Red Hat Ceph Storage 4

Dashboard Guide

Monitoring Ceph Cluster with Ceph Dashboard

Last Updated: 2022-11-04

Red Hat Ceph Storage 4 Dashboard Guide

Monitoring Ceph Cluster with Ceph Dashboard

Legal Notice

Abstract

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

Table of Contents

CHAPTER 1. CEPH DASHBOARD OVERVIEW

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

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

1.1. PREREQUISITES

- System administrator level experience.

1.2. DASHBOARD COMPONENTS

The functionality of the dashboard is provided by multiple components.

- The Ansible automation 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 node of the storage cluster.
- The Grafana platform to provide monitoring user interface and alerting.

Additional Resources

- For more information, see the Ansible website
- For more information, see the Prometheus website.
- For more information, see the Grafana website.

1.3. DASHBOARD FEATURES

The Ceph dashboard provides multiple features.

Management features

- **View cluster hierarchy**: You can view the CRUSH map, for example, to determine which node a specific OSD ID is running on. This is helpful if there is an issue with an OSD.
- **Configure manager modules**: You can view and change parameters for ceph manager modules.
- **View and filter logs**: You can view event and audit cluster logs and filter them based on priority, keyword, date, or time range.

- **Toggle dashboard components**: You can enable and disable dashboard components so only the features you need are available.
- **Manage OSD settings**: You can set cluster-wide OSD flags using the dashboard.
- **Viewing Alerts**: The alerts page allows you to see details of current alerts.
- **Quality of Service for images** You can set performance limits on images, for example limiting IOPS or read BPS burst rates.

Monitoring features

- **Username and password protection**: You can access the dashboard only by providing a configurable user name and password.
- **SSL and TLS support**: All HTTP communication between the web browser and the dashboard is secured via SSL. A self-signed certificate can be created with a built-in command, but it is also possible to import custom certificates signed and issued by a Certificate Authority (CA). From Red Hat Ceph Storage 4.2, **dashboard_protocol** is set to **https** and Ansible generates the dashboard and grafana certificate. To plot data points and graphs, update the TLS handshake manually as:
  - Alert manager API host - http://grafana_node:9093
  - Prometheus API host - http://grafana_node:9092
  - Grafana API Host - https://grafana_node:3000
- **Overall cluster health**: Displays the overall cluster status, storage utilization (For example, number of objects, raw capacity, usage per pool), a list of pools and their status and usage statistics.
- **Hosts**: Provides a list of all hosts associated with the cluster along with the running services and the installed Ceph version.
- **Performance counters**: Displays detailed statistics for each running service.
- **Monitors**: Lists all Monitors, their quorum status and open sessions.
- **Configuration Reference**: Lists all available configuration options, their description and default values.
- **Cluster logs**: Display and filter the cluster's event and audit logs.
- **View storage cluster capacity**: You can view raw storage capacity of the Red Hat Ceph Storage cluster in the *Capacity* panels of the Ceph dashboard.
- **Pools**: Lists and manages all Ceph pools and their details. For example: applications, placement groups, replication size, EC profile, CRUSH ruleset, etc.
- **OSDs**: Lists and manages all OSDs, their status and usage statistics as well as detailed information like attributes (OSD map), metadata, performance counters and usage histograms for read/write operations.
- **iSCSI**: Lists all hosts that run the tcmu-runner service, displays all images and their performance characteristics, such as read and write operations or traffic.

- **Images**: Lists all RBD images and their properties such as size, objects, and features. Create, copy, modify and delete RBD images. Create, delete, and rollback snapshots of selected images, protect or unprotect these snapshots against modification. Copy or clone snapshots, flatten cloned images.

NOTE

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

- **Mirroring**: Lists all active sync daemons and their status, pools and RBD images including their synchronization state.
- **Filesystems**: Lists all active Ceph file system (CephFS) clients and associated pools, including their usage statistics.
- **Object Gateway (RGW)**: Lists all active object gateways and their performance counters. Displays and manages (adds, edits, deletes) object gateway users and their details, for example quotas, as well as the users' buckets and their details, for example, owner or quotas.

Additional Resources

- See *Toggling dashboard components on or off* in the *Red Hat Ceph Storage Dashboard Guide* for more information.

1.3.1. Toggling dashboard features on or off

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

Available features:

- Ceph Block Devices:
  - Image management, **rbd**
  - Mirroring, **mirroring**
  - iSCSI gateway, **iscsi**
- Ceph Filesystem, **cephfs**
- Ceph Object Gateway, **rgw**

NOTE

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

NOTE

You can disable multiple features at once.

IMPORTANT

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

Prerequisites

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

Procedure

1. To toggle the dashboard features from the dashboard web interface:
   a. From the navigation bar on the dashboard page, navigate to *Cluster*, then *Manager Modules*, then click on *Dashboard*. This opens the *Edit Manager module* page.
   b. From the *Edit Manager module* page, you can enable or disable the dashboard features by checking or unchecking the selection box next to the feature name.

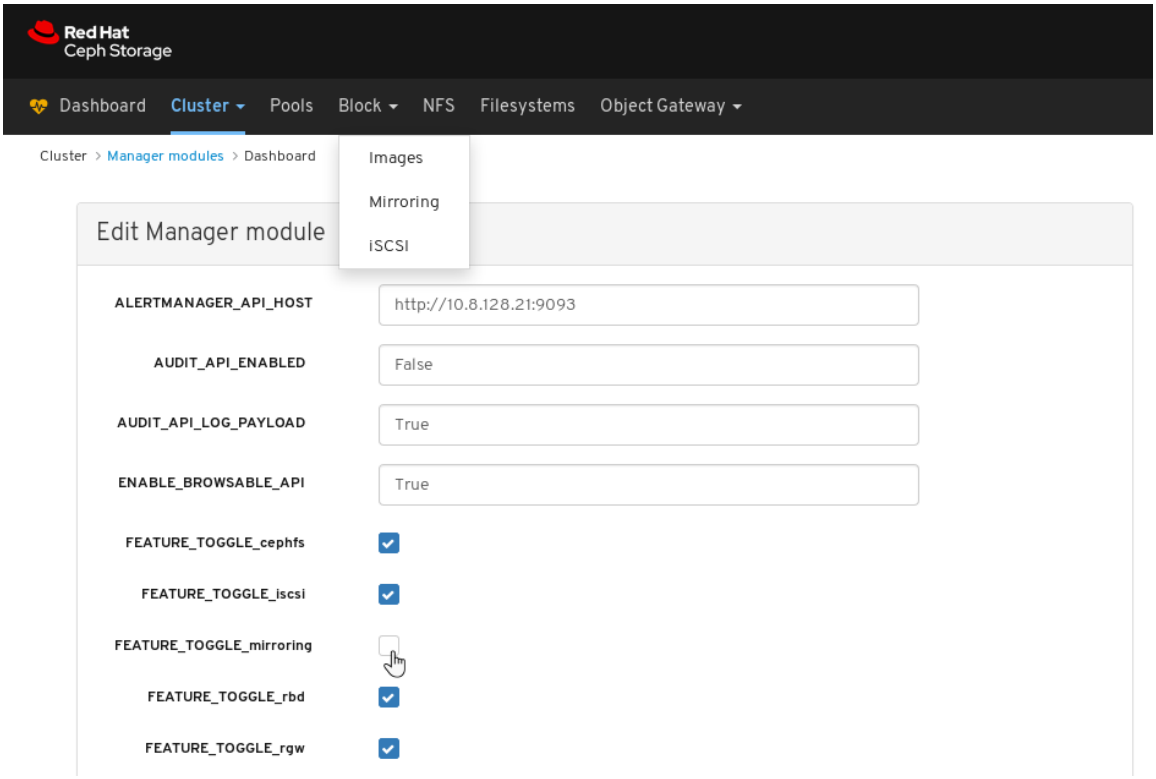

   c. Once the selections have been made, click on the *Update* button at the bottom of the page.

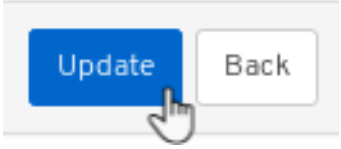

2. To toggle the dashboard features from the command-line interface:

a. Log in to the Ceph Manager node.

b. List the feature status:

```
[user@mon ~]$ ceph dashboard feature status
```

c. Disable a feature:

```
[user@mon ~]$ ceph dashboard feature disable iscsi
```

This example disables the Ceph iSCSI gateway feature.

d. Enable a feature:

```
[user@mon ~]$ ceph dashboard feature enable cephfs
```

This example enables the Ceph Filesystem feature.

1.4. DASHBOARD ARCHITECTURE

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

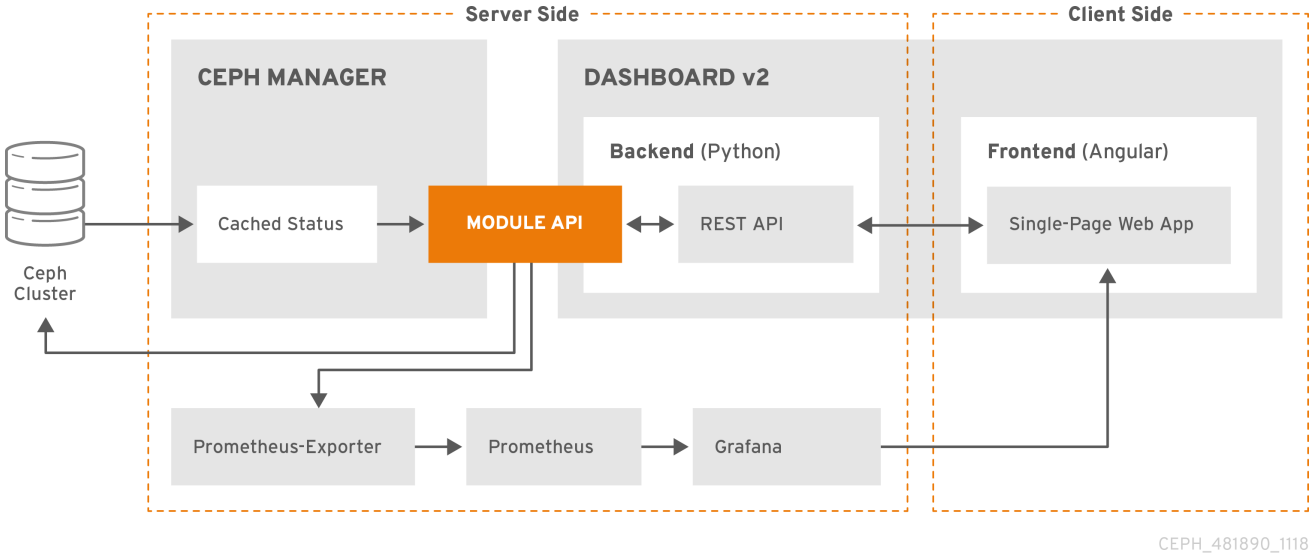

CHAPTER 2. CEPH DASHBOARD INSTALLATION AND ACCESS

As a system administrator, you can install dashboard and access it for the first time.

Red Hat Ceph Storage is installed graphically using the Cockpit web interface, or on the command line using the Ansible playbooks provided by the **ceph-ansible** RPM. Cockpit uses the same Ansible playbooks to install Ceph. Those playbooks install dashboard by default. Therefore, whether you directly use the Ansible playbooks, or use Cockpit to install Ceph, dashboard will be installed.

IMPORTANT

Change the default dashboard password. By default, the password for dashboard is **p@ssw0rd**, which is insecure. You can change the default password before installing Ceph by updating **dashboard_admin_password** in the **all.yml** Ansible playbook before using the playbooks to install Ceph, or after install using the same playbook, or dashboard itself. For more information, see the Install Guide, Changing the dashboard password using the dashboard, or Changing the dashboard password using Ansible.

2.1. INSTALLING DASHBOARD USING COCKPIT

Dashboard is installed by default when using the Cockpit web interface to install Red Hat Ceph Storage. You must set a host with the *Metrics* role for Grafana to be installed on.

Prerequisites

- Consult the Installation Guide for full prerequisites. This procedure only highlights the steps relevant to the dashboard install.

Procedure

1. On the *Hosts* page, add a host and set the *Metrics* role.

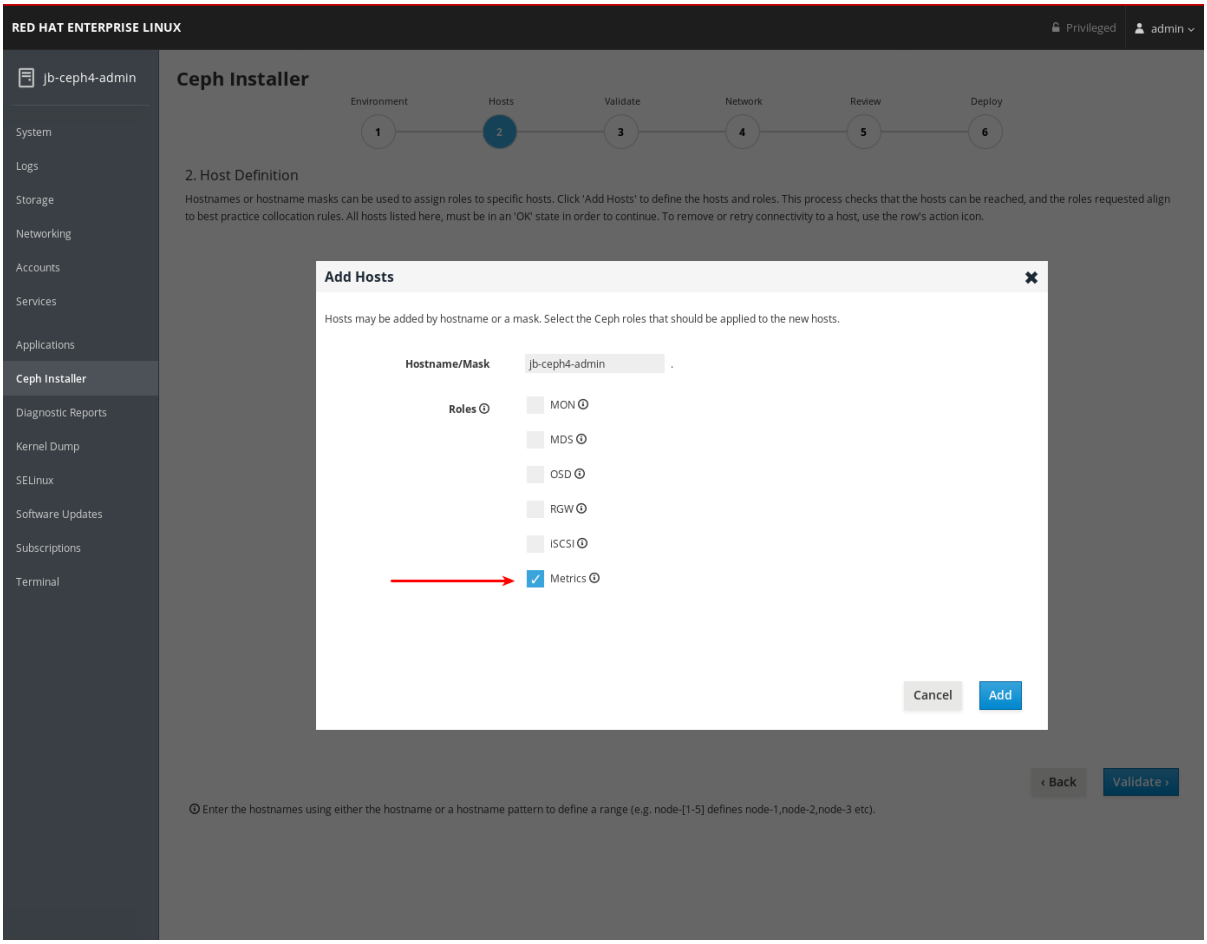

2. Click Add.

3. Complete the remaining Cockpit Ceph Installer prompts.

4. After the deploy process finishes, click the *Complete* button at the bottom right corner of the page. This opens a window which displays the output of the command **ceph status**, as well as dashboard access information.

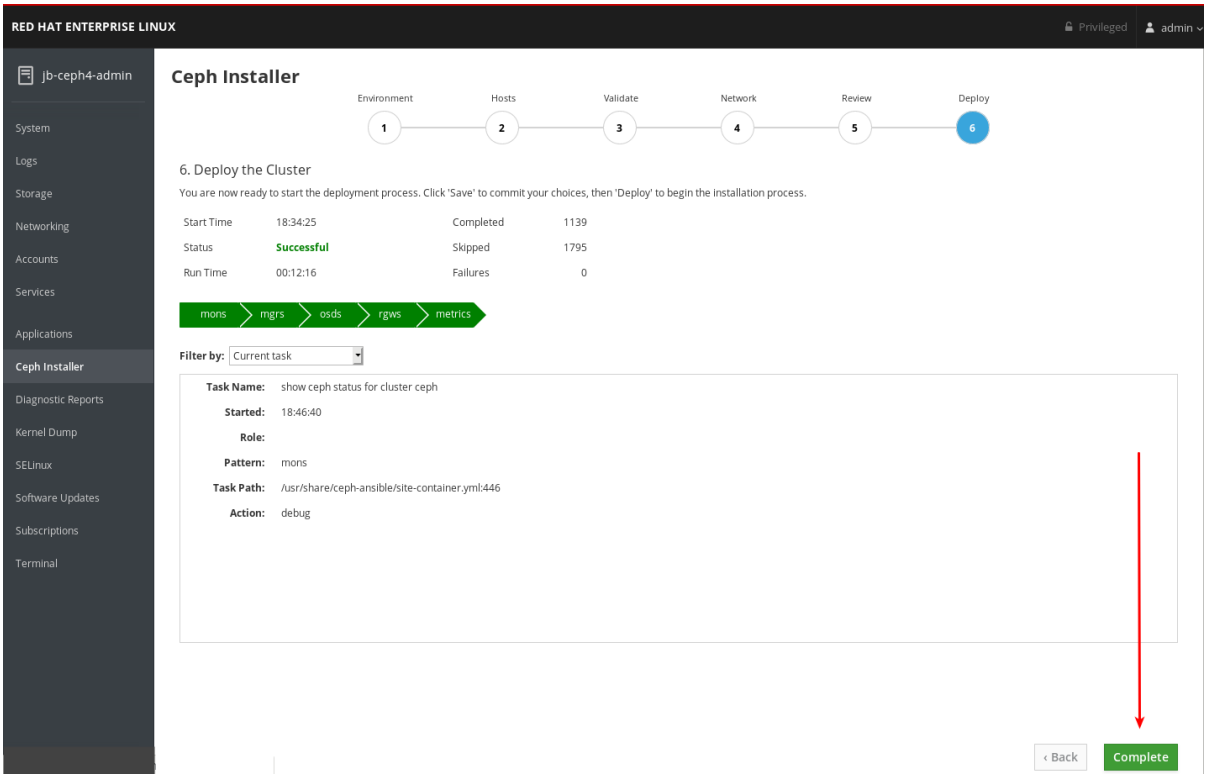

5. At the bottom of the Ceph Cluster Status window, the dashboard access information is displayed, including the URL, user name, and password. Take note of this information.

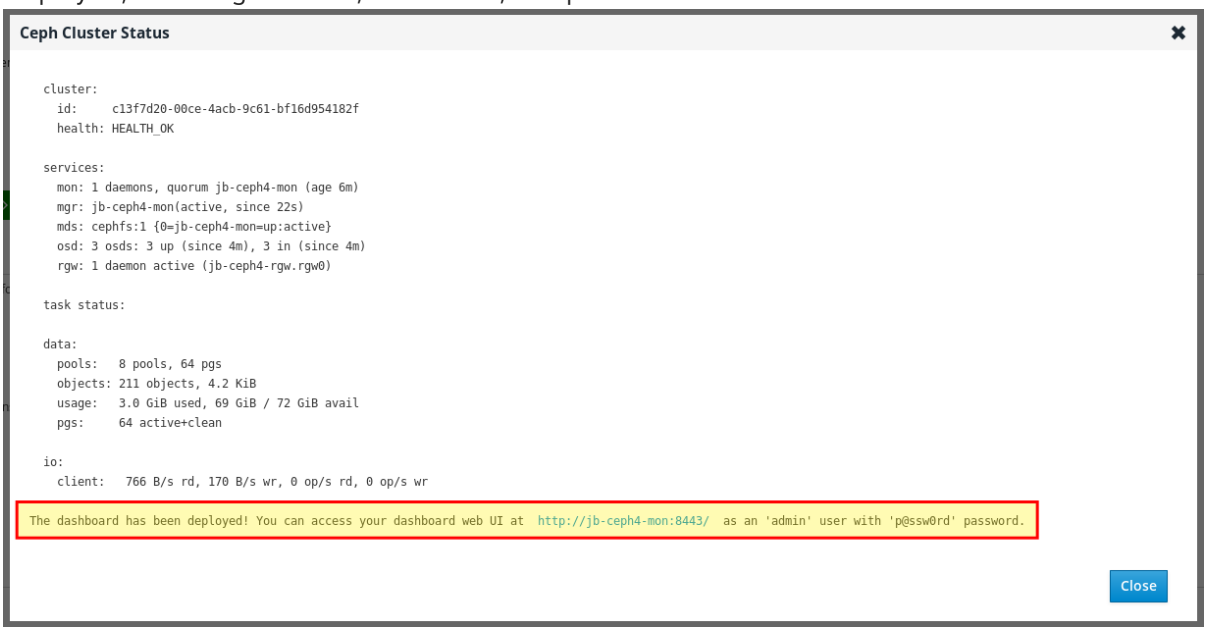

- For more information, see Installing Red Hat Ceph Storage using the Cockpit Web User Interface in the Installation Guide.

2.2. INSTALLING DASHBOARD USING ANSIBLE

Dashboard is installed by default when installing Red Hat Ceph Storage using the Ansible playbooks provided by the **ceph-ansible** RPM.

Prerequisites

- Consult the Installation Guide for full prerequisites. This procedure only highlights the steps relevant to the dashboard install.

Procedure

1. Ensure a **[grafana-server]** group with a node defined under it exists in the Ansible inventory file. Grafana and Prometheus are installed on this node.

```
[root@jb-ceph4-admin ~]# grep grafana-server -A 1 /etc/ansible/hosts
[grafana-server]
jb-ceph4-mon
```

2. In the **all.yml** Ansible playbook, ensure **dashboard_enabled:** has not been set to **False**. There should be a comment indicating the default setting of **True**.

```
[root@jb-ceph4-admin ~]# grep "dashboard_enabled" /usr/share/ceph-
ansible/group_vars/all.yml
#dashboard_enabled: True
```

3. Complete the rest of the steps necessary to install Ceph as outlined in the Installation Guide.

4. After running **ansible-playbook site.yml** for bare metal installs, or **ansible-playbook site-docker.yml** for container installs, Ansible will print the dashboard access information. Find the dashboard URL, username, and password towards the end of the playbook output:

```
2019-12-13 15:31:17,871 p=11421 u=admin |  TASK [ceph-dashboard : print dashboard
URL] ************************************************************
2019-12-13 15:31:17,871 p=11421 u=admin |  task path: /usr/share/ceph-ansible/roles/ceph-
dashboard/tasks/main.yml:5
2019-12-13 15:31:17,871 p=11421 u=admin |  Friday 13 December 2019  15:31:17 -0500
(0:00:02.189)       0:04:25.380 *******
2019-12-13 15:31:17,934 p=11421 u=admin |  ok: [jb-ceph4-mon] =>
  msg: The dashboard has been deployed! You can access your dashboard web UI at
http://jb-ceph4-mon:8443/ as an 'admin' user with 'p@ssw0rd' password.
```

Take note of the output **You can access your dashboard web UI at http://jb-ceph4-mon:8443/ as an 'admin' user with 'p@ssw0rd' password.**

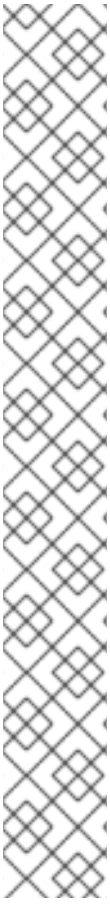

NOTE

The Ansible playbook does the following:

- Enables the Prometheus module in **ceph-mgr**.
- Enables the dashboard module in **ceph-mgr** and opens TCP port 8443.
- Deploys the Prometheus **node_exporter** daemon to each node in the storage cluster.
  - Opens TCP port 9100.
  - Starts the **node_exporter** daemon.
- Deploys Grafana and Prometheus containers under Docker/systemd on the node under **[grafana-server]** in the Ansible inventory file.
  - Configures Prometheus to gather data from the ceph-mgr nodes and the node-exporters running on each Ceph host
  - Opens TCP port 3000.
  - Creates the dashboard, theme, and user accounts in Grafana.
  - Displays the Ceph Dashboard login page URL.

- For more information, see *Installing a Red Hat Ceph Storage cluster* in the *Red Hat Ceph Storage Installation Guide*.
- To remove the dashboard, see *Purging the Ceph Dashboard using Ansible* in the *Red Hat Ceph Storage Installation Guide*.

2.3. NETWORK PORT REQUIREMENTS

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

Table 2.1. TCP Port Requirements

Port	Use	Originating Node	Destination Node
8443	The dashboard web interface	IP addresses that need access to Ceph Dashboard UI and the node under **[grafana-server]** in the Ansible inventory file, since the AlertManager service can also initiate connections to the Dashboard for reporting alerts.	The Ceph Manager nodes.
3000	Grafana	IP addresses that need access to Grafana Dashboard UI and all Ceph Manager hosts and **[grafana-server]**.	The node under **[grafana-server]** in the Ansible inventory file.
9090	Default Prometheus server for basic Prometheus graphs	IP addresses that need access to Prometheus UI and all Ceph Manager hosts and **[grafana-server]** or Hosts running Prometheus.	The node under **[grafana-server]** in the Ansible inventory file.
9092	Prometheus server for basic Prometheus graphs	IP addresses that need access to Prometheus UI and all Ceph Manager hosts and **[grafana-server]** or Hosts running Prometheus.	The node under **[grafana-server]** in the Ansible inventory file.
9093	Prometheus Alertmanager	IP addresses that need access to Alertmanager Web UI and all Ceph Manager hosts and **[grafana-server]** or Hosts running Prometheus.	All Ceph Manager nodes and the node under **[grafana-server]** in the Ansible inventory file.
9094	Prometheus Alertmanager for configuring a highly available cluster made from multiple instances	All Ceph Manager nodes and the node under **[grafana-server]** in the Ansible inventory file.	Prometheus Alertmanager High Availability (peer daemon sync), so both **src** and **dst** should be nodes running Prometheus Alertmanager.

Port	Use	Originating Node	Destination Node
9100	The Prometheus **node-exporter** daemon	Hosts running Prometheus that need to view Node Exporter metrics Web UI and all Ceph Manager nodes and **[grafana-server]** or Hosts running Prometheus.	All storage cluster nodes, including MONs, OSDS, **[grafana-server]** host.
9283	Ceph Manager Prometheus exporter module	Hosts running Prometheus that need access to Ceph Exporter metrics Web UI and **[grafana-server]**.	All Ceph Manager nodes.
9287	Ceph iSCSI gateway data	All Ceph Manager hosts and **[grafana-server]**.	All Ceph iSCSI gateway nodes.

Additional Resources

- For more information, see the *Red Hat Ceph Storage Installation Guide*.
- For more information, see *Using and configuring firewalls* in *Configuring and managing networking*.

2.4. CONFIGURING DASHBOARD PORTS

The Red Hat Ceph Storage Dashboard, by default, binds to a TCP/IP address and TCP port.

By default, the **ceph-mgr** daemon hosting the dashboard binds to TCP port 8443 or 8080 when SSL is disabled. If no specific address is configured, the web app binds to **::**, which corresponds to all the available IP4 and IP6 addresses.

You can change the IP address and the port using the configuration key facility on a cluster-wide level.

Prerequisites

- A Red Hat Ceph Storage cluster.
- Installation of the Red Hat Ceph Storage Dashboard.
- Root-level access to all the nodes.

Procedure

1. Get the URL for accessing the dashboard:

   Example

```
[root@admin ~]# ceph mgr services
```

2. Get the current IP and port configuration of the **ceph-mgr** daemon:

   Example

   ```
   [root@admin ~]# netstat -ntlp
   ```

3. Set the IP address and the port:

   Syntax

   ```
   ceph config set mgr mgr/dashboard/server_addr IP_ADDRESS
   ceph config set mgr mgr/dashboard/server_port PORT
   ceph config set mgr mgr/dashboard/ssl_server_port PORT
   ```

   Example

   ```
   [root@mon ~]# ceph config set mgr mgr/dashboard/server_addr 192.168.0.120
   [root@mon ~]# ceph config set mgr mgr/dashboard/server_port 8443
   [root@mon ~]# ceph config set mgr mgr/dashboard/ssl_server_port 8443
   ```

4. Optional: Since the **ceph-mgr** hosts its own instance of the dashboard, you can configure them separately. Change the IP address and port for a specific manager instance:

   Syntax

   ```
   ceph config set mgr mgr/dashboard/NAME/server_addr IP_ADDRESS
   ceph config set mgr mgr/dashboard/NAME/server_port PORT
   ceph config set mgr mgr/dashboard/NAME/ssl_server_port PORT
   ```

   Replace: *NAME* with the ID of the **ceph-mgr** instance hosting the dashboard.

   Example

   ```
   [root@mon ~]# ceph config set mgr mgr/dashboard/mgrs-0/server_addr 192.168.0.120
   [root@mon ~]# ceph config set mgr mgr/dashboard/mgrs-0/server_port 8443
   [root@mon ~]# ceph config set mgr mgr/dashboard/mgrs-0/ssl_server_port 8443
   ```

Additional Resources

- See the Knowledgebase article *How to update the IP address or Port of the Ceph-dashboard* for more details.

2.5. ACCESSING DASHBOARD

Accessing the dashboard allows you to administer and monitor your Red Hat Ceph Storage cluster.

Prerequisites

- Successful installation of Red Hat Ceph Storage Dashboard.

- NTP is synchronizing clocks properly.

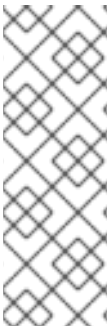

NOTE

A time lag can occur between the dashboard node, cluster nodes, and a browser, when the nodes are not properly synced. Ensure all nodes and the system where the browser runs have time synced by NTP. By default, when Red Hat Ceph Storage is deployed, Ansible configures NTP on all nodes. To verify, for Red Hat Enterprise Linux 7, see Configuring NTP Using ntpd, for Red Hat Enterprise Linux 8, see Using the Chrony suite to configure NTP. If you run your browser on another operating system, consult the vendor of that operating system for NTP configuration information.

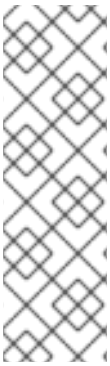

NOTE

When using OpenStack Platform (OSP) with Red Hat Ceph Storage, to enable OSP Safe Mode, use one of the following methods. With Ansible, edit the **group_vars/all.yml** Ansible playbook, set **dashboard_admin_user_ro: true** and re-run **ansible-playbook** against **site.yml**, or **site-container.yml**, for bare-metal, or container deployments, respectively. To enable OSP Safe Mode using the **ceph** command, run **ceph dashboard ac-user-set-roles admin read-only**. To ensure the changes persist if you run the **ceph-ansible** Ansible playbook, edit **group_vars/all.yml** and set **dashboard_admin_user_ro: true**.

Procedure

1. Enter the following URL in a web browser:

   ```
   http://HOST_NAME:PORT
   ```

   Replace:

   - *HOST_NAME* with the host name of the dashboard node.
   - *PORT* with port **8443**
     For example:

     ```
     http://dashboard:8443
     ```

2. On the login page, enter the username **admin** and the default password **p@ssw0rd** if you did not change the password during installation.

Figure 2.1. Ceph Dashboard Login Page

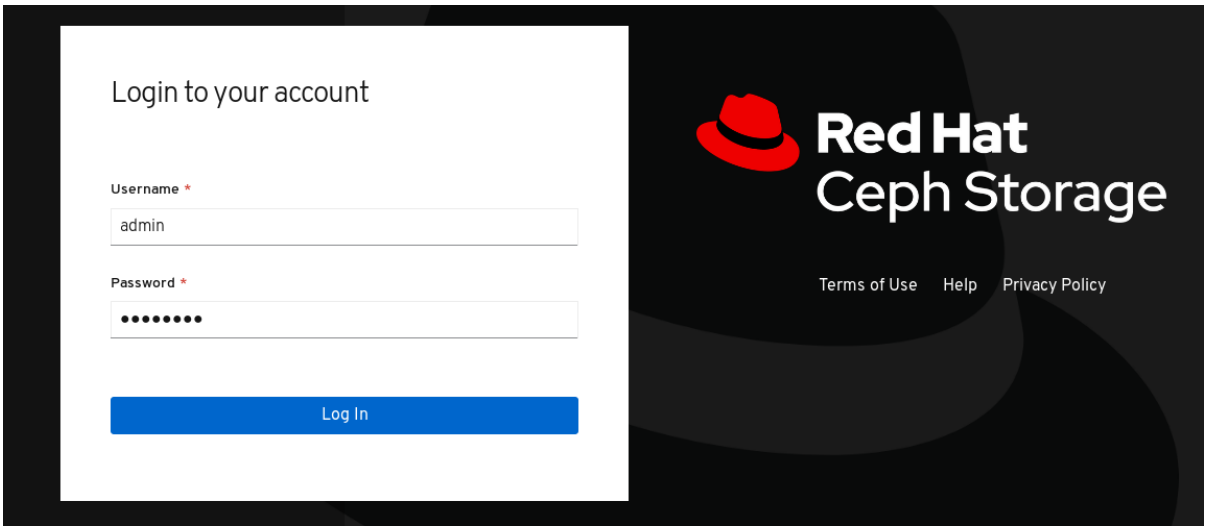

3. After logging in, the dashboard default landing page is displayed, which provides a high-level overview of status, performance, and capacity metrics of the Red Hat Ceph Storage cluster.

Figure 2.2. Ceph Dashboard Default Landing Page

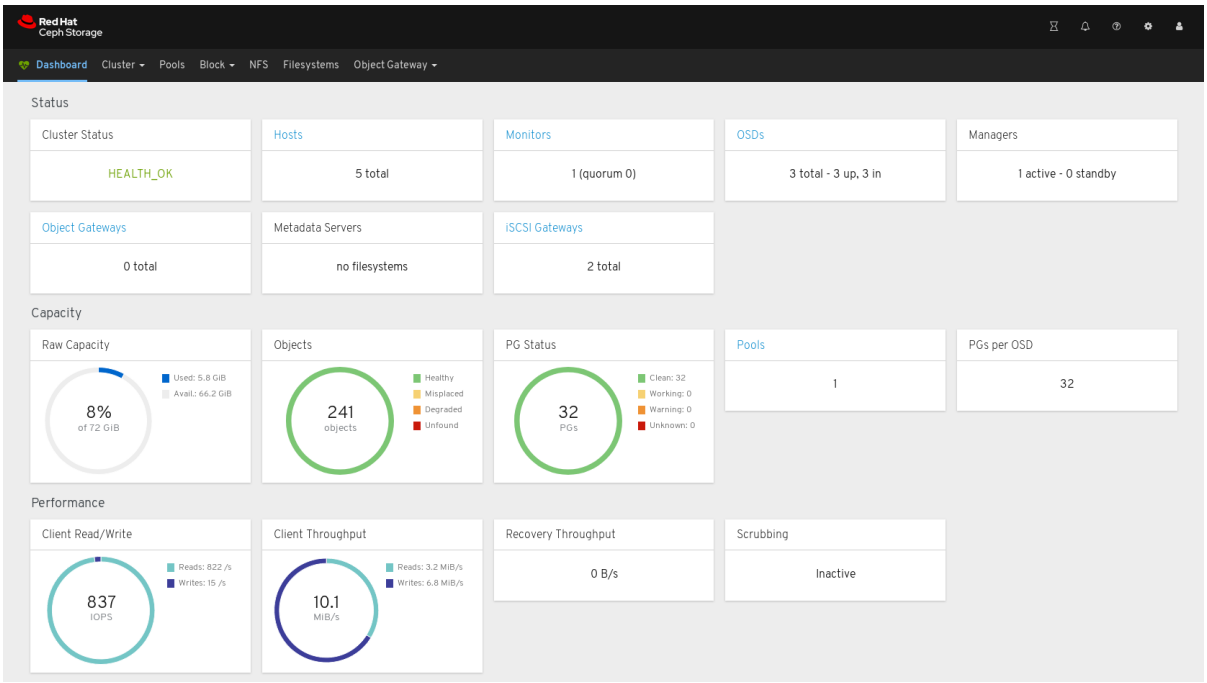

Additional Resources

- For more information, see Changing the dashboard password using the dashboard in the Dashboard guide.
- For more information, see Changing the dashboard password using Ansible in the Dashboard guide.

2.6. CHANGING THE DASHBOARD PASSWORD USING ANSIBLE

By default, the password for accessing dashboard is set to **p@ssw0rd**.

IMPORTANT

For security reasons, change the password after installation.

You can change the dashboard password using Ansible.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Access to the Ansible administration node.

Procedure

1. Open the Ansible playbook file **/usr/share/ceph-ansible/group_vars/all.yml** for editing.
2. Uncomment and update the password on this line:

   ```
   #dashboard_admin_password: p@ssw0rd
   ```

   to:

   ```
   dashboard_admin_password: NEW_PASSWORD
   ```

   Replace *NEW_PASSWORD* with your preferred password.
3. Rerun the Ansible playbook file which deploys or updates the Ceph cluster.
   a. For bare metal installs, use the **site.yml** playbook:

      ```
      [admin@admin ceph-ansible]$ ansible-playbook -v site.yml
      ```

   b. For container installs, use the **site-docker.yml** playbook:

      ```
      [admin@admin ceph-ansible]$ ansible-playbook -v site-docker.yml
      ```
4. Log in using the new password.

Additional Resources

- For more information, see Changing the dashboard password using the dashboard in the Dashboard guide.

2.7. CHANGING THE DASHBOARD PASSWORD USING THE DASHBOARD

By default, the password for accessing dashboard is set to **p@ssw0rd**.

IMPORTANT

For security reasons, change the password after installation.

To change the password using the dashboard, also change the dashboard password setting in Ansible to ensure the password does not revert to the default password if Ansible is used to reconfigure the Red Hat Ceph Storage cluster.

Prerequisites

- A running Red Hat Ceph Storage cluster.

Procedure

1. Update the password in the **group_vars/all.yml** file to prevent the password from being reset to **p@ssw0rd** when Ansible is used to reconfigure the Ceph cluster.
   a. Open the Ansible playbook file **/usr/share/ceph-ansible/group_vars/all.yml** for editing.
   b. Uncomment and update the password on this line:

      ```
      #dashboard_admin_password: p@ssw0rd
      ```

      to:

      ```
      dashboard_admin_password: NEW_PASSWORD
      ```

      Replace *NEW_PASSWORD* with your preferred password.
2. Change the password in the dashboard web user-interface.
   a. Log in to the dashboard:

      ```
      http://HOST_NAME:8443
      ```

   b. At the top right hand side toolbar, click the dashboard settings icon and then click **User management**.

      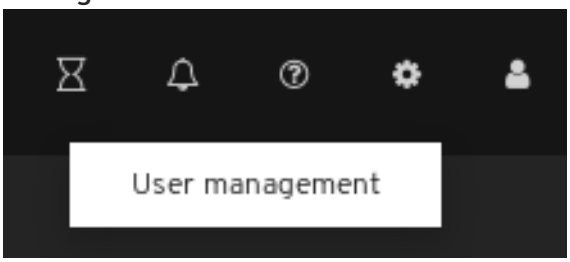

   c. Locate the **admin** user in the **Username** table and click on **admin**.

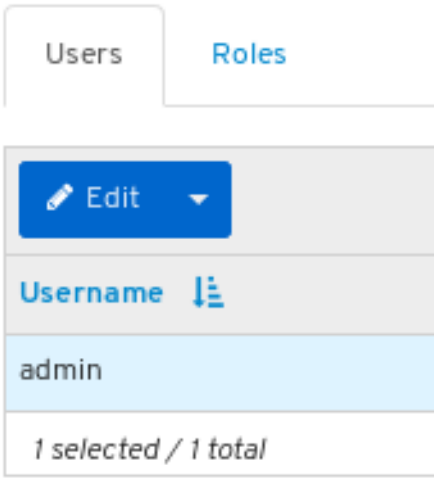

d. Above the table title **Username**, click on the **Edit** button.

e. Enter the new password and confirm it by reentering it and click **Edit User**.

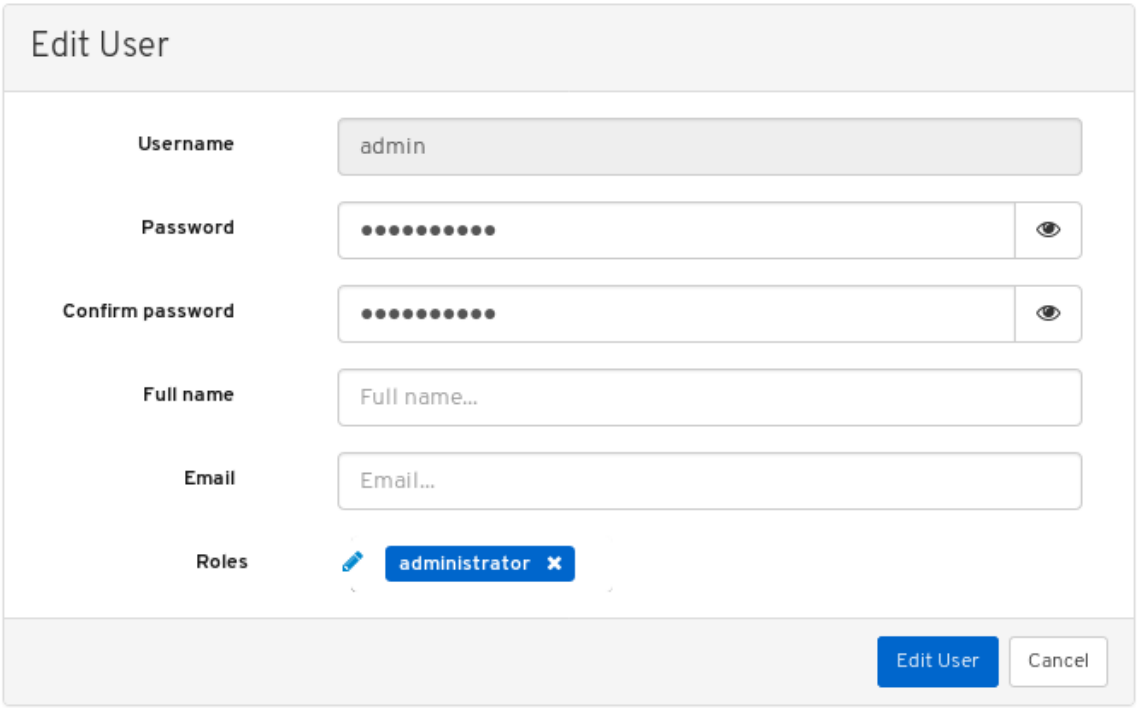

You will be logged out and taken to the log in screen. A notification will appear confirming the password change.

3. Log back in using the new password.

Additional Resources

- For more information, see Changing the dashboard password using Ansible in the Dashboard guide.

2.8. CHANGING THE GRAFANA PASSWORD USING ANSIBLE

By default, the password for Grafana, used by dashboard, is set to **admin**. Use this procedure to change the password.

IMPORTANT

For security reasons, change the password from the default.

Prerequisites

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

Procedure

1. Optional: If you do not know which node the Grafana container is running on, find the node listed under **[grafana-server]** in the Ansible hosts file, usually located at **/etc/ansible/hosts**:

   Example

   ```
   [grafana-server]
   grafana
   ```

2. On the node where the Grafana container is running, change the password:

   Syntax

   ```
   podman exec CONTAINER_ID grafana-cli admin reset-admin-password --homepath
   "/usr/share/grafana" NEW_PASSWORD
   ```

   Change *CONTAINER_ID* to the ID of the Grafana container. Change *NEW_PASSWORD* to the desired Grafana password.

   Example

   ```
   [root@grafana ~]# podman exec 3f28b0309aee grafana-cli admin reset-admin-password --
   homepath "/usr/share/grafana" NewSecurePassword
   t=2020-10-29T17:45:58+0000 lvl=info msg="Connecting to DB" logger=sqlstore
   dbtype=sqlite3
   t=2020-10-29T17:45:58+0000 lvl=info msg="Starting DB migration" logger=migrator

   Admin password changed successfully ✔
   ```

3. On the Ansible administration node, use **ansible-vault** to encrypt the Grafana password, and then add the encrypted password to **group_vars/all.yml**.

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

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

   b. Run **ansible-vault** and create a new vault password:

      Example

      ```
      [admin@admin ceph-ansible]$ ansible-vault encrypt_string --stdin-name
      'grafana_admin_password_vault'
      New Vault password:
      ```

c. Re-enter the password to confirm it:

Example

```
[admin@admin ceph-ansible]$ ansible-vault encrypt_string --stdin-name
'grafana_admin_password_vault'
New Vault password:
Confirm New Vault password:
```

d. Enter the Grafana password, press enter, and then enter CTRL+D to complete the entry:

Syntax

```
ansible-vault encrypt_string --stdin-name 'grafana_admin_password_vault'
New Vault password:
Confirm New Vault password:
Reading plaintext input from stdin. (ctrl-d to end input)
NEW_PASSWORD
```

Replace *NEW_PASSWORD* with the Grafana password that was set earlier.

Example

```
[admin@admin ceph-ansible]$ ansible-vault encrypt_string --stdin-name
'grafana_admin_password_vault'
New Vault password:
Confirm New Vault password:
Reading plaintext input from stdin. (ctrl-d to end input)
NewSecurePassword
```

e. Take note of the output that begins with **grafana_admin_password_vault: !vault |** and ends with a few lines of numbers, as it will be used in the next step:

Example

```
[admin@admin ceph-ansible]$ ansible-vault encrypt_string --stdin-name
'grafana_admin_password_vault'
New Vault password:
Confirm New Vault password:
Reading plaintext input from stdin. (ctrl-d to end input)
NewSecurePassword
grafana_admin_password_vault: !vault |
          $ANSIBLE_VAULT;1.1;AES256

38383639646166656130326666633262643836343930373836376331326437353032376165306234

31613863346166326535303832313166316364623637616 0a3733383733346634343638 65356633

66383963323033303662333765383938353630623433346565363534636434643634364330643438
```

```
6134306662646365370a343135316633303830653565633736303466636261326361333766613462

393533653431373231633439376364646635343832343265316661393765616635 32
Encryption successful
```

f. Open for editing **group_vars/all.yml** and paste the output from above into the file:

Example

```
grafana_admin_password_vault: !vault |
          $ANSIBLE_VAULT;1.1;AES256

383836396461666561303266666332626438363439303738363763313264373530323761653062 34

316138633461663265353038323131663163646236376166 0a373338373334663434363865356633

663839633230333036623337653839383536306234333465653635346364346436343643038643438

6134306662646365370a343135316633303830653565633736303466636261326361333766613462

393533653431373231633439376364646635343832343265316661393765616635 32
```

g. Add a line below the encrypted password with the following:

Example

```
grafana_admin_password: "{{ grafana_admin_password_vault }}"
```

NOTE

Using two variables as seen above is required due to a bug in Ansible that breaks the string type when assigning the vault value directly to the Ansible variable.

h. Save and close the file.

4. Re-run **ansible-playbook**.

a. For container based deployments:

Example

```
[admin@node1 ceph-ansible]$ ansible-playbook --ask-vault-pass -v site-container.yml -i hosts
```

Note that **-i hosts** is only necessary if you are not using the default Ansible hosts file location of **/etc/ansible/hosts**.

b. For bare-metal, RPM based deployments:

Example

```
[admin@node1 ceph-ansible]$ ansible-playbook --ask-vault-pass -v site.yml -i hosts
```

Note that **-i hosts** is only necessary if you are not using the default Ansible hosts file location of **/etc/ansible/hosts**.

2.9. SYNCING USERS USING RED HAT SINGLE SIGN-ON FOR THE DASHBOARD

Administrators can provide access to users on Red Hat Ceph Storage Dashboard using Red Hat Single Sign-on (SSO) with Lightweight Directory Access Protocol (LDAP) integration.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Admin level access to the dashboard.
- Users are added to the dashboard.
- Root-level access on all the nodes.
- Red hat Single Sign-On installed from a ZIP file. See the *Installing Red Hat Single Sign-On from a zip file* for additional information.

Procedure

1. Download the *Red Hat Single Sign-On 7.4.0 Server* on the system where Red Hat Ceph Storage is installed.
2. Unzip the folder:

```
[root@cephuser]# unzip rhsso-7.4.0.zip
```

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

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

4. Replace three instances of **localhost** and two instances of **127.0.0.1** with the IP address of the machine where Red Hat Single Sign-On is installed.
5. Optional: For Red Hat Enterprise Linux 8, users might get Certificate Authority (CA) issues. Import the custom certificates from CA and move them into the keystore with exact java version.

Example

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

6. To start the server from the **bin** directory of **rh-sso-7.4** folder, run the **standalone** boot script:

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

7. Create the admin account in http:_IP_ADDRESS_:8080/auth with a username and password:

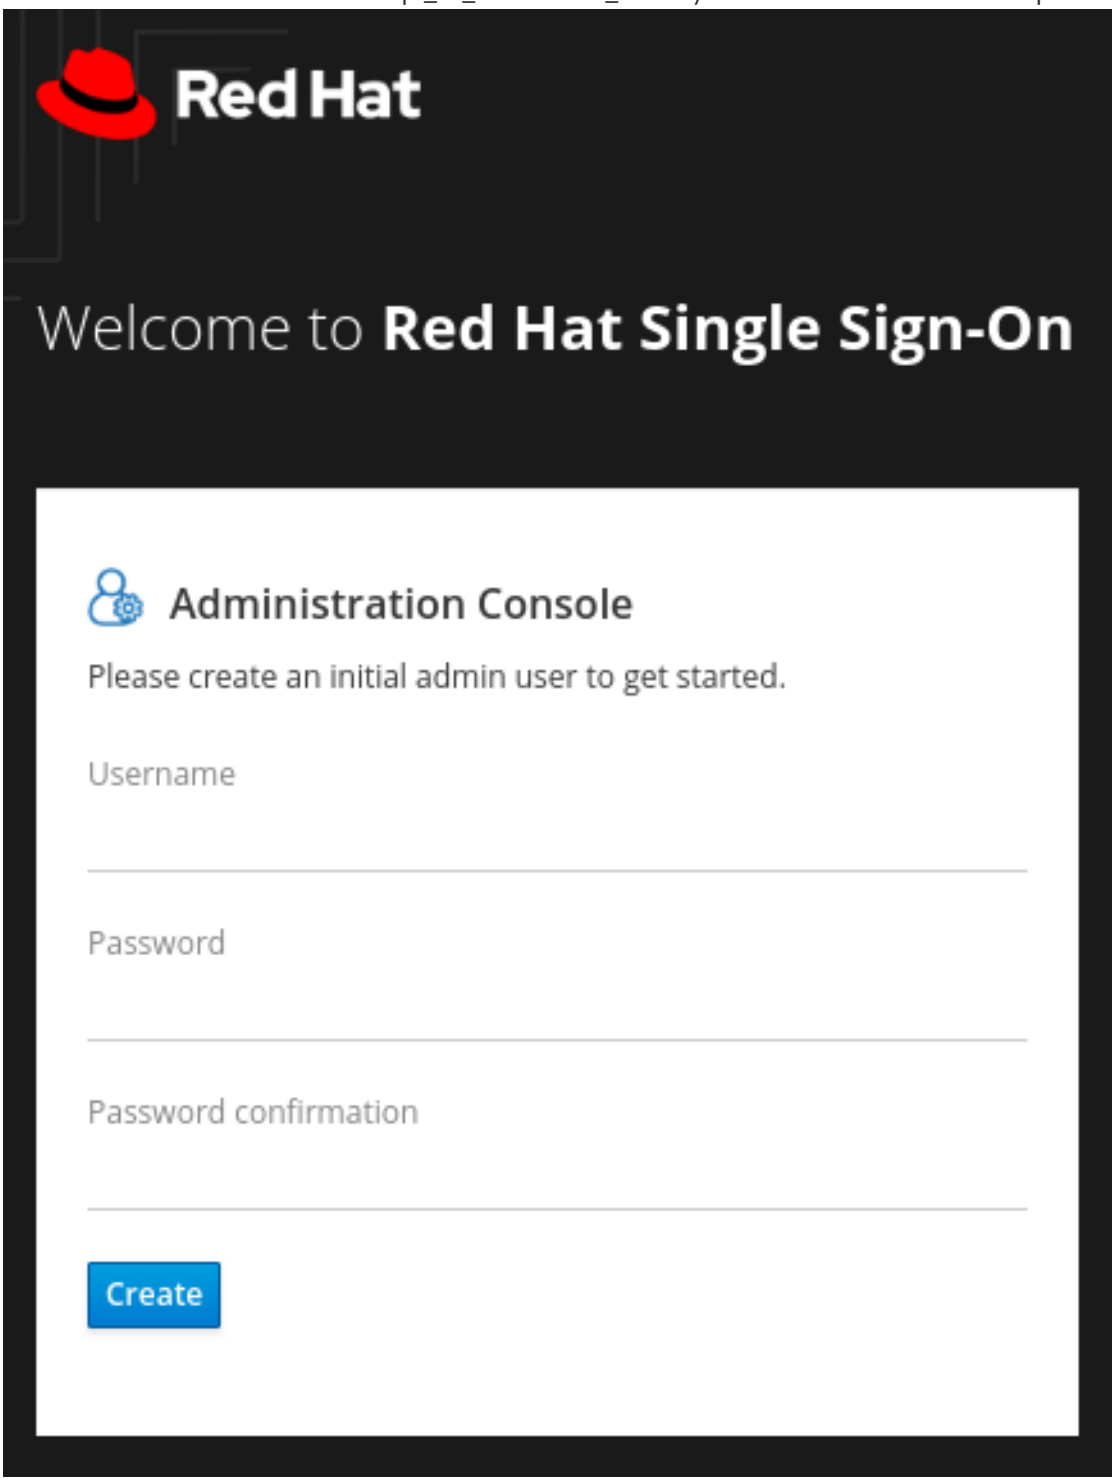

NOTE

The admin account has to be created only the first time you log into the console.

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

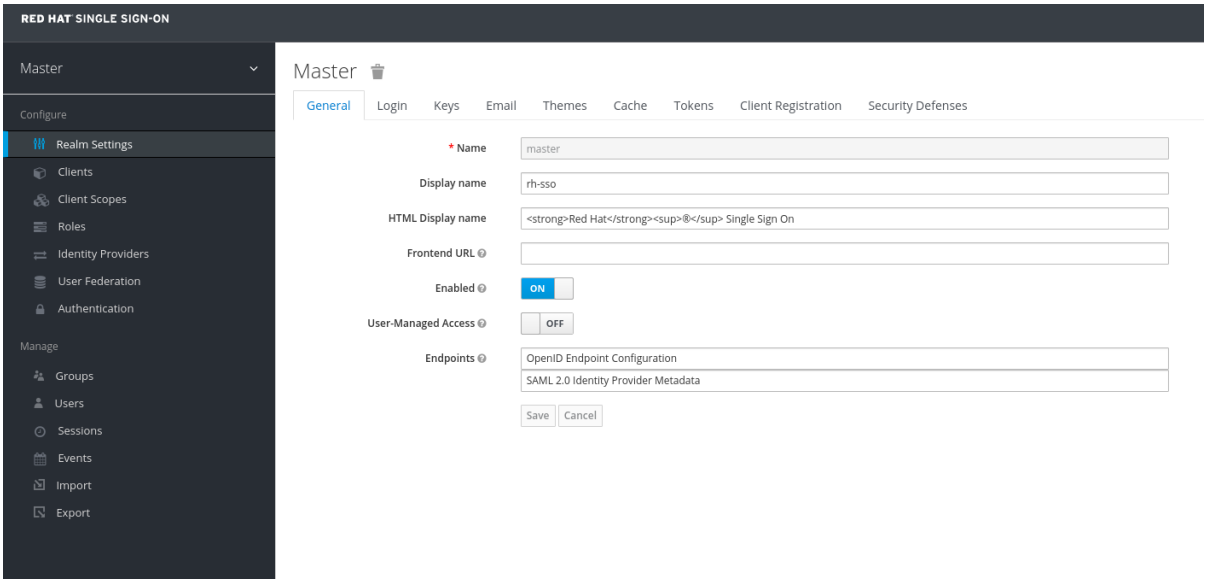

9. To create a realm, click the *Master* drop-down. In this realm, administrators provide access to users and applications.

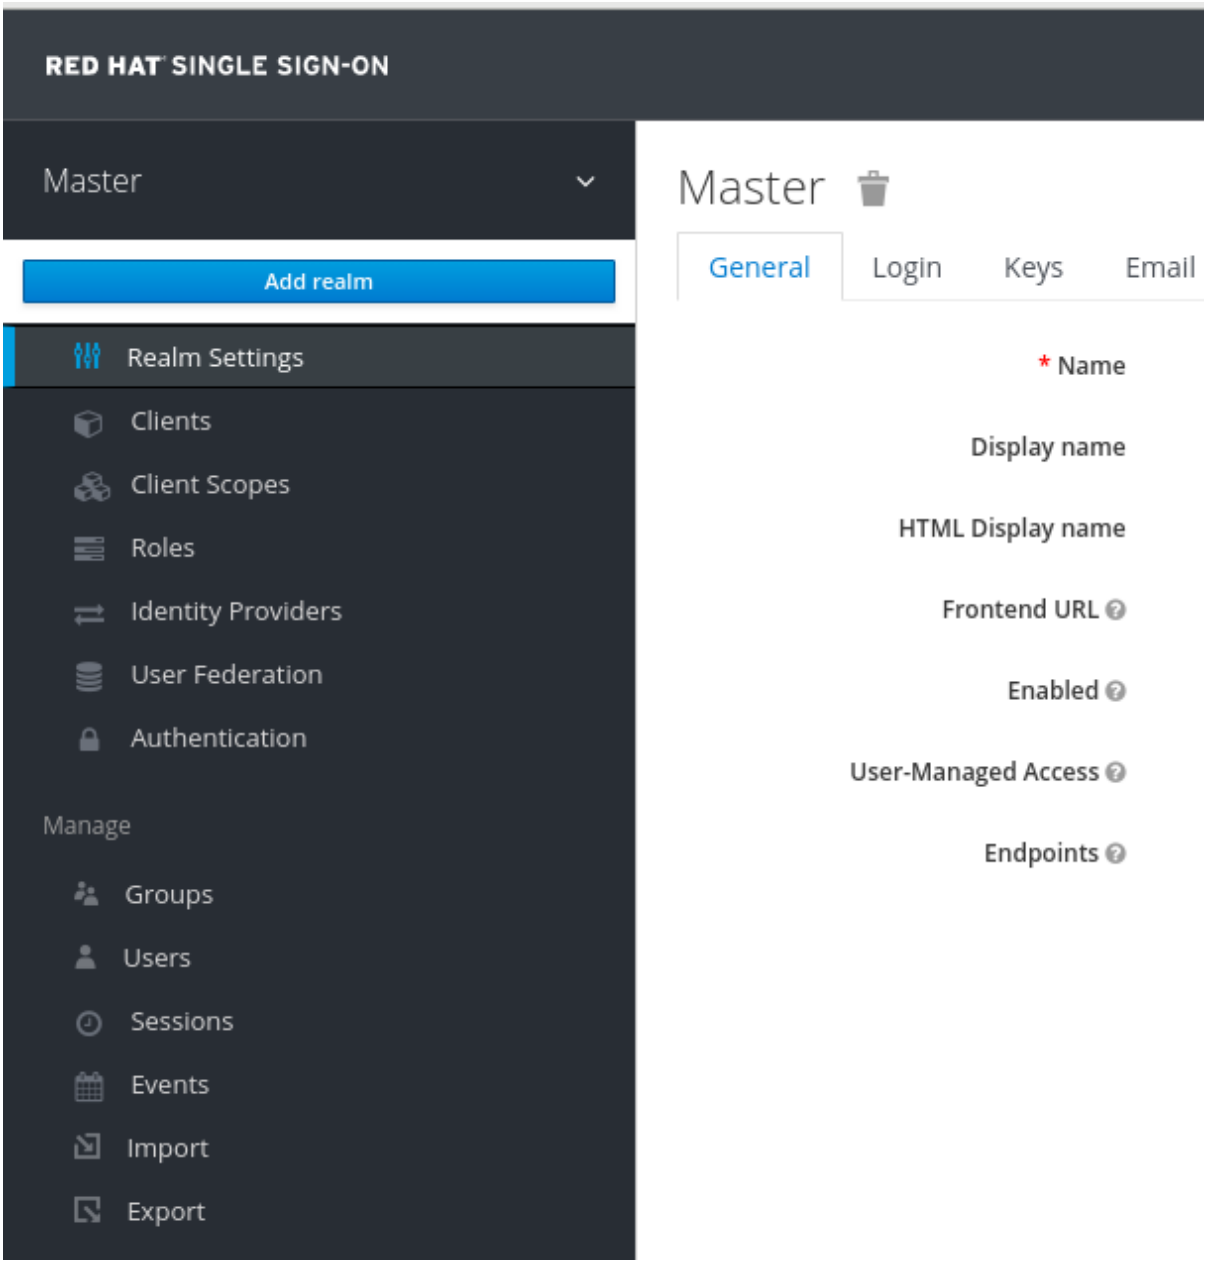

10. In the *Add Realm* window, enter a name for the realm and set the parameter *Enabled* to ON and click Create:

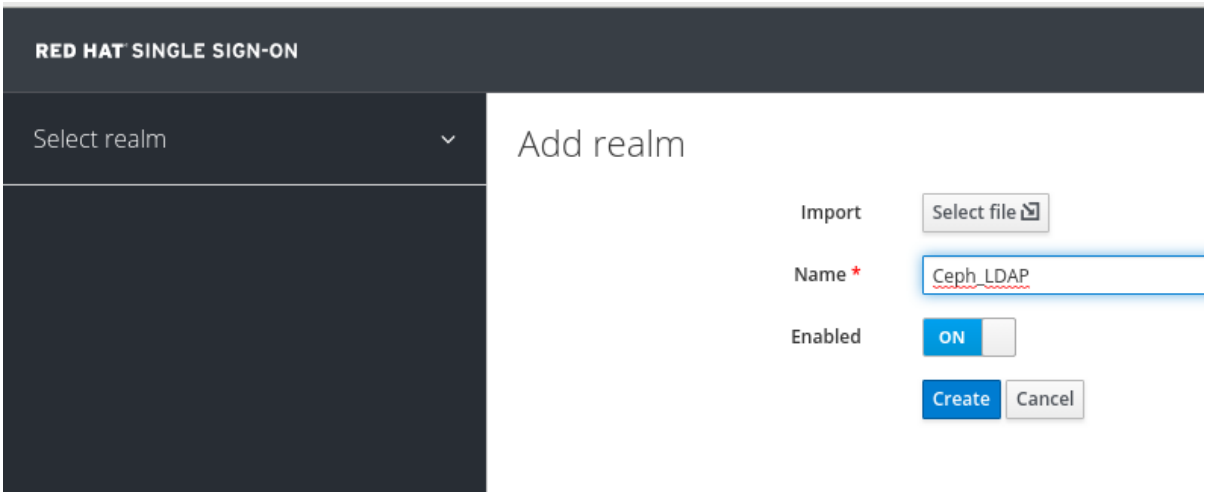

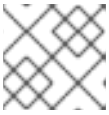

NOTE

The realm name is case-sensitive.

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

    a. Enabled - ON

    b. User-Managed Access - ON

    c. Copy the link address of SAML 2.0 Identity Provider Metadata

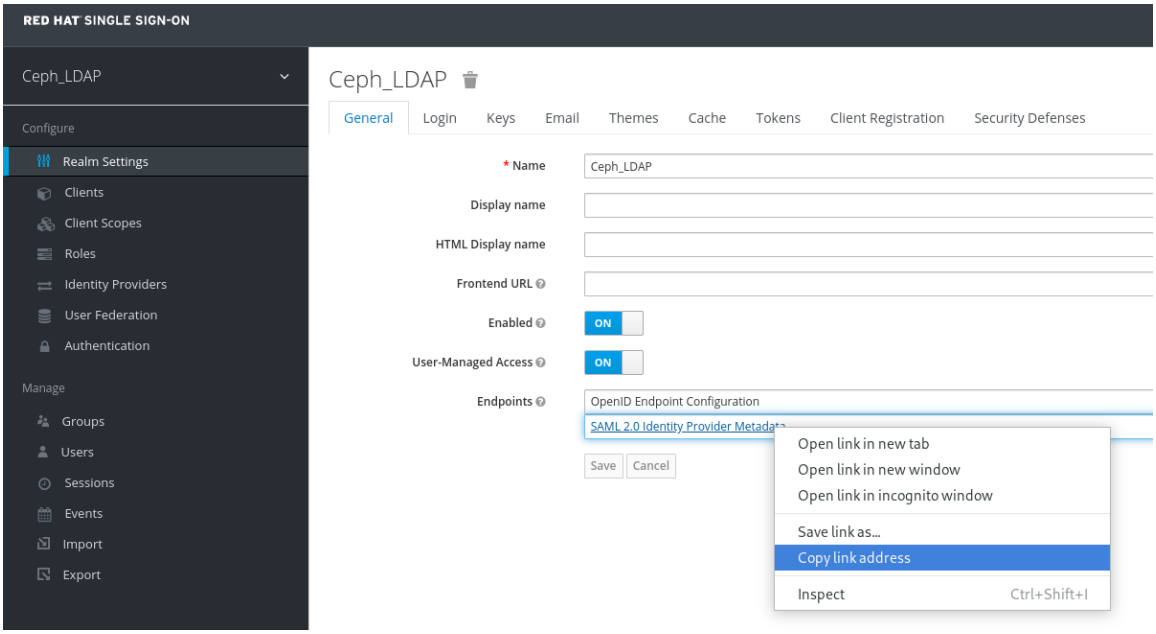

12. In the *Clients* tab, click Create:

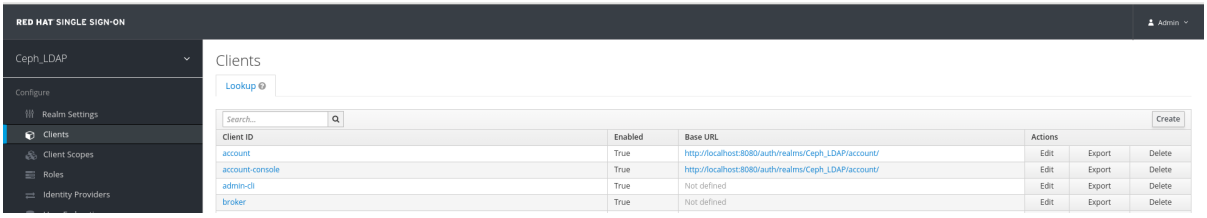

13. In the *Add Client* window, set the following parameters and click Save:

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

Example

https://magna082.ceph.redhat.com:8443/auth/saml2/metadata

b. Client Protocol - saml

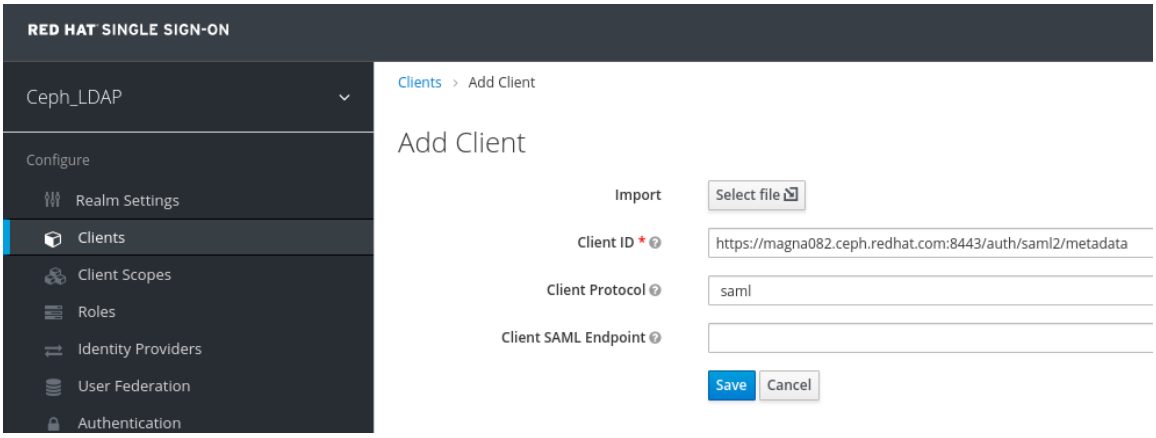

14. In the *Clients* window, under *Settings* tab, set the following parameters and click Save:

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

Example

https://magna082.ceph.redhat.com:8443/auth/saml2/metadata

b. Enabled - ON

c. Client Protocol - saml

d. Include AuthnStatement - ON

e. Sign Documents - ON

f. Signature Algorithm - RSA_SHA1

g. SAML Signature Key Name - KEY_ID

h. Valid Redirect URLs - BASE_URL:8443/*

Example

https://magna082.ceph.redhat.com:8443/*

i. Base URL - BASE_URL:8443

Example

https://magna082.ceph.redhat.com:8443/

j. Master SAML Processing URL - http://localhost:8080/auth/realms/*REALM_NAME*/protocol/saml/descriptor

Example

http://localhost:8080/auth/realms/Ceph_LDAP/protocol/saml/descriptor

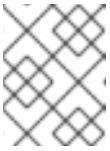

NOTE

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

Under Fine Grain SAML Endpoint Configuration, set the parameters:

k. Assertion Consumer Service POST Binding URL - BASE_URL:8443/#/dashboard

Example

https://magna082.ceph.redhat.com:8443/#/dashboard

l. Assertion Consumer Service Redirect Binding URL - BASE_URL:8443/#/dashboard

Example

https://magna082.ceph.redhat.com:8443/#/dashboard

m. Logout Service Redirect Binding URL - BASE_URL:8443/

Example

https://magna082.ceph.redhat.com:8443/

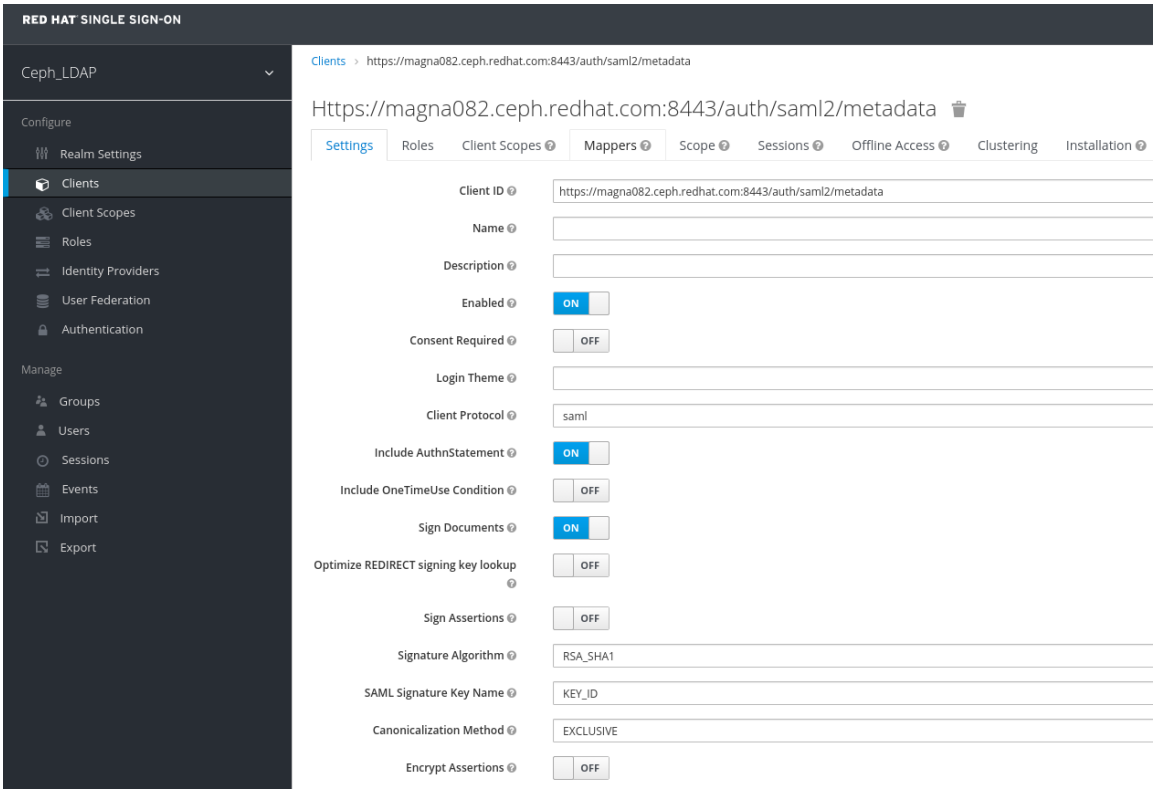

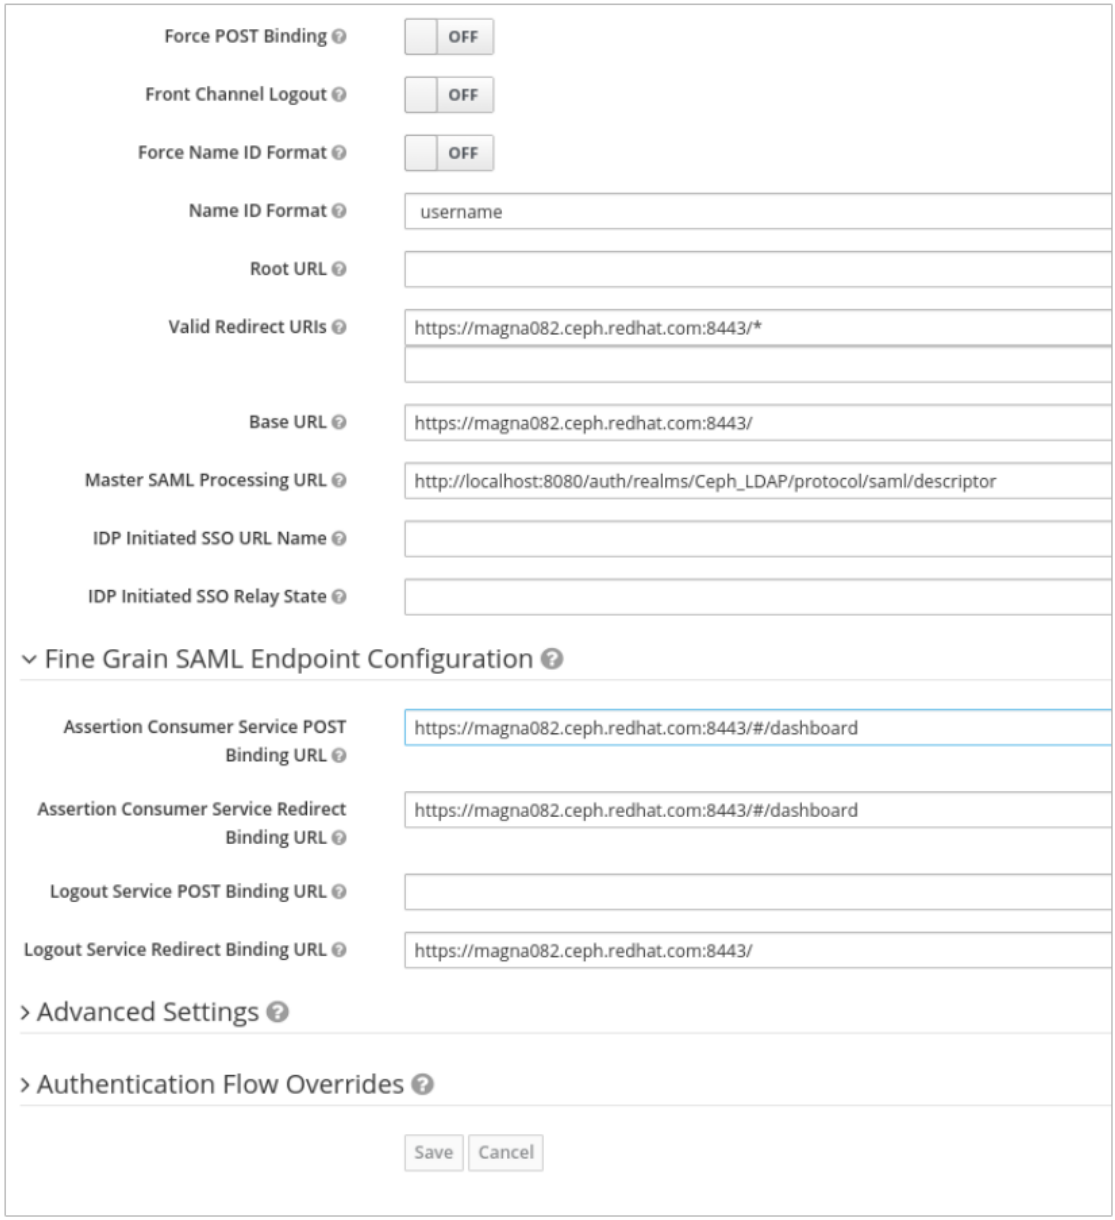

15. In the *Clients* window, *Mappers* tab, set the following parameters and click Save:

    a. Protocol - saml

    b. Name - username

    c. Mapper Property - User Property

    d. Property - username

    e. SAML Attribute name - username

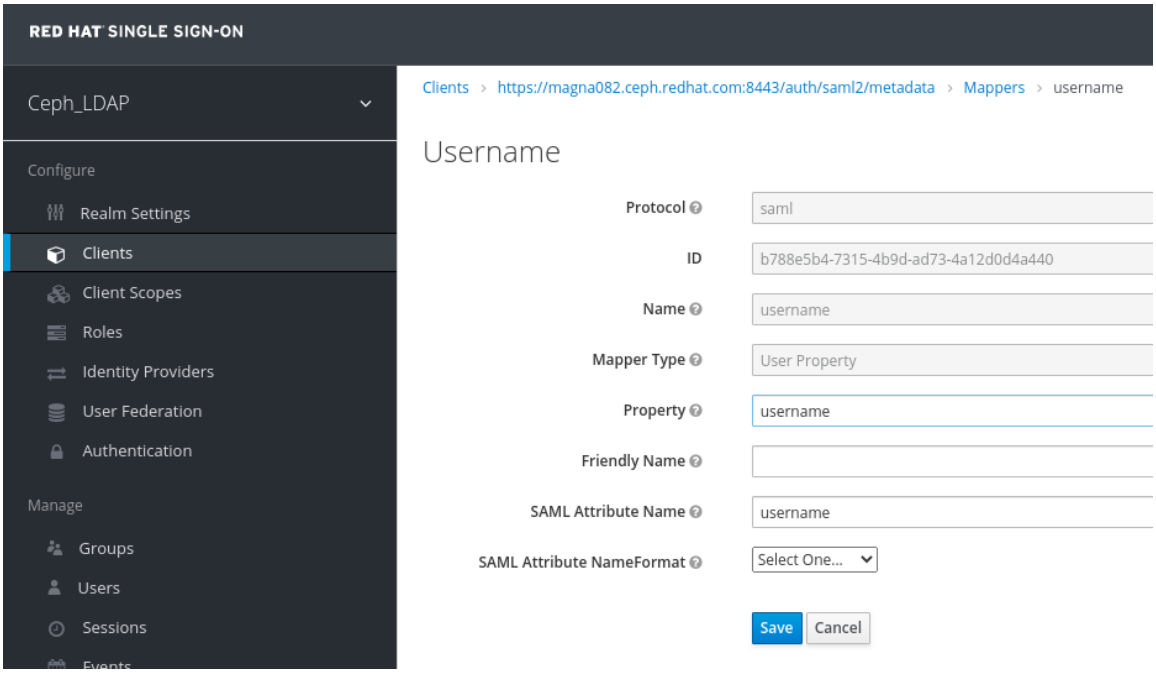

16. In the *Clients Scope* tab, select *role_list*:

    a. In *Mappers* tab, select *role list*, set the *Single Role Attribute* to ON.

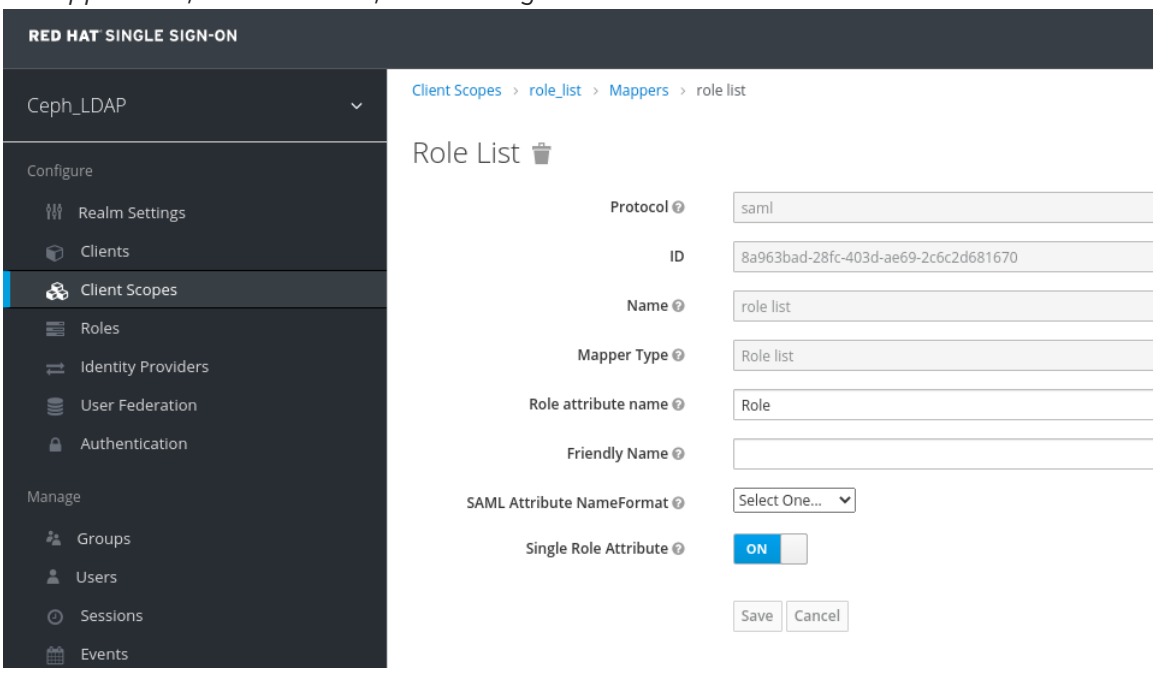

17. Select *User_Federation* tab:

    a. In *User Federation* window, select *ldap* from the drop-down:

User Federation

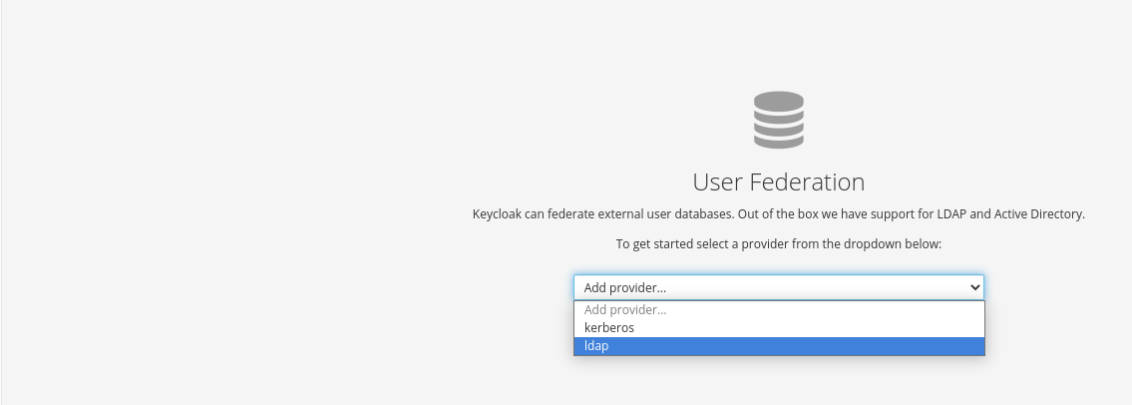

18. In *User_Federation* window, *Settings* tab, set the following parameters and click Save:

    a. Console Display Name - rh-ldap

    b. Import Users - ON

    c. Edit_Mode - READ_ONLY

    d. Username LDAP attribute - username

    e. RDN LDAP attribute - username

    f. UUID LDAP attribute - nsuniqueid

    g. User Object Classes - inetOrgPerson, organizationalPerson, rhatPerson

    h. Connection URL - ldap::://myldap.example.com

    Example

    ldap://ldap.corp.redhat.com

    Click *Test Connection*.

    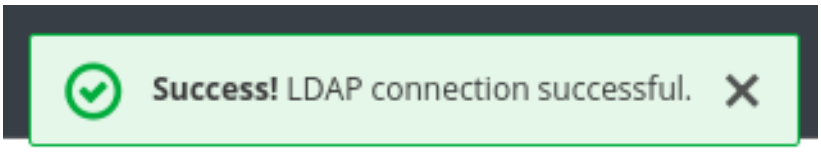

    You will get a notification that the LDAP connection is successful.

    i. Users DN - ou=users, dc=example, dc=com

    Example

    ou=users,dc=redhat,dc=com

    j. Bind Type - simple

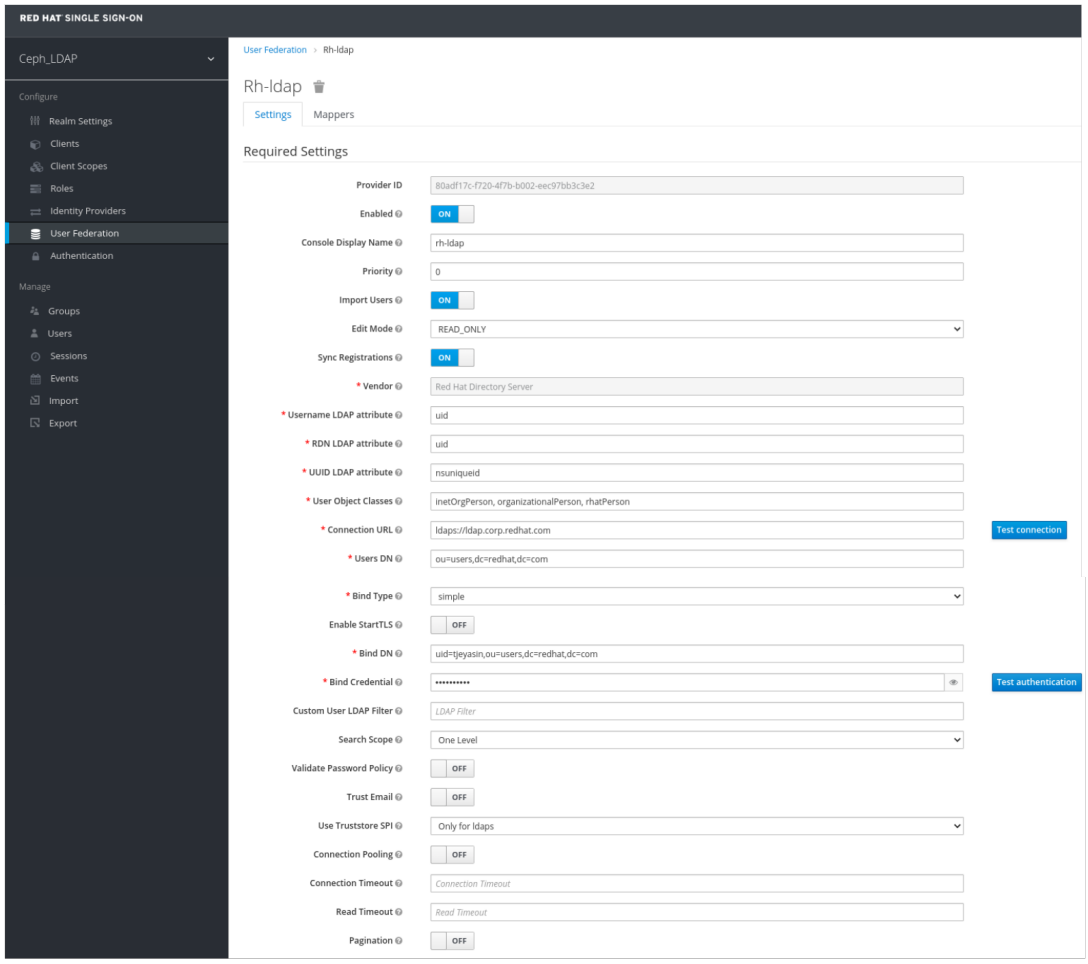

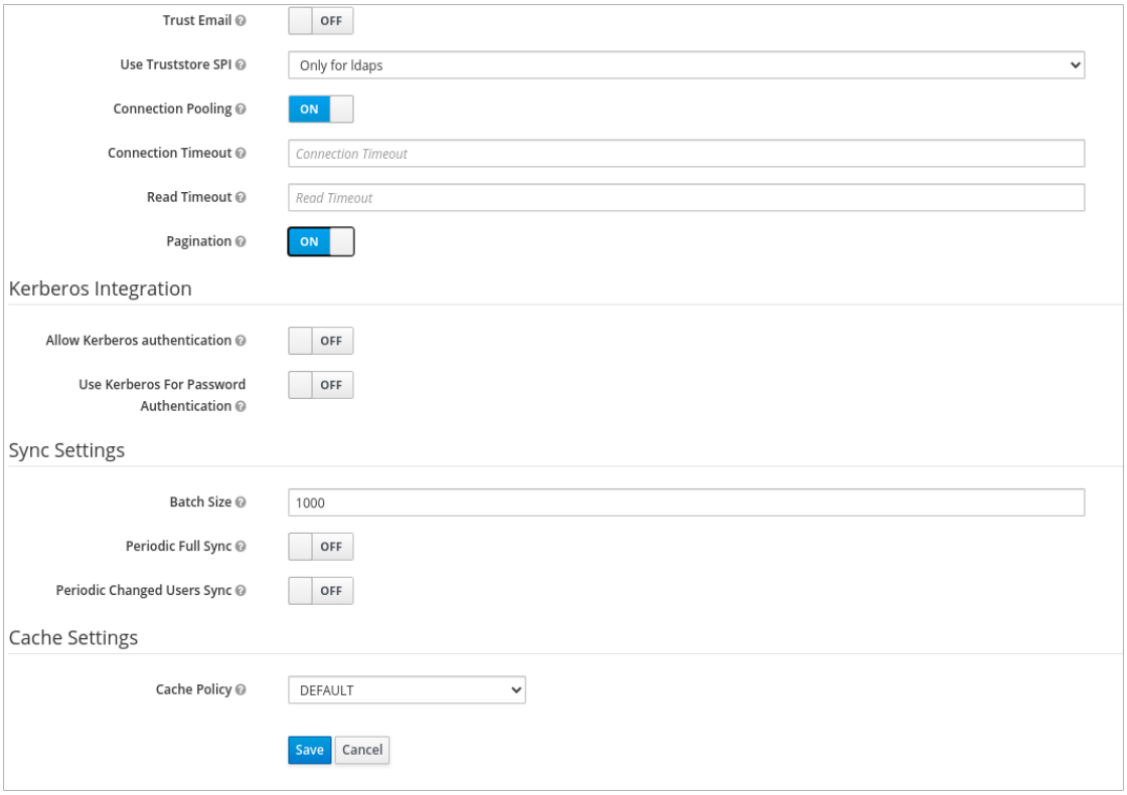

k. Click *Test authentication*.

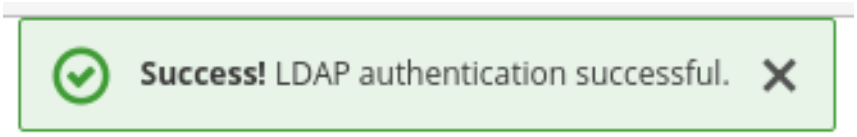

You will get a notification that the LDAP authentication is successful.

