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

Due to known issue [BZ #2168121](#), the updates are not retrieved properly during the `ceph orch upgrade ls` command that is being used. To upgrade, use the **Upgrade using custom image** option. Get the latest custom image name from [Custom image details for the monitoring stack in Configuring a private registry for a disconnected installation](#).

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

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

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

AFTER COMPLETING THE TASK

After verifying that the upgrade is complete, unset the `noout`, `noscrub`, and `nodeep-scrub` flags.

```
ceph osd unset noout
ceph osd unset noscrub
ceph osd unset nodeep-scrub
```

For example,

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

Managing access and single sign-on with the dashboard

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

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

To connect your SAML IdP with the Ceph Dashboard, use the following Service Provider (SP) configuration details. In each of the following configuration URLs, replace `DASHBOARD_HOST` with the fully qualified domain name (FQDN) and port of your dashboard. For example,

`ceph-dashboard.example.com:8443`

Entity ID (Audience URI)

The unique identifier for the Red Hat Ceph Storage Dashboard service provider.
`https://DASHBOARD_HOST/auth/saml2/metadata`

Assertion Consumer Service (ACS)

The endpoint where the IdP sends the SAML assertion (using POST binding).
`https://DASHBOARD_HOST/auth/saml2/`

Default Relay State (Start URL)

The page that users are directed to after a successful login.
`https://DASHBOARD_HOST/#/dashboard`

Target URL

`https://DASHBOARD_HOST/#/dashboard`

Note: Some IdPs can automatically configure these fields from the service provider metadata. To provide the metadata, provide the following URL to your IdP:
`https://DASHBOARD_HOST/auth/saml2/metadata`

OAuth2 SSO uses the `oauth2-proxy` service to work with the Ceph Management gateway (`mgmt-gateway`), providing unified access and improved user experience. For more information about the Ceph Management gateway and the OAuth2 Proxy service, see [Using the Ceph Management gateway \(`mgmt-gateway`\)](#) and [Using the OAuth2 Proxy \(`oauth2-proxy`\) service](#).

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

For more information about IBM Security Verify, see [IBM Security Verify documentation](#). For a free trial, see [IBM Security Verify \(SaaS\)](#).

Creating an admin account for syncing users to the Ceph dashboard

Create an admin account to synchronize users to the Ceph dashboard.

PREREQUISITE

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

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

For more information, see the following:

- [Creating roles](#).
 - [Creating users](#).
1. Download the *Red Hat Single Sign-On 7.4.0 Server* on the system where Red Hat Ceph Storage is installed. For download package and information, see [Software Details for Red Hat Single Sign-On 7.4.0 Server](#) on the [Red Hat Customer Portal](#).
 2. Extract the folder.

```
unzip rhssso-7.4.0.zip
```

For example,

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

3. Go to the `standalone/configuration` directory and open the `standalone.xml` file for editing. For example,

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

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

```
./add-user-keycloak.sh -u admin
```

For example,

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

5. Replace the **localhost** and 127.0.0.1 with the IP address of the server where the Red Hat SSO is installed.
6. Start the server.
From the bin directory of rh-ssso-7.4 folder, run the standalone boot script.

```
./standalone.sh
```

For example,

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

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

Note: You must create an admin account only the first time that you log into the console.

8. Log in to the admin console with the credentials created.

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.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.
- Ceph Dashboard installed.
- Admin-level access to the dashboard.
- Users added to the dashboard. See [“Creating users” on page 34](#).
- Root-level access on all the hosts.
- Admin account created for syncing users. See [“Creating an admin account for syncing users to the Ceph dashboard” on page 20](#).

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

1. To create a realm, select **Master** from the drop-down list.
In this realm, you can provide access to users and applications.
2. In the **Add Realm** window, enter a case-sensitive realm name, and set the parameter **Enabled** to ON and click **Create**.
3. In the **Realm Settings** tab, set the parameters and click **Save**.
Set the parameters as follows:
 - Enabled - ON
 - User-Managed - ON

Note 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 client ID parameter, and click **Save**.

BASE_URL:8443/auth/saml2/metadata

For example,

```
https://example.ceph.redhat.com:8443/auth/saml2/metadata ..  
Client Protocol - saml
```

6. In the **Client** window, **Settings** tab, set the parameters and click **Save**.
Use the parameters that are listed in “[Client Settings tab](#)” on page 22.

<i>Client Settings tab</i>		
Name of the parameter	Syntax	Example
Client ID	BASE_URL:8443/auth/saml2/metadata	https://example.ceph.redhat.com:8443/auth/saml2/metadata
Enabled	ON	ON
Client Protocol	saml	saml
Include AuthnStatement	ON	ON
Sign Documents	ON	ON
Signature Algorithm	RSA_SHA1	RSA_SHA1
SAML Signature Key Name	KEY_ID	KEY_ID
Valid Redirect URLs	BASE_URL:8443/*	https://example.ceph.redhat.com:8443/
Base URL	BASE_URL:8443	https://example.ceph.redhat.com:8443/
Master SAML Processing URL	https://localhost:8080/auth/realms/REALM_NAME/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.		

7. In **Fine Grain SAML Endpoint Configuration**, set the parameters and click **Save**.
Use the parameters that are listed in “[Fine Grain SAML Endpoint configuration](#)” on page 22.

<i>Fine Grain SAML Endpoint configuration</i>		
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
Logout Service Redirect Binding URL	BASE_URL:8443/	https://example.ceph.redhat.com:8443/

8. From **Clients > Mappers** tab, set the parameters, and click **Save**.
Use the parameters that are listed in “[Client Mappers tab](#)” on page 23.

<i>Client Mappers tab</i>	
Name of the parameter	Value
Protocol	saml
Name	username
Mapper Property	User Property
Property	username
SAML Attribute name	username

9. In the **Client Scope** tab, select the **role_list**.
 - a. In the **Mappers** tab, select the **role list**, set the **Single Role Attribute** to ON.
10. Select the **User_Federation** tab.
11. In the **User Federation** window, select **ldap** from the drop-down list.
12. From the **User Federation** window, **Settings** tab, set the parameters and click **Save**.
Use the parameters that are listed in “[User Federation Settings tab](#)” on page 23.

<i>User Federation Settings tab</i>	
Name of the parameter	Value
Console display name	rh-ldap
Import Users	ON
Edit_Mode	READ_ONLY
Username LDAP attribute	UID
RDN LDAP attribute	UID
UUID LDAP attribute	nsuniqueid
User Object Classes	inetOrgPerson, organizationalPerson, rhatPerson
Connection URL	Example: ldap://ldap.corp.redhat.com
Users DN	ou=users, dc=example, dc=com
Bind Type	None

13. Click **Test authentication**.
A notification displays with a message that the LDAP authentication is successful.
14. Click **Test Connection**.
A notification displays that the LDAP connection is successful.
15. In the **Mappers** tab, select the **first name** row, edit the **LDAP Attribute** parameter, and click **Save**.

- LDAP Attribute - uid

16. In the **User Federation Settings** tab, click **Synchronize all users**.

A notification displays with a message that the sync of users is finished successfully.

17. In the **Users** tab, search for the user added to the dashboard and click the Search icon. To view the user, click the specific row. The federation link is displayed as the name provided for the **User Federation**.

Important: LDAP does not synchronize users who are added manually. Click **Delete** if you have added the user manually.

Note: If Red Hat SSO is used within your work environment, be sure to first enable SSO. For more information, see [“Enabling single sign-on with SAML 2.0”](#) on page 24.

AFTER COMPLETING THE TASK

Users that are added to the realm and the dashboard can access the Ceph dashboard with their email address and password.

For example,

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

Related information

[Creating roles](#)

[Creating an admin account for syncing users to the Ceph dashboard](#)

Enabling single sign-on with SAML 2.0

Enable single sign-on (SSO) for the Ceph Dashboard with SAML 2.0.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Installation of the Ceph Dashboard.
 - Root-level access to the Ceph Manager hosts.
1. Configure SAML 2.0 SSO on the Ceph Dashboard.

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

Replace the following variables:

- *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.

For example,

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

2. Verify the SAML 2.0 configuration.

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

For example,

```
[ceph: root@host01 /]# ceph dashboard ceph dashboard sso show saml2
```

3. Enable SSO.

```
cephadm shell ceph dashboard sso enable saml2
```

For example,

```
[ceph: root@host01 /]# ceph dashboard ceph dashboard sso enable saml2
SSO is "enabled" with "SAML2" protocol.
```

AFTER COMPLETING THE TASK

1. Open your dashboard URL.
`https://dashboard_hostname.ceph.redhat.com:8443`
2. On the SSO page, enter the login credentials.
The SSO redirects to the dashboard web interface.

Note: Check the SSO status at any time with the `cephadm shell ceph dashboard sso status` command. For example:

```
[ceph: root@host01 /]# ceph dashboard ceph dashboard sso status
SSO is "enabled" with "SAML2" protocol.
```

Configuring Red Hat SSO as the identity provider using OAuth2 Protocol

You can enable single sign-on (SSO) for the Red Hat Ceph Storage Dashboard by configuring Red Hat SSO as the Identity Provider (IdP).

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- Installation of the Ceph Dashboard.
 - Enable the Ceph Management gateway (mgmt-gateway) service. For more information, see [Enabling the Ceph Management gateway](#).
 - Enable the OAuth2 Proxy service (oauth2-proxy). For more information, see [Enabling the OAuth2 Proxy service](#).

Ensure that oauth2 proxy and identity provider are configured before you can configure the Identity Provider.

1. Create a realm in Red Hat SSO.
 - a. Click the **Master** drop-down menu and select **Add Realm**.
 - b. Enter a case-sensitive name for the realm.
 - c. Set **Enabled** to **ON** and click **Create**.
2. Configure realm settings.
 - a. Go to the **Realm Settings** tab.
 - b. Set **Enabled** to **ON** and **User-Managed Access** to **ON**.
3. Click **Save**.
4. Copy the **OpenID Endpoint Configuration** URL. Use the issuer endpoint from this configuration for OAuth2 Proxy setup.
For example,

```
Issuer: http://<your-ip>:<port>/realms/<realm-name>
```

5. Create and configure the OIDC client.
 - a. Create a client of type **OpenID Connect**.
 - b. Enable client authentication and authorization.
 - c. Set redirect URLs.
 - **Valid redirect URL:** `https://<host_name|IP_address>/oauth2/callback`
 - **Valid logout redirect URL:** `https://<host_name|IP_address>/oauth2/sign_out`
 - d. Go to the **Credentials** tab and note the values for `client_id` and `client_secret`. These credentials are required to configure the oauth2-proxy service.
6. Create a role for dashboard access.
 - a. Navigate to the **Roles** tab.
 - b. Create a role and assign it to the OIDC client.
 - c. Use a role name supported by the Red Hat Ceph Storage dashboard.
7. Create a user and assign roles.
 - a. Create a new user.
 - b. Define user credentials.
 - c. Open the **Role Mappings** tab and assign the previously created role to the user.
8. Enable SSO for the dashboard.

```
ceph dashboard sso enable
```

SSO is now enabled for the Ceph management stack. When users access the dashboard, they are redirected to the Red Hat SSO login page. After successful authentication, users are granted access across all monitoring stacks, including dashboard, Grafana, and Prometheus, based on the roles assigned in Red Hat SSO.

Configuring IBM Security Verify as the identity provider using OAuth2 protocol

You can enable single sign-on (SSO) for the Red Hat Ceph Storage Dashboard by configuring IBM Security Verify as the Identity Provider (*IdP*).

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Installation of the Ceph Dashboard.
 - Enable the Ceph Management gateway (mgmt-gateway) service. For more information, see [Enabling the Ceph Management gateway](#).
 - Enable the OAuth2 Proxy service (oauth2-proxy). For more information, see [Enabling the OAuth2 Proxy service](#).
1. From the IBM Security Verify application, navigate to the **Applications > Application** and click **Add An Application**.
 2. Under Application Type, select **Custom Application** and click **Add Application**.
 3. Add the company name under the **General** tab.
 4. Navigate to **Sign-on** tab and select **Open ID Connect 1.0** from the **Sign-on Method** drop-down list.
 5. Enter the mgmt-gateway URL under the **Application URL** section.
For example,
`https://<HOSTNAME/IP_OF_THE_MGMT-GATEWAY_HOST>`
 6. Clear the **PKCE verification** checkbox.
 7. From the **Redirect URLs** section, fill in the redirect the URLs.
Use the following redirect URL format.
 - `https://<HOSTNAME/IP_OF_THE_MGMT-GATEWAY_HOST>/oauth2/callback`
 - `https://<HOSTNAME/IP_OF_THE_MGMT-GATEWAY_HOST>/oauth2/sign_out`
 8. From the **Token settings**, select **JWT** as the access token format.
For each of the following attributes, enter the details in the respective fields and then select the corresponding attribute source from the drop-down list:
 - `name` | `preferred_username`
 - `email` | `email`
 9. From the **Consent Settings**, select **Do not ask for consent** from the drop-down list.
 10. From the **Custom Scopes**, clear the **Restrict Custom Scopes** checkbox and click **Save**.
 11. Add a user.
 - a. Navigate to **Directory > Users & Groups** and click **Add User**.
 - b. Provide the desired username and a valid email address.
The initial user credentials are sent to the email address provided.
 - c. Click **Save**.
 12. Navigate to the **Directory > Attributes** and click **Add Attribute**.
 - a. Select **Custom Attribute**, and then select the **Single Sign On (SSO)** checkbox.
 - b. Click **Next**.
 - c. In the **Attribute Name** field, enter roles.

- d. Under **Data Type**, select **Multivalue String** from the drop-down list.
 - e. Click **Additional settings**, select **Lowercase** from the **Apply Transformation** drop-down list and click **Next**.
 - f. Click **Add Attribute**.
13. Navigate to **Directory > Users & Groups** and edit the newly added user.
 - a. Select **View Extended Profile**.
 - b. From **Custom User Attributes**, add your desired dashboard role.
For example, administrator, read-only, block-manager, and so on. This user inherit the dashboard roles automatically.
 - c. Click **Save settings**.
14. Go to **Applications > Application** and click the **Settings** icon for the custom application created in step “2” on page 27.
 - a. From the **Sign-on** tab, go to **Token settings**.
 - b. For the roles attribute, enter the detail in the respective fields and then select the corresponding attribute source from the drop-down list:
 - roles | roles
 - c. From the **Entitlements** tab add the newly created user and click **Save**.

Enabling OAuth2 single sign-on

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

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- Installation of the Ceph Dashboard.
- Root-level access to the Ceph Manager hosts.
- You should either have an admin account with Red Hat Single-Sign-On 7.6.0 with OAuth2 protocol running or IBM Security Verify with OAuth2 protocol running. For more information, see [“Configuring IBM Security Verify as the identity provider using OAuth2 protocol” on page 27](#).
- Enable the Ceph Management gateway (mgmt-gateway) service. For more information, see [Enabling the Ceph Management gateway](#).
- Enable the OAuth2 Proxy service (oauth2-proxy). For more information, see [Enabling the OAuth2 Proxy service](#).

1. Enable Ceph Dashboard OAuth2 SSO access.

```
ceph dashboard sso enable oauth2
```

For example,

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

2. Set the valid redirect URL.
`https://HOST_NAME|IP_ADDRESS/oauth2/callback`

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

3. Configure a valid user role.

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

AFTER COMPLETING THE TASK

1. Open your dashboard URL.
<https://<hostname/ip> address of mgmt-gateway-host>:<mgmt-gateway-port>
2. On the SSO page, enter the Identity Provider (IdP) credentials.
The SSO redirects to the dashboard web interface.

Note: Check the SSO status at any time with the `cephadm shell ceph dashboard sso status` command. For example:

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

Disabling single sign-on for the Ceph Dashboard

You can disable single sign-on for Ceph Dashboard.

PREREQUISITE

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

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

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

1. Optional: View the status of the SSO.

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

For example,

```
[ceph: root@host01 /]# ceph dashboard ceph dashboard sso status
SSO is "enabled" with "SAML2" protocol.
```

2. Disable SSO.

```
cephadm shell ceph dashboard sso disable
```

For example,

```
[ceph: root@host01 /]# ceph dashboard ceph dashboard sso disable
SSO is "disabled".
```

Managing roles

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.

User roles and permissions

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

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

- `cephfs` - Includes all features related to CephFS management.
- `config-opt` - Includes all features related to management of Ceph configuration options.
- `dashboard-settings` - Allows to edit the dashboard settings.
- `grafana` - Include all features related to Grafana proxy.
- `hosts` - Includes all features related to the Hosts menu entry.
- `log` - Includes all features related to Ceph logs management.
- `manager` - Includes all features related to Ceph manager management.
- `monitor` - Includes all features related to Ceph monitor management.
- `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

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.
- `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.

Note: you need to provide `rgw-manager` access to the users for all Ceph Object Gateway operations.

Figure: Displaying the user roles

User management / Roles

Users Roles

Search

Create +

Name	Description	System Role
administrator	allows full permissions for all security scopes	✓
block-manager	allows full permissions for rbd-image, rbd-mirroring, and iscsi scopes	✓
cephfs-manager	allows full permissions for the cephfs scope	✓
cluster-manager	allows full permissions for the hosts, osd, mon, mgr, and config-opt scopes	✓
ganesha-manager	allows full permissions for the nfs-ganesha scope	✓
pool-manager	allows full permissions for the pool scope	✓
read-only	allows read permission for all security scope except dashboard settings and config-opt	✓
rgw-manager	allows full permissions for the rgw scope	✓

Items per page: 10 1-8 of 8 items 1 of 1 page

For more information, see the following:

- [Creating Ceph Object Gateway](#)
- [Creating roles on the Ceph dashboard](#)

Creating roles

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

Prerequisites

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

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level of 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** form, enter the **Name**, **Description**, and select the **Permissions** for this role, and then click **Create Role**.

Create Role

Name (required)

NFS-RGW

Description

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

Permissions

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

Cancel

Create Role



A notification displays 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.

Reference

For more information, see the following:

- [“User roles and permissions” on page 30](#)
- [“Creating users” on page 34](#)

Editing roles

The dashboard allows you to edit roles on the dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin-level of 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, select the role to edit and click **Edit**.
4. In the **Edit Role** form, edit the parameters, and then click **Edit Role**.
A notification displays that the role was updated successfully.

Cloning roles

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 of access to the dashboard.
- Roles are created on the dashboard.

Procedure

1. From the dashboard menu, click the **Dashboard Settings** icon and then click **User management**.
2. On **Roles** tab, click the role you want to clone.
3. Select **Clone** from action drop-down.
4. In the **Clone Role** dialog box, enter the details for the role, and then click **Clone Role**.
Once you clone the role, you can customize the permissions as per the requirements.

Reference

For more information, see the following:

- [Creating roles](#).

Deleting roles

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

PREREQUISITE

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

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

- Admin-level of access to the dashboard.
- A custom role is created on the dashboard.

Note: You cannot delete the system roles of the Ceph Dashboard.

1. Log in to the Dashboard.
2. Click the **Dashboard Settings** icon and then select **User management**.
3. On the **Roles** tab, click the role you want to delete and select **Delete** from the action drop-down.
4. In the **Delete Role** notification, select **Yes, I am sure** and then click **Delete Role**.

Related information

[Creating roles](#)

Managing users

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 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.

Creating users

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 of 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.
2. Click the Dashboard Settings icon and then click **User management**.
3. On **Users** tab, click **Create**.
4. In the **Create User** window, set the **Username** and other parameters including the roles, and then click **Create User**.

Figure: Create User

Create User

Username *

dashboard_user

✓

Password ?

✓

👁

Confirm password

✓

👁

Password expiration date ?

Password expiration date...

✓

✕

Full name

Dashboard user

Email

dashboarduser@example.com

Roles

✎

There are no roles.

✓

Enabled

✓

User must change password at next login

Cancel

Create User

You get a notification that the user was created successfully.

Reference

For more information, see the following:

- [Creating roles.](#)
- [User roles and permissions.](#)

Editing users

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 of 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, from the **Users** tab, select the user row and click **Edit**.
4. In the **Edit User** form, edit parameters like password and roles, and then click **Edit User**.

Edit User

Username

Password ?

👁

Confirm password

👁

Password expiration date ?

✕

Full name

Email

Roles

🔑 administrator ✕

Cancel

Edit User

Note: If you want to disable any user's access to the Ceph dashboard, you can clear the **Enabled** option.

A notification displays that the user was created successfully.

Deleting users

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.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - Admin-level of access to the Dashboard.
 - User created on the dashboard.
1. Log in to the Dashboard.
 2. Click the **Dashboard Settings** icon and then select **User management**.
 3. On the **Users** tab, click the user you want to delete and select **Delete** from the action drop-down.
 4. In the **Delete User** dialog, select **Yes, I am sure** and then click **Delete User**.
In the **Delete User** dialog, select **Yes, I am sure** and then click **Delete User**.

Related information

[Creating users](#)

Managing 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 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.

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.

Reference

- [Access capabilities](#)

Accessing capabilities

Understand 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 run 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.

Reference

- [User capabilities](#)

Creating user capabilities

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

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.

- Dashboard is installed.
- Admin-level access to the dashboard.

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

1. From the dashboard navigation, go to **Administration > Ceph Users**.
2. Click **Create**.
3. Fill in the **Create User** form.
 - a. **User Entity:** Enter as `TYPE.ID`
 - b. **Entity:** This can be `mon`, `mgr`, `osd`, or `mds`.
 - c. **Entity Capabilities:** Enter the capabilities that you can 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.

Administration > Ceph Users > Create

Create User

User entity *

client.test1

✓

Capabilities

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

Cancel

Create User

4. Click **Create User**.
A notification displays that the user is created successfully.

Editing user capabilities

Edit the roles of users or clients on the dashboard.

PREREQUISITE

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

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

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

1. From the dashboard navigation, go to **Administration > Ceph Users**.
2. Click **Edit**.
3. Fill in the **Edit User** form.
 - a. Edit the entity and entity capabilities.

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

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

Importing user capabilities

Import the roles of users or clients from 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.

For 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. Go to **Administration > Ceph Users**.
4. Select the user whose roles you want to import and click **Import** from the action drop-down.
5. In the **Import User** window, click **Choose File** and select the appropriate file.
6. Click **Import User**.
A notification displays that the keys are imported successfully.

Exporting user capabilities

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

PREREQUISITE

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

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

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

1. From the dashboard navigation, go to **Administration > Ceph Users**.
2. Select the user whose roles you want to export.
3. Select **Export** from the action drop-down.
4. From the **Ceph user export data** dialog, click **Copy to Clipboard**.
A notification displays that the keys are successfully copied.
5. Click **Close**.
6. On your local system, create a keyring file and paste the keys.

```
[localhost:~]$ cat exported.keyring  
  
[client.test11]  
key = AQD9S29kmjgJFxAakvhFar6Af3AWKDY2DsULRg==  
caps mds = "allow *"  
caps mgr = "allow *"  
caps mon = "allow *"  
caps osd = "allow r"
```

Deleting user capabilities

Delete the roles of users or clients on the dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - Admin-level access to the dashboard.
1. From the dashboard navigation, go to **Administration > Ceph Users**.
 2. Select the user that you want to delete and select **Delete** from the action drop-down.
 3. In the **Delete user** dialog, select **Yes, I am sure** and click **Delete user**.
A notification displays that the user is deleted successfully.

Managing Ceph services

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

Note: In the Ceph Dashboard, daemon processes are referred to as services. While this documentation may reference daemons when describing underlying behavior, all management actions in the dashboard are performed on services.

Running Ceph services

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

PREREQUISITE

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

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

Note: In the Ceph Dashboard, Ceph daemon processes are managed as services. While the UI and this documentation may refer to daemons to reflect underlying Ceph components, all management actions in the dashboard are performed at the service level

Note: Daemon actions are supported on all daemons except monitor and manager daemons.

Managing daemons can be done in the following ways:

- [“Managing daemons from Services” on page 43](#)
- [“Managing daemons from Hosts” on page 44](#)

Managing daemons from Services

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

Note: The row can be collapsed at any time.

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

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

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

Figure: Managing daemons from Services

Administration / Services

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

Daemons Service Events

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

Items per page: 10 1-3 of 3 items

▼ crash * 3 / 3 3 minutes ago

Managing daemons from Hosts

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

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

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

Figure: Managing daemons from Hosts

The screenshot shows the Red Hat Ceph Storage Dashboard. The top navigation bar includes 'Cluster / Hosts', 'Hosts List', and 'Overall Performance'. The 'Hosts List' tab is active, displaying a table of hosts. The first host is 'ceph-vpap-i80-9bx9bs-node1-installer (10.0.66.213)'. It has several service instances: 'mon: 1', 'mgr: 1', 'ceph-exporter: 1', 'crash: 1', 'node-exporter: 1', 'grafana: 1', 'prometheus: 1', 'alertmanager: 1', and 'osd: 6'. The host status is 'available'. 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'. The daemons listed are 'alertmanager.ceph-vpap-i80-9bx9bs-node1-installer', 'ceph-exporter.ceph-vpap-i80-9bx9bs-node1-installer', 'crash.ceph-vpap-i80-9bx9bs-node1-installer', 'grafana.ceph-vpap-i80-9bx9bs-node1-installer', and 'mgr.ceph-vpap-i80-9bx9bs-node1-installer.suhhdm'. All daemons are in a 'running' status. A context menu is open over the 'alertmanager' daemon, showing options: 'Start', 'Stop', 'Restart', and 'Redeploy'.

Managing hosts

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.

In addition to removing hosts, you can also remove hosts using *Start Drain* and *Remove* options in the Red Hat Ceph Storage 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.

Entering maintenance mode

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.

PREREQUISITE

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

- 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.
1. From the dashboard navigation, go to **Cluster > Hosts**.
 2. From the **Hosts List**, select the host to put in maintenance.
 3. Click **Enter Maintenance** from the action drop-down.

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.

A notification is displayed that the host has entered maintenance successfully and the maintenance label is seen in the **Status** column.

Note: If the maintenance mode fails, you get a notification indicating the reasons for failure.

Exiting maintenance mode

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

PREREQUISITE

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

- 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.
1. From the dashboard navigation, go to **Cluster > Hosts**.
 2. From the **Hosts List**, select the host that is in maintenance.

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

3. Click **Exit Maintenance** from the action drop-down.
After exiting the maintenance mode, you need to create the required services on the host by default-crash and the node-exporter gets deployed.
A notification is displayed that the host has been moved out of maintenance successfully and the maintenance label is removed from the **Status** column.

Removing hosts

To remove a host from a Ceph cluster, you can use `Start DrainandRemoveoptions` in Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

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

2. Select the host that is to be removed, and click **Start Drain** from the action drop-down. 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. To stop the draining of daemons from the host, click **Stop Drain** option from the action drop-down menu.
3. Check that all the daemons are removed from the host.
 - a. Expand the host row.
 - b. Go to the **Daemons** tab.
No daemons should be listed.
 4. Select the host that is to be removed, and click **Remove** from the action drop-down.
 5. In the **Remove Host** dialog, select **Yes, I am sure** and click **Remove Host**.
A notification displays that the host is removed successfully.

Managing pools

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

Creating pools

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.

PREREQUISITE

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

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

For more information about pools, see [Pools](#).

1. From the dashboard navigation, go to **Cluster > Pools**.
2. Click **Create**.
3. Fill out the **Create Pool** form.
The form changes based off of selection. Not all fields are mandatory.
 - Set the name of the pool .
 - Select the **Pool type**, either `replicated` or `erasure`. Erasure is referred to as Erasure Coded (EC).
 - Optional: Select if the the **PG Autoscale** is `on`, `off`, or `warn`.
 - Enter the number of Placement groups (PGs).
 - 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 **Information** icon.
 - Optional: Add a new EC profile by clicking the **Add** icon.
 - Optional: Click the **Edit** icon to select an application for the pool.
 - Optional: Set the CRUSH rule, if applicable.
 - Optional: If compression is required, select `passive`, `aggressive`, or `force`.
 - Optional: Set the Quotas.

- Optional: Set the Quality of Service configuration.
4. Save the changes and complete creating the pool, by clicking **Create Pool**.
A notification displays that the pool was created successfully.

Editing pools

Edit the pools on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool is created
1. From the dashboard navigation, go to **Cluster > Pools**.
 2. Select the pool that needs to be edited, and click **Edit**.
 3. In the **Edit Pool** form, edit the required parameters and click **Edit Pool**.
A notification displays that the pool was updated successfully.

Deleting pools

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.

PREREQUISITE

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

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

For more information about pools, see the following:

- [Pools](#)
- [Pool values](#)

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

Note: If needed, clear filters and search for the configuration.

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

Managing Ceph OSDs from the Ceph 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.

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

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

Managing the OSDs

Manage Ceph OSD lifecycle and maintenance tasks by using the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

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

Use the Red Hat Ceph Storage Dashboard to centrally manage Object Storage Daemons (OSDs) and perform common operational and maintenance tasks.

From the dashboard, you can complete the following OSD management actions:

- Create new OSDs.
- Edit the device class assigned to an OSD.
- Set OSD flags, such as **No Up**, **No Down**, **No In**, or **No Out**.
- Run scrub and deep-scrub operations on OSDs.
- Reweight OSDs to rebalance data placement.
- Change OSD state to **Out**, **In**, **Down**, or **Lost**.
- Purge, destroy, or delete OSDs as part of decommissioning workflows.

To access OSD management, from the dashboard navigation, click **Cluster > OSDs**.

Creating an OSD

1. From the **OSDs List** page, click **Create**.

Note: Ensure you have an available host and a few available devices. You can check for available devices in **Cluster > Physical Disks**.