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

    a. LDAP Attribute - givenName

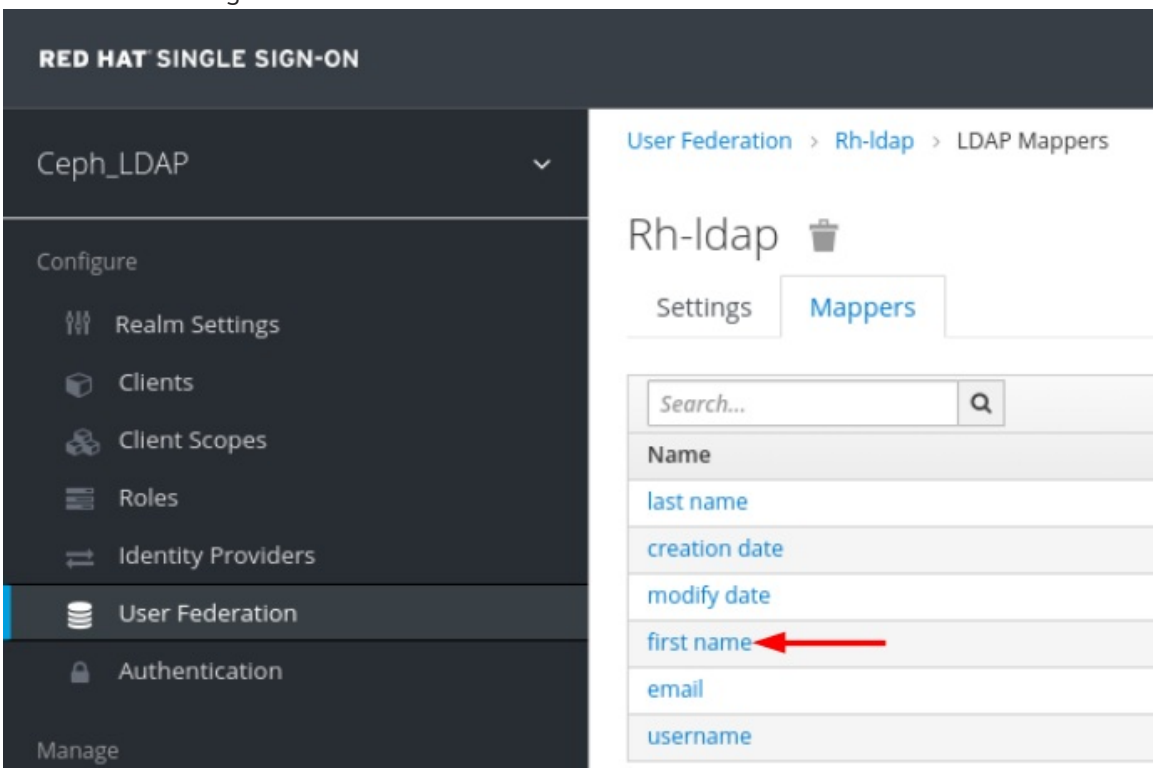

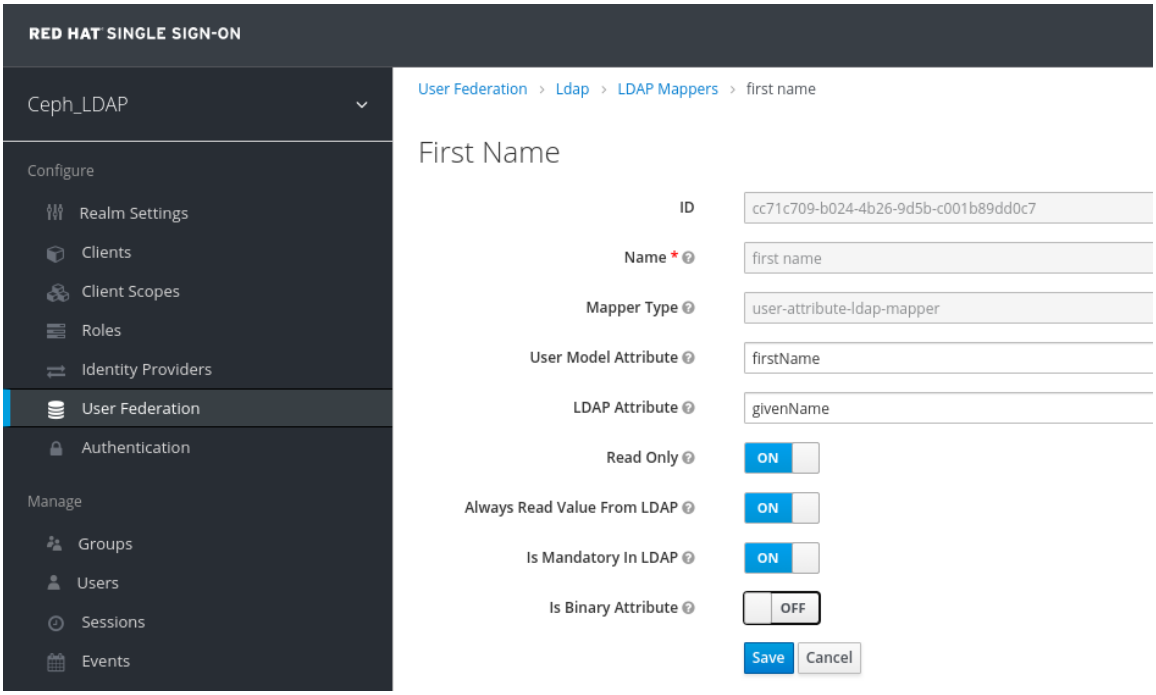

20. In *User_Federation* tab, Settings tab, Click *Synchronize all users*:

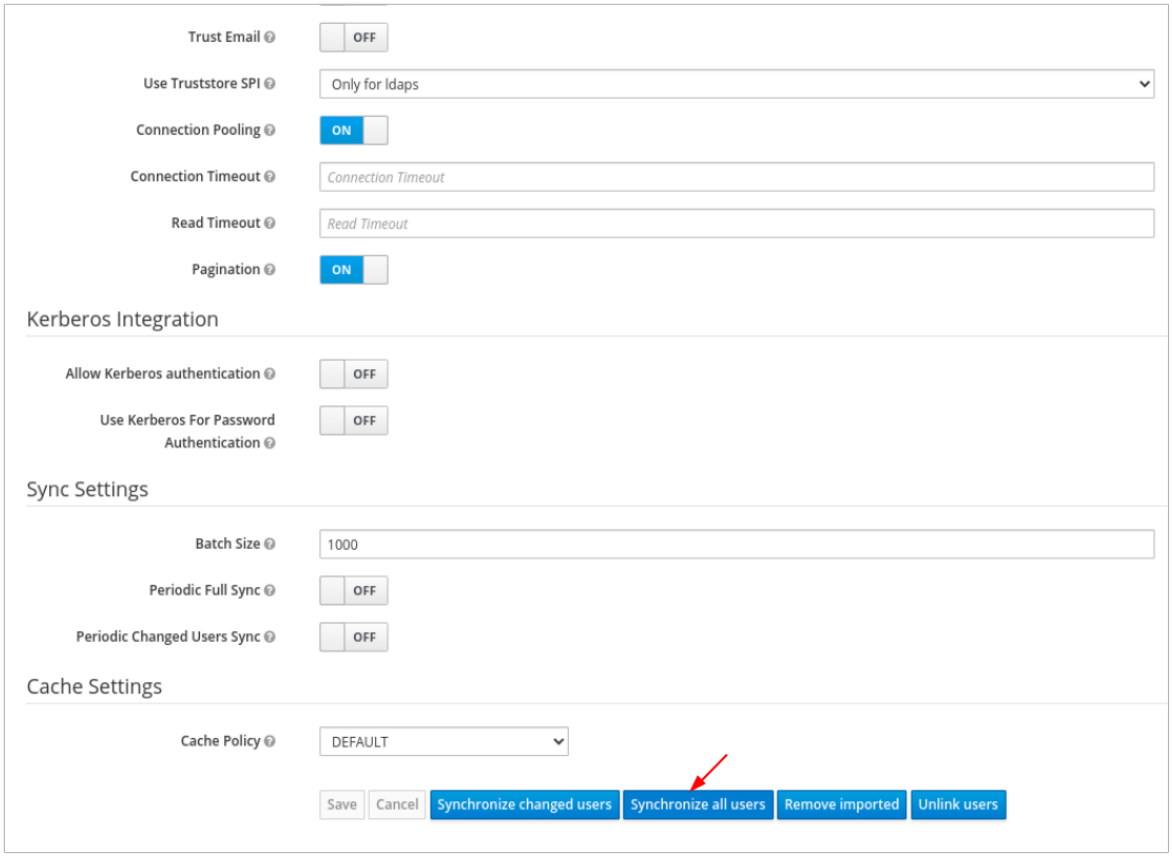

You will get a notification that the sync of users are updated successfully.

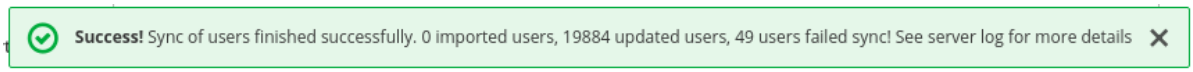

21. In the *Users* tab, search for the user added to the dashboard and click the Search icon:

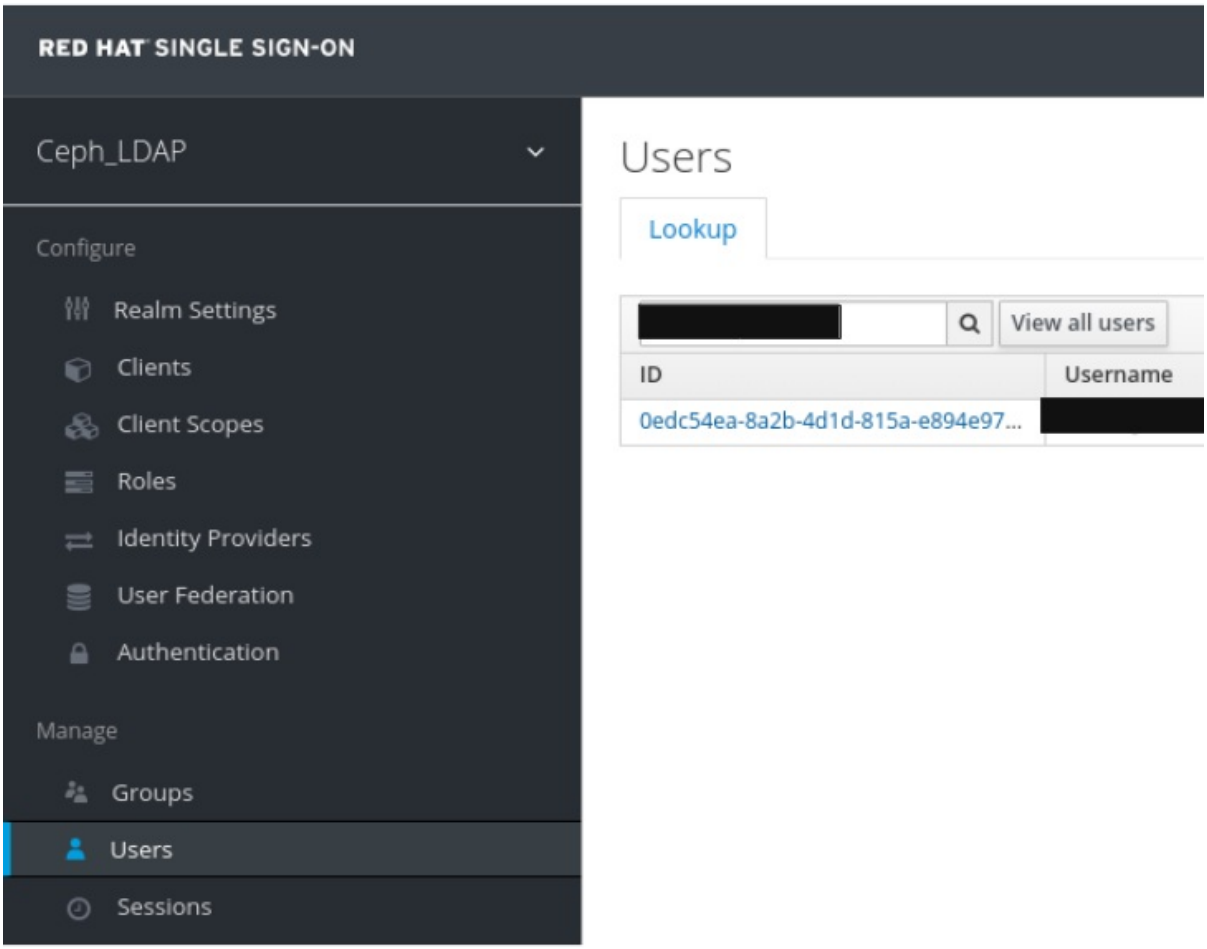

22. To view the user , click it's row. You should see the federation link as the name provided for the *User Federation*.

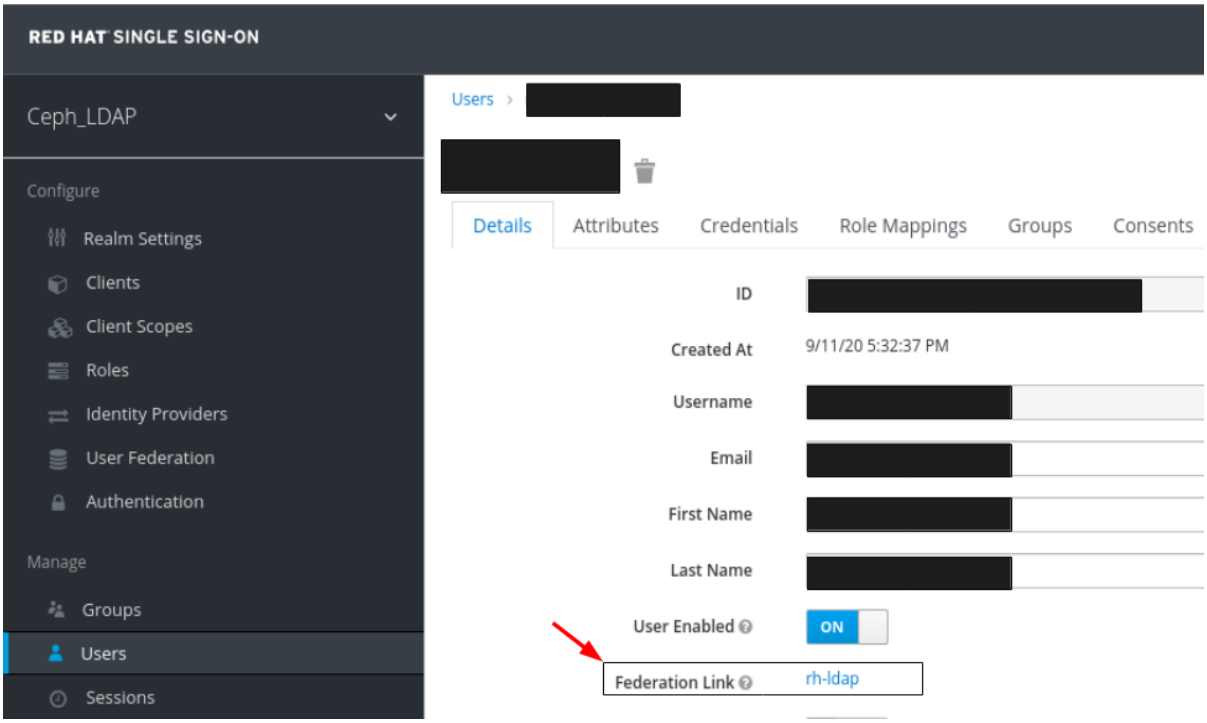

IMPORTANT

Do not add users manually. If added manually, delete the user by clicking *Delete*.

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

    Example

    https://magna082.ceph.redhat.com:8443

Additional Resources

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

2.10. ENABLING SINGLE SIGN-ON FOR THE CEPH DASHBOARD

The Ceph Dashboard supports external authentication of users with the Security Assertion Markup Language (SAML) protocol. Before using single sign-On (SSO) with the Ceph dashboard, create the dashboard user accounts and assign the desired roles. The Ceph Dashboard performs authorization of the users and the authentication process is performed by an existing Identity Provider (IdP). Red Hat uses Keycloak to test the dashboard SSO feature.

Prerequisites

- A running Red Hat Ceph Storage cluster.

- Installation of the Ceph Dashboard software.
- Launch the Dashboard.
- Root-level access to the Ceph Manager nodes.
- Installation of the following library packages on the Ceph Manager nodes:
  - **python3-saml**
  - **python3-defusedxml**
  - **python3-isodate**
  - **python3-xmlsec**

Procedure

1. To configure SSO on Ceph Dashboard, run the following command:
   a. **Bare-metal** deployments:

      Syntax

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

      Example

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

   b. **Container** deployments:

      Syntax

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

      Example

      ```
      [root@mon ~]# podman exec ceph-mgr-hostname ceph dashboard sso setup saml2
      http://dashboard_hostname.ceph.redhat.com:8443 idp-metadata.xml username
      http://10.70.59.125:8080/auth/realms/realm_name /home/certificate.txt /home/private-
      key.txt
      ```

   Replace

   - *CEPH_MGR_NODE* with Ceph **mgr** node. For example, **ceph-mgr-hostname**
   - *CEPH_DASHBOARD_BASE_URL* with the base URL where Ceph Dashboard is accessible.

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

2. Verify the current SAML 2.0 configuration:
   a. **Bare-metal** deployments:

      Syntax

      ```
      ceph dashboard sso show saml2
      ```

   b. **Container** deployments:

      Syntax

      ```
      podman exec CEPH_MGR_NODE ceph dashboard sso show saml2
      ```

3. To enable SSO, run the following command:
   a. **Bare-metal** deployments:

      Syntax

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

   b. **Container** deployments:

      Syntax

      ```
      podman exec CEPH_MGR_NODE ceph dashboard sso enable saml2
      SSO is "enabled" with "SAML2" protocol.
      ```

4. Open your dashboard URL. For example:

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

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

Additional Resources

- To disable single sign-on, see *Disabling Single Sign-on for the Ceph Dashboard* in the *Red Hat Ceph StorageDashboard Guide*.

2.11. DISABLING SINGLE SIGN-ON FOR THE CEPH DASHBOARD

You can disable single sign on for Ceph Dashboard.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Installation of the Ceph Dashboard software.
- Launch the Dashboard.
- Root-level access to the Ceph Manager nodes.
- Single sign-on enabled for Ceph Dashboard
- Installation of the following library packages on the Ceph Manager nodes:
  - **python3-saml**
  - **python3-defusedxml**
  - **python3-isodate**
  - **python3-xmlsec**

Procedure

1. To view status of SSO, run the following command:
   a. **Bare-metal** deployments:

      #### Syntax

      ```
      ceph dashboard sso status
      SSO is "enabled" with "SAML2" protocol.
      ```

   b. **Container** deployments:

      #### Syntax

      ```
      podman exec CEPH_MGR_NODE ceph dashboard sso status
      SSO is "enabled" with "SAML2" protocol.
      ```

      #### Replace

      - *CEPH_MGR_NODE* with Ceph **mgr** node. For example, **ceph-mgr-hostname**

2. To disable SSO, run the following command:
   a. **Bare-metal** deployments:

      #### Syntax

```
ceph dashboard sso disable
SSO is "disabled".
```

b. **Container** deployments:

Syntax

```
podman exec CEPH_MGR_NODE ceph dashboard sso disable
SSO is "disabled".
```

Replace

- *CEPH_MGR_NODE* with Ceph **mgr** node. For example, **ceph-mgr-hostname**

Additional Resources

- To enable single sign-on, see *Enabling Single Sign-on for the Ceph Dashboard* in the *Red Hat Ceph StorageDashboard Guide*.

CHAPTER 3. MANAGING ROLES ON DASHBOARD

As a storage administrator, you can create, view, edit, and delete roles on the dashboard. You can give permissions to the roles and you can assign specific roles for users.

3.1. CREATING ROLES ON DASHBOARD

The dashboard allows you to create roles on the dashboard which can be assigned to the users.

Prerequisites

- 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. On the upper right side of the Dashboard, click the gear icon and select *User management*:

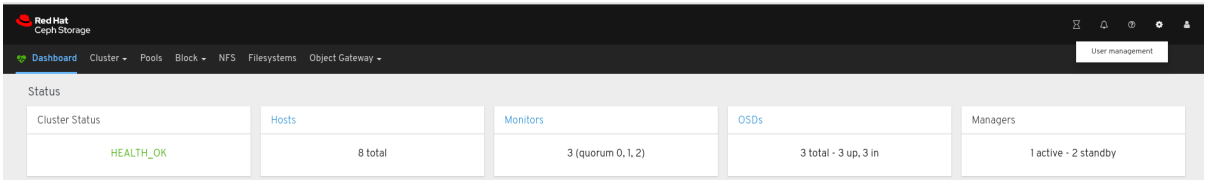

3. On *Roles* tab, click the *Create* button:

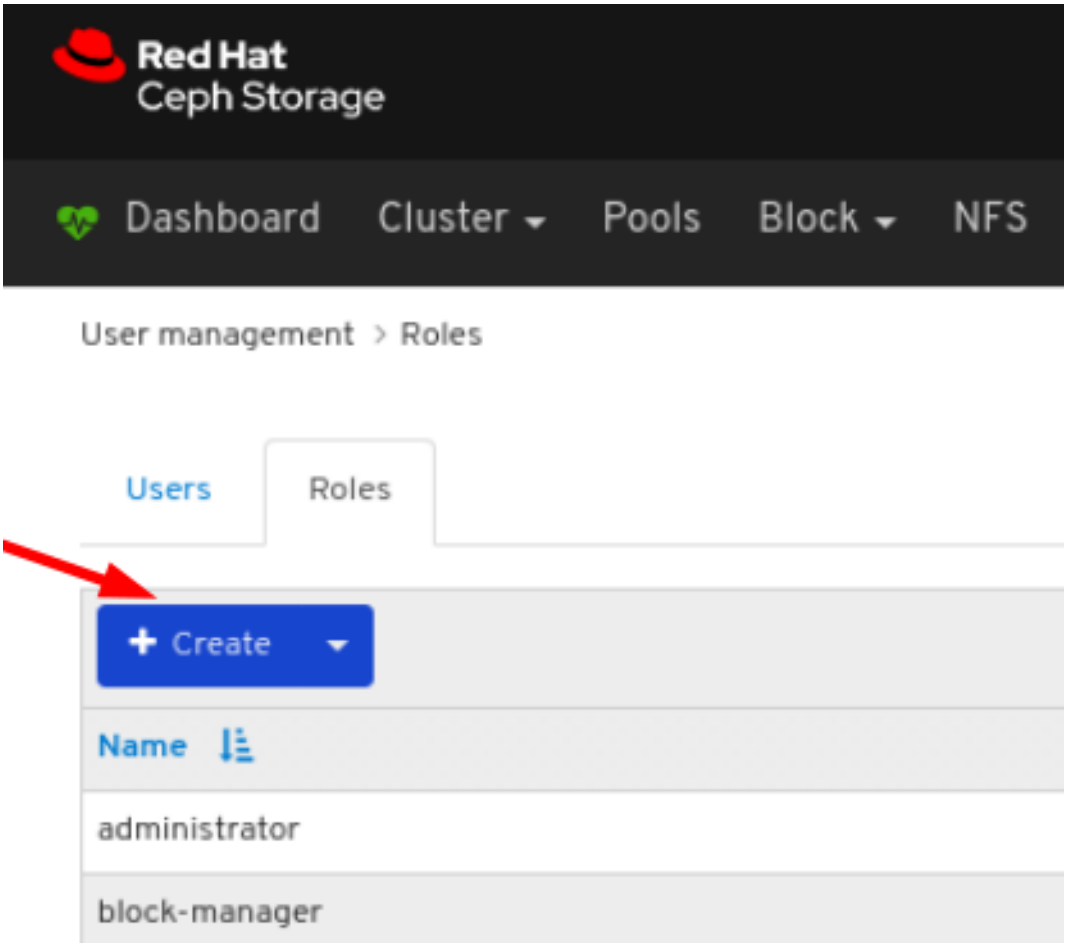

4. In the *CreateRole* window, set the *Name* and select the *Permissions* for this role, and then click the *CreateRole* button:

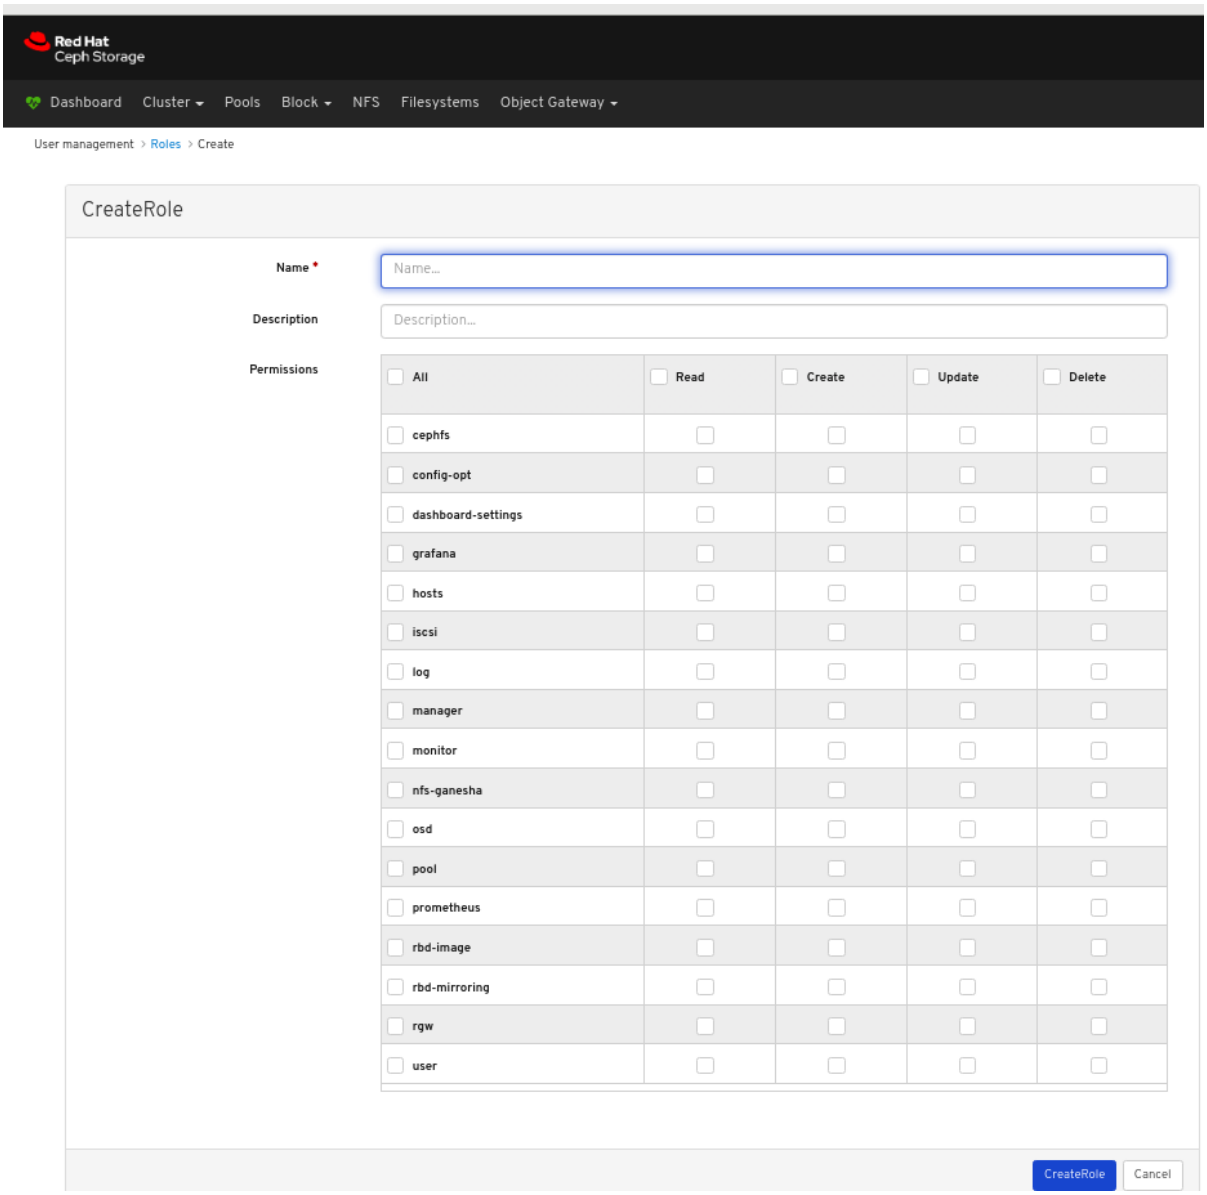

5. A notification towards the top right corner of the page indicates the role was created successfully.

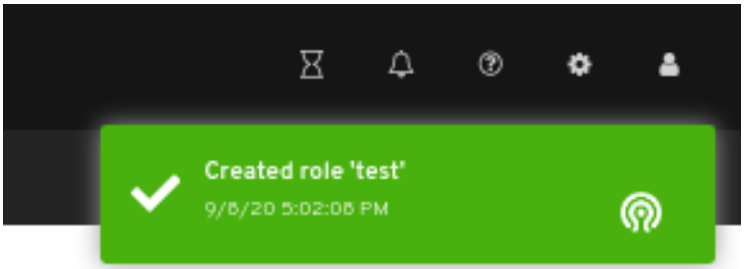

3.2. VIEWING ROLES ON DASHBOARD

The dashboard allows you to view the details of the roles on the 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. Log in to the Dashboard.
2. On the upper right side of the Dashboard, click the gear icon and select *User management*:

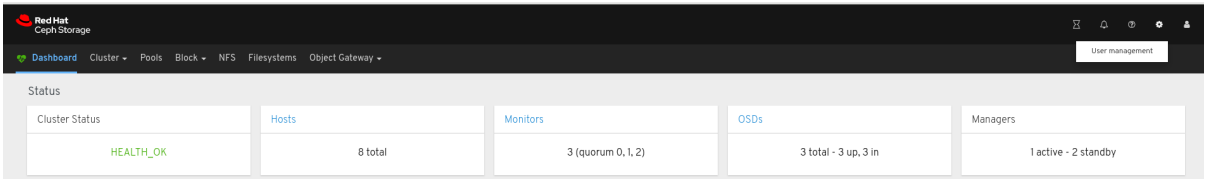

3. To view the details of the role, click the row:

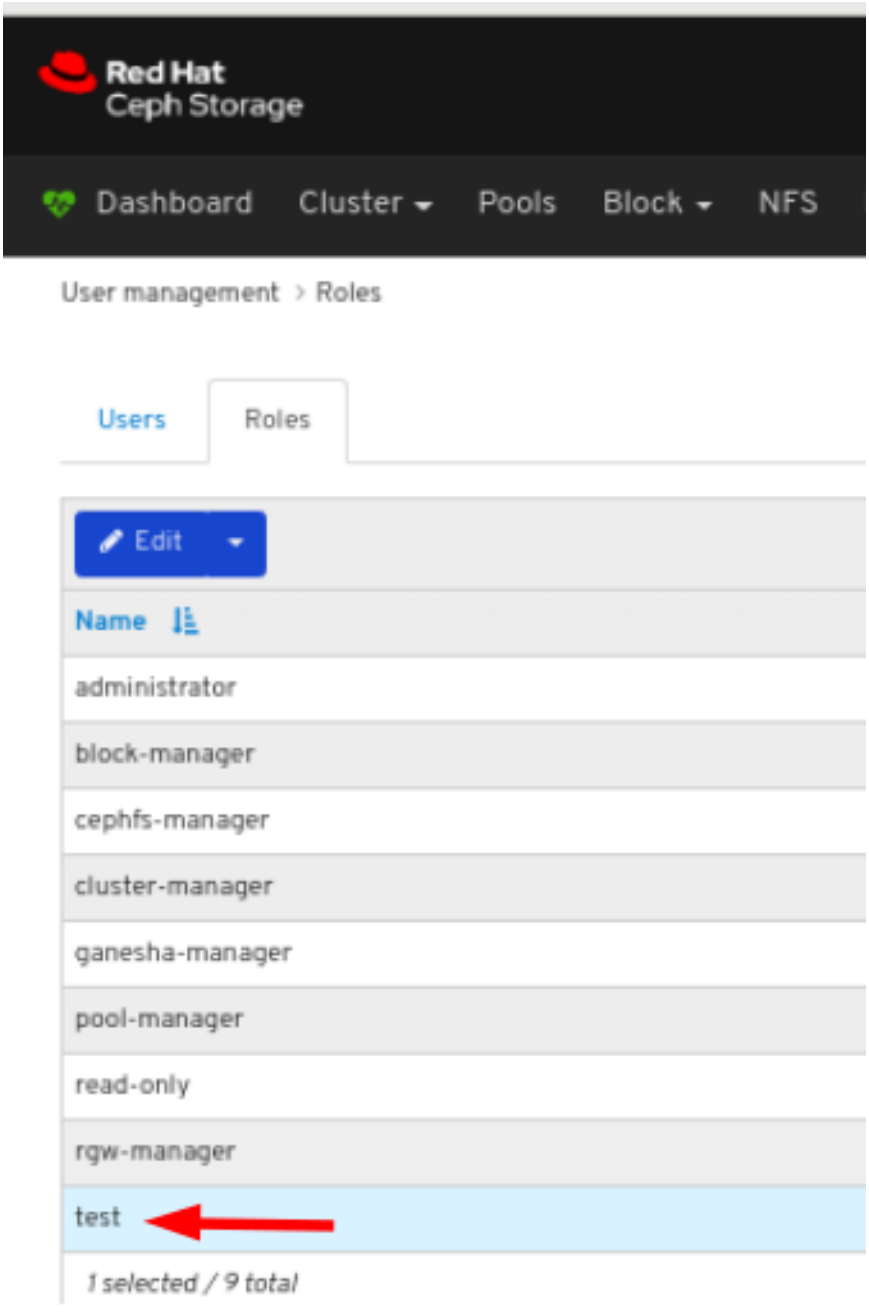

4. You can see the details of the permissions provided for the roles.

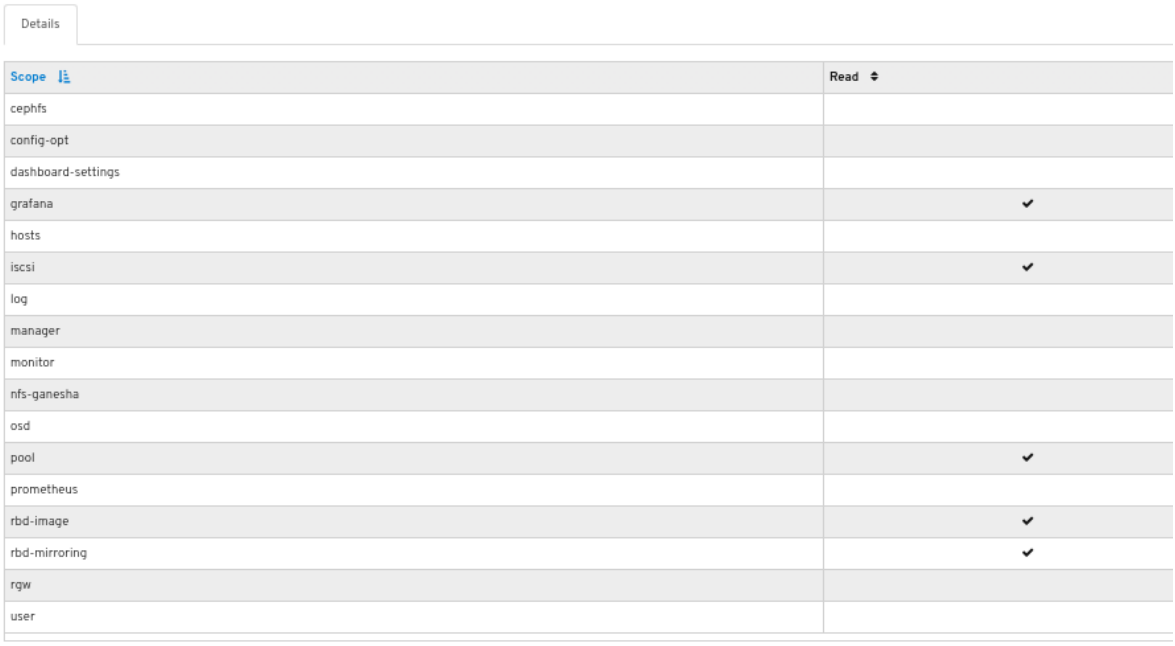

Additional Resources

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

3.3. EDITING ROLES ON DASHBOARD

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. On the upper right side of the Dashboard, click the gear icon and select *User management*:

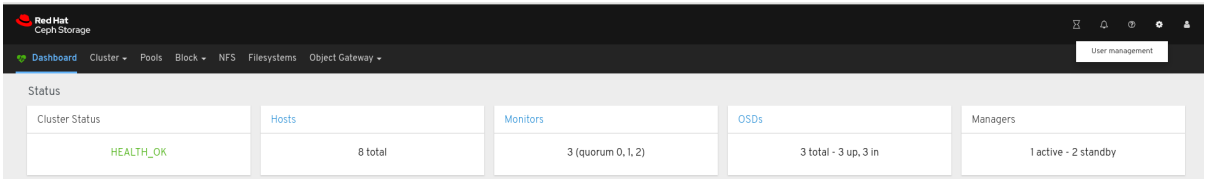

3. To edit the role, click the row:

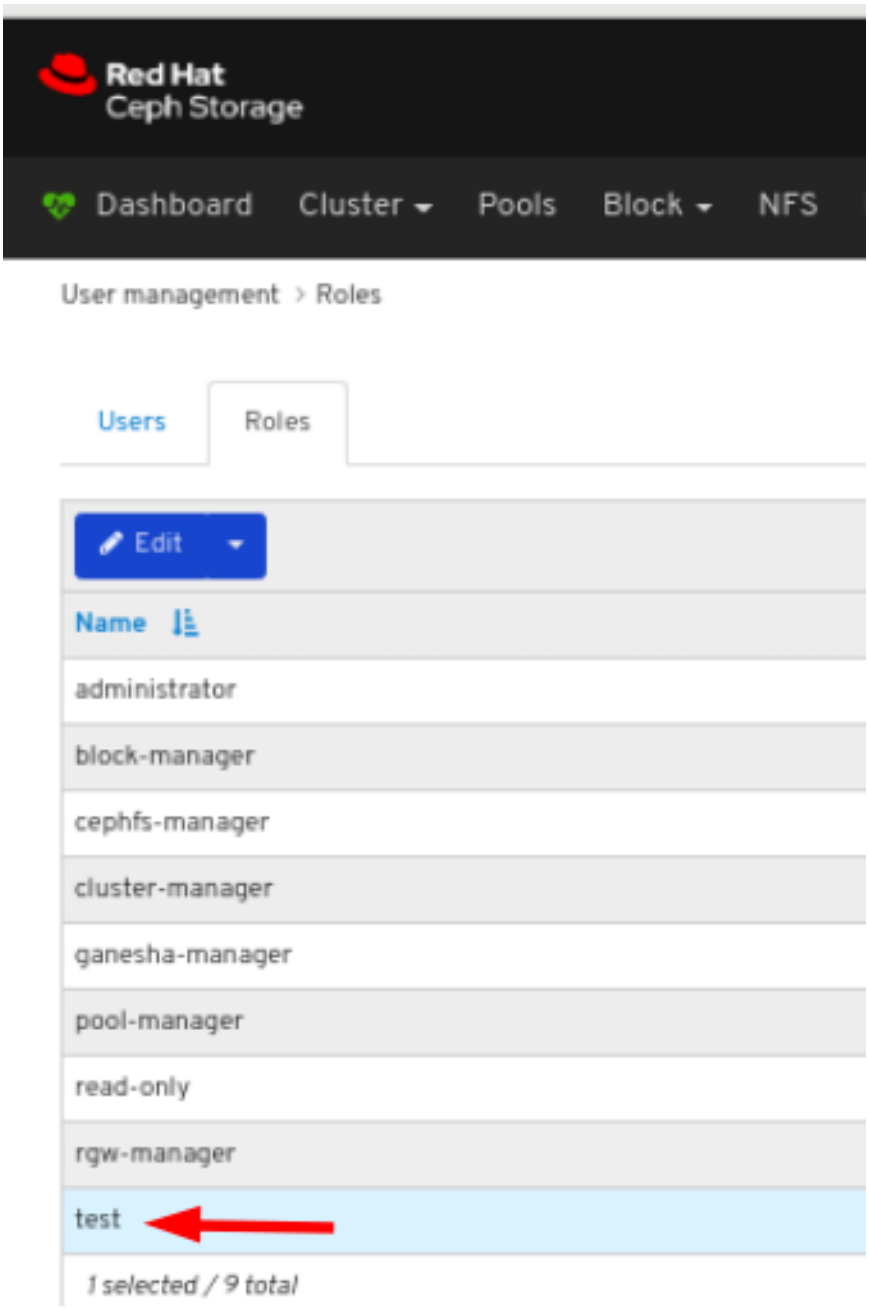

4. On the *Roles* tab, select *Edit* from the *Edit* dropdown menu:

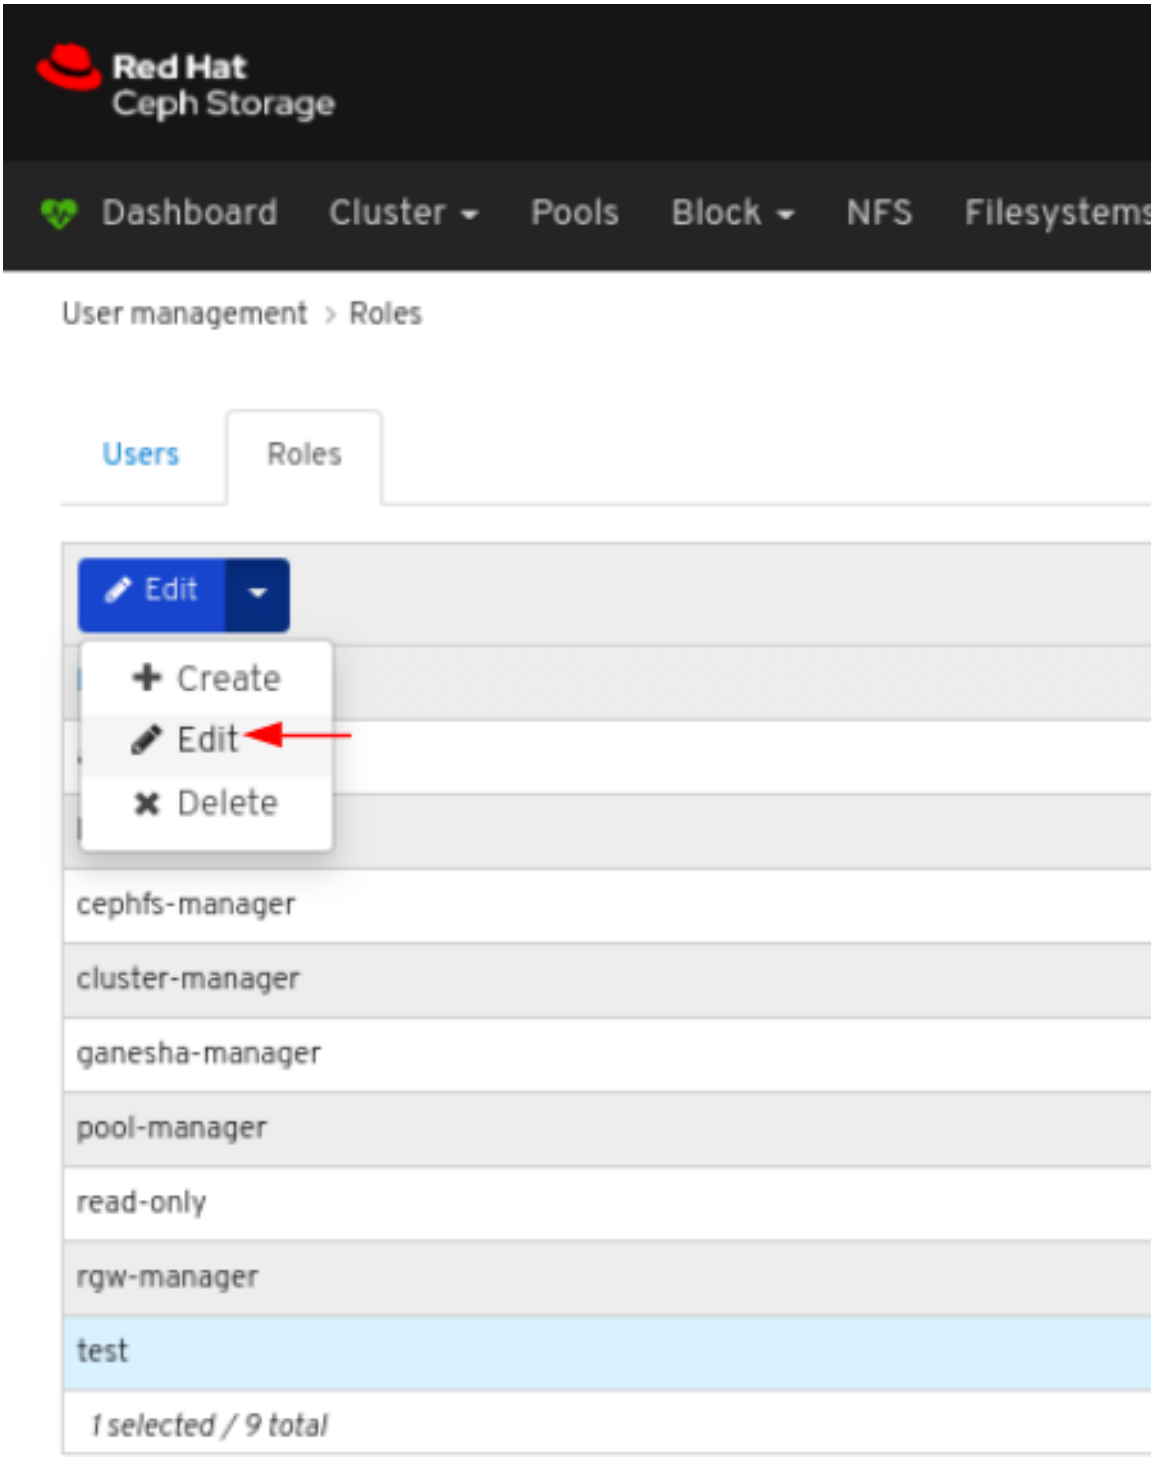

5. In the *EditRole* window, edit parameters including, and then click the *EditRole* button:

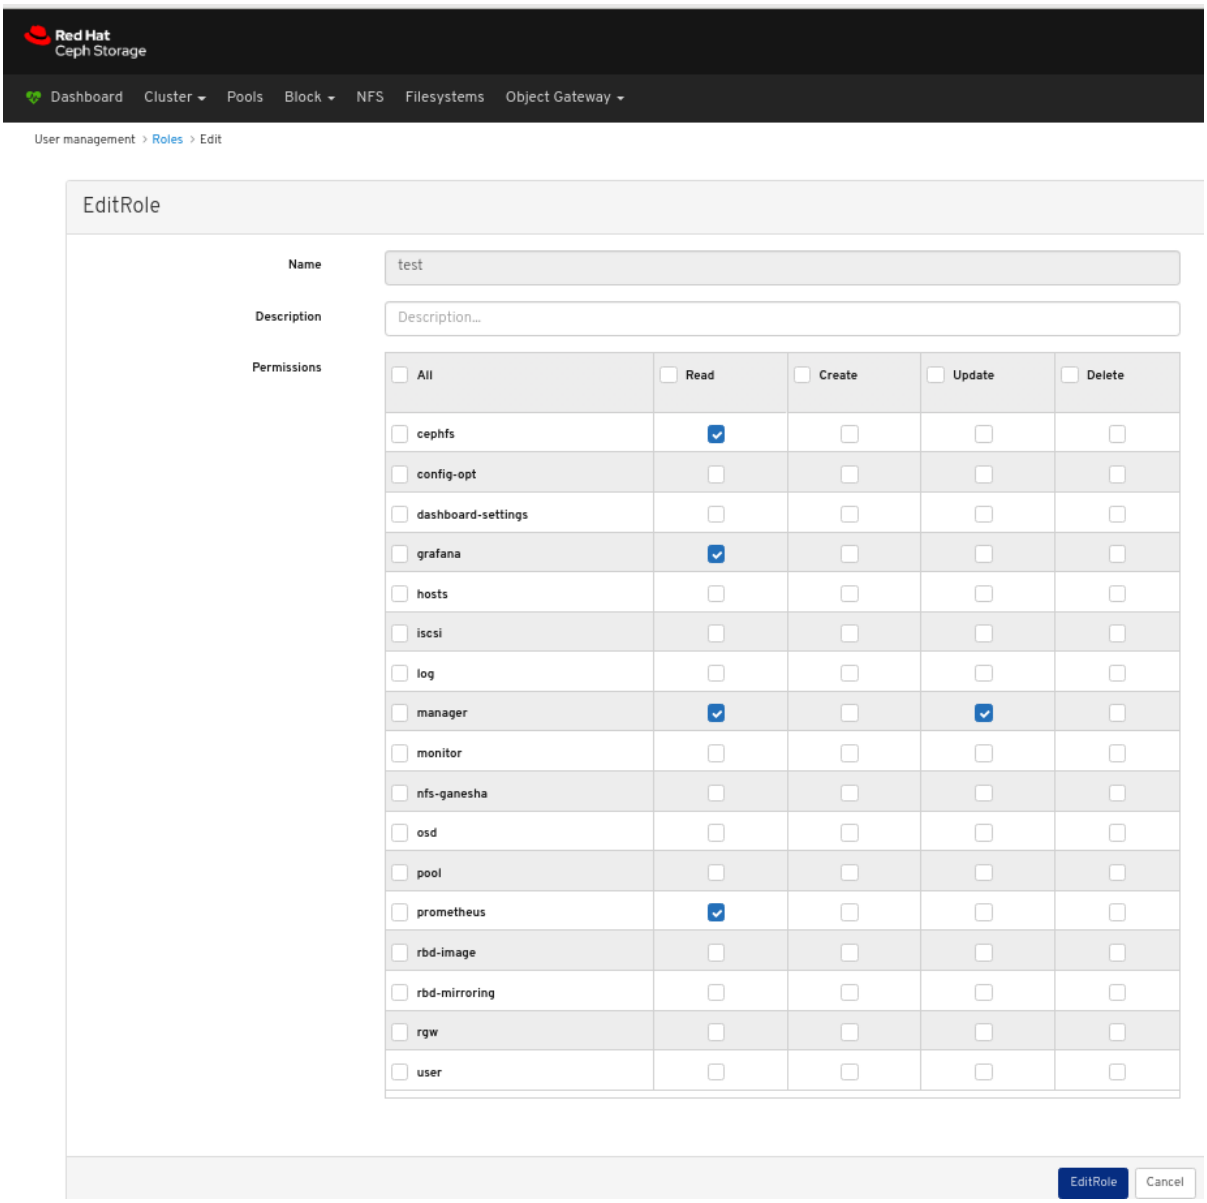

6. A notification towards the top right corner of the page indicates the role was updated successfully.

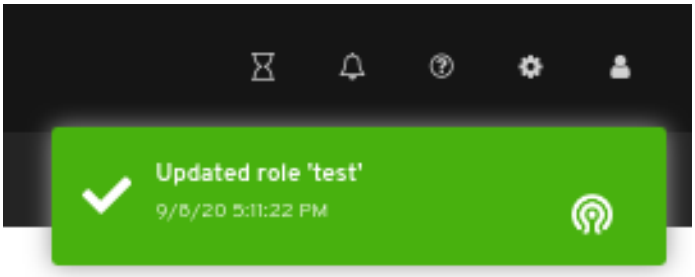

Additional Resources

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

3.4. DELETING ROLES ON DASHBOARD

The dashboard allows you to delete roles on the dashboard.

NOTE

You can only delete the roles that you have created.

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. On the upper right side of the Dashboard, click the gear icon and select *User management*:

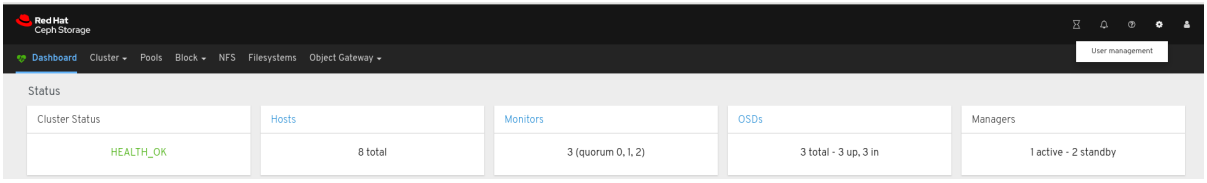

3. To delete the role, click the row:

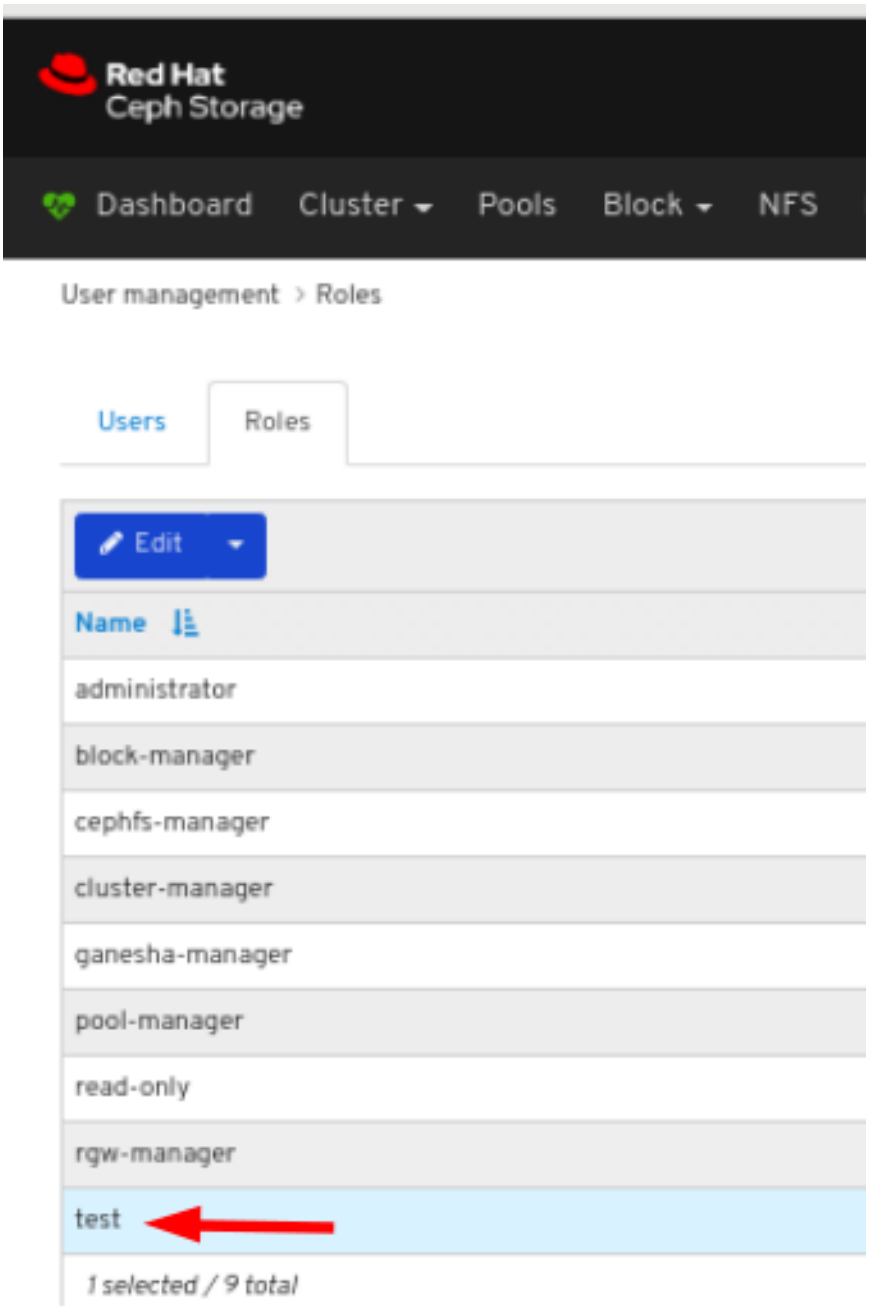

4. On *Roles* tab, select *Delete* from the *Edit* dropdown menu:

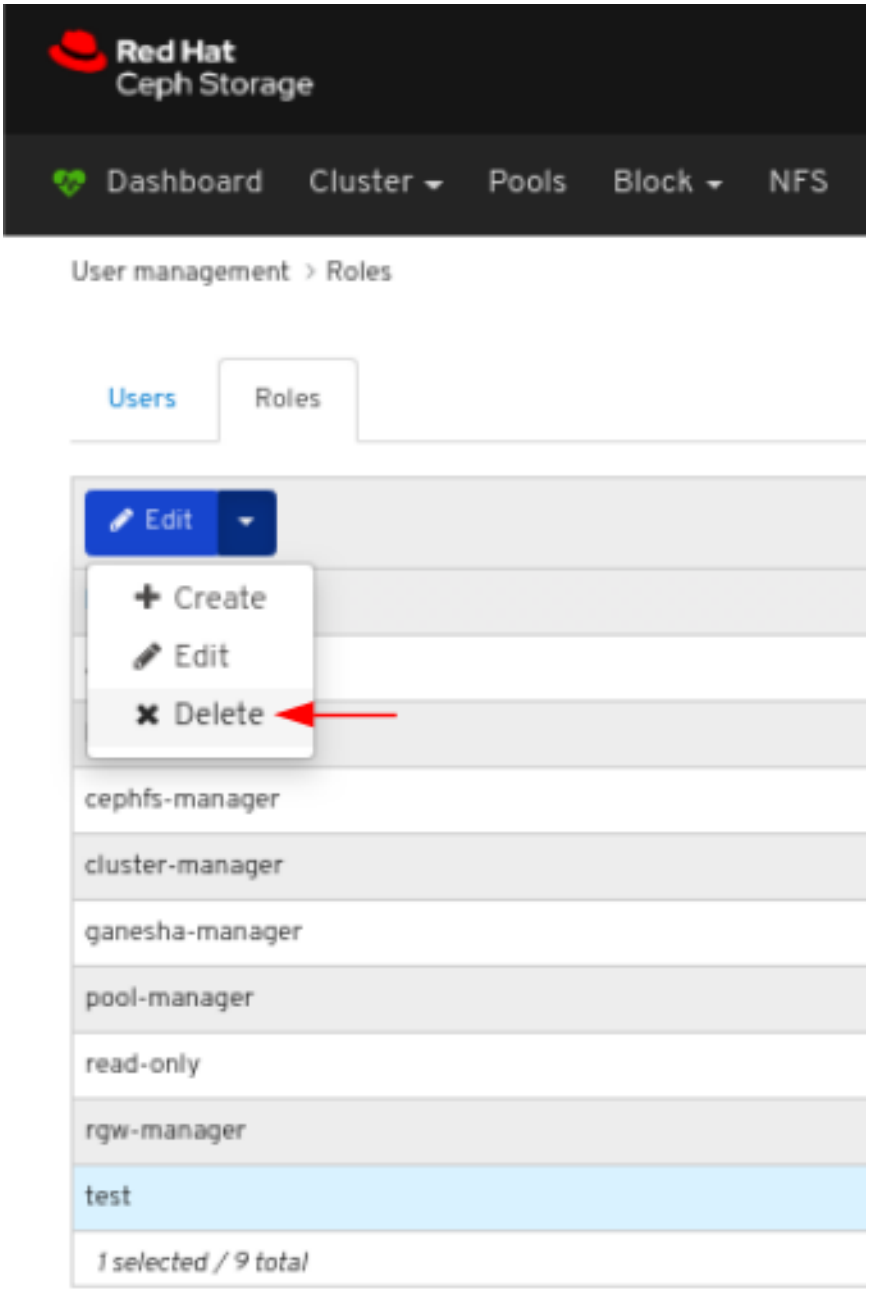

5. In the *Delete Role* dialog window, Click the *Yes, I am sure* box and then Click *Delete Role* to save the settings:

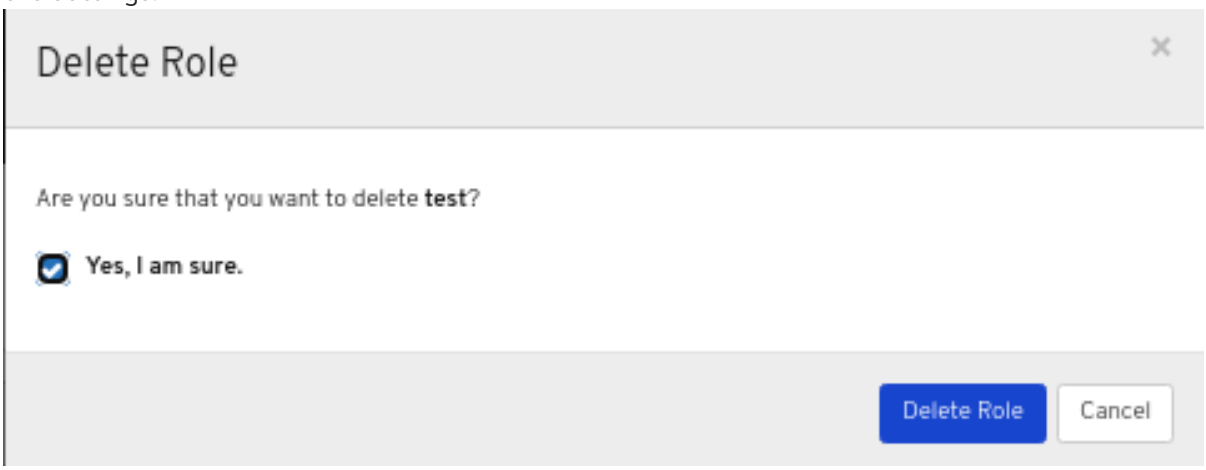

Additional Resources

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

CHAPTER 4. MANAGING USERS ON DASHBOARD

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

4.1. CREATING USERS ON DASHBOARD

The dashboard allows you to create users on the dashboard.

Prerequisites

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

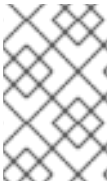

NOTE

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. On the upper right side of the Dashboard, click the gear icon and select *User management*:

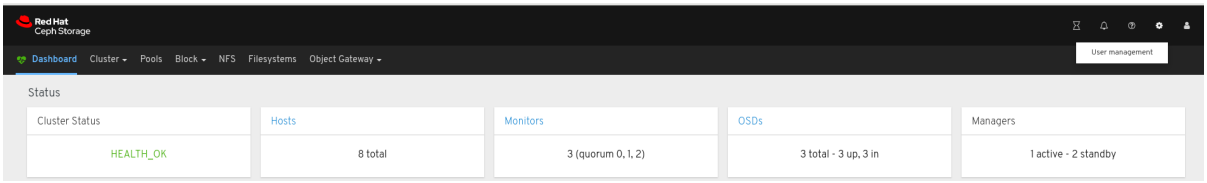

3. On *Users* tab, click the *Create* button:

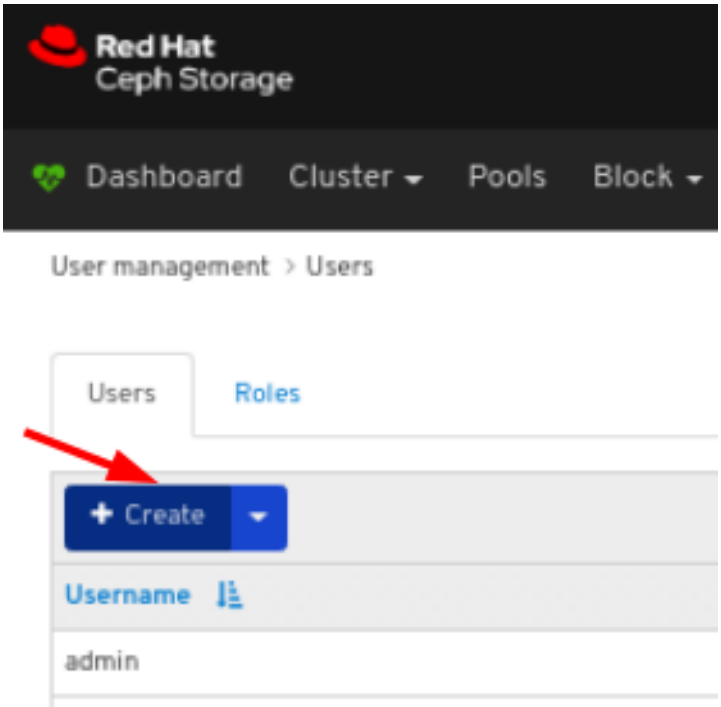

4. In the *CreateUser* window, set the *Username* and other parameters including the roles, and then click the _CreateUser_button:

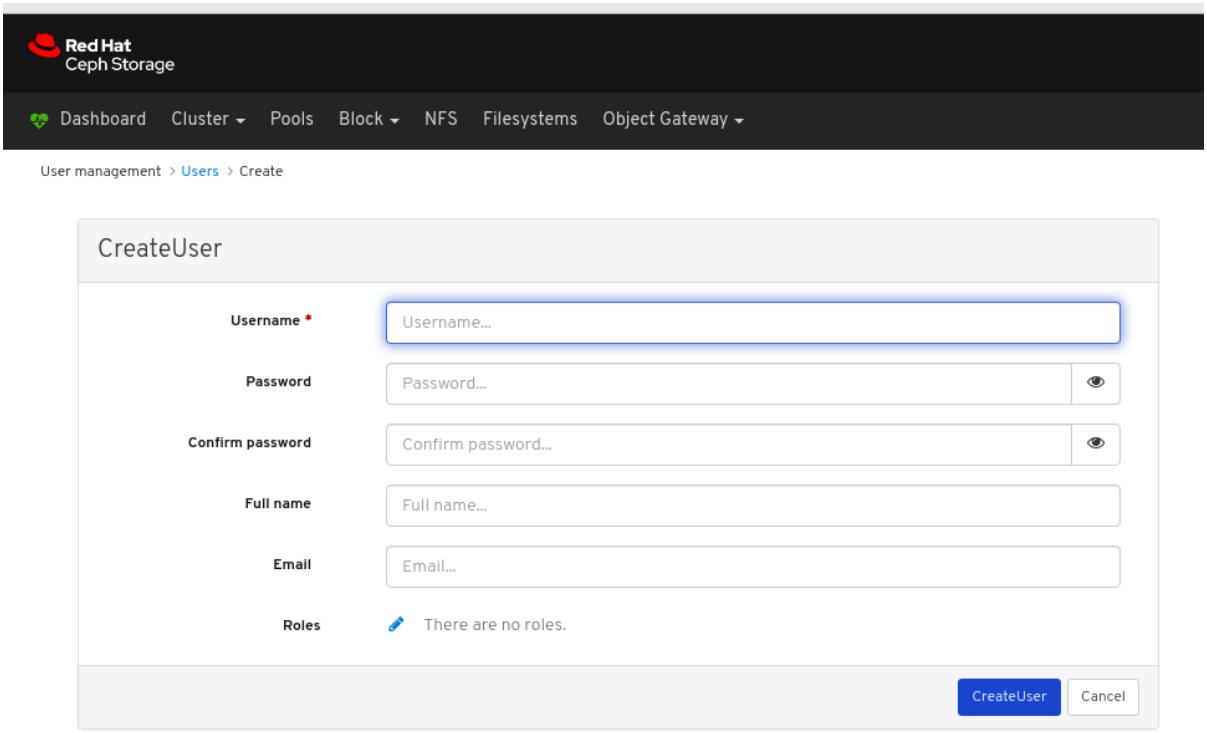

5. A notification towards the top right corner of the page indicates the user was created successfully.

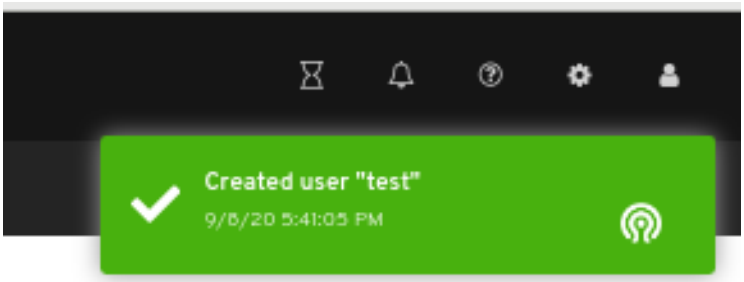

Additional Resources

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

4.2. EDITING USERS ON DASHBOARD

The dashboard allows you to edit users on the dashboard.

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. On the upper right side of the Dashboard, click the gear icon and select *User management*:

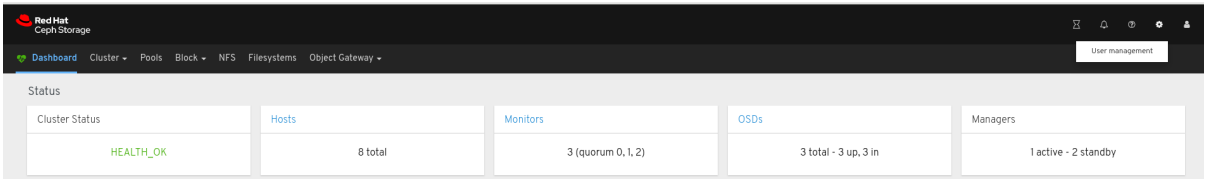

3. To edit the user, click the row:

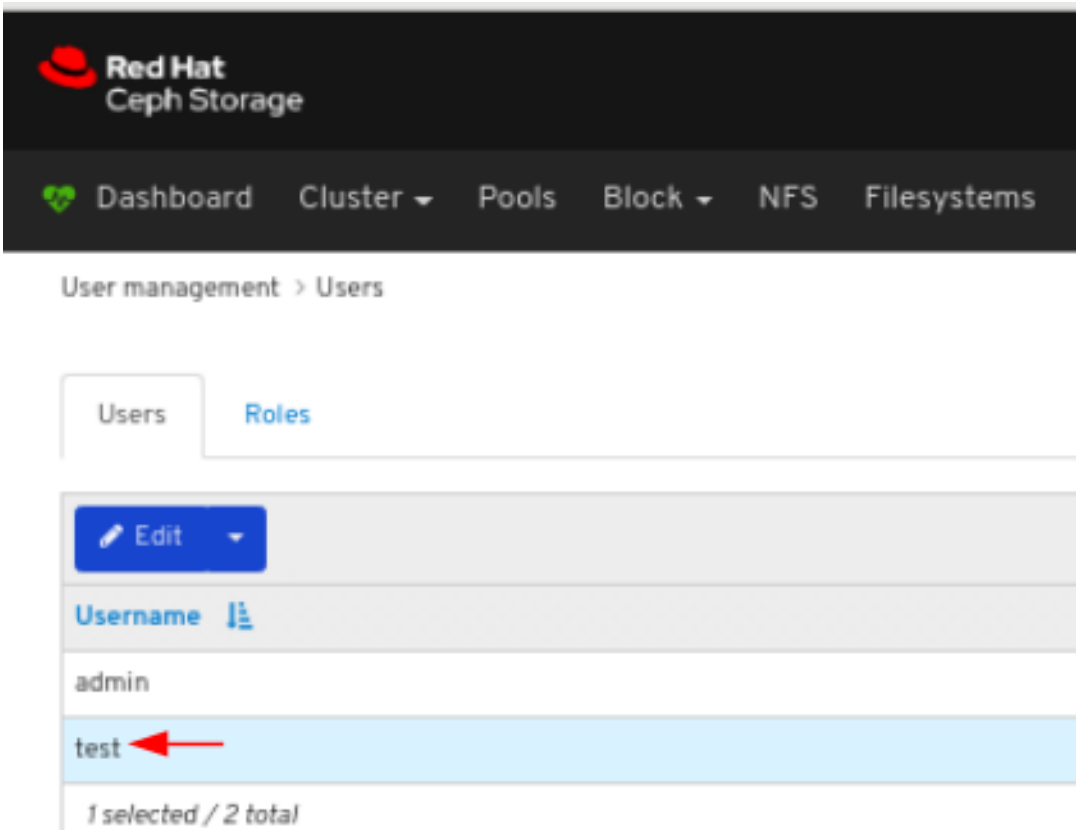

4. On *Users* tab, select *Edit* from the *Edit* dropdown menu:

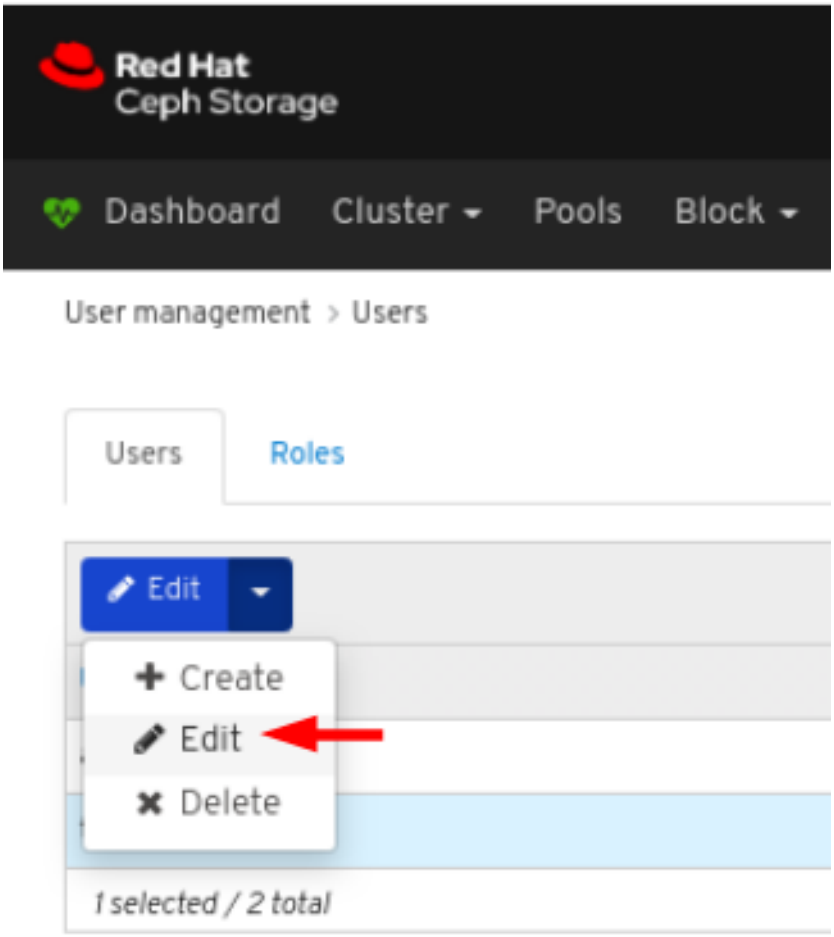

5. In the *EditUser* window, edit parameters including, and then click the *EditUser* button:

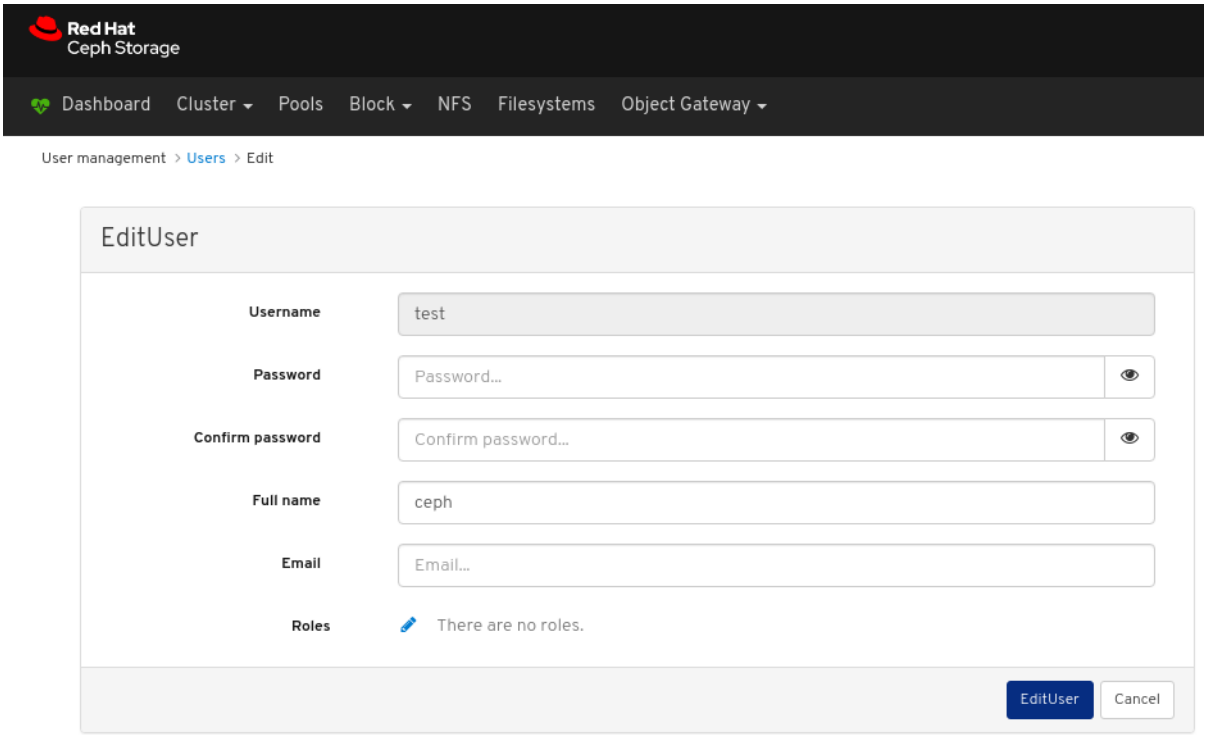

6. A notification towards the top right corner of the page indicates the user was updated successfully.

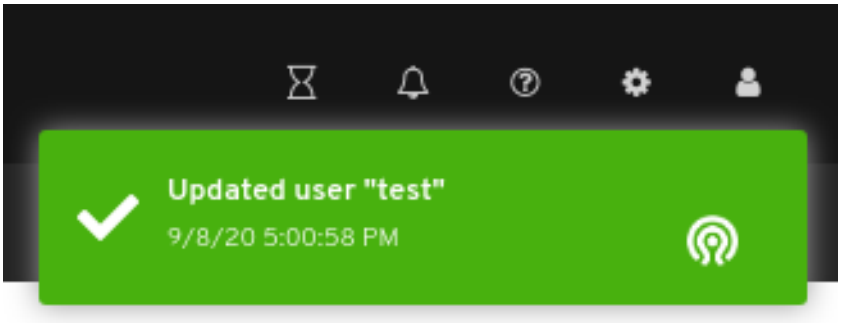

Additional Resources

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

4.3. DELETING USERS ON DASHBOARD

The dashboard allows you to delete users on the dashboard.

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. On the upper right side of the Dashboard, click the gear icon and select *User management*:

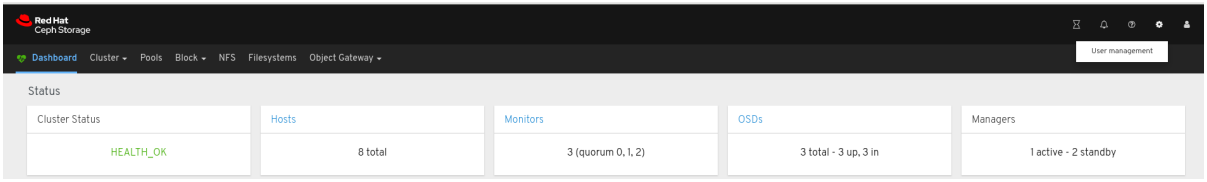

3. To delete the user, click the row:

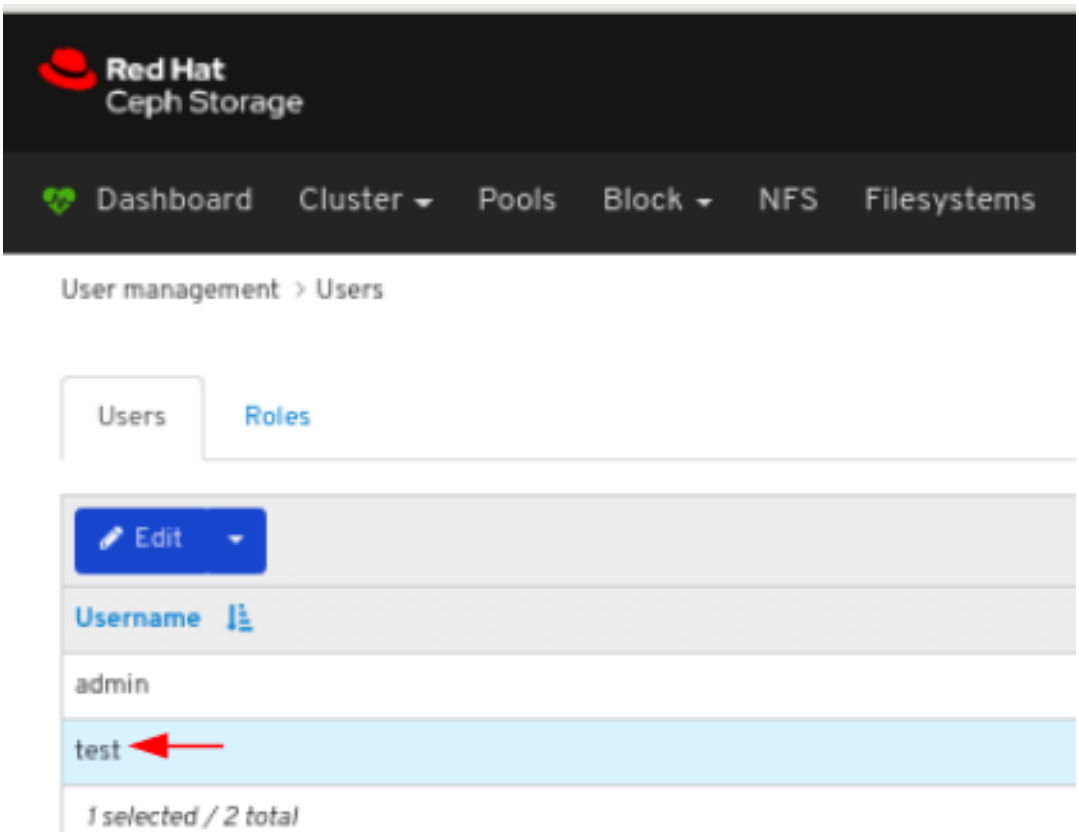

4. On *Users* tab, select *Delete* from the *Edit* dropdown menu:

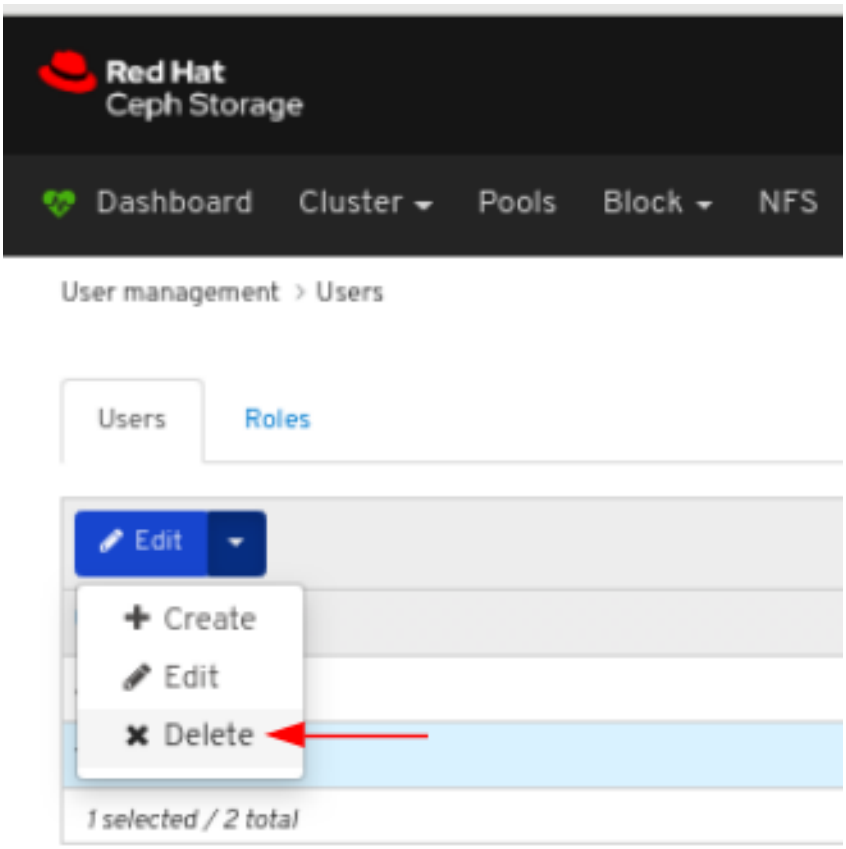

5. In the *Delete User* dialog window, Click the *Yes, I am sure* box and then Click *Delete user* to save the settings:

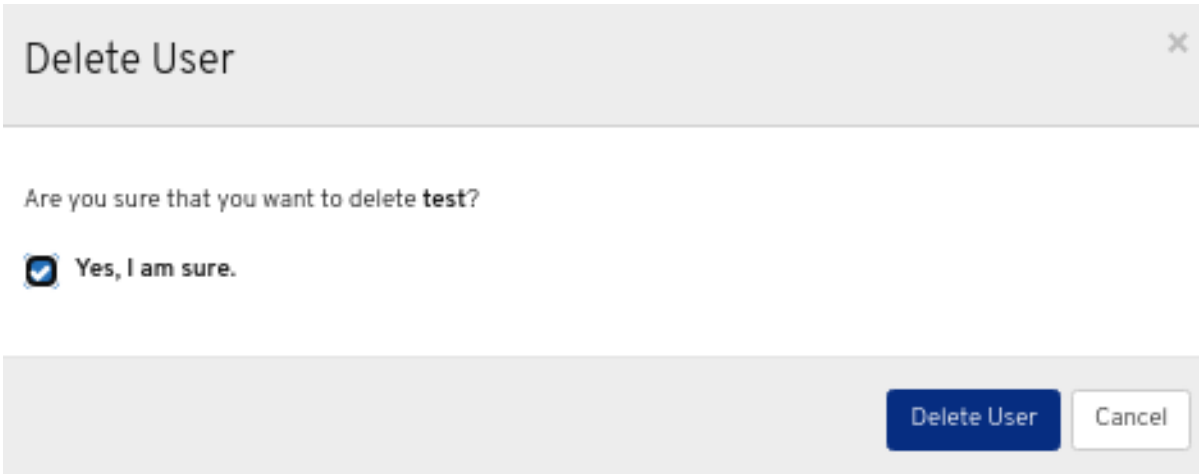

Additional Resources

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

CHAPTER 5. MONITORING THE CLUSTER

The monitoring functions of the dashboard provide different web pages which update regularly to indicate various aspects of the storage cluster. You can monitor the overall state of the cluster using the landing page, or you can monitor specific functions of the cluster, like the state of block device images.

Additional Resources

- For more information, see Accessing the landing page in the Dashboard guide.
- For more information, see Understanding the landing page in the Dashboard guide.
- For more information, see Monitoring specific functions in the Dashboard guide.

5.1. ACCESSING THE LANDING PAGE

After you log in to the dashboard, the landing page loads.

Prerequisites

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

Procedure

1. Log in to the Dashboard:

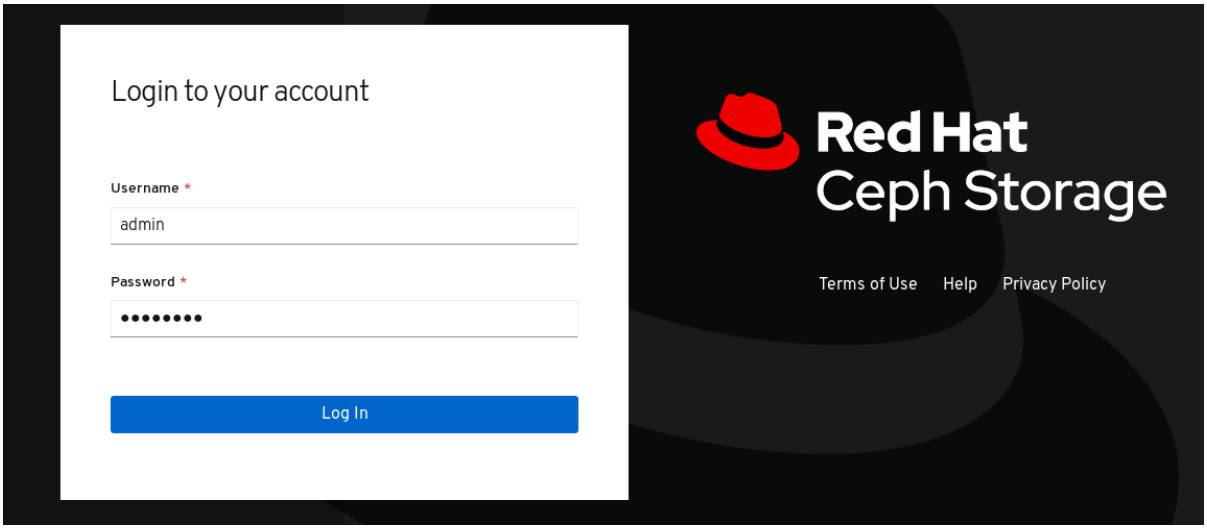

2. After you log in to the dashboard, the landing page loads:

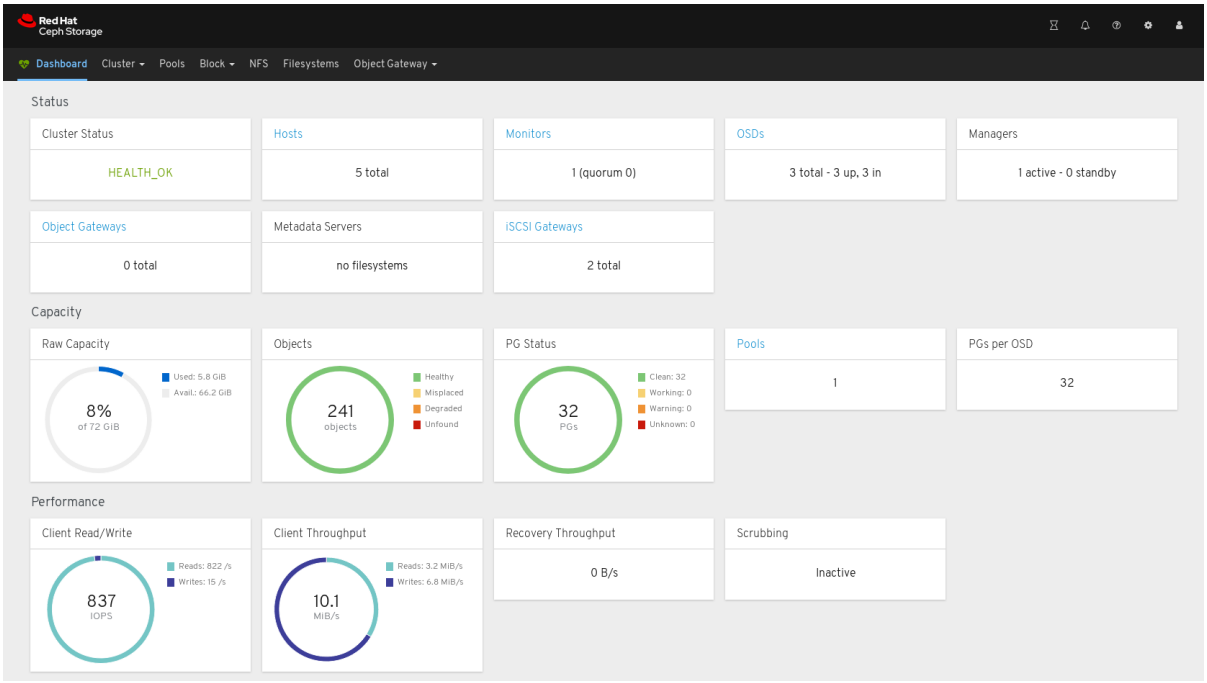

3. To return to the landing page after viewing other dashboard pages, click *Dashboard* towards the top left corner of the page:

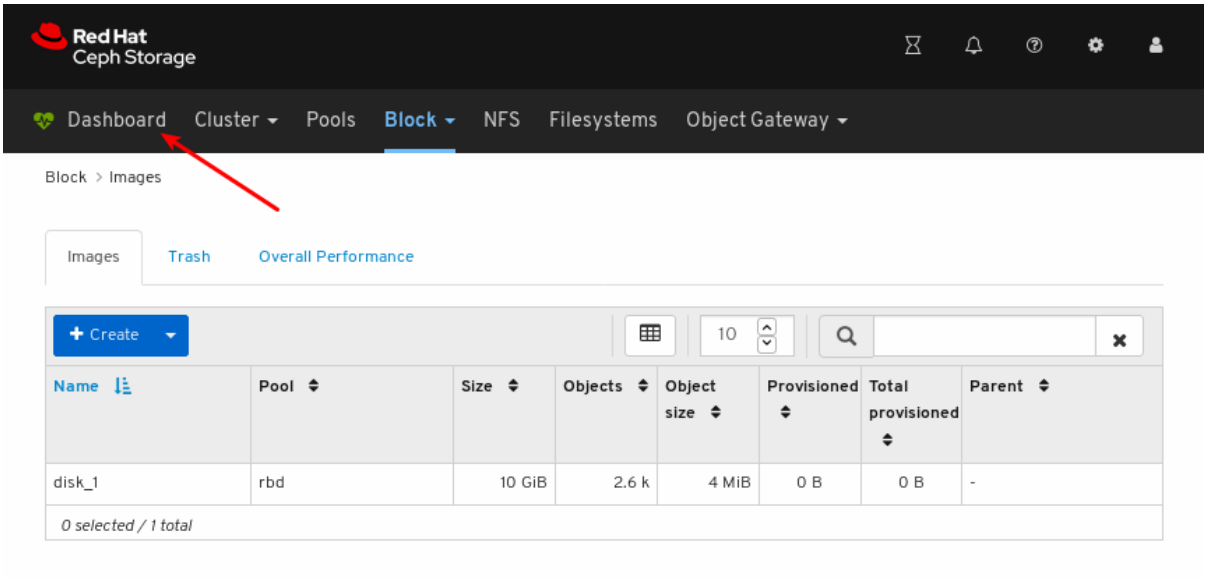

Additional Resources

- For more information, see Understanding the landing page in the Dashboard guide.
- For more information, see Monitoring specific functions in the Dashboard guide.

5.2. UNDERSTANDING THE LANDING PAGE

The landing page displays an overview of the entire Ceph cluster using individual panels. Each panel displays specific information about the state of the cluster.