2. In the **Create OSDs** form, from the **Deployment Options** section, select a deployment option.
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.

- Optional: Add additional devices in the **Advanced Mode** section, 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**.

- For security, select **Encryption** in the **Features** section.
- Click **Create**.
A notification displays that the OSD was created successfully and the OSD status changes from **in** and **down** to **in** and **up**.

Editing an OSD

- Select an OSD in the **OSDs List**, and click **Edit**.
- In the **Edit OSD** dialog, edit the device class and click **Edit OSD**.
A notification displays that the OSD was updated successfully.

Marking OSD flags

- Select the OSD in the **OSDs List**, and click **Flags** from the action drop-down.
- Select the flags needed and click **Update**.
A notification displays that the OSD flags updated successfully.

Scrubbing OSDs

- Select the OSD in the **OSDs List**, and click **Scrub** from the action drop-down.
- In the **OSDs Scrub** dialog, click **Update**.
A notification displays that the OSD scrubbing was initiated successfully.

Deep scrubbing

- Select the OSD in the **OSDs List**, and click **Deep Scrub** from the action drop-down.
- In the **OSDs Deep Scrub** dialog, click **Update**.
A notification displays that the OSD deep scrubbing was initiated successfully.

Reweighting OSDs

- Select the OSD in the **OSDs List**, and click **Reweight** from the action drop-down.
- In the **Reweight OSD** dialog, enter a value between 0 and 1.
- Click **Reweight**.

Marking OSDs Out

- Select the OSD in the **OSDs List**, and click **Mark Out** from the action drop-down.
- In the **Mark OSD out** notification, click **Mark out**.
The OSD status changes to **Out**.

Marking OSDs In

1. Select the OSD in the **OSDs List**, and click **Mark In** from the action drop-down.
2. In the **Mark OSD in** notification, click **Mark in**.
The OSD status changes to In.

Marking OSDs Down

1. Select the OSD in the **OSDs List**, and click **Mark Down** from the action drop-down.
2. In the **Mark OSD down** notification, click **Mark Down**.
The OSD status changes to Down.

Marking OSDs Lost

1. Select the OSD that is in out and down status from the **OSDs List**.
2. Click **Mark Lost** from the action drop-down.
3. In the **Mark OSD Lost** dialog, select **Yes, I am sure**, and click **Mark Lost**.

Purging OSDs

1. Select the OSD that is in the down status from the **OSDs List**.
2. Click **Purge** from the action drop-down.
3. In the **Purge OSDs** dialog, select **Yes, I am sure**, and click **Purge OSD**.
All the flags are reset and the OSD is back in in and up status.

Destroying OSDs

1. Select the OSD that is in the down status from the **OSDs List**.
2. Click **Destroy** from the action drop-down.
3. In the **Destroy OSDs** dialog, select **Yes, I am sure**, and click **Destroy OSD**.
The status of the OSD changes to destroyed.

Deleting OSDs

1. Select the OSD that is in the down status from the **OSDs List**.
2. Click **Delete** from the action drop-down.
3. In the **Delete OSDs** dialog, select **Yes, I am sure**, and click **Delete OSD**.

Note: You can preserve the OSD_ID when you have to replace the failed OSD.

Replacing the failed OSDs

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.

PREREQUISITE

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


- 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
1. From the dashboard, identify the failed OSDs.
The failed OSDs can be identified in any of 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.

You can also view the LED blinking lights on the physical drive if one of the OSDs is down.

“[Figure: Health status of OSDs on dashboard](#)” on [page 51](#) shows one down and one out OSD from the Ceph Dashboard landing page.

Figure: Health status of OSDs on dashboard



2. From **Cluster > OSDs** on the **OSDs List** table, select the out and down OSD.
 - a. Click **Flags** from the action drop-down, select **No Up** in the **Individual OSD Flags** form, and click **Update**.
 - b. Click **Delete** from the action drop-down. In the **Delete OSD** dialog, select **Preserve OSD ID(s) for replacement** and **Yes, I am sure** and click **Delete OSD**.
 - c. Wait until the status of the OSD changes to **out and destroyed**.
3. Optional: To change the **No Up** flag for the entire cluster, from the **Cluster-wide configuration** menu, select **Flags**.
 - a. In the **Cluster-wide configuration** form, select **No Up** and click **Update**.
4. Optional: If the OSDs are down due to a hard disk failure, replace the physical drive.
 - If the drive is hot-swappable, replace the failed drive with a new one.
 - If the drive is not hot-swappable and the host contains multiple OSDs, you might have to shut down the whole host and replace the physical drive. Consider preventing the cluster from backfilling. For more information, see [Stopping and Starting Rebalancing](#).
 - 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:

```
ceph orch device zap HOST_NAME PATH --force
```

For example,

```
ceph orch device zap ceph-adm2 /dev/sdc --force
```


5. From **Cluster > OSDs** on the **OSDs List** table, click **Create**.
6. In the **Create OSDs** from **Advanced Mode** section, add a primary device.
 - a. In the **Primary devices** dialog, select a **Hostname** filter.
 - b. Select a device type from the list.

Note: You have to select the Hostname first and then at least one filter to add the devices. For example, from **Hostname** list, select **Type** and then **hdd**.

- c. Select **Vendor** from the device list, select **ATA**.
- d. Click **Add**.

“Figure: Using Primary devices filter” on page 52 shows how to select the **hdd** type in the Primary devices filter.

Figure: Using Primary devices filter



7. In the **Create OSDs** from, click **Preview**.
8. In the **OSD Creation Preview** dialog, click **Create**.
A notification displays that the OSD is created successfully and the OSD changes to be in the **outanddown** status.
9. Select the newly created OSD that has **outanddown** status.
 - a. From the action drop-down, click **Mark In**.
 - b. In the **Mark OSD in** notification, click **Mark In**.
The OSD status changes to **in**.
 - c. Click **Flags** from the action drop-down.
 - d. Clear the **No Up** selection and click **Update**.
10. Optional: If you have changed the **No Up** flag before for cluster-wide configuration, in the **Cluster-wide configuration** menu, select **Flags**.
 - a. In **Cluster-wide OSDs Flags** form, clear the **No Up** selection and click **Update**.

AFTER COMPLETING THE TASK

Verify that the OSD that was destroyed is created on the device and the OSD ID is preserved.

Managing alerts

As a storage administrator, you can see the details of alerts and create silences for them on the Red Hat Ceph Storage Dashboard.

The following table lists the following pre-defined alerts.

Alert name	Severity	Description
CephHealthError	Critical	This alert is triggered when the cluster remains in HEALTH_ERROR state for more than 5 minutes. This indicates a serious issue affecting cluster stability or data availability.
CephHealthWarning	Warning	This alert is triggered when the cluster remains in HEALTH_WARN state for more than 15 minutes. This indicates degraded performance or potential risks that require attention but are not yet critical.
CephMonDownQuorumAtRisk	Critical	This alert is triggered when monitor quorum is at risk. Quorum requires a majority of monitors (for example, 3 of 5) to be active. Without quorum, the cluster becomes inoperable and all services and clients are affected.
CephMonDown	Warning	This alert is triggered when one or more monitors are down, but quorum is still intact. If another monitor goes down, quorum will be lost and the cluster will become inoperable.
CephMonDiskSpaceCritical	Critical	This alert is triggered when the free space available to a monitor's store is critically low. If space runs out, monitor services may stop functioning, which can affect quorum and cluster stability. The default store directory is: • /var/lib/ceph/mon-*/data/store.db for traditional deployments • /var/lib/rook/mon-*/data/store.db on the mon pod's worker node in Rook
CephMonDiskSpaceLow	Warning	This alert is triggered when the space available to a monitor's store is approaching full (default threshold: >70%). If space usage continues to grow, monitor services may fail, putting quorum and cluster stability at risk. The default store directory is: • /var/lib/ceph/mon-*/data/store.db for traditional deployments • /var/lib/rook/mon-*/data/store.db on the mon pod's worker node in Rook
CephMonClockSkew	Warning	This alert is triggered when Ceph monitors detect a time drift (clock skew). Monitors rely on synchronized time to maintain quorum and cluster consistency. If time drift is too high, quorum may be lost and cluster operations disrupted.
CephOSDDownHigh	Critical	This alert is triggered when more than 10% of OSDs are down. If too many OSDs fail, the cluster may lose data redundancy and client I/O performance will be impacted.
CephOSDHostDown	Warning	This alert is triggered when one or more OSD hosts are offline. When a host goes down, all OSDs on that host also become unavailable, which may reduce data redundancy and impact client performance.
CephOSDDown	Warning	This alert is triggered when one or more OSDs have been down for more than 5 minutes. If OSDs remain down, data redundancy and cluster performance may be impacted.

Alert name	Severity	Description
CephOSDNearFull	Warning	This alert is triggered when one or more OSDs reach the NEARFULL threshold. This condition indicates that available disk space is running low and action is required to prevent OSDs from becoming full, which could impact data placement and cluster operations.
CephOSDFull	Critical	This alert is triggered when one or more OSDs reach the FULL threshold. When this occurs, writes to pools that use the affected OSDs are blocked. This can halt client I/O and place the cluster at risk until the condition is resolved.
CephOSDBackfillFull	Warning	This alert is triggered when one or more OSDs reach the BACKFILL FULL threshold. At this point, backfill and rebalance operations cannot complete, which may prevent the cluster from recovering from failures or redistributing data.
CephOSDTooManyRepairs	Warning	This alert is triggered when reads from an OSD frequently rely on a secondary placement group (PG) to return data. This behavior indicates that the OSD may be experiencing a high number of read errors, which is often a sign of a failing drive.
CephOSDTimeoutsPublicNetwork	Warning	This alert is triggered when OSD heartbeats on the public (frontend) network are running slow. This condition may indicate network latency or packet loss, which can affect OSD communication and client I/O performance.
CephOSDTimeoutsClusterNetwork	Warning	This alert is triggered when OSD heartbeats on the cluster (backend) network are slow. This condition may indicate network latency or packet loss on the backend subnet, which can disrupt replication, recovery, and overall cluster performance.
CephOSDInternalDiskSizeMismatch	Warning	This alert is triggered when one or more OSDs show an internal inconsistency between their metadata and the actual size of the underlying device. This mismatch can cause instability and may lead to OSD crashes in the future.
CephDeviceFailurePredicted	Warning	This alert is triggered when the device health module predicts that one or more devices will fail soon. To confirm, review device status with <code>ceph device ls</code> or inspect a specific device with <code>ceph device info <dev id></code> . To remediate, mark the affected OSD out to allow data migration, then replace the failing device and redeploy the OSD.
CephDeviceFailurePredictionTooHigh	Critical	This alert is triggered when the device health module predicts multiple device failures, but remediation cannot proceed automatically. Too many OSDs would need to be removed, which risks cluster performance and availability. To prevent data integrity issues, add new OSDs so data can be safely relocated.
CephDeviceFailureRelocationIncomplete	Warning	This alert is triggered when the device health module predicts one or more device failures, but the relocation of data from those devices is blocked. This usually happens when the cluster does not have enough free space to migrate data, or when the balancer is disabled. Ensure that the cluster has available capacity or enable the balancer to allow migration.
CephOSDFlapping	Warning	This alert is triggered when an OSD repeatedly flaps between up and down states within a short interval (marked down and back up multiple times in 5 minutes). This behavior usually indicates network-related issues such as latency, packet loss, or MTU mismatches on the cluster network (or the public network if no cluster network is deployed). Administrators should check the network statistics on the affected host(s).
CephOSDReadErrors	Warning	This alert is triggered when an OSD encounters read errors. The OSD is able to recover by retrying the reads, but the errors may point to underlying hardware issues (such as a failing disk) or kernel-level problems. Administrators should investigate hardware logs, kernel logs, and device health status to identify the cause.

Alert name	Severity	Description
CephPGImbalance	Warning	This alert is triggered when the number of placement groups (PGs) assigned to one or more OSDs deviates by more than 30% from the cluster average. Such imbalance may lead to uneven resource utilization and performance issues. Investigate the cluster to ensure PG autoscaler is active and pool configurations are correct.
CephFilesystemDamaged	Critical	This alert is triggered when CephFS metadata becomes corrupted, which may render data inaccessible. Administrators should analyze metrics from the MDS daemon's admin socket or escalate the issue to support for recovery actions.
CephFilesystemOffline	Critical	This alert is triggered when all MDS (Metadata Server) ranks are unavailable. With no active MDS daemons, the CephFS filesystem cannot serve metadata requests and becomes offline.
CephFilesystemDegraded	Critical	This alert is triggered when one or more CephFS Metadata Server (MDS) ranks are failed or in a damaged state. This condition can lead to partial filesystem availability, or in severe cases, complete failure.
CephFilesystemMDSRanksLow	Warning	This alert is triggered when the number of active CephFS Metadata Server (MDS) daemons is less than the configured <code>max_mds</code> value for the filesystem. This can reduce performance and affect the scalability of CephFS operations.
CephFilesystemInsufficientStandby	Warning	This alert is triggered when the number of standby MDS daemons is less than the value configured by <code>standby_count_wanted</code> . Insufficient standby daemons reduce fault tolerance and may delay recovery if an active MDS fails.
CephFilesystemFailureNoStandby	Critical	This alert is triggered when an active MDS daemon fails and no standby MDS is available to take over. This leaves the filesystem at risk of metadata service unavailability.
CephFilesystemReadOnly	Critical	This alert is triggered when the CephFS filesystem switches to <code>read-only</code> mode due to unexpected write errors in the metadata pool. This condition prevents metadata updates and may cause disruptions in filesystem operations.
CephMgrModuleCrash	Critical	This alert is triggered when one or more Ceph manager (mgr) modules have crashed and remain unacknowledged by an administrator. A crashed module may impact cluster functionality. Use the <code>ceph crash</code> command to identify and archive the failure.
CephMgrPrometheusModuleInactive	Critical	This alert is triggered when the <code>mgr/prometheus</code> module becomes unreachable. This may occur if the module is disabled or if the Ceph manager daemon itself is down. Without this module, cluster metrics and alerts will no longer function. Administrators should check module status with <code>ceph mgr module ls</code> and re-enable it if needed.
CephPGsInactive	Critical	This alert is triggered when one or more placement groups (PGs) remain inactive for more than 5 minutes in a pool. Inactive PGs cannot serve read/write requests, which impacts data availability. Administrators should investigate OSD and pool status to identify and resolve the cause.
CephPGsUnclean	Warning	This alert is triggered when one or more placement groups (PGs) remain in an <code>unclean</code> state for more than 15 minutes. Unclean PGs have not fully recovered from a failure, which may impact data reliability and availability. Administrators should check cluster health and recovery status to identify and fix the underlying issue.
CephPGsDamaged	Critical	This alert is triggered when one or more placement groups (PGs) are flagged as damaged or inconsistent during a data consistency check (scrub). Such issues may prevent reliable data access. Administrators should identify the affected PGs using <code>rados list-inconsistent-pg <pool></code> and repair them with <code>ceph pg repair <pg_num></code> .

Alert name	Severity	Description
CephPGRecoveryAtRisk	Critical	This alert is triggered when one or more OSDs have reached or exceeded the full threshold , putting data redundancy at risk. Recovery or backfill operations cannot proceed, which may compromise cluster resilience. To resolve, add storage capacity, restore down/out OSDs, or delete unnecessary data.
CephPGUnavailableBlockingIO	Critical	This alert is triggered when one or more placement groups (PGs) become unavailable, reducing data availability and blocking I/O requests in the cluster. This condition prevents clients from reading or writing data until resolved.
CephPGBackfillAtRisk	Critical	This alert is triggered when one or more OSDs have reached the <i>backfillfull</i> threshold, preventing data backfill operations. This puts data redundancy at risk since rebalancing cannot proceed. To resolve, add capacity to the cluster or delete unnecessary data.
CephPGNotScrubbed	Warning	This alert is triggered when one or more placement groups (PGs) have not been scrubbed within the expected interval. Scrubbing validates metadata integrity and ensures consistency across replicas, protecting against bit-rot. Missed scrubs may indicate an insufficient scrub window or that PGs were not in a <i>clean</i> state. You can manually trigger a scrub with <code>ceph pg scrub <pgid></code> .
CephPGsHighPerOSD	Warning	This alert is triggered when the number of placement groups (PGs) assigned per OSD exceeds the configured <code>mon_max_pg_per_osd</code> limit. An excessive PG count can impact cluster performance and stability. Verify that the PG autoscaler is enabled and properly configured using <code>ceph osd pool autoscale-status</code> . Adjust pool parameters such as <code>target_size_ratio</code> if necessary, or manually tune <code>pg_num</code> values to restore balance.
CephPGNotDeepScrubbed	Warning	This alert is triggered when one or more placement groups (PGs) have not undergone a deep scrub within the expected interval. Deep scrubs are critical for detecting and correcting data inconsistencies (bit-rot) by comparing all data replicas. Missed deep scrubs may indicate insufficient scrub windows or PGs not being in a clean state when scheduled.
CephNodeRootFilesystemFull	Critical	This alert is triggered when the root filesystem on a Ceph node has less than 5% free space remaining. A nearly full root volume can cause node instability, daemon crashes, and prevent essential services from running. Immediate action is required to free up space or extend the root volume.
CephNodeRootFilesystemFull	Critical	This alert is triggered when the root filesystem on a Ceph node has less than 5% free space remaining. A nearly full root volume can cause node instability, daemon crashes, and prevent essential services from running. Immediate action is required to free up space or extend the root volume.
CephNodeNetworkPacketDrops	Warning	This alert is triggered when a node's network interface reports packet drops greater than 0.5% or more than 10 packets per second. Persistent packet drops can indicate network congestion, faulty cables, or NIC issues, potentially impacting cluster performance and client I/O.
CephNodeNetworkPacketErrors	Warning	This alert is triggered when a node's network interface reports packet errors greater than 0.01% or more than 10 packets per second. Such errors may indicate faulty hardware, driver issues, or misconfigured NIC settings, which can impact cluster communication and performance.
CephNodeNetworkBondDegraded	Warning	This alert is triggered when a bonded network interface has fewer active slave interfaces than expected, indicating degraded redundancy or performance. A degraded bond can reduce fault tolerance and throughput on the affected node.

Alert name	Severity	Description
CephNodeDiskpaceWarning	Warning	This alert is triggered when a host filesystem mountpoint is predicted to run out of free space within 5 days, based on the trailing 48-hour growth rate. Running out of space can impact node stability and Ceph services.
CephNodeInconsistentMTU	Warning	This alert is triggered when one or more Ceph nodes have a network interface MTU (Maximum Transmission Unit) size that differs from the cluster median. Inconsistent MTU values can cause packet fragmentation, latency, or dropped connections across the cluster.
CephPoolGrowthWarning	Warning	This alert is triggered when a pool is projected to reach full capacity within 5 days based on the average data growth rate observed over the past 48 hours. Continuous growth without adding capacity can lead to pool saturation and write failures.
CephPoolBackfillFull	Warning	This alert is triggered when a pool is nearing the <i>backfillfull</i> threshold, preventing recovery or backfill operations from completing. This condition can impact data redundancy and recovery processes. Add storage capacity or free up space to resolve the issue.
CephPoolFull	Critical	This alert is triggered when a pool has reached its maximum quota or when the OSDs supporting the pool have reached the <i>FULL</i> threshold. In this state, all writes to the affected pool are blocked to prevent data loss. To resolve the issue, add storage capacity or increase the pool's quota after expanding the cluster.
CephPoolNearFull	Warning	This alert is triggered when a pool exceeds the <i>nearfull</i> threshold or when the supporting OSDs reach the warning capacity limit. Although writes can continue, the pool is at risk of becoming read-only if additional capacity is not added. Review pool usage and increase capacity or quota as needed.
CephSlowOps	Warning	This alert is triggered when one or more OSD operations take longer than the configured complaint threshold (<i>osd_op_complaint_time</i>). Prolonged slow operations may indicate high latency, overloaded OSDs, or network issues affecting cluster performance.
CephDaemonSlowOps	Warning	This alert is triggered when a Ceph daemon reports operations taking longer than the defined complaint threshold. Persistent slow operations may indicate resource contention, network latency, or backend storage performance issues.
CephadmUpgradeFailed	Critical	This alert is triggered when the Ceph cluster upgrade process initiated by cephadm fails, leaving the cluster in an incomplete or unstable state. Review the cephadm logs to identify the cause and take corrective action before retrying the upgrade.
CephadmDaemonFailed	Critical	This alert is triggered when a daemon managed by cephadm becomes inactive or fails to respond. This may impact the cluster's overall functionality depending on the failed service. Use <code>ceph health detail</code> to identify the affected daemon and restart it using <code>ceph orch daemon start <daemon_id></code> .
CephadmPaused	Warning	This alert is triggered when cluster management operations have been manually paused in cephadm. While paused, the orchestrator stops performing automated service reconciliation and health management tasks. Resume operations using <code>ceph orch resume</code> if the pause was unintentional.
HardwareStorageError	Critical	This alert is triggered when one or more storage devices report hardware errors. Such issues can lead to data inaccessibility or degraded cluster performance. Check <code>ceph health detail</code> and system logs to identify and replace the affected devices.

Alert name	Severity	Description
HardwareMemoryError	Critical	This alert is triggered when memory (DIMM) errors are detected on one or more cluster nodes. Such errors can cause instability, crashes, or data corruption. Check <code>ceph health detail</code> and system hardware logs to identify and replace faulty memory modules.
HardwareProcessorError	Critical	This alert is triggered when processor (CPU) hardware errors are detected on one or more cluster nodes. Such issues can degrade performance, cause instability, or impact daemon operations. Check <code>ceph health detail</code> and system logs to identify the affected CPU and replace faulty components.
HardwareNetworkError	Critical	This alert is triggered when one or more network-related hardware errors are detected on Ceph cluster nodes. Such issues may lead to packet loss, latency, or communication failures between daemons. Check <code>ceph health detail</code> and system network logs to identify and resolve the fault.
HardwarePowerError	Critical	This alert is triggered when one or more power supply units report hardware errors. Such failures can lead to unexpected node shutdowns, degraded redundancy, or hardware instability. Check <code>ceph health detail</code> and system hardware logs to identify and replace the faulty power supply.
HardwareFanError	Critical	This alert is triggered when one or more cooling fans report hardware errors. Faulty or failed fans can lead to overheating, degraded performance, or potential hardware damage. Check <code>ceph health detail</code> and system logs to identify and replace the affected fan units.
PrometheusJobMissing	Critical	This alert is triggered when the Prometheus scrape job responsible for collecting Ceph metrics is missing or undefined. As a result, Ceph metrics and alerts will not be available. Review and restore the ceph job configuration in the <code>prometheus.yml</code> file to re-enable monitoring.
CephObjectMissing	Critical	This alert is triggered when the latest version of a RADOS object cannot be found, even though all OSDs are up. Client I/O requests for the missing object will hang until it is resolved. Recovery may require manually rolling back the object to a previous version and verifying data integrity.
CephDaemonCrash	Critical	This alert is triggered when one or more Ceph daemons have recently crashed and the crashes have not yet been acknowledged. These crashes may affect cluster stability or functionality. Use <code>ceph crash ls</code> to list unacknowledged crashes and <code>ceph crash archive <id></code> to acknowledge them.
CephRBDMirrorImagesPerDaemonHigh	Critical	This alert is triggered when the number of RBD image replications handled by a single mirroring daemon exceeds the threshold of 100. A high replication count can lead to performance degradation and replication delays. Consider deploying additional mirroring daemons or redistributing the replication workload.
CephRBDMirrorImagesNotInSync	Critical	This alert is triggered when local and remote RBD mirror images are out of sync. This condition can lead to data inconsistencies between primary and secondary sites, impacting disaster recovery reliability. Check network connectivity, daemon status, and mirroring configuration to ensure replication is functioning correctly.
CephRBDMirrorImagesNotInSyncVeryHigh	Critical	This alert is triggered when more than 10% of RBD mirror images are out of sync between the local and remote clusters. This may indicate serious replication delays or connectivity issues that could compromise data consistency across sites. Investigate network performance, mirror daemon health, and synchronization logs to restore replication stability.

Alert name	Severity	Description
CephRBDMirrorImageTransferBandwidthHigh	Warning	This alert is triggered when RBD replication bandwidth usage exceeds 80% for a sustained period (30 minutes). While not always an error, sustained high bandwidth may indicate multiple concurrent replications or potential network congestion. Review the replication workload and network capacity to ensure optimal performance.

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

Enabling monitoring stack

Alerts for the monitoring stack of the Red Hat Ceph Storage are viewed on the dashboard. However, enabling the monitoring stack is done through the command-line interface. The monitoring stack includes Prometheus, Alertmanager, and Grafana. You can use the Prometheus and Alertmanager API to manage alerts and silences.

Prerequisites

- 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
```

Configuring a Grafana certificate

The `cephadm` deploys Grafana that uses the certificate that is 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.

Prerequisites

- A running Red Hat Ceph Storage.

Procedure

1. Log in to 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, run the following commands.

- 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 before 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
```

Reference

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

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.

Prerequisites

- 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/3PYDB0sIofE.3PYDB0sIofE",
"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/3PYDB0sIofE"
},
"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"

```

Viewing alerts

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

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A running simple mail transfer protocol (SMTP) configured.
 - An alert fired.
1. From the dashboard navigation, go to **Observability > Alerts**.
 2. View the active Prometheus alerts from the **Active Alerts** tab.
 3. View all alerts from the **Alerts** tab.
 4. To view the source of an alert, click on its row and then click **Source**.

Creating a silence

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

PREREQUISITE

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

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

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

Figure: Creating a silence

Observability > Alerts > Silences > Create Silence

Create Silence

Creator *

admin

✓

Comment *

creating a silence

✓

Start time * ⓘ

2024-05-13 08:36

✓

Duration *

2h

✓

End time *

2024-05-13 10:36

✓

Matchers *

alertname

=

CephadmPaused

✎

🗑

+

Add matcher

Matches 1 rule with no active alerts.

Cancel

Create Silence

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

Recreating a silence

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

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - An alert fired.
 - A silence created for the alert.
1. From the dashboard navigation, go to **Observability > Alerts**.
 2. On the **Silences** tab, select the row with the alert that you want to recreate, and click **Recreate** from the action drop-down.

3. Edit any needed details, and click **Recreate Silence** button.
A notification displays indicating that the silence was edited successfully and the status of the silence is now active.

Editing a silence

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

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - An alert fired.
 - A silence created for the alert.
1. From the dashboard navigation, go to **Observability > Alerts**.
 2. On the **Silences** tab, select the row with the alert that you want to edit, and click **Edit** from the action drop-down.

Note: To edit a matcher, click the **Edit** icon.

A notification displays indicating that the silence was edited successfully.

Expiring a silence

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

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - An alert fired.
 - A silence created for the alert.
1. From the dashboard navigation, go to **Observability > Alerts**.
 2. On the **Silences** tab, select the row with the alert that you want to expire, and click **Expire** from the action drop-down.
 3. In the **Expire Silence** notification, select **Yes, I am sure** and click **Expire Silence**.
A notification displays indicating that the silence was expired successfully and the Status of the alert is **expired**, in the Silences table.

Managing Ceph Object Gateway

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.

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
```

4. 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
```

5. Optional: If the Ceph 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
```

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 either through the command-line interface or by using the dashboard. Use this information to learn how to create the Ceph Object Gateway service by using the dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- SSL key from Certificate Authority (CA).

For more information, see [Configuring SSL for Beast](#).

PREREQUISITE

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.

1. From the dashboard navigation, go to **Administration > Services**.
2. Click **Create**.
3. Fill in the **Create Service** form.
 - a. Select **rgw** from the **Type** service list.
 - b. Enter the ID that is used in `service_id`.
 - c. Select **SSL**.
 - d. Click **Choose File** and upload the SSL certification .pem format.

Figure: Creating a Ceph Object Gateway services with SSL

Create Service

×

Type *

rgw

▼

Id ⓘ *

rgw.ssl

✓

Realm

▼

Zone Group

default

▼

Zone

default

▼

☐ Unmanaged ⓘ

Placement

Hosts

▼

Hosts

✎

ceph-vpap-doc-s3lpr0-node1-installer

✕

Count ⓘ

Port

443

✓

☒ SSL

Certificate ⓘ

```

-----BEGIN RSA PRIVATE KEY-----
MIICXAIBAAKBgQCSws3GV/yiiuwf519LHMrDgC25Z1zbEWqh3
XhAmxYPFgr64x1Y01JL+jumespJz/+swjyzdpbIX7awp3912k

```

Cancel

Create Service

4. Click **Create Service**.
5. Check that the Ceph Object Gateway service is up and running.

Managing Ceph Object Gateway users

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.
- Ceph Object Gateway login credentials are added to the dashboard.

Creating Ceph Object Gateway users

You can create Ceph Object Gateway users on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Ceph Object Gateway is installed.
- Ceph Object Gateway login credentials are added to the dashboard.