Categories

The landing page orgnanizes panels into the following three categories:

1. **Status**

2. **Capacity**

3. **Performance**

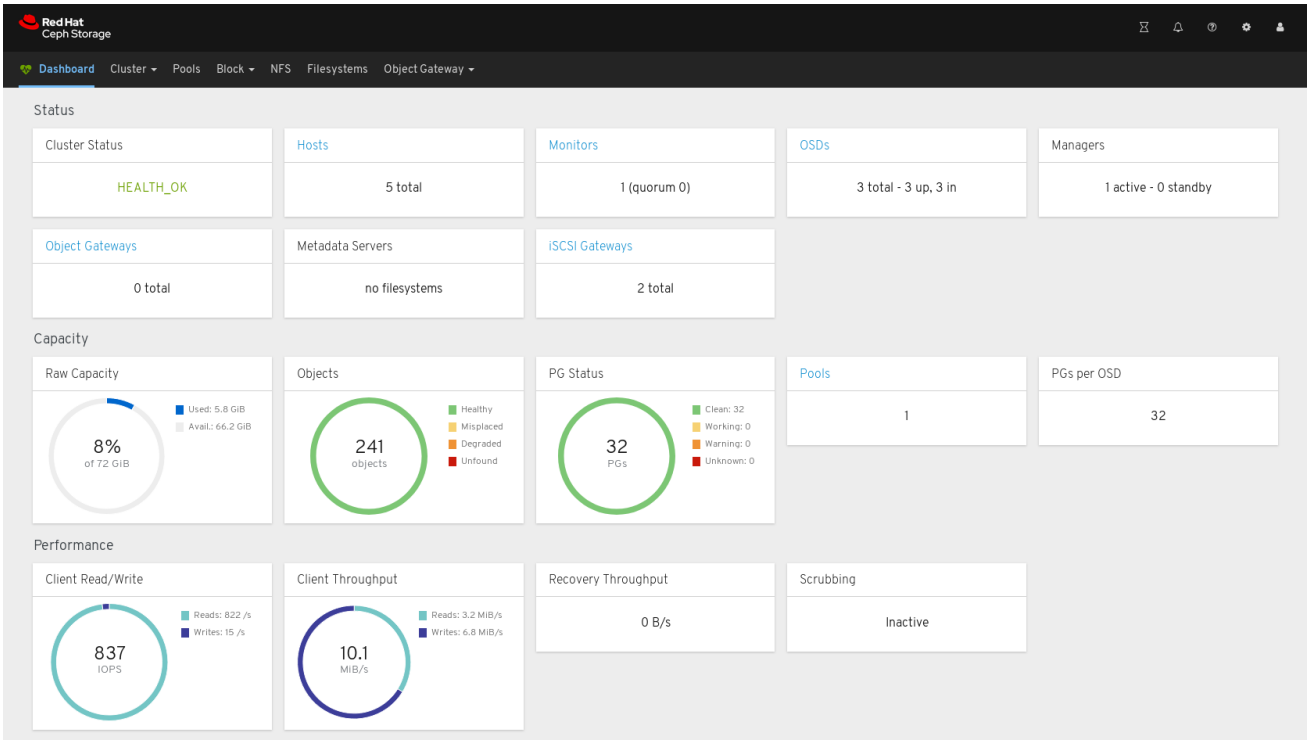

Status panels

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

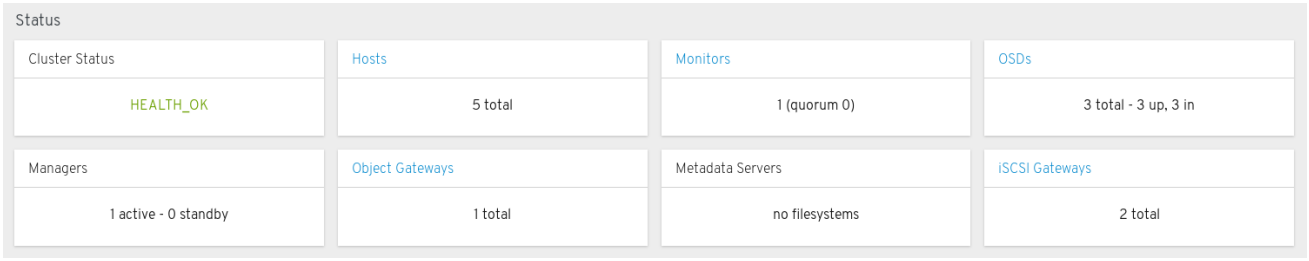

Cluster Status: Displays the current health status of the Ceph cluster.

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

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

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

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

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

Metadata Servers: Displays the number and status of metadata servers for Ceph Filesystems.

iSCSI Gateways: Displays the number of iSCSI Gateways in the Ceph storage cluster.

Capacity panels

The capacity panels display storage usage metrics.

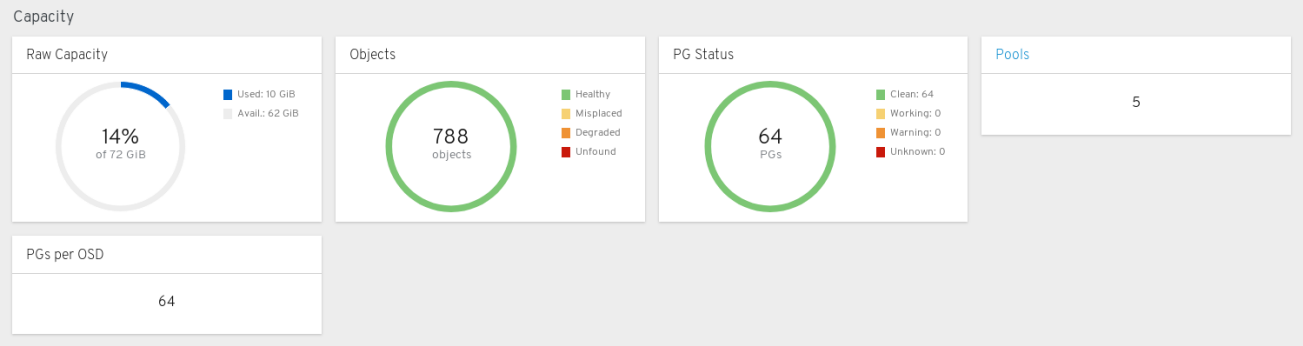

Raw Capacity: Displays the utilization and availability of the raw storage capacity of the cluster.

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

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

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

- activating
- backfill_wait
- backfilling
- creating
- deep
- degraded
- forced_backfill
- forced_recovery
- peering
- peered
- recovering
- recovery_wait
- repair
- scrubbing
- snaptrim
- snaptrim_wait

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

- backfill_toofull

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

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

PGs per OSD: Displays the number of Placement Groups per OSD.

Performance panels

The performance panels display information related to data transfer speeds.

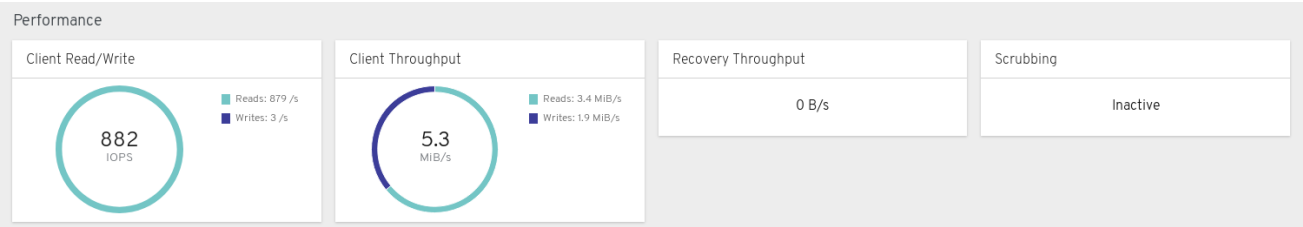

Client Read/Write: Displays total input/output opreations per second, reads per second, and writes per second.

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

Recovery Throughput Displays the Client recovery rate.

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

Additional Resources

- For more information, see Accessing the landing page in the Dashboard guide.
- For more information, see Monitoring specific functions in the Dashboard guide.

CHAPTER 6. MONITORING SPECIFIC FUNCTIONS

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

6.1. MONITORING MONITORS

The Red Hat Ceph Storage Dashboard allows you to view various details about Monitor nodes.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster* and then click *Monitors*.
3. The *Monitors* overview page displays information about the overall monitor status as well as tables of *in Quorum* and *Not in quorum* Monitor nodes:

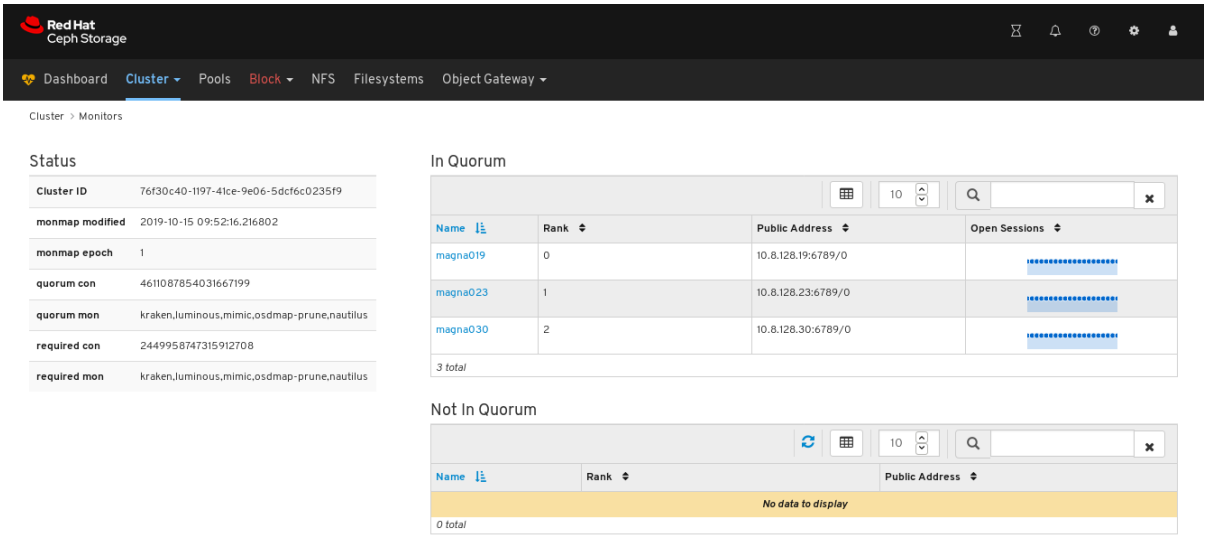

4. To see the number of open sessions, hover the cursor over the blue dotted trail:

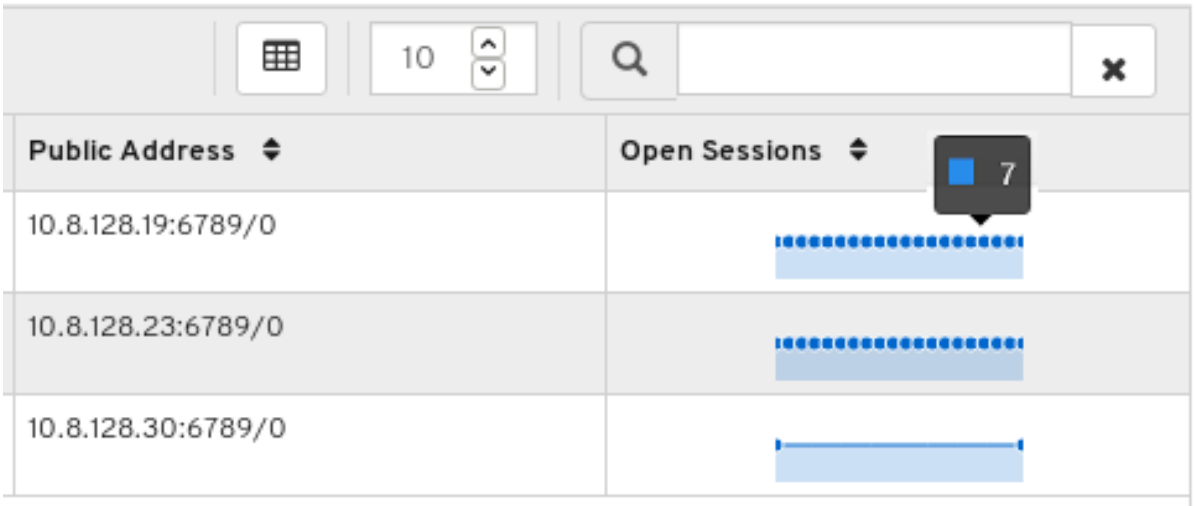

5. To see performance counters for any monitor, click its host name:

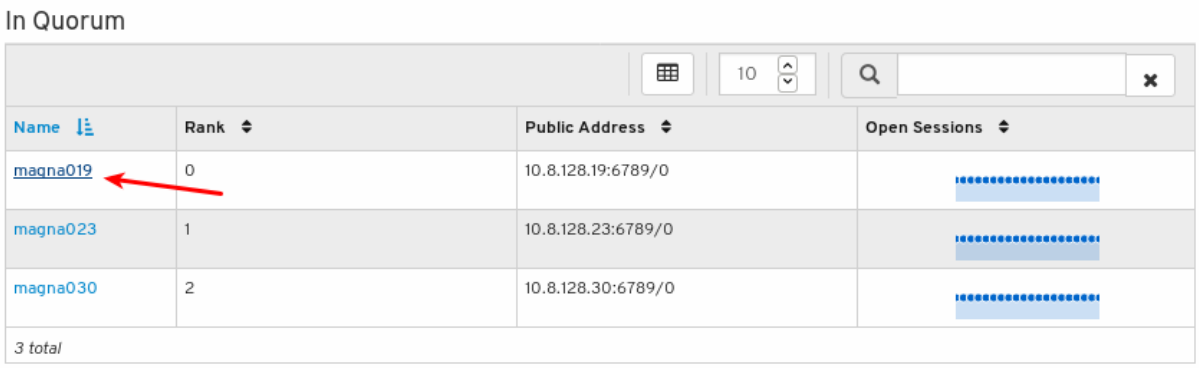

6. View the performance counters:

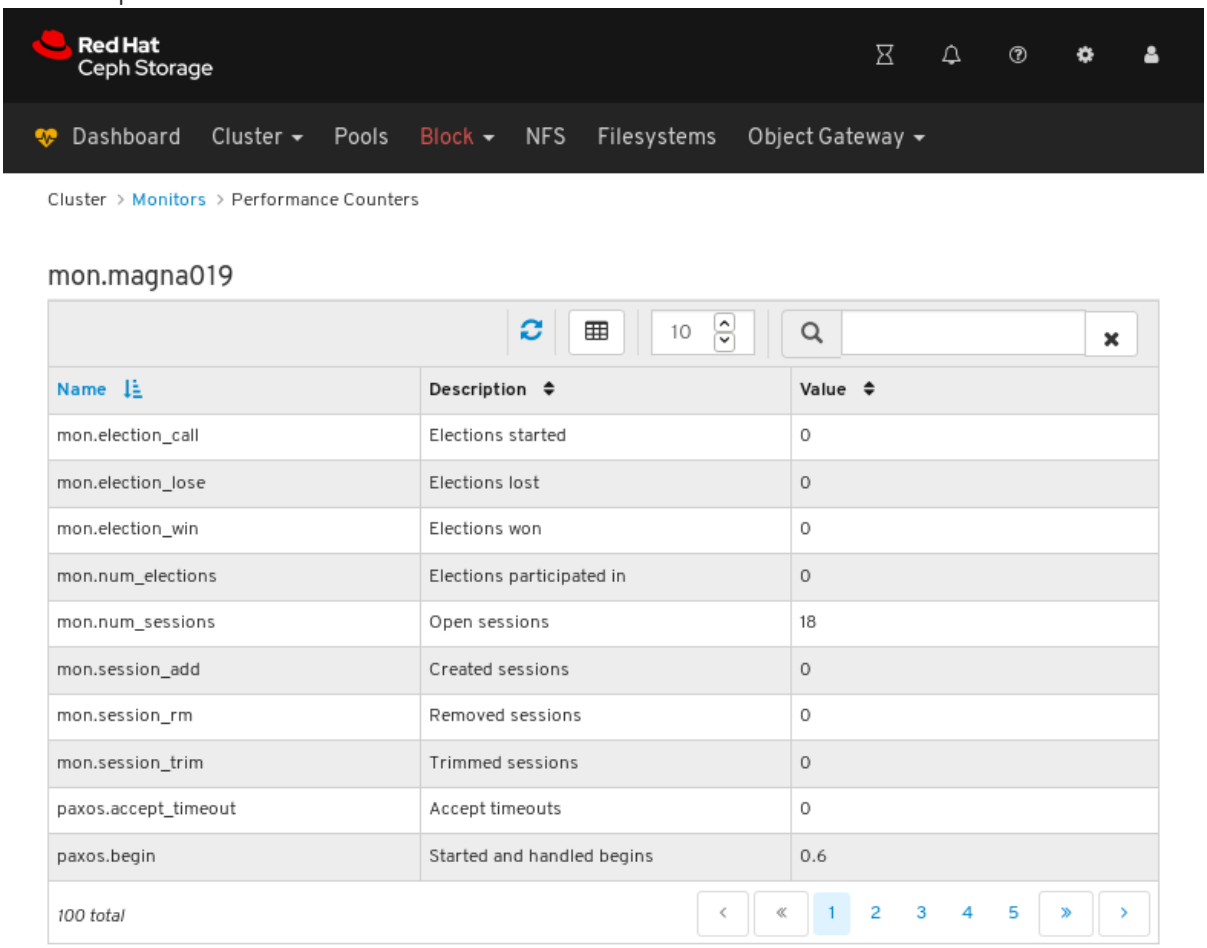

Additional Resources

- For more information about Monitors, see Ceph monitors in the Operations guide.
- For more information about performance counters, see Performance counters in the Administration Guide

6.2. MONITORING HOSTS

The Red Hat Ceph Storage Dashboard allows you to view various details about hosts.

Prerequisites

- A running Red Hat Ceph Storage cluster.

- Dashboard is installed.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster* and then click *Hosts*.
3. To view the list of hosts, click the *Hosts List* tab:

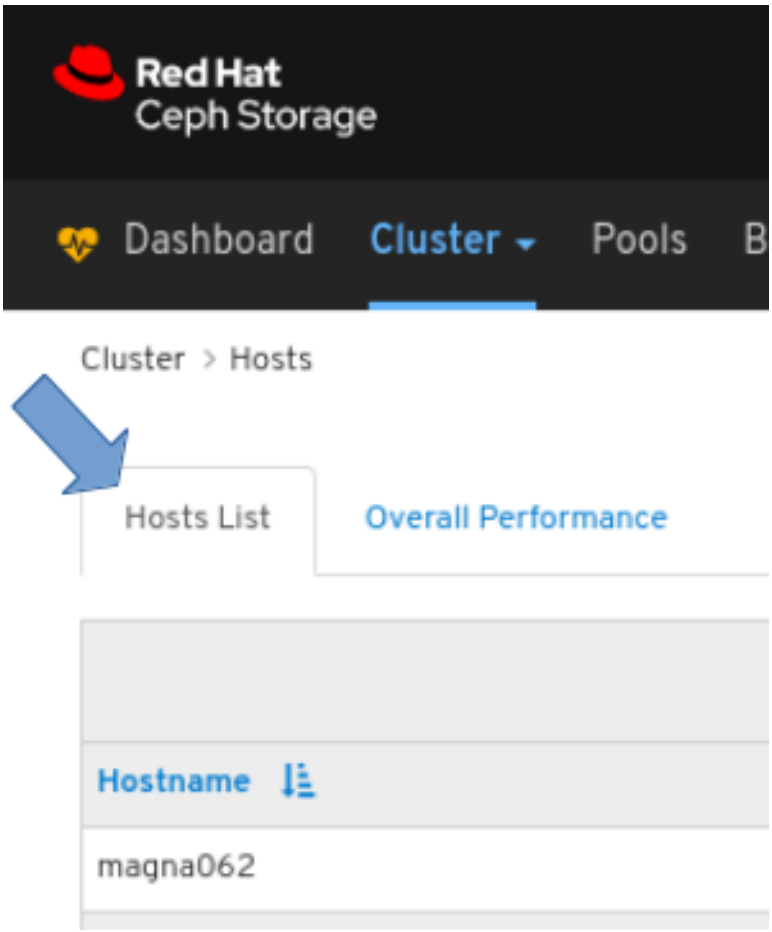

4. To view the *Performance Details* of a host, in the *Host* tab, click its row and select the time range from the *Grafana Time Picker* drop-down menu:

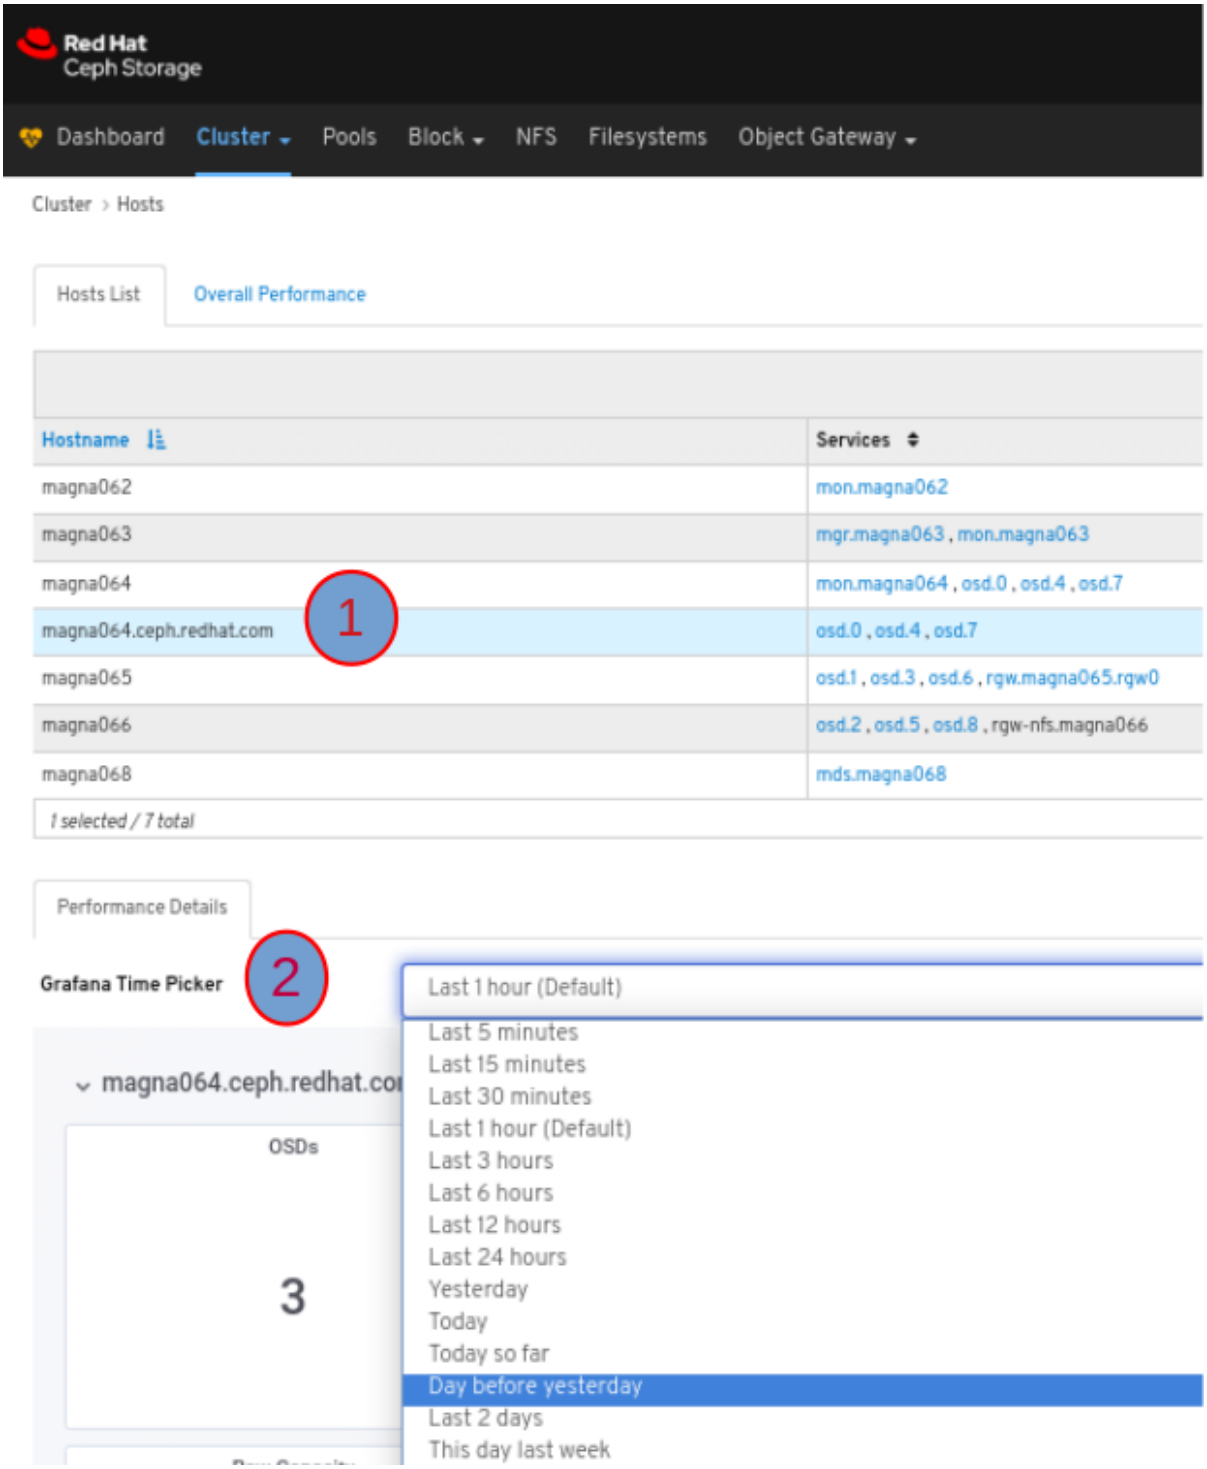

5. To view the performance counters for a specific service on a host, click the service:

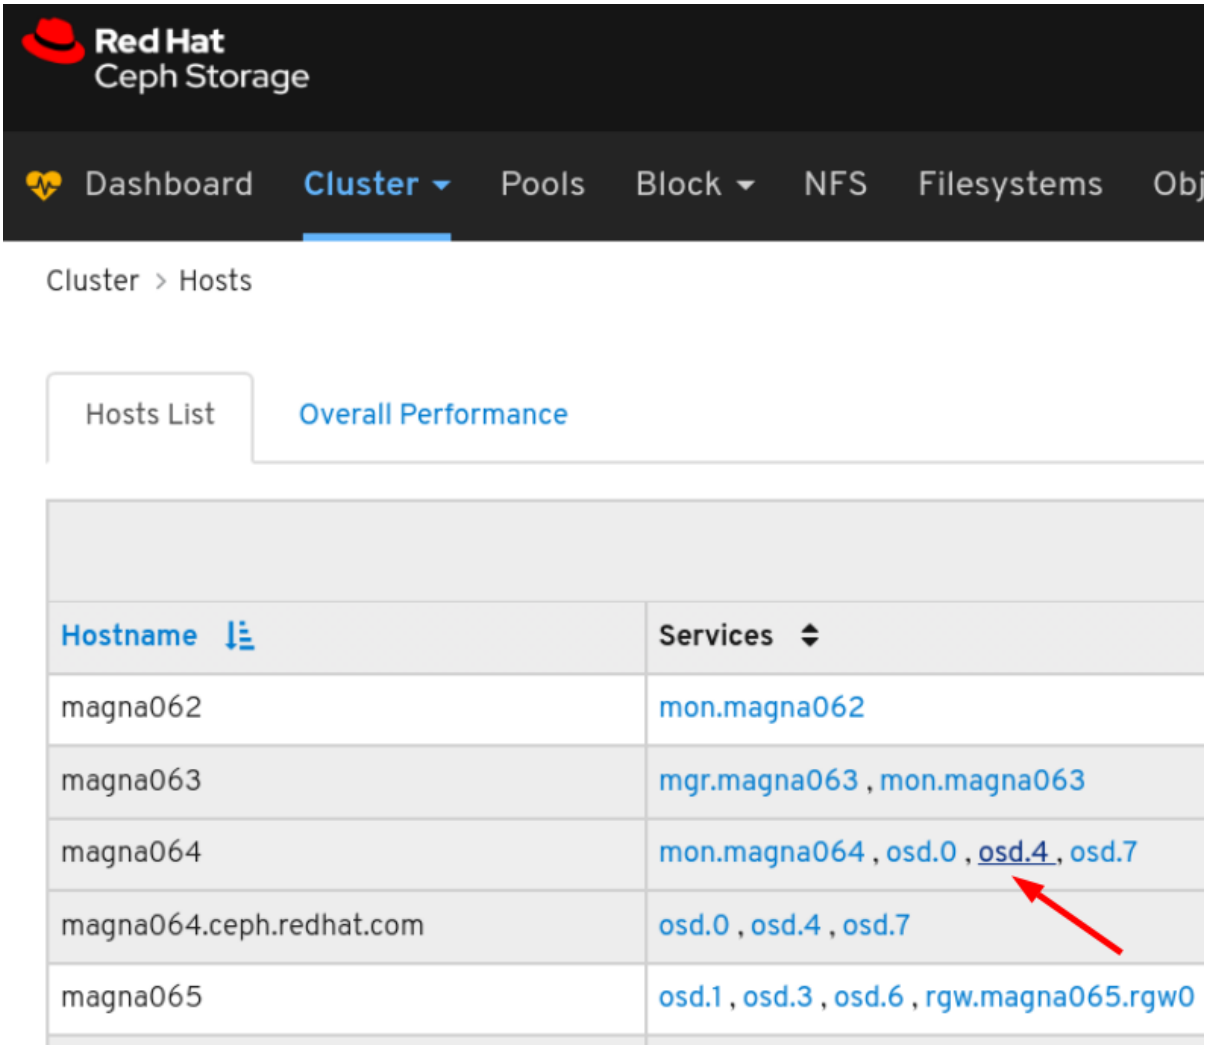

6. View the performance counters:

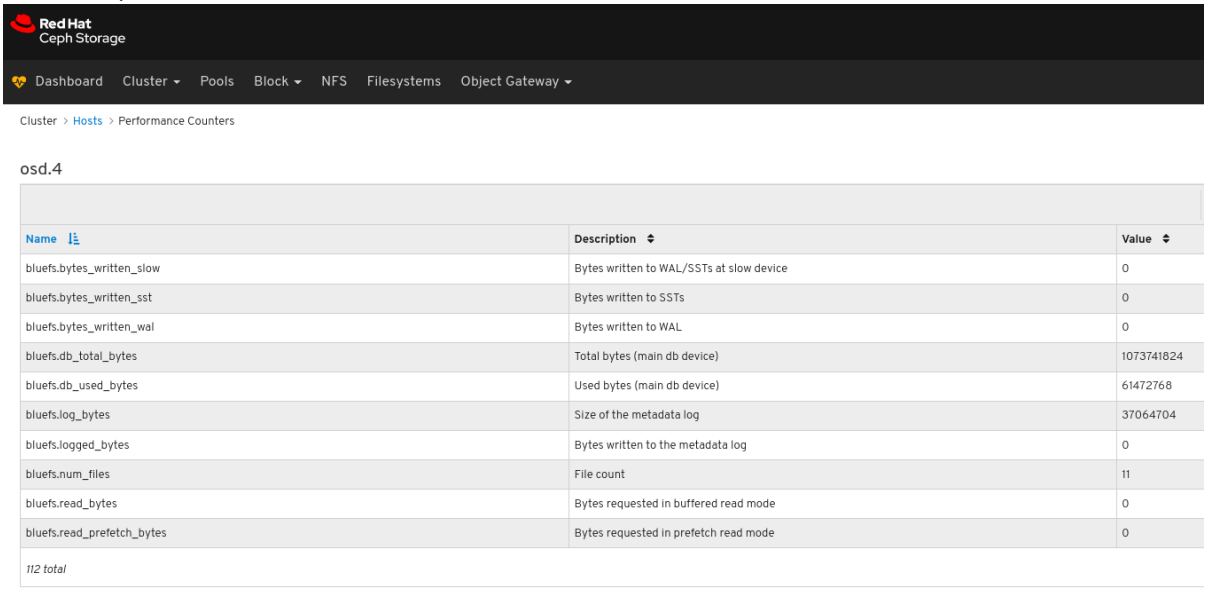

7. To view performance data for all the hosts:

   a. Click the *Overall Performance* tab towards the top left of the page:

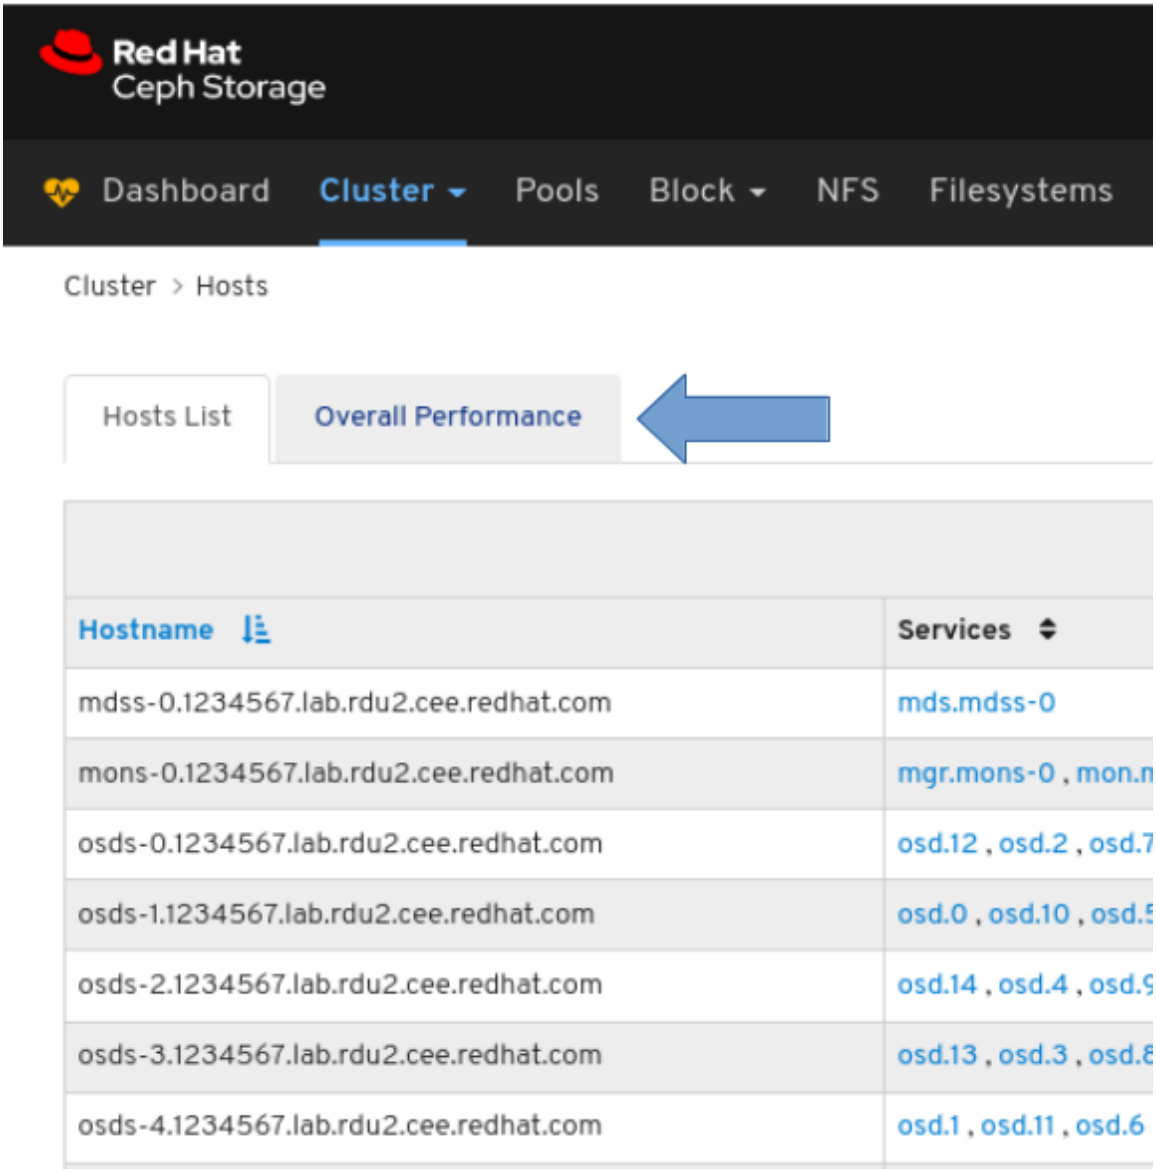

b. Select the time range from the *Grafana Time Picker* drop-down:

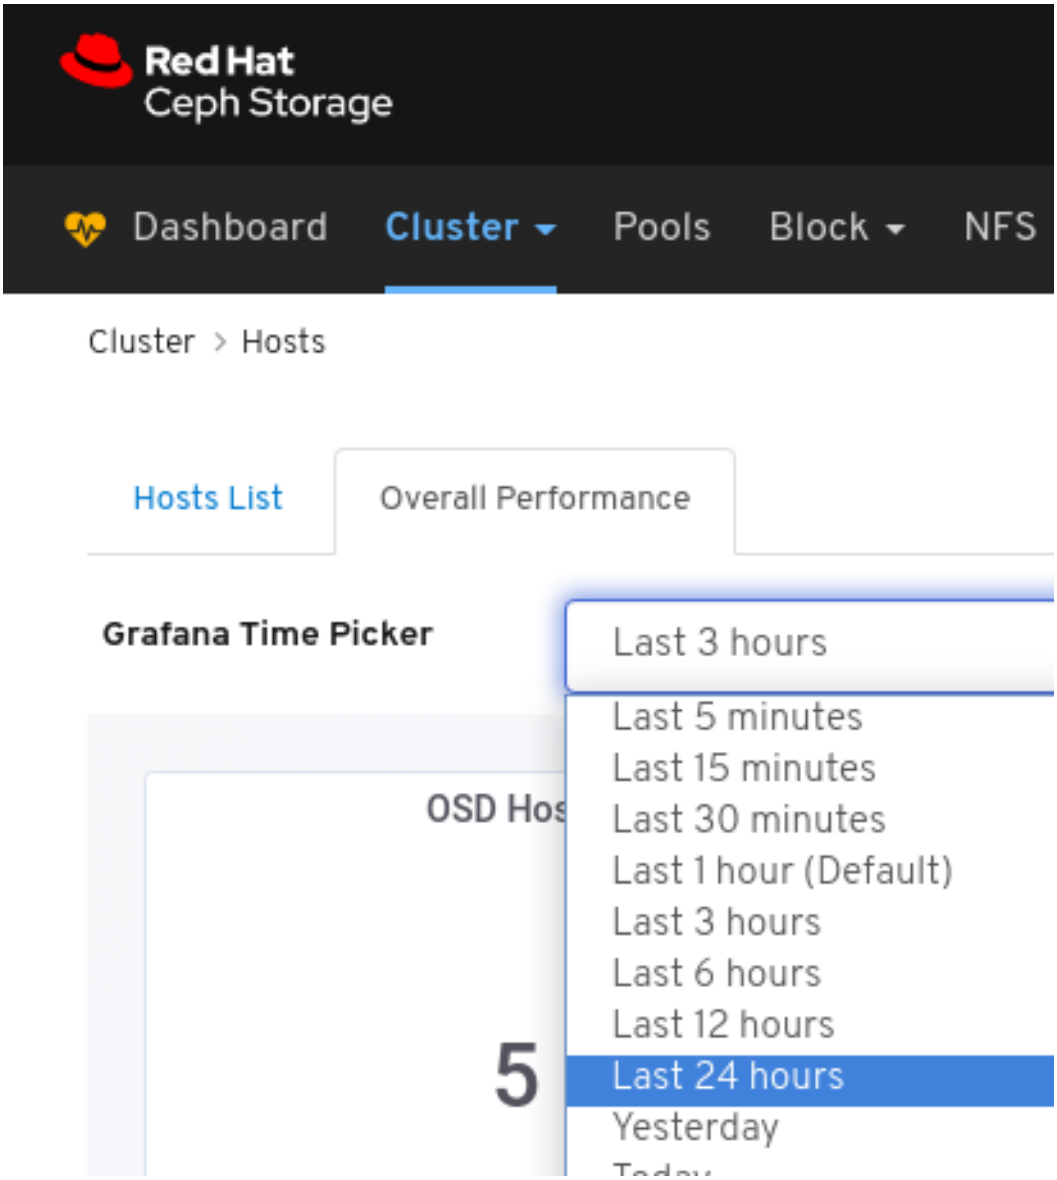

c. View the *Overall Performance* page:

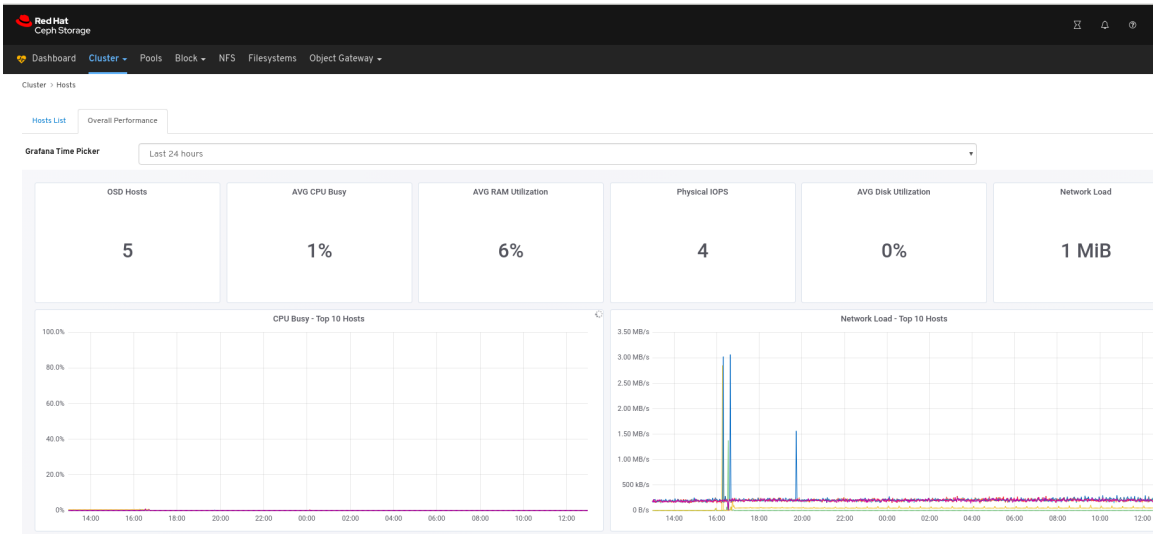

Additional Resources

- See the *Performance counters* in the *Red Hat Ceph Storage Administration Guide* for more details.

6.3. MONITORING OSDS

The Red Hat Ceph Storage Dashboard allows you to view various details about OSDs.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster* and then click *OSDs*.
3. To view the list of OSDs, click the *OSDs List* tab:

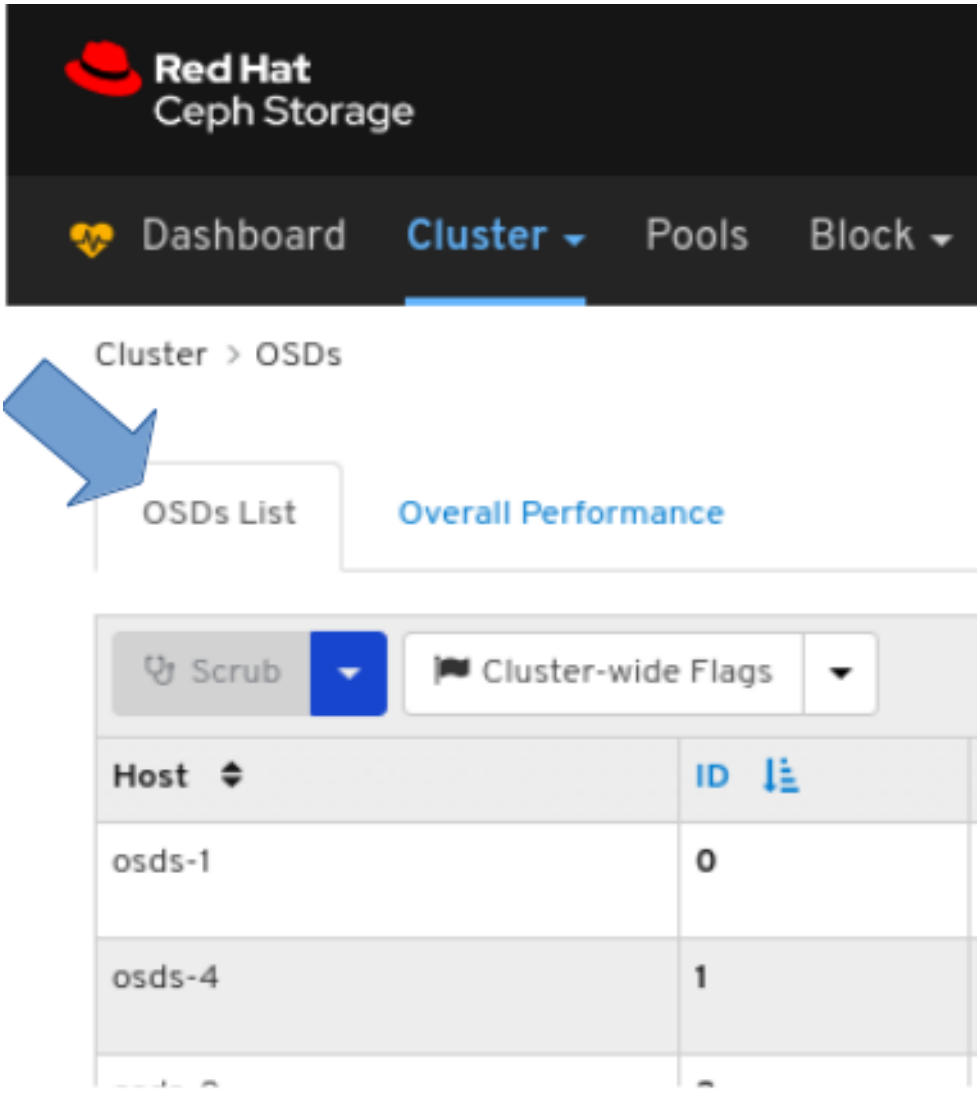

4. To view the attributes of an OSD, on the *OSDS List* tab, click its row:

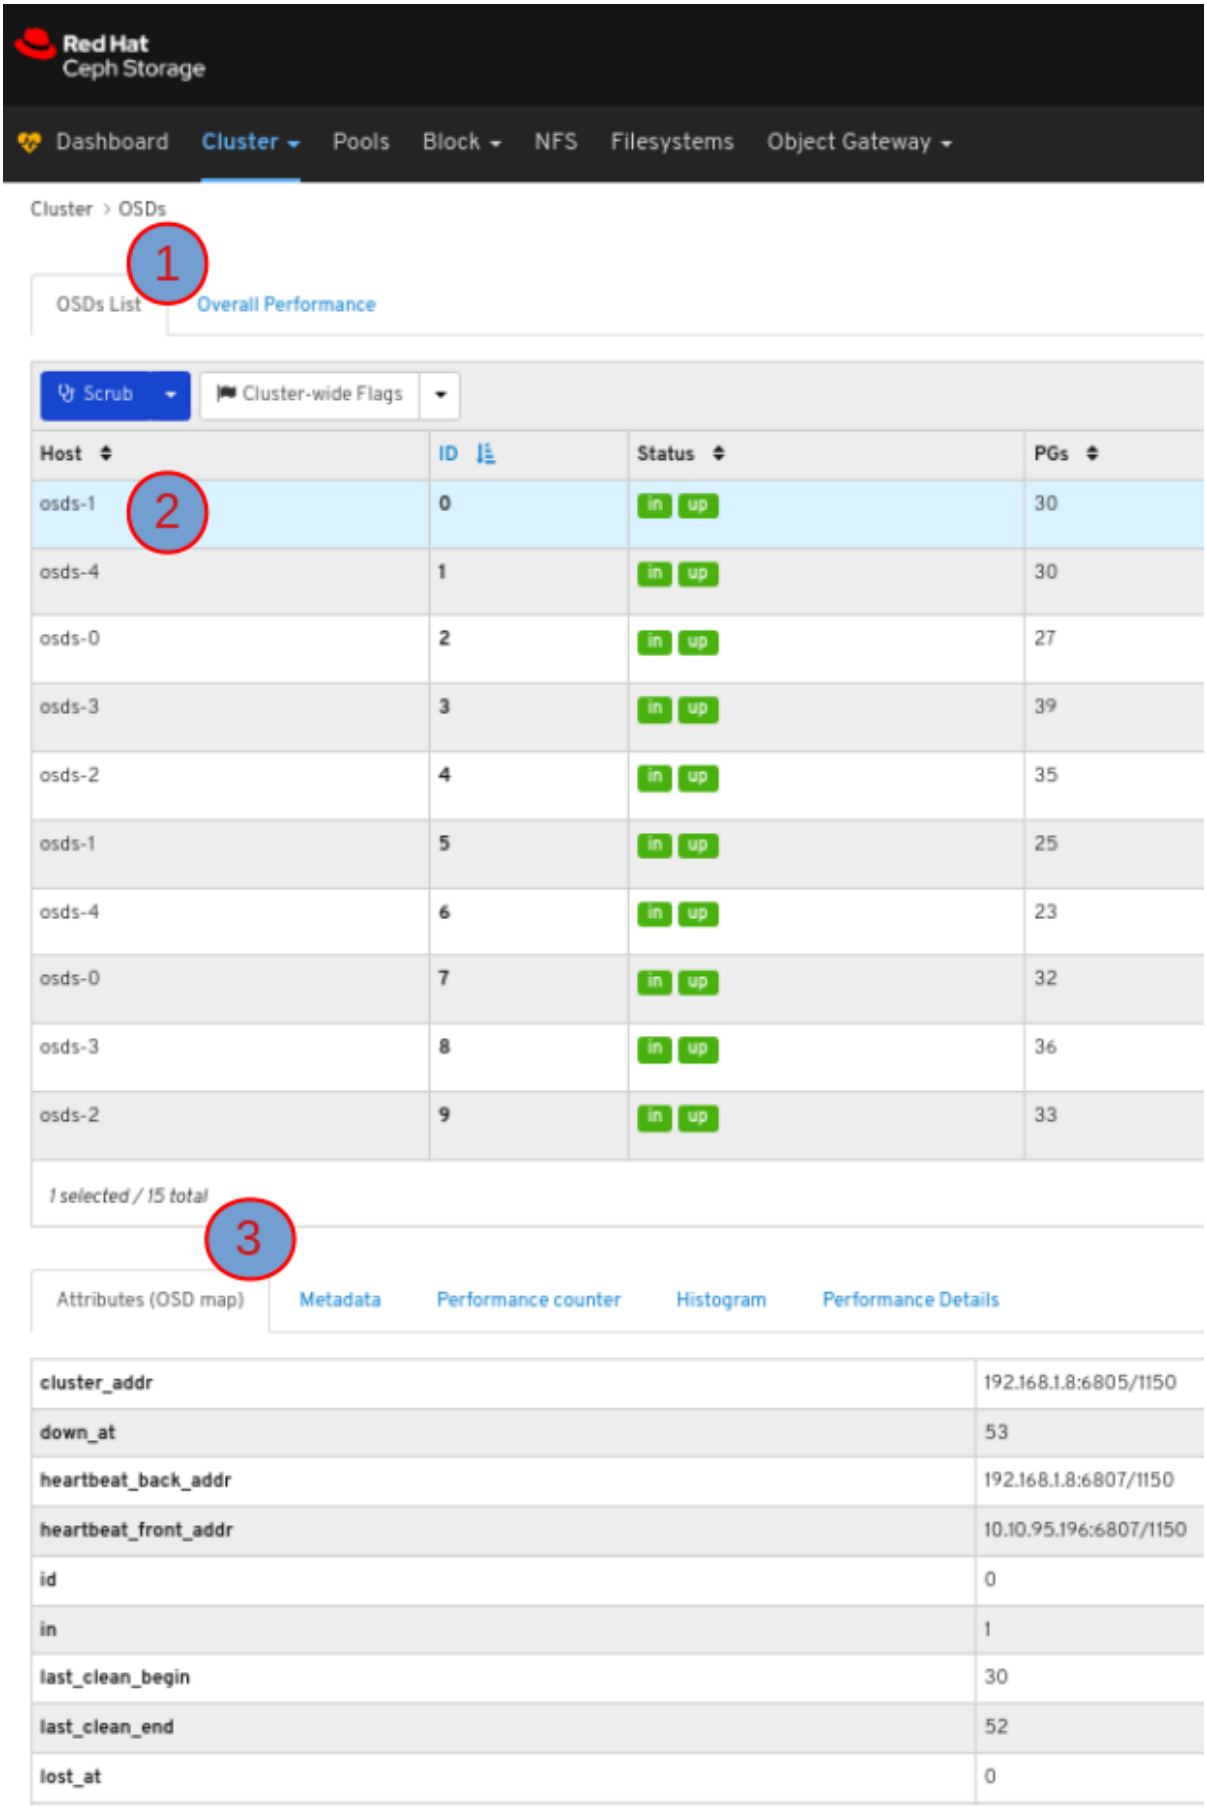

5. To view the metadata of the OSD, in the *OSDs* tab, click its row and click the *Metadata* tab:

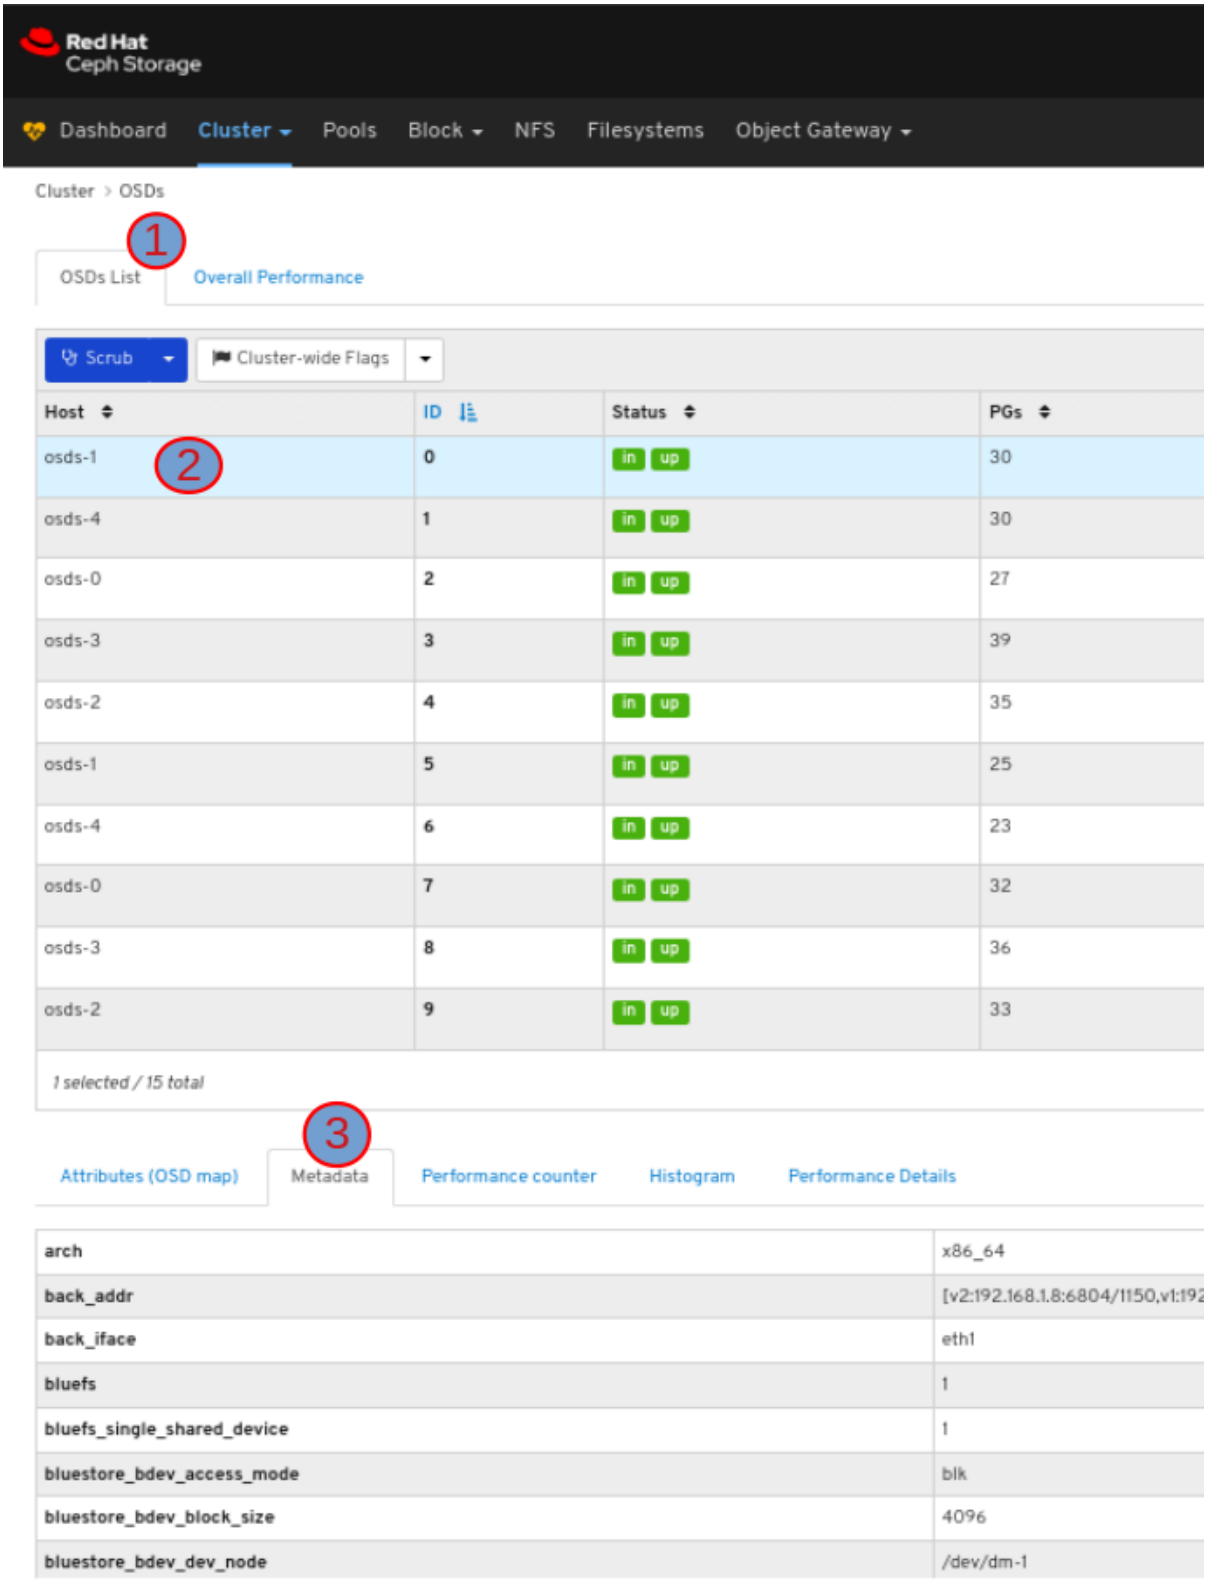

6. To view the performance counter of the OSD, in the *OSDs* tab, click its row and click the *Performance counter* tab:

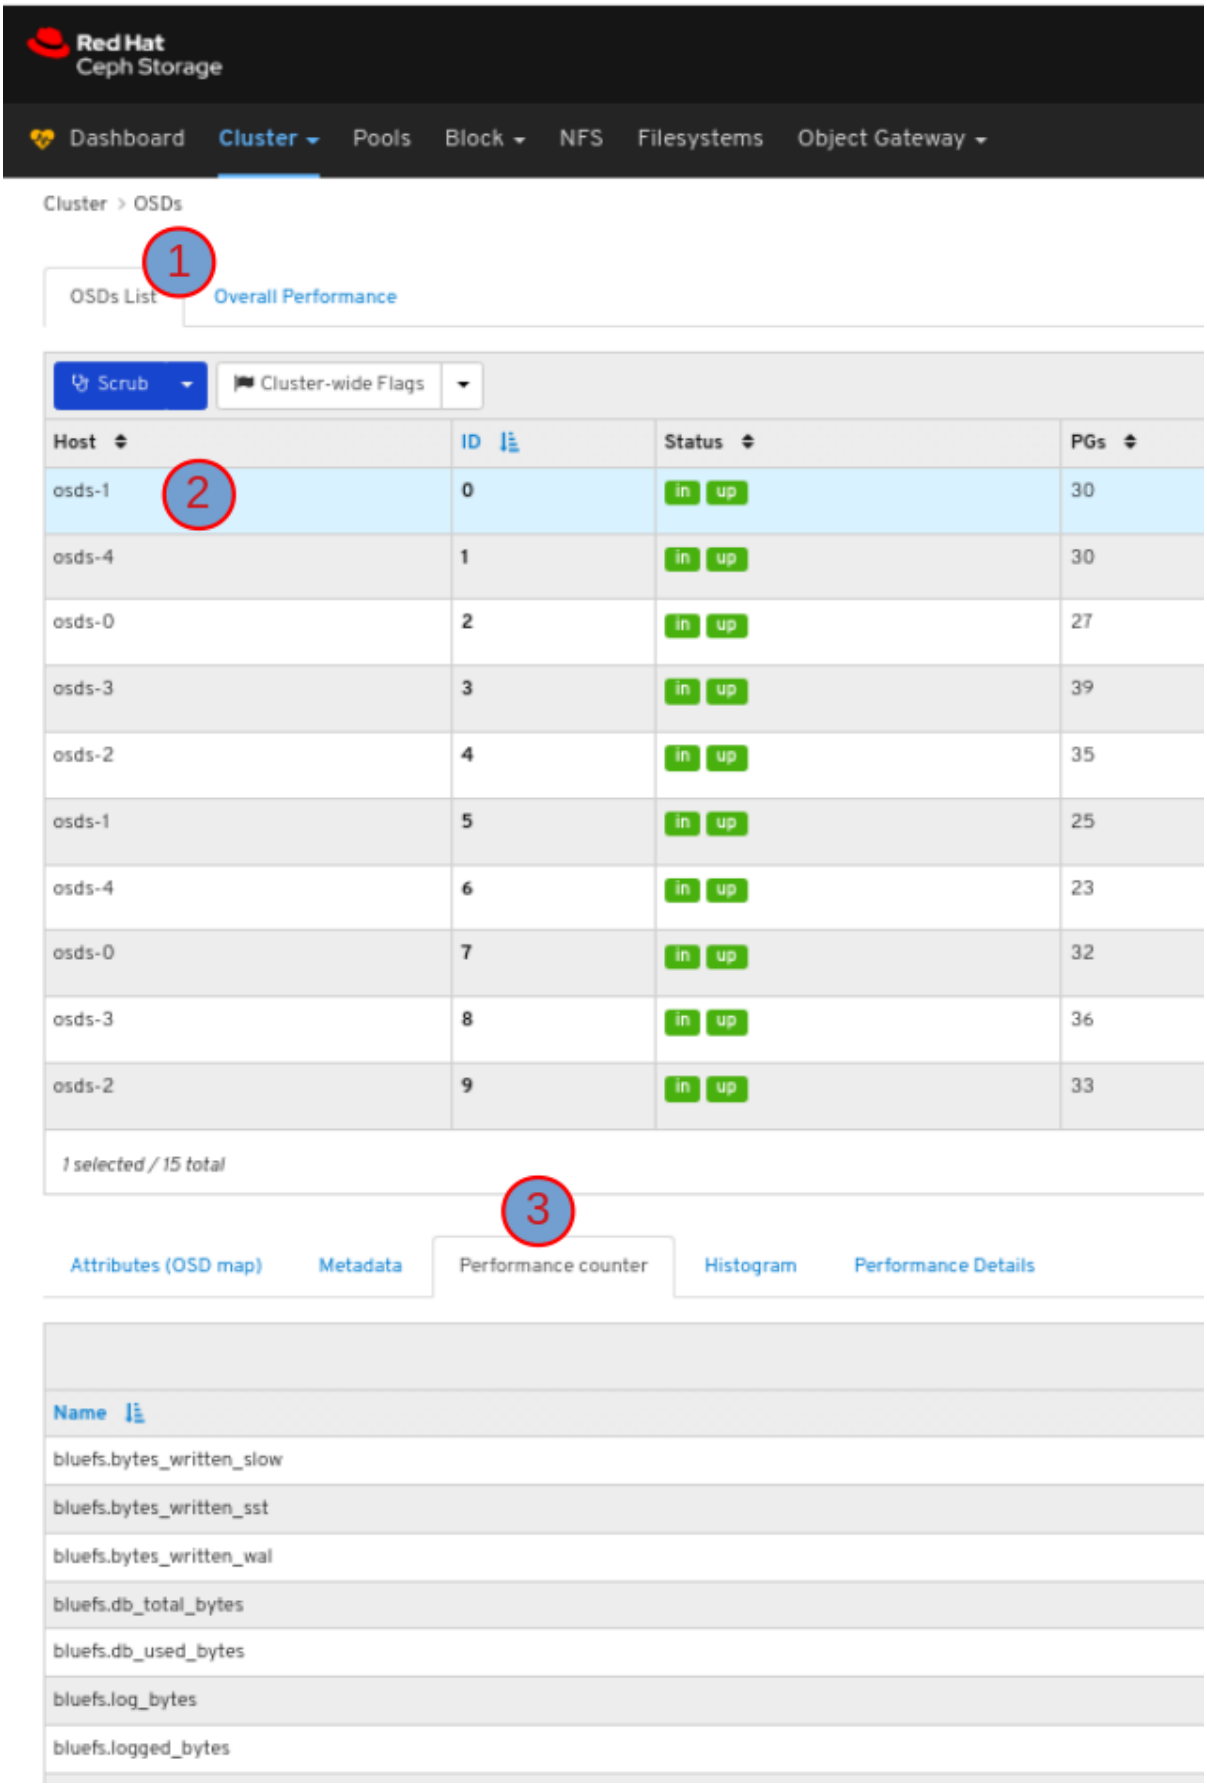

7. To view the histogram of the OSD, in the *OSDs* tab, click its row and click the *Histogram* tab:

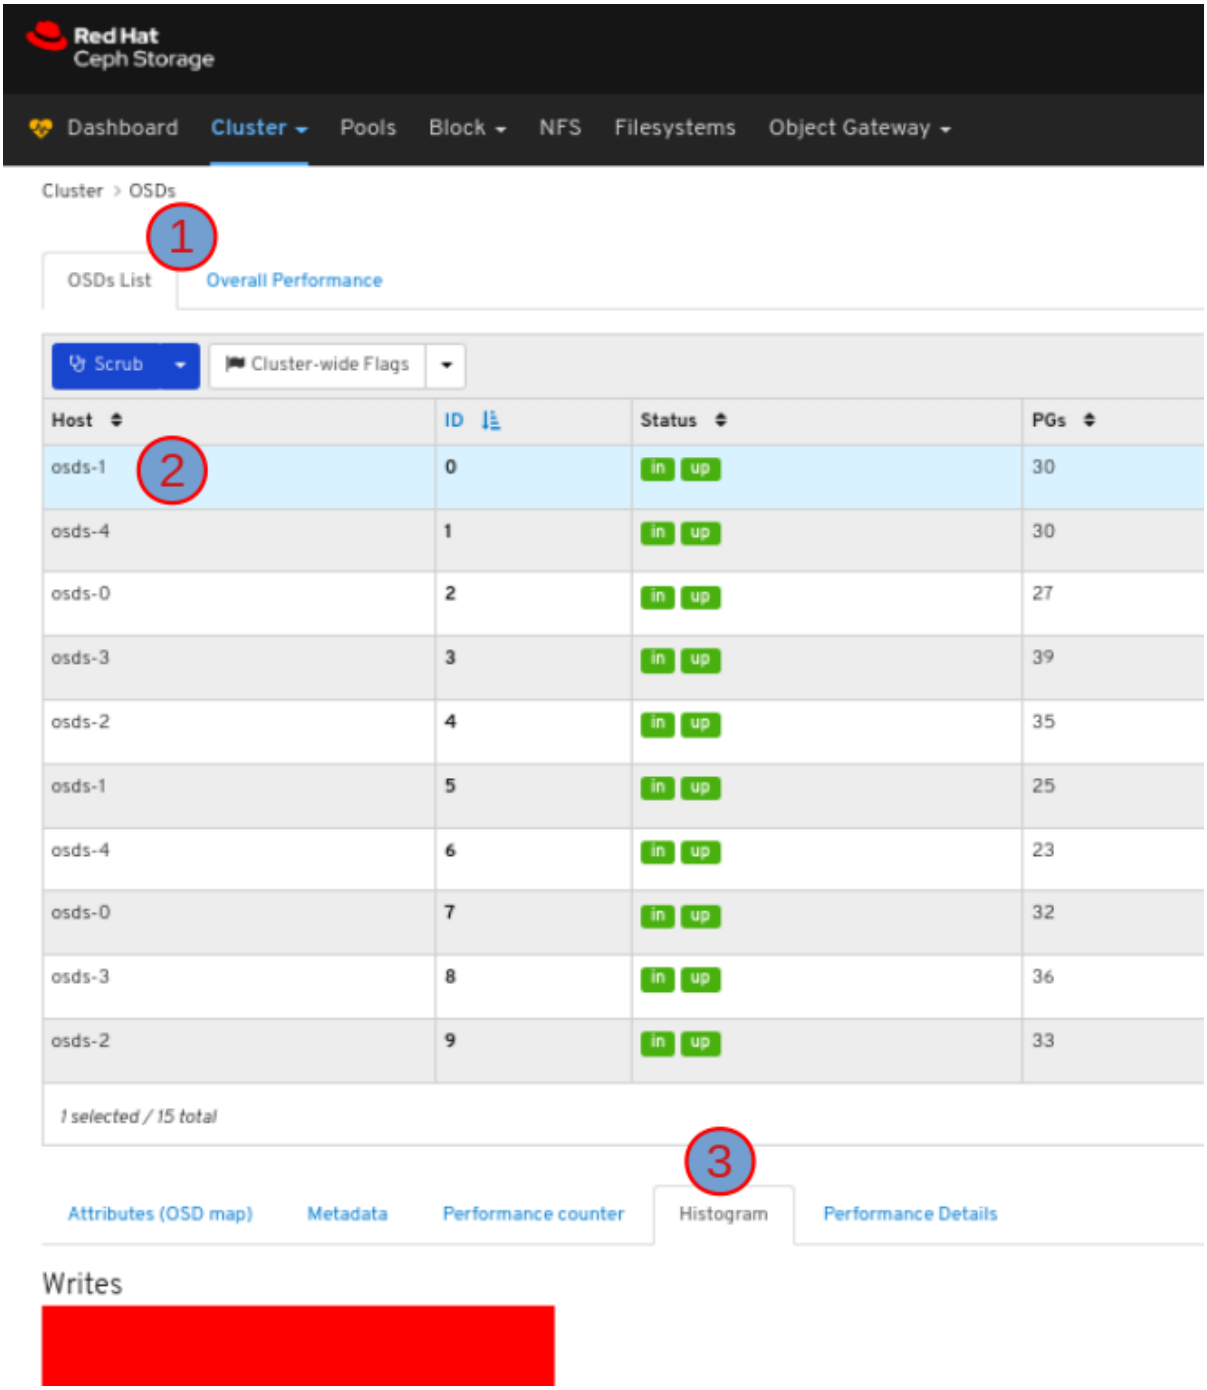

8. To view the performance details of the OSD:

   a. In the *OSDs* tab, click its row, click the *Performance Details* tab:

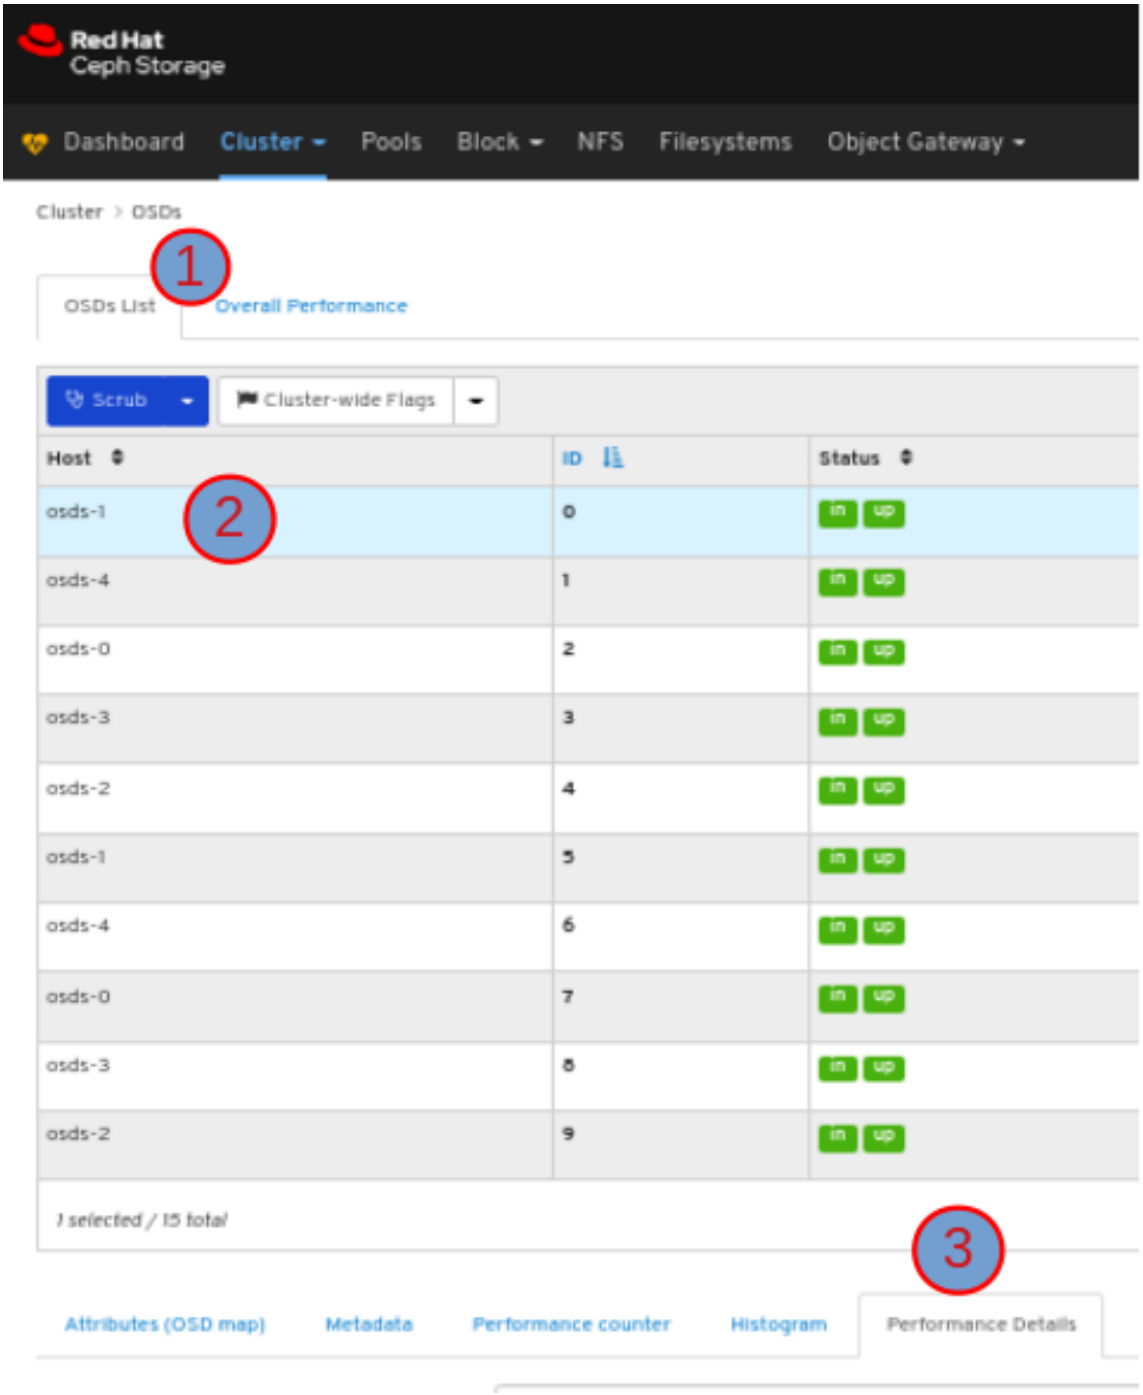

b. Select the time range from the *Grafana Time Picker* drop-down menu:

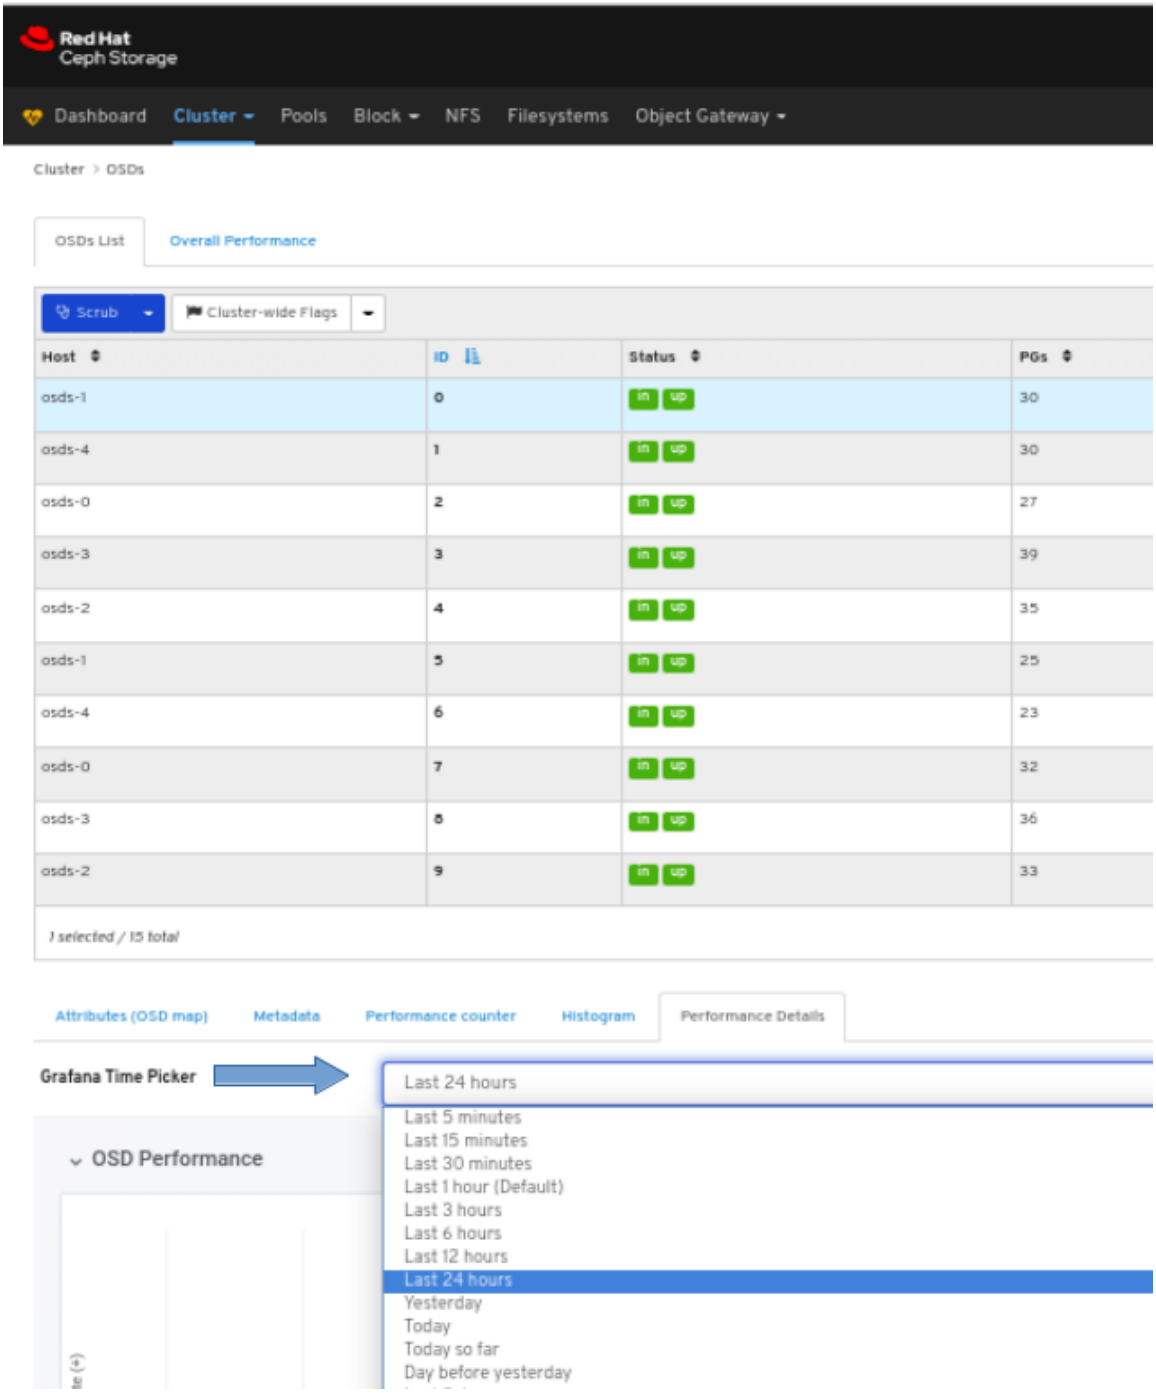

c. View the performance details of the OSD:

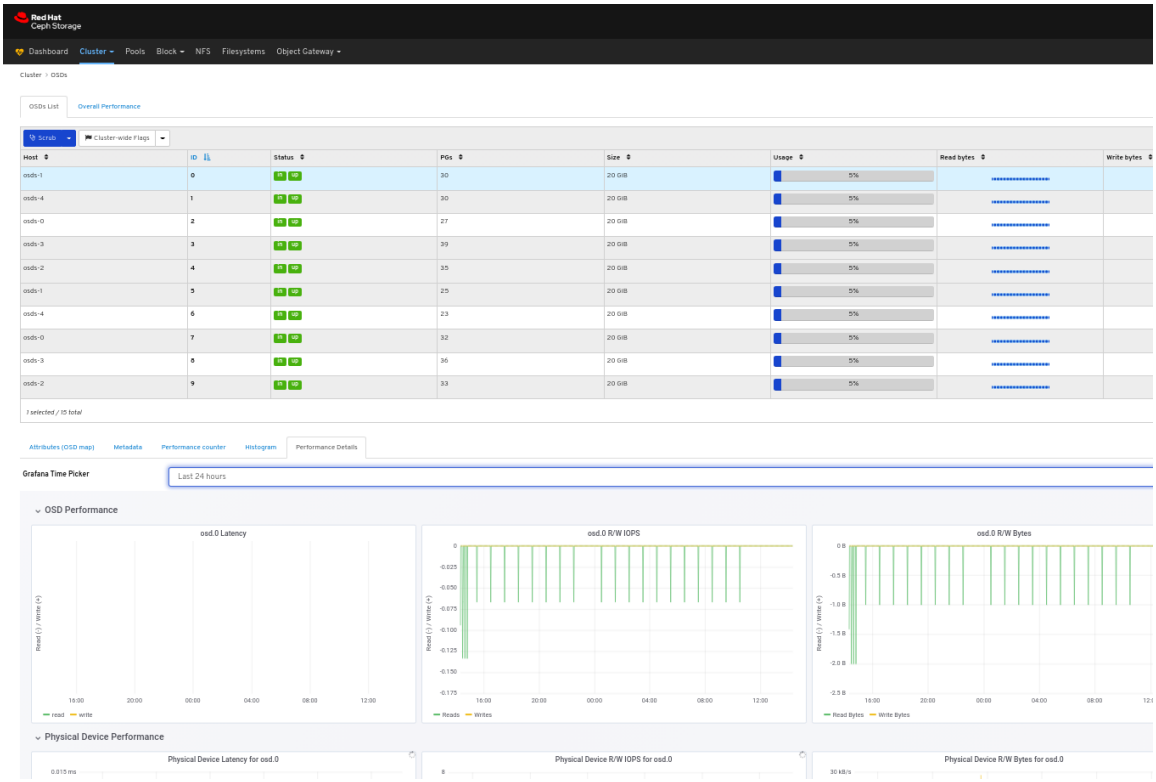

9. To view the overall performance of all the OSDs:

   a. Click the *Overall Performance tab* towards the top left of the page:

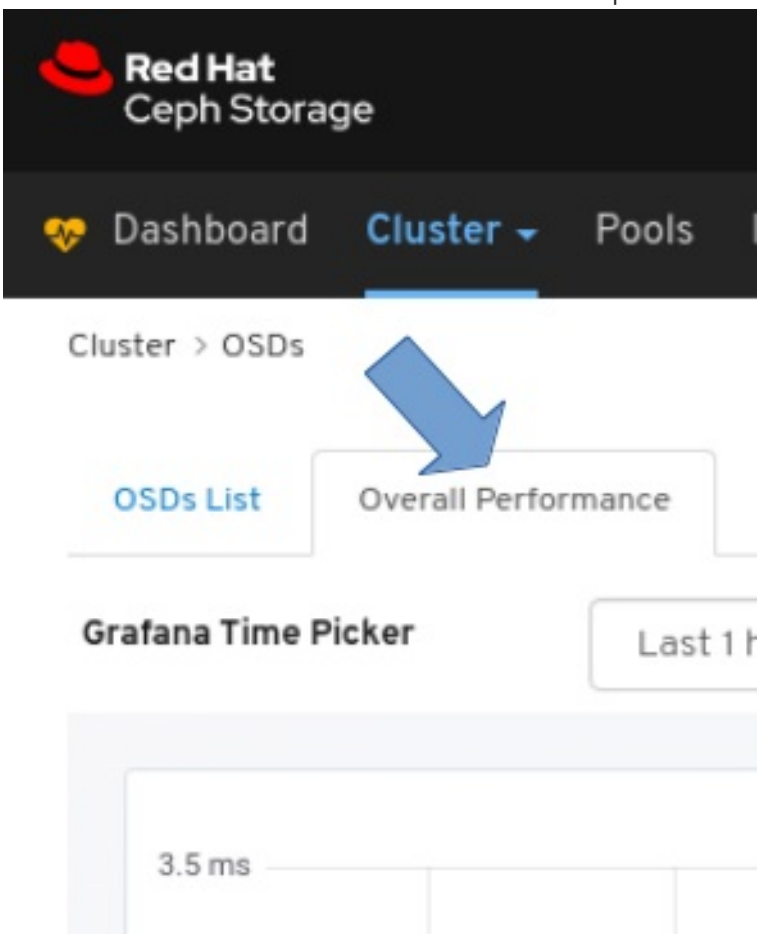

   b. Select the time range from the *Grafana Time Picker* drop-down:

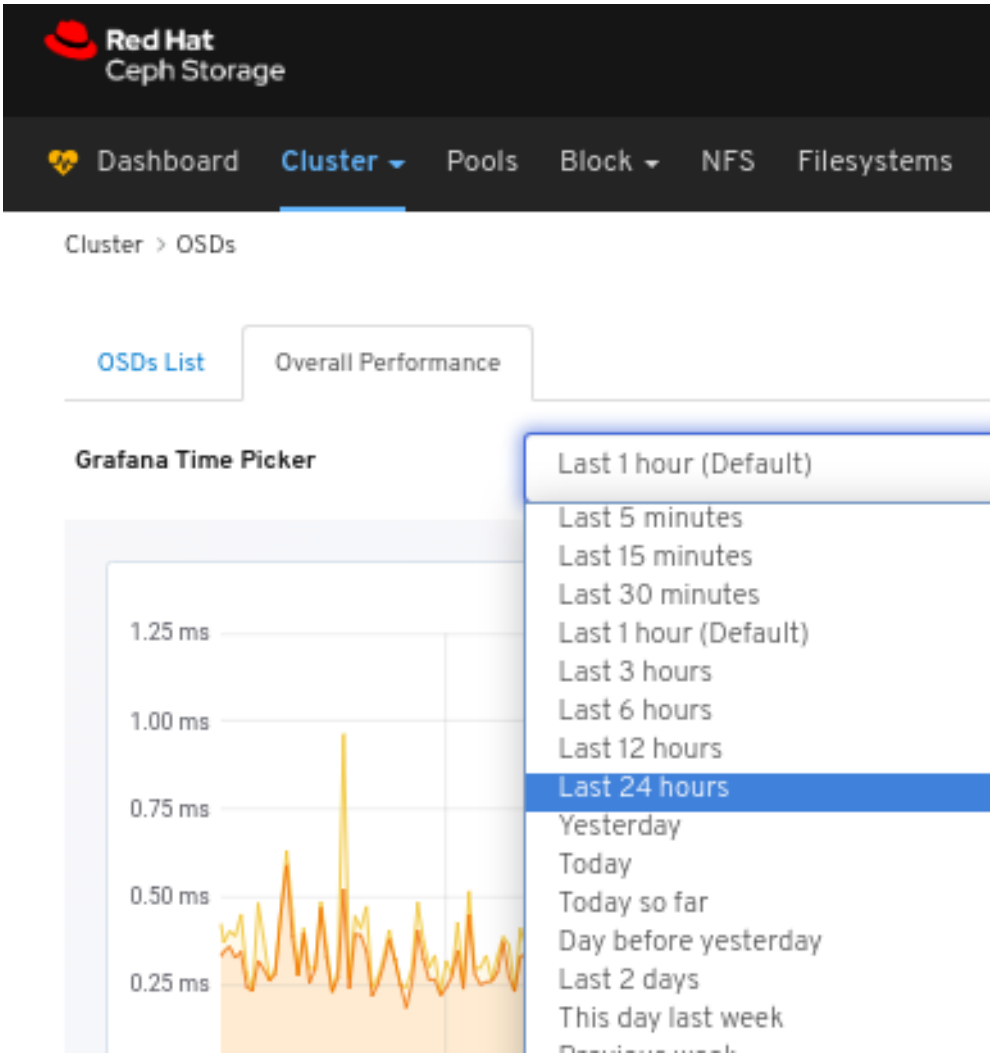

10. View the *Overall Performance* page:

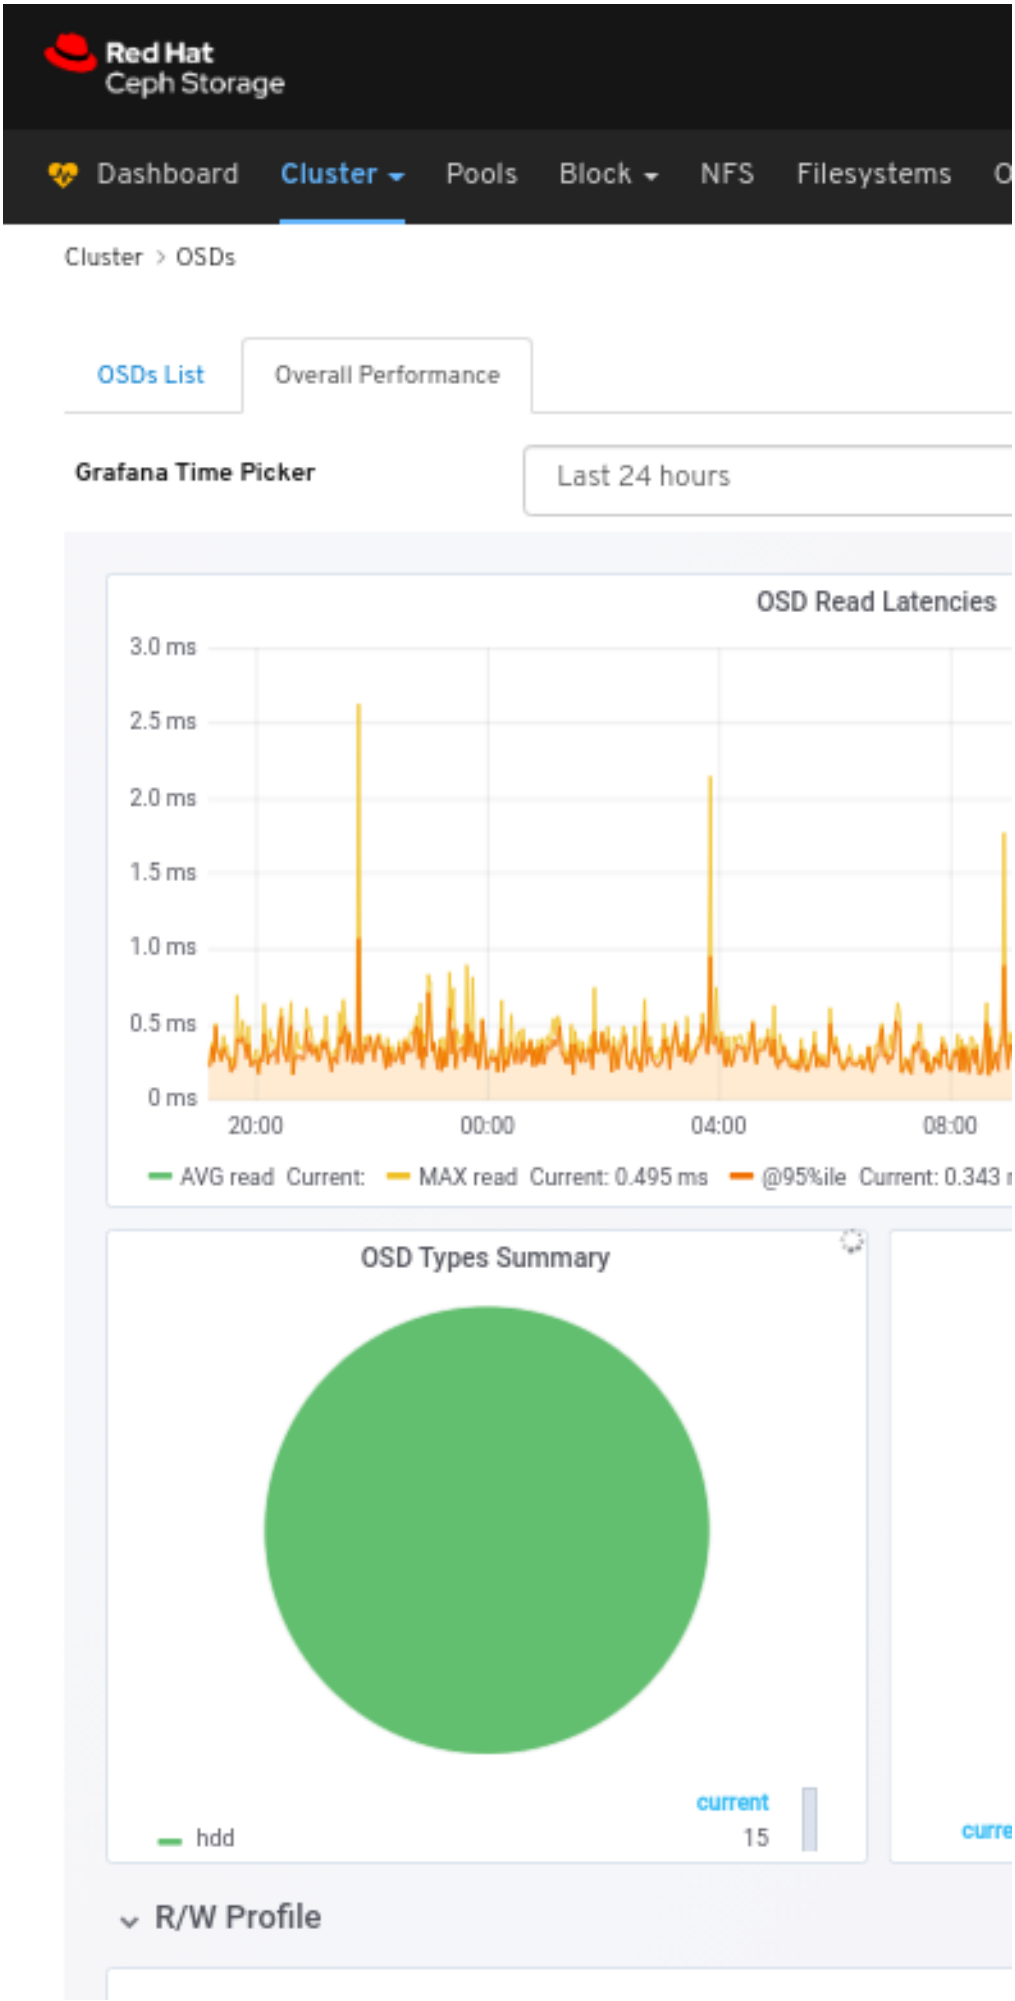
Red Hat
Ceph Storage
Dashboard
Cluster
Pools
Block
NFS
Filesystems
Cluster > OSDs
OSDs List
Overall Performance
Grafana Time Picker
Last 24 hours
OSD Read Latencies
3.0 ms
2.5 ms
2.0 ms
1.5 ms
1.0 ms
0.5 ms
0 ms
20:00
00:00
04:00
08:00
AVG read Current:
MAX read Current: 0.495 ms
@95%ile Current: 0.343
OSD Types Summary
current
hdd
15
R/W Profile

Additional Resources

- See the *Performance counters* in the *Red Hat Ceph Storage Administration Guide* for more details.

6.4. MONITORING POOLS

The Red Hat Ceph Storage Dashboard allows you to view various details about pools in the cluster.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Pools*:

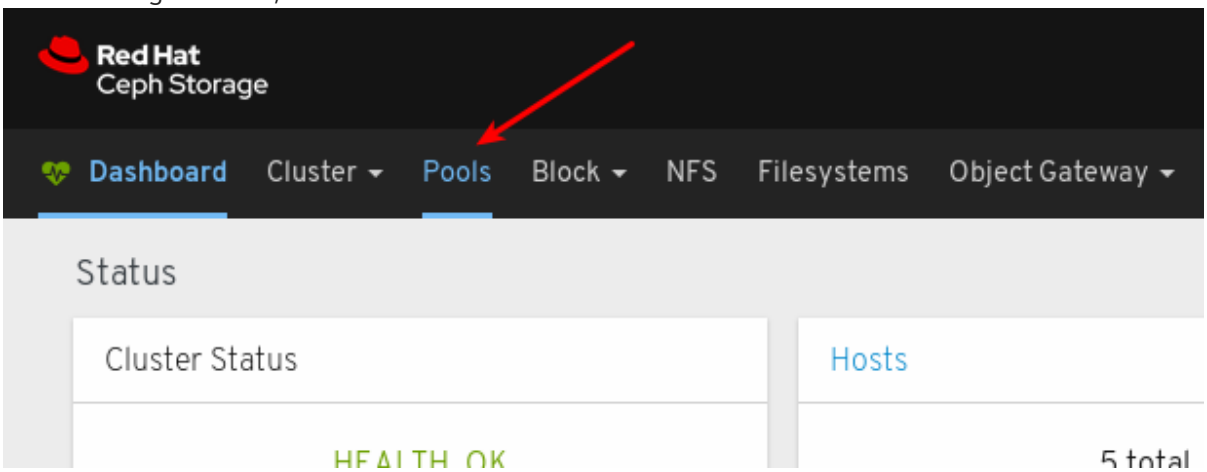

3. View the pools list:

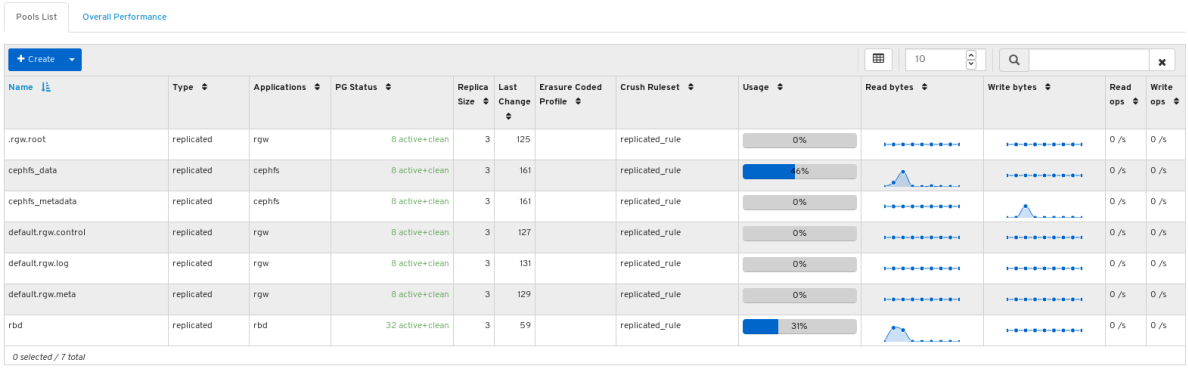

4. Mouse over a usage bar graph to view the actual used and free space:

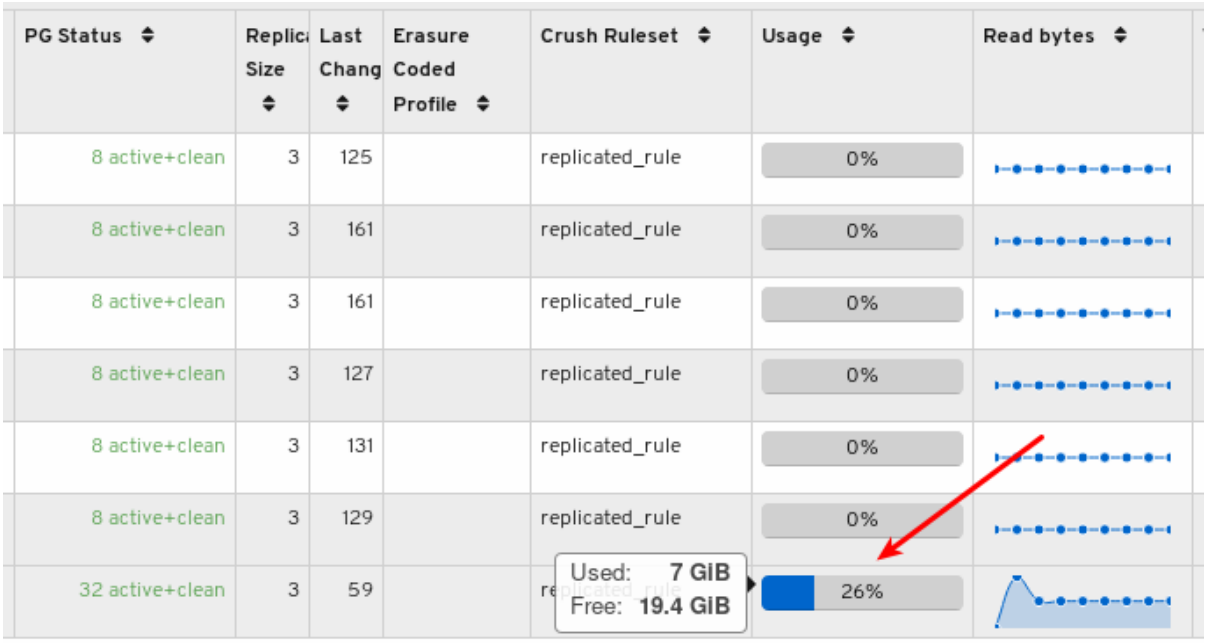

5. To view more information about a pool, select it by clicking on its row:

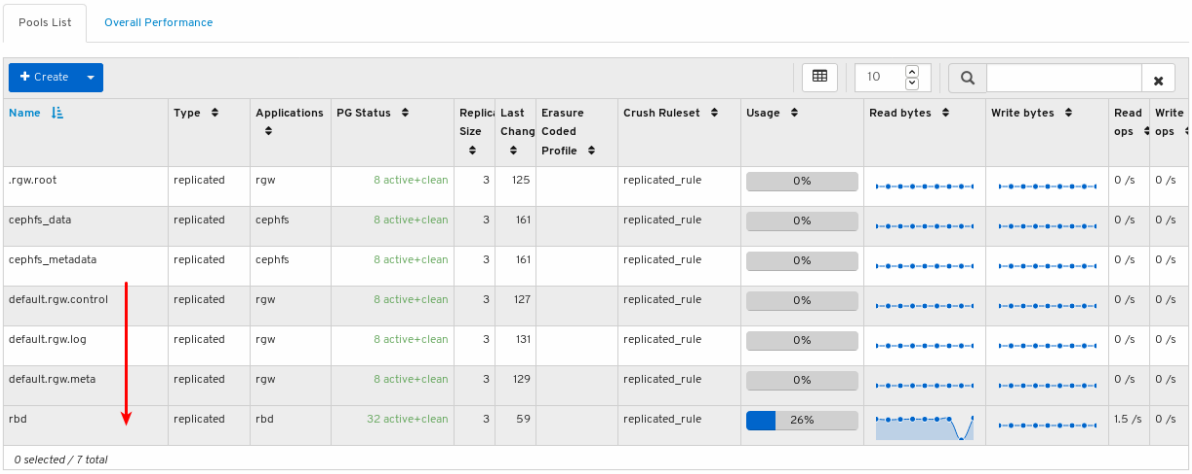

6. View the details of the pool. To view performance details and configuration data for the pool, click on the associated tabs.

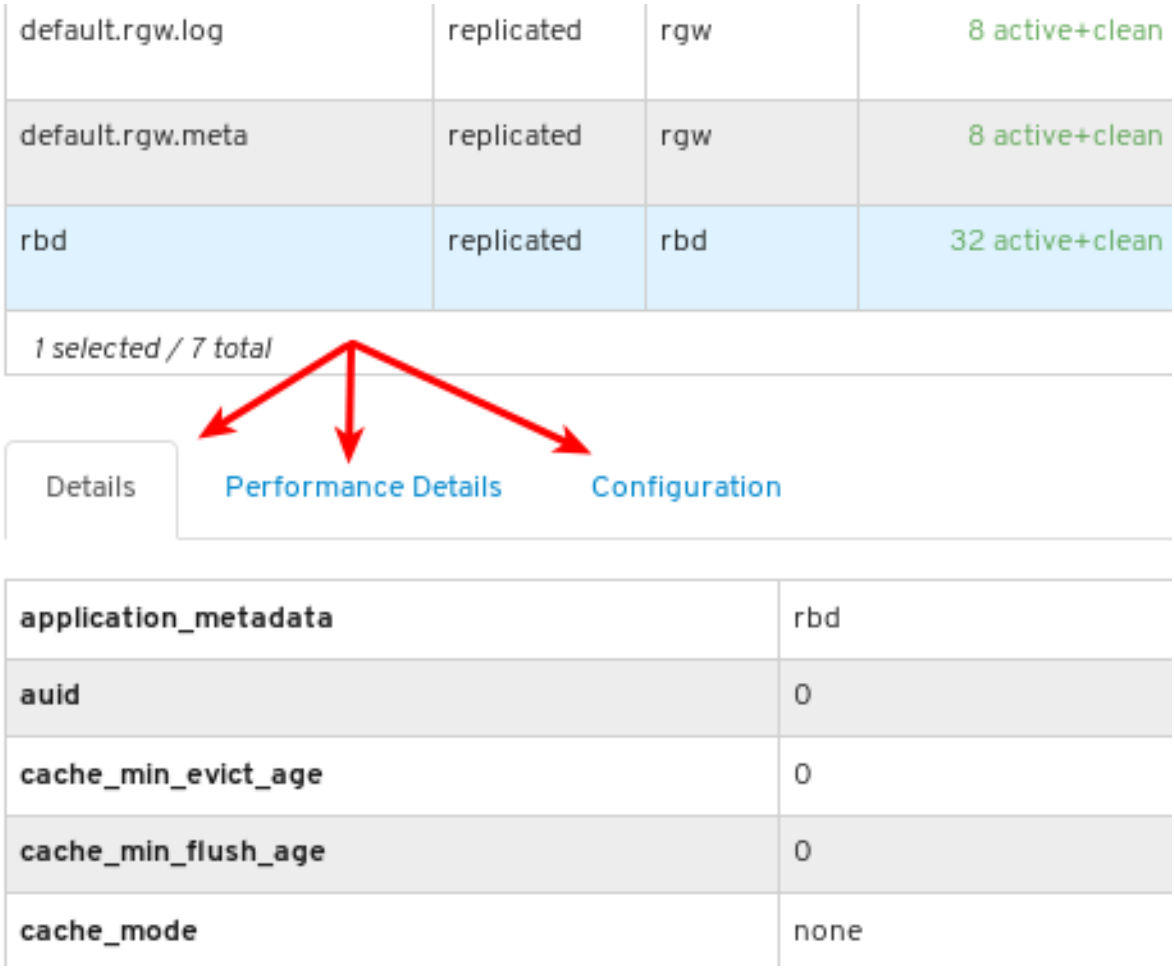

7. To view performance data for all the pools, click the *Overall Performance* tab towards the top left of the page:

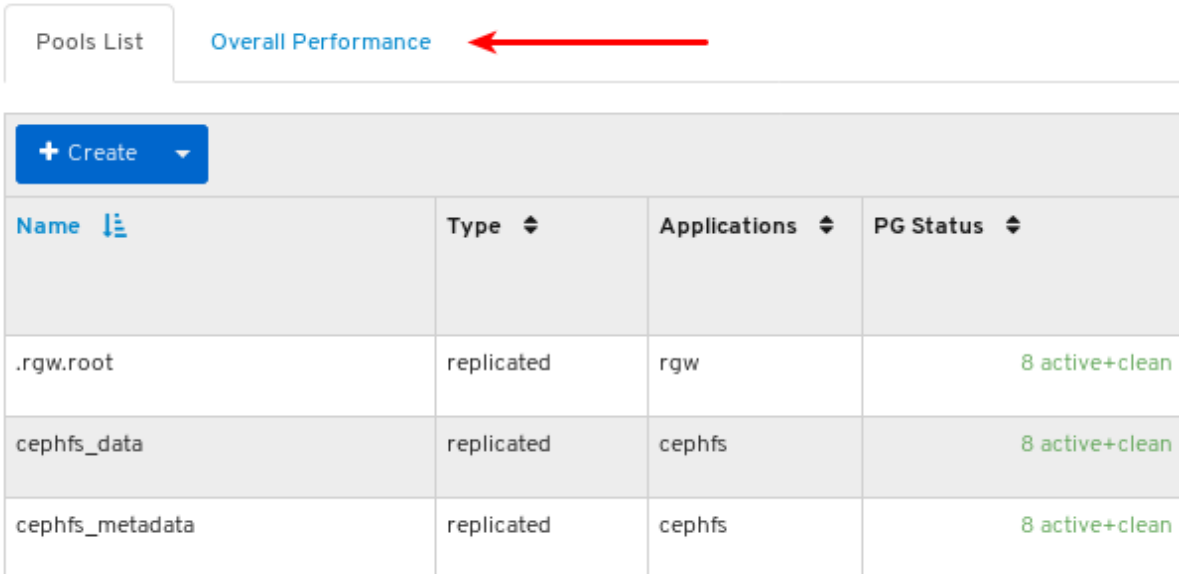

8. View the *Overall Performance* page:

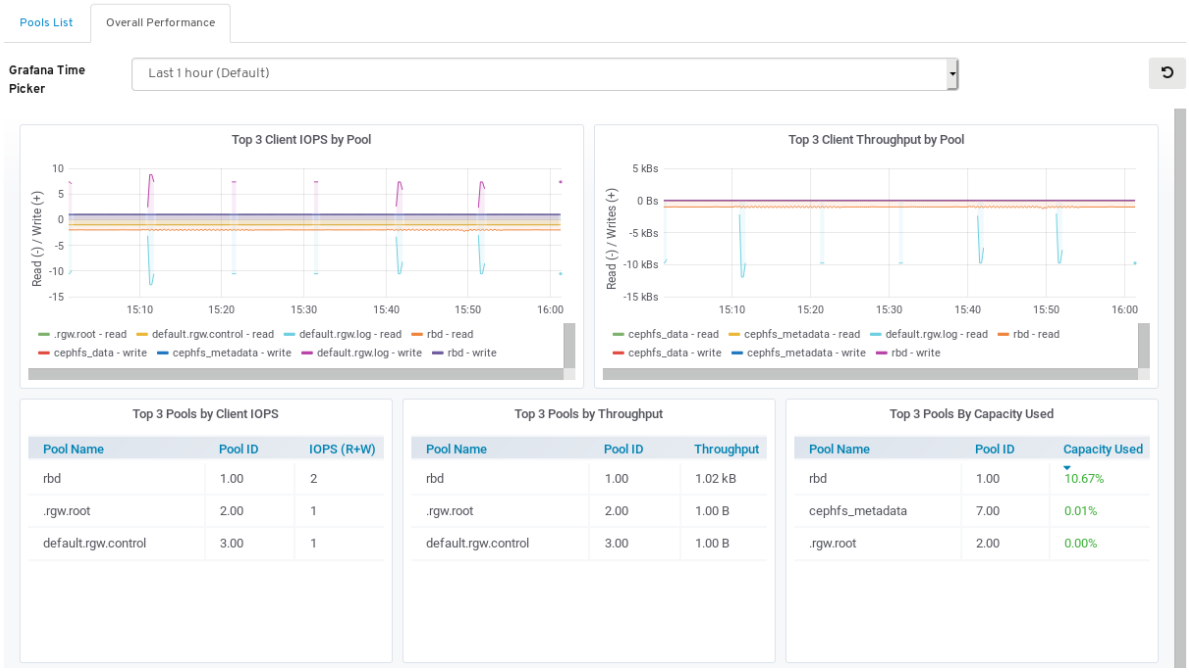

Additional Resources

- For more information about pools, see Pools in the Architecture guide.

6.5. MONITORING CEPH FILE SYSTEMS

As a storage administrator, you can use the Red Hat Ceph Storage Dashboard to monitor Ceph File Systems (CephFS) and related components.

Prerequisites

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

Procedure

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

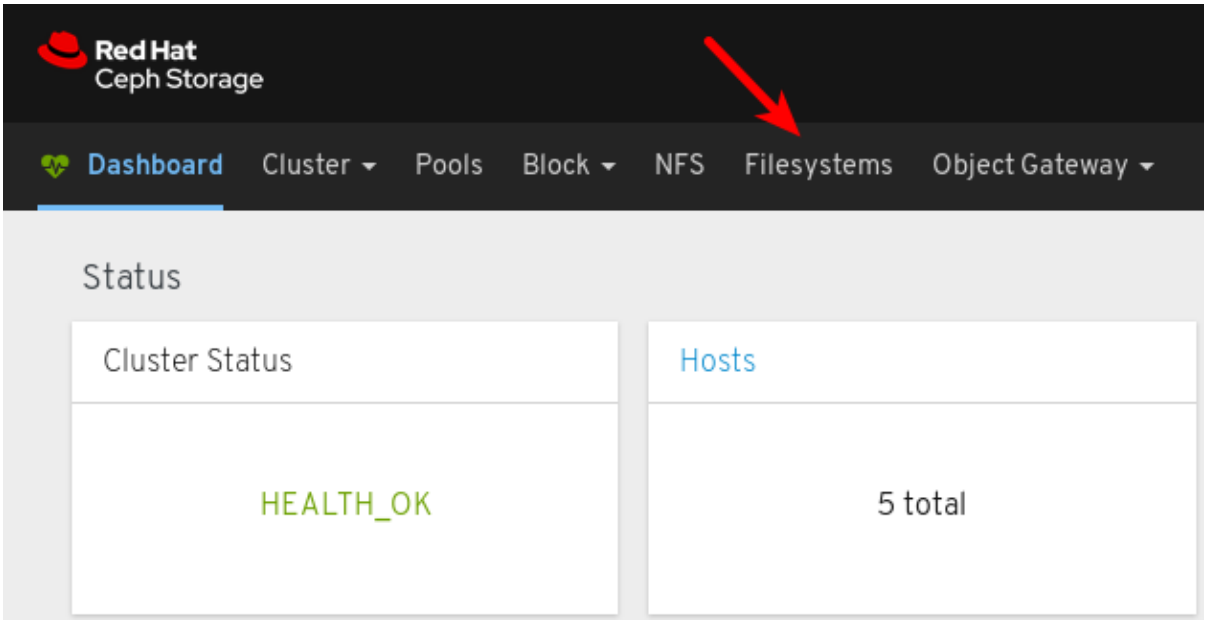

3. In the example below, you can see the *cephfs* file system.

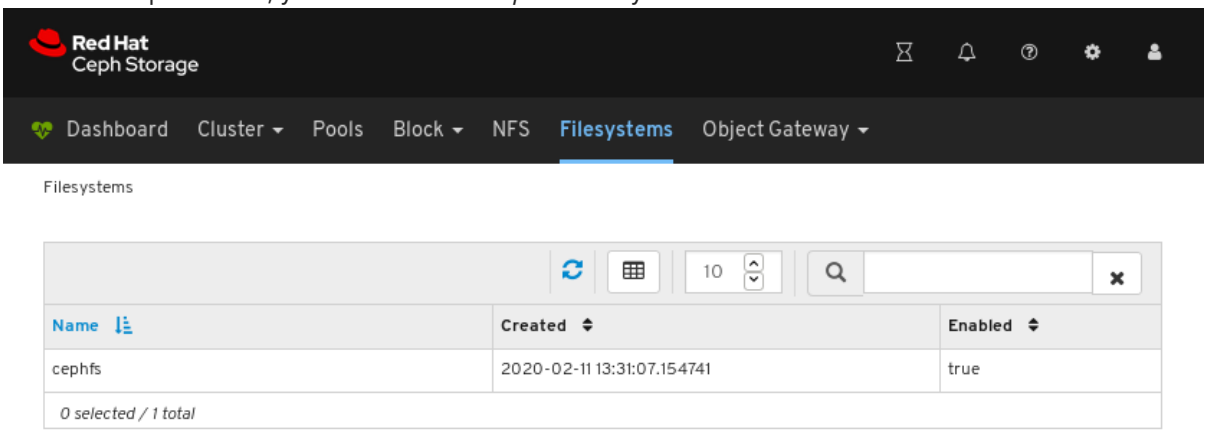

4. To view details for the file system, click the row for *cephfs*.

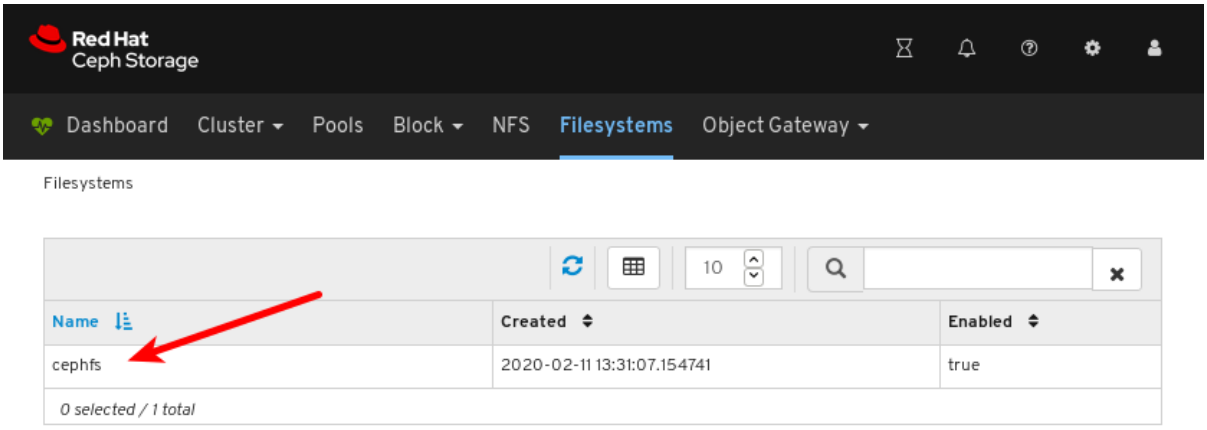

5. On the *Details* tab, you can see metadata servers and their rank plus any standby daemons, at **1**, pools and their usage, at **2**, and performance counters at **3**.

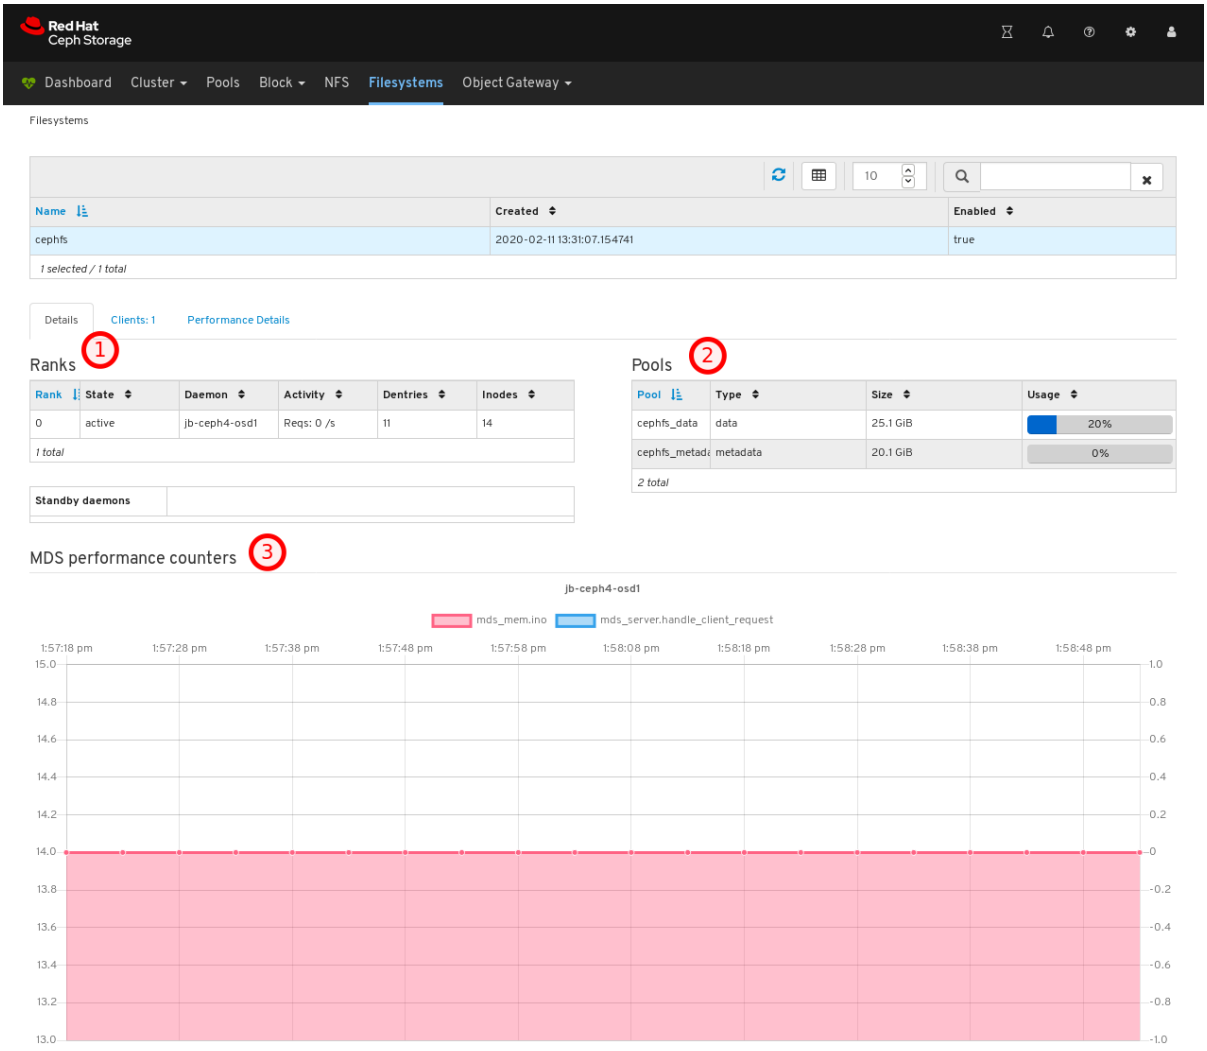

6. To view the list of clients which have mounted the file system, click the *Clients* tab.

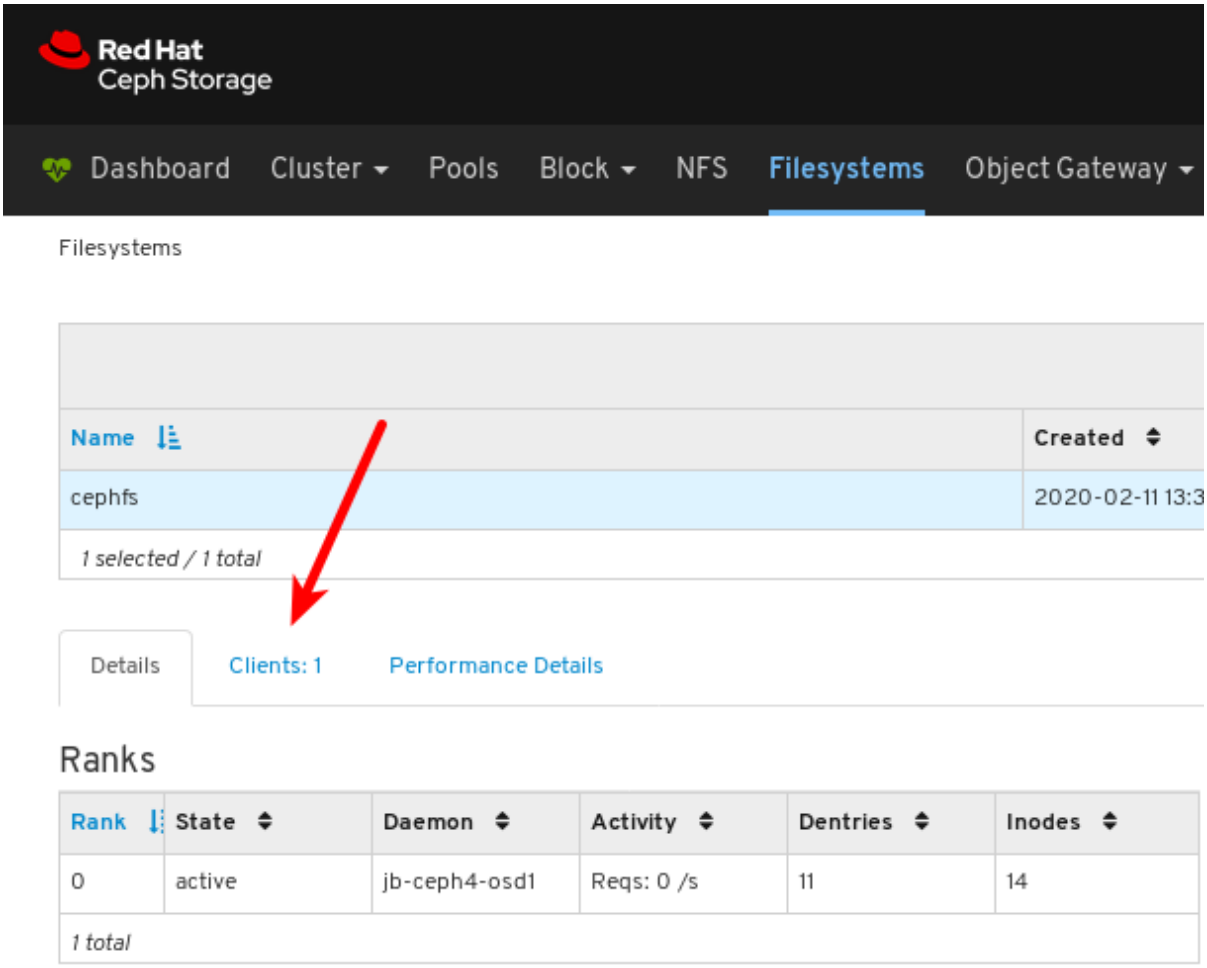

7. In the example below, you can see the **jb-ceph4-client** host has the *cephfs* file system opened.

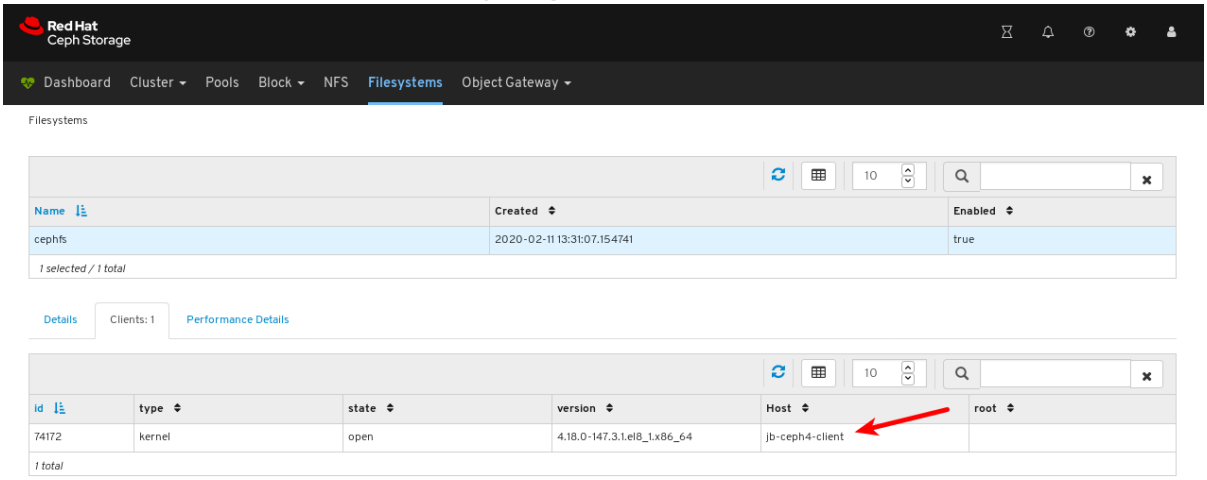

8. To view the performance of the file system, click the *Performance Details* tab.

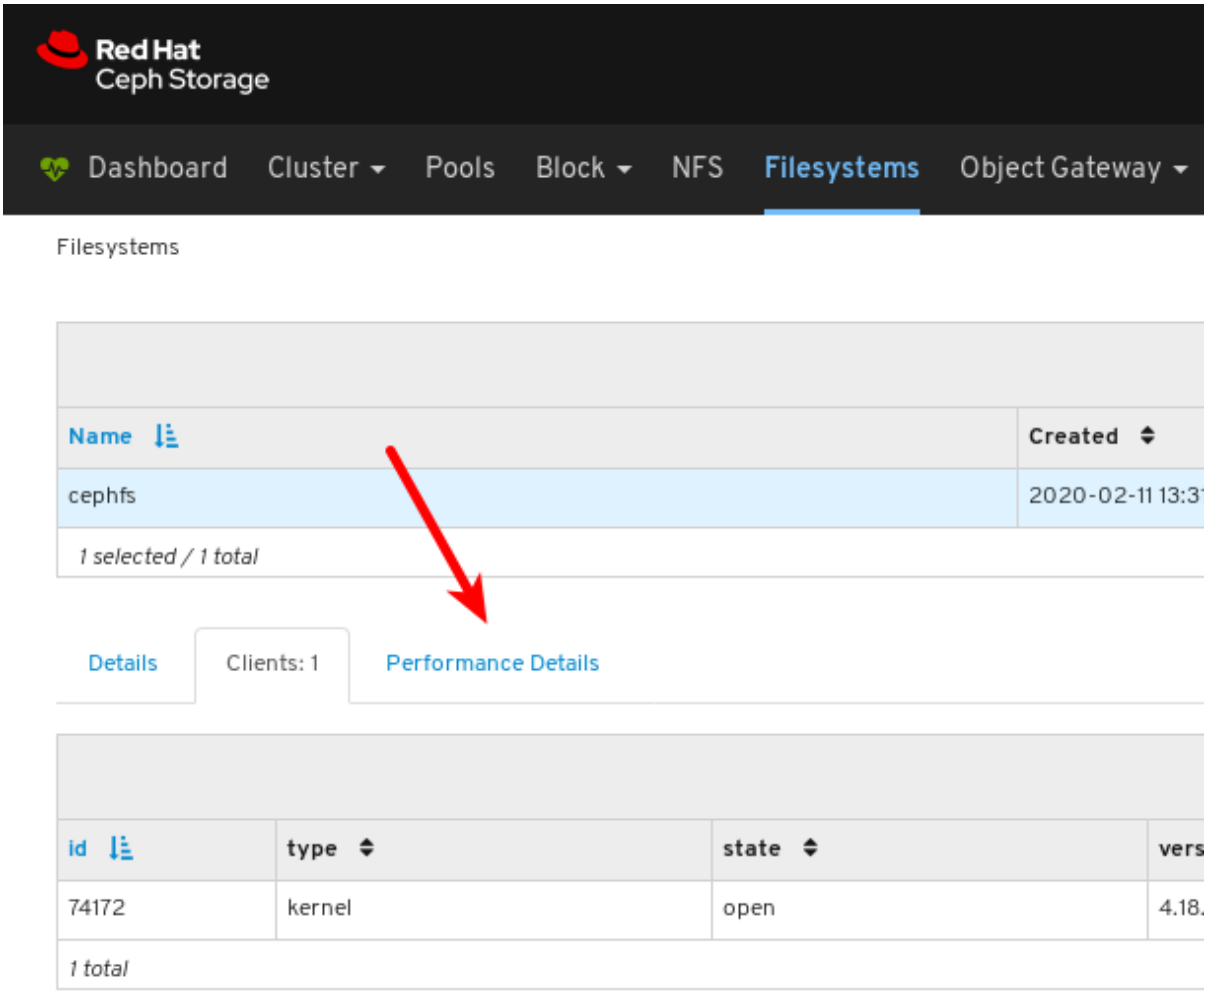

9. In the example below, you can see the read and write operations, at **1**, client requests, at **2**, and you can change the time range at **3**.

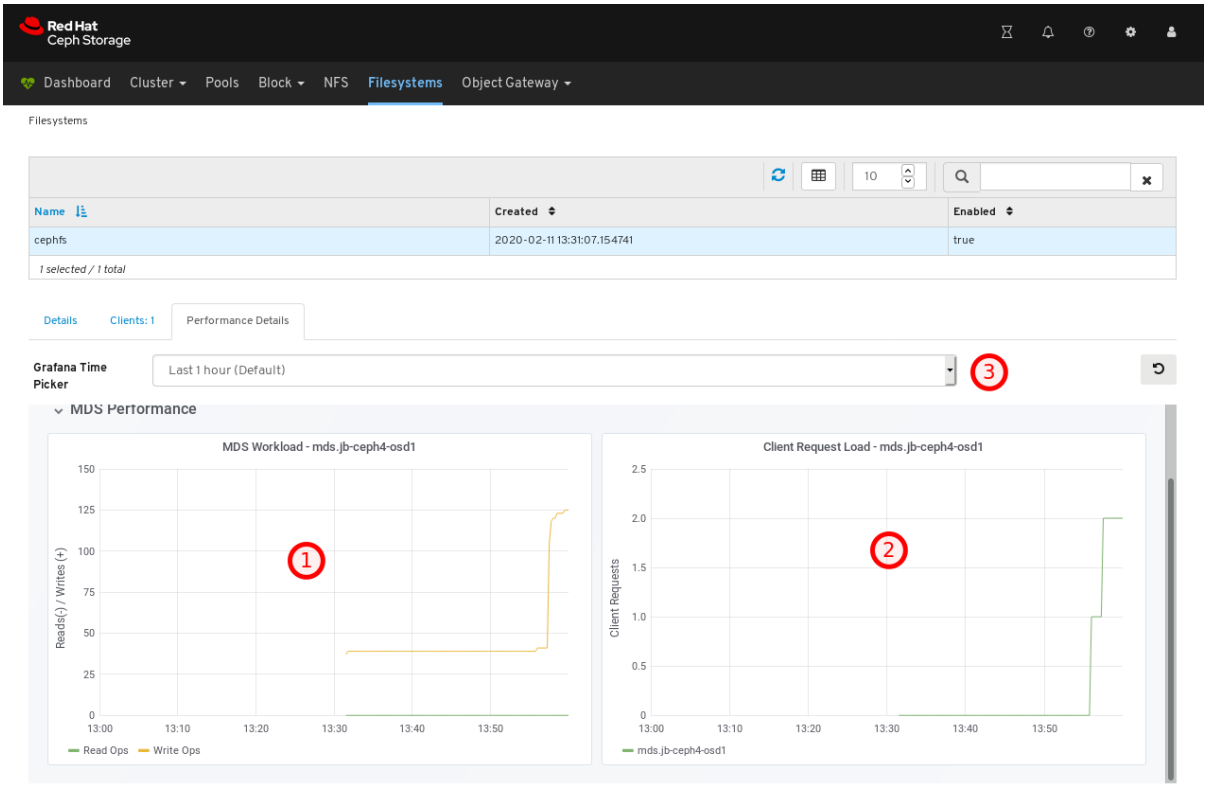

Additional Resources

- For more information, see Installing Metadata servers in the Installation Guide.
- For more information, see the File System Guide.

6.6. MONITORING OVERALL PERFORMANCE OF IMAGES

The Red Hat Ceph Storage Dashboard allows you to monitor the overall performance of the images in Blocks.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- Rados Block Device (RBD) pool is created.
- Image is created.
- Manager module **Prometheus** exporter is enabled for the pool.

Figure 6.1. Example

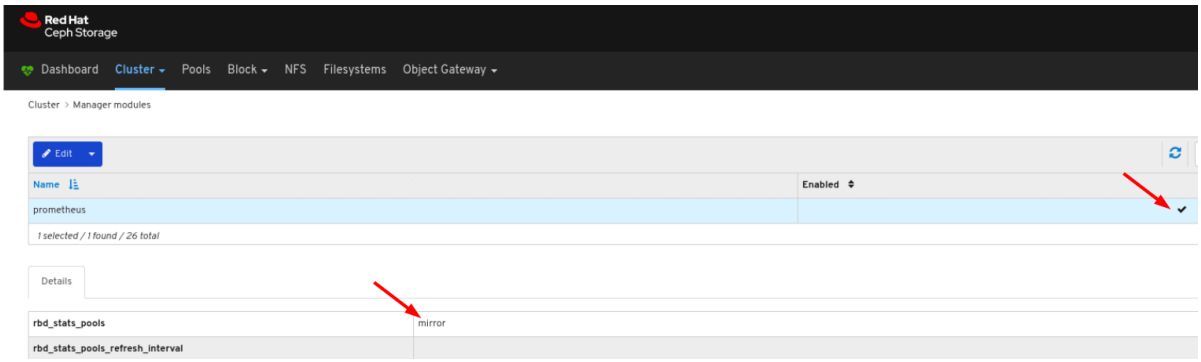

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

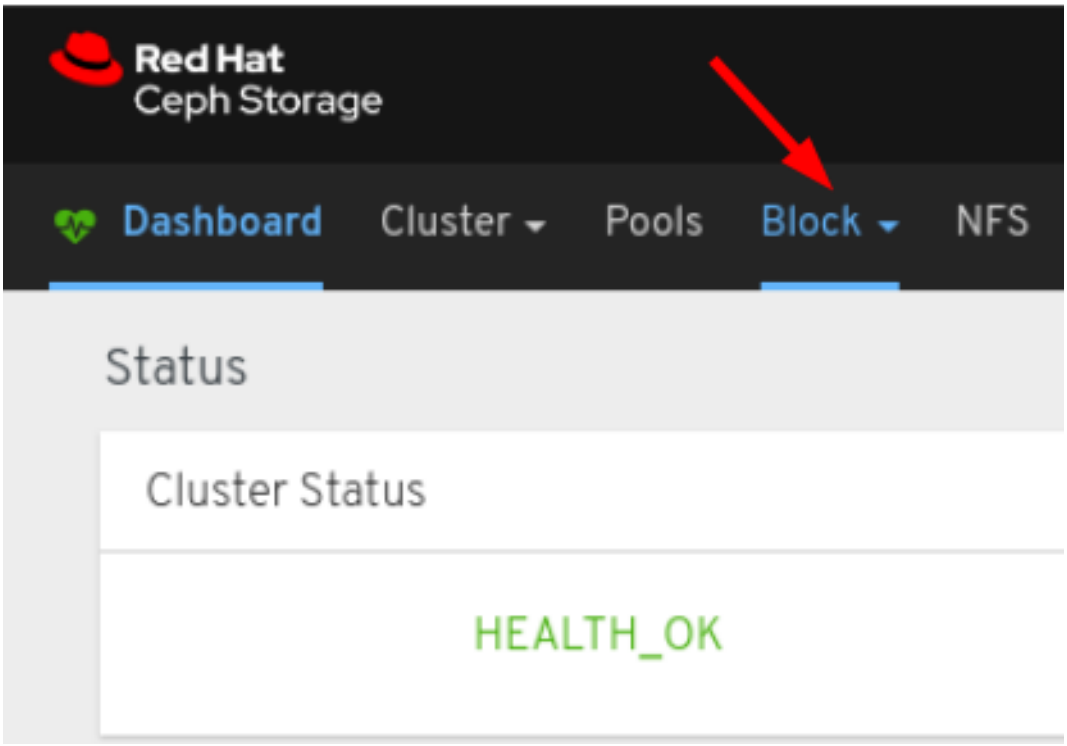

3. Select *Images* from the drop-down:

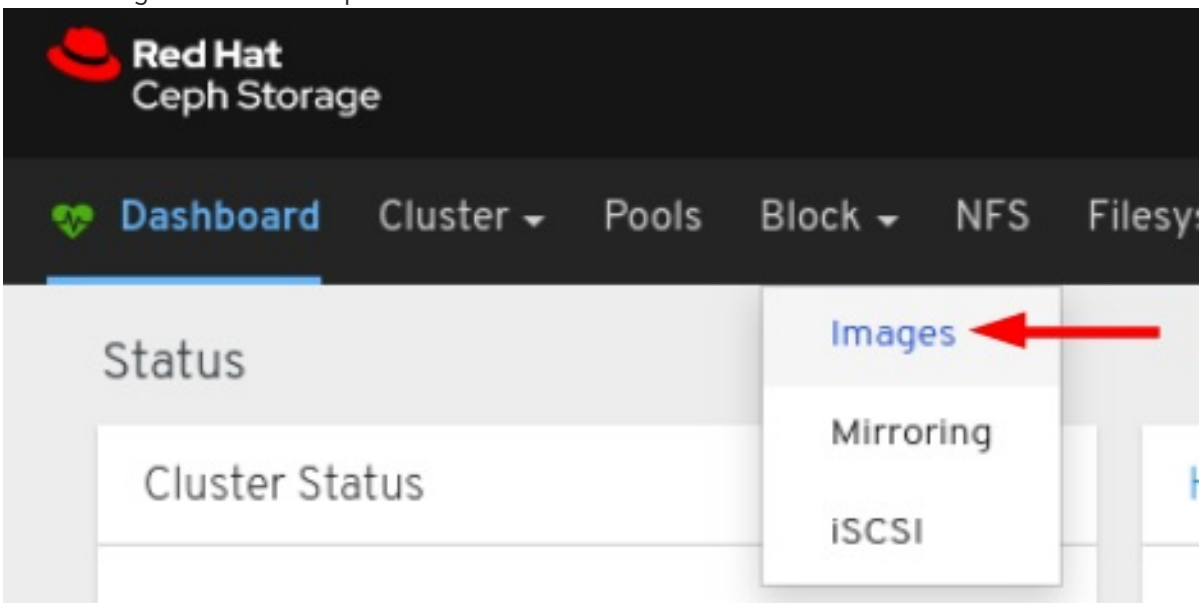

4. Select the *Overall Performance* tab:

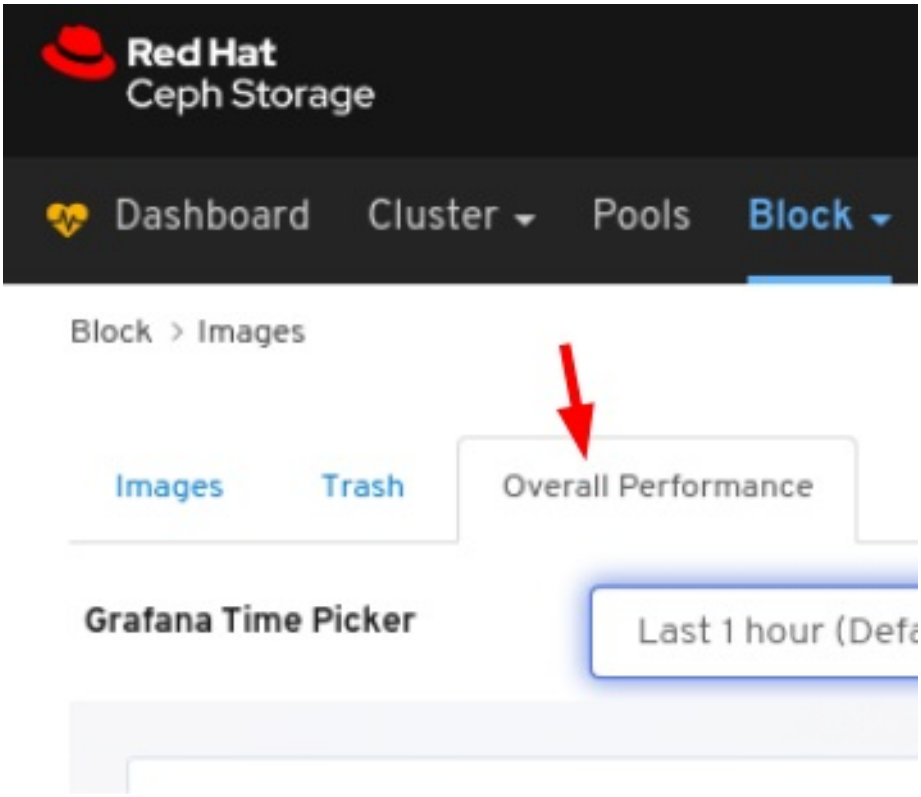

5. Select the time range from the *Grafana Time Picker* drop-down:

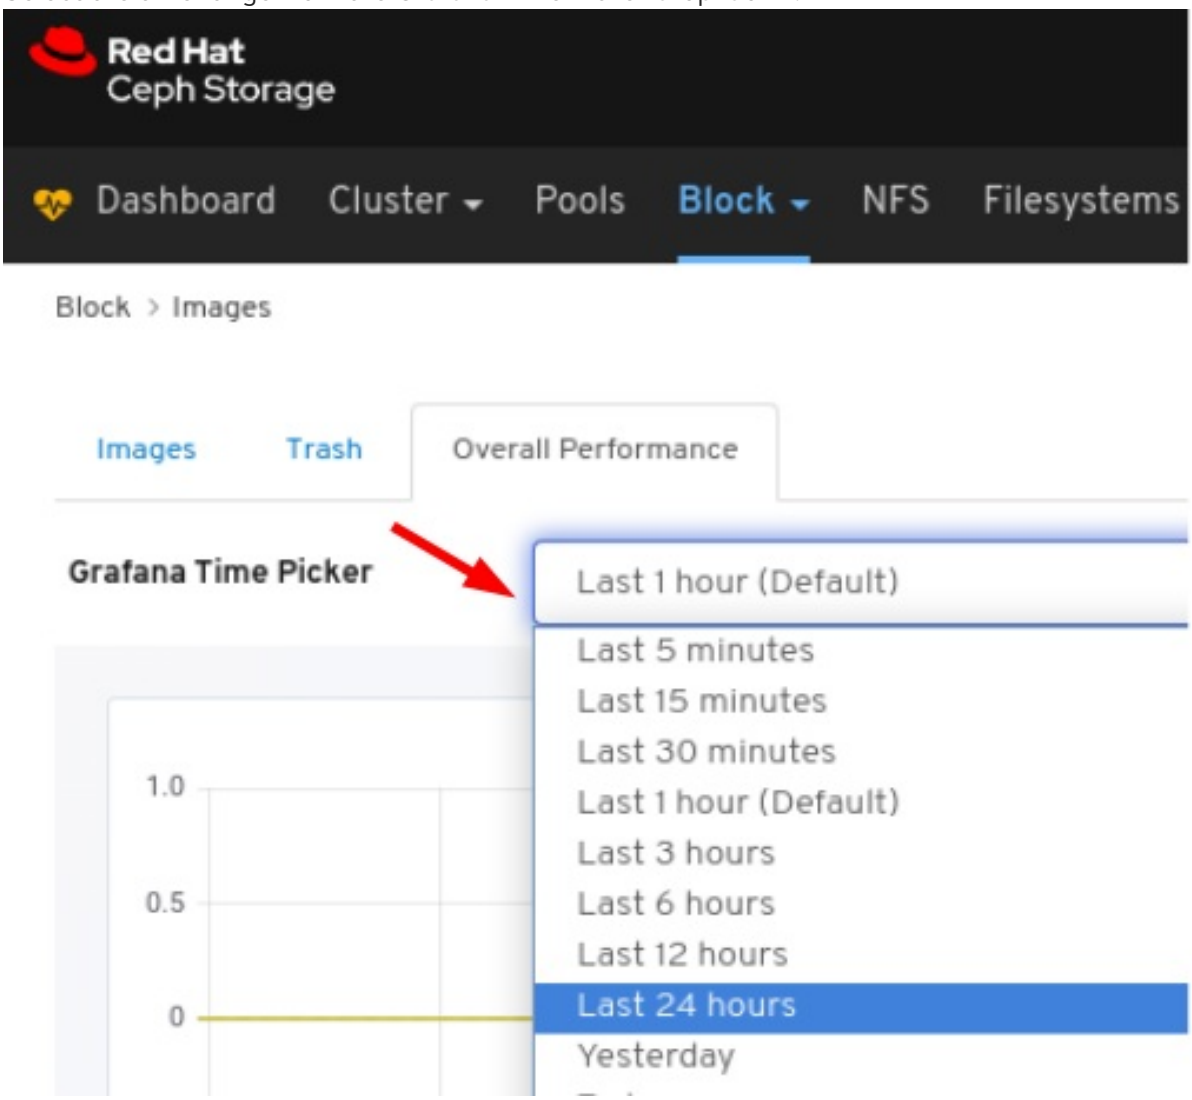

6. View the *Overall Performance* page:

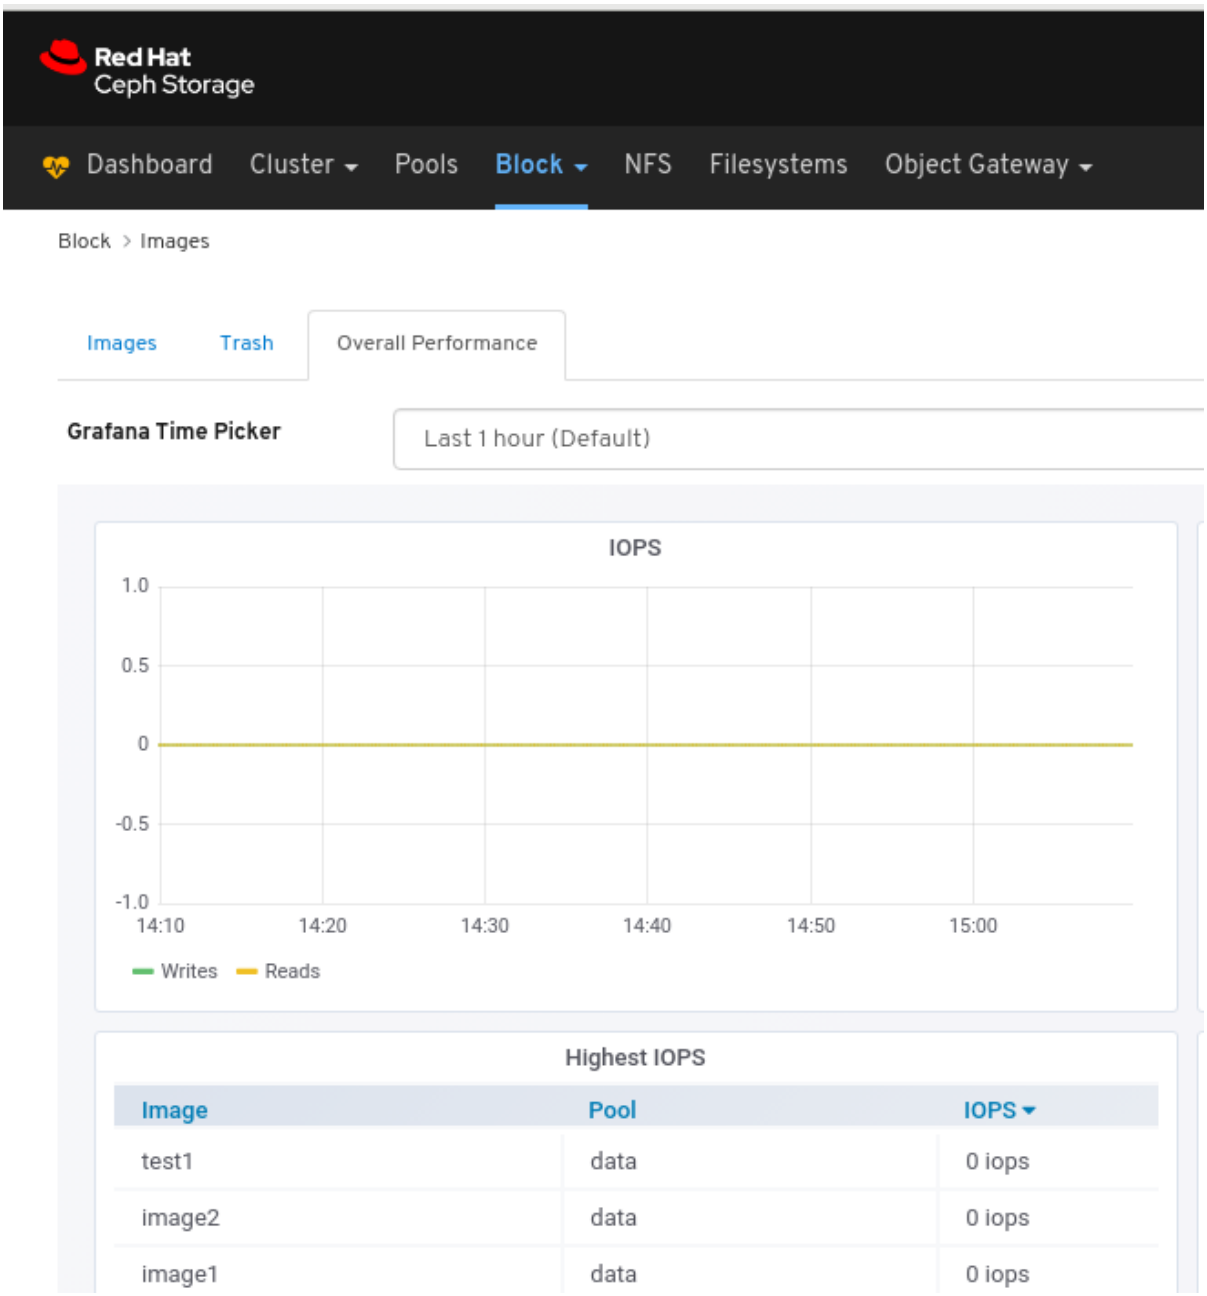
Red Hat
Ceph Storage
Dashboard
Cluster
Pools
Block
NFS
Filesystems
Object Gateway
Block > Images
Images
Trash
Overall Performance
Grafana Time Picker
Last 1 hour (Default)
IOPS
1.0
0.5
0
-0.5
-1.0
14:10
14:20
14:30
14:40
14:50
15:00
Writes
Reads
Highest IOPS
Image
Pool
IOPS
test1
data
0 iops
image2
data
0 iops
image1
data
0 iops

CHAPTER 7. MANAGING THE CLUSTER

7.1. MANAGING THE CLUSTER

The management functions of the dashboard allow you to view and modify configuration settings, and manage cluster resources.

7.2. VIEWING THE CRUSH MAP

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

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

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster*.
3. Click *CRUSH map*.

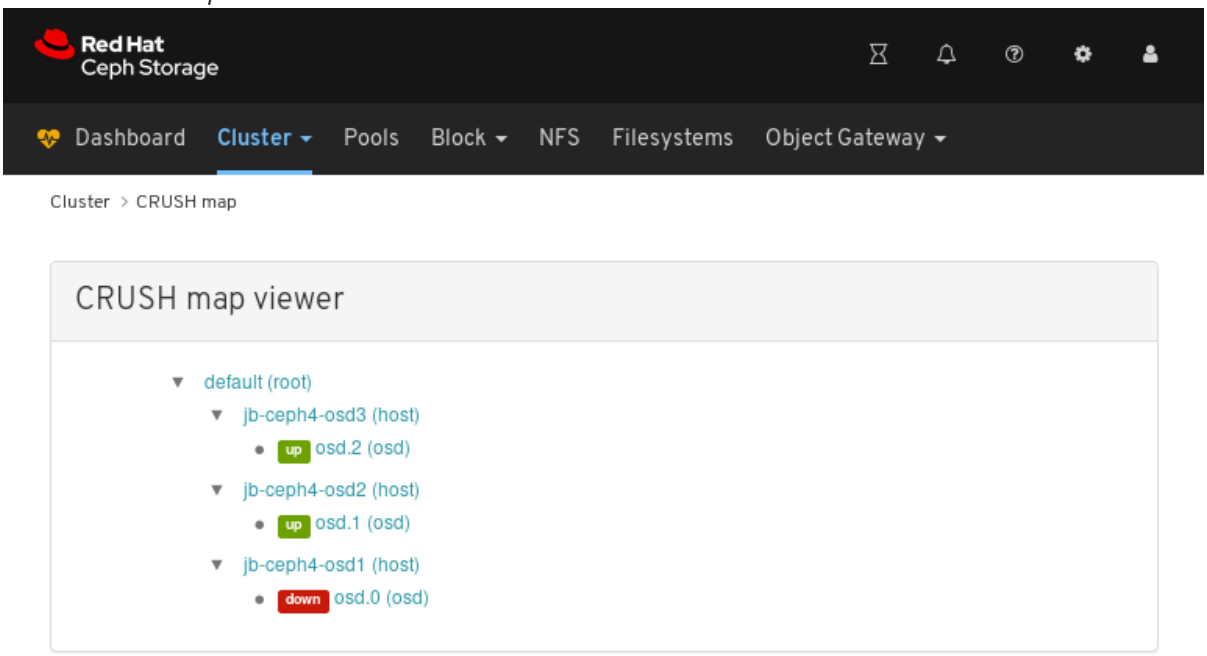

In the above example, you can see the default CRUSH map, three nodes, and OSDs running on two of the three nodes.

4. Click on the CRUSH map name, nodes, or OSDs, to view details about each object.

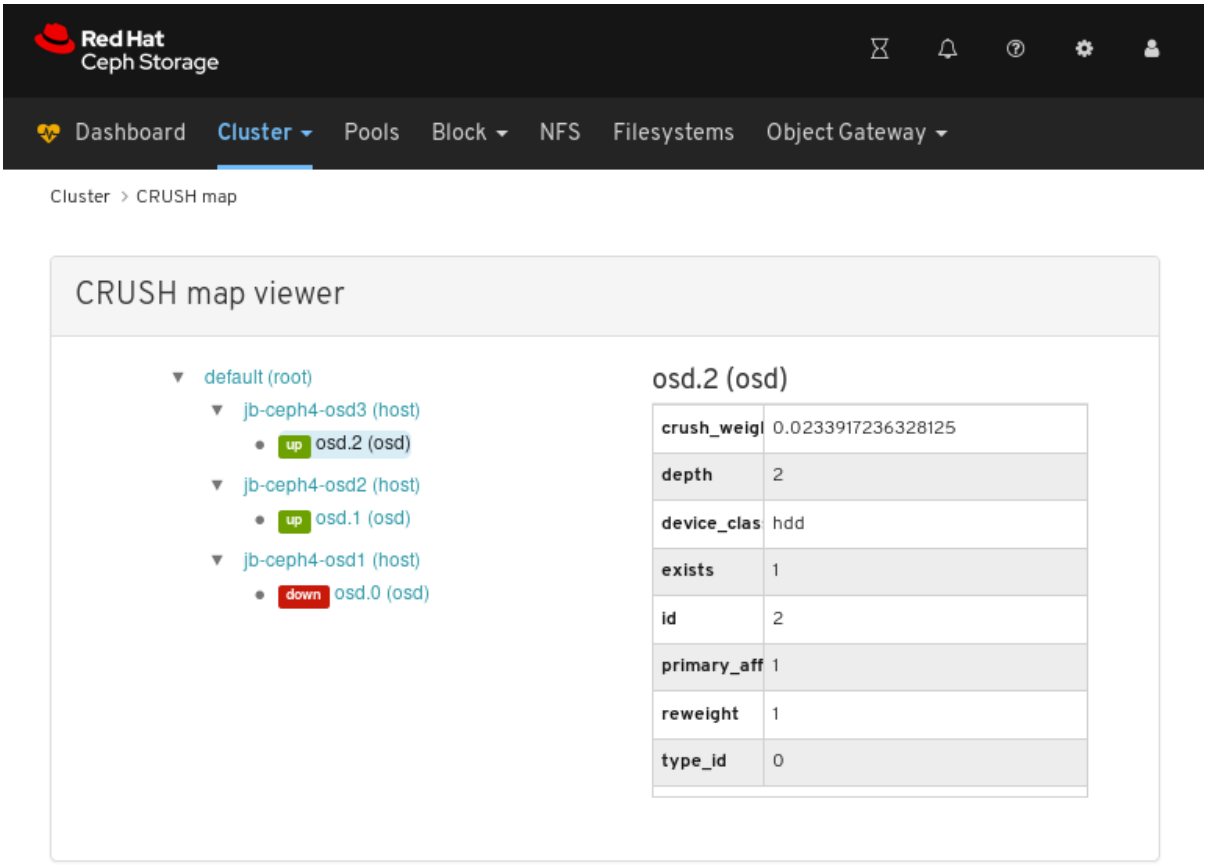

In the above example, you can see the values of variables related to an OSD running on the **jb-rhel-osd3** node. In particular, note the **id** is **2**.

Additional Resources

- For more information about the CRUSH map, see CRUSH administration in the Storage strategies guide.

7.3. CONFIGURING MANAGER MODULES

The Red Hat Ceph Storage dashboard allows you to view and configure manager module parameters.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster*.
3. Click *Manager modules*:

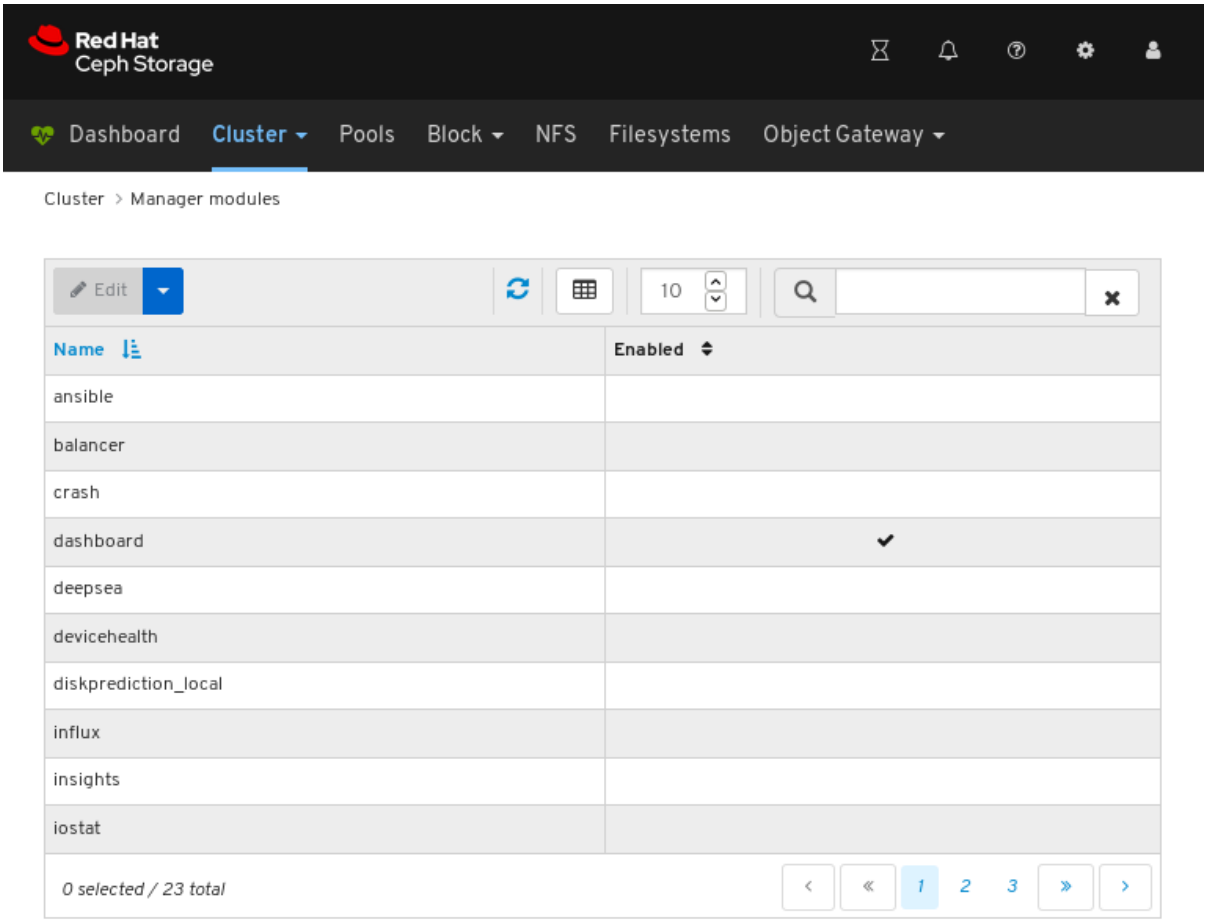

The above screenshot shows the first of three pages of manager modules.

4. Click on a row for a module you want to configure:

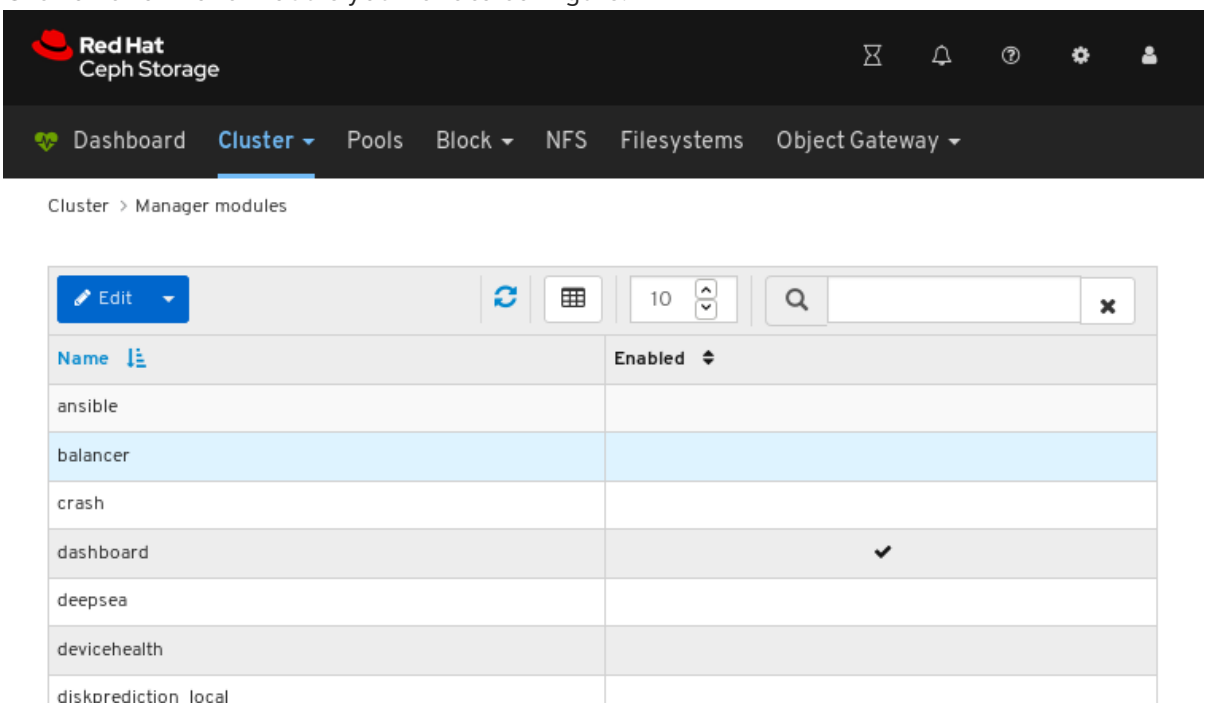

Not all modules have configurable parameters. If a module is not configurable, the *Edit* button is disabled.

5. Towards the upper left of the page, click the *Edit* button to load the page with the configurable parameters.

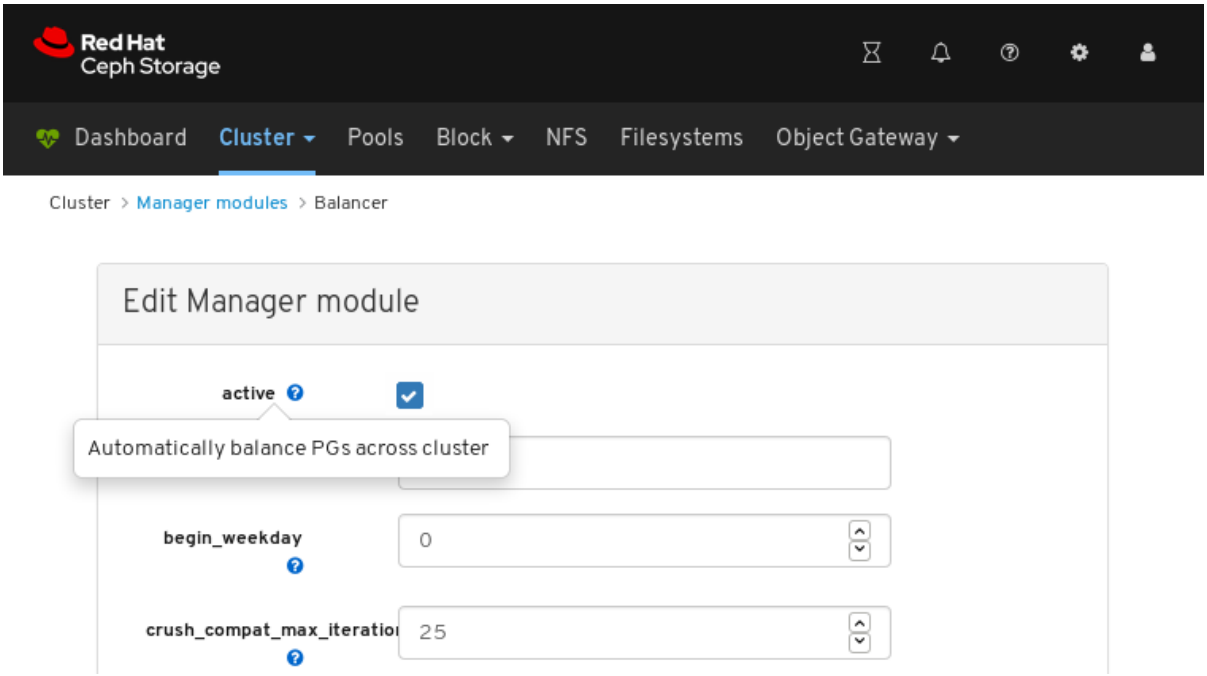

The above screenshot shows parameters that you can change for the balancer module. To display a description of a parameter, click the question mark button.

6. To change a parameter, modify the state of the parameter and click the *Update* button at the bottom of the page:

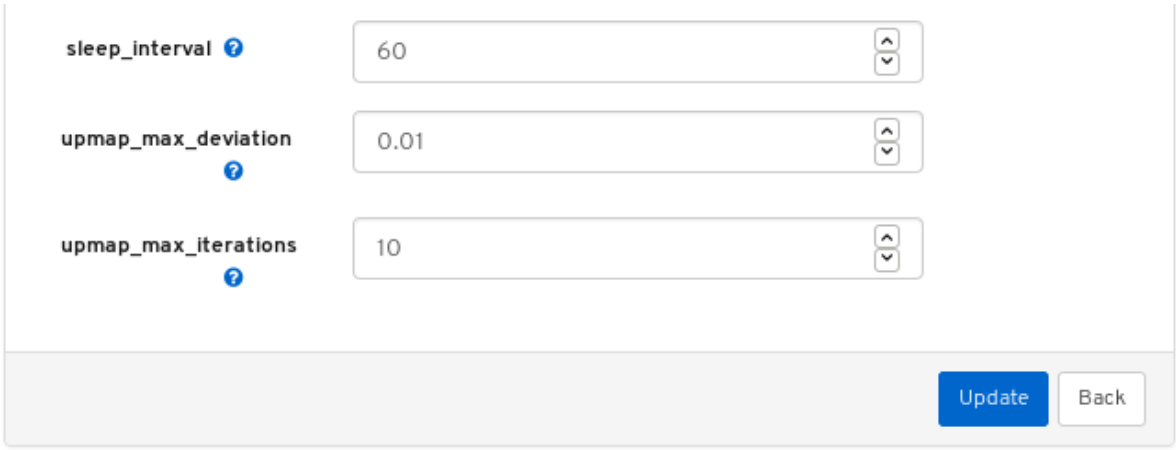

A notification confirming the change appears in the upper-right corner of the page:

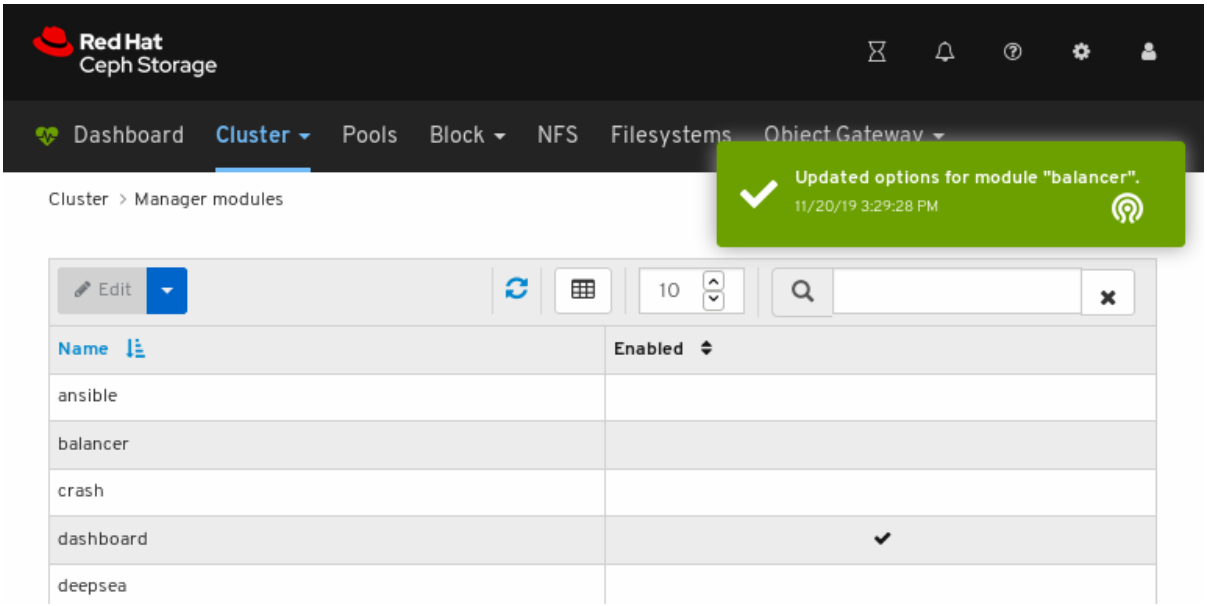

Additional Resources

- See *Using the Ceph Manager balancer module* in the *Red Hat Ceph Storage Operations Guide*.

7.4. FILTERING LOGS

The Red Hat Ceph Storage Dashboard allows you to view and filter logs based on several criteria. The criteria include *priority*, *keyword*, *date*, and *time range*.

Prerequisites

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

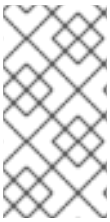

NOTE

The Dashboard logging feature only displays the thirty latest high level events. The events are stored in memory by the Monitor. The entries disappear after restarting the Monitor. If you need to review detailed or older logs, refer to the file based logs. See *Additional Resources* below for more information about file based logs.

Procedure

1. Log in to the Dashboard.
2. Click the **Cluster** drop-down menu in the top navigation bar.
3. Click **Logs** in the drop-down menu.
4. View the last thirty unfiltered log entries.

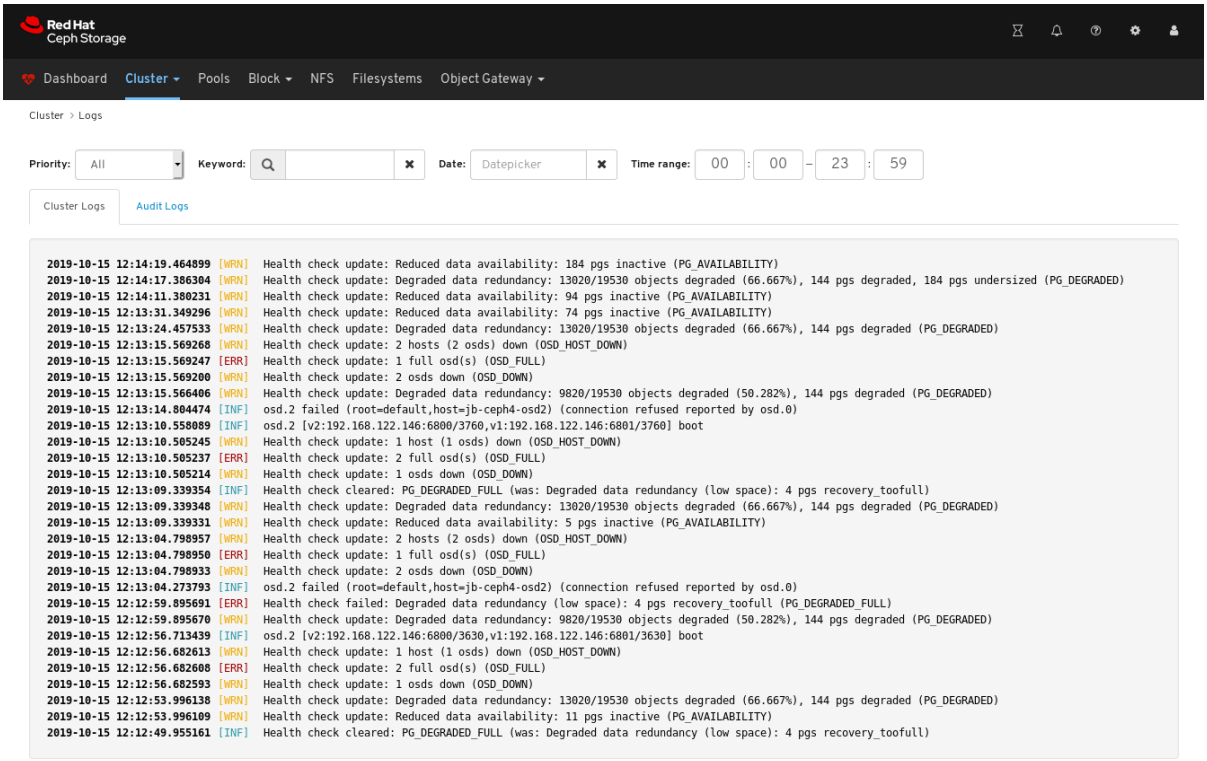

a. To filter by priority, click the **Priority** drop-down menu and select either **Info**, **Warning**, or **Error**. The example below only shows log entries with the priority of **Error**.

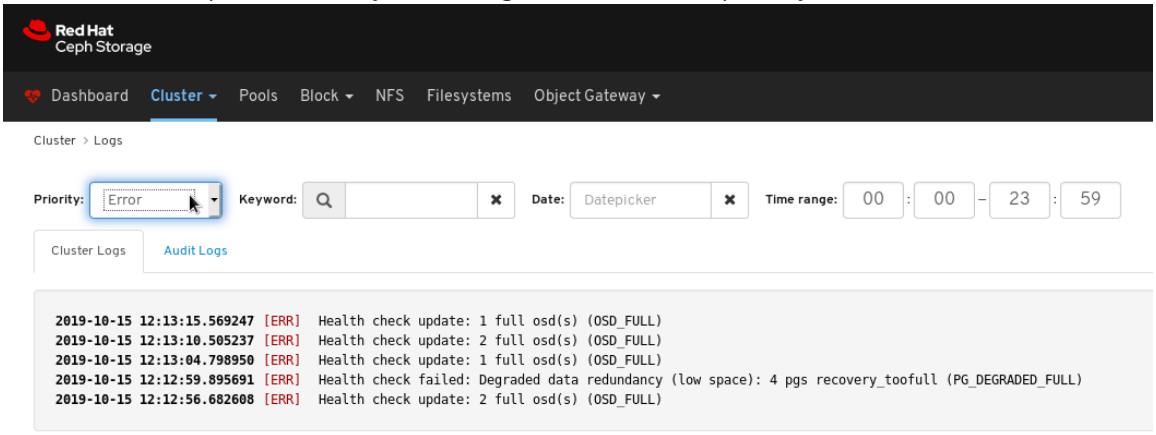

b. To filter by keyword, enter text into the **Keyword** form. The example below only shows log entries that include the text **osd.2**.

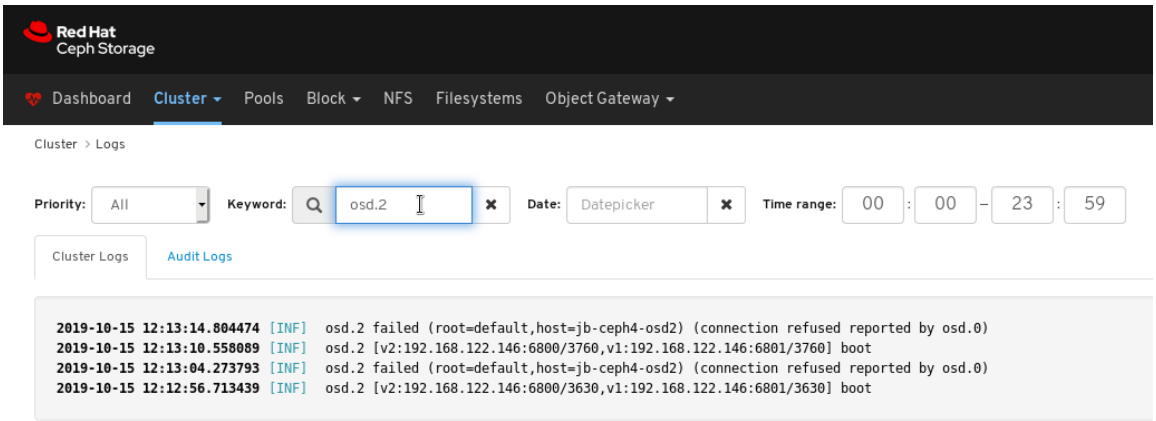

c. To filter by date, click the **Date** form and either use the date picker to select a date from the menu, or enter a date in the form of *YYYY-MM-DD*. The example below only shows log entries with the date of **2019-10-15**.

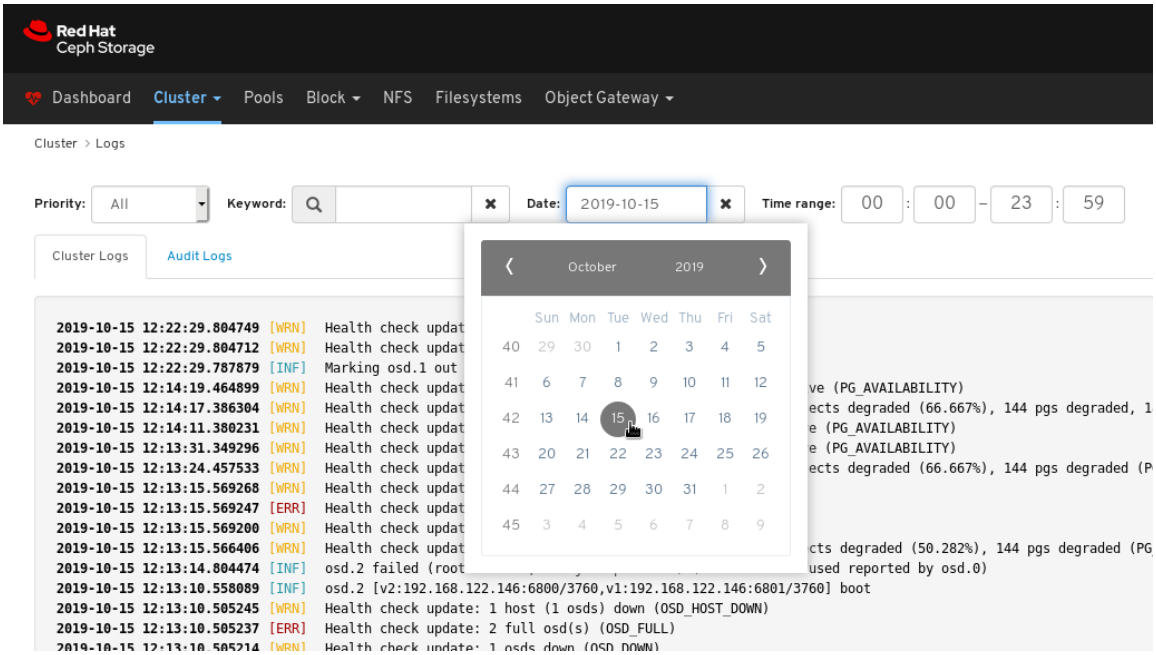

d. To filter by time, enter a range in the **Time range** fields using the *HH:MM - HH:MM* format. Hours must be entered using numbers **0** to **23**. The example below only shows log entries from **12:14** to **12:23**.

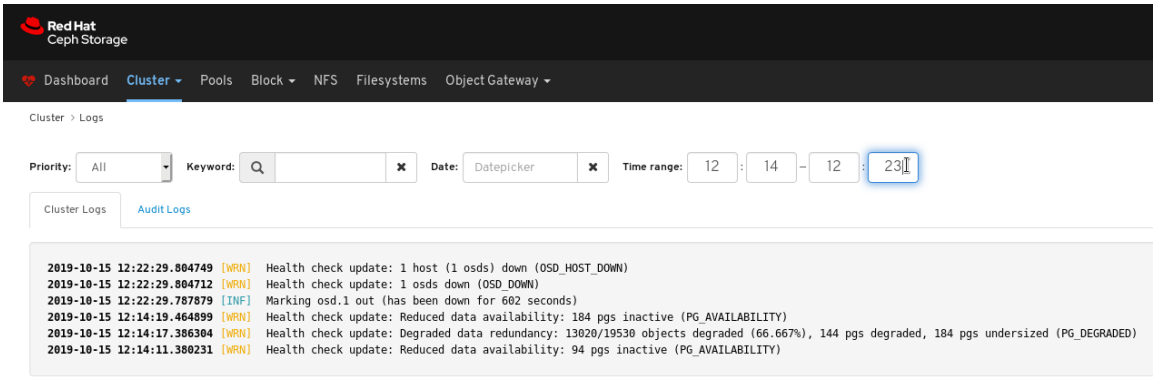

e. To combine filters, set two or more filters. The example below only shows entries that have both a Priority of **Warning** and the keyword of **osd**.

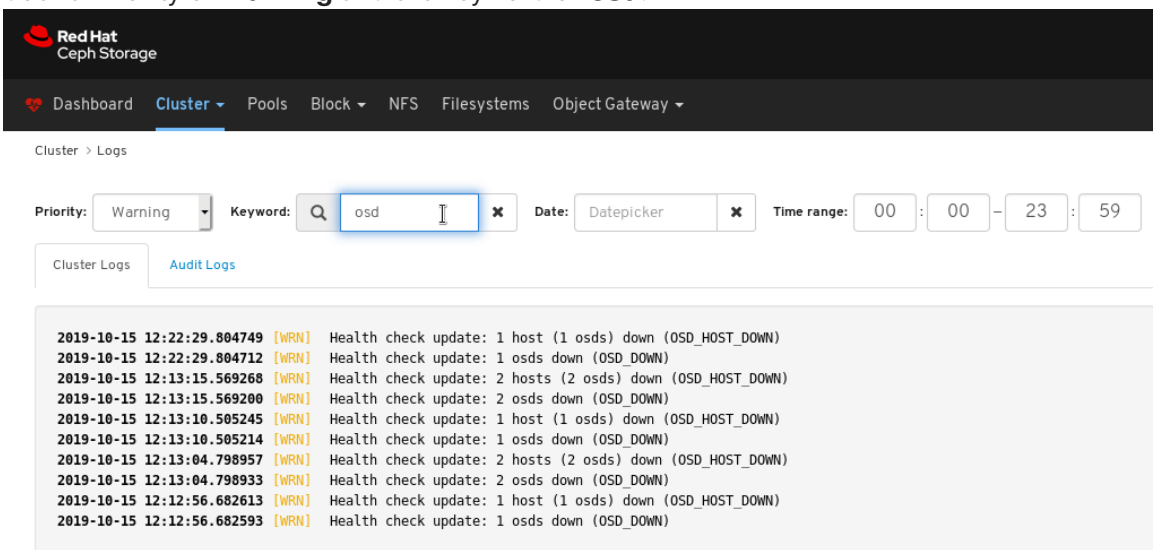

Additional Resources

- See the Configuring Logging section in the Troubleshooting Guide for more information.
- See the Understanding Ceph Logs section in the Troubleshooting Guide for more information.

7.5. CONFIGURING OSD RECOVERY SETTINGS

As a storage administrator, you can change the OSD recovery priority and customize how the cluster recovers. This allows you to influence your cluster's rebuild performance or recovery speed.

Prerequisites

- A Red Hat Ceph Storage cluster.
- The dashboard is installed.

Procedure

1. Log in to the dashboard.
2. Click the *Cluster* drop-down menu in the top navigation bar.
3. Click *OSDs* in the drop-down menu.
4. Click the *Cluster-Wide Flags* drop-down menu.