For more information, see the following:

- [“Manually adding Ceph Object Gateway login credentials to the dashboard” on page 66.](#)
 - [Ceph Object Gateway.](#)
1. From the dashboard navigation, go to **Object > Users**.
 2. On the **Users** tab, click **Create**.
 3. Set the parameters in the **Create User** form.
 - a. Enter the **User ID** and **Full name**
 - b. If required, edit the maximum number of buckets.
 - c. Fill in an **Email address**.
 - d. Select if the user is **Suspended** or a **System user**.
 - e. In the **S3 key** section, set a custom access key and secret key by clearing the **Auto-generate key** selection.
 - f. In the **User quota** section, select the user quota.
The options for the user quota can be **Enabled**, **Unlimited size**, or has **Unlimited objects**. If there is a limited size enter the maximum size. If there are limited objects, enter the maximum objects.
 - g. In the **Bucket quota** section, select the bucket quota.
The options for the bucket quota can be **Enabled**, **Unlimited size**, or has **Unlimited objects**. If there is a limited size enter the maximum size. If there are limited objects, enter the maximum objects.
 4. Click **Create User**.
A notification displays that the user was created successfully.
Figure: Create user

Create User

User ID *

rgw-test-11

✓

☐ Show Tenant

Full name *

rgw-test

✓

Email address

rgw-test@mail.com

✓

Max. buckets

Custom

▼

1000

☐ Suspended

S3 key

☒ Auto-generate key

User quota

☐ Enabled

Bucket quota

☐ Enabled

Cancel

Create User

Creating Ceph Object Gateway subusers

A subuser is associated with a user of the S3 interface. You can create a subuser 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.
- Ceph Object Gateway login credentials are added to the dashboard.
- Ceph Object Gateway user is created.

Procedure

1. From the dashboard navigation, go to **Object > Users**.
2. On the **Users** tab, select a user and click **Edit**.
3. In the **Edit User** form, from the **Subusers** section, click **Create Subuser**.

4. In the **Create Subuser** dialog, enter the **Subuser** and select the appropriate permissions.
5. Select **Auto-generate secret** and click **Create Subuser**.

Note: Selecting **Auto-generate-secret** automatically generates the Ceph Object Gateway secret key.

Figure: Create Subuser form

Create Subuser [X]

Username: rgw-test-11

Subuser*: rgw-test-sub-user ✓

Permission*: read, write ✓ ▾

Swift key

☒ Auto-generate secret

Cancel Create Subuser

6. In the **Edit User** form, save your changes by clicking **Edit user**.
A notification displays that the Object Gateway user was updated successfully.

Adding roles to Ceph Object Gateway users

You can add a role to a specific Ceph object gateway user on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

- Ceph Object Gateway is installed.
 - Ceph Object gateway login credentials are added to the dashboard.
 - Ceph Object gateway user is created.
1. From the dashboard navigation, go to **Object > Users** and click on the **Roles** tab.
 2. Click **Create**.
 3. In the **Create Role** form, configure **Role name**, **Path**, and **Assume Role Policy Document**.

Figure: Create Role form

Create Role

Role name *

s3access

Path *

/

Assume Role Policy Document * ?

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::user/kvm"
        ]
      },
      "Action": [
        "sts:AssumeRole",
        "s3:PutObject"
      ]
    }
  ]
}
```

Cancel

Create Role

- Click **Create Role**.
A notification displays that the role was created successfully.

Editing Ceph Object Gateway users on the dashboard

Edit Ceph Object Gateway users on the Red Hat Ceph Storage, after the credentials are set-up using the CLI.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - The Ceph Object Gateway is installed.
 - Ceph Object Gateway login credentials are added to the dashboard.
 - A Ceph Object Gateway user is created.
- From the dashboard navigation, go to **Object > Users**.
 - From the **Users** tab, select the user row and click **Edit**.
 - In the **Edit User** form, edit the required parameters and click **Edit User**.
A notification displays that the user was updated successfully.

Related information

[Manually adding Ceph Object Gateway login credentials to the dashboard](#)

Deleting Ceph Object Gateway users

Delete Ceph Object Gateway users on the Red Hat Ceph Storage, after 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.

- Ceph Object Gateway login credentials are added to the dashboard.
- A Ceph Object Gateway user is created.

Procedure

1. From the dashboard navigation, go to **Object > Users**.
2. Select the **Username** to delete, and click **Delete** from the action drop-down.
3. In the **Delete user** notification, select **Yes, I am sure** and click **Delete User**.
The user is removed from the **Users** table.

Reference

For more information, see the following:

- [“Manually adding Ceph Object Gateway login credentials to the dashboard” on page 66.](#)

Managing user accounts

Accounts provide authenticated identities for accessing Object Gateway resources. A user can operate independently or belong to an account, where it inherits shared ownership, quotas, and centrally managed policies.

A user represents an identity within the Object Gateway. It can exist on its own, or it can be linked to an account for shared resource ownership and unified policy management. When a user operates independently (traditional Ceph Object Gateway user), they own their buckets and objects directly. When linked to an account, the ownership shifts, all buckets and objects created by the user are owned by the account rather than by the individual.

Linking a user to an account provides the following benefits:

Aggregated quotas and usage

Resource consumption is tracked and enforced at the account level.

Shared visibility

Users and roles within the same account can view and interact with account-owned buckets.

This structure helps create consistent identity boundaries and improves multitenancy governance.

Managing user accounts allows administrators to:

- Control who can access Object Gateway resources
- Maintain consistency across all identities under an account
- Track usage and enforce quotas
- Support multitenant application environments

Roles and responsibilities for user accounts

Different types of user identities exist under an account.

[“Roles and responsibilities of users linked to an account” on page 73](#) lists these user types and their permissions.

<i>Roles and responsibilities of users linked to an account</i>		
User access	Description	Permission
User with root access	user has default permission owned by the account.	Full access to all account resources.
User without root access	user can select the required permissions.	Require explicit managed policies (S3 read-only and S3 full access).

Creating an account

Create an account to define a managed identity boundary and configure quotas, limits, and access controls for its users and buckets.

PREREQUISITE

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

- You must have administrator access to the Red Hat Ceph Storage Dashboard.
 - Ensure the account email address is unique.
 - Plan the quotas, resource modes, and identity limits for the account.
1. Go to **Object > User management > Accounts > Create**.
 2. Enter the required account information.
 3. Configure resource modes and limits, including users, roles, groups, and access keys.
 4. Optional: Set account-level quota.
 5. Optional: Set bucket-level quota.
 6. Click **Create Account**.

Result

The account is created and appears in the **Accounts** list. The configured quotas, limits, and access settings are applied to the account, and the account is available for linking with user.

AFTER COMPLETING THE TASK

After creating an account, the next step is to create an account user.

Editing a user account

Update a user account's attributes, permissions, or access settings.

PREREQUISITE

Before you begin, make sure that you have administrator permission to edit user accounts.

Note: Account-linked users cannot be removed at the account level. To modify their attributes or policies, use the **User** page.

1. Go to **Object > User management > Accounts**.
2. In the accounts list, open the **⋮** (three-dot) menu for the account that you want to modify.
3. Select **Edit** and update the required identity details, limits, or managed policies.
4. Save the changes.

Result

The user account is updated with the new configuration. The changes take effect immediately and appear in the **User Details** panel.

Deleting a user account

PREREQUISITE

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

- Ensure that the user does not own any buckets.
- You must have administrator permission to delete user accounts.

Note: An Account cannot be deleted until all user linked the accounts are deleted.

1. Go to **Object > User management > Accounts**.
2. In the accounts list, open the **⋮ (three-dot)** menu for the account that you want to remove.
3. Select **Delete** and confirm the action.

Result

The user account is deleted from the system. All associated credentials are deleted, and the user no longer appears in the Users list.

Linking a user to an account

Associate a user with an account so they inherit the account's quotas, policies, and resource ownership. Account membership is permanent. When a user is linked to an account, all existing buckets owned by that user are automatically reassigned to the account.

PREREQUISITE

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

- At least one account must exist.
 - Determine whether the user requires root-level access within the account
 - Ensure the appropriate managed policies are available if the user needs predefined access permissions.
1. Go to **Object > Users**.
 2. Associate a user with an account.
Choose one of the following ways to associate a user with an account:
 - Create a new user and link it to an account during creation
 - Edit an existing user and link it to an account.
 3. In the **Link Account** field, select the account to associate with the user.
 4. Optional: Select **Account Root User** if the user requires administrative privileges to manage other users or roles.
 5. Optional: Apply **managed policies** if the user needs predefined full-access or read-only permissions.
 6. Save the user configuration.

Result

The user is linked to the selected account. The user now inherits the account's quotas, ownership rules, and applicable managed policies. If configured as a root account user, the user receives elevated management capabilities, and this status is indicated in the Users list.

Link bucket with an account

You can link a bucket to an account so the bucket becomes part of the account's ownership scope and follows the account's quota, visibility, and policy rules.

PREREQUISITE

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

- The user must have sufficient permissions, through a managed policy, to create or modify bucket ownership.

1. Choose one of the following methods to associate a bucket with an account:
 - Create a new bucket under a user who belongs to an account, and specify the account user as the bucket owner.
 - Edit an existing bucket created by a normal Ceph Object Gateway user and update its ownership to an account user.
2. In the bucket creation or edit form, select the account user from the owner list.
3. Save the bucket configuration.

Result

The bucket is linked to the account and appears in the bucket list with the updated ownership. Account users can view and manage the bucket according to their assigned permissions, and the bucket now follows the account's quota and policy settings.

Managing Ceph Object Gateway buckets

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.
- Ceph Object Gateway login credentials are added to the dashboard.

Creating Ceph Object Gateway buckets

Use the dashboard to create Ceph Object Gateway buckets on the Red Hat Ceph Storage once the credentials are set up using the CLI.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - The Ceph Object Gateway is installed.
 - Ceph Object Gateway login credentials are added to the dashboard.
 - Ceph Object Gateway user is created and not suspended.
1. From the dashboard navigation, go to **Object > Buckets**.
 2. Click **Create**.
The **Create Bucket** displays.
 3. Fill in the **Create Bucket** form.
 4. Enter a **Name** for the bucket.
 5. Select a **Owner**. The owner is a user that is not suspended.
 6. Select a **Placement target**.

Important: A bucket's placement target cannot be changed after creation.

7. Optional: In the **Locking** section, select **Enabled** to enable locking for the bucket objects.

Important: Locking can only be enabled while creating a bucket and cannot be changed after creation.

- a. Select the **Mode**, either **Compliance** or **Governance**.
 - b. In the **Days** field, select the default retention period that is applied to new objects placed in this bucket.
8. Optional: In the **Security** section, select **Security** to encrypt objects in the bucket.
- a. Set the configuration values for SSE-S3. Click the **Encryption** information icon and then **Click here**.

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

- b. In the **Update RGW Encryption Configurations** dialog, ensure that **SSE-S3** is selected as the **Encryption Type**.
 - c. Fill in the other required information.
 - d. Click **Submit**.
- Figure: Encrypt objects in the bucket*

Update RGW Encryption Configurations

Encryption Type:

☒ SSE-S3 Encryption
 ☐ SSE-KMS Encryption

Key management service provider *

vault

Authentication Method *

agent

Secret Engine *

transit

Secret Path

/v1/transit

Namespace

Vault Address *

Role *

CA Certificate ?

Choose file

Client Certificate ?

Choose file

Client Private Key ?

Choose file

Cancel

Submit

9. In the **Tags** section, click **Add** to add bucket tags.
These tags are equivalent to the S3 PutBucketTagging. Enter the tag **Key** and tag **Value** to categorize your storage buckets.
10. Set the bucket policies in the **Policies** section.
 - a. Enter the **Bucket policy**.
Use the **Policy generator** or **Policy examples** buttons to help create the bucket policies, as needed.
Enter or modify the policy in JSON format.

Use the following links from within the form to help create your bucket policy. These links open a new tab in your browser.

- **Policy generator** is an external tool from AWS to generate a bucket policy. For more information, see [AWS Policy Generator](#).

Note: You can use the policy generator with the **S3 Bucket Policy** type as a guideline for building your Ceph Object Gateway bucket policies.

- **Policy examples** takes you to AWS documentation with examples of bucket policies.

For more information about managing bucket policies through the dashboard, see [“Managing Ceph Object Gateway bucket policies” on page 86](#).

- Set the Access Control Lists (ACL) grantee and permission information.

ACL user operations		
Permission	Bucket	Object
READ	Grantee can list the objects in the bucket.	Grantee can read the object.
WRITE	Grantee can write or delete objects in the bucket.	N/A
FULL_CONTROL	Grantee has full permissions for object in the bucket.	Grantee can read or write to the object ACL.

- Click **Create Bucket**.
A notification displays that the bucket was created successfully.

Editing Ceph Object Gateway buckets

You can edit Ceph Object Gateway buckets on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - The Ceph Object Gateway is installed.
 - Ceph Object Gateway login credentials are added to the dashboard.
 - Ceph Object Gateway user is created and not suspended.
 - A Ceph Object Gateway bucket created.
- From the dashboard navigation, go to **Object > Buckets**.
 - Select the bucket row that needs to be updated, and click **Edit**.
The **Edit Bucket** displays.
 - Optional: Enable **Versioning** to enable a versioning state for all 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.
 - Optional: Select **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.

5. Optional: As needed, update the **Tags** and **Policies**.
Updating the **Policies** includes updating the **Bucket policy** and Access Control Lists (ACL) grantee and permission information. For more information, see steps “9” on page 78 and “10” on page 78 in “Creating Ceph Object Gateway buckets” on page 76.
6. Click **Edit Bucket** to save the changes.
A notification displays that the bucket was updated successfully.

Deleting Ceph Object Gateway buckets

You can delete Ceph Object Gateway buckets on the Red Hat Ceph Storage once the credentials are set-up using the CLI.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - The Ceph Object Gateway is installed.
 - Ceph Object Gateway login credentials are added to the dashboard.
 - Ceph Object Gateway user is created and not suspended.
 - A Ceph Object Gateway bucket created.
1. From the dashboard navigation, go to **Object > Buckets**.
 2. Select the bucket to be deleted, and click **Delete** from the action drop-down.
 3. In the **Delete bucket** notification, select **Yes, I am sure** and click **Delete bucket**.
The bucket is deleted from the **Buckets** table.

Figure: Deleting a Ceph Object Gateway bucket

Delete bucket

×

Are you sure that you want to delete **rgw-test**?

☒ **Yes, I am sure.**

Cancel

Delete bucket

Managing Ceph Object Gateway notification destinations

Learn how notification destinations in Red Hat Ceph Storage Object Gateway define where S3 bucket event messages are delivered.

A notification destination represents an external endpoint that receives S3 bucket event messages. These destinations define *how* and *where* events are sent—for example to HTTPS endpoints, Kafka topics, AMQP queues, or custom services. Notification destinations are created once and can be reused across multiple bucket notifications.

Why notification destinations are important

Typical consumers include:

- Stream processing services
- Data ingestion pipelines
- Automation tools
- Audit and compliance systems
- Analytics applications

Key components of a notification destination

A notification destination includes the following components:

Destination name

A unique identifier

Endpoint type

HTTPS, Kafka, AMQP, or other supported mechanisms

Endpoint details

URL, port, credentials, TLS settings, timeout, and retry configuration

Delivery format

How event data is packaged and transmitted

Associations

Bucket notifications that reference the destination

How notification destinations fit into the event architecture

The following sequence explains how destinations participate in event routing:

1. The administrator creates one or more notification destinations.
2. Bucket notifications reference these destinations.
3. The Object Gateway monitors bucket events.
4. When an event occurs, the gateway identifies the mapped destination.
5. The event is delivered to the configured endpoint.

Use cases

Notification destinations support scenarios such as:

- Sending new-object events to Kafka
- Triggering workflows through HTTPS webhooks
- Forwarding delete events to audit systems
- Integrating with AMQP or message queues
- Driving real-time data ingestion

Creating a notification destination

Create a new notification destination to receive S3 event messages.

PREREQUISITE

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

- You have administrative permissions.

- You have the endpoint information, including URL, protocol, and authentication details.
1. Go to **Object Gateway > Bucket notifications** and select the **Destinations** tab.
 2. Click **Create destination**.
 3. Enter the destination details, including the name, endpoint type, and required configuration fields.
 4. Click **Save**.
 5. Verify that the destination appears in the list.

Result

A new notification destination is created and available for bucket notification mapping.

Editing a notification destination

Modify the configuration of an existing notification destination.

PREREQUISITE

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

- The destination is not used by critical or production bucket notifications.
 - You have the updated endpoint or configuration details.
1. Go to **Object Gateway > Bucket notifications** and select the **Destinations** tab.
 2. Select the destination and click **Edit**.
 3. Update the required fields.
 4. Click **Save**.
 5. Verify that the updated details appear in the list.

Result

The notification destination is updated..

Deleting a notification destination

Delete a notification destination that is no longer required.

PREREQUISITE

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

- No active bucket notifications reference the destination.
 - You have confirmed that deleting the destination will not impact downstream workflows.
1. Go to **Object Gateway > Bucket notifications** and select the **Destinations** tab.
 2. Click the **Destinations** tab.
 3. Select the destination you want to delete.
 4. Click **Delete**.
 5. Confirm the deletion when prompted.
 6. Verify that the destination is removed from the list.

Result

The selected notification destination is deleted.

Listing notification destinations

PREREQUISITE

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

- The Dashboard service is running.
 - You are signed in with a user ID that has permission to view notification destinations.
1. In the Dashboard navigation pane, select **Object Gateway > Bucket notifications** and select the **Destinations** tab
 2. Review the list of configured notification destinations.
 3. Optional. Use the search box or filters to refine the list.

Result

All available notification destinations are displayed..

Managing bucket notifications in Ceph Object Gateway

Learn how bucket notifications map S3 events to configured topics to enable event-driven workflows.

Bucket notifications define which S3 events generate messages and where those messages should be delivered. Each rule includes the bucket, event types, and a target topic. When the specified events occur, the Object Gateway sends structured event records to the mapped topics. Bucket notifications provide the dynamic part of the notification framework: topics define the destination, and notifications define the triggering conditions.

Bucket notifications have the following characteristics:

- Multiple notifications per bucket: Each bucket can have multiple event mappings.
- Reusability: A topic can be referenced by multiple notifications.
- Granularity: Notifications can be based on specific events or object paths.
- Dependency: A topic must exist before creating a bucket notification.

How bucket notifications work

The following workflow describes how bucket notifications function:

1. An administrator configures the notification in the Dashboard.
2. The notification references a previously created topic.
3. The Object Gateway tracks events on the bucket.
4. When an event occurs, the gateway generates an event record.
5. The gateway delivers the event to the topic endpoint.

Use cases

Bucket notifications support several real-world use cases, including:

- Triggering pipelines when new objects are uploaded.
- Synchronizing metadata to external systems.
- Automating cleanup or archival workflows
- Generating alerts when objects are deleted
- Integrating bucket events with third-party analytics tools

Creating a bucket notification

Create a bucket notification rule that maps S3 events to a notification destination.

PREREQUISITE

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

- At least one notification destination exists.
 - You know the bucket name and the required event types.
1. In the **Notifications** tab, click **Create notification**.
 2. Enter the bucket name.
 3. Select the event types to monitor.
 4. Select the target notification destination.
 5. Click **Save**.

Result

A new bucket notification is created and associated with the selected destination.

Editing a bucket notification

Modify an existing bucket notification rule.

PREREQUISITE

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

- You have the updated event types or destination information.
1. Open the **Notifications** tab.
 2. Select the notification and click **Edit**.
 3. Update the required fields.
 4. Click **Save**.
 5. Verify the updates in the table.

Result

The bucket notification is updated with the new configuration..

Deleting a bucket notification

Delete a bucket notification that is no longer required.

PREREQUISITE

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

- The notification is not required by production services.
1. Go to the **Notifications** tab.
 2. Select the bucket notification.
 3. Click **Delete**.
 4. Confirm the deletion.

Result

The selected bucket notification is deleted..

Listing bucket notifications

View the list of bucket notifications configured for an Object Gateway bucket.

PREREQUISITE

Before you begin, make sure that you have permission to view bucket notifications.

1. Select **Object Gateway** > **Bucket notifications**.
2. Click the **Notifications** tab.
3. Review the table of existing bucket notifications.
4. Optional: Apply filters or search.

Result

All configured bucket notifications are displayed.

Viewing Performance Metrics for Multi-Site Sync Deltas in Ceph Object Gateway

You can view the replication performance of each shard between zones in the Grafana dashboard.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.
- ceph-exporter, Grafana, and Prometheus are installed.
- Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- Object gateway users and buckets are created and not suspended.

1. Log in to the Grafana URL.