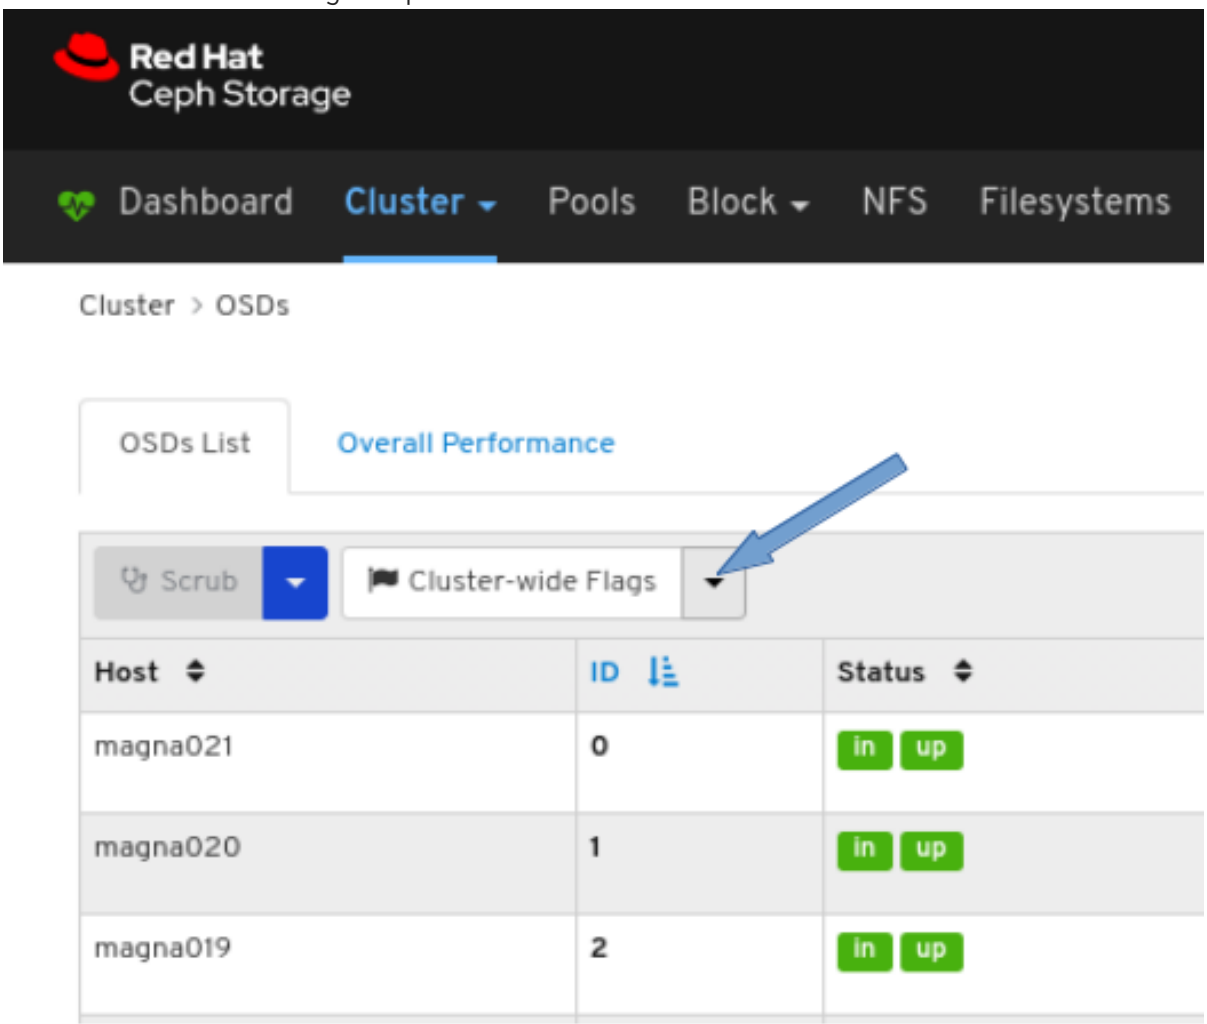

5. Select *Cluster-wide Recovery Priority* in the drop-down.

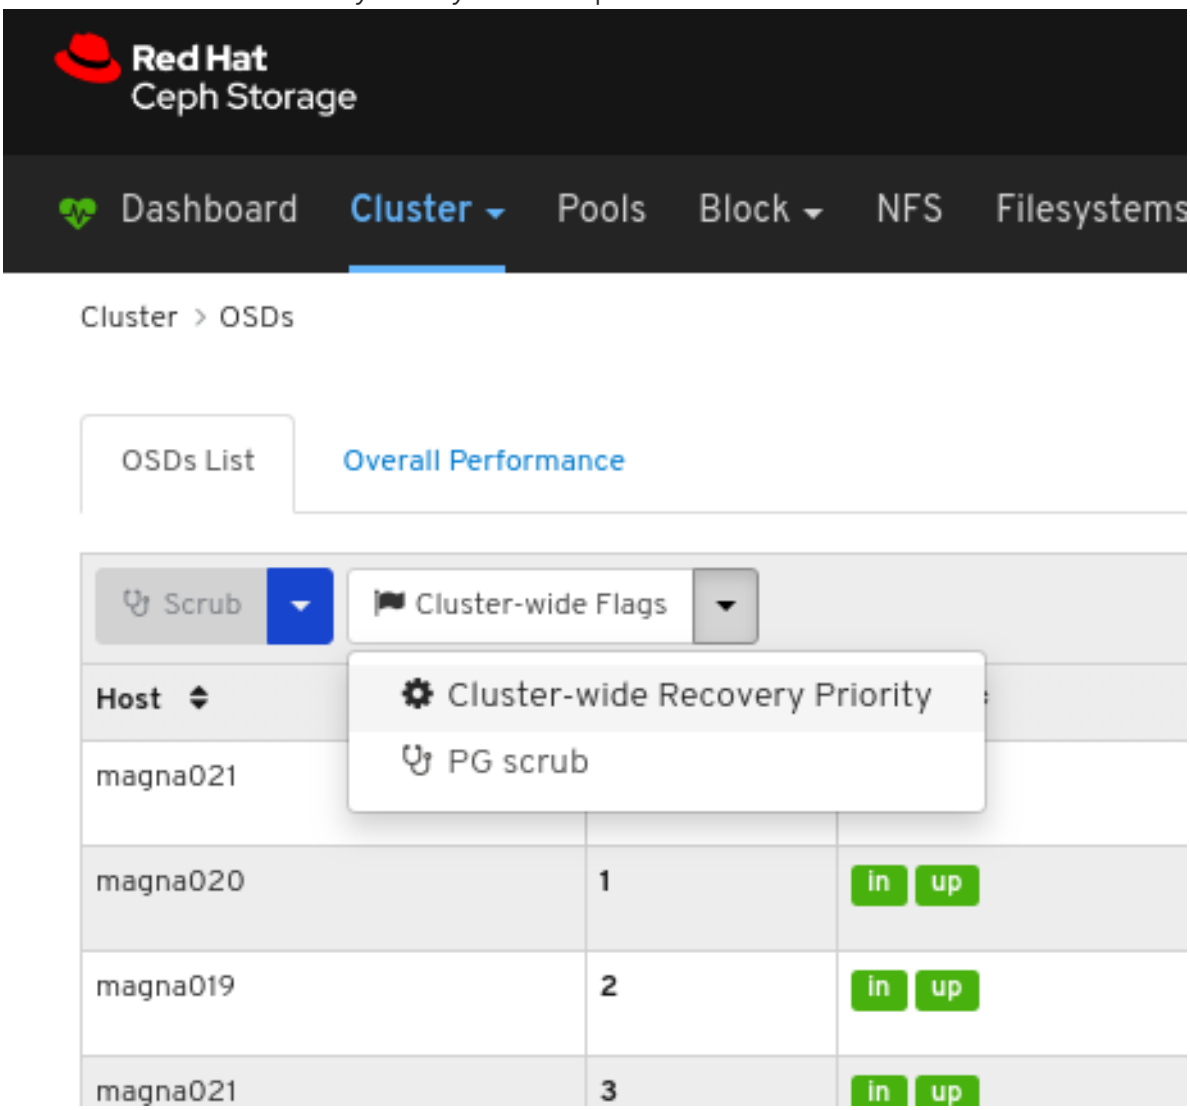

6. Optional: Select **Priority** in the drop-down menu , and then click the *Submit* button.

NOTE

There are 3 predefined options: **Low**, **Default**, **High**

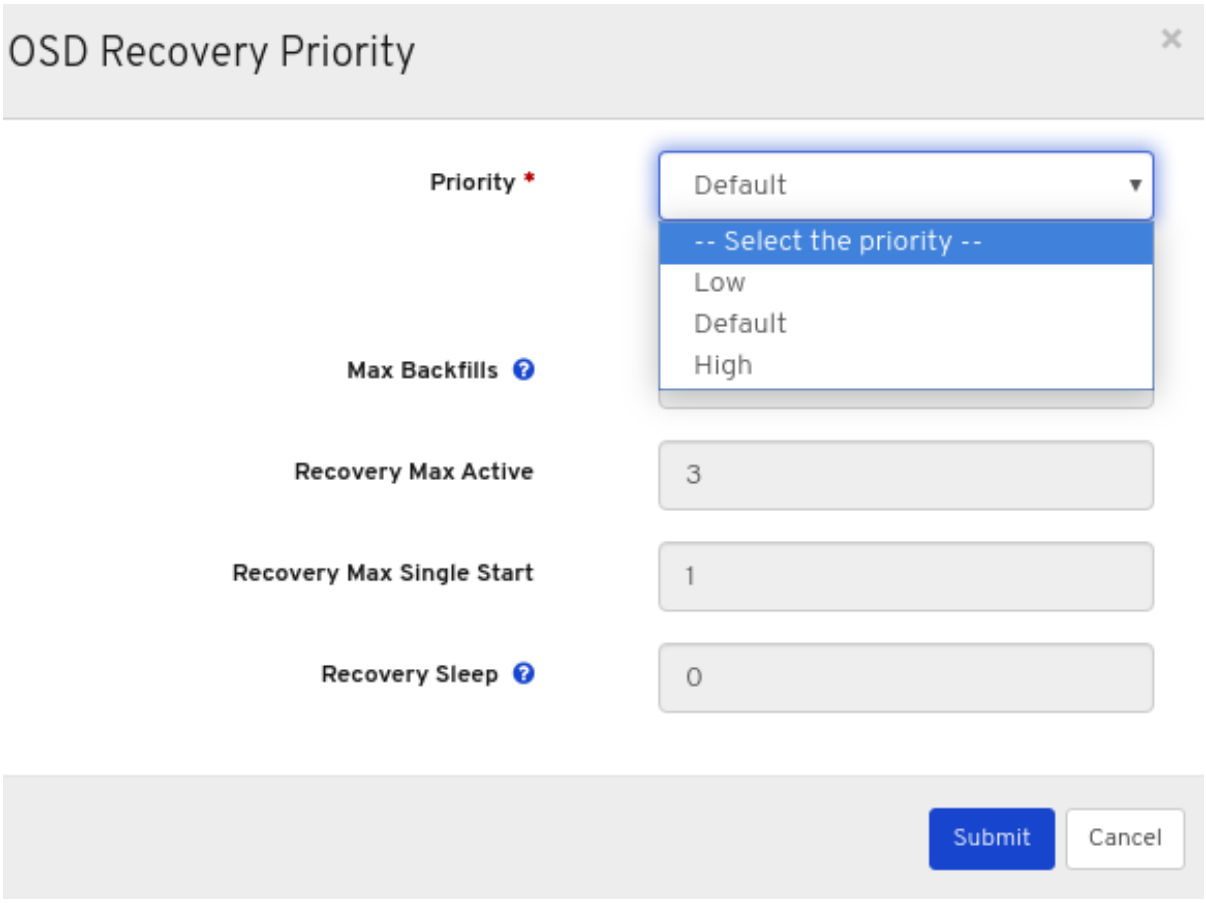

7. Optional: Click *Customize priority values*, make the required changes, and then click the *Submit* button.

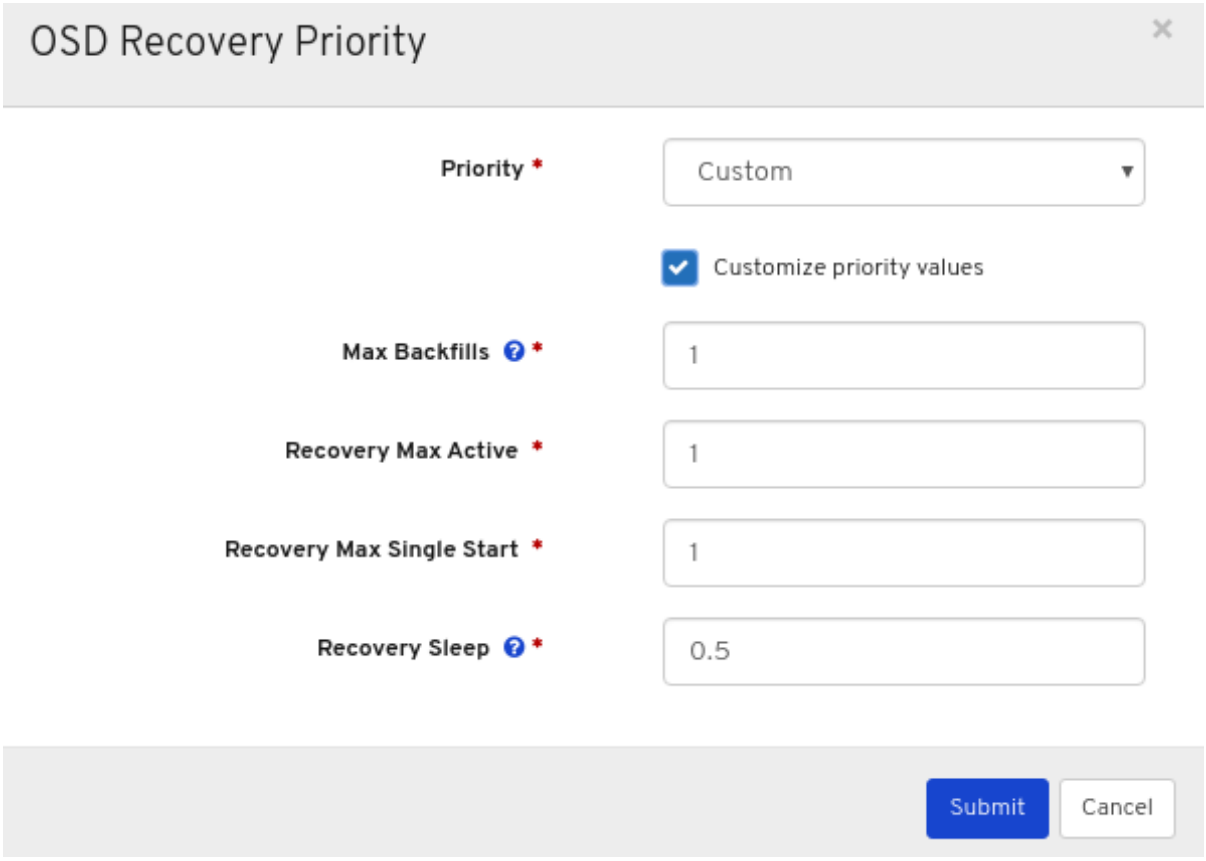

8. A notification towards the top right corner of the page pops up indicating the flags were updated successfully.

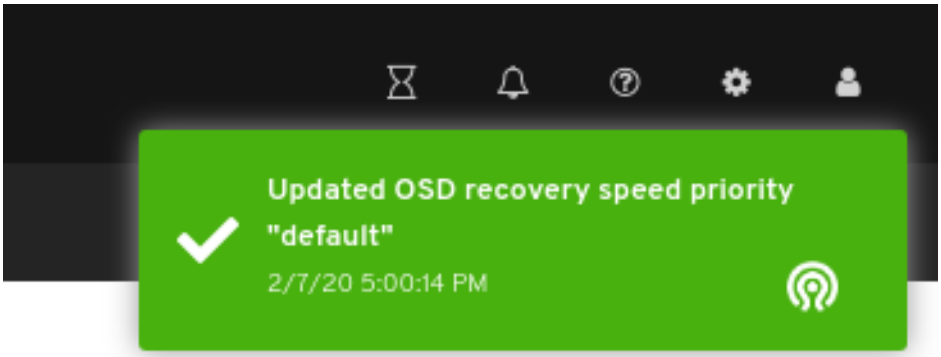

Additional Resources

- For more information on OSD recovery, see OSD Recovery in the Configuration Guide.

7.6. VIEWING AND MONITORING CONFIGURATION

The Red Hat Ceph Storage Dashboard allows you to view the list of all configuration options for the Ceph cluster. You can also edit the configuration on the Dashboard.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster* and then click *Configuration*.
3. To view the details of the configuration, click its row:

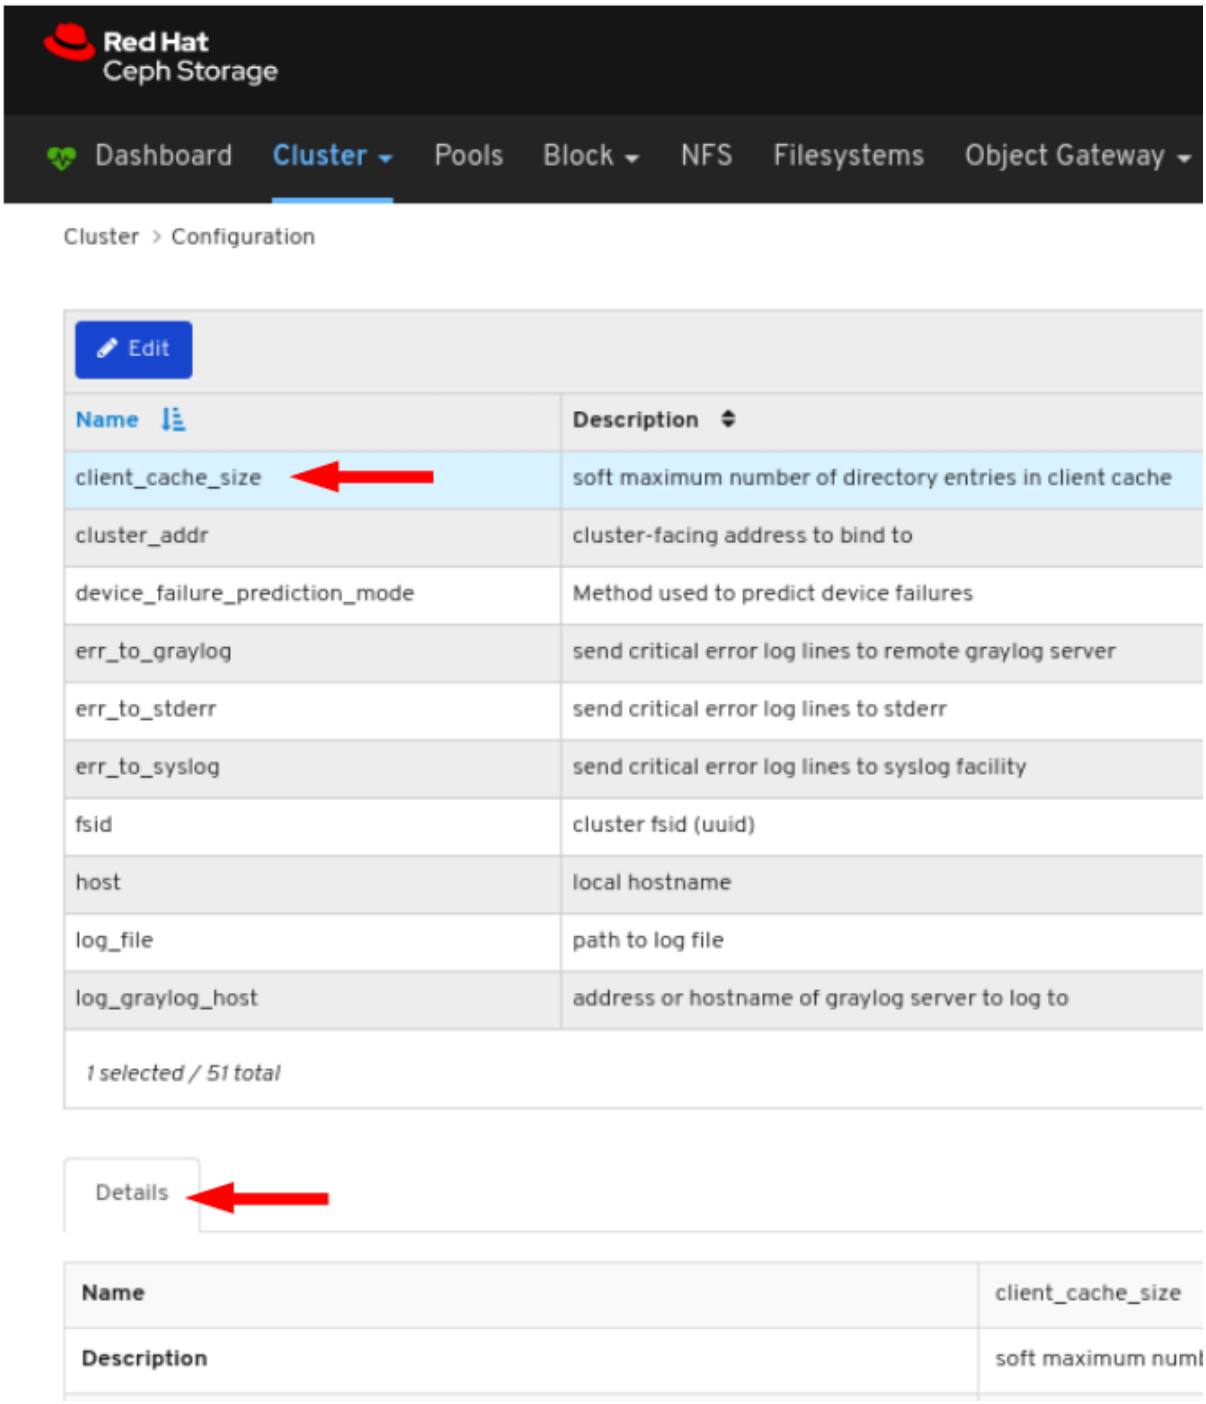

- You can search for the configuration using the *Search* box:

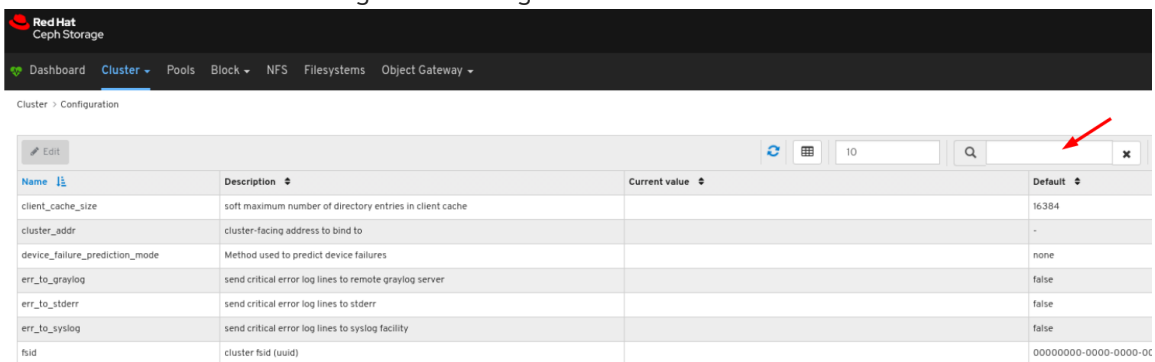

- You can filter for the configuration using *Level*, *Service* or *Source* drop-down:

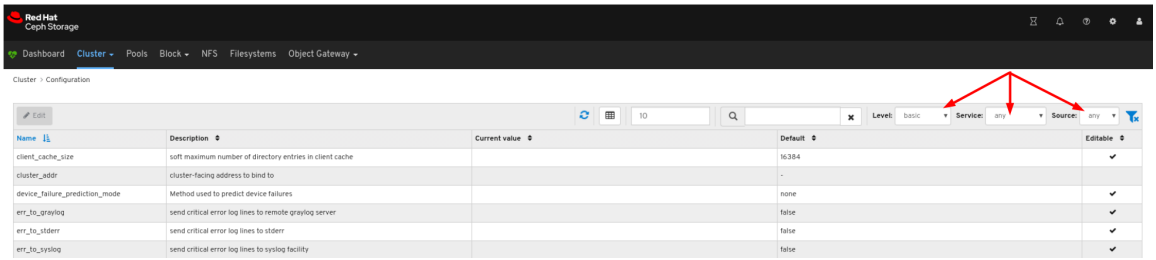

4. To edit a configuration, click its row and click the *Edit* button:

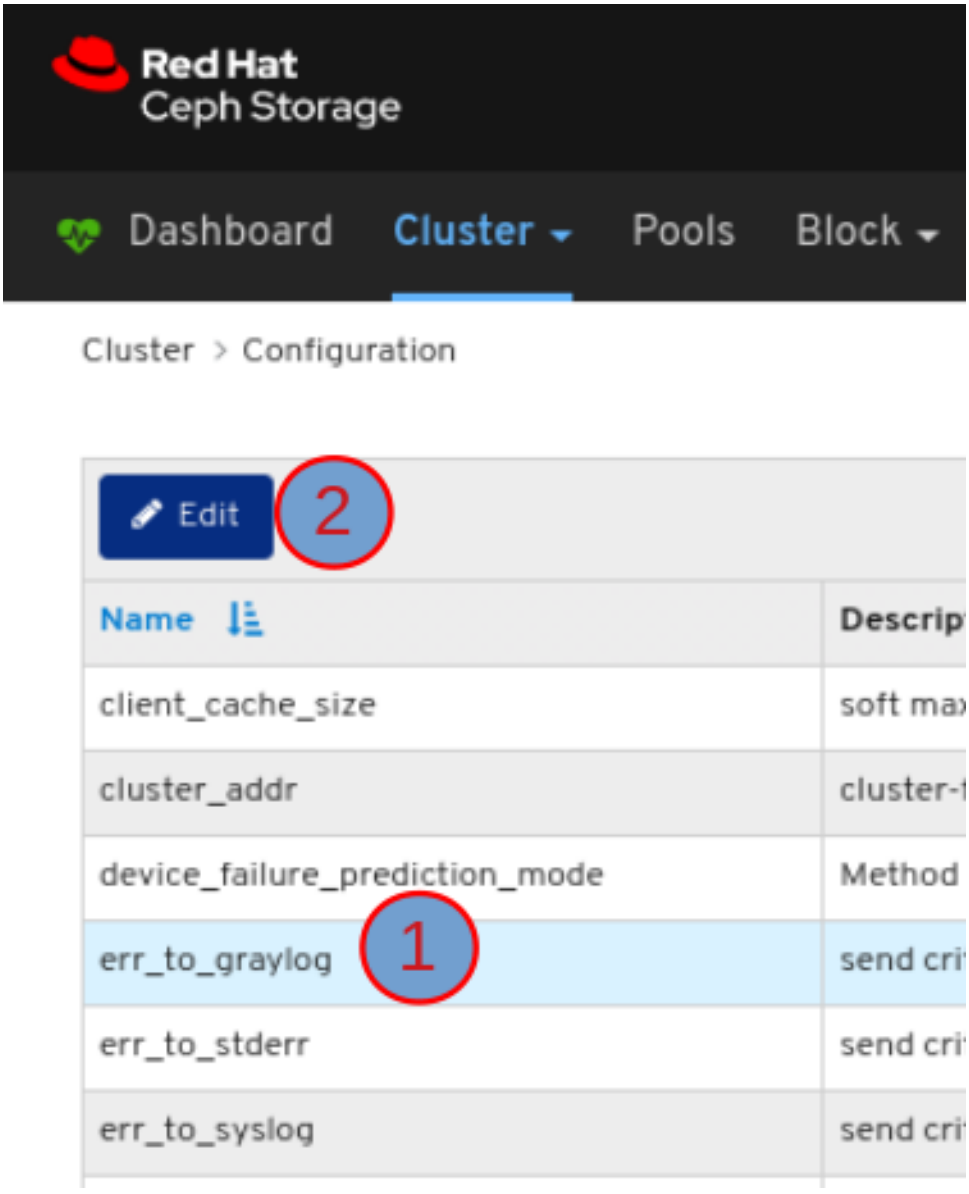

   a. In the dialog window, edit the required parameters and Click the *Save* button:

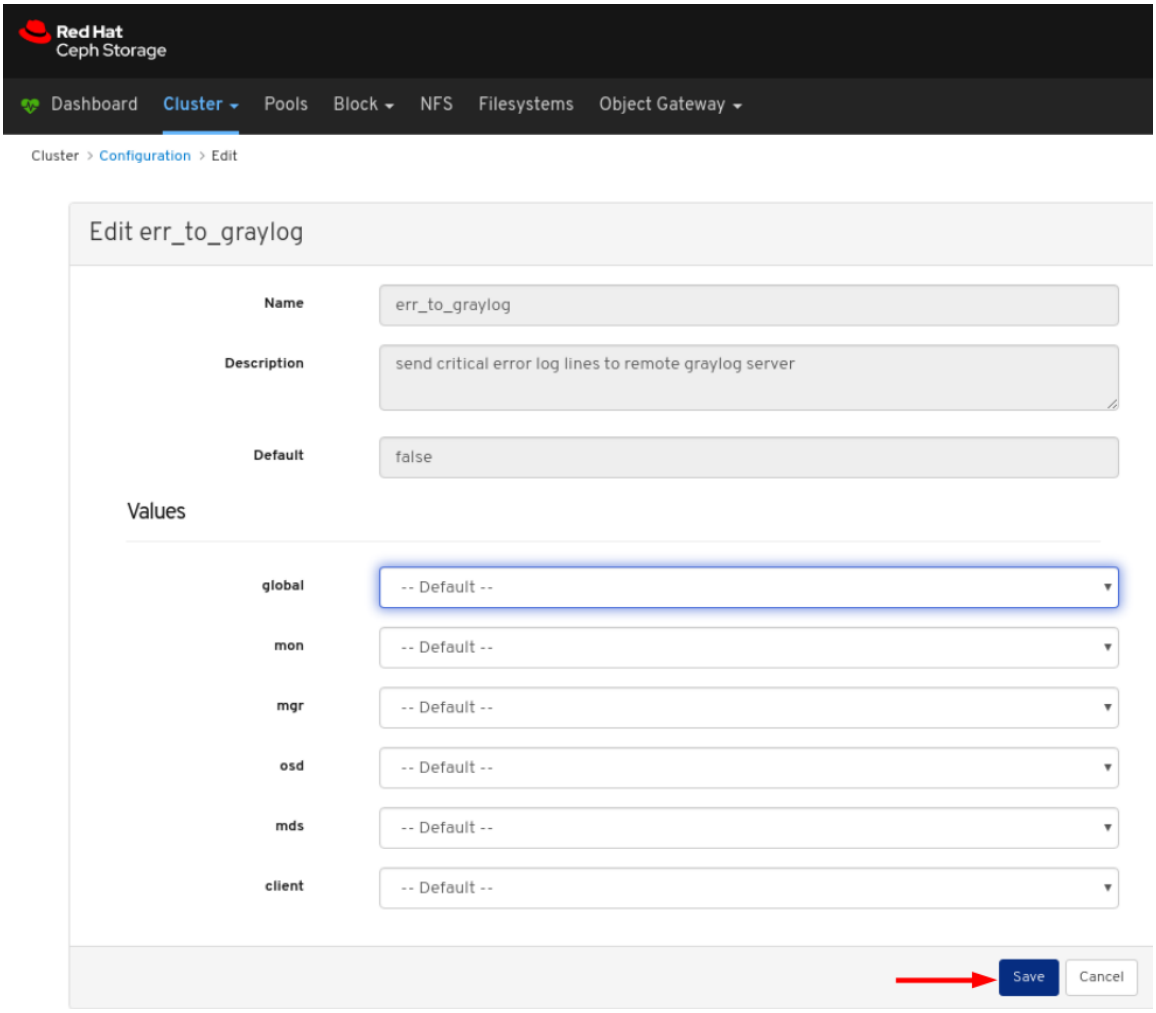

A notification confirming the change appears in the upper-right corner of the page.

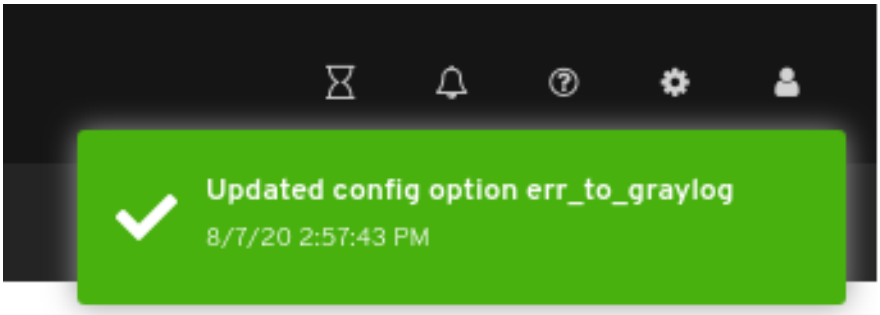

Additional Resources

- See the *Ceph Network Configuration* chapter in the *Red Hat Ceph Storage Configuration Guide* for more details.

7.7. MANAGING THE PROMETHEUS ENVIRONMENT

To monitor a Ceph storage cluster with Prometheus you can configure and enable the Prometheus exporter so the metadata information about the Ceph storage cluster can be collected.

Prerequisites

- A running Red Hat Ceph Storage 3.1 or higher cluster.

- Installation of the Red Hat Ceph Storage Dashboard.
- Root-level access to the Red Hat Ceph Storage Dashboard node.

Procedure

1. Open and edit the **/etc/prometheus/prometheus.yml** file.
   a. Under the **global** section, set the **scrape_interval** and **evaluation_interval** options to 15 seconds.

      Example

      ```
      global:
        scrape_interval:     15s
        evaluation_interval: 15s
      ```

   b. Under the **scrape_configs** section, add the **honor_labels: true** option, and edit the **targets**, and **instance** options for each of the **ceph-mgr** nodes.

      Example

      ```
      scrape_configs:
        - job_name: 'node'
          honor_labels: true
          static_configs:
          - targets: [ 'node1.example.com:9100' ]
            labels:
              instance: "node1.example.com"
          - targets: ['node2.example.com:9100']
            labels:
              instance: "node2.example.com"
      ```

      

      NOTE

      Using the **honor_labels** option enables Ceph to output properly-labelled data relating to any node in the Ceph storage cluster. This allows Ceph to export the proper **instance** label without Prometheus overwriting it.

   c. To add a new node, simply add the **targets**, and **instance** options in the following format:

      Example

      ```
      - targets: [ 'new-node.example.com:9100' ]
        labels:
          instance: "new-node"
      ```

      

      NOTE

      The **instance** label has to match what appears in Ceph's OSD metadata **instance** field, which is the short host name of the node. This helps to correlate Ceph stats with the node's stats.

2. Add Ceph targets to the **/etc/prometheus/ceph_targets.yml** file in the following format.

   Example

```
[
    {
        "targets": [ "cephnode1.example.com:9283" ],
        "labels": {}
    }
]
```

3. Enable the Prometheus module:

```
[root@mon ~]# ceph mgr module enable prometheus
```

7.8. RESTORING GRAFANA-SERVER AND PROMETHEUS

The grafana-server includes the Grafana UI, Prometheus, the containers, and the Red Hat Ceph Storage configuration. When the grafana-server crashes or is faulty, you can restore it by taking a back-up of the files and restoring it using the backed-up files. For Prometheus, you can take an external back-up and then restore the data.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Root-level access to the Grafana nodes.

Procedure

1. Take the back-up of the Grafana database:

   a. On the grafana-server node, stop the Grafana service:

      Example

```
[root@node04 ~]# systemctl stop grafana-server.service
[root@node04 ~]# systemctl status grafana-server.service
```

   b. Take the back-up of the Grafana database:

      Example

```
[root@node04 ~]# cp /var/lib/grafana/grafana.db /var/lib/grafana/grafana_backup.db
```

   c. On the grafana-server node, restart the Grafana service:

      Example

```
[root@node04 ~]# systemctl restart grafana-server.service
```

2. Restore the grafana-server:

a. On the grafana-server node, if the Grafana service is running, stop the service:

Example

```
[root@node04 ~]# systemctl stop grafana-server.service
[root@node04 ~]# systemctl status grafana-server.service
```

b. Move the backed-up **grafana.db** file to **/var/lib/grafana/** directory:

Example

```
[root@node04 ~]# mv /var/lib/grafana/grafana_backup.db /var/lib/grafana/
```

c. On the grafana-server node, restart the Grafana service:

Example

```
[root@node04 ~]# systemctl restart grafana-server.service
```

3. For the Prometheus alerts, you have to take external back-up of **prometheus_data_dir** directory, a Ceph-Ansible setting which by default is **var/lib/prometheus** directory and restore the service using the backed-up directory.

a. On the grafana-server node, stop the Prometheus service:

Example

```
[root@node04 ~]# systemctl stop prometheus.service
[root@node04 ~]# systemctl status prometheus.service
```

b. Take the back-up of the default Prometheus directory:

Example

```
[root@node04 ~]# cp /var/lib/prometheus/ /var/lib/prometheus_backup/
```

c. Replace the **prometheus_data_dir** directory with the backed-up directory:

Example

```
[root@node04 ~]# mv /var/lib/prometheus_backup/ /var/lib/prometheus_data_dir
```

d. On the grafana-server node, restart the prometheus service:

Example

```
[root@node04 ~]# systemctl restart prometheus.service
[root@node04 ~]# systemctl status prometheus.service
```

NOTE

If you have made changes to the Prometheus parameters in **group_vars/all.yml** file, then you have to rerun the playbook.

4. Optional: If the changes do not reflect on the Red Hat Ceph Storage Dashboard, then you have to disable and then enable the dashboard:

   Example

```
[root@node04 ~]# ceph mgr module disable dashboard
[root@node04 ~]# ceph mgr module enable dashboard
```

7.9. VIEWING AND 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. This includes the following pre-defined alerts:

- OSD(s) Down
- Ceph Health Error
- Ceph Health Warning
- Cluster Capacity Low
- Disk(s) Near Full
- MON(s) Down
- Network Errors
- OSD Host Loss Check
- OSD Host(s) Down
- OSD(s) with High PG Count
- PG(s) Stuck
- Pool Capacity Low
- Slow OSD Responses

7.9.1. Viewing alerts

After an alert has fired, you can view it on the Red Hat Ceph Storage Dashboard. You can also enable the dashboard to send an email about the alert.

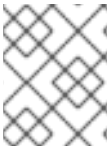

NOTE

Simple mail transfer protocol (SMTP) and SSL is not supported in Red Hat Ceph Storage 4 cluster.

Prerequisite

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

Procedure

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

Figure 7.1. Edit Manager module for alerts

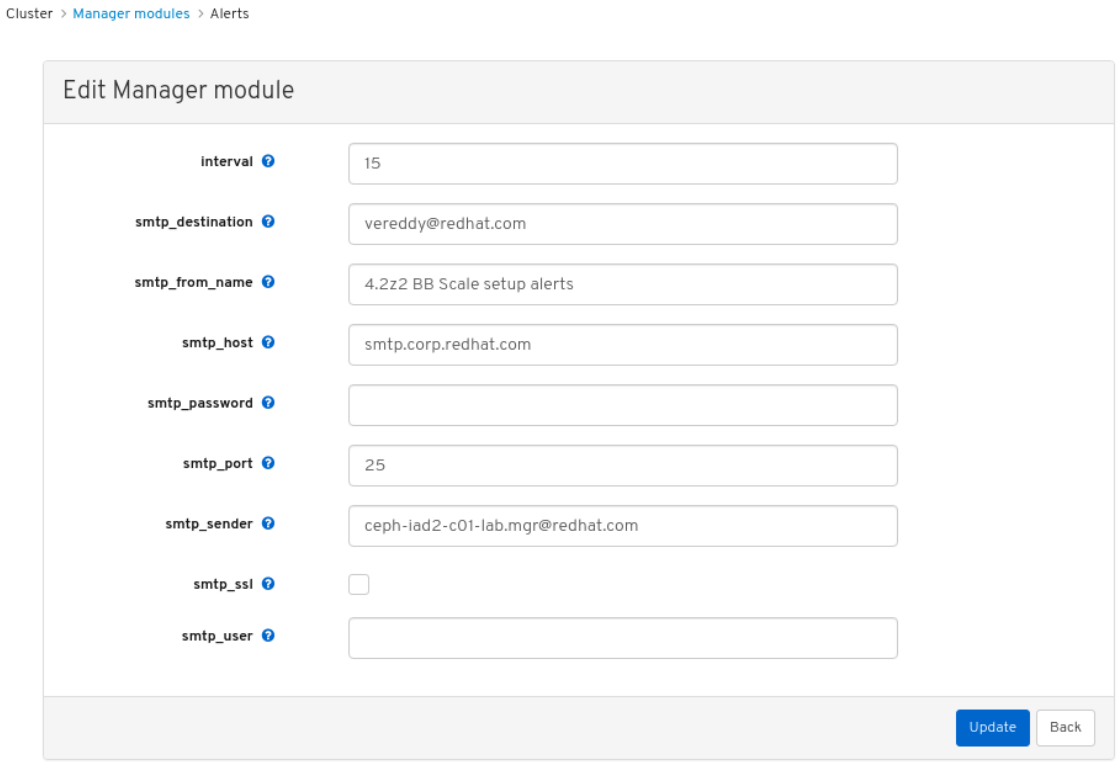

3. On the navigation bar, click *Cluster*.
4. Select *Monitoring* from the drop-down menu.
5. To view details about the alert, click on its row:

Figure 7.2. Alert Details

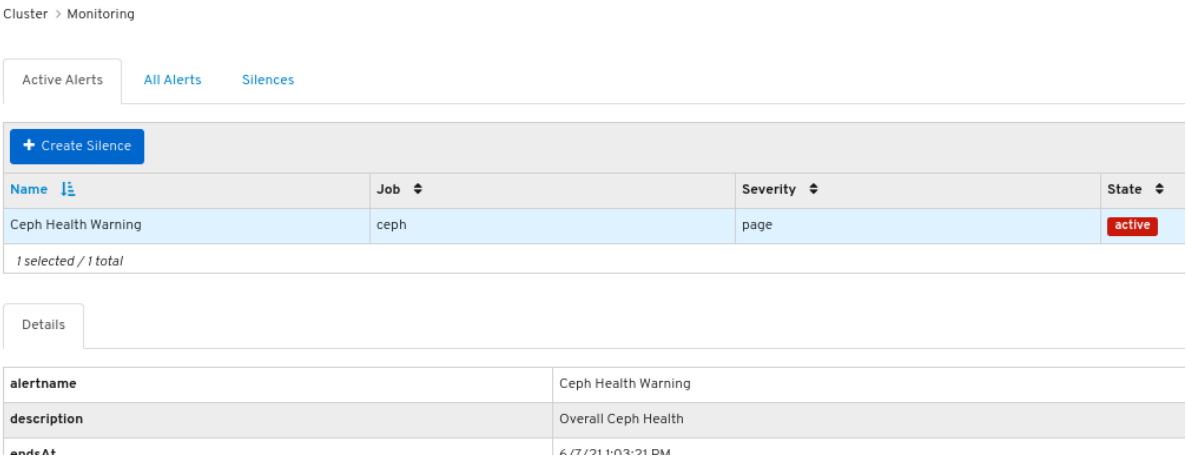

6. To view the source of an alert, click on its row, and then click **Source**.

7.9.2. 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

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster*.
3. Select *Monitoring* from the drop-down menu.
4. Click on the row for the alert and then click *+Create Silence*.
5. In the *CreateSilence* window, Add the details for the *Duration* and click *Create Silence*.

Figure 7.3. Create Silence

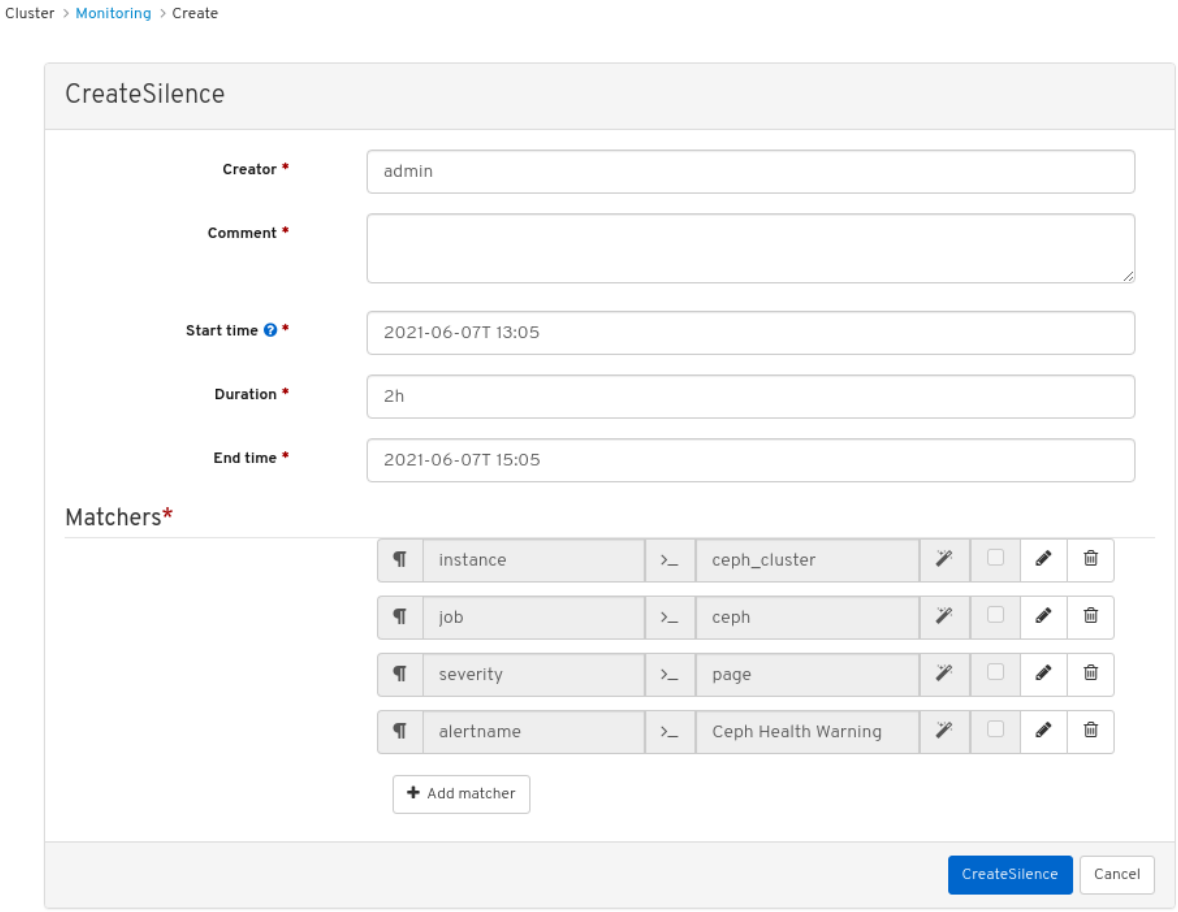

6. You get a notification that the silence was created successfully.

7.9.3. Re-creating a silence

You can re-create a silence from an expired silence on the Red Hat Ceph Storage Dashboard.

Prerequisite

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

Procedure

1. Log in to the Dashboard.
2. Select *Monitoring* from the drop-down menu.
3. Click the *Silences* tab.
4. Click on the row for the expired silence.
5. Click the *Recreate* button.

6. In the *RecreateSilence* window, add the details and click *RecreateSilence*.

Figure 7.4. Recreate silence

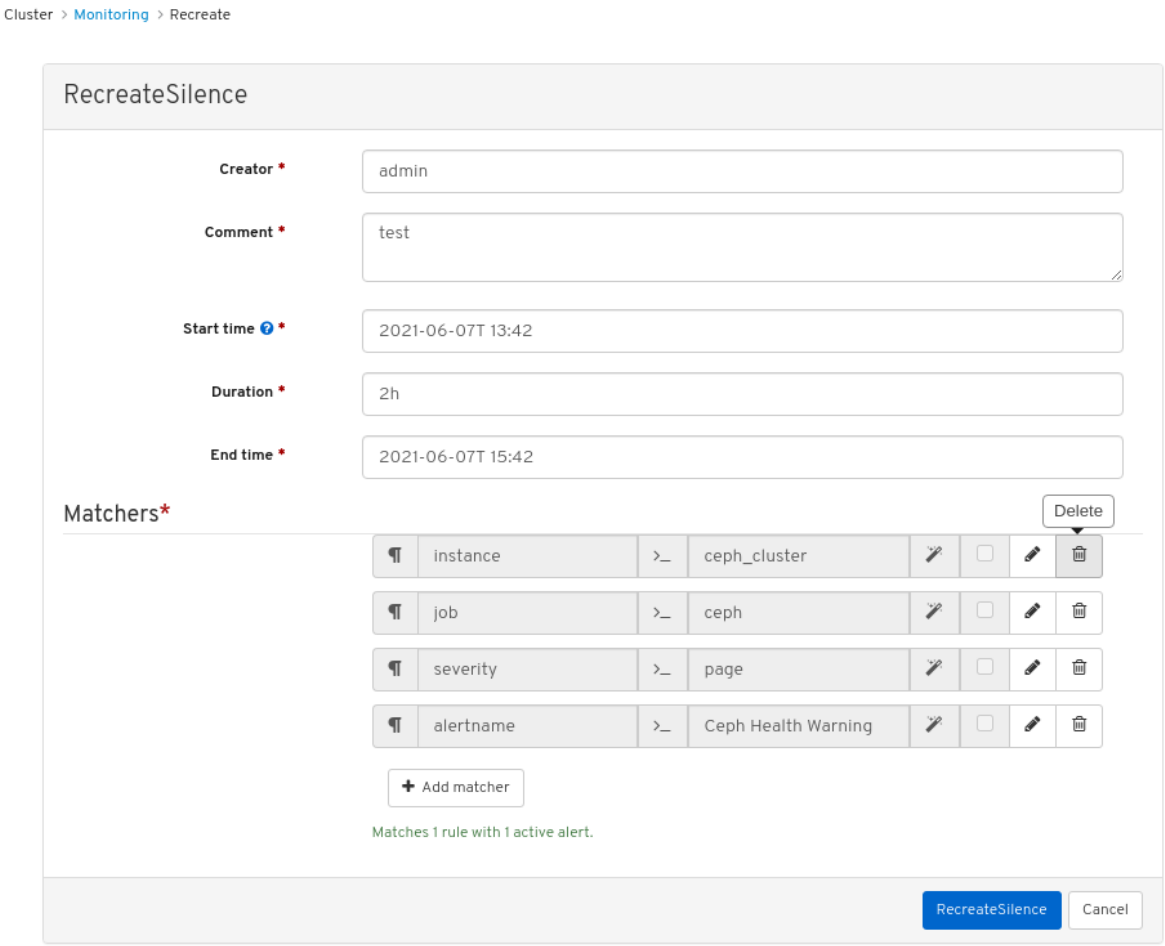

7. You get a notification that the silence was recreated successfully.

7.9.4. 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

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster*.
3. Select *Monitoring* from the drop-down menu.

4. Click the *Silences* tab.
5. Click on the row for the silence.
6. In the *Edit* drop-down menu, select *Edit*.
7. In the *EditSilence* window, update the details and click *Edit Silence*.

Figure 7.5. Edit silence

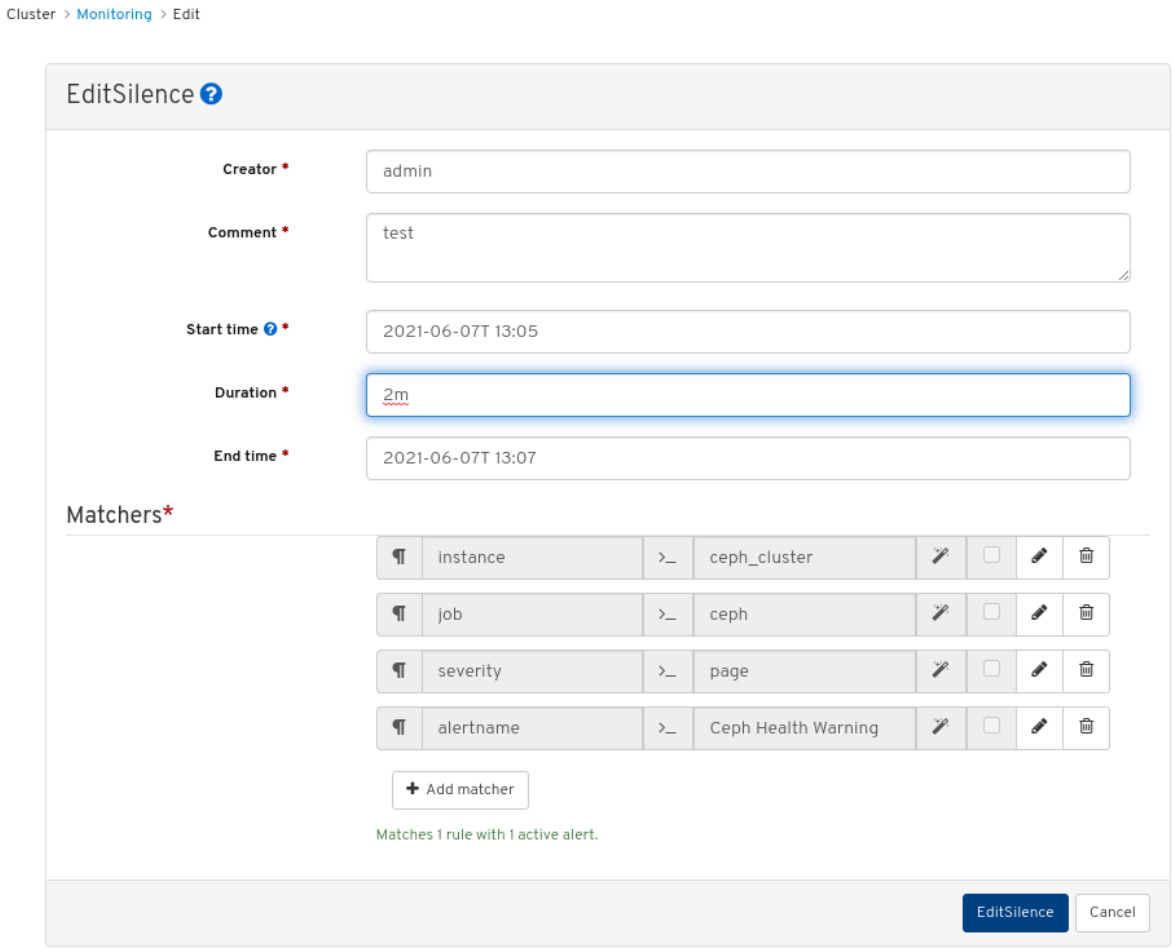

8. You get a notification that the silence was updated successfully.

7.9.5. Expiring a silence

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

Prerequisite

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Cluster*.
3. Select *Monitoring* from the drop-down menu.
4. Click the *Silences* tab.
5. Click on the row for the silence.
6. In the *Edit* drop-down menu, select *Expire*.
7. In the *Expire Silence* dialog box, select *Yes, I am sure* , and then click *Expire Silence*.

Figure 7.6. Expire Silence

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

7.9.6. Additional Resources

- For more information, see the *Red Hat Ceph StorageTroubleshooting Guide*.

7.10. MANAGING POOLS

As a storage administrator, you can create, delete, and edit pools.

Prerequisites

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

7.10.1. Creating pools

You can create pools to logically partition your storage objects.

Prerequisites

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

Procedure

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

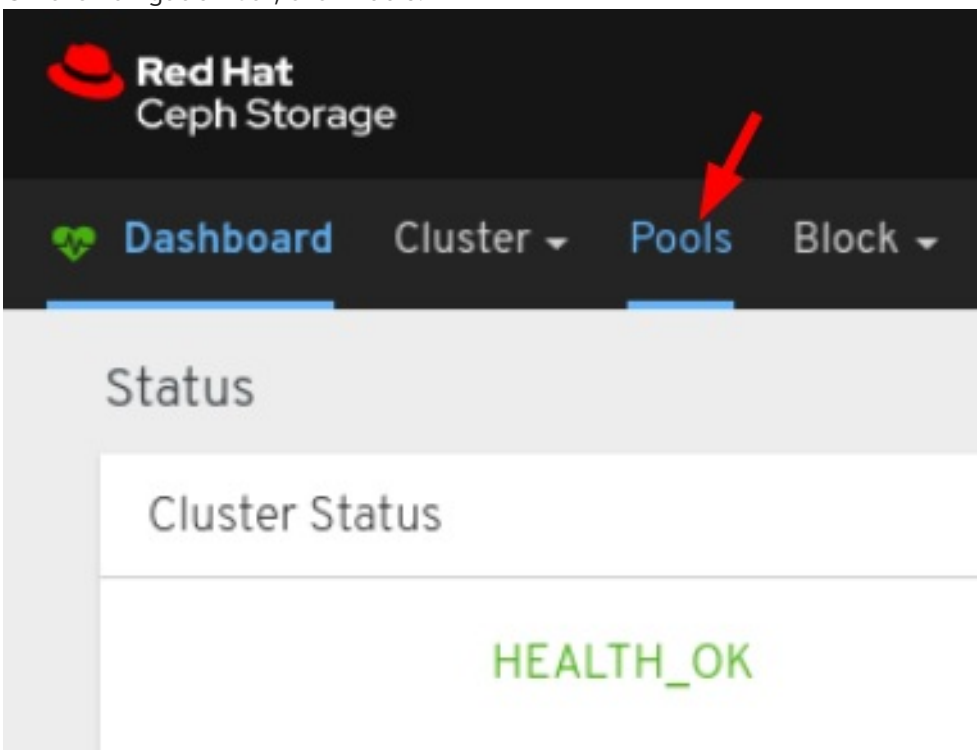

3. Click the *Create* button towards the top left corner of the page.

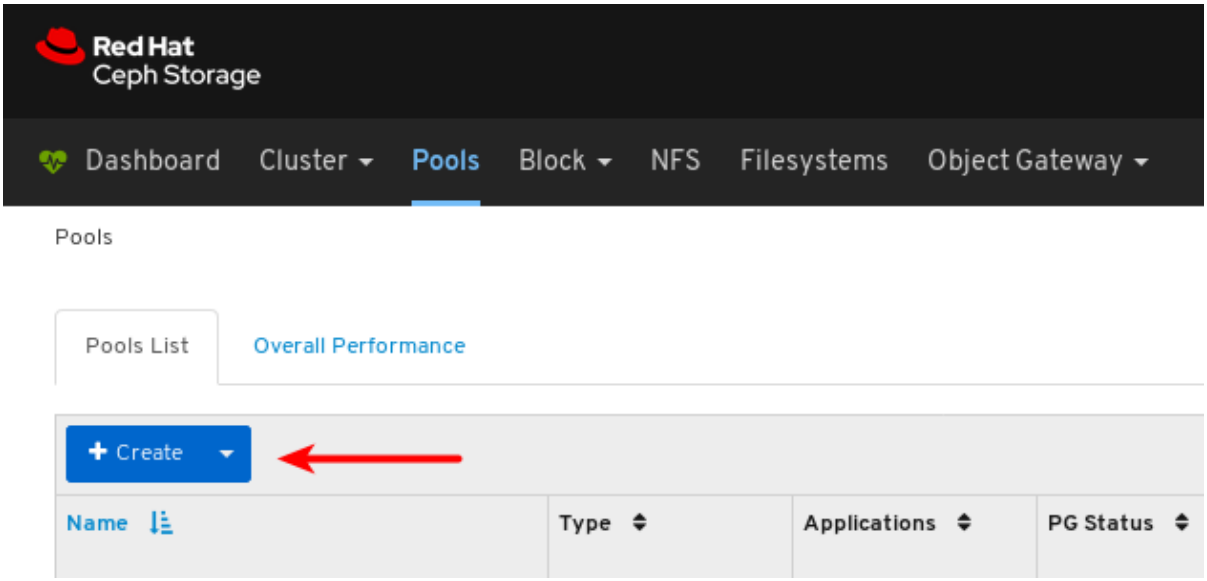

4. In the dialog window, set the name.

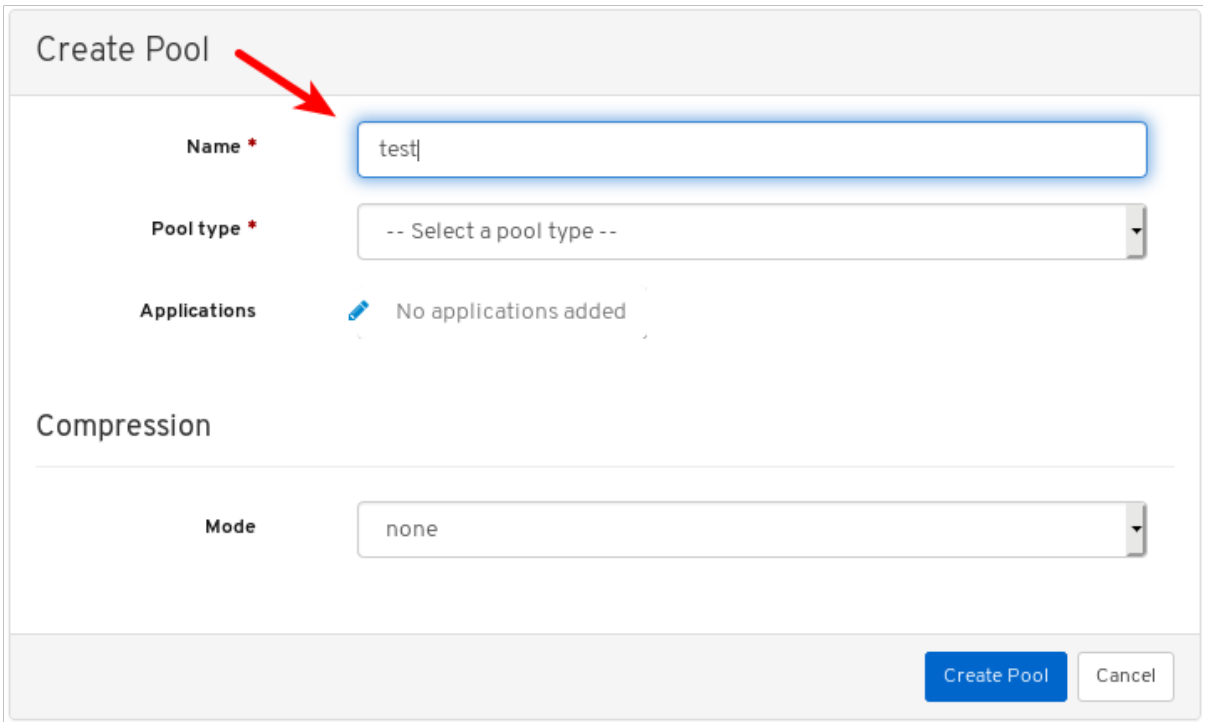

5. Select either replicated or Erasure Coded (EC) pool type.

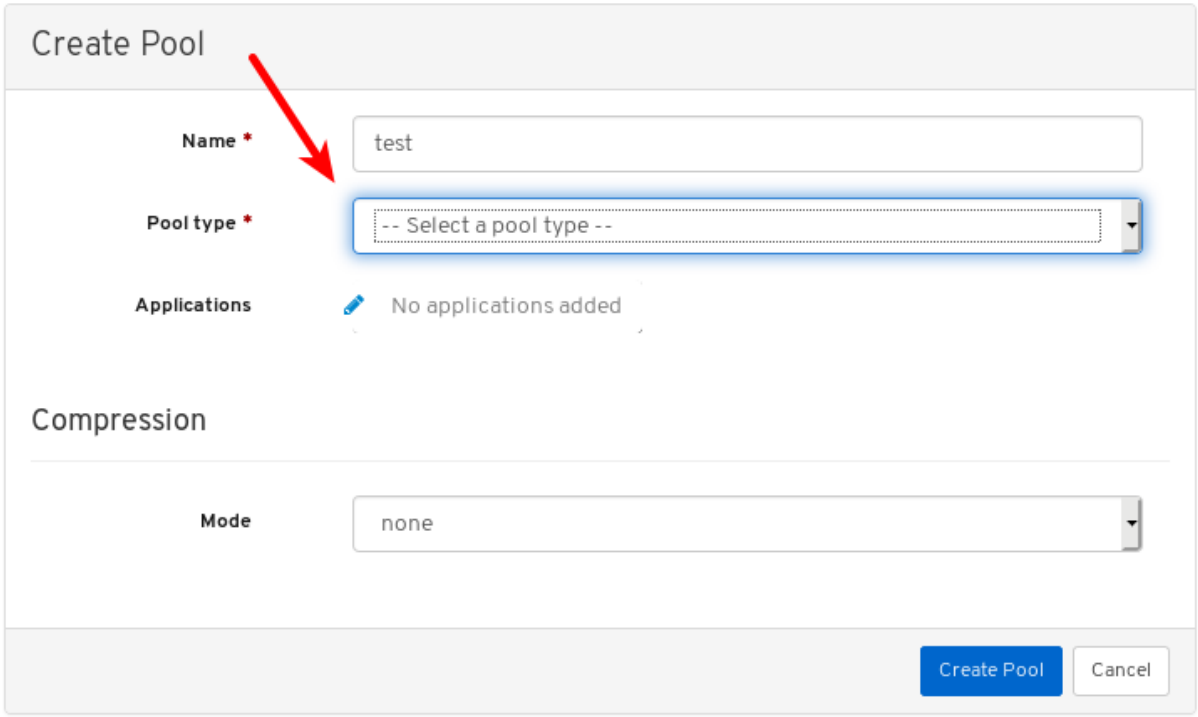

6. Set the Placement Group (PG) number.

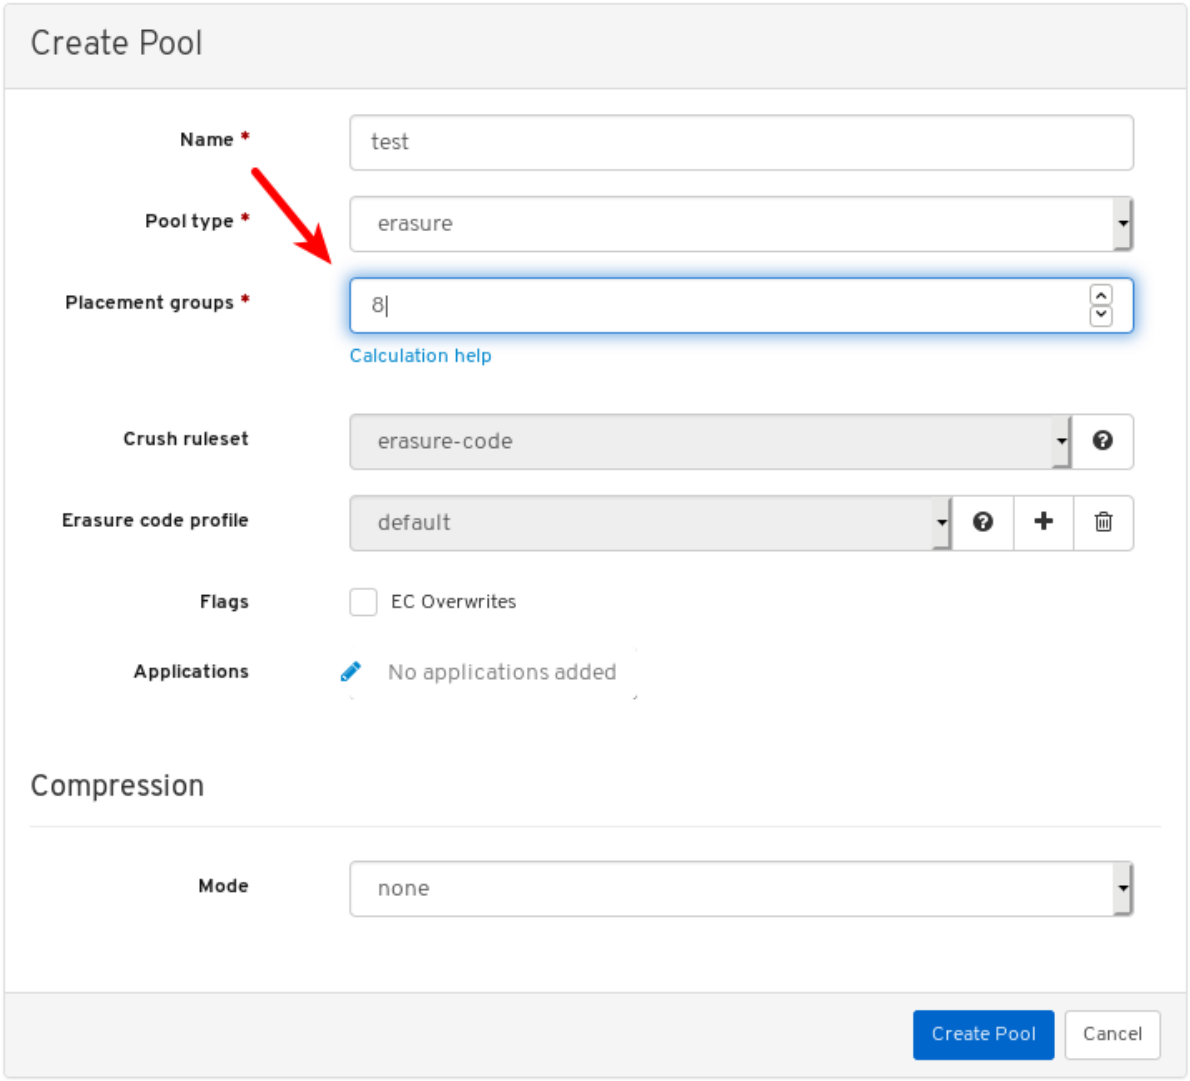

For assistance in choosing the PG number, use the PG calculator. Contact Red Hat Technical Support if unsure.

7. Optional: If using a replicated pool type, set the replicated size.

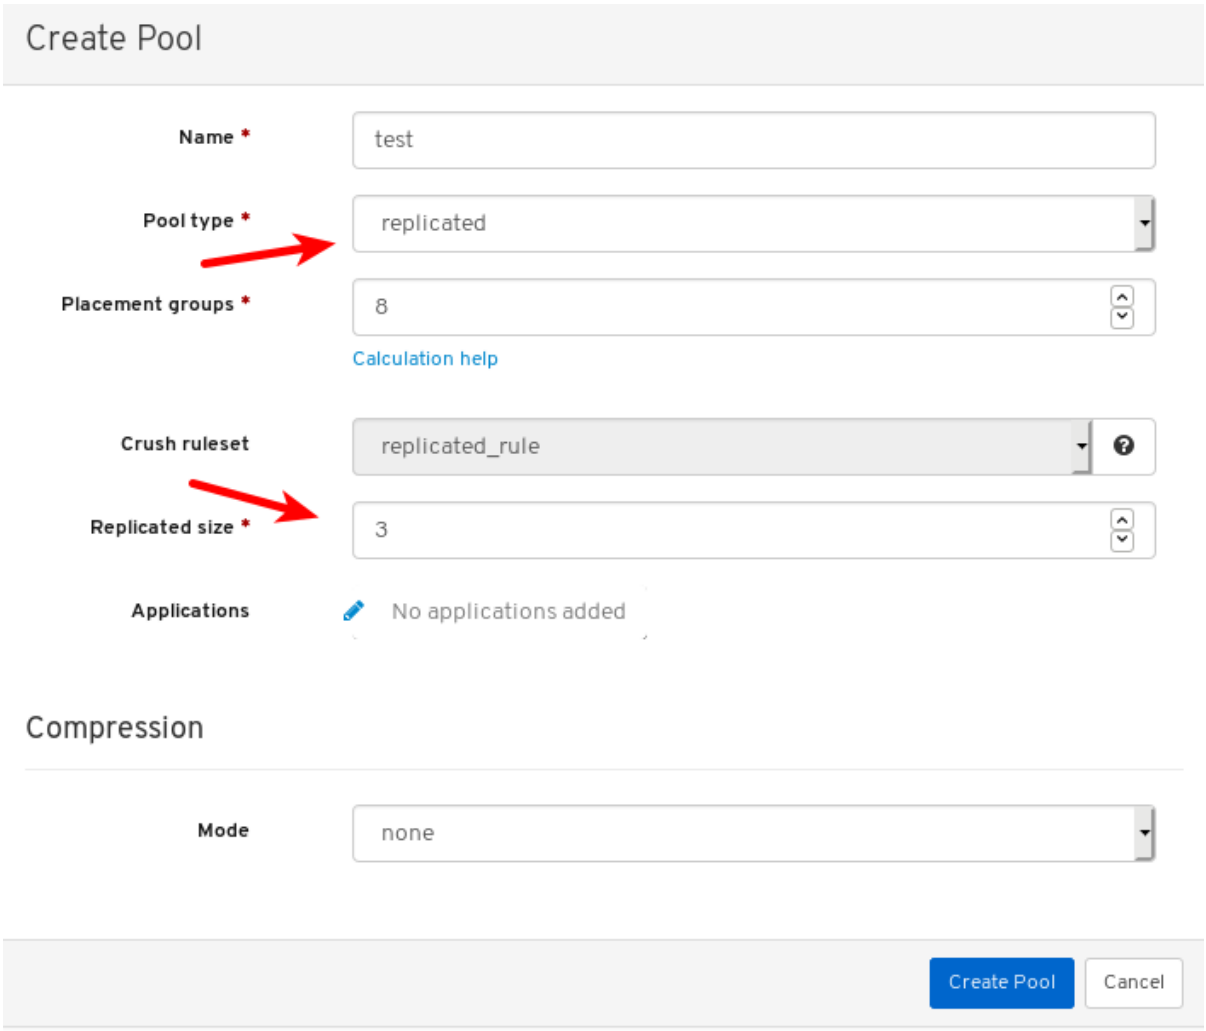

8. Optional: If using an EC pool type configure the following additional settings.

   a. Optional: To see the settings for the currently selected EC profile, click the question mark.

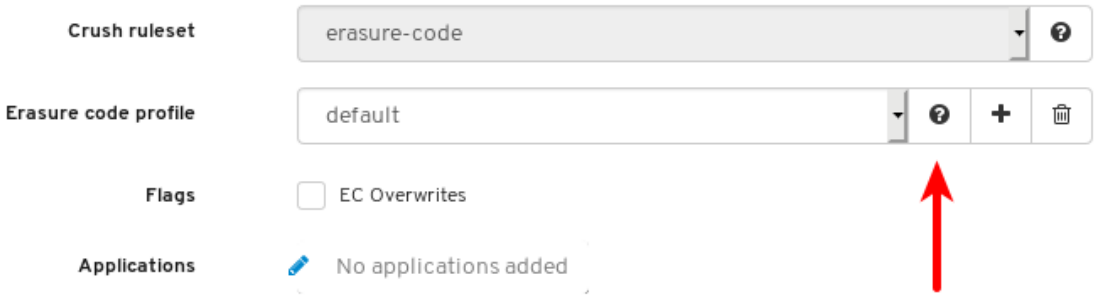

   i. A table of the settings for the selected EC profile is shown.

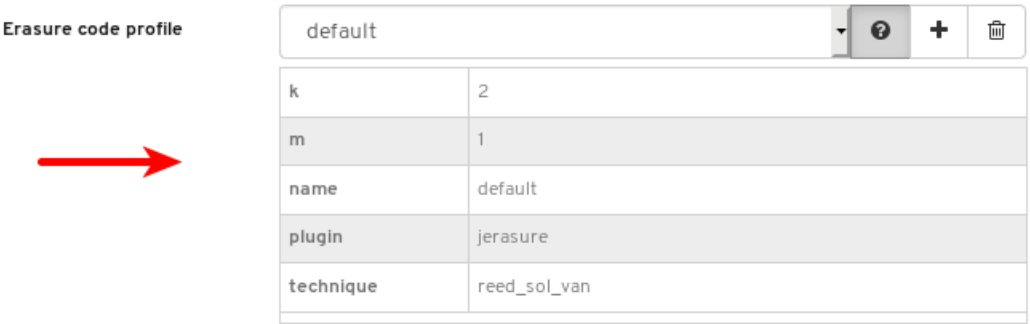

b. Optional: Add a new EC profile by clicking the plus symbol.

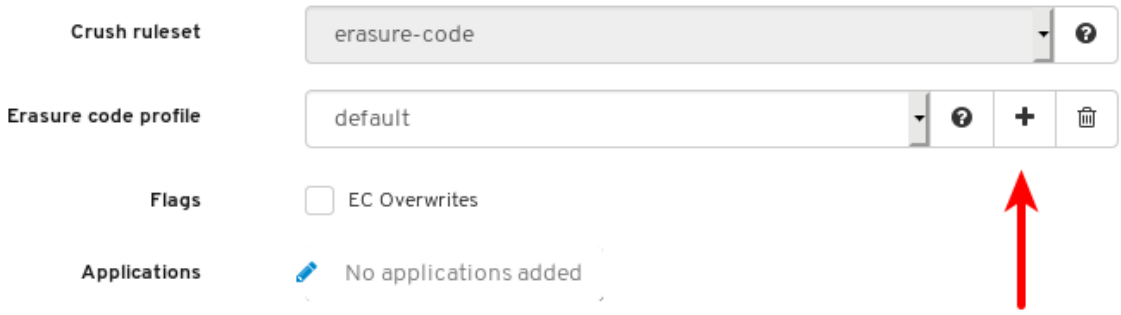

i. Set the name of the new EC profile, at **1**, click any question mark symbol for info about that setting, at **2**, and after modifying all required settings, click *Create EC Profile*, at **3**.

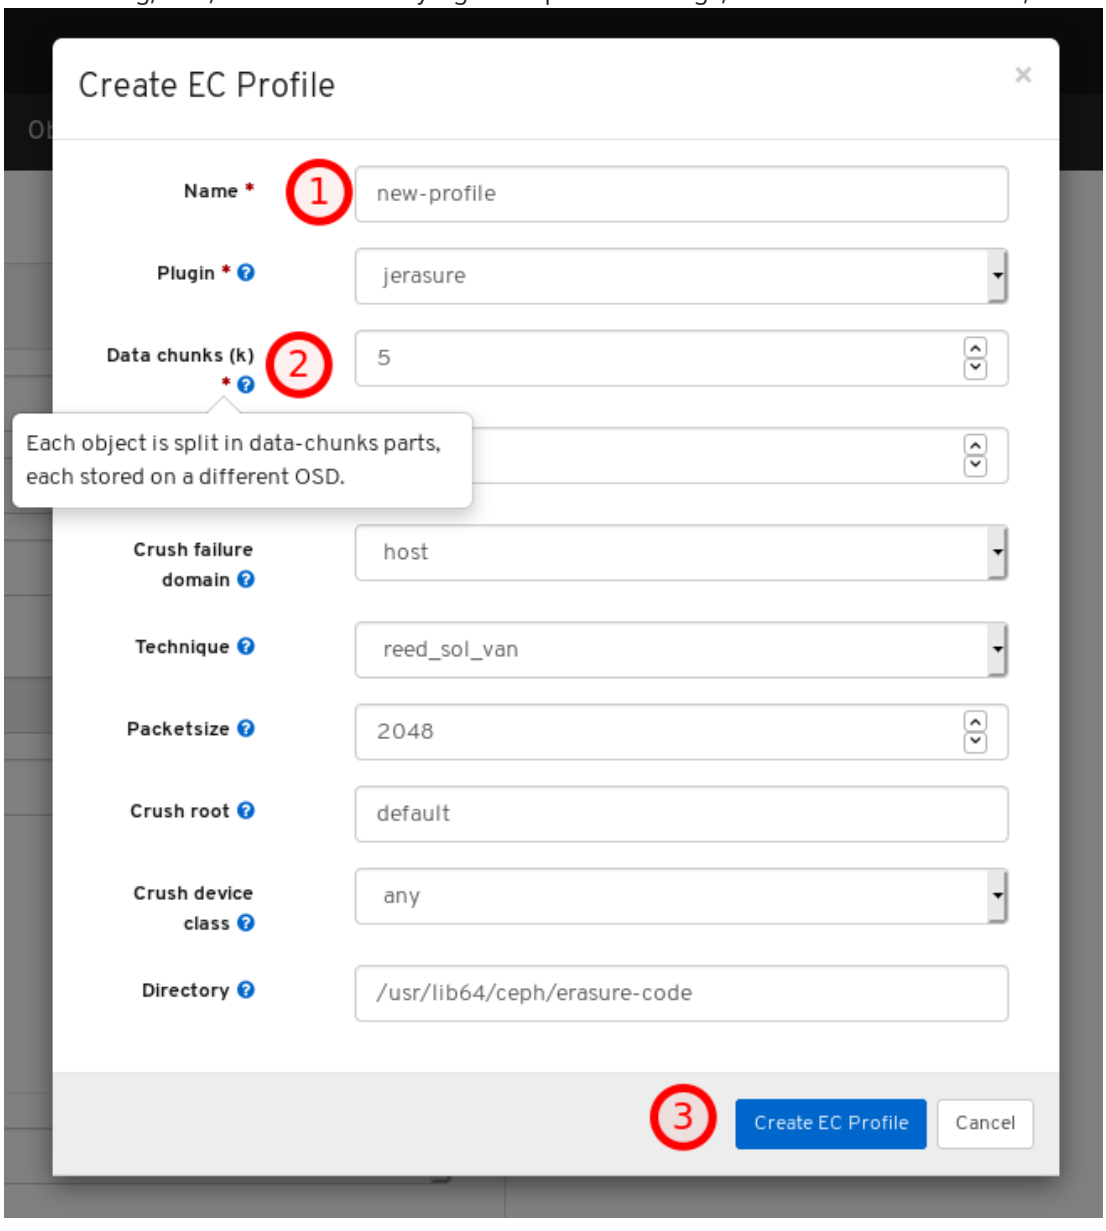

ii. Select the new EC profile.

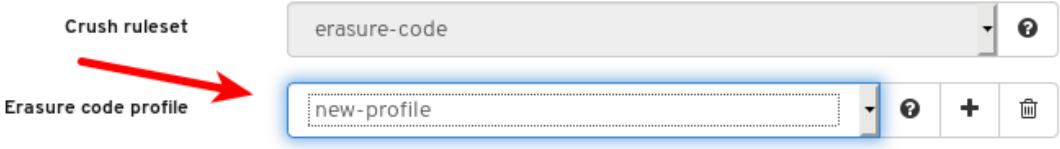

c. Optional: If EC overwrites are required, click its button.

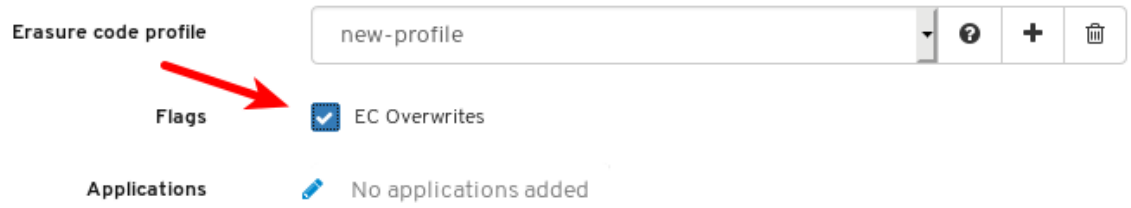

9. Optional: Click the pencil symbol to select an application for the pool.

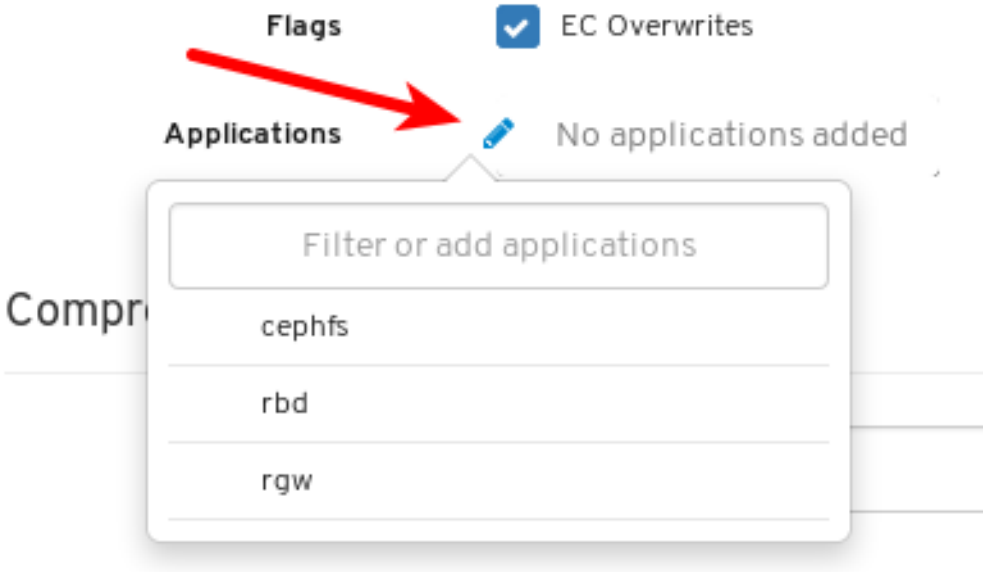

10. Optional: If compression is required, select *passive*, *aggressive*, or *force*.

11. Click the *Create Pool* button.

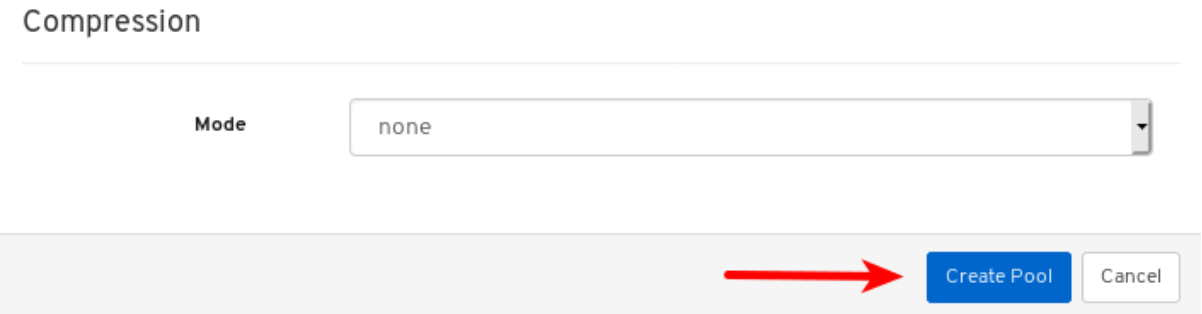

12. Notifications towards the top right corner of the page indicate the pool was created successfully.

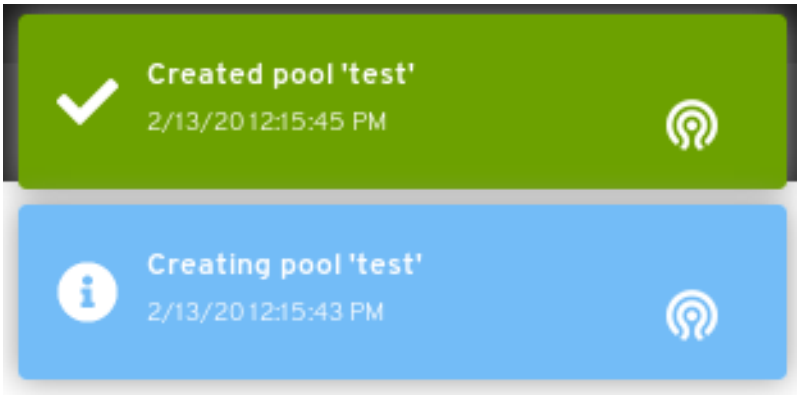

Additional Resources

- For more information, see Ceph pools in the Architecture Guide.

7.10.2. Editing pools

The Red Hat Ceph Storage Dashboard allows editing of pools.