```
https://<ip-where-grafana-deployed>:3000
```

Example

```
https://<ip-where-grafana-deployed>:3000
```

2. Go to the **Dashboard** tab and search for **RGW Sync Overview**.
3. To view the performance metrics of the multi-site sync deltas, see the **Replication(Time) Delta per shard** panel

Figure: Replication (Time) Delta per shard



The panel displays the time taken for each shard in the cluster to replicate.

Viewing Ceph Object Gateway per-user and per-bucket performance counters

You can view the Ceph Object Gateway performance counters per user per bucket in the Grafana dashboard.


1. Log in to the Grafana URL.

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


Example

`https://dashboard_url:3000`

- Go to the **Dashboard** tab and search for **RGW S3 Analytics**.
- To view per-bucket Ceph Object gateway operations, select the 'Bucket' window:
*Figure: **Bucket operations counter** form*



- To view per-user Ceph Object gateway operations, select the 'User' window:
*Figure: **User operations counter** form*



Note: The output of per-bucket and per-user get operation count command increases by two for each get operation that is run from client: s3cmd. This behaviour is a known issue.

Managing Ceph Object Gateway bucket policies

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and manage Ceph Object Gateway bucket policies.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- At least one Ceph Object Gateway user is created.
- Ceph Object Gateway login credentials are added to the dashboard.
- At least one Ceph Object Gateway bucket. For information about creating a bucket, see [“Creating Ceph Object Gateway buckets” on page 76](#).

Creating and editing Ceph Object Gateway bucket policies

Use the dashboard to create and edit a Ceph Object Gateway bucket policy on the Red Hat Ceph Storage dashboard.

PREREQUISITE

Before creating or editing Ceph Object Gateway bucket policies, ensure the following are available:

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - The Ceph Object Gateway is installed.
 - At least one Ceph Object Gateway user is created.
 - Ceph Object Gateway login credentials are added to the dashboard.
 - At least one Ceph Object Gateway bucket. For information about creating a bucket, see [“Creating Ceph Object Gateway buckets” on page 76](#).
1. From the dashboard, go to **Object > Buckets**.
 2. Create or modify a bucket policy for an existing bucket.

Note: To create a bucket policy during bucket creation, click **Create** and fill in the bucket policy information in the **Policies** section of the **Create Bucket** form.

Select the bucket for which the bucket policy will be created or modified, and then click **Edit**.

3. In the **Edit Bucket** form, go to **Policies**.
4. Enter or modify the policy in JSON format.
Use the following links from within the form to help create your bucket policy. These links open a new tab in your browser.
 - **Policy generator** is an external tool from AWS to generate a bucket policy. For more information, see [AWS Policy Generator](#).

Note: You can use the policy generator with the **S3 Bucket Policy** type as a guideline for building your Ceph Object Gateway bucket policies.

- **Policy examples** takes you to AWS documentation with examples of bucket policies.
5. To save the bucket policy, click **Edit Bucket**.

Note: When creating a bucket policy during an initial bucket creation, click **Create Bucket**.

When the bucket policy is saved, the Updated Object Gateway bucket `bucketname` notification is displayed.

Deleting Ceph Object Gateway bucket policies

Use the dashboard to delete Ceph Object Gateway bucket policies on the Red Hat Ceph Storage dashboard.

PREREQUISITE

Before deleting Ceph Object Gateway bucket policies, ensure the following are available:

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - The Ceph Object Gateway is installed.
 - At least one Ceph Object Gateway user is created.
 - Ceph Object Gateway login credentials are added to the dashboard.
 - At least one Ceph Object Gateway bucket. For information about creating a bucket, see [“Creating Ceph Object Gateway buckets”](#) on page 76.
1. From the dashboard, go to **Object > Buckets**.
 2. Select the bucket for which the bucket policy will be deleted, and then click **Edit**.
 3. In the **Edit Bucket** form, go to **Policies**.
 4. Click **Clear**.
 5. To complete the bucket policy deletion, click **Edit Bucket**.
When the bucket policy is deleted, the Updated Object Gateway bucket `bucketname` notification is displayed.

Managing S3 bucket lifecycle policies

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and manage Ceph Object Gateway bucket lifecycle policies.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- At least one Ceph Object Gateway user is created.
- Ceph Object Gateway login credentials are added to the dashboard.
- At least one Ceph Object Gateway bucket. For information about creating a bucket, see [“Creating Ceph Object Gateway buckets”](#) on page 76.

Applying and viewing bucket lifecycle policies

Use the dashboard to apply and manage a Ceph Object Gateway bucket lifecycle policy on the Red Hat Ceph Storage dashboard.

1. From the dashboard, go to **Object > Buckets**.
2. Select the bucket for which the lifecycle policy needs to be applied and click **Edit**

Note: Bucket lifecycle policy cannot be applied during the creation of the bucket. They can be applied only after a bucket is created.

3. In the **Edit Bucket** form, go to **Policies** and apply the lifecycle rule in the **Lifecycle** field in JSON format.

4. To save the bucket lifecycle policy, click **Edit Bucket**.

Figure: Apply bucket LC rule

Enable ☐

Enables encryption for the objects in the bucket. To enable encryption on a bucket you need to set the configuration values for SSE-S3 or SSE-KMS. To set the configuration values [Click here](#)

Replication

Enable ☐

Enables replication for the objects in the bucket.

Tags

Tagging provides a way to categorize storage

+ Add tag

Policies

Bucket policy

{ }

✕ Clear

Policy examples

Policy generator

Lifecycle ?

{
 "LifecycleConfiguration": {
 "Rule": {
 "ID": "Delete lc",
 "Prefix": "",
 "Status": "Enabled",
 "Expiration": {
 "Days": "1"
 }
 }
 }
}

✕ Clear

Policy examples

ACL ?

Grantee ?

Owner

▼

Permissions ?

Full control

▼

Cancel

Edit Bucket

5. After the bucket lifecycle policy is applied, it can viewed in the bucket listing screen by expanding the relevant bucket entry.

Figure: View bucket LC rule

Object » Buckets

	Name	Owner	Used Capacity	Capacity Limit %	Objects	Object Limit %
☐	> test-ck1	ckulal	0 B	No Limit	0	No Limit
☐	> test-ck2	ckulal	0 B	No Limit	0	No Limit
☐	> testbkt3	test1	0 B	No Limit	0	No Limit
☐	▼ testbkt4	ckulal	0 B	No Limit	0	No Limit

Details
Policies

Bucket policy

{}

Lifecycle

JSON XML

```
{
  "LifecycleConfiguration": {
    "Rule": {
      "ID": "Rule 1",
      "Prefix": "",
      "Status": "Enabled",
      "Expiration": {
        "Date": "2022-04-05"
      }
    }
  }
}
```

Replication policy