Prerequisites

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

Procedure

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

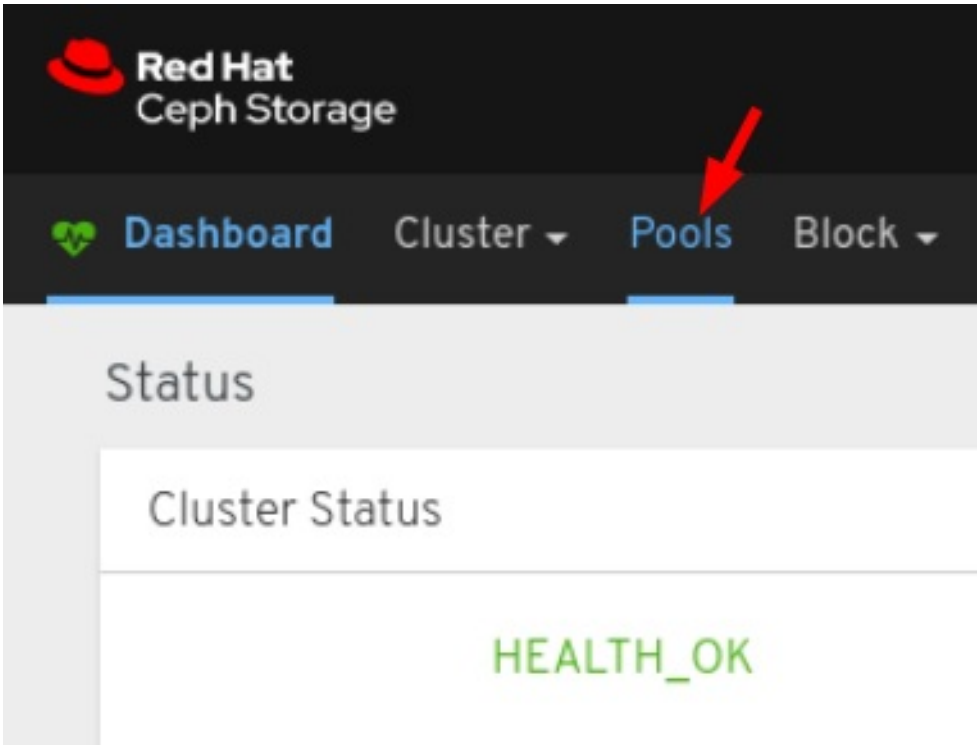

3. To edit the pool, click its row:

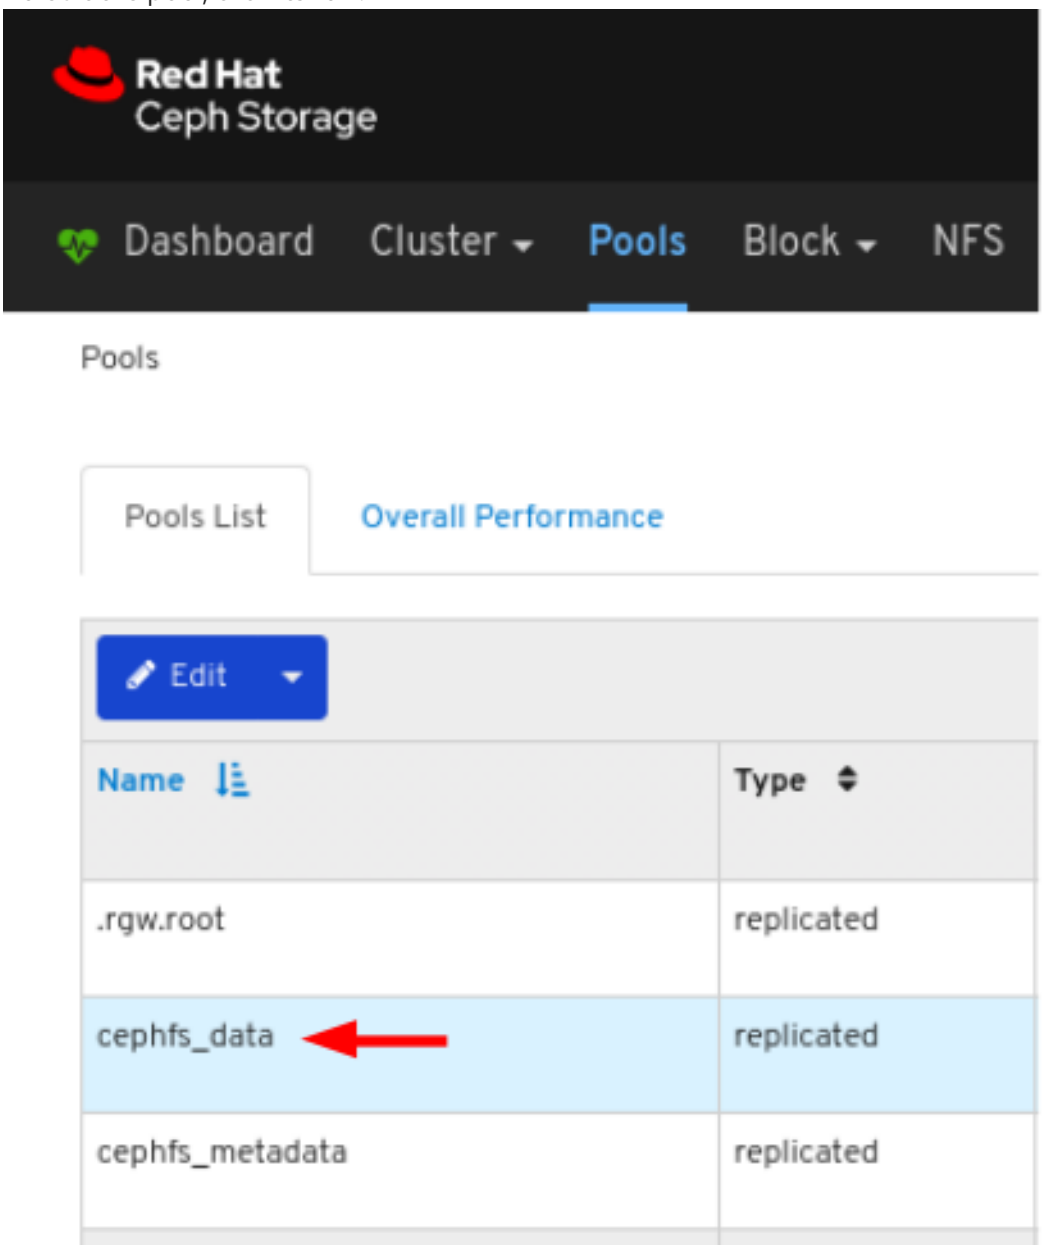

4. Select *Edit* In the *Edit* drop-down:

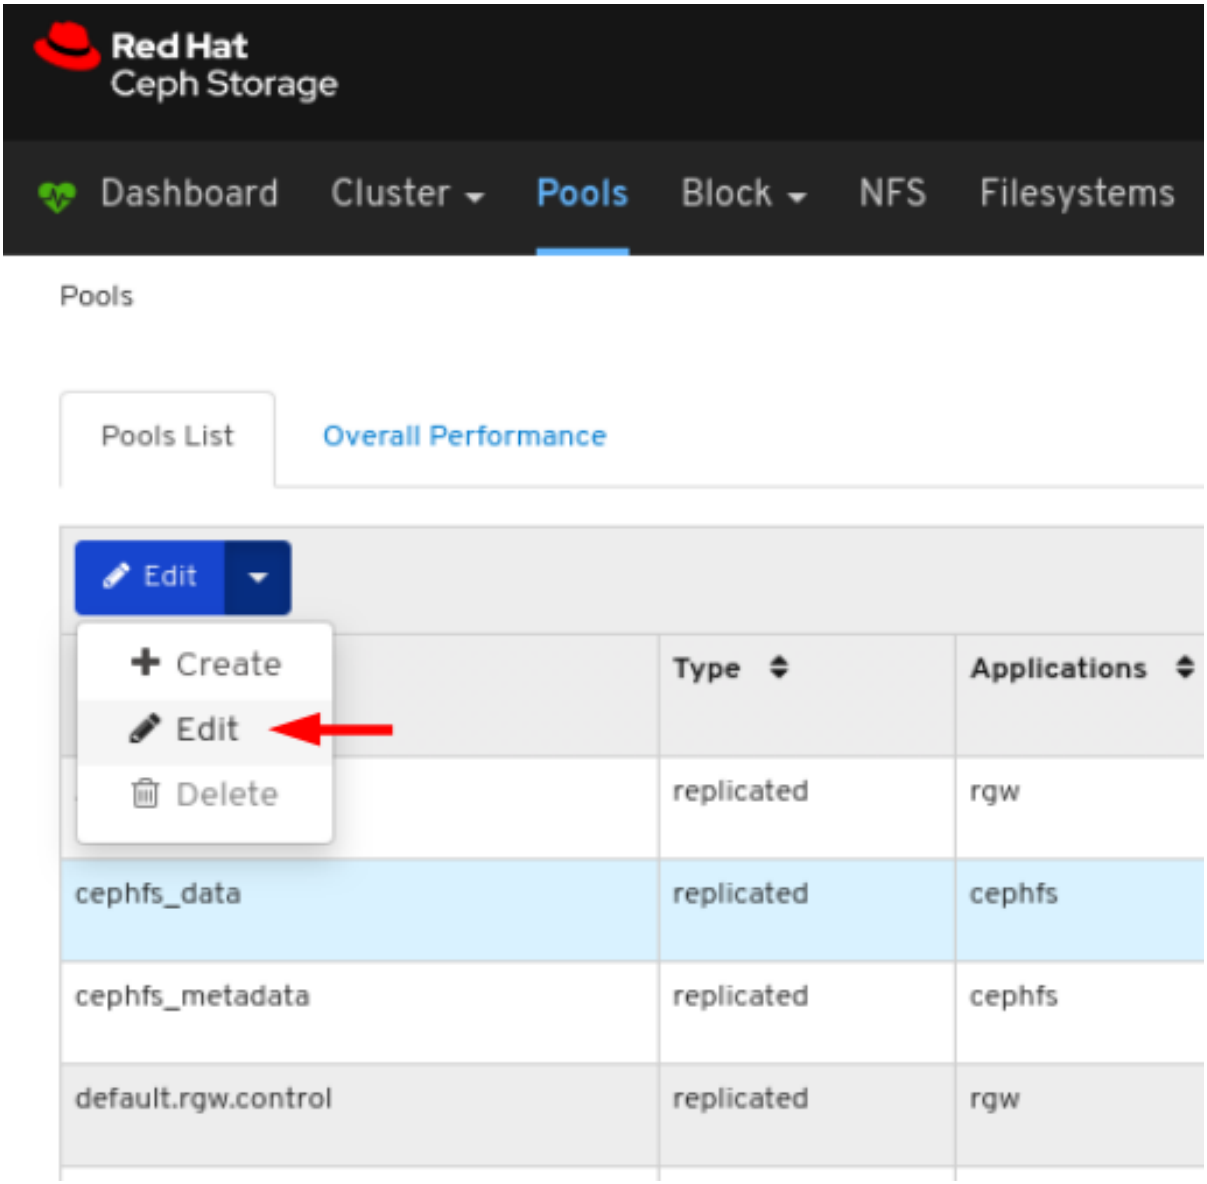

5. In the dialog window, edit the required parameters and click the *EditPool* button:

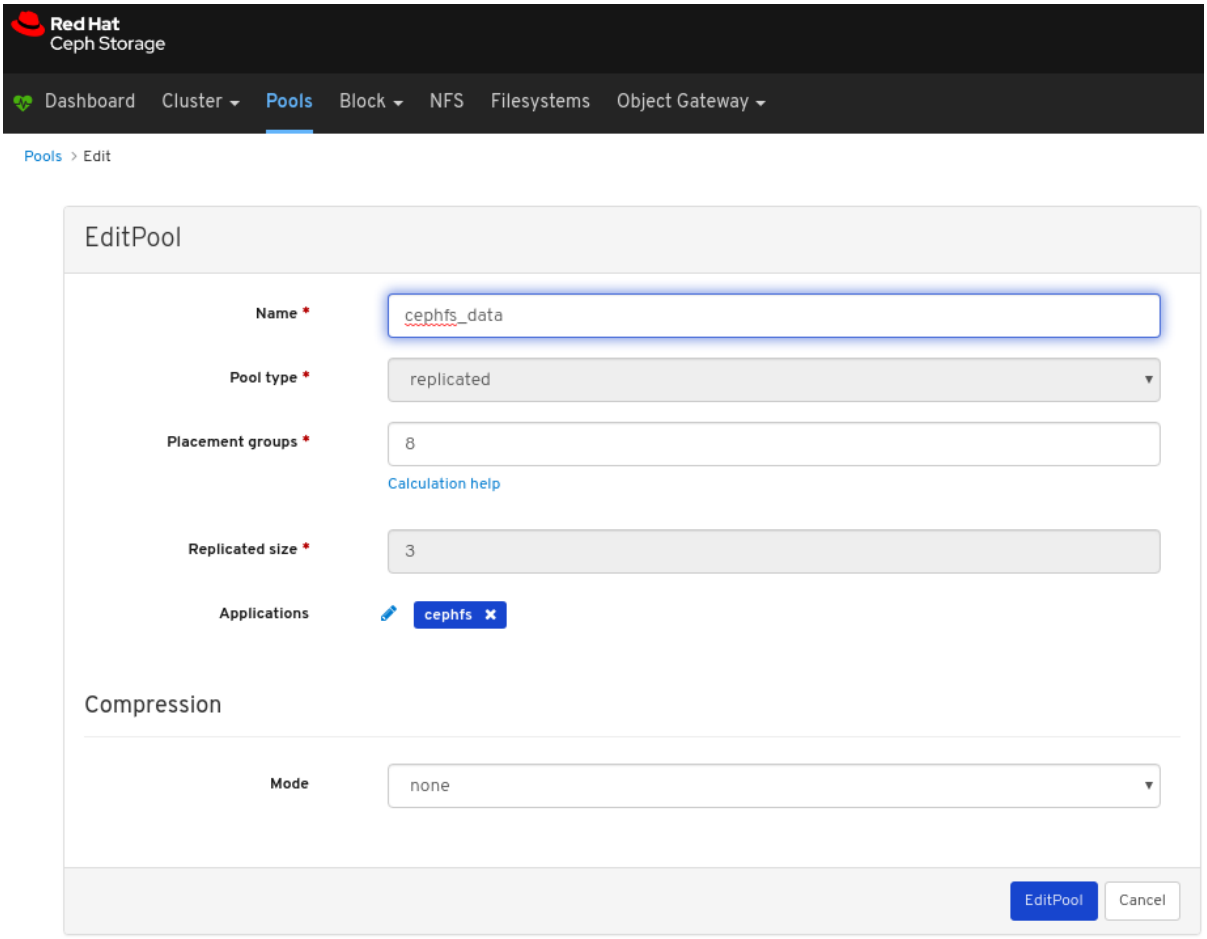

6. A notifications towards the top right corner of the page indicates the pool was updated successfully.

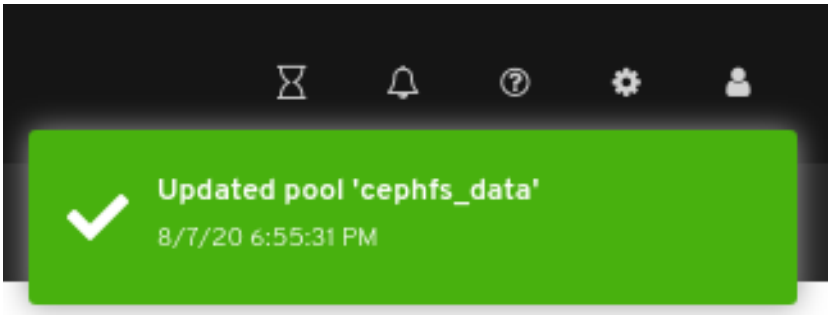

Additional Resources

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

7.10.3. Deleting pools

The Red Hat Ceph Storage Dashboard allows deletion of pools.

Prerequisites

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

- A pool is created.

Procedure

1. Log in to the dashboard.
2. Ensure the values of **mon_allow_pool_delete** is set to **true**:
   a. On the navigation bar, click *Cluster* and then click *Configuration*.
   b. In the *Level* drop-down menu, select **Advanced**:

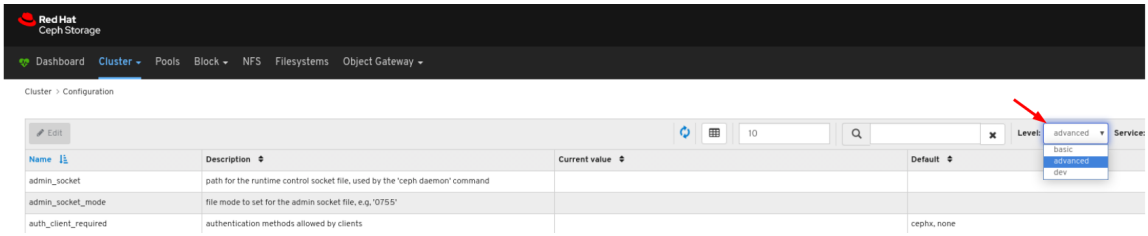

   c. Search for **mon_allow_pool_delete** and set the values to **true**

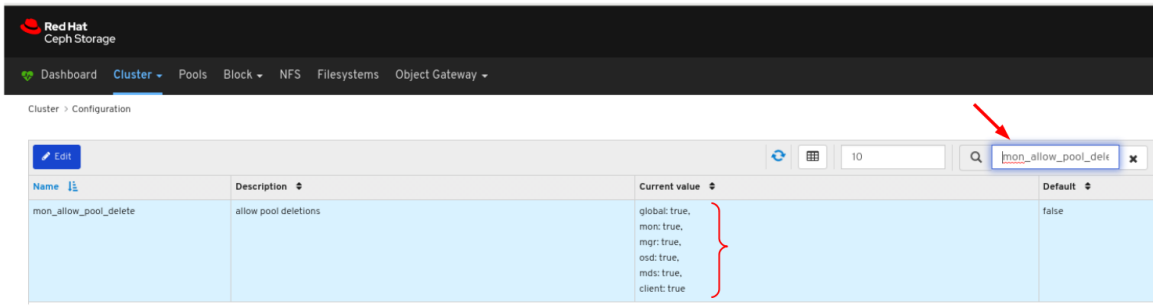

3. On the navigation bar, click *Pools*:

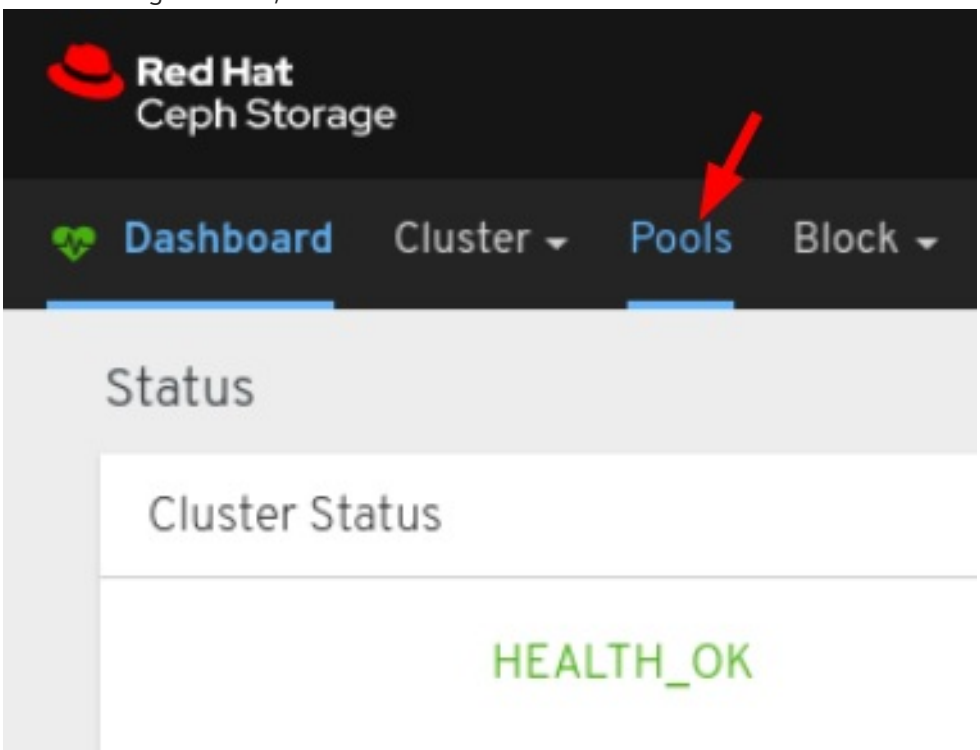

4. To delete the pool, click on its row:

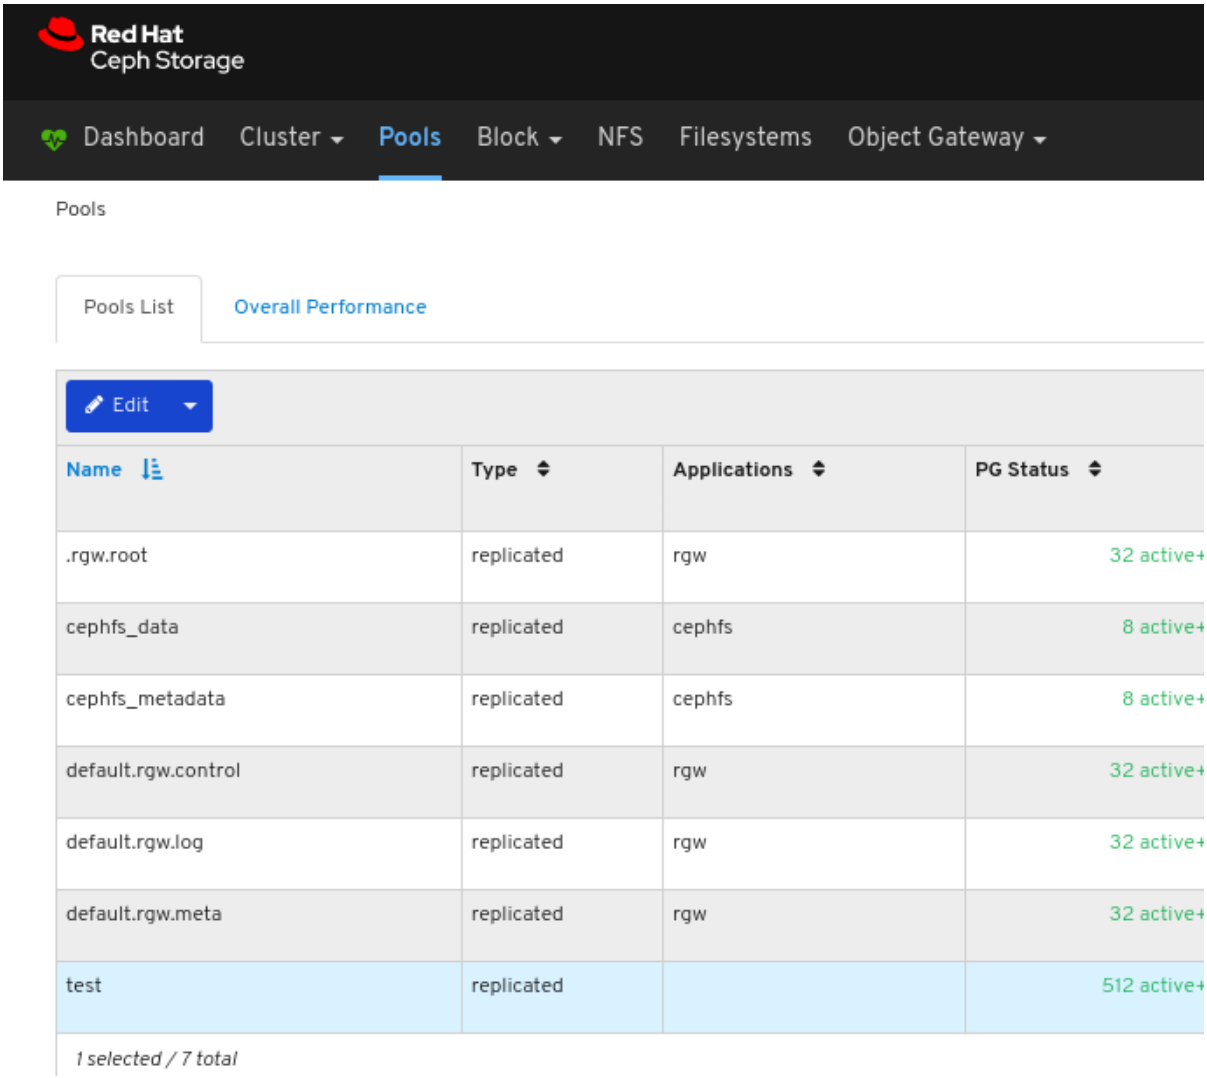

5. Select *Delete* in the *Edit* drop-down:

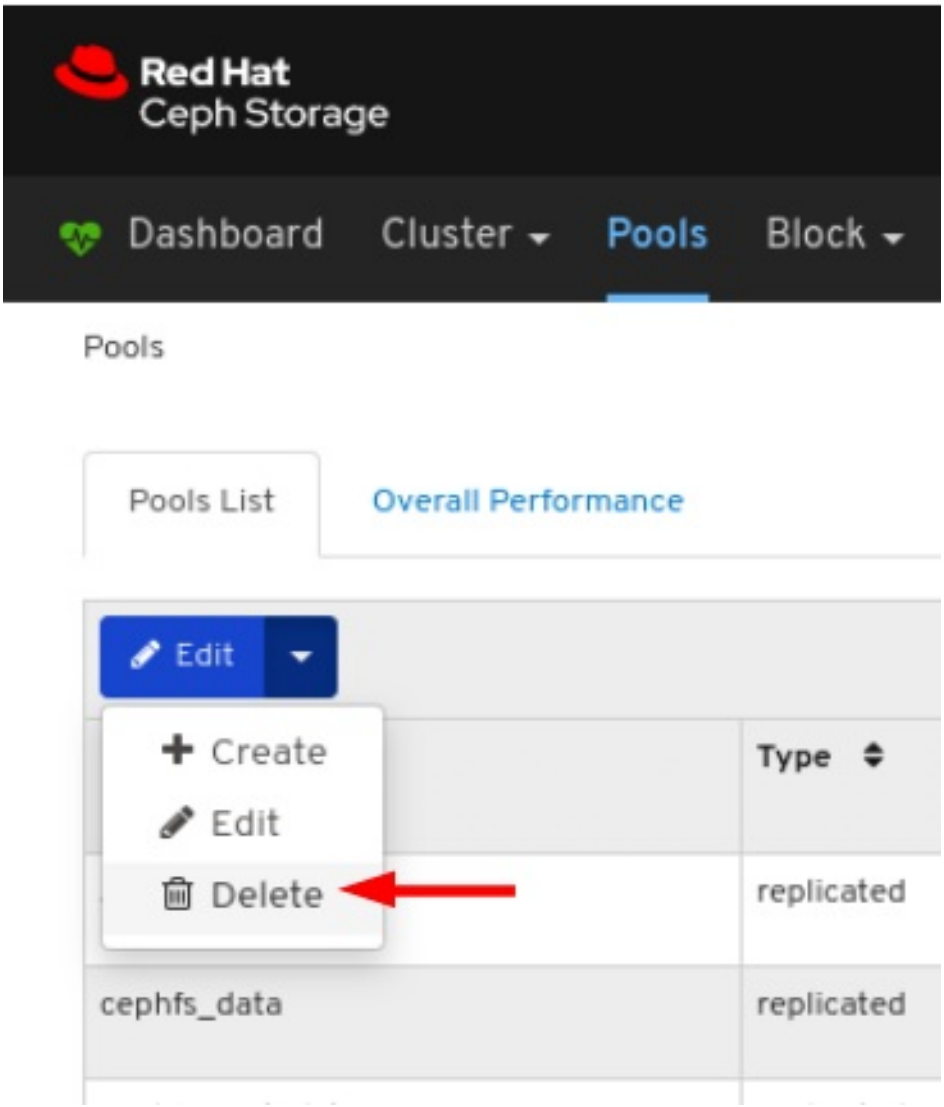

6. In the *Delete Pool* dialog window, Click the *Yes, I am sure* box and then Click *Delete Pool* to save the settings:

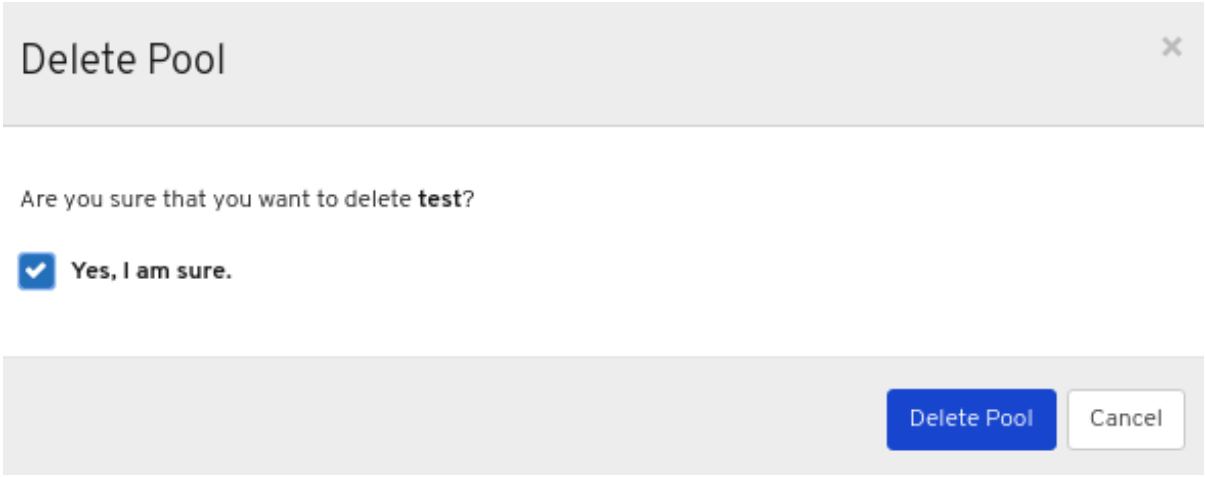

7. A notification towards the top right corner of the page indicates the pool was deleted successfully.

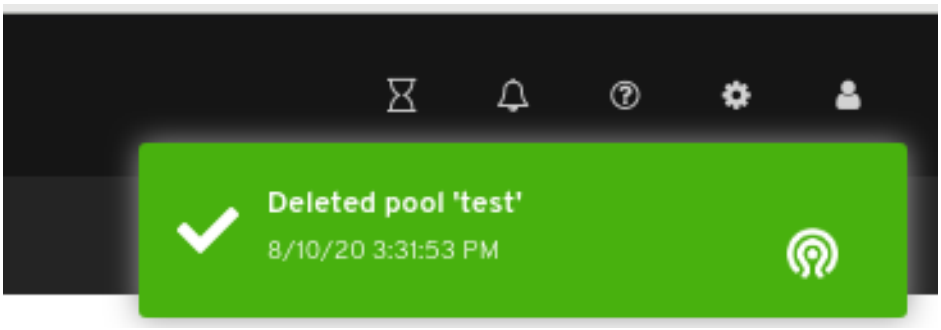

Additional Resources

- See the *Ceph pools* in the *Red Hat Ceph Storage Architecture Guide* for more information.
- See the *Monitoring Configuration* in the *Red Hat Ceph Storage Dashboard Guide* for more inforamtion.
- See the *Pool values* in the *Red Hat Ceph Storage Storage Strategies Guide* for more information on Compression Modes.

CHAPTER 8. OBJECT GATEWAY

As a storage administrator, the object gateway functions of the dashboard allow you to manage and monitor the Ceph Object Gateway.

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.

Object gateway functions are divided between daemon functions, user functions, and bucket functions.

8.1. PREREQUISITES

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

8.2. OBJECT GATEWAY DAEMON FUNCTIONS

As a storage administrator, the Red Hat Ceph Storage Dashboard allows you to view and monitor information about Ceph Object Gateway daemons.

8.2.1. Prerequisites

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

8.2.2. Viewing object gateway daemons

The dashboard allows you to view a list of all Ceph Object Gateway daemons.

Prerequisites

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

Procedure

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

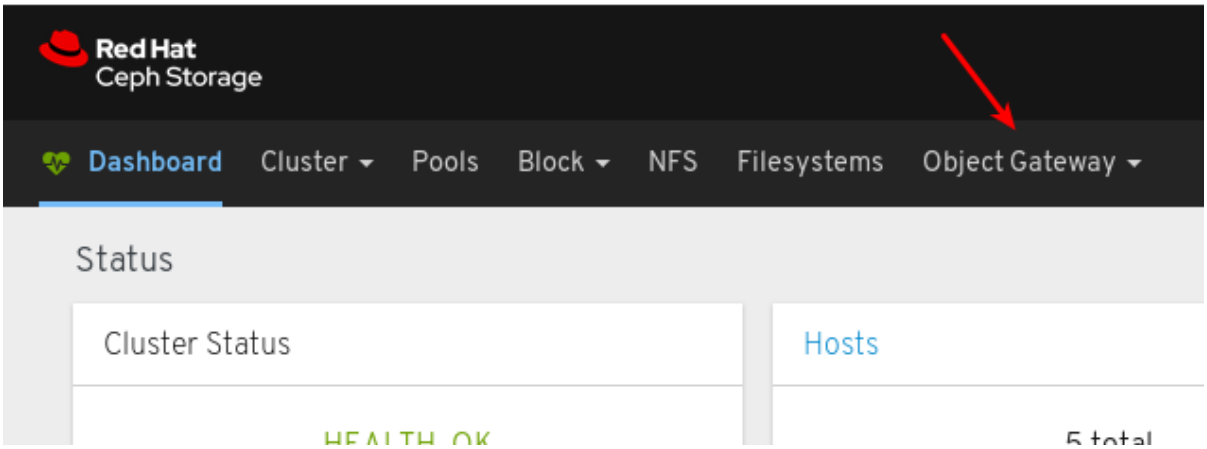

3. Click *Daemons*.

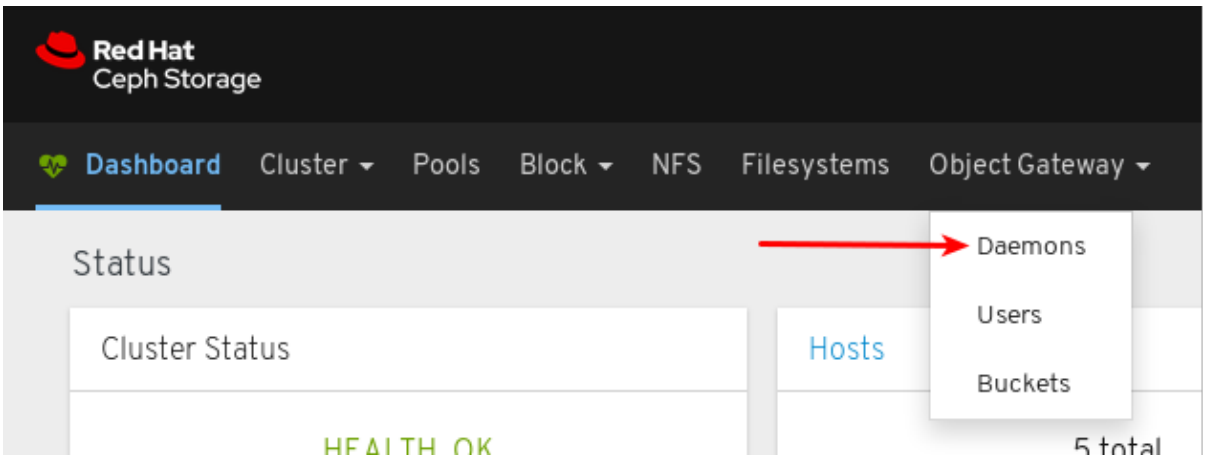

4. In the example below, you can see a daemon with the *ID* **jb-ceph4-rgw.rgw0** in the *Daemons List*.

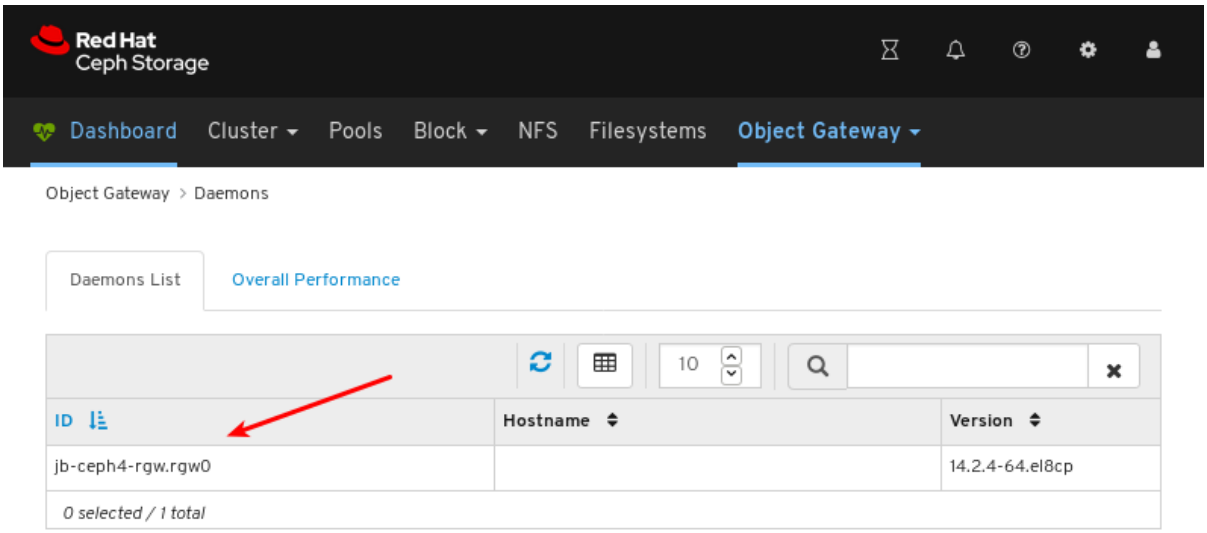

5. To view details, select the daemon by clicking the row for **jb-ceph4-rgw.rgw0**:

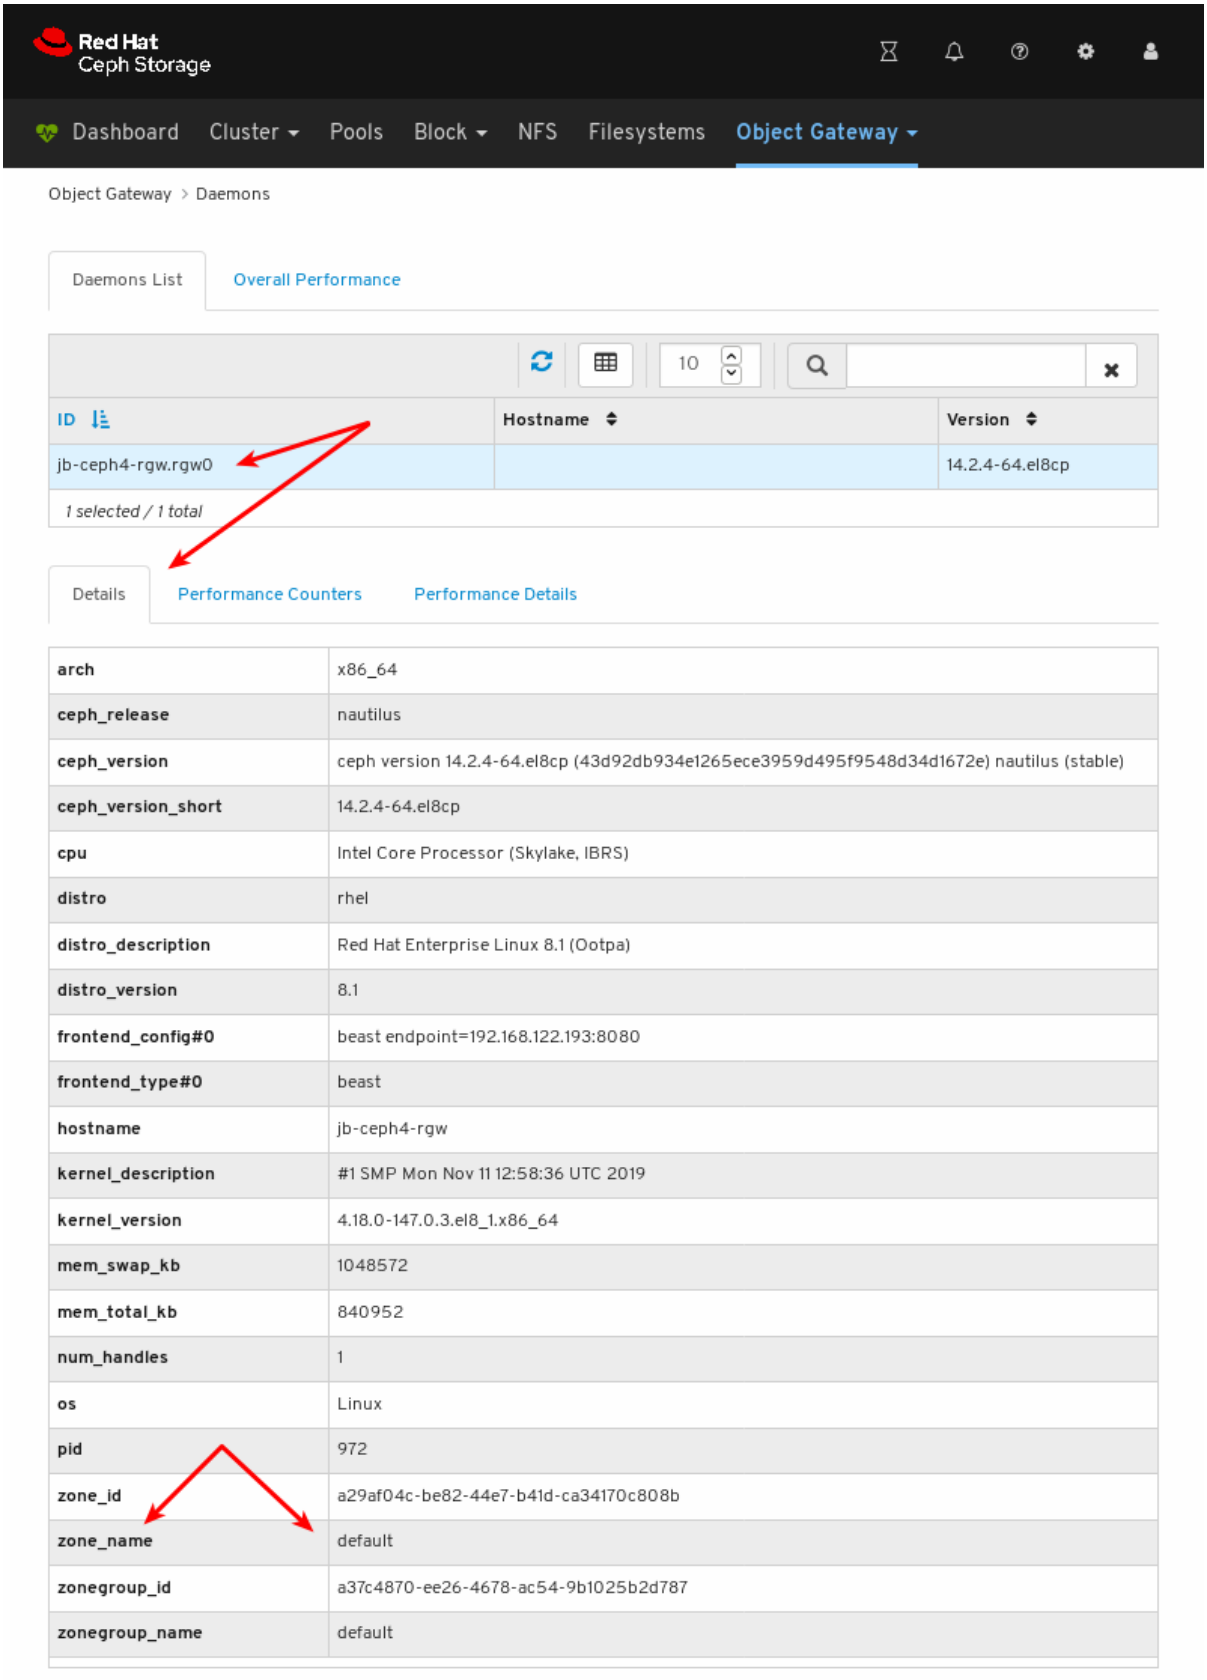

You can see the zone name the daemon is serving is **default**.

Additional Resources

- For information on how to install the Ceph Object Gateway, see Installing the Ceph Object Gateway in the Installation Guide.

- For information on how to add object gateway login credentials to the dashboard, see Adding object gateway login credentials to the dashboard in the Dashboard guide.

8.3. OBJECT GATEWAY USER FUNCTIONS

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

8.3.1. Prerequisites

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

8.3.2. Viewing object gateway users

The dashboard allows you to view a list of all Ceph Object Gateway users.

Prerequisites

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

Procedure

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

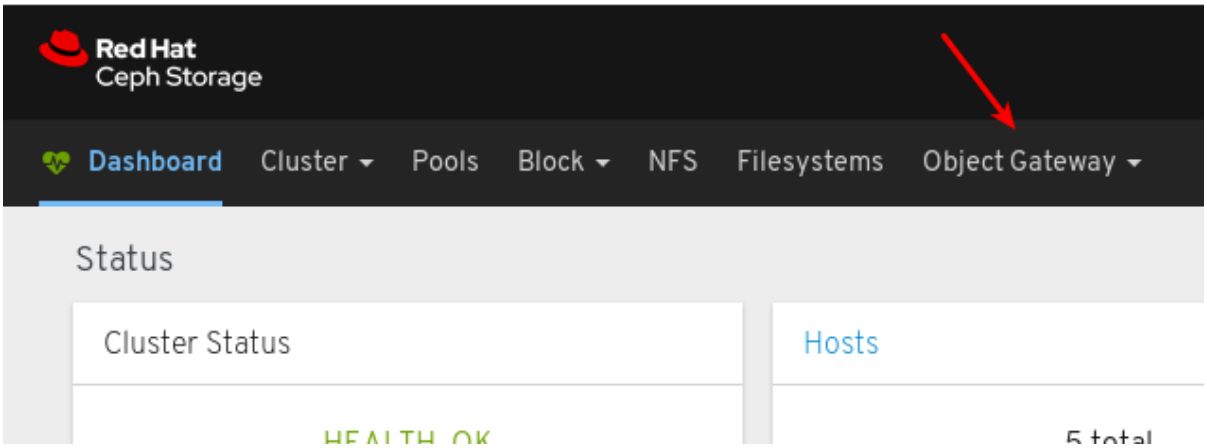

3. Click *Users*.

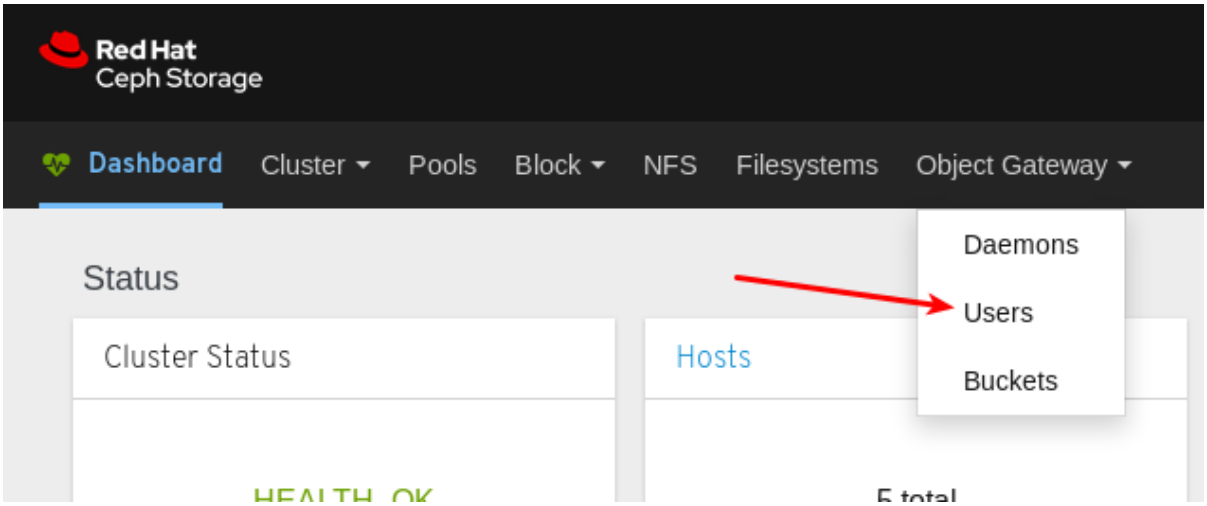

4. In the example below, you can see a user named **rgw-user** in the table.

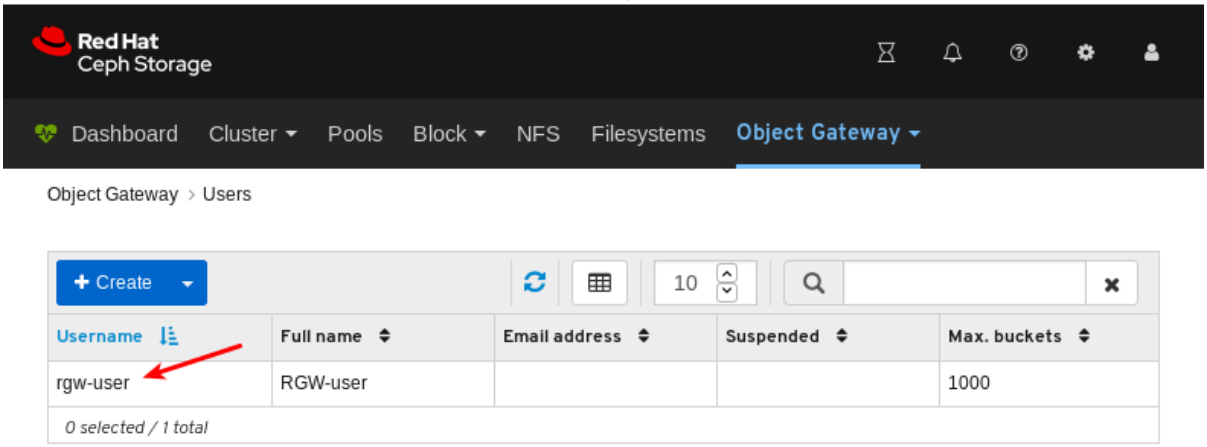

5. To view details, select the user by clicking the row for **rgw-user**:

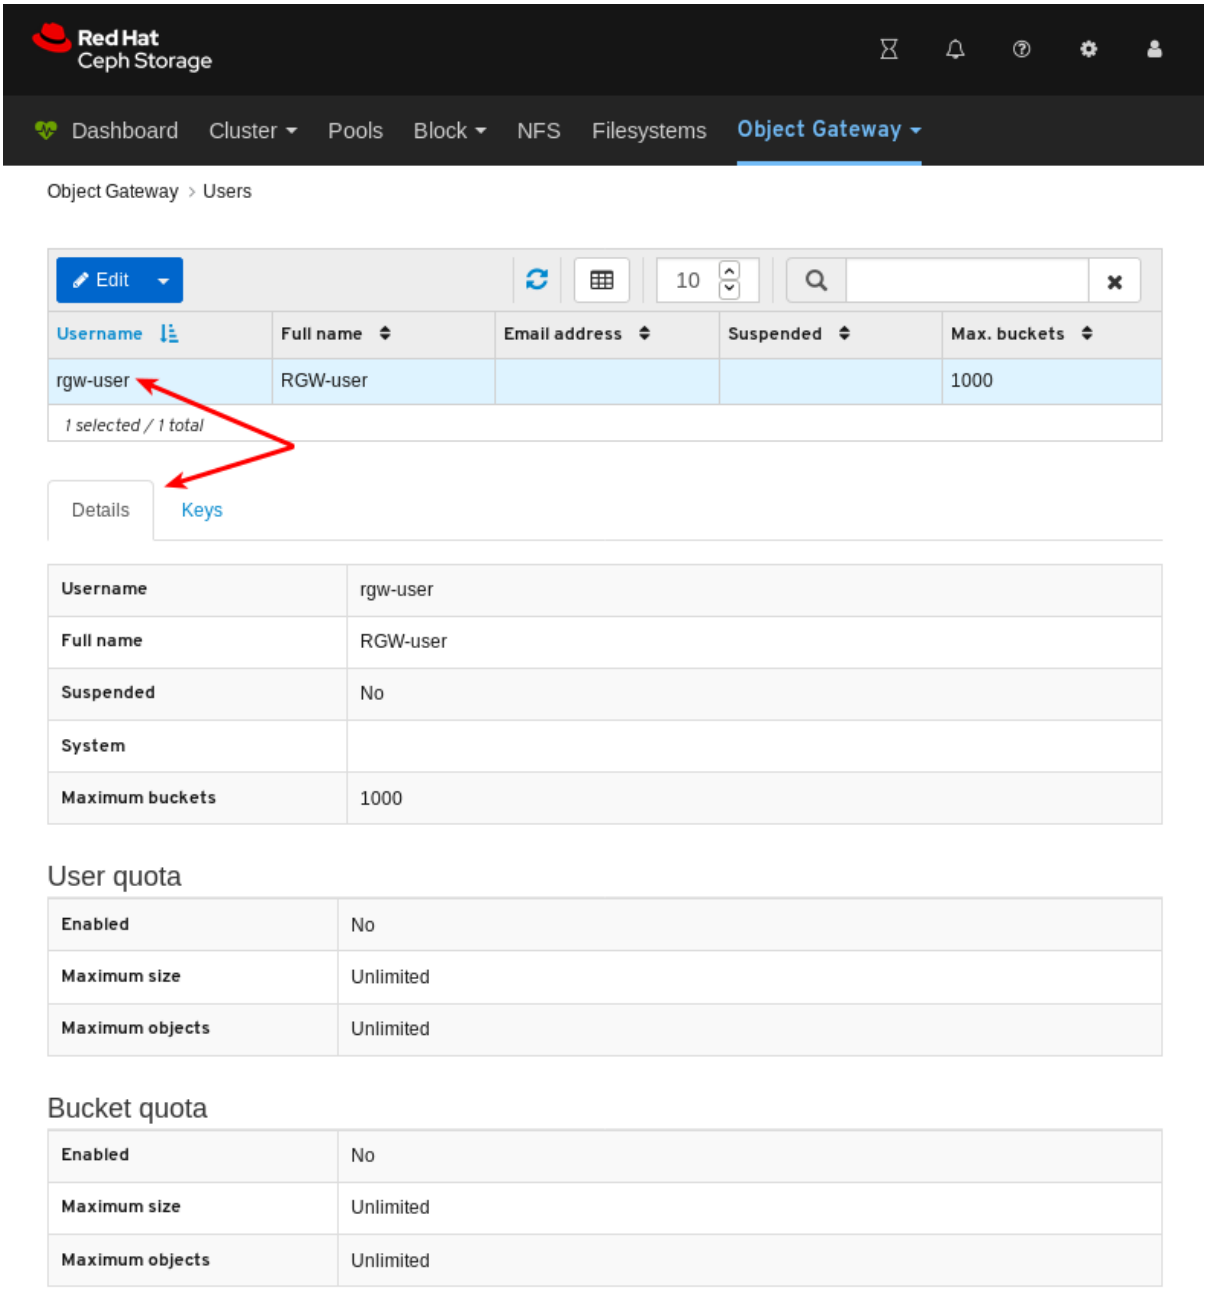

Additional Resources

- For information on how to install the Ceph Object Gateway, see Installing the Ceph Object Gateway in the Installation Guide.
- For information on how to add object gateway login credentials to the dashboard, see Adding object gateway login credentials to the dashboard in the Dashboard guide.
- For more information on the Ceph Object Gateway, see the Object Gateway Configuration and Administration Guide.

8.3.3. Creating object gateway users

The dashboard allows you to create Ceph Object Gateway users.

Prerequisites

- A running Red Hat Ceph Storage cluster.

- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.

Procedure

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

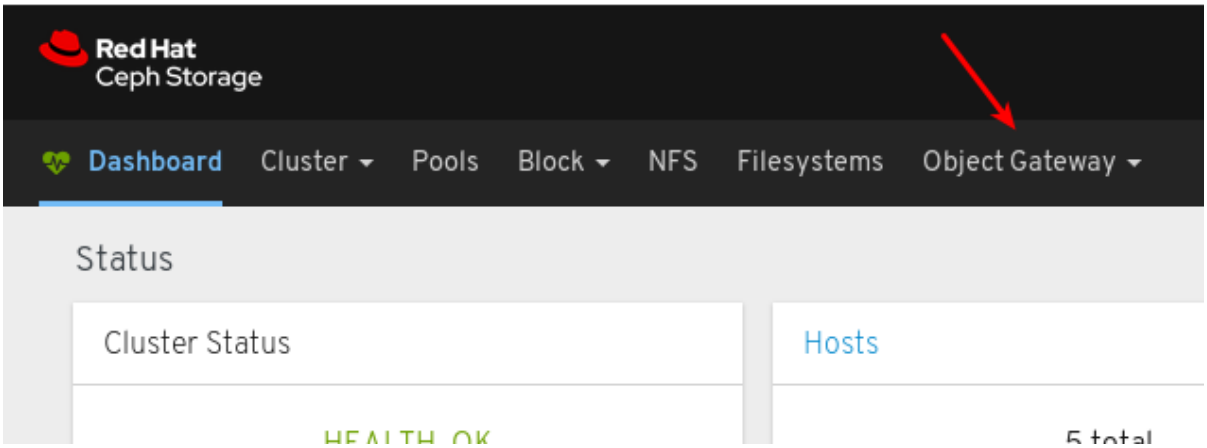

3. Click *Users*.

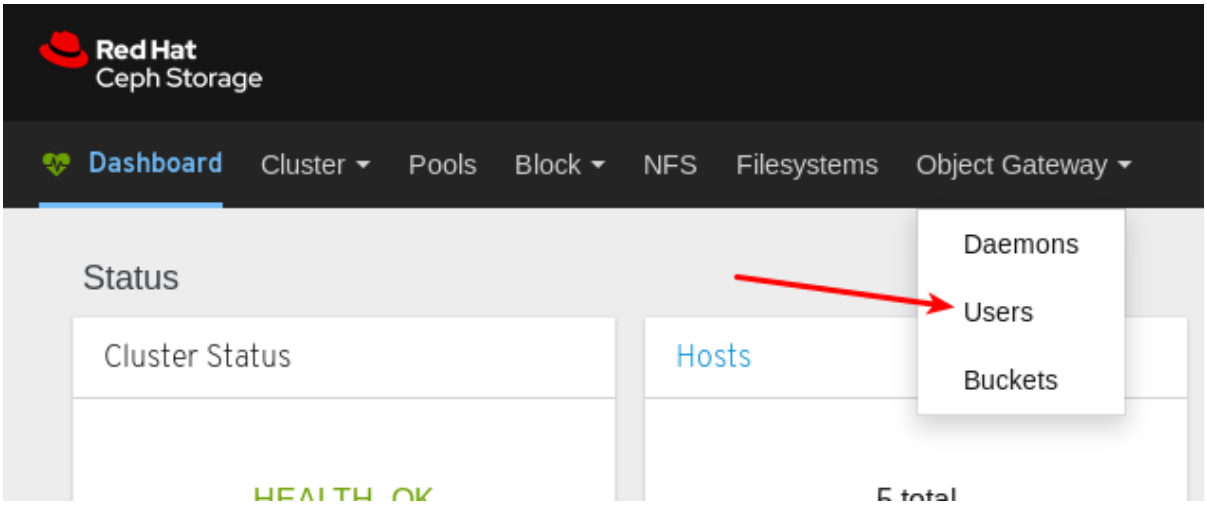

4. Click *Create*.

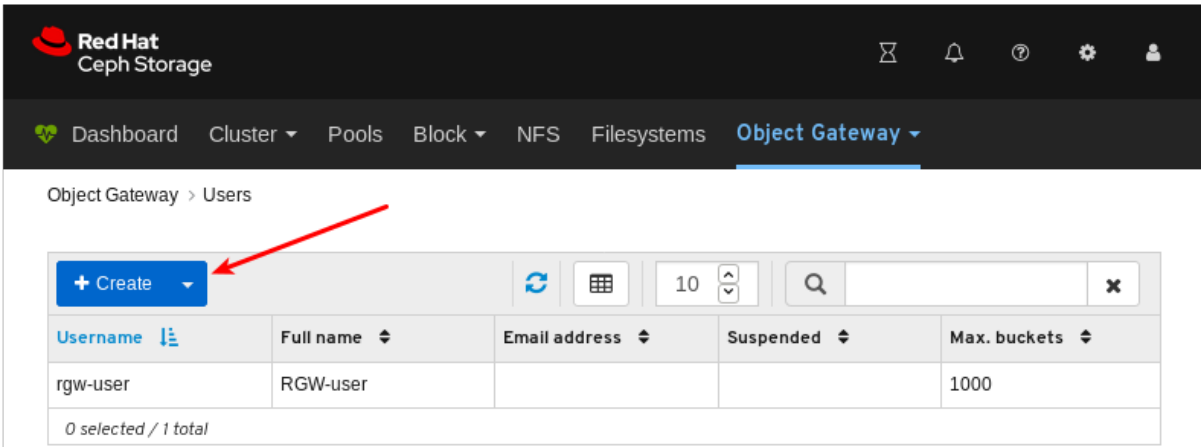

5. Set the user name, full name, and edit the maximum number of buckets if required.

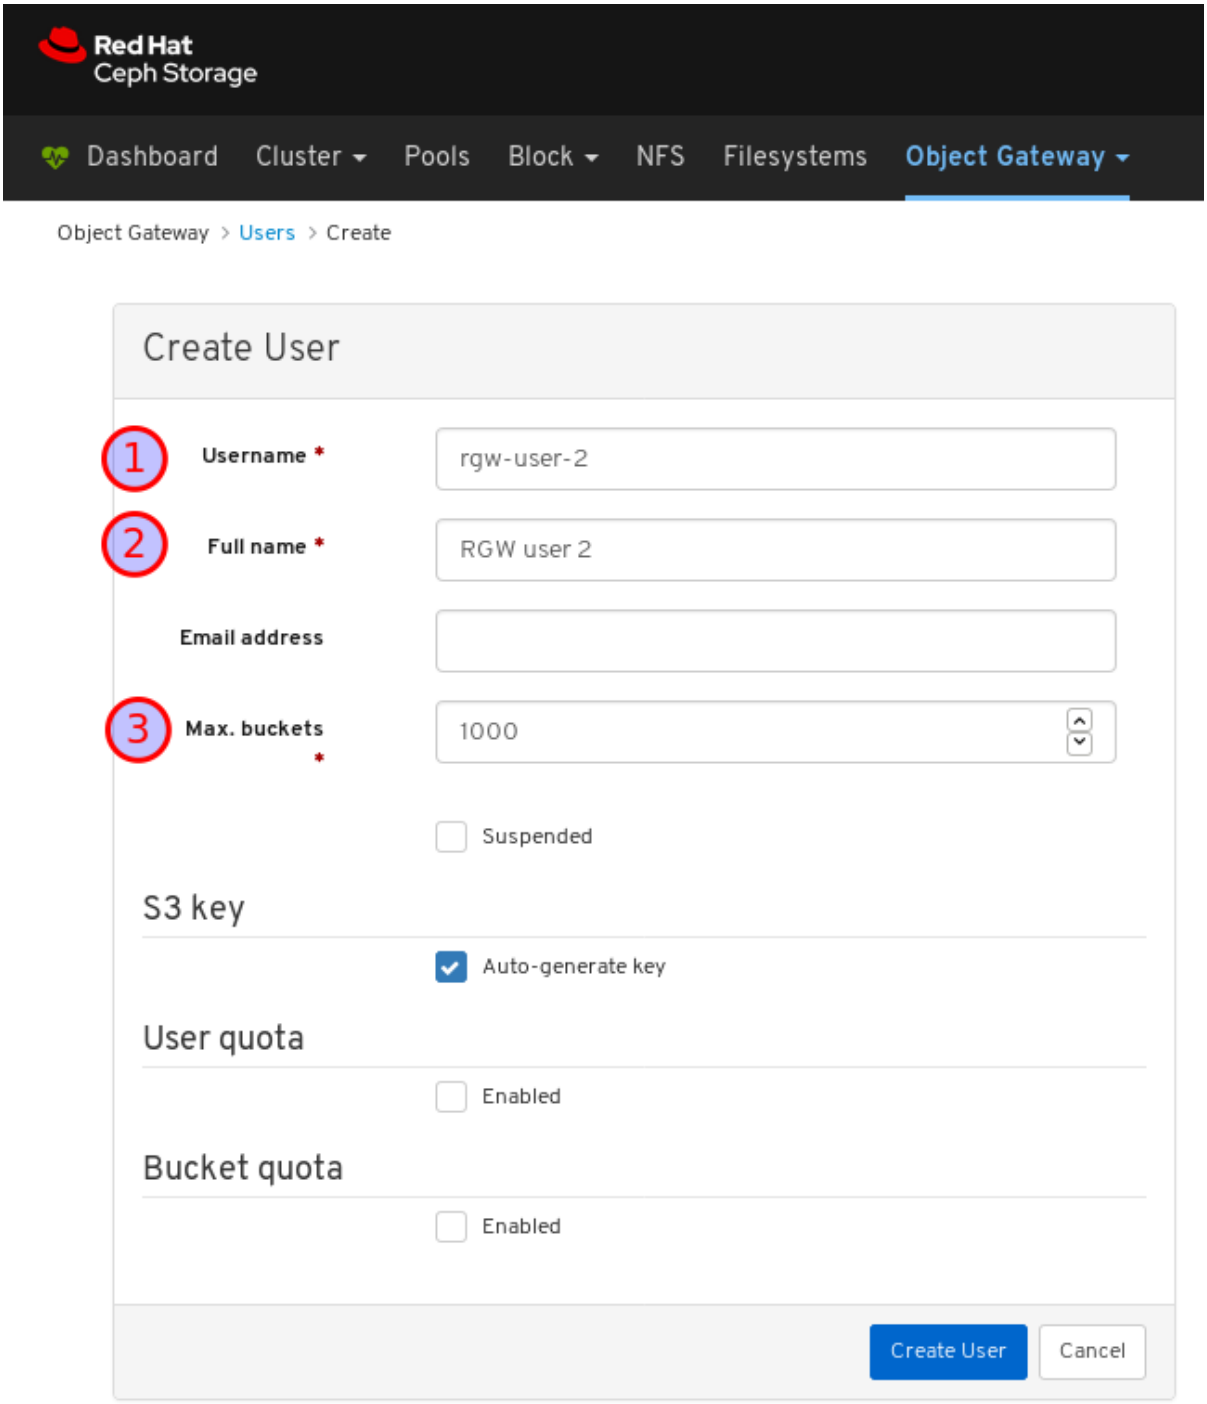

6. Optional: Set an email address or suspended status.

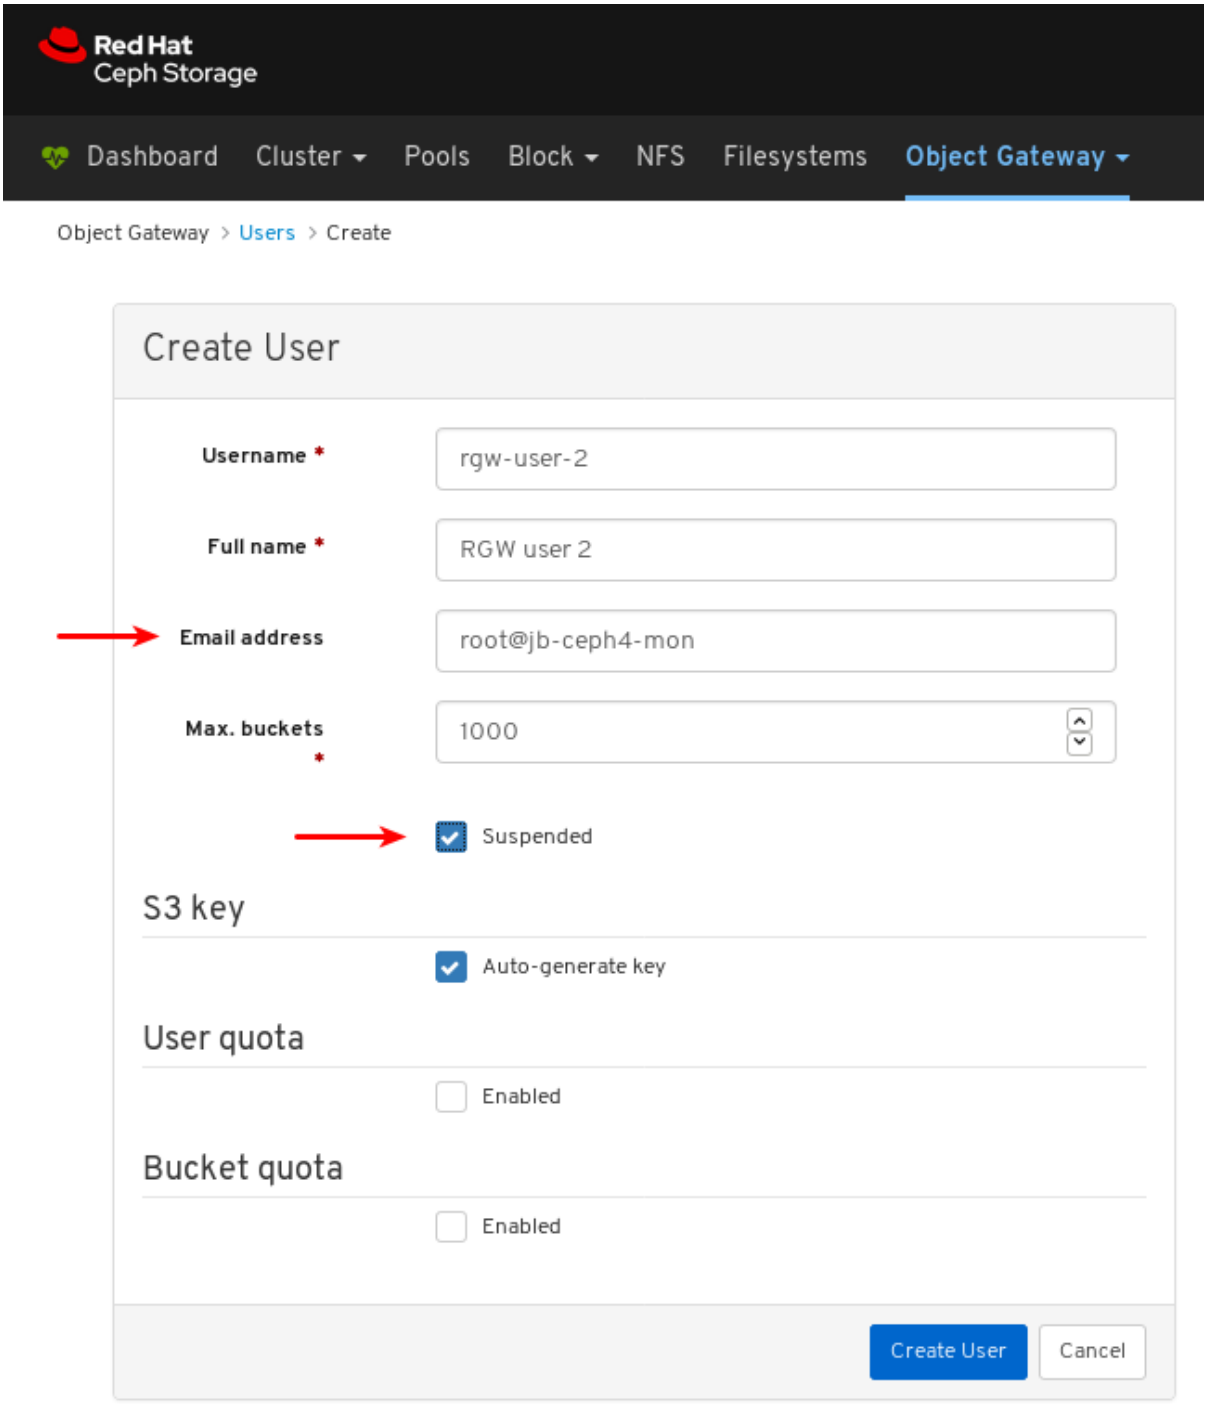

7. Optional: Set a custom access key and secret key by unchecking *Auto-generate key*.

   a. Uncheck *Auto-generate key*:

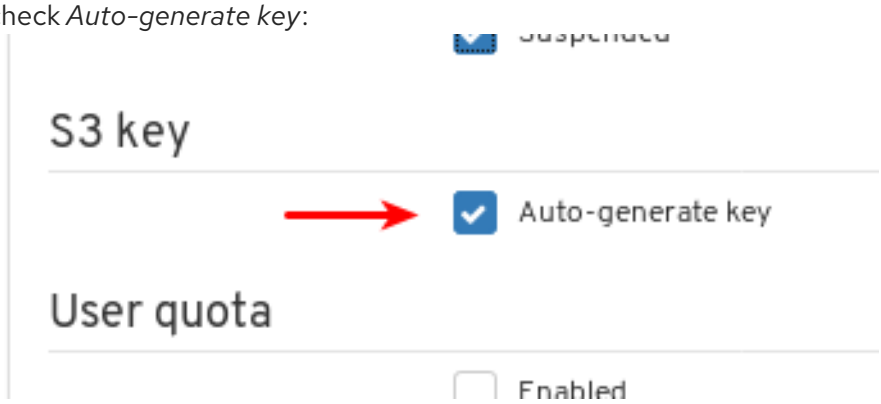

   b. Set the access key and secret key:

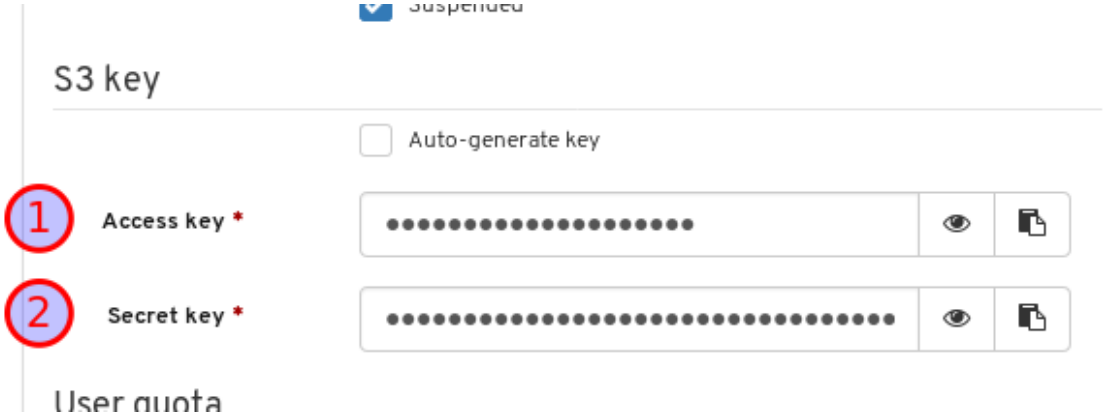

8. Optional: Set a user quota.

   a. Check *Enabled* under *User quota*:

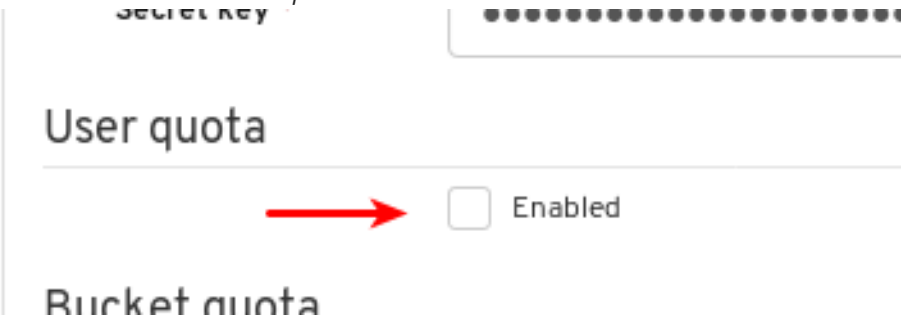

   b. Uncheck *Unlimited size* or *Unlimited objects*:

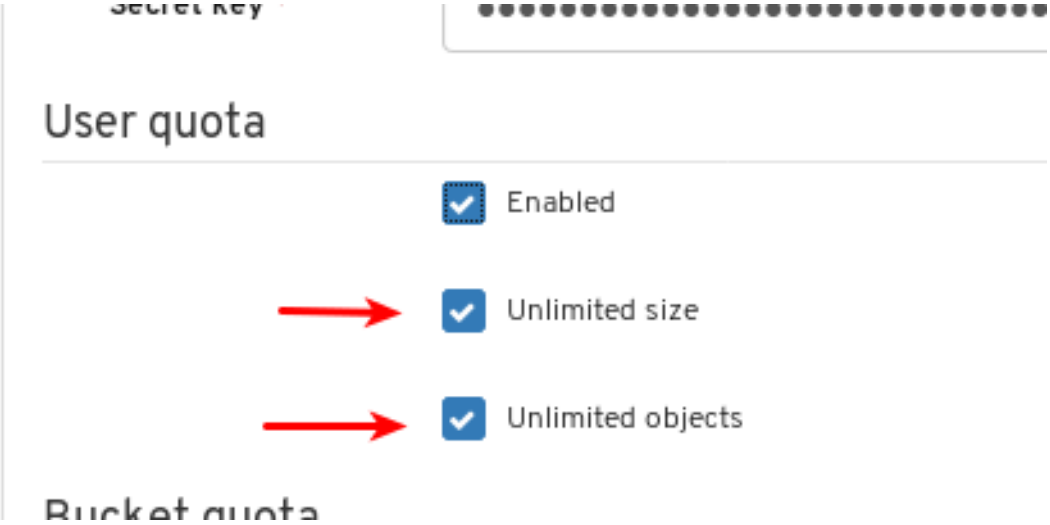

   c. Enter the required values for *Max. size* or *Max. objects*:

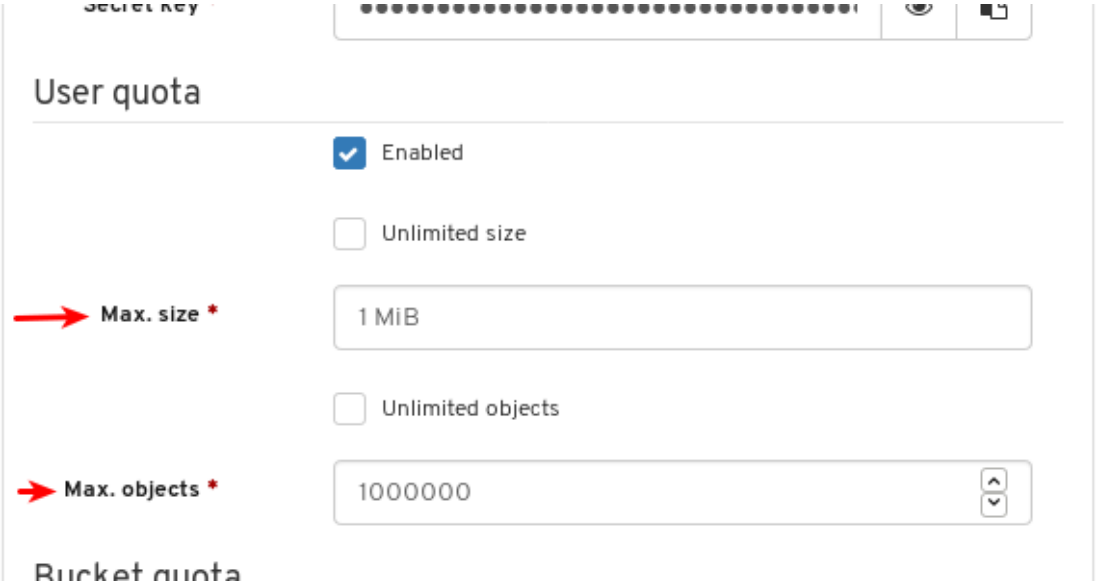

9. Optional: Set a bucket quota.

   a. Check *Enabled* under *Bucket quota*:

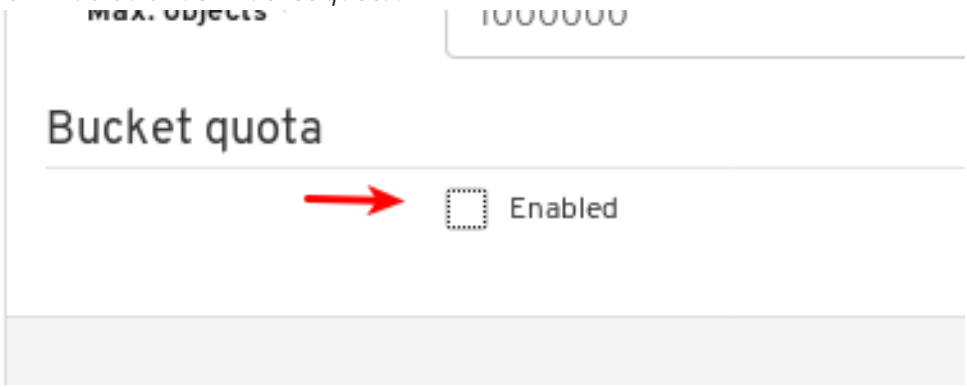

   b. Uncheck *Unlimited size* or *Unlimited objects*:

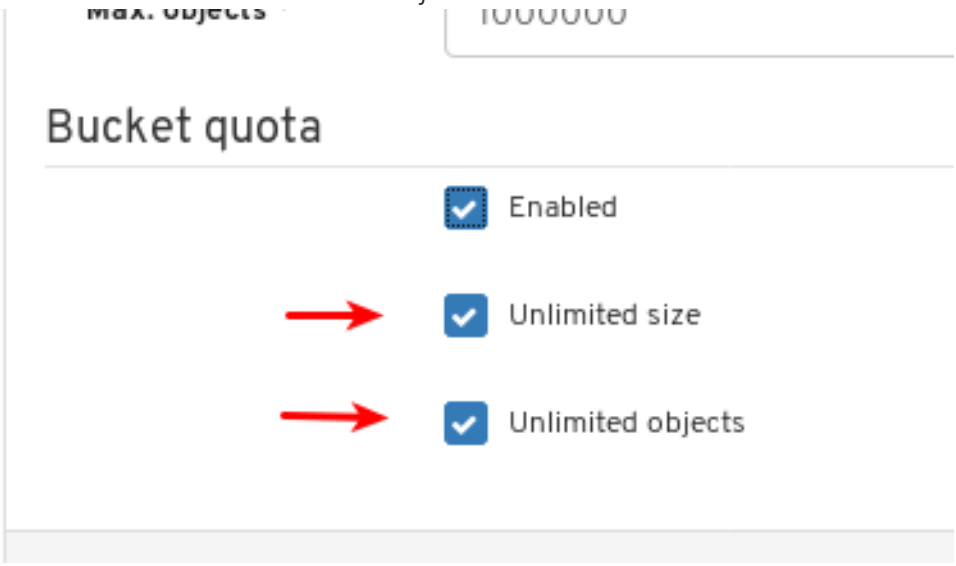

   c. Enter the required values for *Max. size* or *Max. objects*:

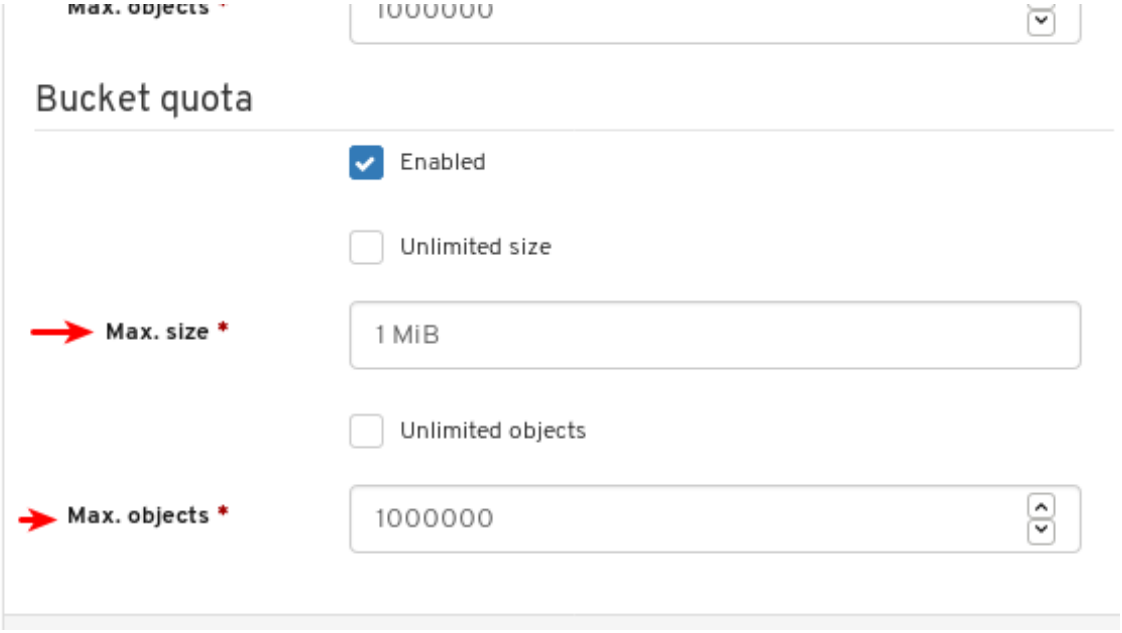

10. Finalize the user creation by clicking *Create User*.

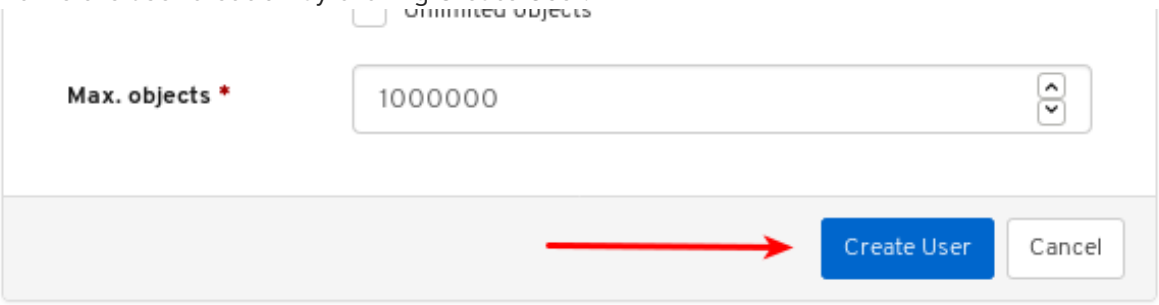

11. Verify the user creation was successful. A notification confirms the user was created and the user can be seen in the table of users.

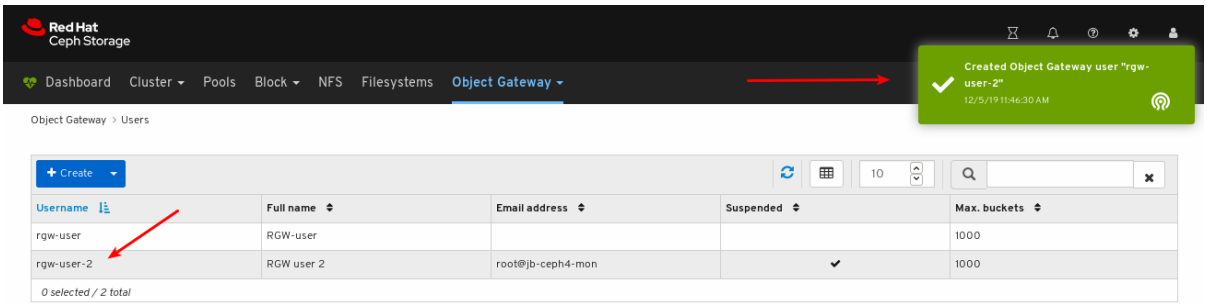

Additional Resources

- For information on how to install the Ceph Object Gateway, see Installing the Ceph Object Gateway in the Installation Guide.
- For information on how to add object gateway login credentials to the dashboard, see Adding object gateway login credentials to the dashboard in the Dashboard guide.
- For more information on the Ceph Object Gateway, see the Object Gateway Configuration and Administration Guide.

8.3.4. Editing object gateway users

The dashboard allows you to edit Ceph Object Gateway users.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*:

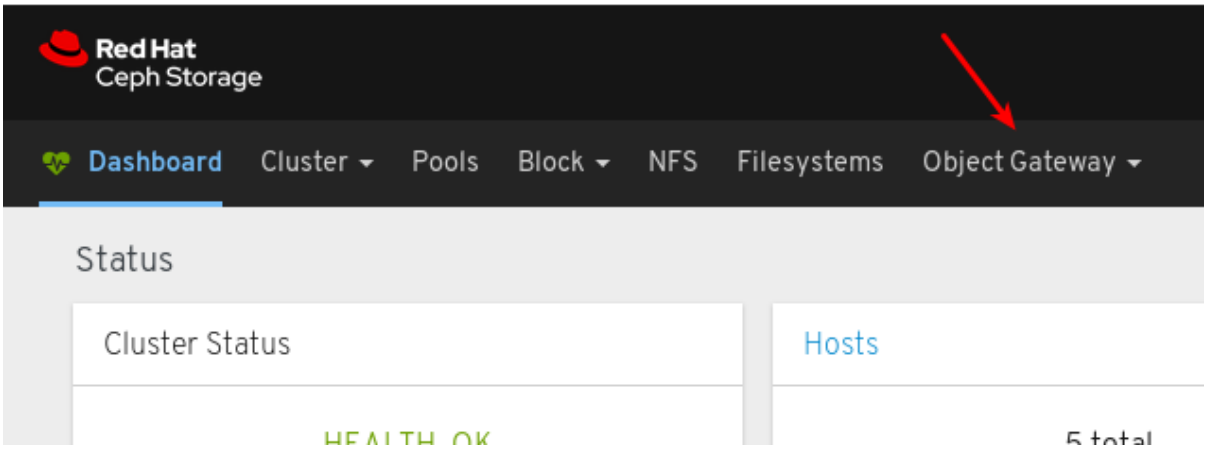

3. Click *Users*:

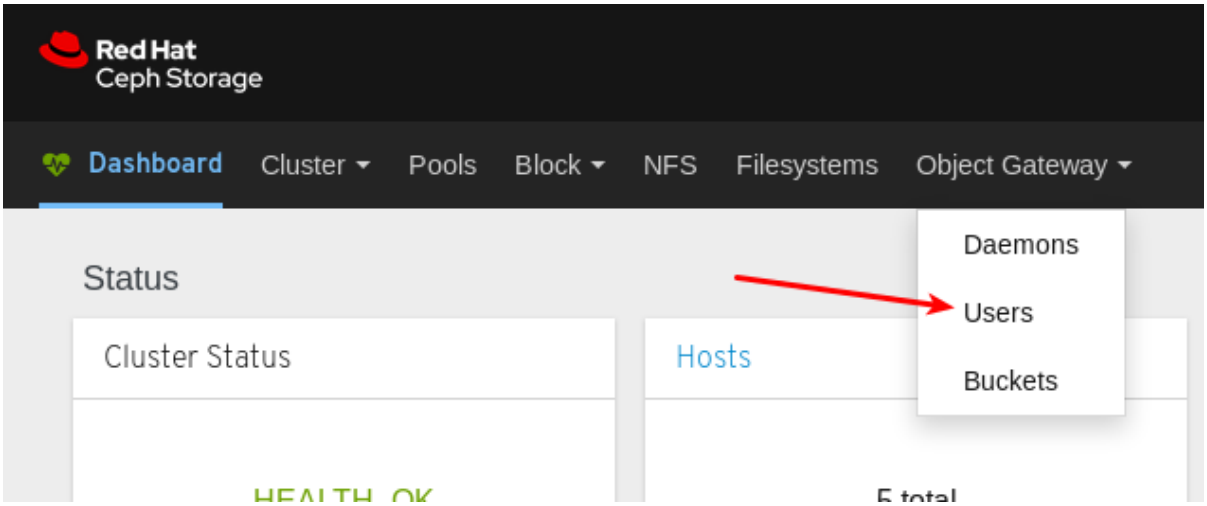

4. To edit the user capabilities, click its row:

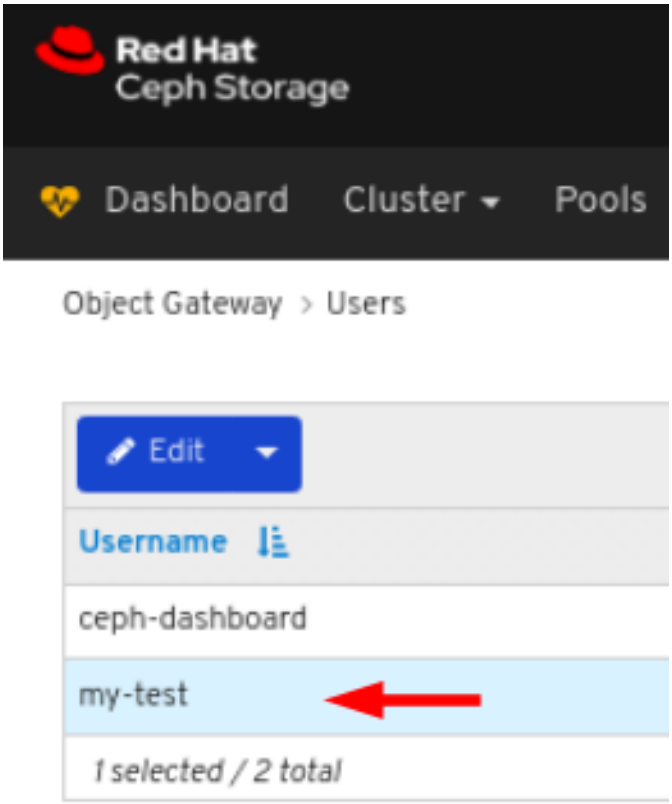

5. Select *Edit* In the *Edit* drop-down:

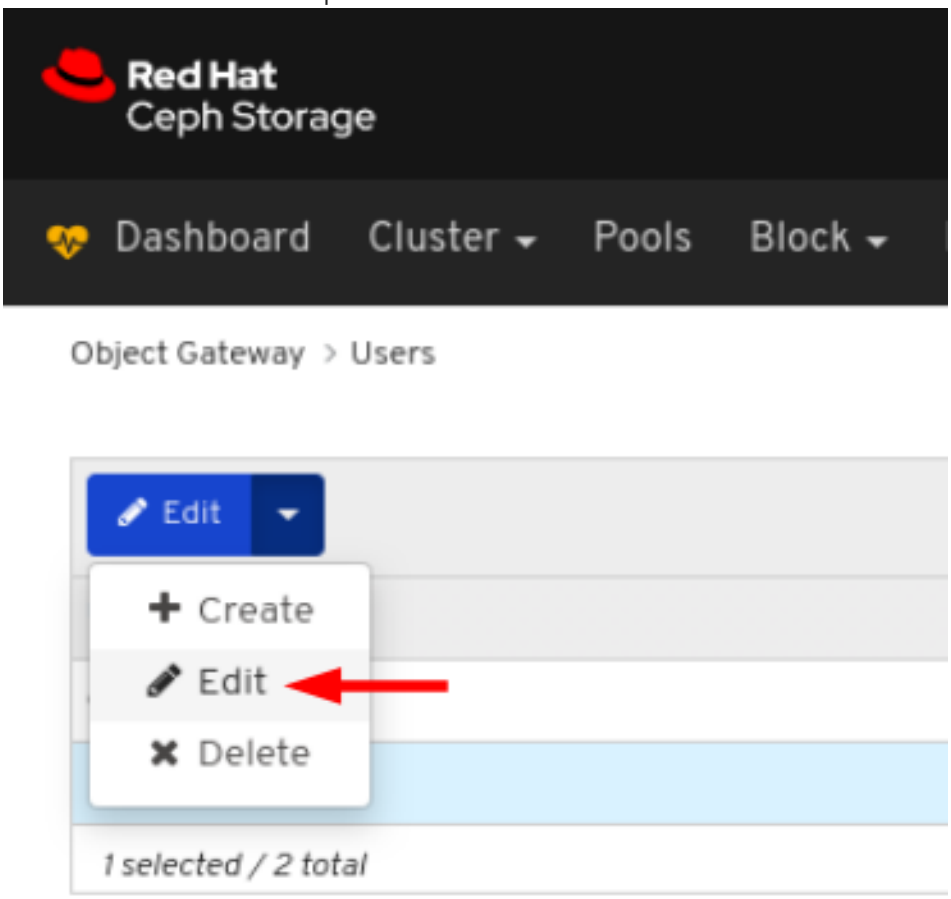

6. In the *EditUser* window, edit the required parameters and click the *EditUser* button:

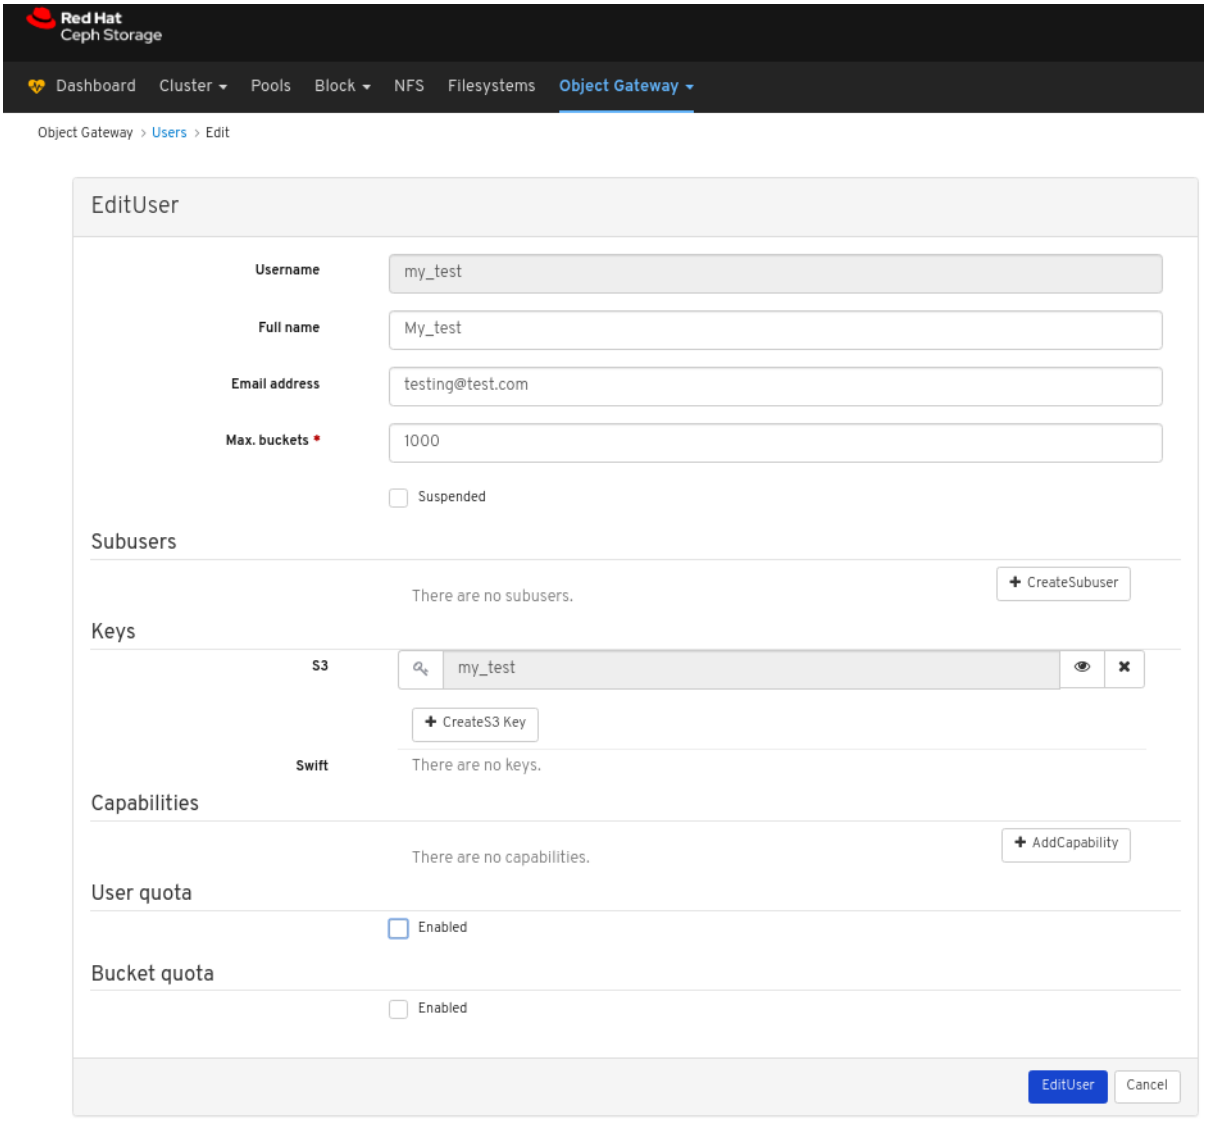

7. A notification towards the top right corner of the page indicates the user was updated successfully.

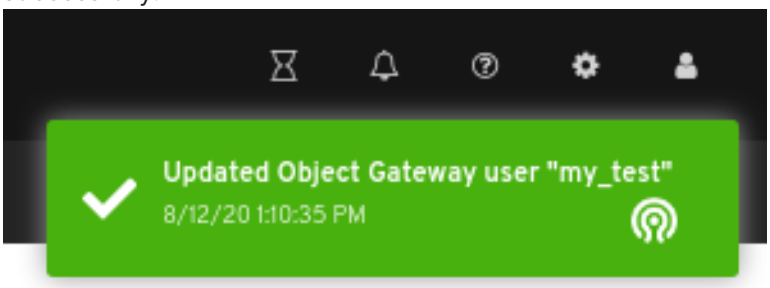

8.3.5. Deleting object gateway users

The dashboard allows you to delete Ceph Object Gateway users.