```
{
  "Role": ""
}
```

ACL

Grantee	Permissions
Bucket Owner	FULL_CONTROL
Everyone	-
Authenticated users group	-

Deleting bucket lifecycle policies

Use the dashboard to delete Ceph Object Gateway bucket lifecycle policies on the Red Hat Ceph Storage dashboard.

1. From the dashboard, go to **Object > Buckets**.
2. Select the bucket for which the lifecycle policy needs to be deleted, and click **Edit**.
3. In the **Edit Bucket** form, go to **Policies**.
4. Click **Clear**.
5. To complete the bucket policy deletion, click **Edit Bucket**.

Creating a tiering storage class

Use the Dashboard to create a tiering storage class for local storage, Cloud S3, or Cloud S3 Glacier targets.

PREREQUISITE

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

- You have administrator privileges to access the Ceph Dashboard.
- The Object Gateway service is running.
- You have the access key and secret key for the target storage system.

1. Go to **Dashboard > Object > Tiering**.
2. Click **Create**.
3. In the **Create Tiering Storage** dialog, select the storage type: Local, Cloud S3, or Cloud S3 Glacier.
4. Enter the required details:

Name

Unique name for the storage class.

Target Region

Region of the storage endpoint.

Target Endpoint

URL endpoint of the remote cloud object storage service where Ceph transitions objects. Use this parameter to define the destination endpoint for data movement to an external S3-compatible cloud, such as AWS S3 or S3 Glacier. The value must include the protocol, host, and optional port number.

Access key and Secret key

Credentials for the target storage.

Target Path

Target path and bucket name in the target storage.

5. Specify **Read Through Restore Days**.
6. Select one of the restore options:
 - **Allow read through (default restore behavior)**
 - **Head object (Stub file)**
7. Select the **Restore Storage Class** from the dropdown list.
8. Select the **Glacier Restore Type**.
9. Expand the **Advanced** pane and configure the optional parameters:

Multipart Sync Threshold

object size that triggers multipart upload.

Multipart Minimum Part Size

Minimum size of each part.

Placement Target

Placement configuration that defines where transitioned objects are stored within a zonegroup. The placement target maps a storage class to a specific storage tier or cloud destination, ensuring objects are directed to the correct backend during transition.

ACL Mapping

Defines how access control list (ACL) entries from the source Ceph zone are translated to the destination cloud service. Use this parameter to maintain user or group permissions when objects transition to a cloud tier.

10. Click **Create** to save the tiering storage class.

Result

The new tiering storage class appears under **Dashboard > Object > Tiering**.

Managing buckets of a multi-site object configuration

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.

For more information about multi-site and Ceph Object Gateway buckets, see:

- [Multi-site configuration and administration](#)
- [“Manually adding Ceph Object Gateway login credentials to the dashboard” on page 66](#)
- [“Managing Ceph Object Gateway buckets” on page 76](#)
- [“User roles and permissions” on page 30](#)

Prerequisites

- At least one running Red Hat Ceph Storage cluster deployed on both the sites.
- Dashboard is installed.
- The multi-site Ceph Object Gateway is configured on the primary and secondary sites.

- Ceph Object Gateway login credentials of the primary and secondary sites are added to the dashboard.
- Ceph Object Gateway users are created on the primary site.
- Ceph Object Gateway buckets are created on the primary site.
- At least `rgw-manager` level of access on the Ceph dashboard.

Monitoring buckets of a multi-site object

Monitor the multi-site sync status of a bucket on the dashboard. You can view the source zones and sync status from the **Multi-site sync status** card on **Object > Overview** page.

The multi-site sync status is divided into two sections.

Primary Source Zone

Displays the default realm, zonegroup, and the zone the Ceph Object Gateway is connected to.

Source Zones

View both the metadata sync status and data sync information progress. When you click the status, a breakdown of the shard syncing is displayed.

The sync status shows the **Last Synced** time stamp with the relative time of the last sync occurrence in relation to the current time. When the sync is complete, this shows as **Up to Date**.

When a sync is not caught up the status shows as Syncing. However, the Last sync shows the number of days the sync is not caught up. By clicking Syncing, it displays the details about shards which are not synced.

Editing buckets of a multi-site Ceph Object Gateway configuration

You can edit and update the details of the buckets of one zone in another zone on the Red Hat Ceph Storage Dashboard in a multi-site Ceph Object Gateway configuration. You can edit the owner, versioning, multi-factor authentication and locking features of the buckets with this feature of the dashboard.

PREREQUISITE

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

- Minimum of one running Red Hat Ceph Storage cluster deployed on both sites.
 - Dashboard is installed.
 - The multi-site Ceph Object Gateway is configured on the primary and secondary sites.
 - Ceph Object Gateway login credentials of the primary and secondary sites are added to the dashboard.
 - Ceph Object Gateway users are created on the primary site.
 - Ceph Object Gateway buckets are created on the primary site.
 - Minimum of `rgw-manager` level of access on the Ceph dashboard.
1. From the dashboard navigation of the secondary site, go to **Object > Buckets**.
From the secondary landing page, view the Object Gateway buckets that were created for Object Gateway users on the primary site.
 2. Click the row of the bucket that you want to edit and click **Edit**.
 3. From the **Edit bucket** form, edit the required parameters.
 4. Click **Edit Bucket** to save the changes.
A notification displays showing that the bucket was updated successfully.

Deleting buckets of a multi-site Ceph Object Gateway configuration

You can delete buckets of secondary sites in primary sites on the Red Hat Ceph Storage Dashboard in a multi-site Ceph Object Gateway configuration.

PREREQUISITE

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

- Minimum of one running Red Hat Ceph Storage cluster deployed on both sites.
- Dashboard is installed.
- The multi-site Ceph Object Gateway is configured on the primary and secondary sites.
- Ceph Object Gateway login credentials of the primary and secondary sites are added to the dashboard.
- Ceph Object Gateway users are created on the primary site.
- Ceph Object Gateway buckets are created on the primary site.
- Minimum of `rgw-manager` level of access on the Ceph dashboard.

Important: Do not delete buckets of primary site from secondary sites.

1. From the dashboard navigation of the primary site, go to **Object > Buckets**.
2. Select the bucket of the secondary site to be deleted, and click **Delete** from the action drop-down.
3. In the **Delete Bucket** notification, select **Yes, I am sure** and click **Delete bucket**.
The bucket is deleted from the **Buckets** table.

Monitoring multi-site Ceph Object Gateway configuration

The Red Hat Ceph Storage dashboard supports monitoring the users and buckets of one zone in another zone in a multi-site Ceph 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.

PREREQUISITE

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

- Minimum of one running Red Hat Ceph Storage cluster deployed on both sites.
 - Dashboard is installed.
 - The multi-site Ceph Object Gateway is configured on the primary and secondary sites.
 - Ceph Object Gateway login credentials of the primary and secondary sites are added to the dashboard.
 - Ceph Object Gateway users are created on the primary site.
 - Ceph Object Gateway buckets are created on the primary site.
1. From the dashboard navigation of the secondary site, go to **Object > Buckets**.
 2. View the Object Gateway buckets on the secondary landing page that were created for the Object Gateway users on the primary site.

Configuring a multi-site

Configure multi-site on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A Red Hat Ceph Storage cluster at both the sites.
 - At least one Ceph Object Gateway service installed at both the sites.
1. Enable the Ceph Object Gateway module on both the sites to enable import/export feature.
 - a. From the dashboard navigation of the primary site, go to **Object > Multi-site**.
 - b. In the **Information** box, click **Enable RGW Module**.
 2. On the primary site dashboard, create a default realm, zonegroup, and zone.
 - a. From the **Object > Multi-site**, click **Create Realm**.

- b. In the **Create Realm** dialog, provide a realm name, and select **Default**.
- c. Click **Create Realm**.
- d. Click **Create Zone Group** from the action drop-down.
- e. In the **Create Zone Group** dialog, provide a zone group name, the Ceph Object Gateway endpoints, and select **Default**.
- f. Click **Create Zonegroup**.
- g. Click **Create Zone** from the action drop-down.
- h. In the **Create Zone** dialog, provide a **Zone Name**, select **Default**, provide the Ceph Object Gateway endpoints of the primary site. For the user, provide the access and secret key of the user with system privileges.

Note: While creating a zone, provide the access key and secret key of the dashboard default user, dashboard.

3. After the zone is created, a warning is emitted to restart the Ceph Object Gateway Service.
 - a. From the **Administration > Services**.
 - b. Select the Ceph Object Gateway (rgw) service row and expand to see the service details.
 - c. From the **Daemons** table, select the host and click **Restart** from the action drop-down.
4. When you go to **Object > Gateways**, you get an error that The Object Gateway Service is not configured. This bug is a [BZ#2231072](#).
 - a. As a workaround, set the Ceph Object Gateway credentials on the command-line interface.

```
ceph dashboard set-rgw-credentials
RGW credentials configured
```

- b. Go to **Object > Gateways**, to verify that you are able to access the Ceph Object Gateway on the dashboard.
5. Create a replication user on the primary site. You can use the following two options:
 - Create user using the CLI.

```
[ceph: root@host01 /]# radosgw-admin user create --uid="uid" --display-name="displayname" --system
```

- Create user from the dashboard and modify the user from the CLI.

```
[ceph: root@host01 /]# radosgw-admin user modify --uid="uid" --system
```

6. From, **Object > Gateways**, on the **Users** table, select the user row created in the previous step.
 - a. Expand the user row.
 - b. Copy the **Access Key** and **Secret Key**.
 - c. From, **Object > Multi-site**, select the zone.
 - d. Click the edit icon and paste the access key and secret key that was copied earlier.
 - e. Click **Edit Zone**.
7. Click **Export**.
 - a. From the **Export Multi-site Realm Token** dialog box, copy the token.
 - b. Close the dialog box.
8. On the secondary site, go to **Object > Multi-site**.
9. Click **Import** to import token from the primary zone.

10. In the **Import Multi-site Token** dialog box, paste the token that was copied earlier, and provide a secondary zone name.
 - a. Under **Service** details, select the placement and the port where the new Ceph Object Gateway service is going to be created.
 - b. Click **Import**.
11. After the token is imported, you get a warning to restart the Ceph Object Gateway Service.
 - a. Go to **Administration > Services**.
 - b. Select the Ceph Object Gateway (rgw) service row.
 - c. From the **Daemons** table, select the host. Verify that the Ceph Object Gateway is configured on the secondary site.
 - d. Click **Restart** from the action drop-down.
12. You need to wait for sometime till the users are synced to the secondary site. Verify that the sync is complete with the following commands.


```
[ceph: root@host01 /]# radosgw-admin sync status
[ceph: root@host01 /]# radosgw-admin user list
```
13. When you go to **Object > Gateways**, you get an error that The Object Gateway Service is not configured. This bug is a [BZ#2231072](#).
 - a. As a workaround, set the Ceph Object Gateway credentials on the command-line interface.


```
ceph dashboard set-rgw-credentials
RGW credentials configured
```
 - b. Go to **Object > Gateways**, to verify that you are able to access the Ceph Object Gateway on the dashboard.
14. On the primary site, under an error displays on the **Multi Site Sync Status** panel of the **Object > Overview** page. This is because on the secondary zone, you can see that the endpoints are of the hostname and not the IP address. This bug is a known issue while configuring multi-site. See [BZ#2242994](#).
 - a. As a workaround, in the secondary site, go to **Object > Multi-site** and select the secondary zone.
 - b. Edit the endpoints to reflect the IP address.
 - c. Click **Edit Zone**.
15. On the primary site and secondary site dashboard, you can see the status displays on the **Multi Site Sync Status** panel of the **Object > Overview** page.

AFTER COMPLETING THE TASK

To verify, create a user on the primary site. You see that the user syncs to the secondary site.

Managing Ceph File Systems (CephFS)

As a storage administrator, you can create, edit, delete, and manage accesses of file systems on the Red Hat Ceph Storage Dashboard.

Configuring CephFS volumes

As a storage administrator, you can configure Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

Creating CephFS volumes

You can create Ceph File System (CephFS) volumes on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

Before you begin, make sure that you have a working Red Hat Ceph Storage cluster with MDS deployed.

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,mds
host07 10.0.98.23 nfs,mds
7 hosts in cluster
```

4. Click **Create Volume**.
A notification displays that the volume was created successfully.

Editing CephFS volumes

You can edit Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.
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.

Removing CephFS volumes

You can remove Ceph File System (CephFS) volumes on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.
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.

Configuring CephFS subvolume groups

As a storage administrator, you can configure Ceph File System (CephFS) subvolume groups on the Red Hat Ceph Storage dashboard.

Creating CephFS subvolume groups

Create Ceph File System (CephFS) subvolume groups on the Red Hat Ceph Storage dashboard. You can use subvolume groups to apply policies across a set of subvolumes.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

You can create subvolume groups to create subvolume on the dashboard.

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:
 - **Name:** Set the name of the subvolume group.
 - **Volume name:** Validate that the correct name of the volume is selected.
 - **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.
 - **Pool:** Set the pool of the subvolume group. By default, data_pool1_layout of the parent directory is selected.
 - **UID:** Set the UID of the subvolume group.
 - **GID:** Set the GID of the subvolume group.
 - **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.

Editing CephFS subvolume groups

Edit Ceph File System (CephFS) subvolume groups on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

You can edit the subvolume groups on the dashboard.

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 subvolume and expand the row.
3. Under the **Subvolume groups** tab, select the row 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.

Removing CephFS subvolume groups

You can remove Ceph File System (CephFS) subvolume groups on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

You can remove a subvolume group from the cluster on the dashboard.

Warning: Ensure to remove the subvolume within the subvolume group before removing the subvolume group.

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.

Configuring CephFS subvolumes

As a storage administrator, you can configure Ceph File System (CephFS) subvolumes on the Red Hat Ceph Storage dashboard.

Creating CephFS subvolume

You can create Ceph File System (CephFS) subvolume on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.

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:
 - **Name:** Set the name of the subvolume group.
 - **Subvolume name:** Set the name of the volume.
 - **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.
 - **Pool:** Set the pool of the subvolume group. By default, `data_pool_layout` of the parent directory is selected.
 - **UID:** Set the UID of the subvolume group.
 - **GID:** Set the GID of the subvolume group.
 - **Mode:** Set the permissions for the directory. By default, the mode is 755 which is `rwxr-xr-x`.
 - **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.

Editing CephFS subvolume

You can edit Ceph File System (CephFS) subvolume on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.
- A subvolume created.

You can edit the subvolumes on the dashboard.

Note: You can only edit the Size of the subvolumes on the dashboard.

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. In the **Edit Subvolume** window, edit the required parameters:
 - **Name:** Set the name of the subvolume group.
 - **Subvolume name:** Set the name of the volume.
 - **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.
 - **Pool:** Set the pool of the subvolume group. By default, data_pool_layout of the parent directory is selected.
 - **UID:** Set the UID of the subvolume group.
 - **GID:** Set the GID of the subvolume group.
 - **Mode:** Set the permissions for the directory. By default, the mode is 755 which is `rwxr-xr-x`.
5. Click **Edit Subvolume**.
A notification that the subvolume was edited successfully is displayed.

Removing CephFS subvolume

You can remove Ceph File System (CephFS) subvolume on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A working Red Hat Ceph Storage cluster with Ceph File System deployed.
 - A subvolume created.
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.

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.

Creating CephFS subvolume snapshots

As a storage administrator, you can create Ceph File System (CephFS) subvolume snapshot on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- A subvolume group with corresponding subvolumes.

You can create an immutable, point-in-time view of a Ceph File System (CephFS) by creating a snapshot.

1. From the dashboard navigation, go to **File > File Systems**.
2. Expand the CephFS where you want to create the subvolume snapshot.
If there are no file systems available, create a new file system. For more information, see [“Creating CephFS volumes” on page 95](#).
3. Go to the **Snapshots** tab. There are three columns - **Groups**, **Subvolumes** and **Create**.
4. From the **Groups** and **Subvolumes** columns, select the group and the subvolume for which you want to create a snapshot.
5. Click **Create** in the third column. The *Create snapshot* form opens.
 - **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.
 - **Volume name:** The volume name is the file system name. It is already added to the form as per your selection.
 - **Subvolume group:** The subvolume group name is already added as per your selection. Alternatively, you select a different subvolume group from the dropdown list.
 - **Subvolume:** The subvolume name is already added as per your selection. Alternatively, you can select a different subvolume from the dropdown list.
6. Click **Create Snapshot**.
A notification is displayed that the snapshot is created successfully.
7. You can verify that the snapshot is created by going to the **Snapshots** tab, select the subvolume and subvolume group for which the snapshot is created.

Deleting CephFS subvolume snapshots

As a storage administrator, you can delete Ceph File System (CephFS) subvolume snapshot on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.
- A subvolume snapshot.

You can now delete Ceph File System (CephFS) subvolume snapshot on the dashboard.

1. From the Ceph Dashboard, go to **File > File Systems**
2. Select the CephFS where you want to delete the subvolume snapshot.
3. Go to the **Snapshots** tab.
4. 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.
5. In the **Delete Subvolume** dialog, select **Yes, I am sure** and then click **Delete Subvolume**.
A notification is displayed that the subvolume snapshot is deleted successfully.

Cloning CephFS subvolume snapshots

As a storage administrator, you can clone Ceph File System (CephFS) subvolume snapshot on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.
- A subvolume snapshot.

You can now clone a Ceph File System (CephFS) subvolume snapshot on the dashboard.

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.
6. Click the **arrow** next to the **Delete** button.
7. Click **Clone**. The *Create clone* form opens.
 - **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.
 - **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.

Creating CephFS volume snapshots

As a storage administrator, you can create Ceph File System (CephFS) volume snapshot on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.
- A subvolume group with corresponding subvolumes.

You can create an immutable, point-in-time view of a Ceph File System (CephFS) volume by creating a snapshot.