Prerequisites

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

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*:

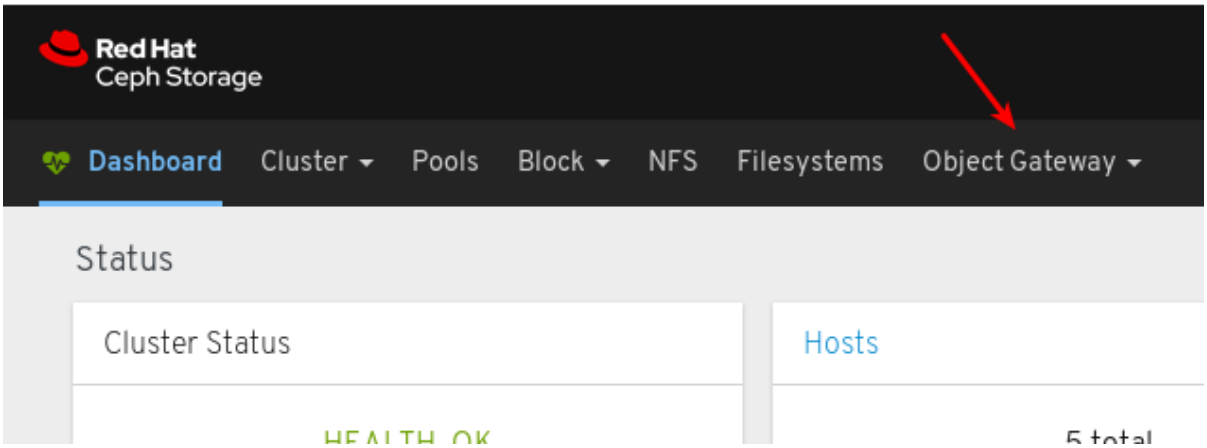

3. Click *Users*:

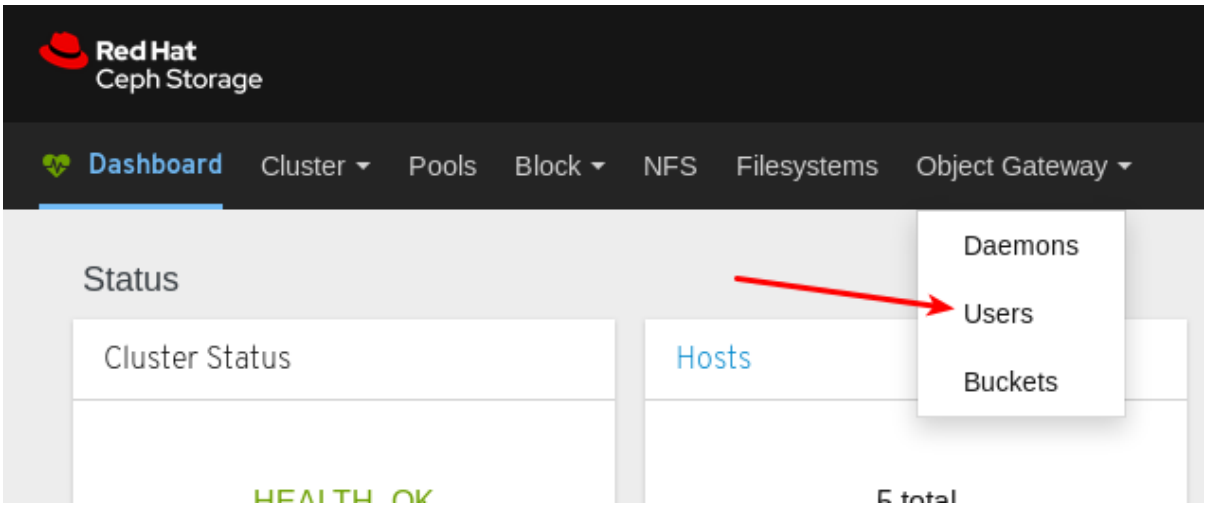

4. To delete the user, click its row:

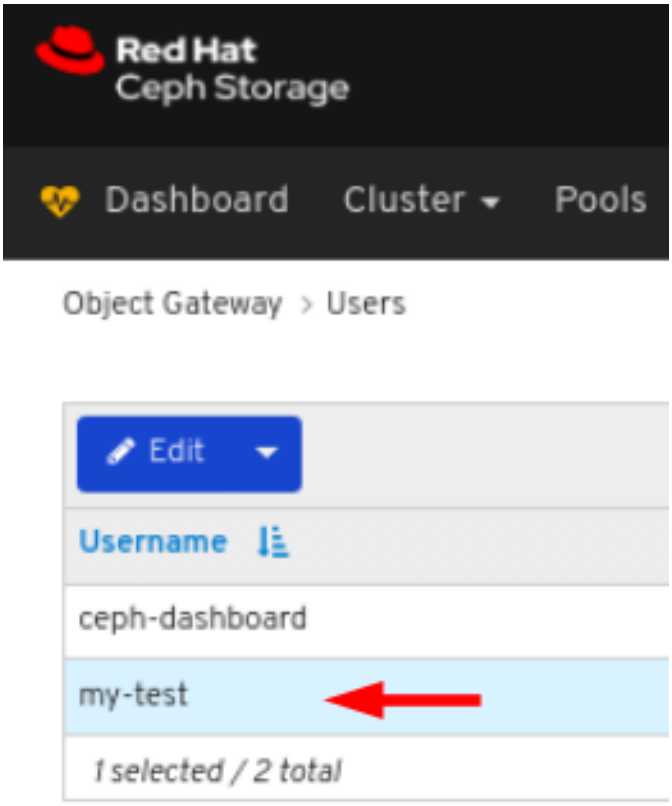

5. Select *Delete* In the *Edit* drop-down:

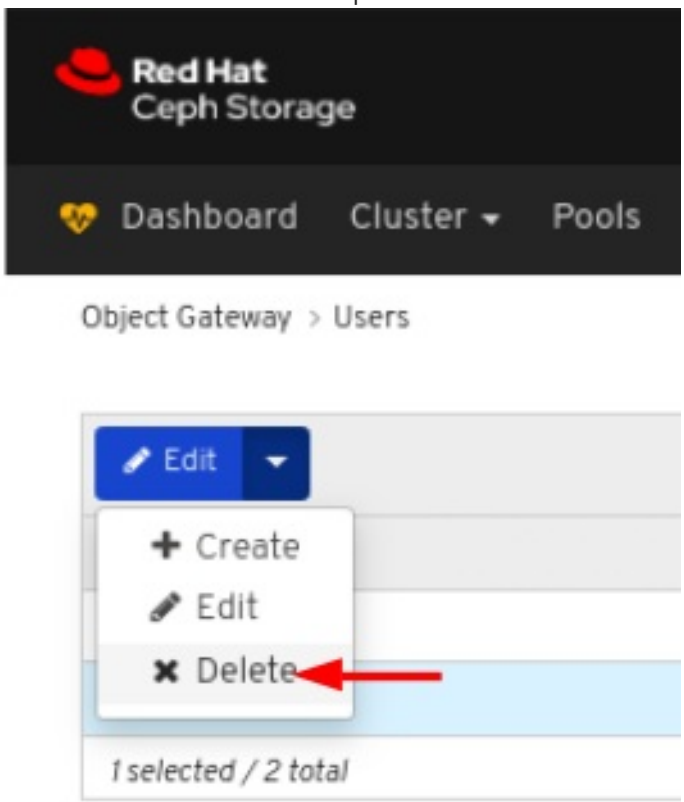

6. In the *Delete User* dialog window, Click the *Yes, I am sure* box and then Click *Delete User* to save the settings:

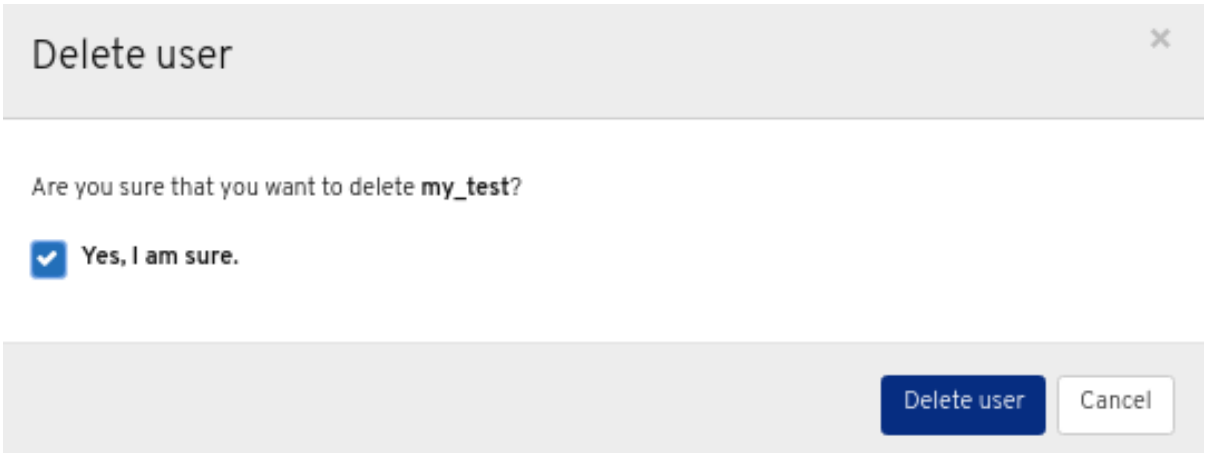

8.3.6. Creating object gateway subusers

A subuser is associated with a user of the S3 interface. The dashboard allows you to create Ceph Object Gateway subusers.

Prerequisites

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

Procedure

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

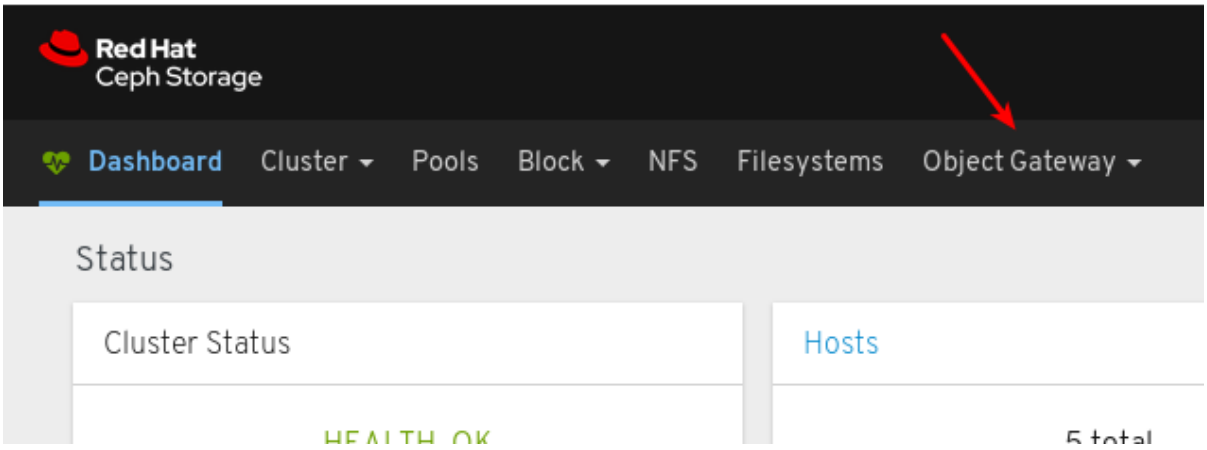

3. Click *Users* from the drop-down menu.

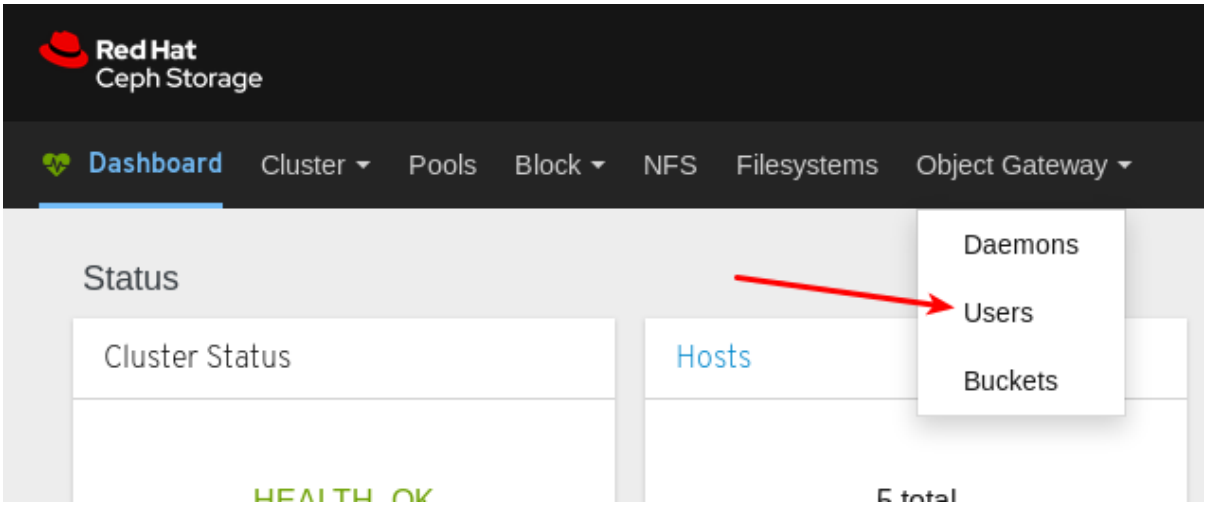

4. In the example below, you can see a user named **ceph-dashboard** in the table. Select the user by clicking its row.

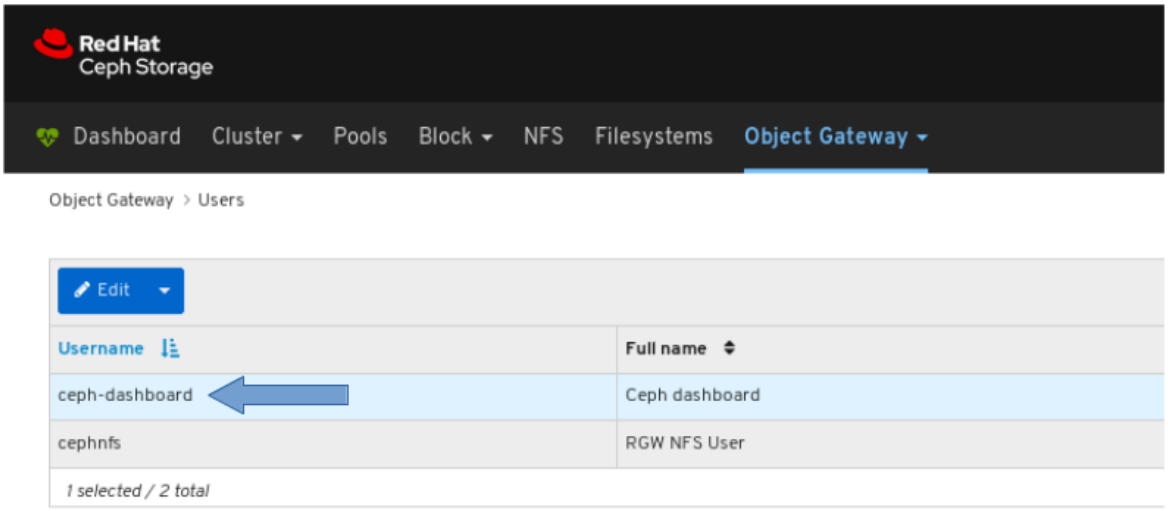

5. Click *Edit* drop-down menu, and then select *Edit*.

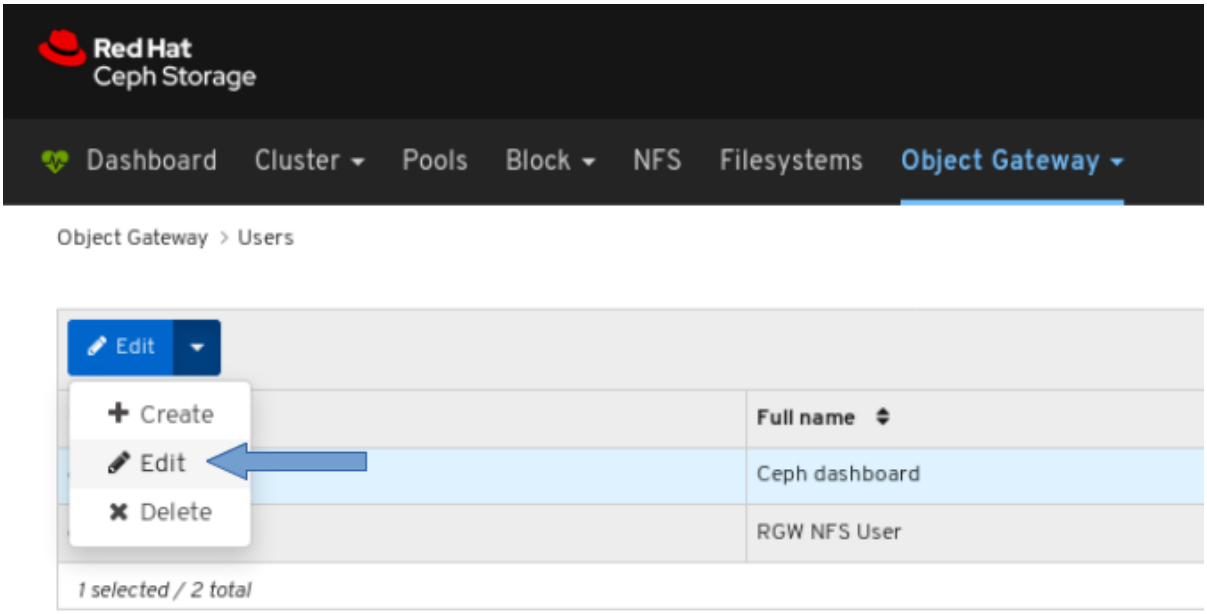

6. Click the *+CreateSubuser* button.

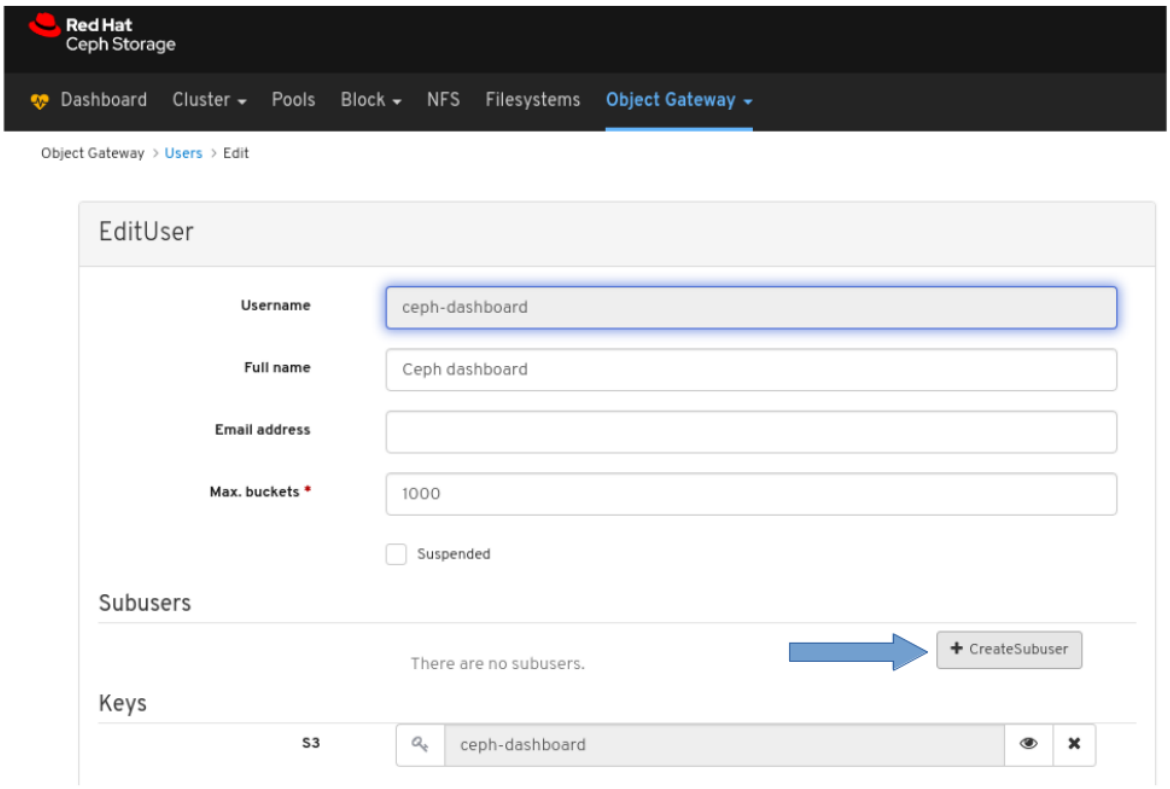

7. Enter a *Subuser* name and select the appropriate permissions.

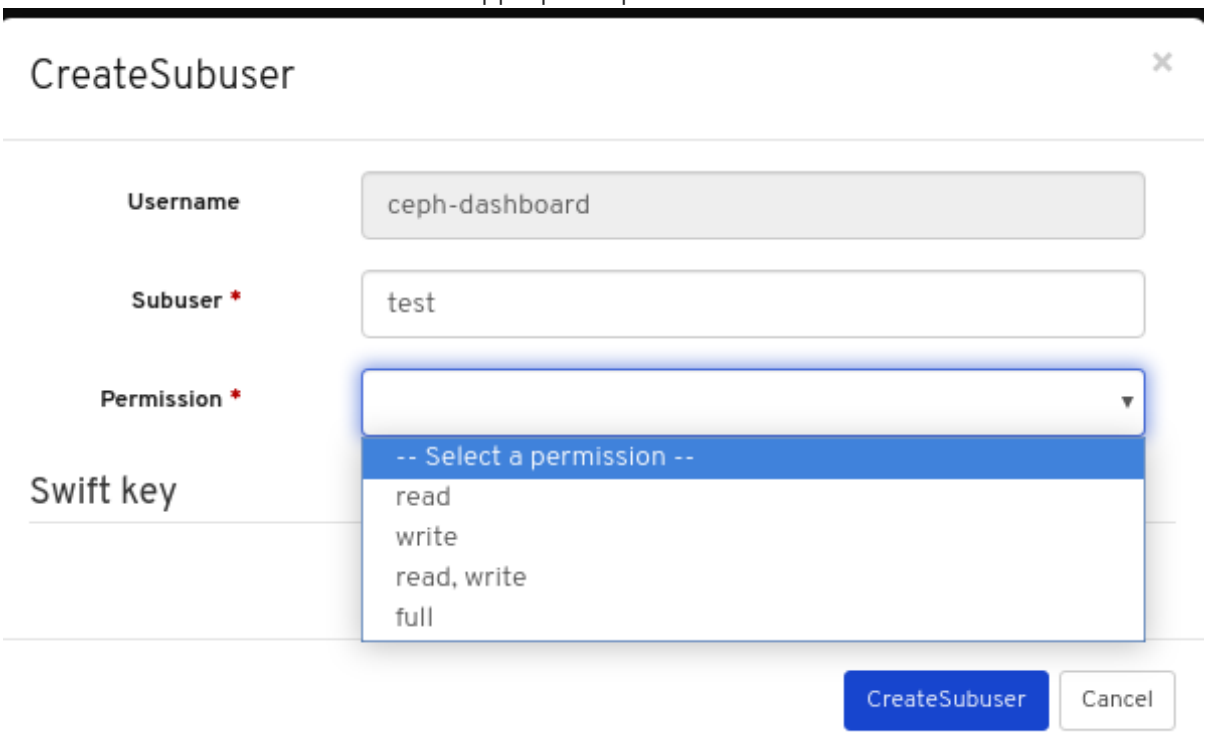

8. Click the *Auto-generate secret* box and then click the *Create Subuser* button.

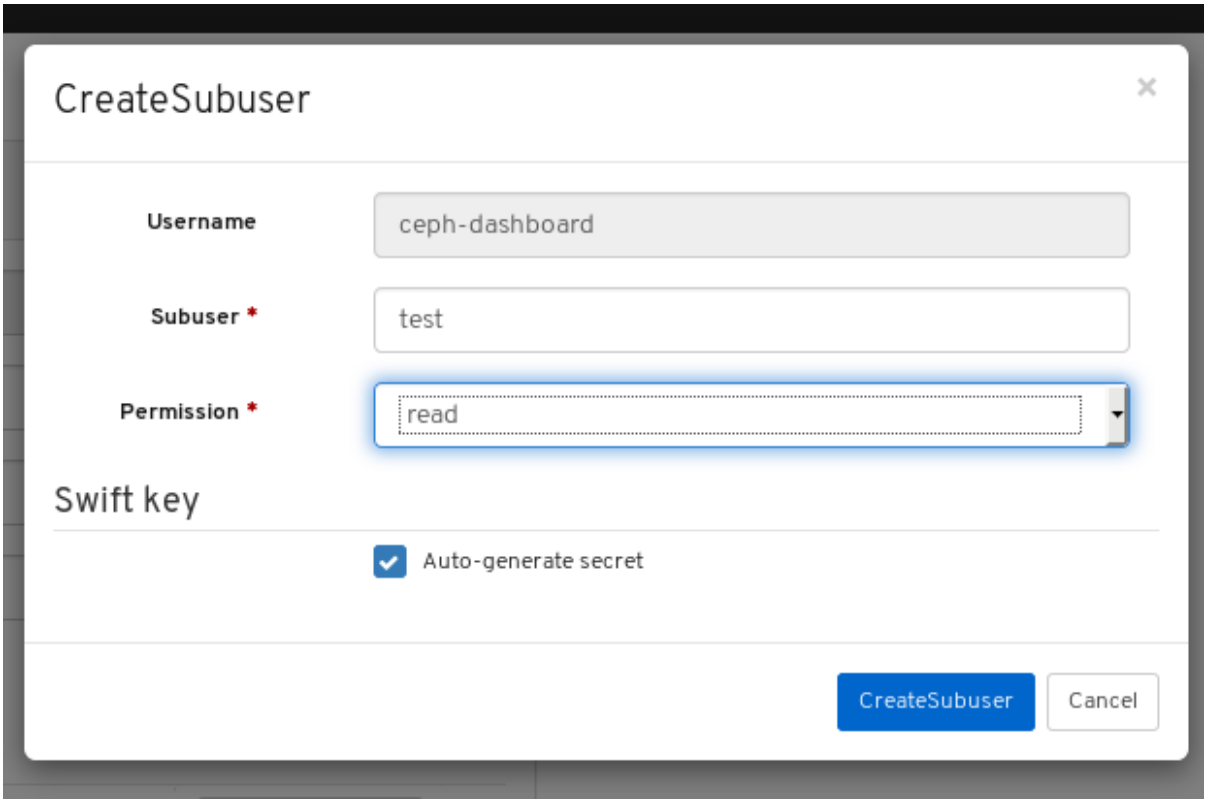

NOTE

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

9. Click the *Edit user* button in the *Edit User* window.

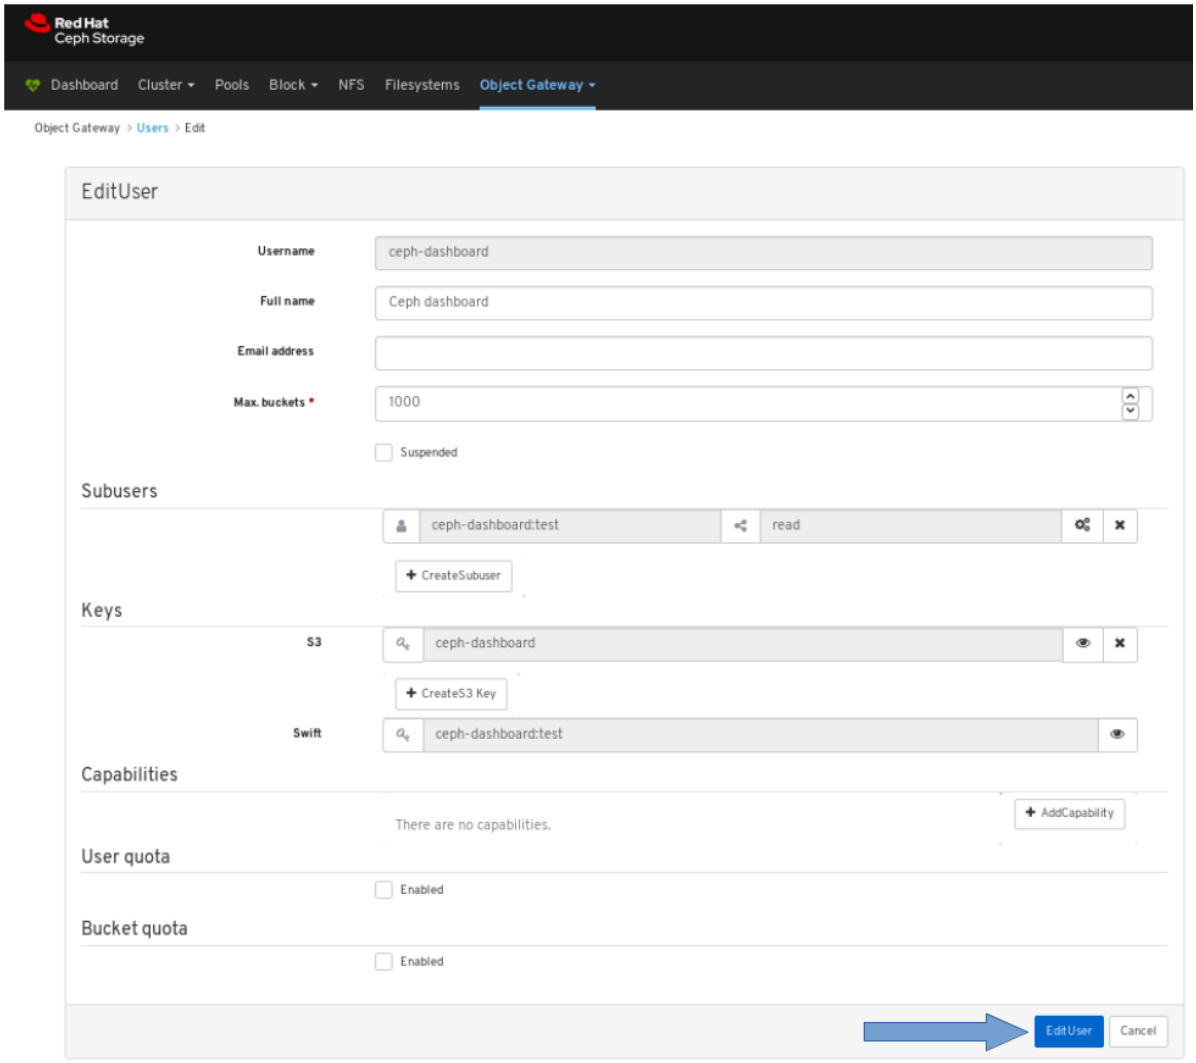

10. Verify the subuser creation was successful. A notification towards the top right corner of the page pops up indicating that the Object Gateway user was updated successfully.

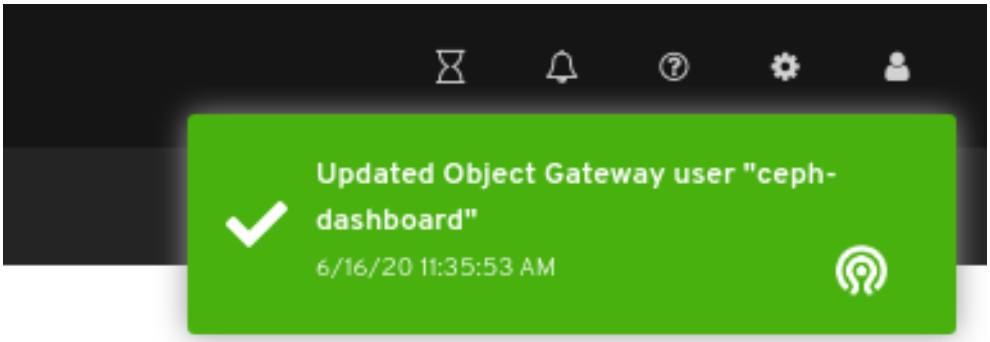

Additional Resources

- For information on how to install the Ceph Object Gateway, see the *Installing the Ceph Object Gateway* in the *Red Hat Ceph Storage Installation Guide*.
- For information on how to manually add object gateway login credentials to the dashboard, see the *Manually adding object gateway login credentials to the dashboard* in the Dashboard guide.
- For more information on the Ceph Object Gateway, see the *Object Gateway Configuration and Administration Guide*.

8.4. OBJECT GATEWAY BUCKET FUNCTIONS

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

8.4.1. Prerequisites

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

8.4.2. Viewing object gateway buckets

The 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.
- Object gateway login credentials are added to the dashboard.
- An object gateway bucket is created.

Procedure

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

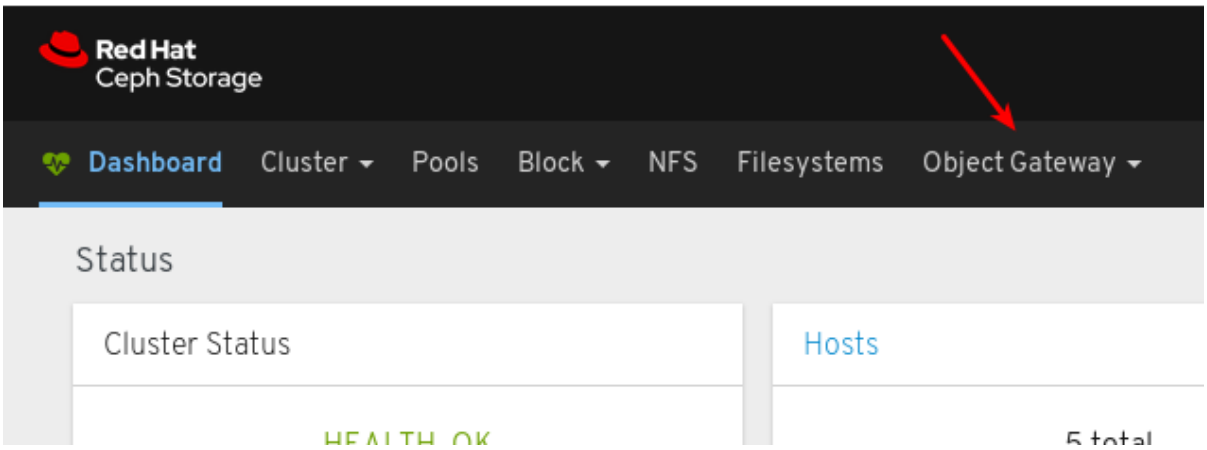

3. Click *Buckets*.

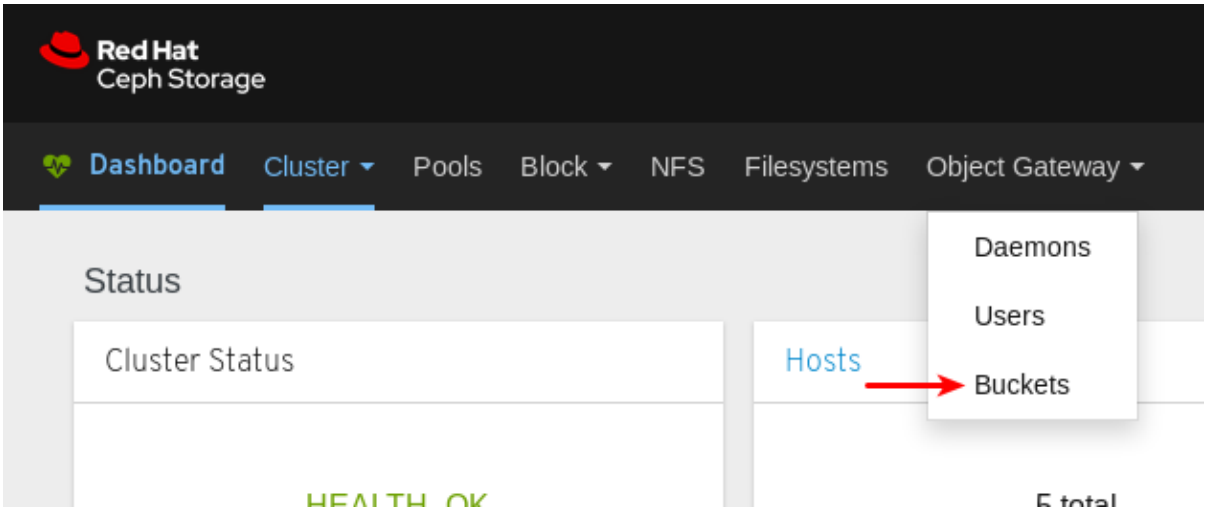

4. In the example below, you can see a bucket named **my-bucket** in the table.

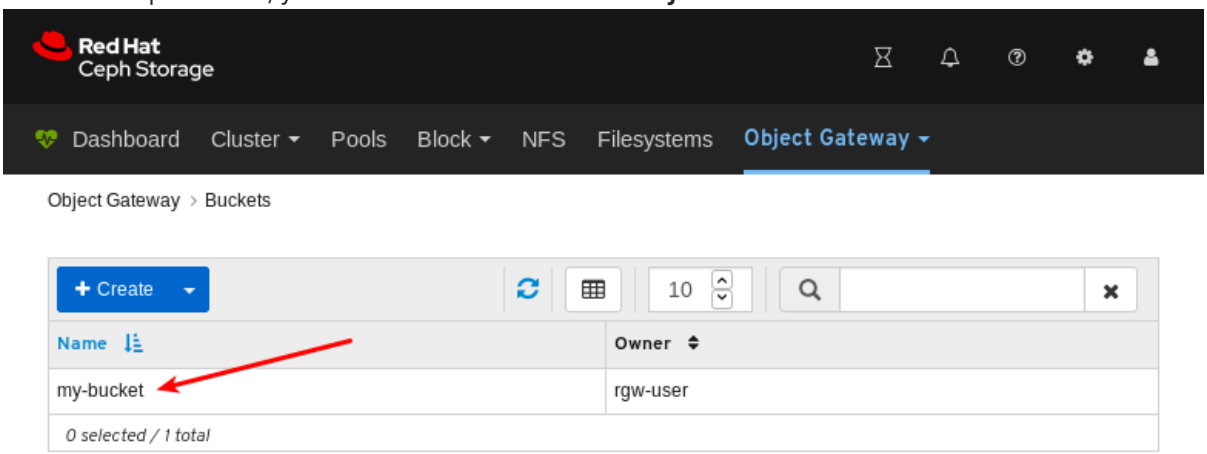

5. To view details, select the bucket by clicking the row for **my-bucket**.

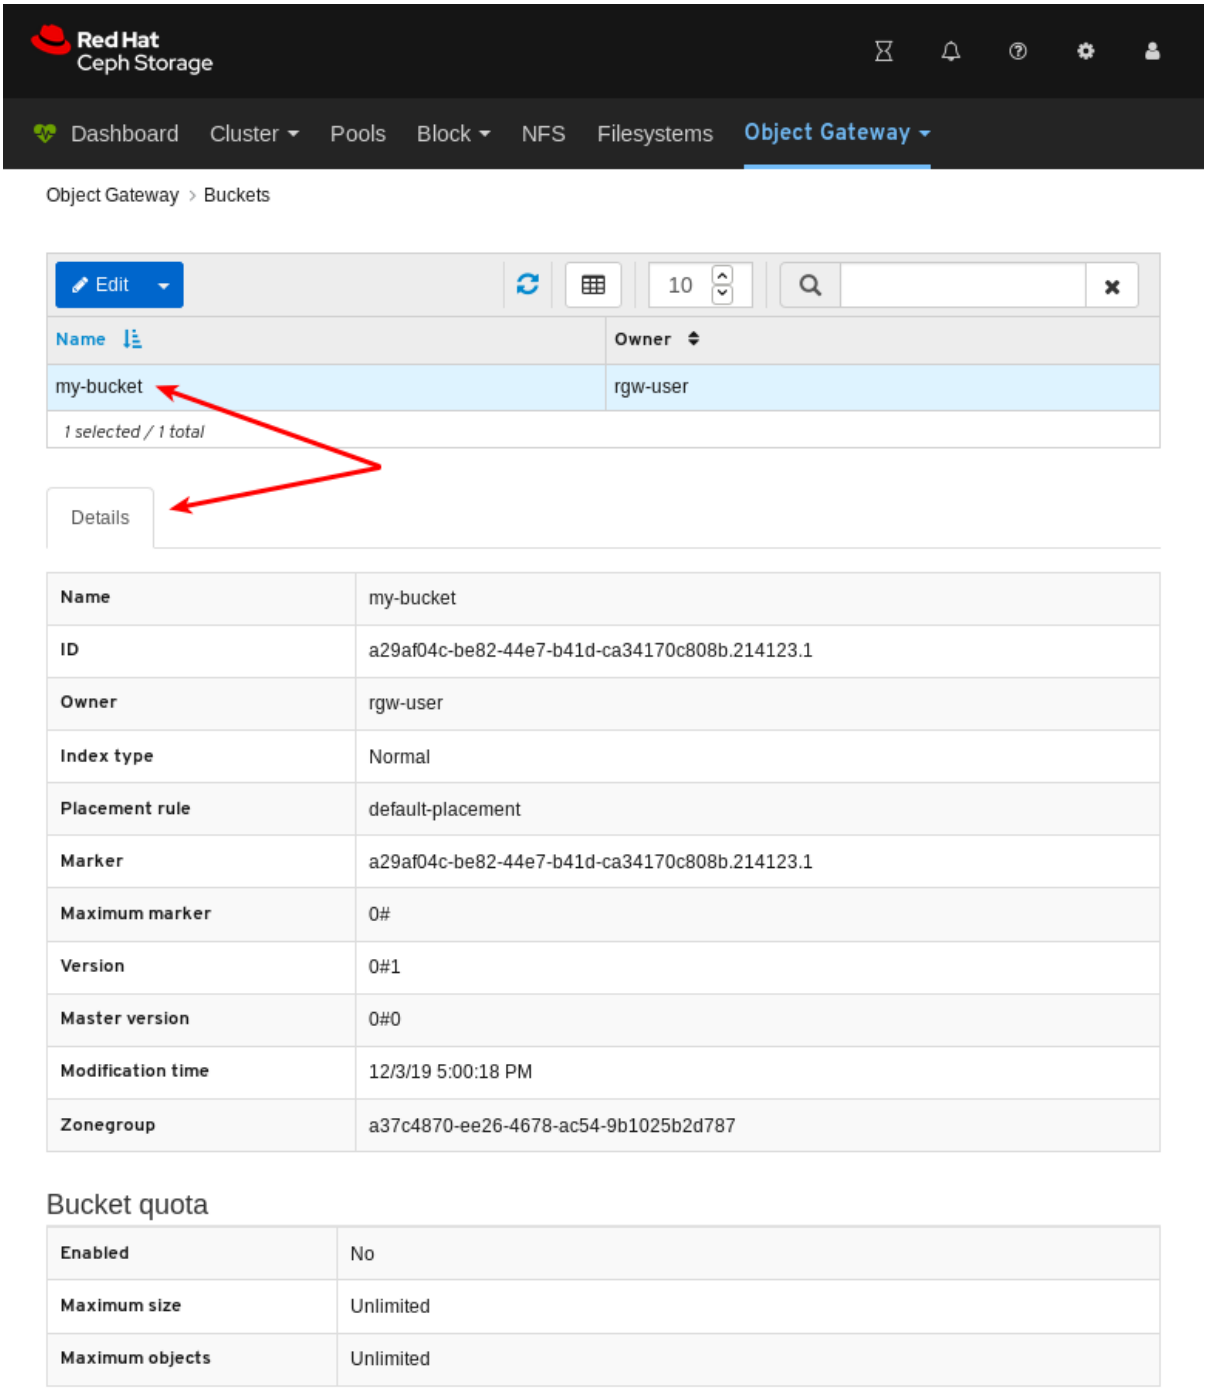

Additional Resources

- For information on how to install the Ceph Object Gateway, see Installing the Ceph Object Gateway in the Installation Guide.
- For information on how to add object gateway login credentials to the dashboard, see Adding object gateway login credentials to the dashboard in the Dashboard guide.
- For more information on the Ceph Object Gateway, see the Object Gateway Configuration and Administration Guide.

8.4.3. Creating object gateway buckets

The dashboard allows you to create Ceph Object Gateway buckets.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- An object gateway user that is not suspended is created.

Procedure

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

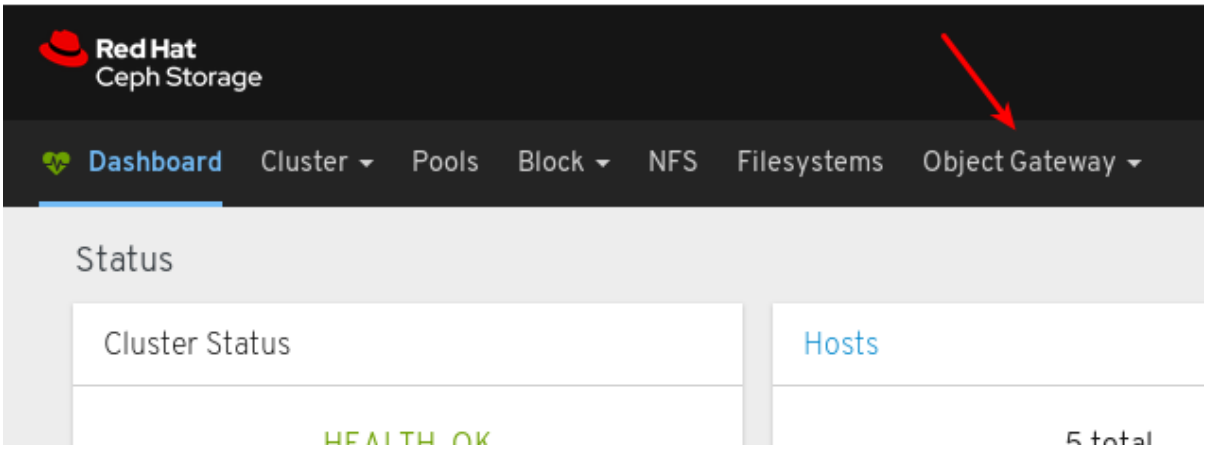

3. Click *Buckets*.

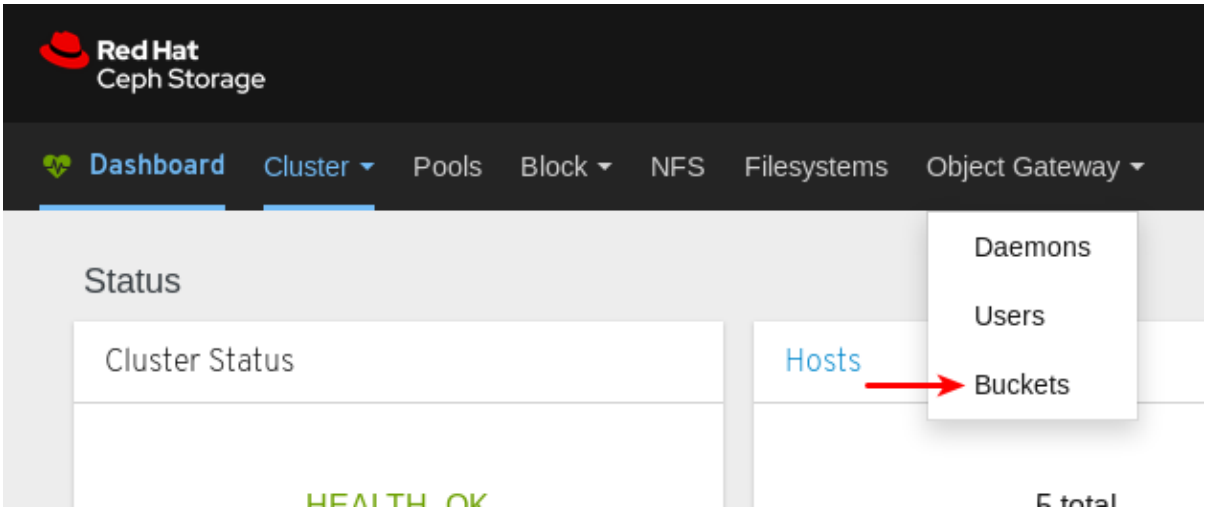

4. Click *Create*.

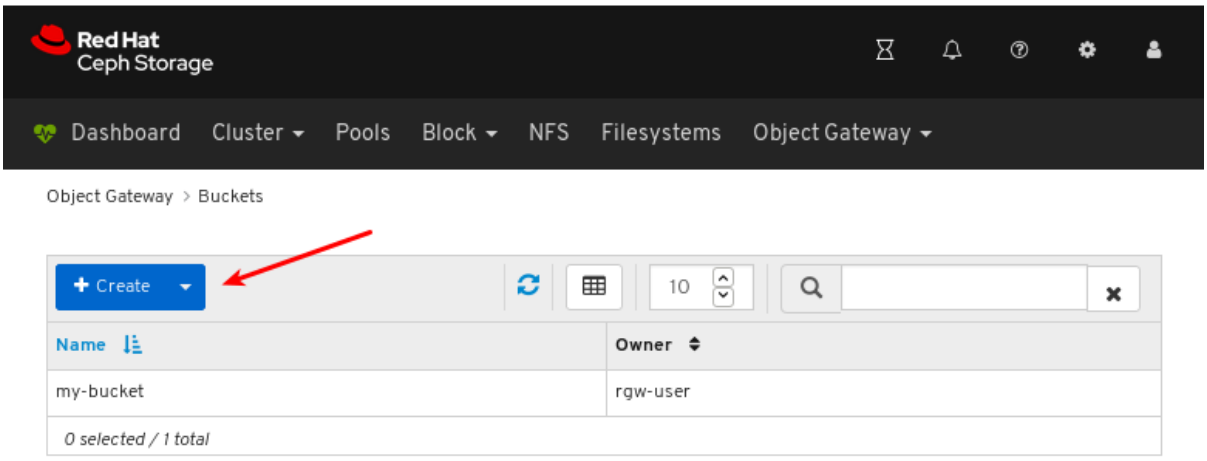

5. Enter a value for *Name* and select a user that is not suspended.

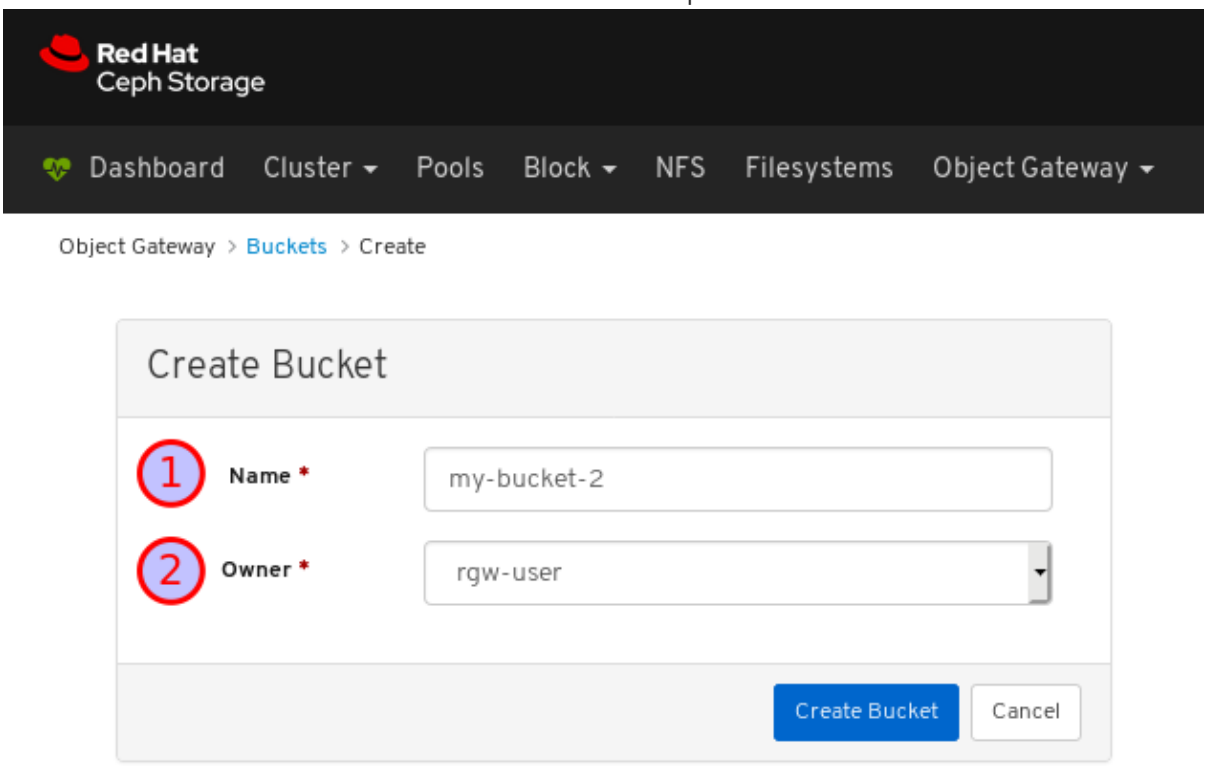

6. Click *Create bucket*.

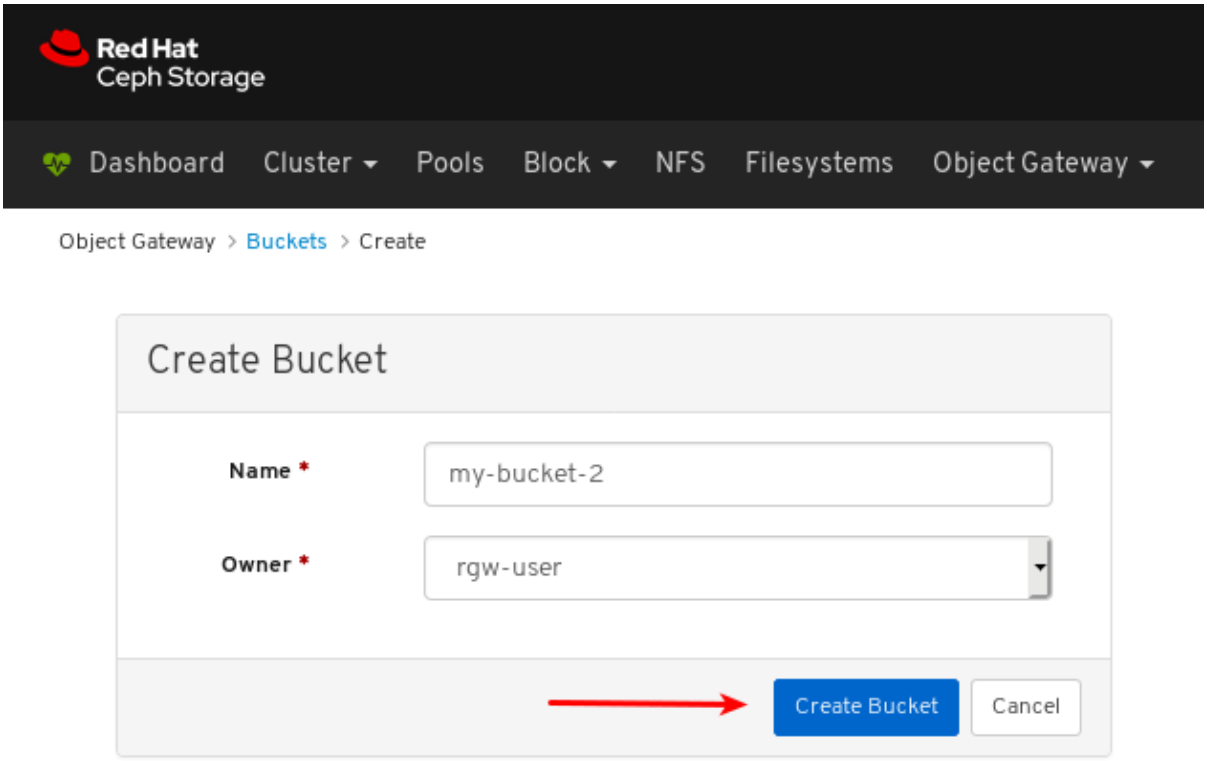

7. Verify the bucket creation was successful. A notification confirms the bucket was created and the bucket can be seen in the table of buckets.

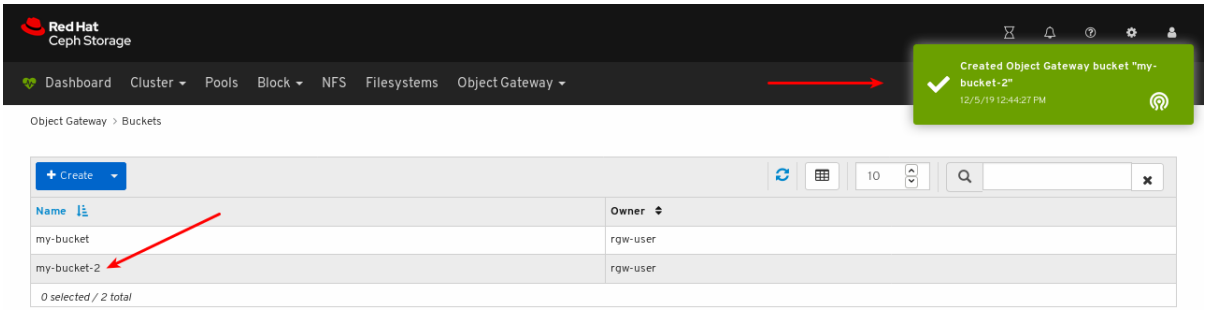

Additional Resources

- For information on how to install the Ceph Object Gateway, see Installing the Ceph Object Gateway in the Installation Guide.
- For information on how to add object gateway login credentials to the dashboard, see Adding object gateway login credentials to the dashboard in the Dashboard guide.
- For more information on the Ceph Object Gateway, see the Object Gateway Configuration and Administration Guide.

8.4.4. Editing object gateway buckets

The dashboard allows you to edit Ceph Object Gateway buckets.

Prerequisites

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

- Object gateway login credentials are added to the dashboard.
- A Ceph Object Gateway user that is not suspended is created.
- A Ceph Object Gateway bucket created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*:

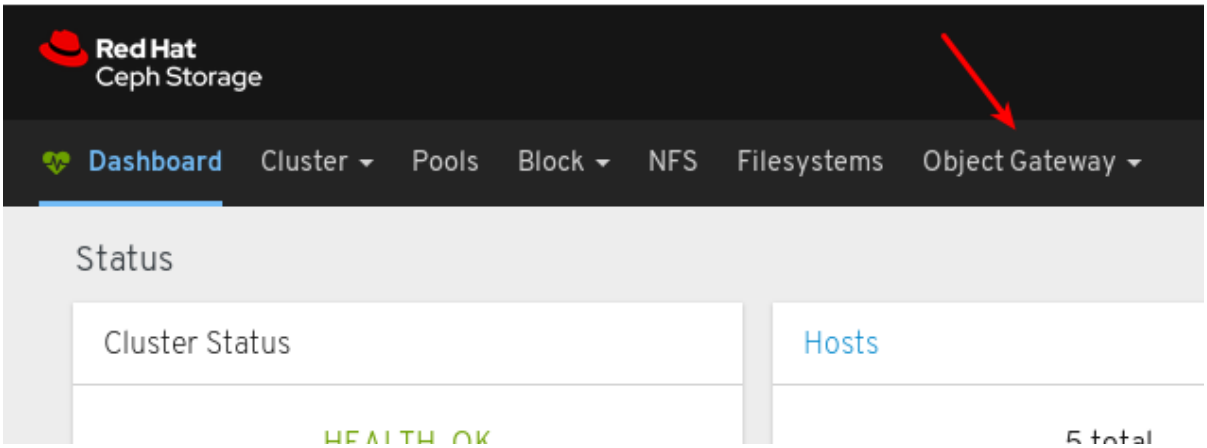

3. Click *Buckets*:

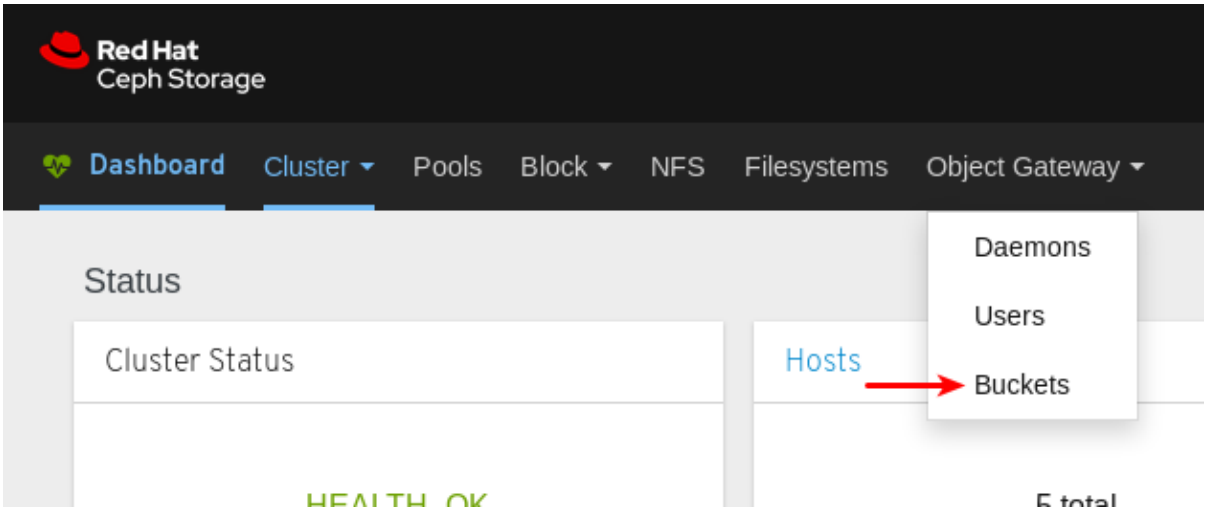

4. To edit the bucket, click its row:

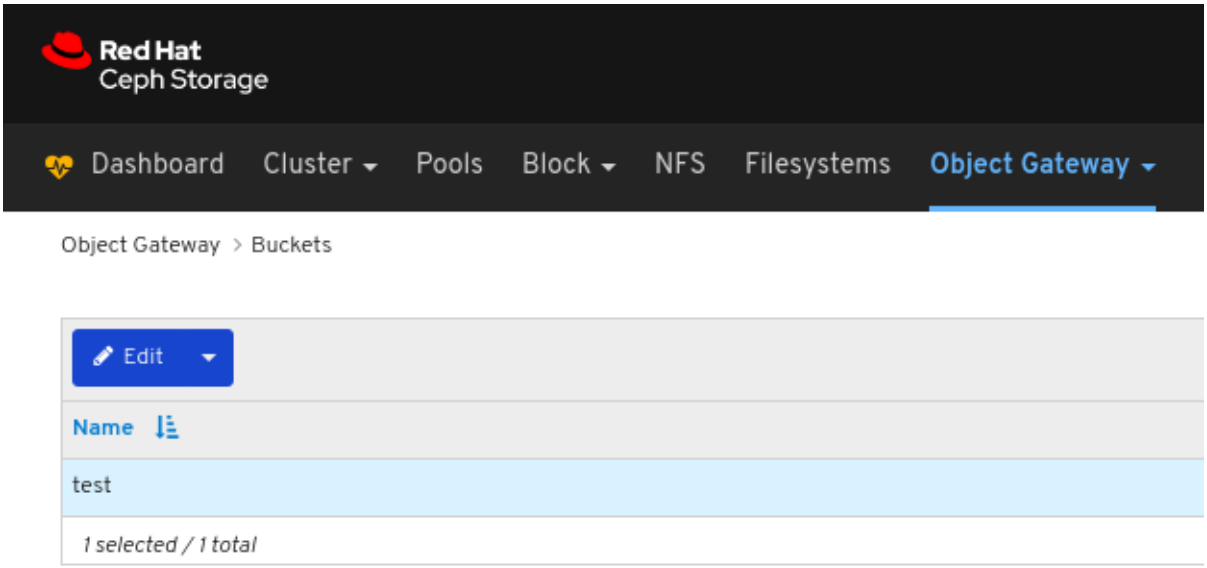

5. Select *Edit* In the *Edit* drop-down:

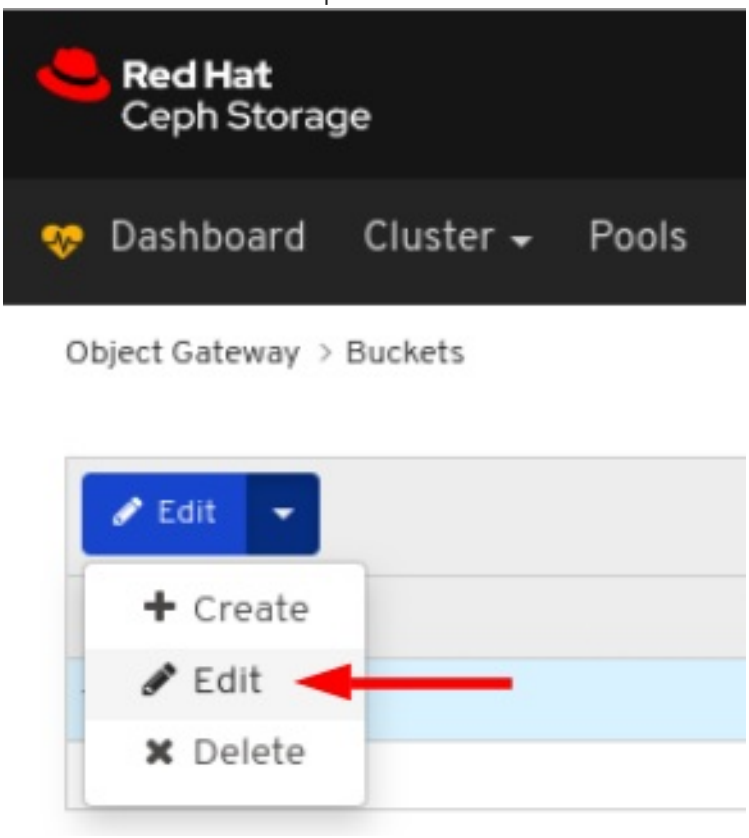

6. In the *EditBucket* window, edit the required parameters and click the *EditBucket* button:

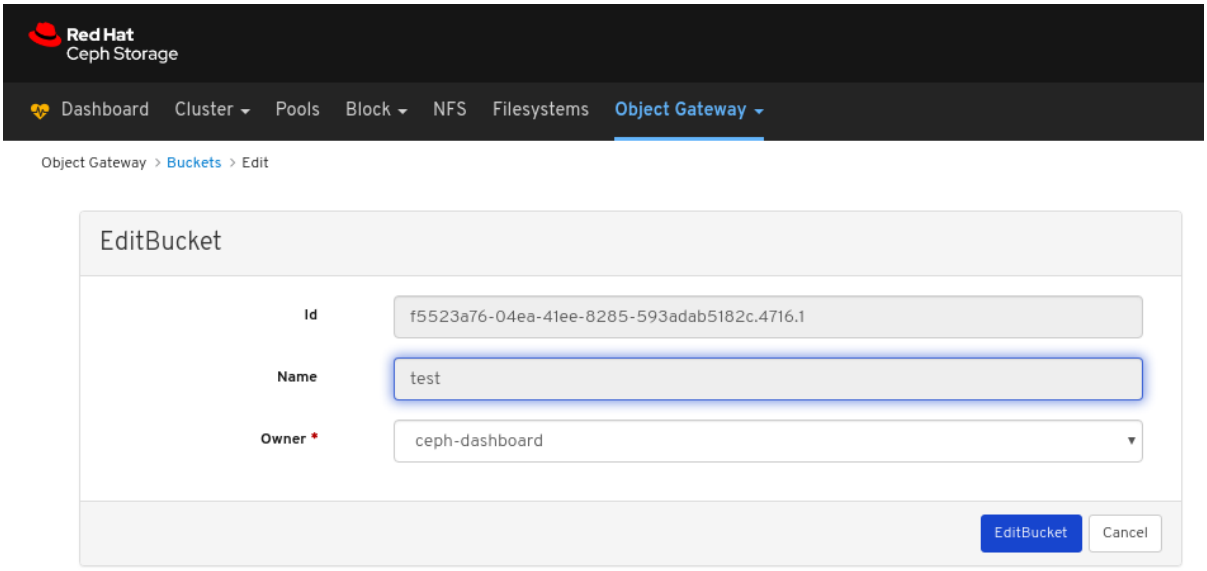

7. A notification towards the top right corner of the page indicates the bucket was updated successfully.

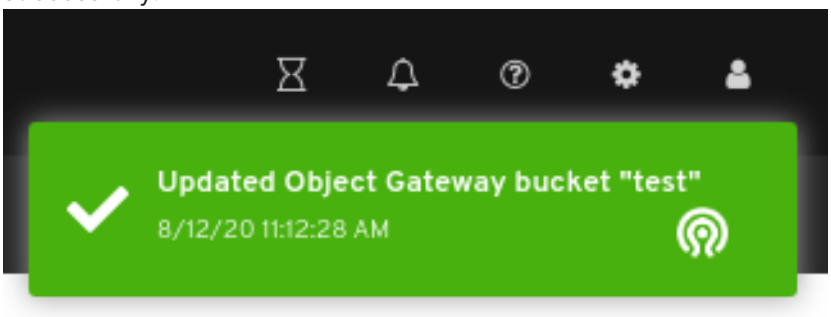

Additional Resources

- See the *Installing the Ceph Object Gateway* section in the *Red Hat Ceph Storage Installation Guide* for more information.
- See the *Adding object gateway login credentials to the dashboard* section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

8.4.5. Deleting object gateway buckets

The dashboard allows you to delete Ceph Object Gateway buckets.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph Object Gateway is installed.
- Object gateway login credentials are added to the dashboard.
- A Ceph Object Gateway user that is not suspended is created.
- A Ceph Object Gateway bucket created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Object Gateway*:

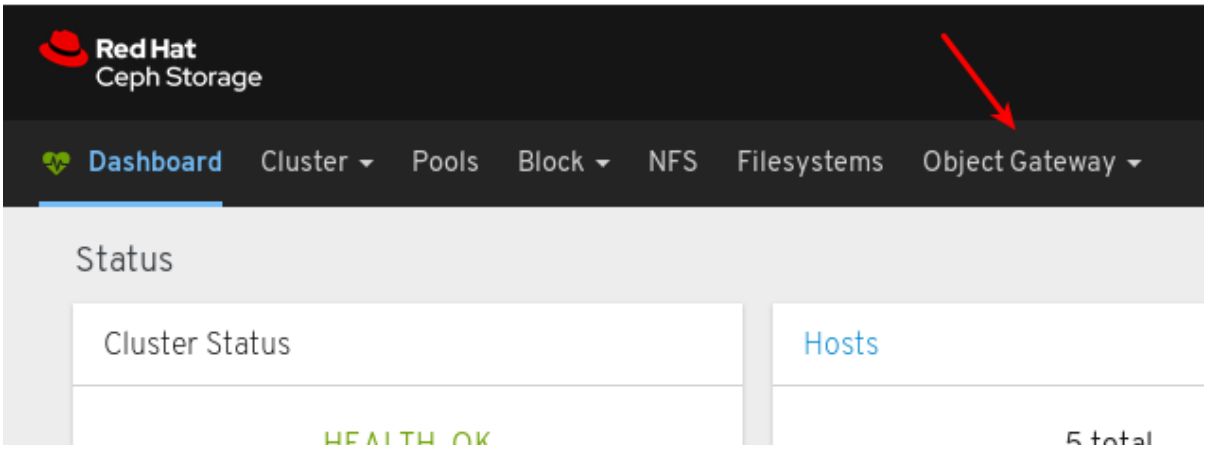

3. Click *Buckets*:

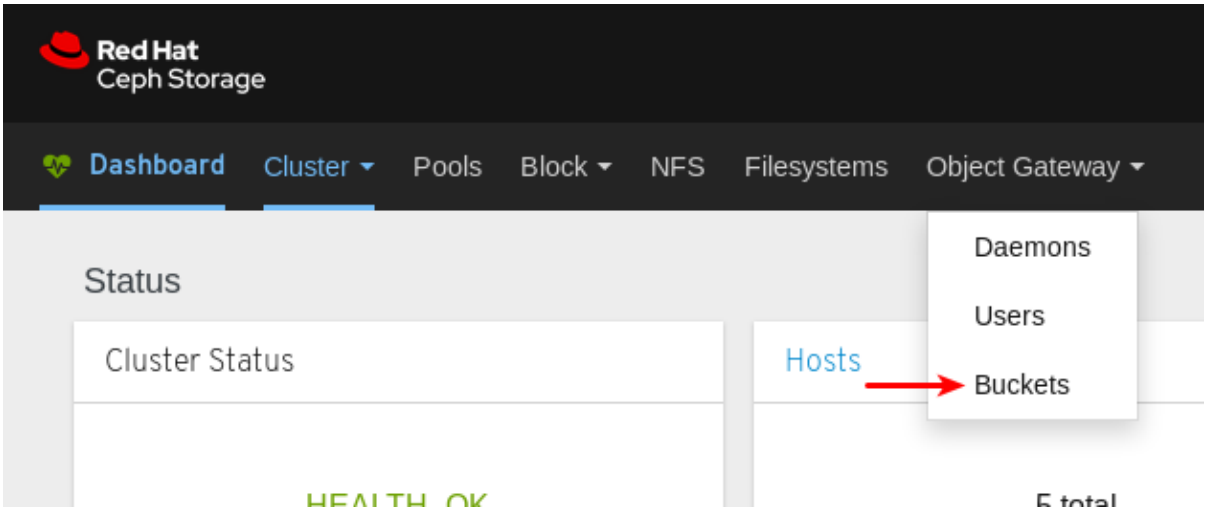

4. To delete the bucket, click its row:

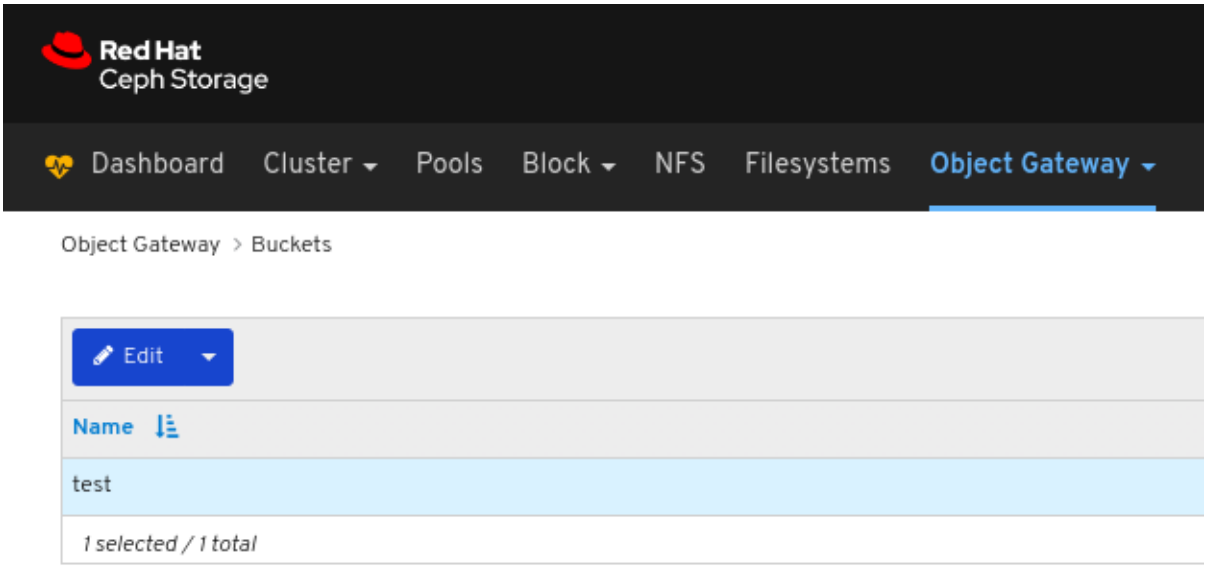

5. Select *Delete* In the *Edit* drop-down:

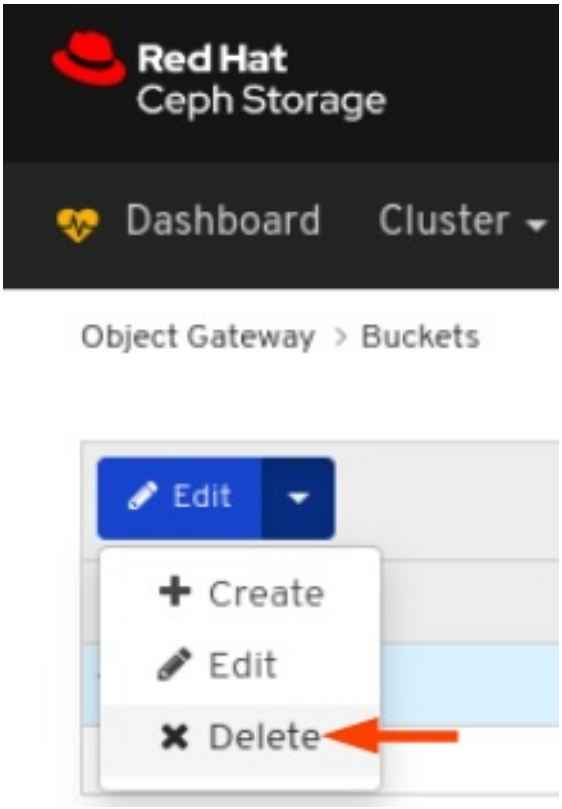

6. In the *Delete Bucket* dialog window, Click the *Yes, I am sure* box and then Click *Delete bucket* to save the settings:

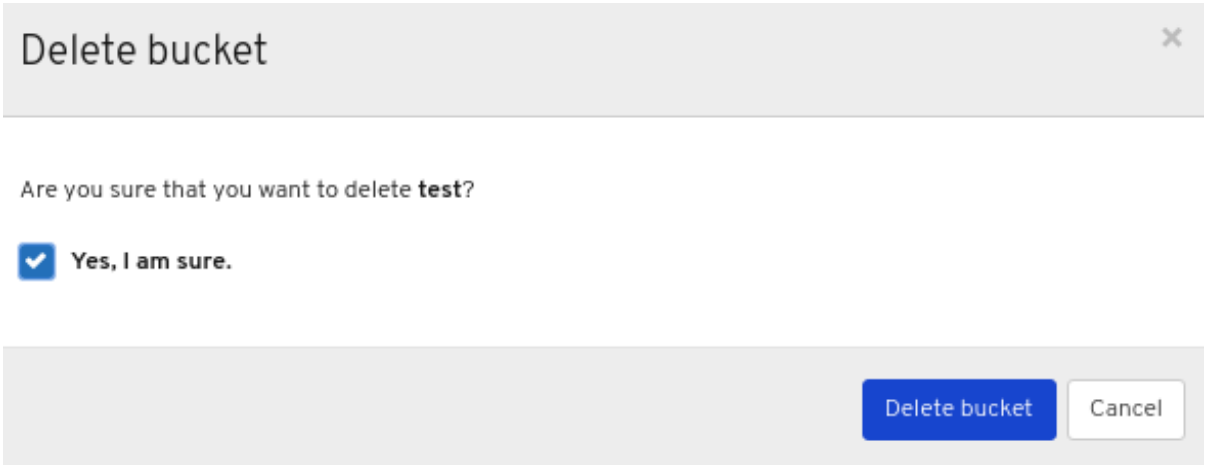

8.5. MANUALLY ADDING OBJECT GATEWAY LOGIN CREDENTIALS TO THE DASHBOARD

The Red Hat Ceph Storage Dashboard can manage the Ceph Object Gateway, also known as the RADOS Gateway, or RGW. To manage the Ceph Object Gateway, the dashboard must connect to it using login credentials of an RGW user with the **system** flag. When the Object Gateway is installed using **ceph-ansible**, it automatically adds the login credentials to the dashboard. It is also possible to set the login credentials manually.

Prerequisites

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

Procedure

1. Obtain the **access_key** and **secret_key** of an RGW user with the **system** flag enabled:

   - If you do not have an RGW user with the **system** flag enabled, create one.

     ```
     radosgw-admin user create --uid=USER_ID --display-name=DISPLAY_NAME --system
     ```

     Example:

     ```
     [root@mon ~]# radosgw-admin user create --uid=rgw-user --display-name=RGW-user --
     system
     {
        "user_id": "rgw-user",
        "display_name": "RGW-user",
        "email": "",
        "suspended": 0,
        "max_buckets": 1000,
        "subusers": [],
        "keys": [
           {
              "user": "rgw-user",
              "access_key": "BYC5SWQQH24A2BFHS2RC",
              "secret_key": "159d94uHK9ADiWZrGsNYWYjRXCDrhL2xVi8PO6kT"
           }
        ],
        "swift_keys": [],
        "caps": [],
        "op_mask": "read, write, delete",
        "system": "true",
        "default_placement": "",
        "default_storage_class": "",
        "placement_tags": [],
        "bucket_quota": {
           "enabled": false,
           "check_on_raw": false,
           "max_size": -1,
           "max_size_kb": 0,
           "max_objects": -1
        },
        "user_quota": {
           "enabled": false,
           "check_on_raw": false,
           "max_size": -1,
           "max_size_kb": 0,
           "max_objects": -1
        },
        "temp_url_keys": [],
        "type": "rgw",
        "mfa_ids": []
     }
     ```

     Take note of the values for **access_key** and **secret_key**. In the example above, **access_key** is **BYC5SWQQH24A2BFHS2RC** and **secret_key** is **159d94uHK9ADiWZrGsNYWYjRXCDrhL2xVi8PO6kT**.

- If an RGW user with the **system** flag enabled is already created, obtain the credentials using the **user info** command of the **radosgw-admin** utility.

```
radosgw-admin user info --uid=USER_ID
```

Example:

```
[root@mon ~]# radosgw-admin user info --uid=rgw-user
{
    "user_id": "rgw-user",
    "display_name": "RGW-user",
    "email": "",
    "suspended": 0,
    "max_buckets": 1000,
    "subusers": [],
    "keys": [
        {
            "user": "rgw-user",
            "access_key": "BYC5SWQQH24A2BFHS2RC",
            "secret_key": "159d94uHK9ADiWZrGsNYWYjRXCDrhL2xVi8PO6kT"
        }
    ],
    "swift_keys": [],
    "caps": [],
    "op_mask": "read, write, delete",
    "system": "true",
    "default_placement": "",
    "default_storage_class": "",
    "placement_tags": [],
    "bucket_quota": {
        "enabled": false,
        "check_on_raw": false,
        "max_size": -1,
        "max_size_kb": 0,
        "max_objects": -1
    },
    "user_quota": {
        "enabled": false,
        "check_on_raw": false,
        "max_size": -1,
        "max_size_kb": 0,
        "max_objects": -1
    },
    "temp_url_keys": [],
    "type": "rgw",
    "mfa_ids": []
}
```

Take note of the values for **access_key** and **secret_key**. In the example above, **access_key** is **BYC5SWQQH24A2BFHS2RC** and **secret_key** is **159d94uHK9ADiWZrGsNYWYjRXCDrhL2xVi8PO6kT**.

2. Provide the **access_key** and **secret_key** credentials to the dashboard:

   a. Provide the **access_key** to the dashboard.

```
ceph dashboard set-rgw-api-access-key ACCESS_KEY
```

Example:

```
[root@mon ~]# ceph dashboard set-rgw-api-access-key BYC5SWQQH24A2BFHS2RC
Option RGW_API_ACCESS_KEY updated
```

b. Provide the **secret_key** to the dashboard.

```
ceph dashboard set-rgw-api-secret-key SECRET_KEY
```

Example:

```
[root@mon ~]# ceph dashboard set-rgw-api-secret-key
159d94uHK9ADiWZrGsNYWYjRXCDrhL2xVi8PO6kT
Option RGW_API_SECRET_KEY updated
```

3. Provide the host name and port of the object gateway to the dashboard:

a. Provide the host name to the dashboard.

```
ceph dashboard set-rgw-api-host HOST_NAME
```

Example:

```
[root@mon ~]# ceph dashboard set-rgw-api-host 192.168.122.193
Option RGW_API_HOST updated
```

b. Provide the port to the dashboard.

```
ceph dashboard set-rgw-api-port PORT
```

Example:

```
[root@mon ~]# ceph dashboard set-rgw-api-port 8080
Option RGW_API_PORT 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.