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.
 - **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.

Deleting CephFS volume snapshots

As a storage administrator, you can delete Ceph File System (CephFS) volume snapshot on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.
- A volume snapshot.

You can now delete Ceph File System (CephFS) volume snapshot on the dashboard.

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. In the **Delete Snapshot** dialog, select **Yes, I am sure** and then click **Delete Snapshot**.
A notification is displayed that the snapshot is deleted successfully.

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.

Creating CephFS snapshot schedule

As a storage administrator, you can create Ceph File System (CephFS) snapshot schedules on the Red Hat Ceph Storage dashboard.

PREREQUISITE

- A running Red Hat Ceph Storage cluster.

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

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.

- a. 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.
 - b. Click **Create**.
The **Create Snapshot schedule** form opens.
 - c. Enter the directory name, start date, start time, and schedule.
 - **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.
 - **Start date:** Enter date on which you want the scheduler to start creating the snapshots. By default the current date is selected.
 - **Start time:** Enter the time at which you want the scheduler to start creating the snapshots. By default the current time is added.
 - **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.

AFTER COMPLETING THE TASK

Verify that the snapshots are created by going to the **Snapshots** tab after the first scheduled snapshot is triggered.

Editing CephFS snapshot schedule

As a storage administrator, you can edit Ceph File System (CephFS) snapshot schedules on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

Edit an existing snapshot schedule using the snapshot scheduler.

You can only edit the retention policy of the snapshot schedule. You can add another retention policy or delete the existing policy.

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:

- 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.

Deleting CephFS snapshot schedule

As a storage administrator, you can delete Ceph File System (CephFS) snapshot schedules on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

You can delete existing snapshot schedules using the snapshot scheduler.

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 delete and click **Delete** from the action drop-down.
4. On the navigation menu, click **File** → **File Systems**.
5. Select the CephFS where you want to edit snapshot schedules.
6. Click **Snapshot schedules**.
7. Select the snapshot schedule that you want to delete.
8. Click the **Edit dropdown menu**.
9. Click **Delete**. A *Delete Snapshot schedule* dialog box appears.
10. Select **Yes, I am sure** to confirm if you want to delete the schedule.
11. Click **Delete snapshot schedule**.
A notification is displayed that the snapshot schedule is deleted successfully.

Deactivating and activating CephFS snapshot schedule

As a storage administrator, you can deactivate and activate Ceph File System (CephFS) snapshot schedules on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

The snapshot schedule is activated by default, it runs according to how it is configured. Deactivating excludes the snapshot from scheduling until it is activated again.

You can deactivate and activate snapshot schedules from the CephFS dashboard using the snapshot scheduler.

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.

Managing block devices

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 that are mirrored across clusters, and set IOPS limits on an image.

Managing block device images

As a storage administrator, you can create, edit, copy, purge, and delete images using the Red Hat Ceph Storage dashboard. In addition, 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.

Creating images

You can create block device images on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

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

For more information, see the following:

- [Block devices](#)
- [Creating pools](#)

1. From the dashboard navigation, go to **Block > Images**.
2. From the **Images** tab, click **Create**.
3. Complete the **Create Images** form.
4. Optional: Click **Advanced** to set advanced parameters, such as **Striping** and **Quality of Service**.
5. Click **Create Image**.
A notification displays that the image was created successfully.

Creating namespaces

Create namespaces for the block device images on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

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

Once the namespaces are created, you can give access to the users for those namespaces.

For more information, see the [How to segregate RBD images within isolated namespaces in ceph?](#) article on the Red Hat Customer Portal.

1. From the dashboard navigation, go to **Block › Images**.
2. From the **Namespaces** tab, click **Create**.
3. In the **Create Namespace** dialog, select the pool and enter a name for the namespace.
4. Click **Create**.
A notification displays that the namespace was created successfully.

Editing images

Edit block device images on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - An image is created.
1. From the dashboard navigation, go to **Block › Images**.
 2. To edit the image, select its row and click **Edit**.
 3. In the **Edit Image** form, edit the required parameters and click **Edit Image**.
Figure: Edit Image image form

Edit Image

Name (required)

test_image1

Pool

test-pool_replicated

☒ Mirroring

Allow data to be asynchronously mirrored between two Ceph clusters



You need to set **mirror mode** in the selected pool to enable mirroring.

[Set mode](#)



Use a dedicated data pool

Use a dedicated pool to store the image data. If not selected, the image data will be stored in the same pool as the image metadata.

Size (required)

5 GiB

Advanced



Cancel

Edit Image



A notification displays that the image was updated successfully.

Copying images

Copy block device images on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

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

For more information, see the following:

- [Block devices](#)
- [“Creating pools” on page 46](#)

1. From the dashboard navigation, go to **Block > Images**.
2. From the **Images** tab, click **Copy** from the image menu.
3. In the **Copy Image** form, set the required parameters and click **Copy Image**.

A notification displays that the image was copied successfully.

Moving images to trash

Move the block device images to trash before it is deleted on the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - An image is created.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, select the image to move to trash, and click **Move to Trash** from the image menu.
 3. In the **Move an image to trash** dialog, change the **Protection expires at** field and click **Move**.
A notification displays that the image was moved to trash successfully.

Purging trash

Purge trash using the Red Hat Ceph Storage dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - An image is trashed.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Trash** tab, click **Purge Trash**.
 3. In the **Purge Trash** dialog, select the pool and click **Purge Trash**.
A notification displays that the pools in the trash were purged successfully.

Restoring images from trash

Restore the images that were trashed and has an expiry date on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - An image is trashed.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Trash** tab, select the row of the image to restore.
 3. Click **Restore** from the image menu.
 4. In the **Restore Image** dialog, enter the new name of the image and click **Restore**.
A notification displays that the image was restored successfully.

Deleting images

Delete the images from the cluster on the dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - An image is created.
1. From the dashboard navigation, go to **Block > Images**.
 2. Select the row to be deleted and click **Delete** from the action drop-down.
 3. In the **Delete RBD** notification, select **Yes, I am sure** and click **Delete RBD**.
A notification displays that the image was deleted successfully.

Deleting namespaces

Delete the namespaces of the images on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - A namespace is created in the pool.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Namespaces** tab, select the namespace and click **Delete** from the action drop-down.
 3. In the **Delete Namespace** notification, select **Yes, I am sure** and click **Delete Namespace**.
A notification displays that the namespace was deleted successfully.

Creating snapshots of images

Take snapshots of the Ceph Block Device images on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

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

For more information, see the following:

- [Creating a block device snapshot](#)
 - [“Creating pools” on page 46](#)
 - [“Creating images” on page 105](#)
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row.

3. In the **Snapshots** tab, click **Create**.
4. In the **Create RBD Snapshot** dialog, enter the snapshot name and click **Create RBD Snapshot**.
A notification displays that the snapshot was created successfully.

Renaming snapshots of images

Rename the snapshots of the Ceph Block Device images on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbt` application enabled is created.
 - An image is created.
 - A snapshot of the image is created.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row.
 3. In the **Snapshots** tab, click **Rename**.
 4. In the **Rename RBD Snapshot** dialog, enter the new name and click **Rename RBD Snapshot**.

Protecting snapshots of images

Protect the snapshots of the Ceph Block Device images on the Red Hat Ceph Storage Dashboard. Protecting snapshots of images is required when you need to clone the snapshots.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbt` application enabled is created.
 - An image is created.
 - A snapshot of the image is created.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row and click on the **Snapshots** tab.
 3. Select the snapshot to protect, and click **Protect** from the action drop-down.
The snapshot updates and the **State** changes from `Unprotected` to `Protected`.

Cloning snapshots of images

Clone the snapshots of images on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

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

- A snapshot of the image is created and protected.

For more information, see the following:

- [Protecting snapshots.](#)
 - [“Protecting snapshots of images” on page 110](#)
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row.
 3. In the **Snapshots** tab, select the snapshot to clone and click **Clone** from the action drop-down.
 4. In the **Clone RBD** form, edit the parameters, as required, and click **Clone RBD**.

Copying snapshots of images

Copy the snapshots of images on the Red Hat Ceph Storage Dashboard.

Prerequisites

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

Procedure

1. From the dashboard navigation, go to **Block > Images**.
2. To protect the snapshot of the image, from the **Images** tab, select the row and then go to the **Snapshots** tab.
3. Click **Copy** from the action drop-down.
4. In the **Copy RBD** form, enter the parameters and click **Copy RBD**.
A notification displays that the snapshot was copied successfully. You can search for the copied image in the *Images* tab.

Reference

For more information, see the following:

- [Creating pools.](#)
- [Creating images.](#)

Unprotecting snapshots of images

You can unprotect the snapshots of the Ceph Block Device images on the Red Hat Ceph Storage Dashboard. Unprotecting snapshots of images is required when you need to delete the snapshots.

PREREQUISITE

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

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

- A snapshot of the image is created and protected.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row and click on the **Snapshots** tab.
 3. Select the protected snapshot, and click **Unprotect** from the action drop-down.
The snapshot updates and the **State** changes from Protected to Unprotected.

Rolling back snapshots of images

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 run 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 preexisting state.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
 - Dashboard is installed.
 - A pool with the `rbd` application enabled is created.
 - An image is created.
 - A snapshot of the image is created.
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row and click on the **Snapshots** tab.
 3. Select the snapshot to rollback, and click **Rollback** from the snapshot menu.
 4. In the **RBD snapshot rollback** dialog, click **Rollback**.

Deleting snapshots of images

Delete the snapshots of the Ceph Block Device images on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the `rbd` application enabled is created.
- An image is created.
- A snapshot of the image is created and is unprotected.

For more information, see the following:

- [Deleting a block device snapshot](#).
 - [“Unprotecting snapshots of images” on page 111](#).
1. From the dashboard navigation, go to **Block > Images**.
 2. From the **Images** tab, expand an image row.
 3. In the **Snapshots** tab, select the snapshot to delete and click **Delete** from the action drop-down.
 4. In the **Delete RBD Snapshot** dialog, select **Yes, I am sure** and click **Delete RBD Snapshot**.
A notification displays that the snapshot was deleted successfully.

Managing mirroring functions

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.

Mirroring view

View the Block device mirroring on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

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

You can view the daemons, the site details, the pools, and the images that are configured for Block device mirroring.

- From the dashboard navigation, go to **Block > Mirroring**.
Figure: Ceph Block Device mirroring view

Block / Mirroring

Site Name
ceph-rbd1

Import Bootstrap Token Create Bootstrap Token

Daemons

Instance	ID	Hostname	Version	Health
No data to display				

Items per page: 10 0-0 of 0 item

Pools

Name	Mode	Leader	# Local	# Remote	Health
rbd	image				Warning
test-pool_replicated	disabled				Disabled
test_pool2	pool				OK

Items per page: 10 1-3 of 3 items

Images

Issues (2) Syncing (0) Ready (0)

Pool	Image	State	Issue
rbd	test_image	Unknown	
rbd	test_image3	Unknown	

Items per page: 10 1-2 of 2 items

Editing mode of pools

Edit mode of the overall state of mirroring functions, which includes pools and images on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

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

- An image is created.
 - Mirroring is configured.
1. From the dashboard navigation, go to **Block > Mirroring**.
 2. From the **Pools** table, select the pool that requires editing, and click **Edit Mode**.
 3. In the **Edit pool mirror mode** dialog, select the mode from the drop-down, and then click **Update**.
A notification displays that the pool was updated successfully.

Adding peer in mirroring

Add storage cluster peer for the rbd-daemon mirror to discover its peer storage cluster on the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- Two healthy running Red Hat Ceph Storage clusters.
- The 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.

For more information, see [Configuring two-way mirroring](#).

Creating a bootstrap token from site A

1. From the dashboard navigation, go to **Block > Mirroring** and click **Create Bootstrap Token**.
2. On the **Navigation** menu, click the **Block** drop-down menu, and click **Mirroring**.
3. Click **Create Bootstrap token**.
 - a. For the provided site name, select the pools to be mirrored.
 - b. For the selected pools, generate a new bootstrap token by clicking **Generate**.
 - c. Click **Copy to Clipboard**.
 - d. Click **Close**.

Figure: Create bootstrap token

4. Enable the pool mirror mode.
 - a. Select the pool.
 - b. Click **Edit Mode**.
 - c. From the **Edit pool mirror mode** dialog, select **Image** from the **Mode** list.
 - d. Click **Update**.

A notification displays that the pool was updated successfully.

Figure: Edit pool mirror mode

Importing the bootstrap token on site B

1. From the dashboard navigation, go to **Block > Mirroring** and click **Import Bootstrap Token** from the action drop-down.

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

2. In the **Import Bootstrap Token** dialog, select the direction, and paste the token copied earlier, from **Site A**.

3. Click **Submit**.
- Figure: Import bootstrap token*

The peer is added and the images are mirrored in the cluster at site B.

- ## Creating a Block Device image with mirroring on site A

- Note:** Select mirroring with either **Journal** or **Snapshot**.

- e. Click **Create Image**.
2. Verify that the image is available at both the sites.
 - From the **Images** table, verify that the image in Site A is set to **primary** and that the image in site B is set to **secondary**.

Edit storage cluster peer for the `rbd-daemon` mirror to discover its peer storage cluster in the Red Hat Ceph Storage Dashboard.

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

- 116

- A pool with the `rbd` application enabled is created.
 - An image is created.
 - Mirroring is configured.
 - A peer is added.
1. From the dashboard navigation, go to **Block > Mirroring**.
 2. From the **Pools** table, click **Edit Peer** from the pool menu.
 3. In the **Edit pool mirror peer** dialog, edit the parameters and click **Submit**.
A notification displays that the peer was updated successfully.

*Figure: Completing the **Edit pool mirror peer** form*

Edit pool mirror peer [X]

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

Cluster Name (required)
cluster1

CephX ID (required)
cluster-id_123

Monitor Addresses
Comma-delimited addresses...

CephX Key
Base64-encoded key...

Cancel Submit

Deleting peer in mirroring

Delete a storage cluster peer for the `rbd-daemon` mirror in the Red Hat Ceph Storage Dashboard.

PREREQUISITE

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

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A pool with the `rbd` application enabled is created.
- An image is created.
- Mirroring is configured.
- A peer is added.

For more information, see [Adding peer in mirroring](#).

1. From the dashboard navigation, go to **Block > Mirroring**.
2. From the **Pools** table, select the pool to edit and click **Delete Peer** from the action drop-down.
3. In the **Delete mirror peer** dialog, select **Yes, I am sure** and click **Delete mirror peer**.
A notification displays that the peer deleted successfully.

Activating and deactivating telemetry

Activate the telemetry module to help Ceph developers understand how Ceph is used and what problems users might be experiencing, helping to 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](#). The public dashboard 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. After telemetry is 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 [Telemetry](#).

1. Activate the telemetry module by going 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 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

Deactivating telemetry

- To deactivate the telemetry module, go to **Settings › Telemetry configuration** and click **Deactivate**.