```
ceph dashboard set-rgw-api-ssl-verify false
```

Example:

```
[root@mon ~]# ceph dashboard set-rgw-api-ssl-verify false
Option RGW_API_SSL_VERIFY updated
```

Additional Resources

- For information on how to install the Ceph Object Gateway, see Installing the Ceph Object Gateway in the Installation Guide.
- For more information on the Ceph Object Gateway, see the Object Gateway Configuration and Administration Guide.

CHAPTER 9. BLOCK DEVICES

9.1. BLOCK DEVICES INTRODUCTION

The block device functions of the dashboard allow you to manage and monitor block device images. The functionality is divided between generic image functions, mirroring functions, iSCSI functions, and Quality of Service configuration. For example, you can create new images, view the state of images mirrored across clusters, manage or monitor iSCSI targets, and set IOPS limits on an image.

9.2. IMAGES FUNCTIONS

The dashboard provides several functions related to managing and monitoring images.

9.2.1. Prerequisites

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

9.2.2. Creating images

The dashboard allows you to create images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click the *Block*:

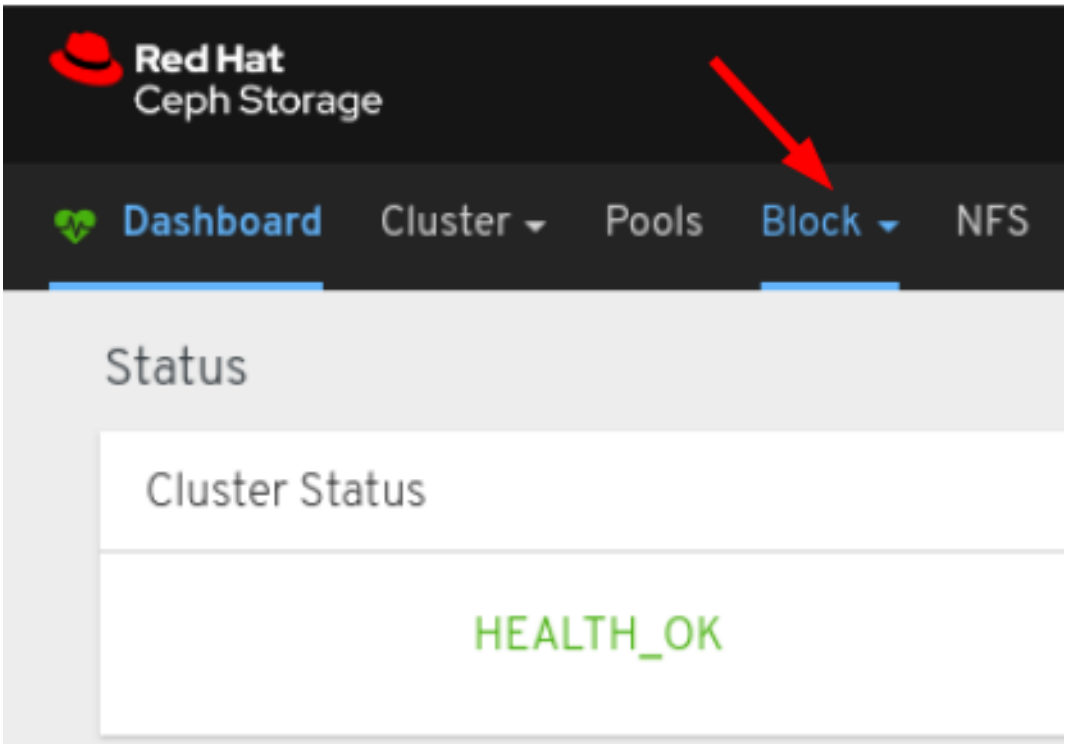

3. Select *Images* from the drop-down:

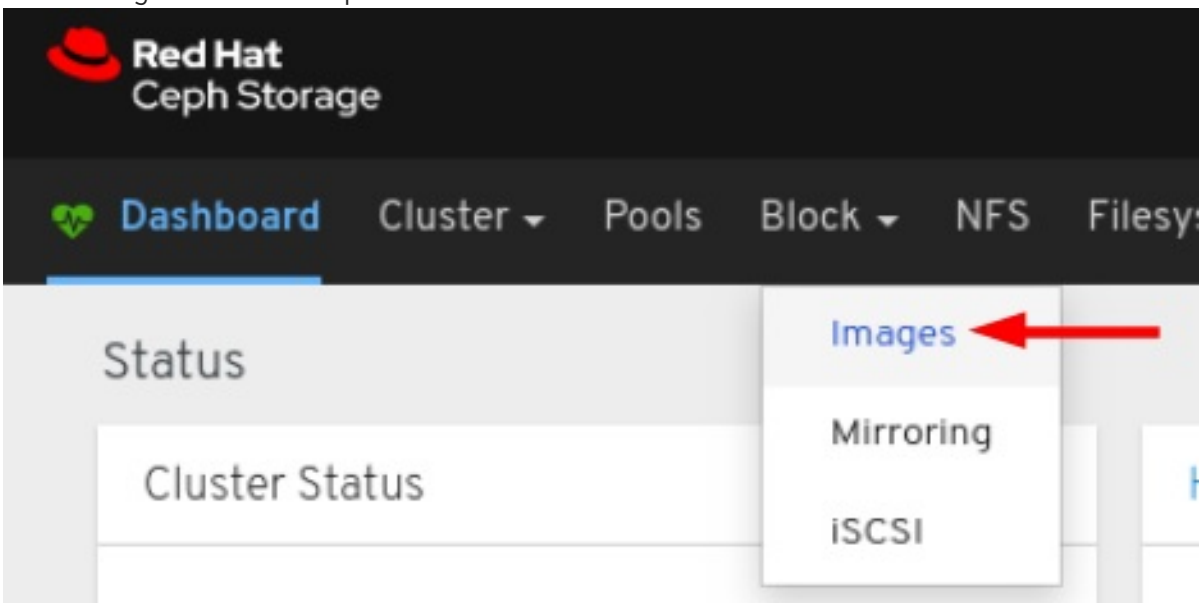

4. In the *Images* tab, Click the *Create* button:

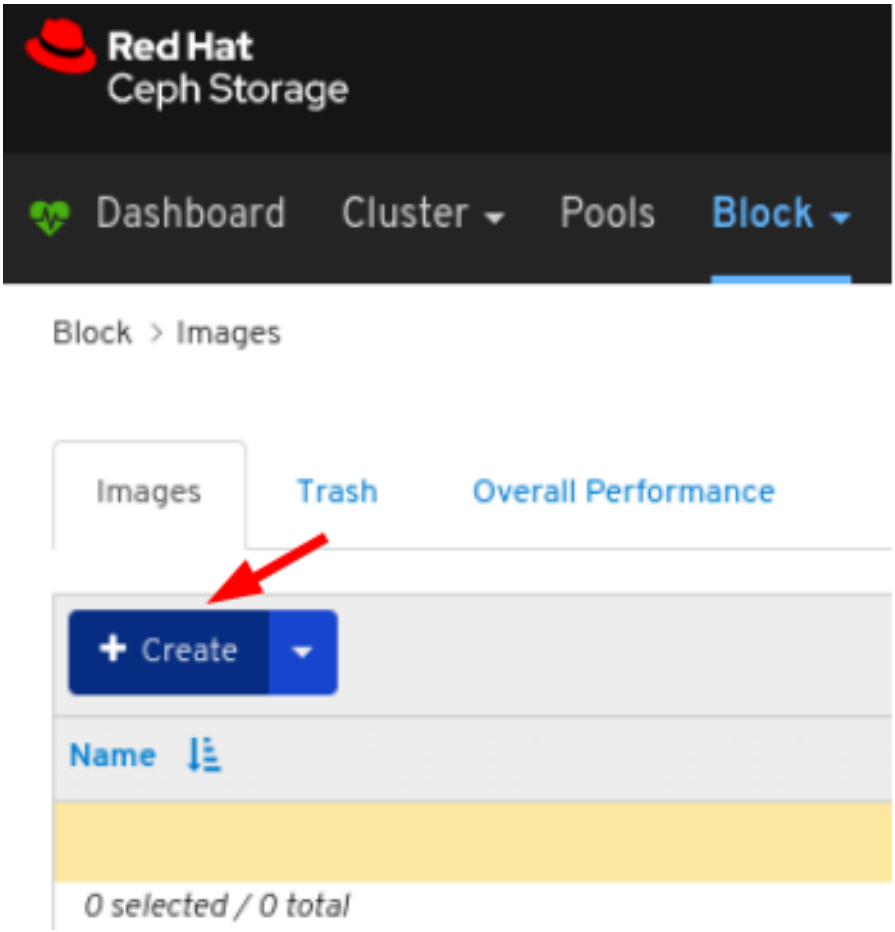

5. In the *CreateRBD* window, enter the parameters:

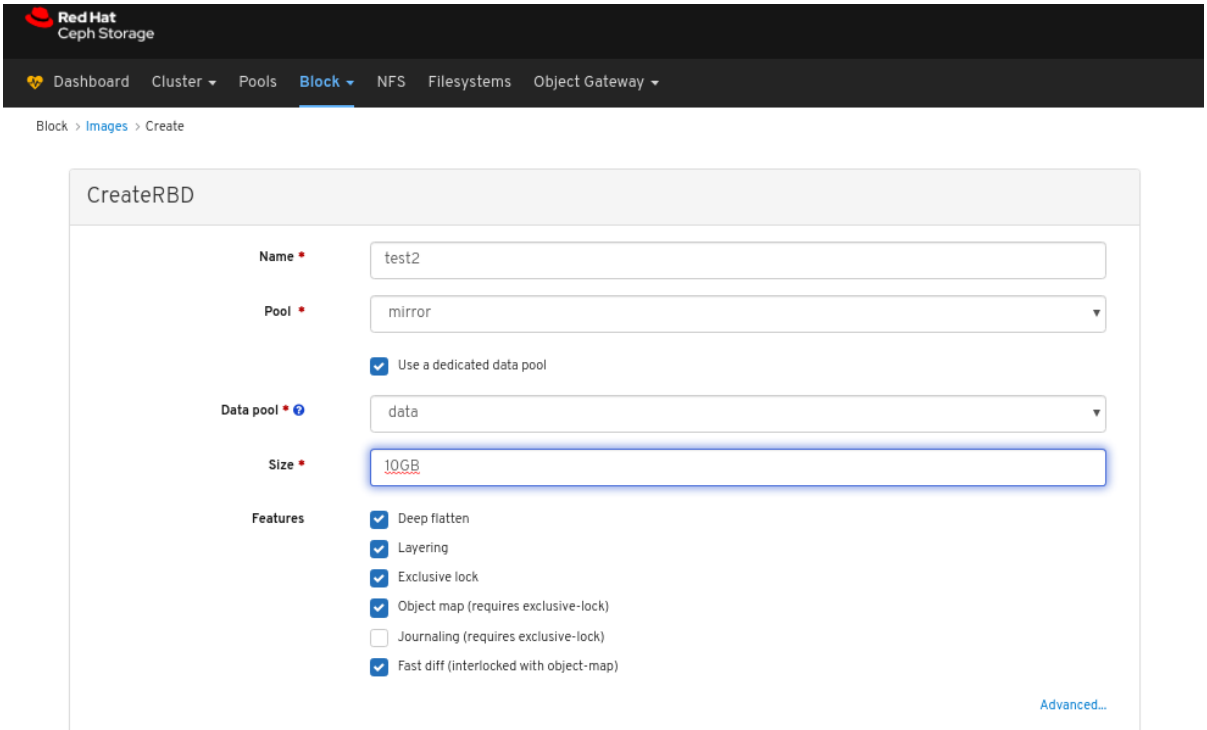

6. Optional: Click *Advanced* and set the parameters:

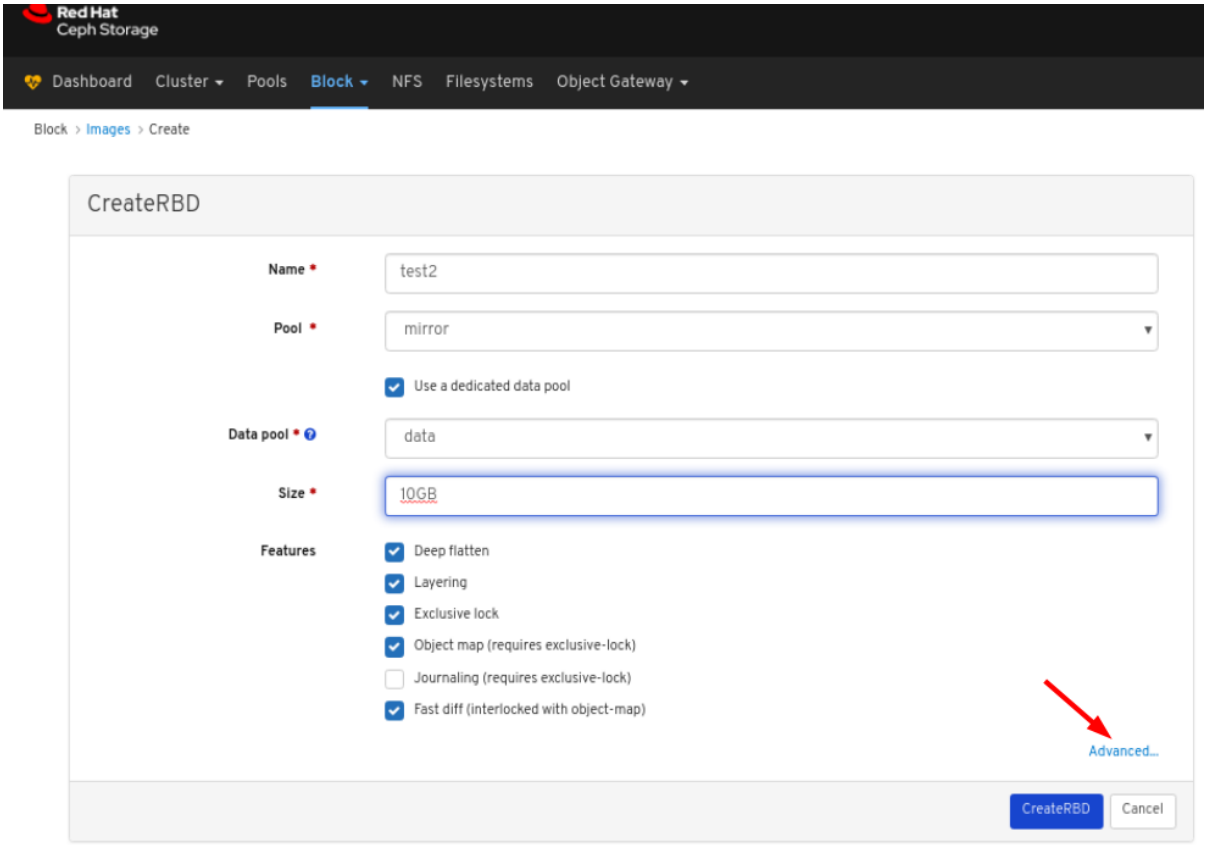

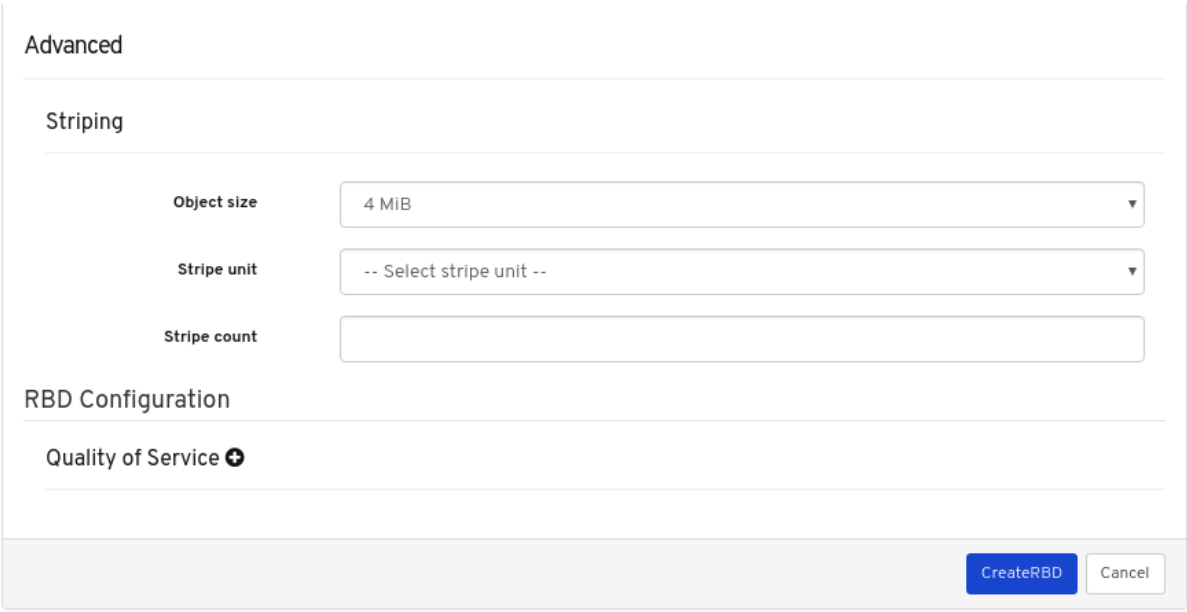

7. Click the *CreateRBD* button:

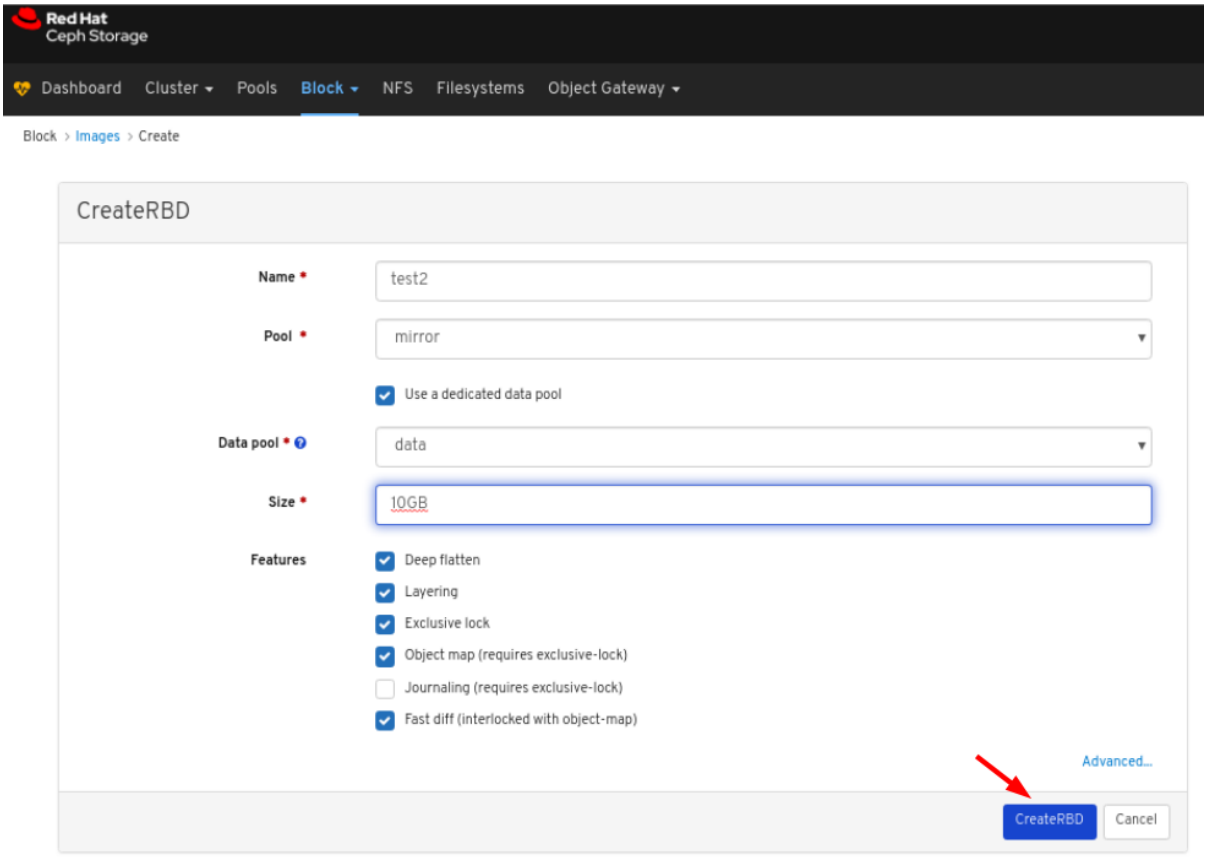

8. Notifications towards the top right corner of the page indicate the image was created successfully.

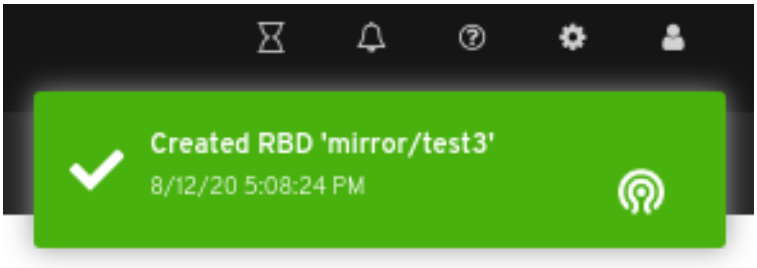

Additional Resources

- See the *Red Hat Ceph Storage Block Device Guide* for more information on Images.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.

9.2.3. Viewing images

The dashboard allows you to view images.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- An image is in the cluster.

Procedure

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.
3. Click *Images*:

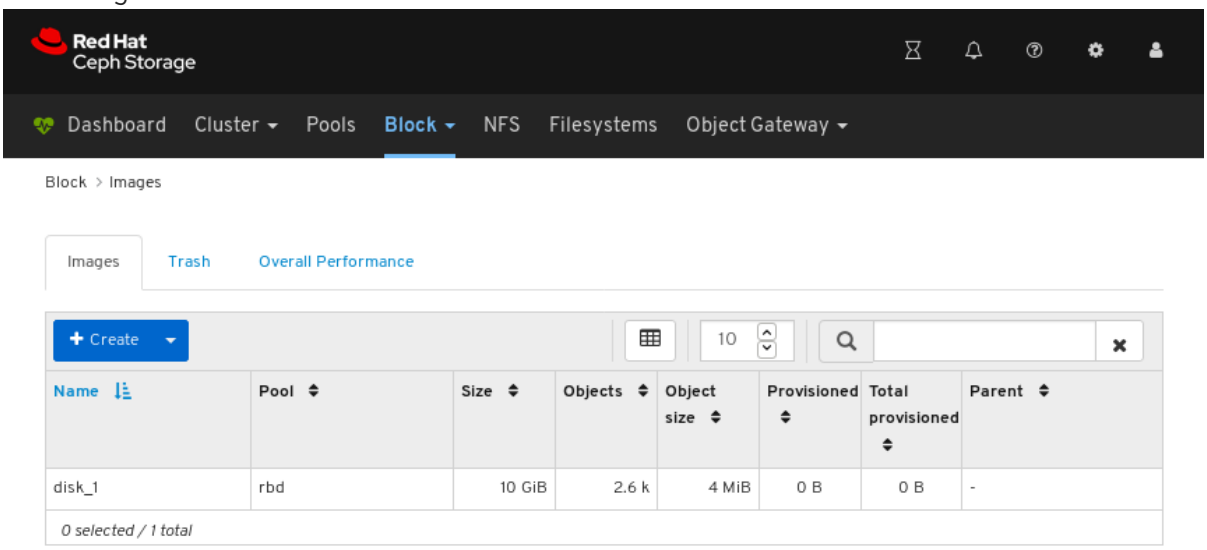

In the above example, you can see a 10 GiB image named *disk_1*.

4. To view details, select the image by clicking the row for *disk_1*:

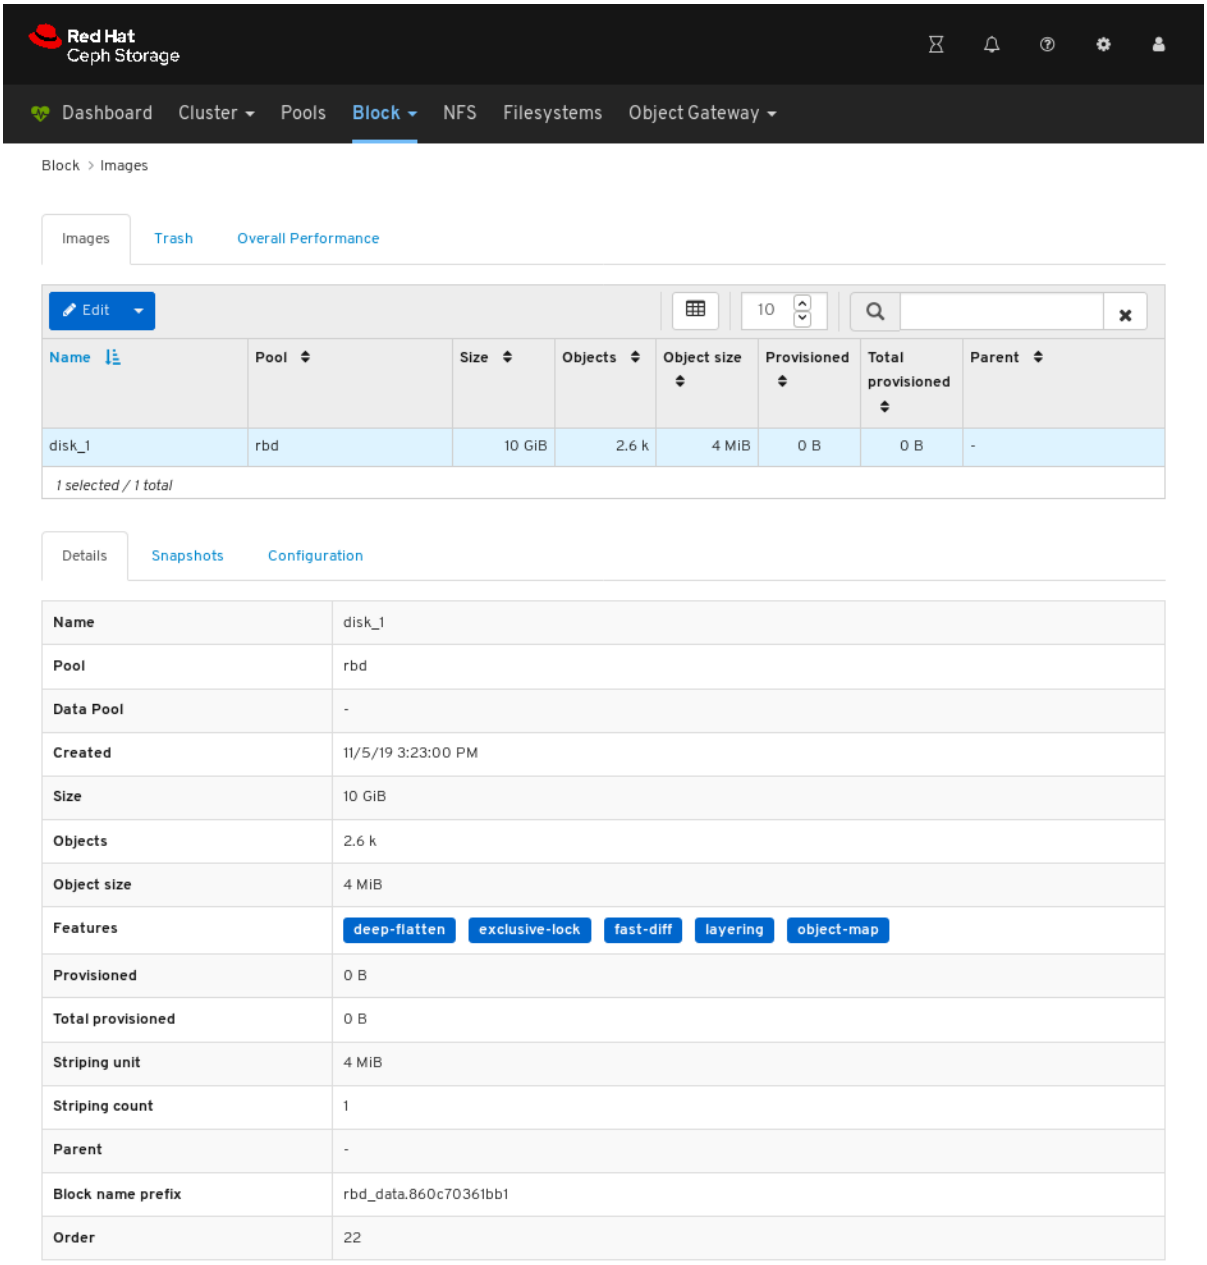

Additional Resources

- For more information on images, see the Block Device Guide.

9.2.4. Editing images

The dashboard allows you to edit images

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

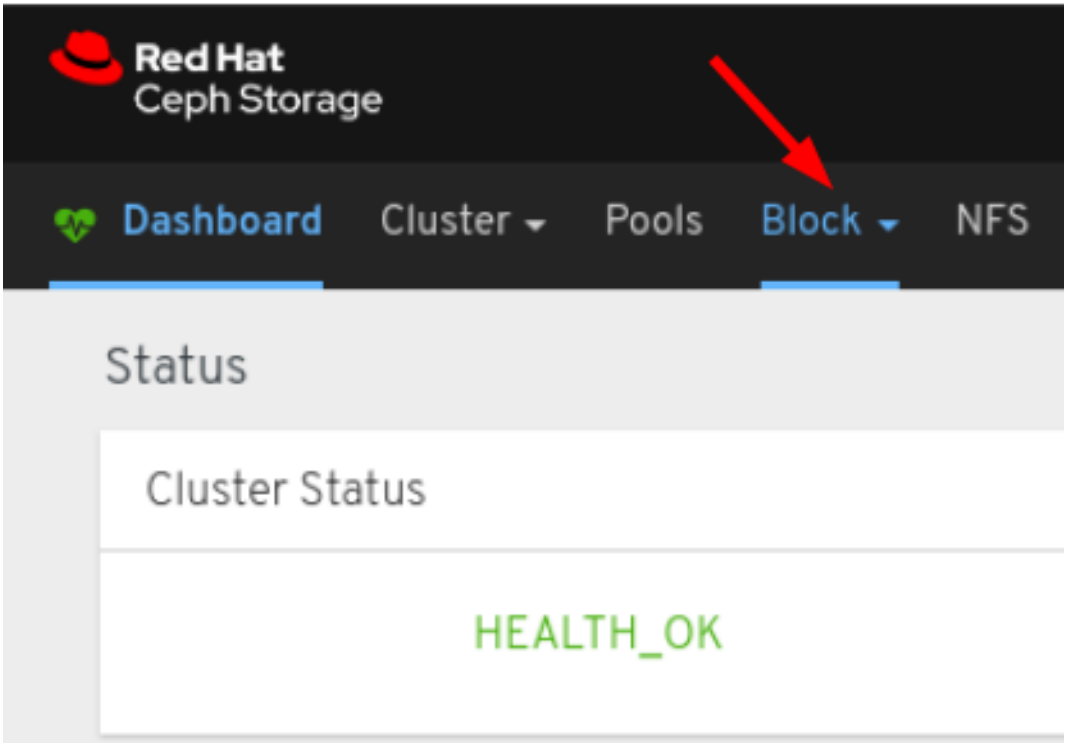

3. Select *Images* from the drop-down:

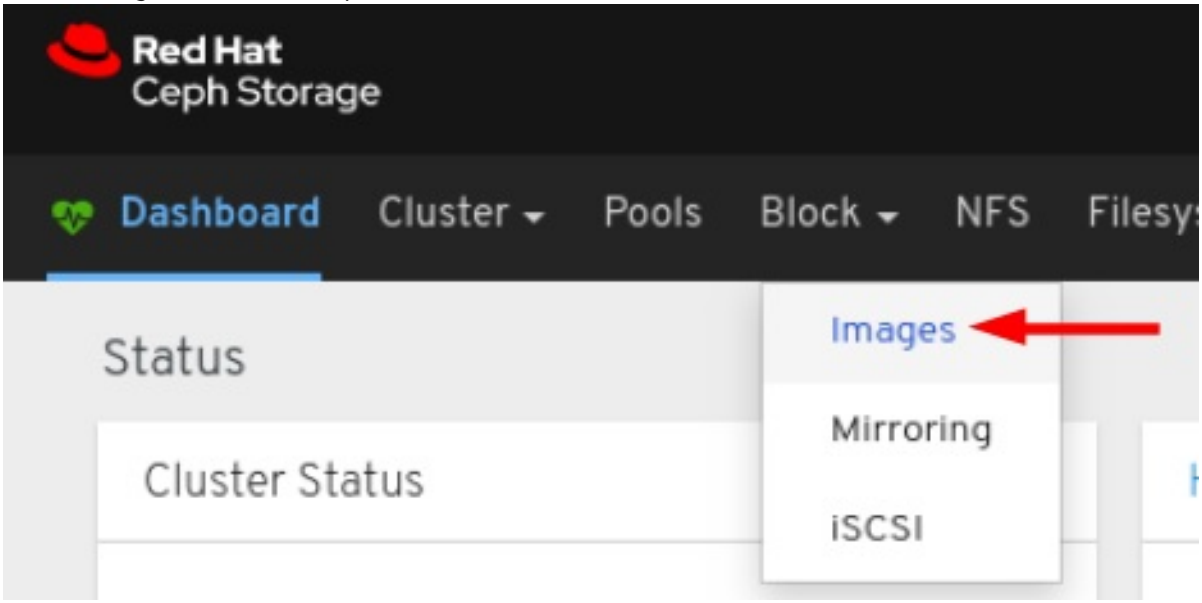

4. To edit the image, in the *Images* tab, click its row:

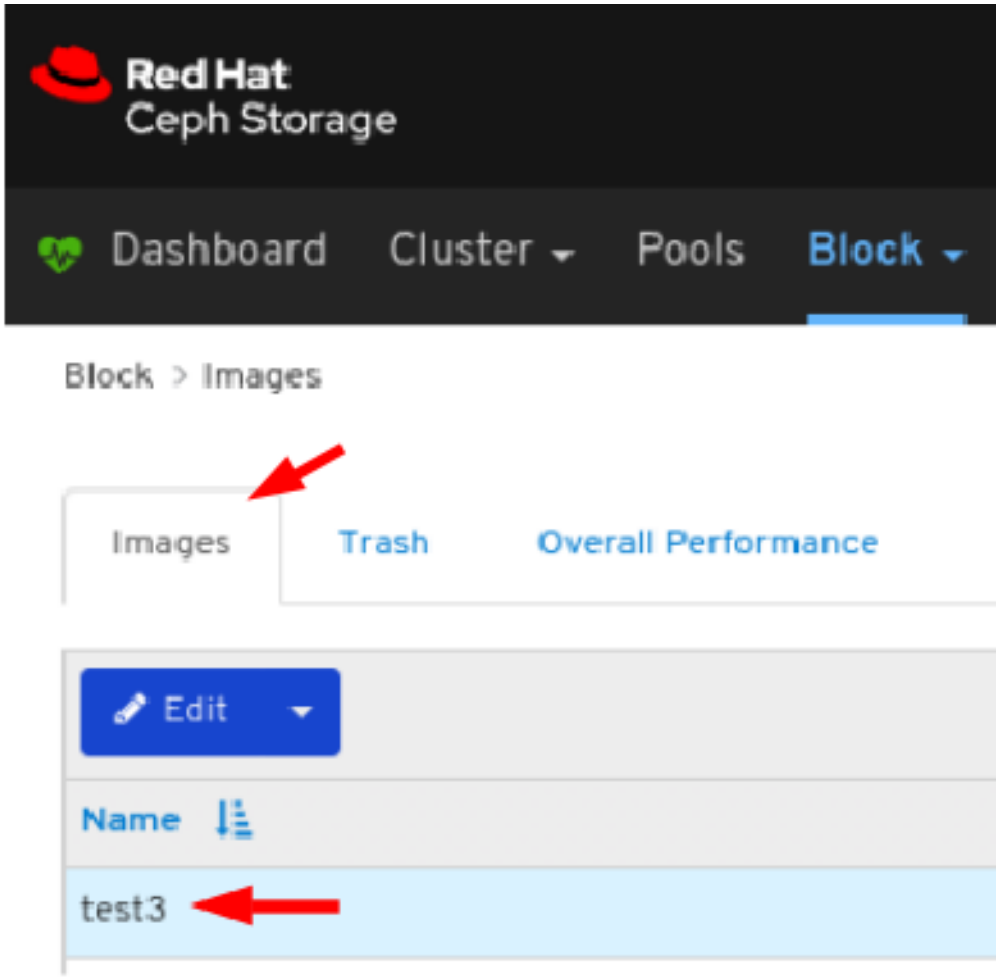

5. Select *Edit* In the *Edit* drop-down:

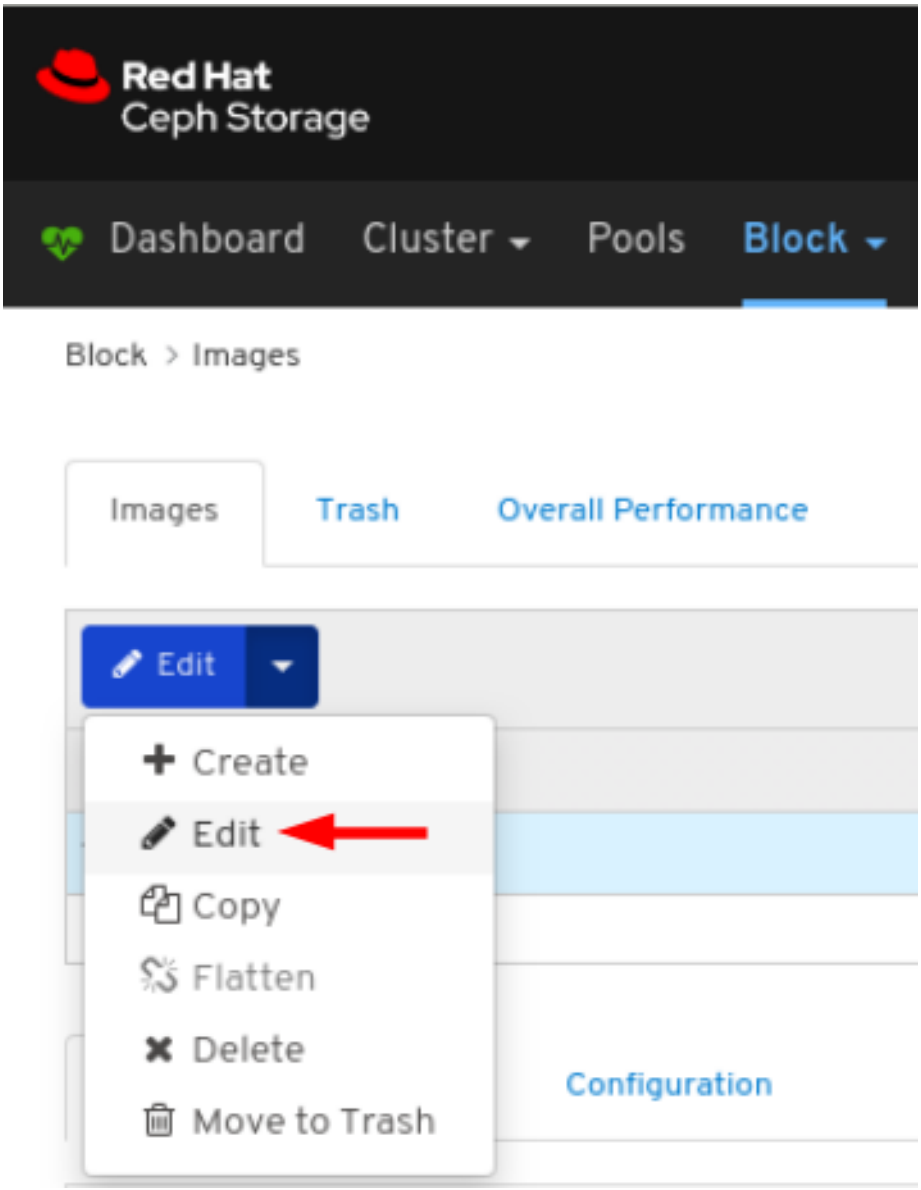

6. In the *EditRBD* dialog, edit the required parameters and click the *EditRBD* button:

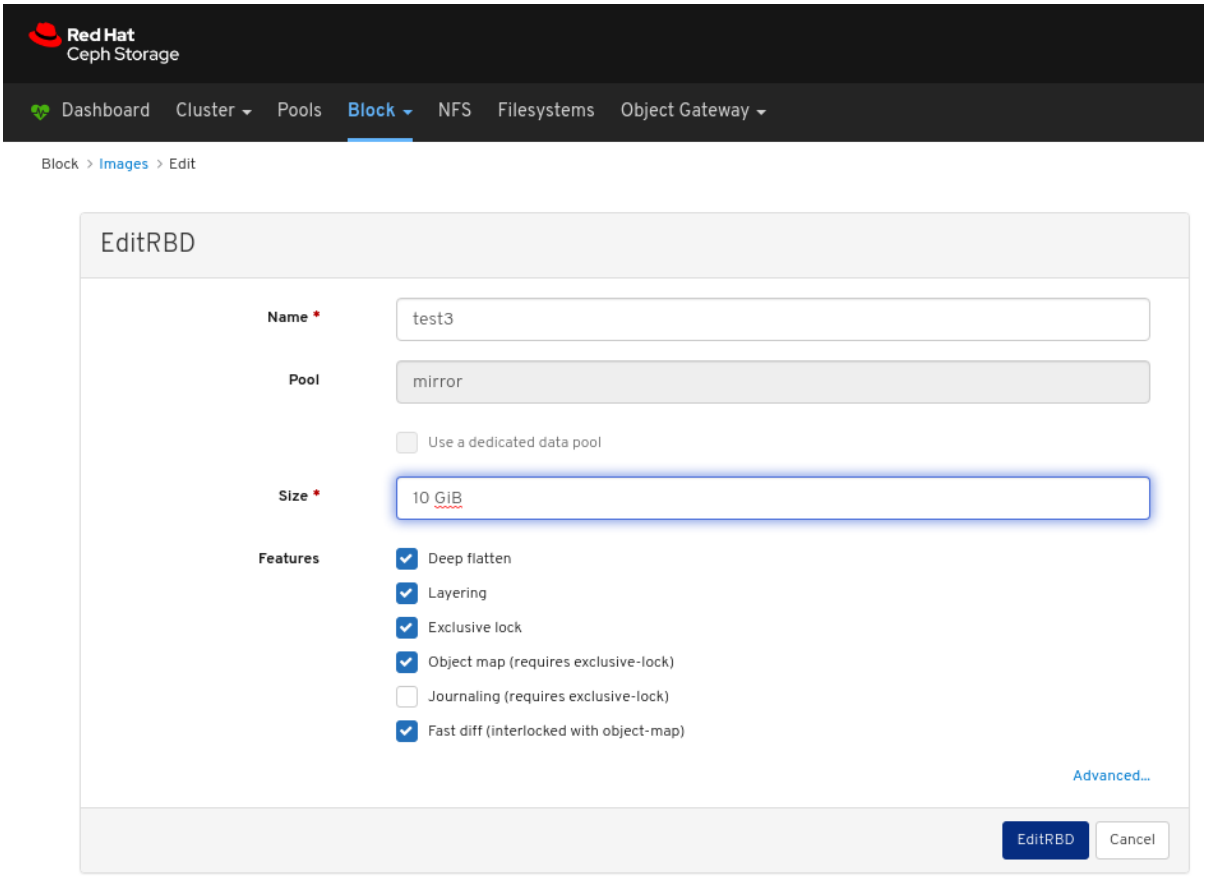

7. A notification towards the top right corner of the page indicates the image was updated successfully.

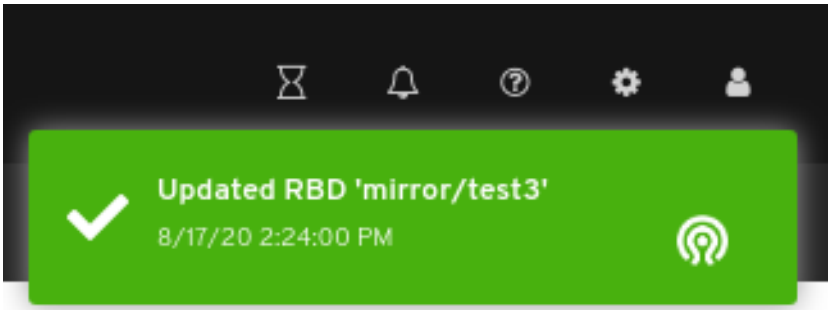

Additional Resources

- See the *Creating Images* section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

9.2.5. Copying images

The dashboard allows you to copy images.

Prerequisites

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

- An image is created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

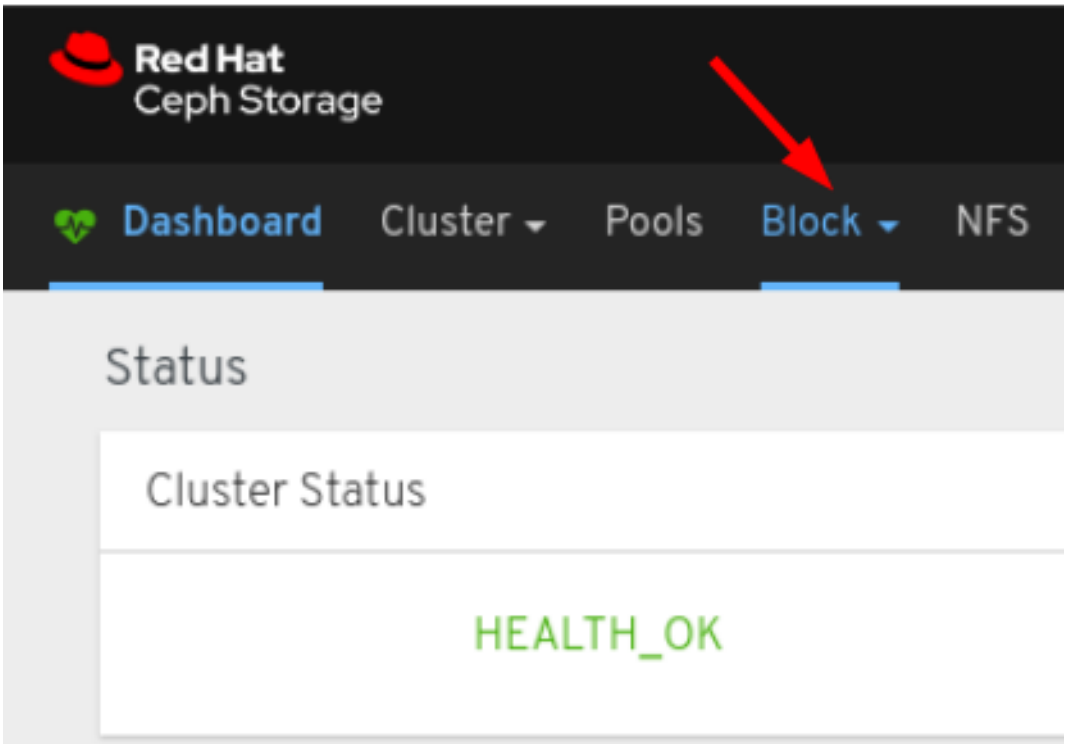

3. Select *Images* from the drop-down:

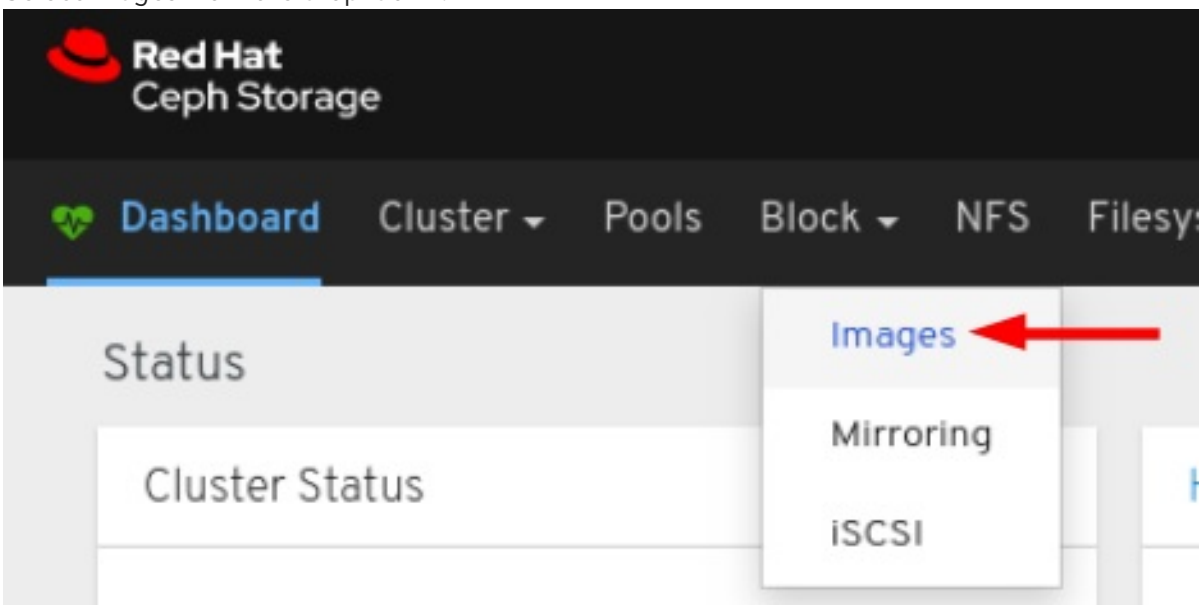

4. To copy the image, in the *Images* tab, click its row:

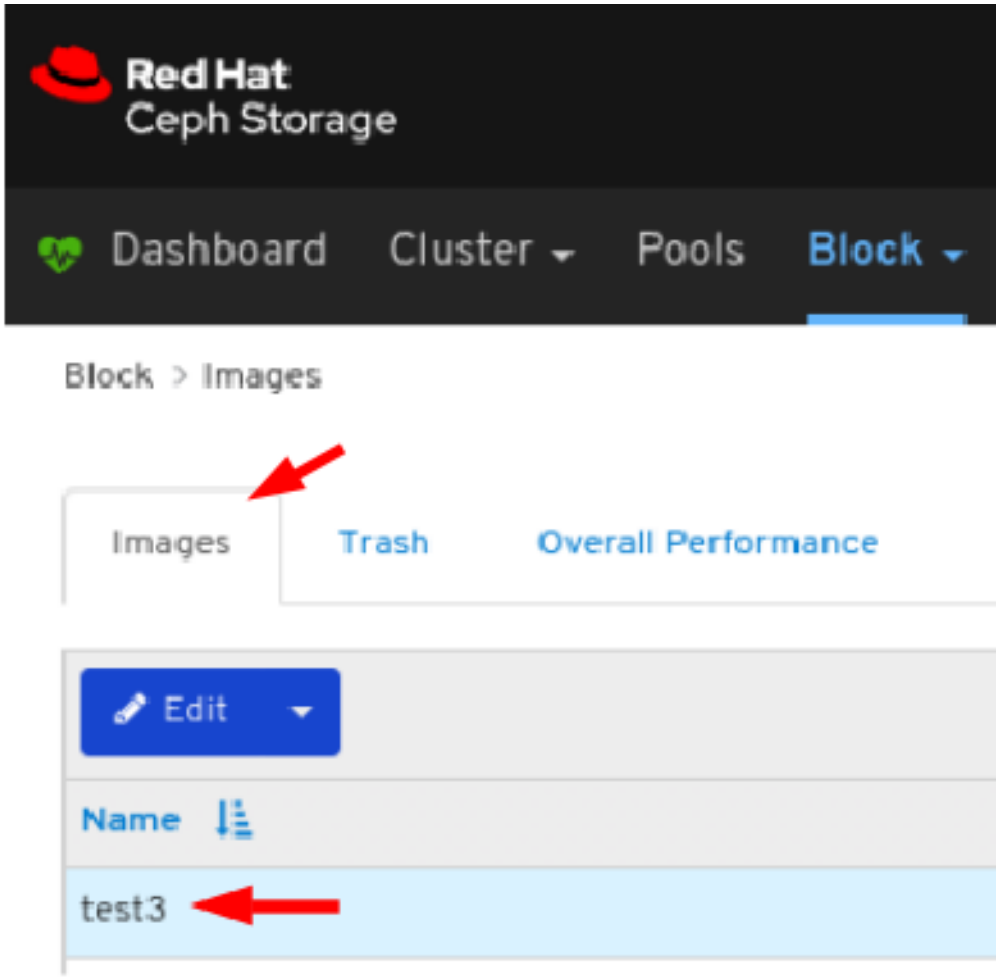

5. Select *Copy* In the *Edit* drop-down:

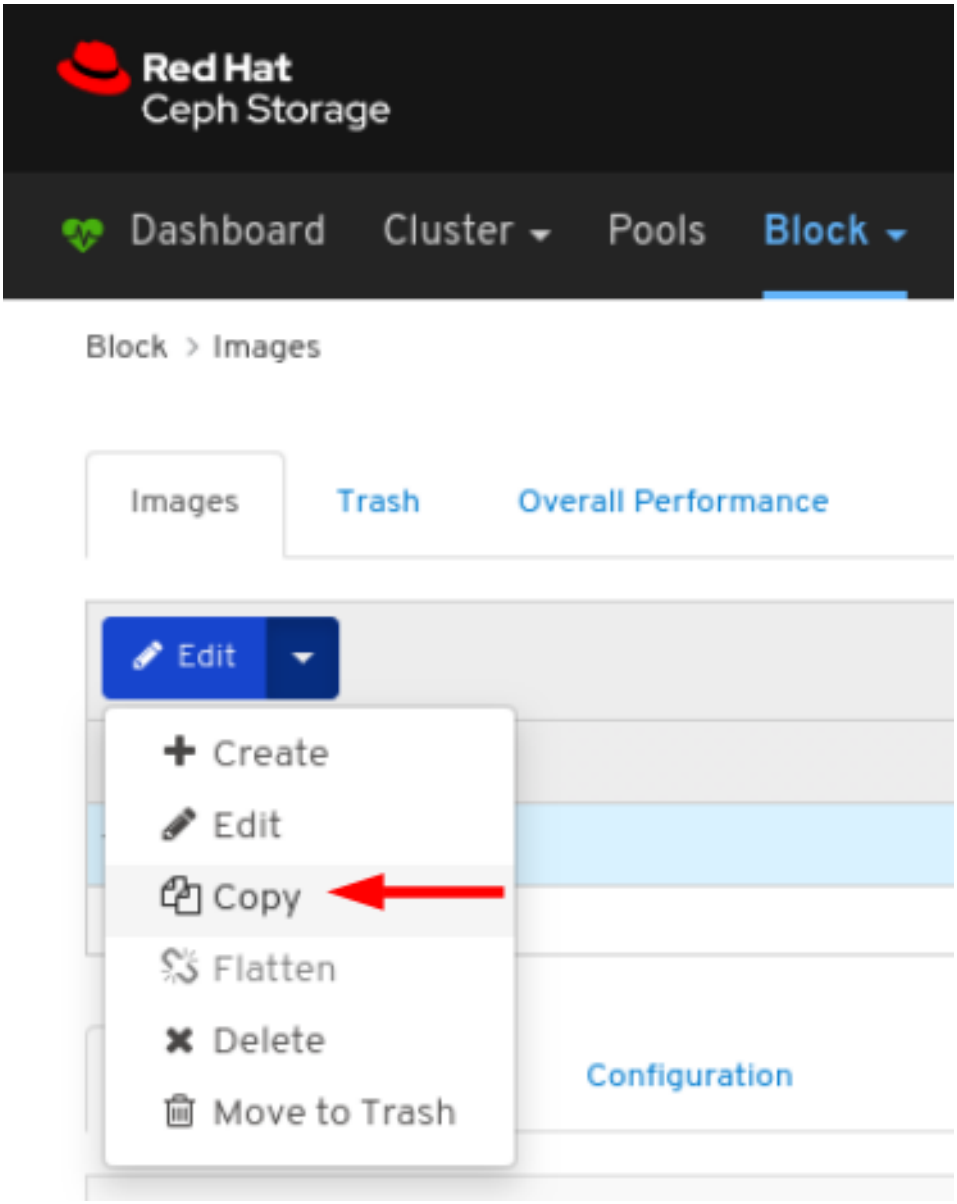

6. In the *CopyRBD* window, edit the required parameters and click the *CopyRBD* button:

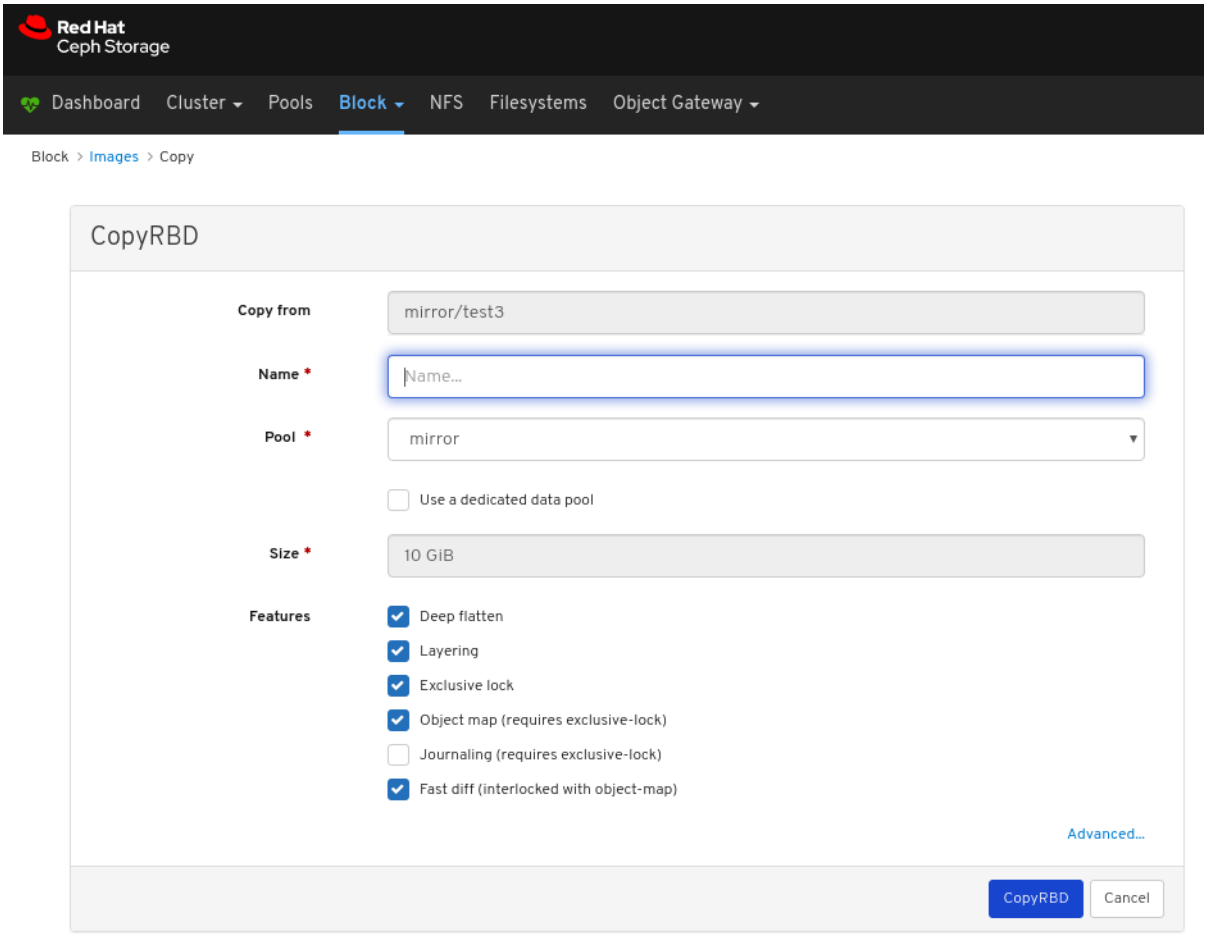

7. A notification towards the top right corner of the page indicates the image was updated successfully.

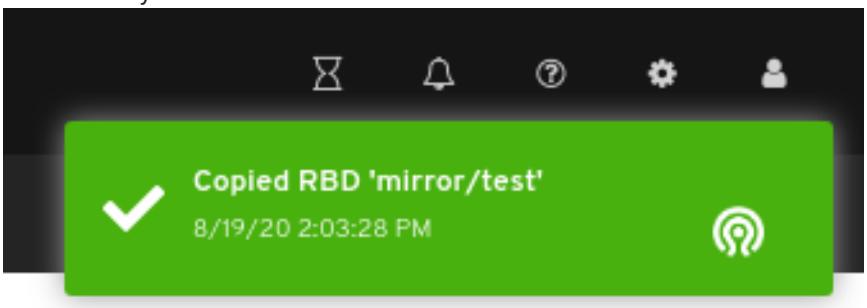

Additional Resources

- See the *Creating Images* section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

9.2.6. Moving images to trash

The dashboard allows you to move images to trash.

Prerequisites

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

- A pool with the *rbd* application enabled is created.
- An image is created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

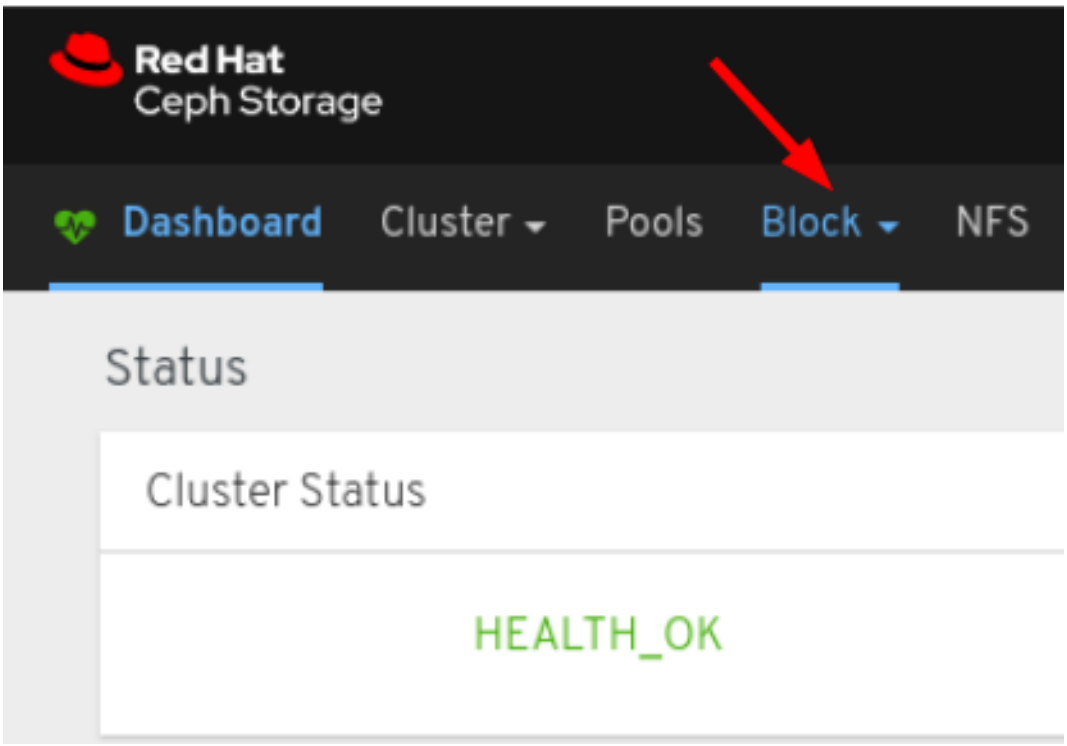

3. Select *Images* from the drop-down:

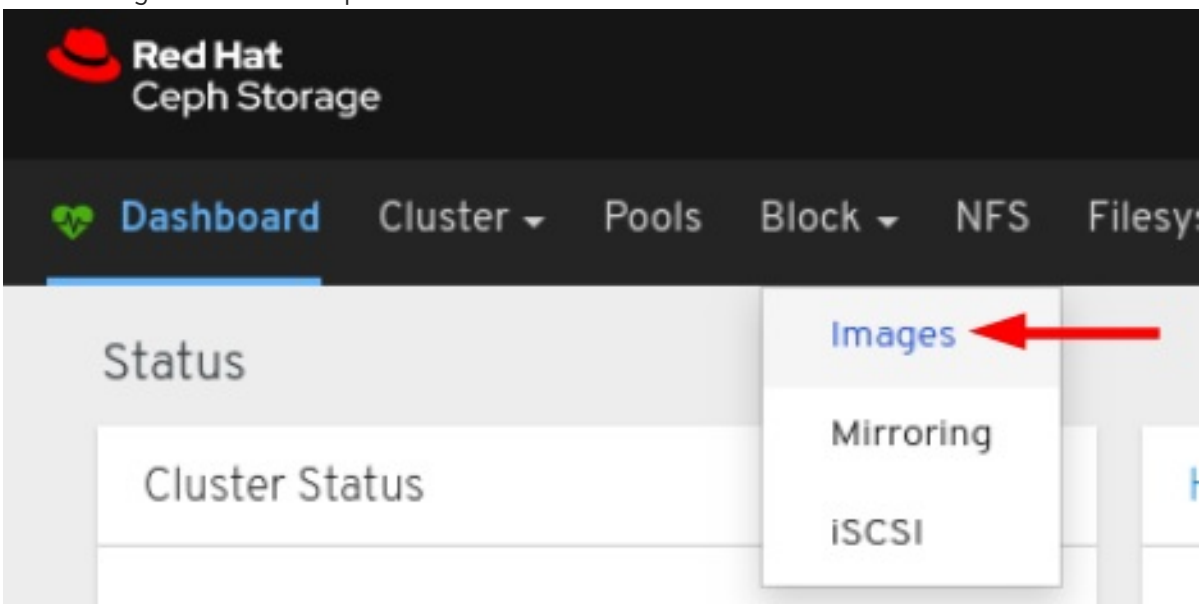

4. To move the image to Trash, in the *Images* tab, click its row:

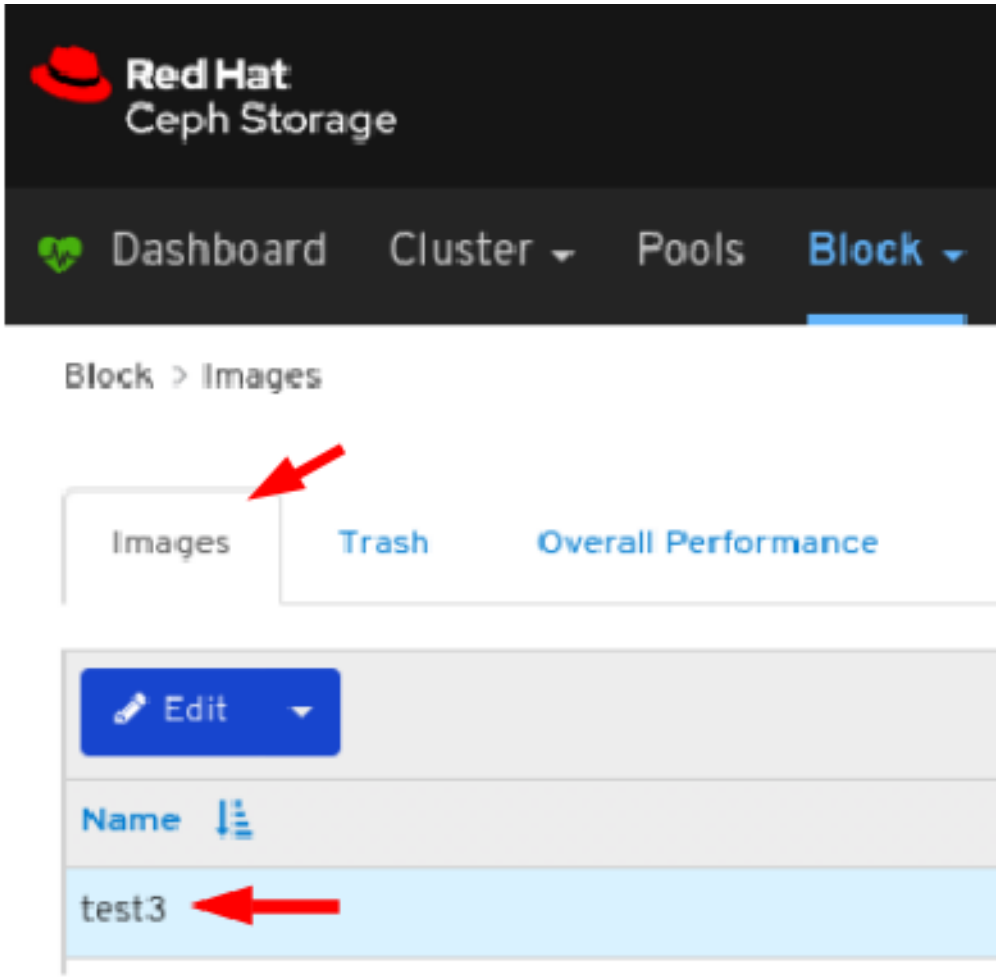

5. Select *Move to Trash* In the *Edit* drop-down:

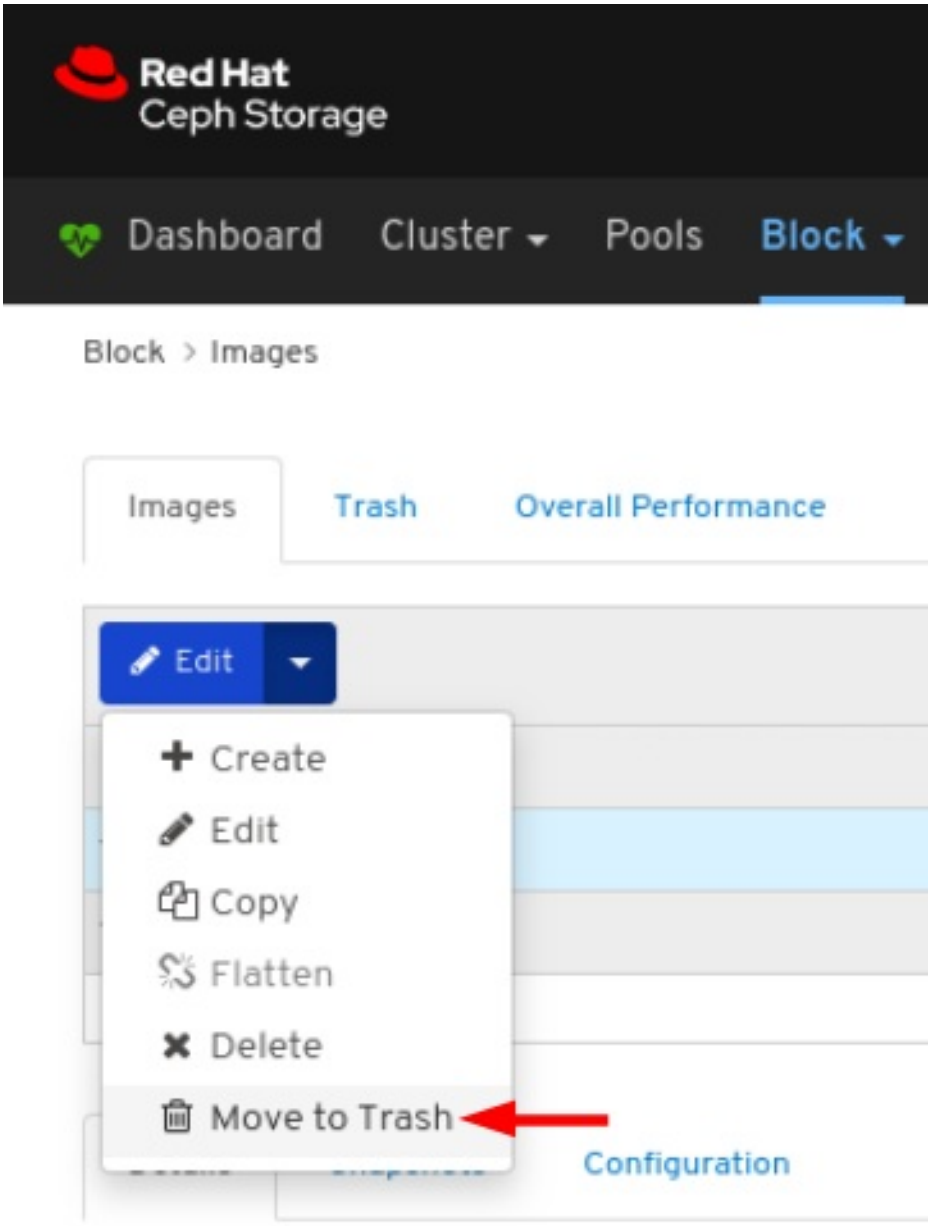

6. In the *Moving an image to trash* window, edit the date till which the image needs protection, and then click the *Move Image* button:

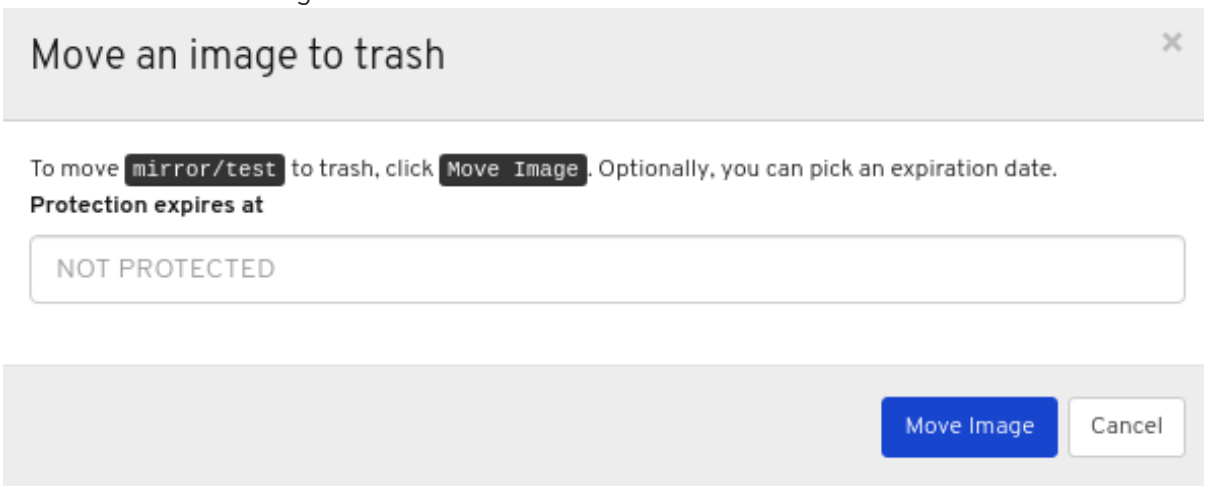

7. A notification towards the top right corner of the page indicates the image was moved to trash successfully.

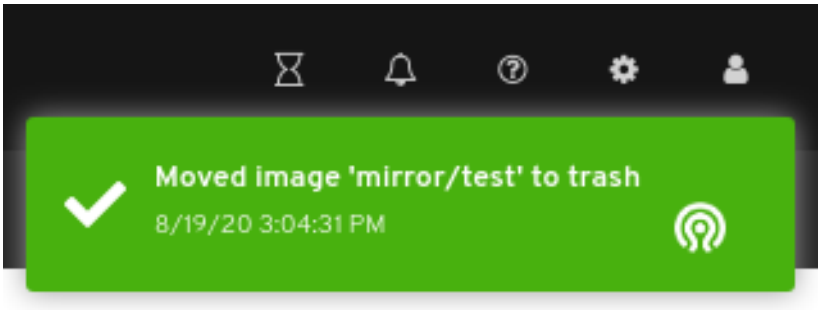

9.2.7. Purging trash

The dashboard allows you to purge trash of images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

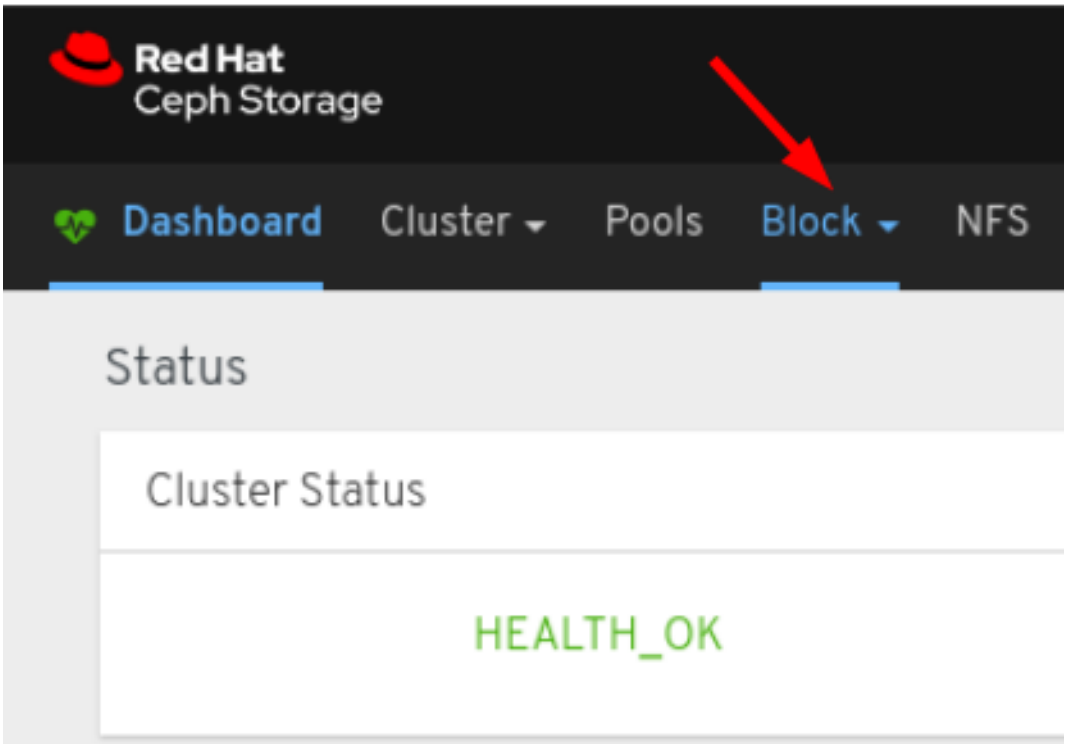

3. Select *Images* from the drop-down:

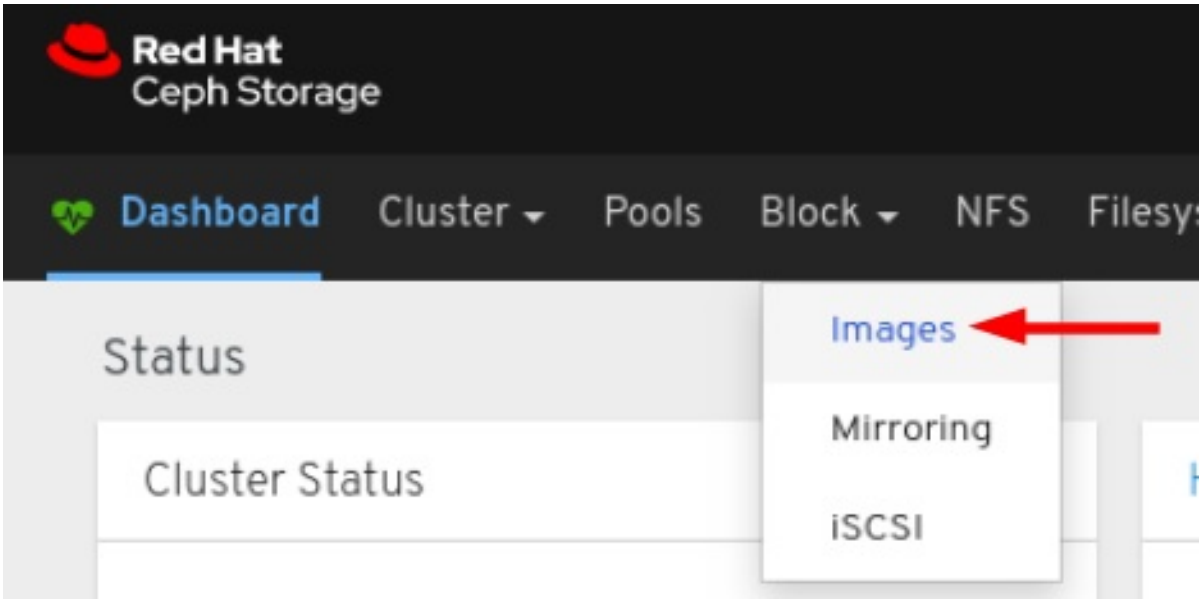

4. In the *Trash* tab, click the *Purge Trash* button:

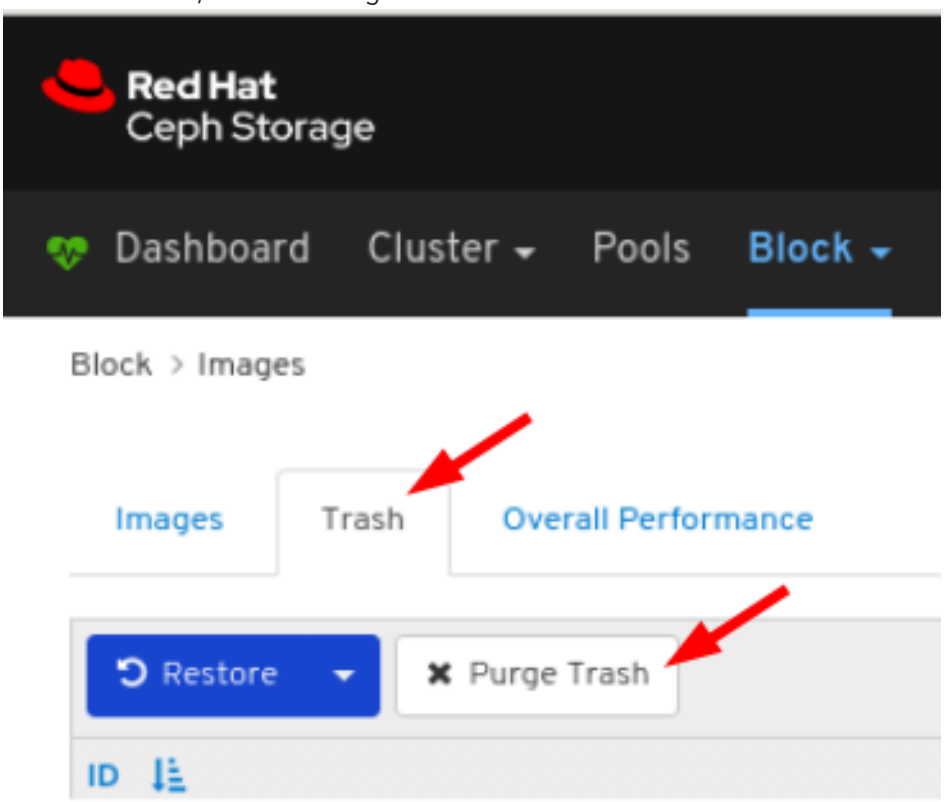

5. In the *Purge Trash* window, select the pool, and then click the *Purge Trash* button:

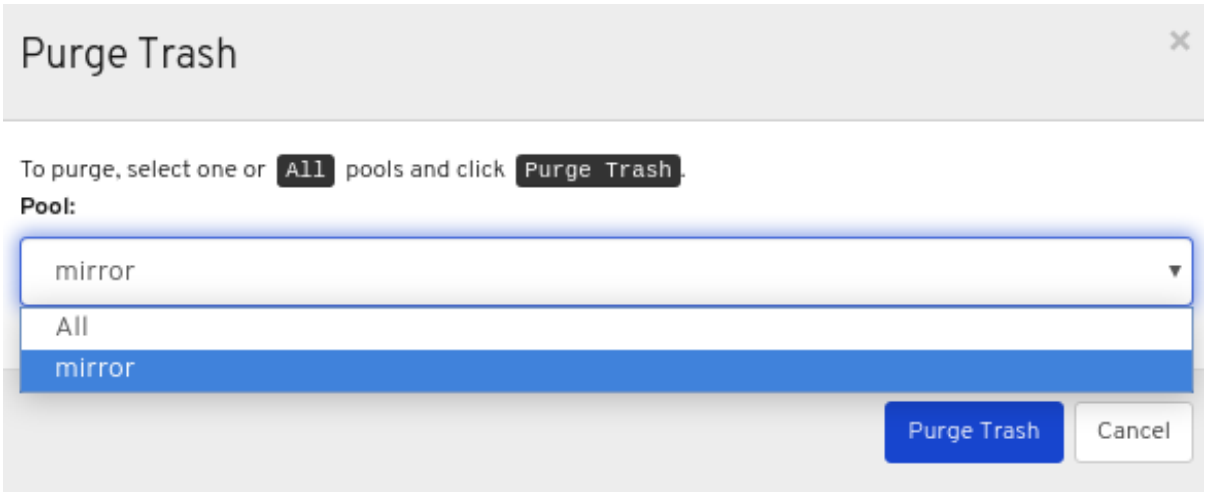

6. A notification towards the top right corner of the page indicates the images were purged from the selected pools successfully.

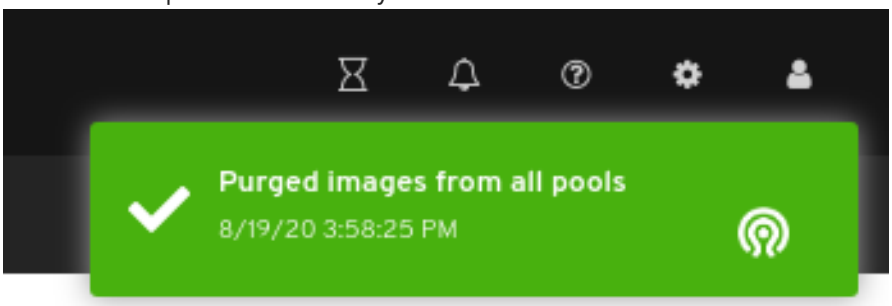

Additional resources

- See the *Purging the Block Device Snapshots* section in the *Red Hat Ceph Storage Block Device Guide* for more details.

9.2.8. Restoring images from trash

The dashboard allows you to restore images from trash.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

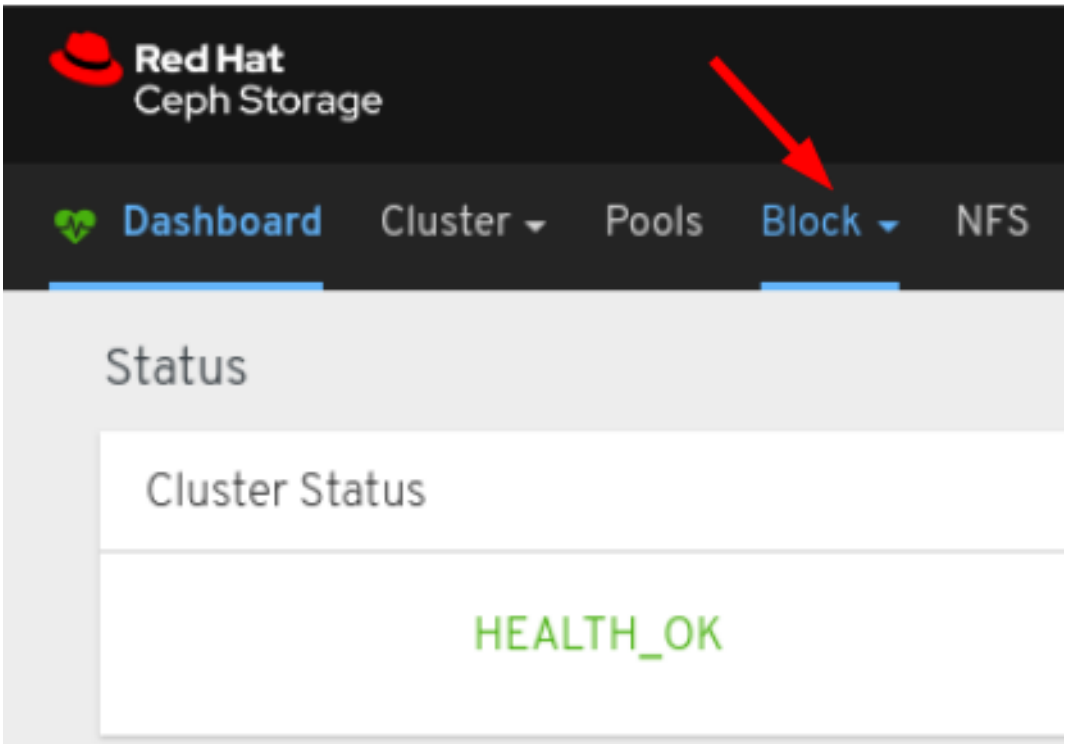

3. Select *Images* from the drop-down:

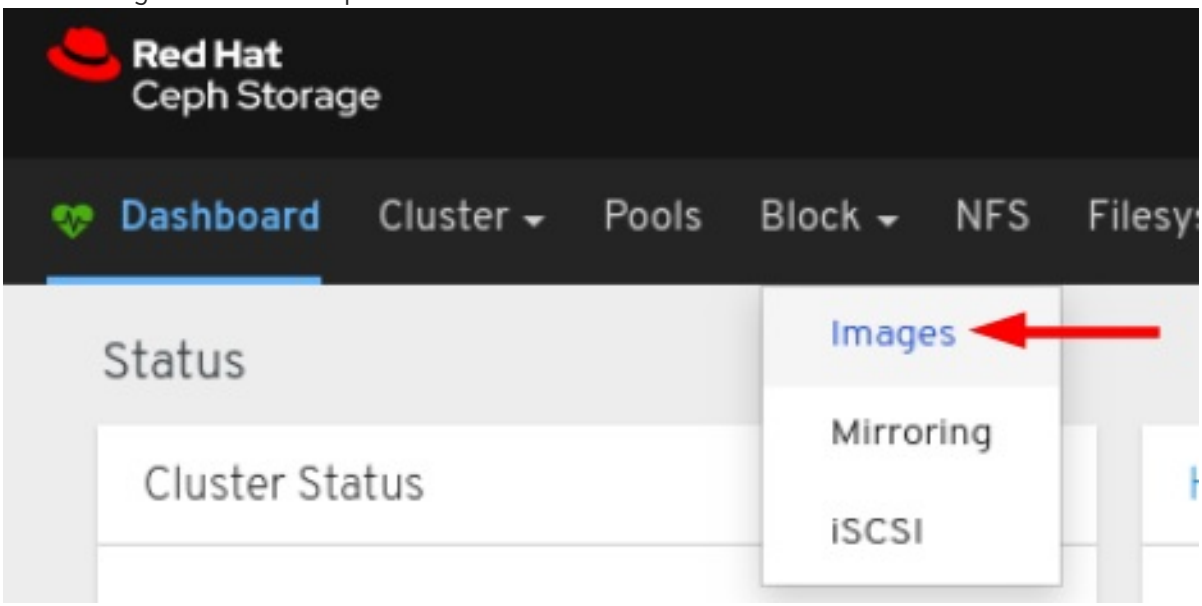

4. To restore the image from Trash, in the *Trash* tab, click its row:

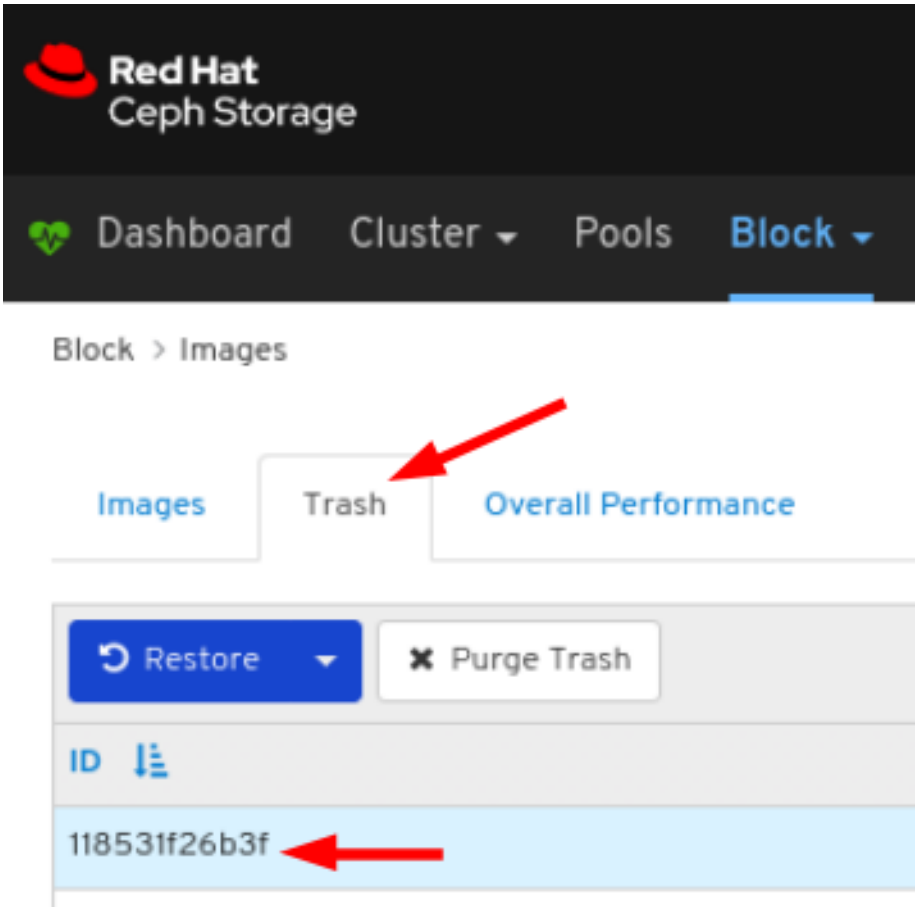

5. Select *Restore* in the *Restore* drop-down:

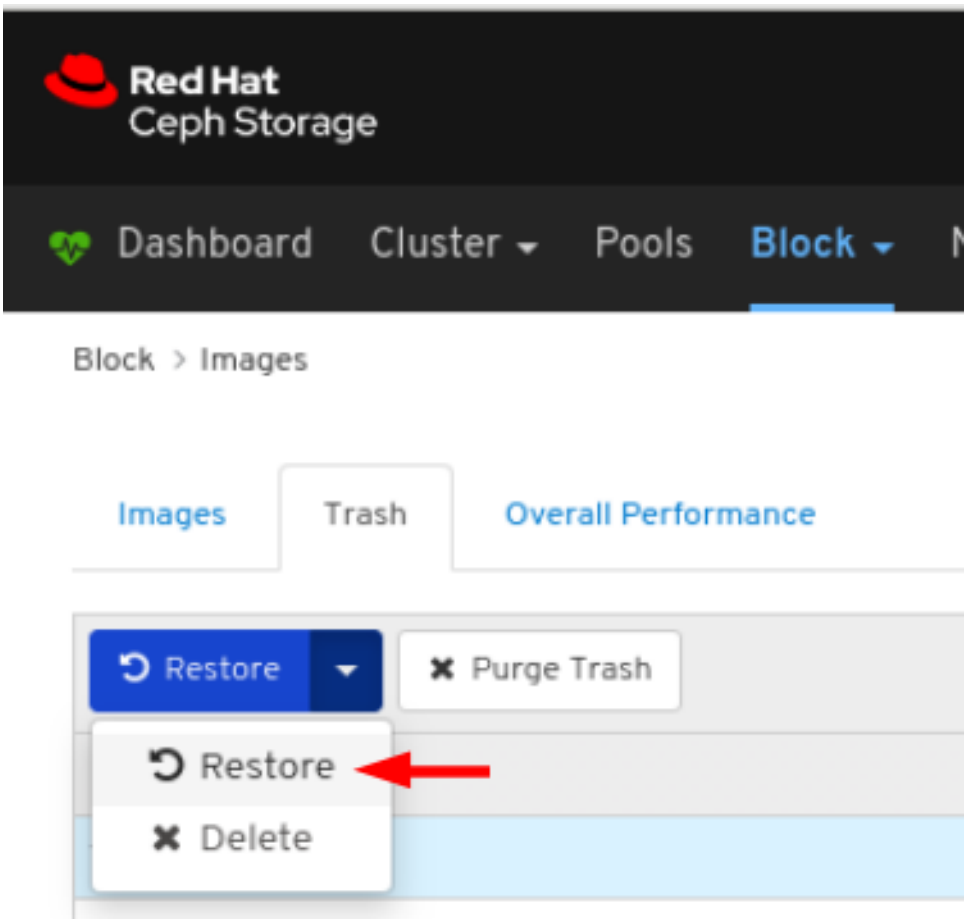

6. In the *Restore Image* window, enter the name of the image , and then click the *Restore Image* button:

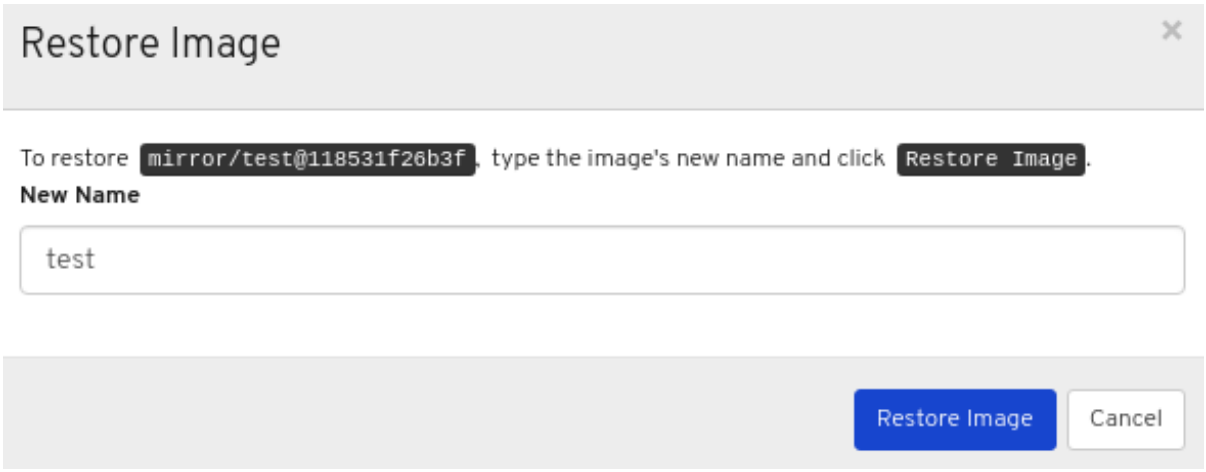

7. Notification towards the top right corner of the page indicate the image was restored from trash successfully.

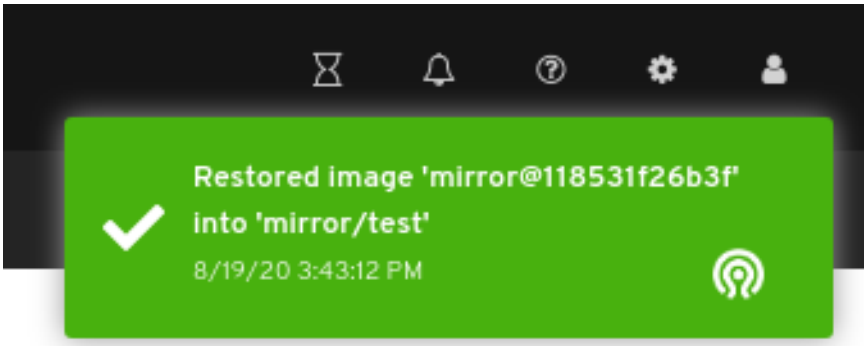

Additional resources

- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.

9.2.9. Deleting images

The dashboard allows you to delete images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

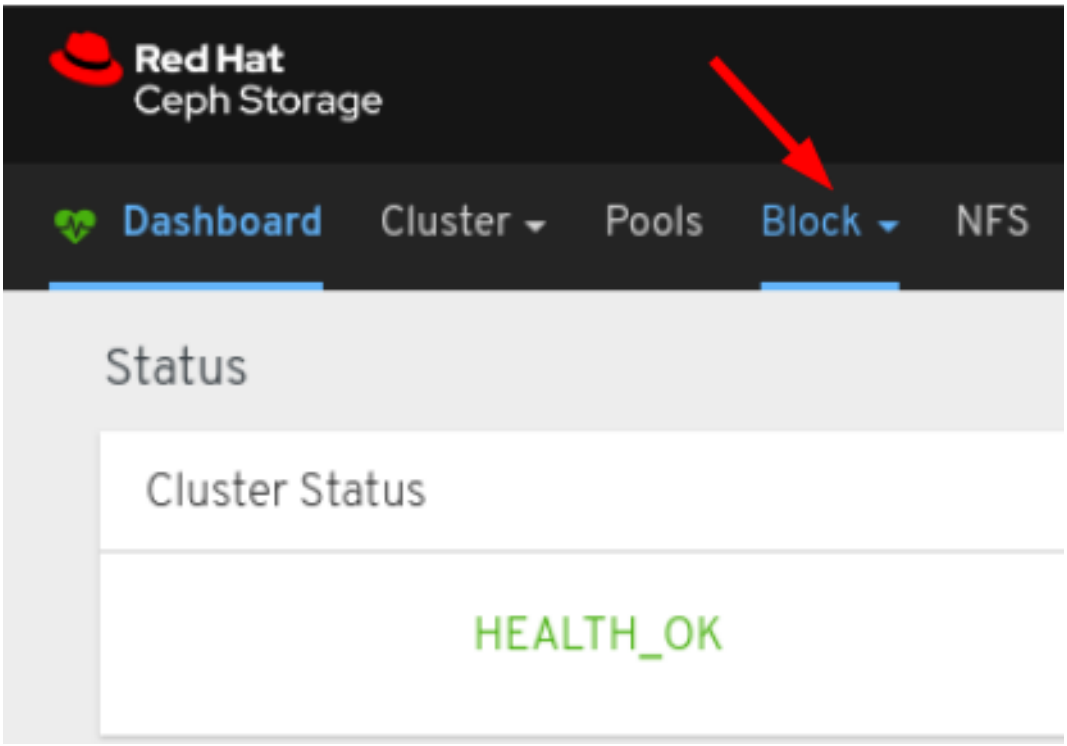

3. Select *Images* from the drop-down:

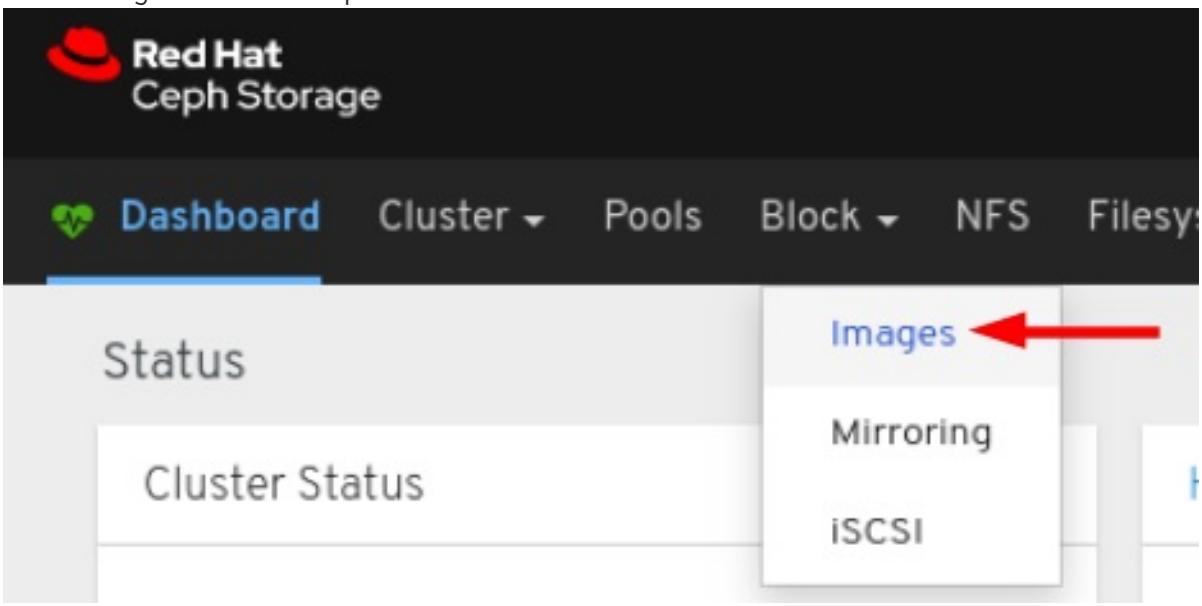

4. To edit the image, in the *Images* tab, click its row:

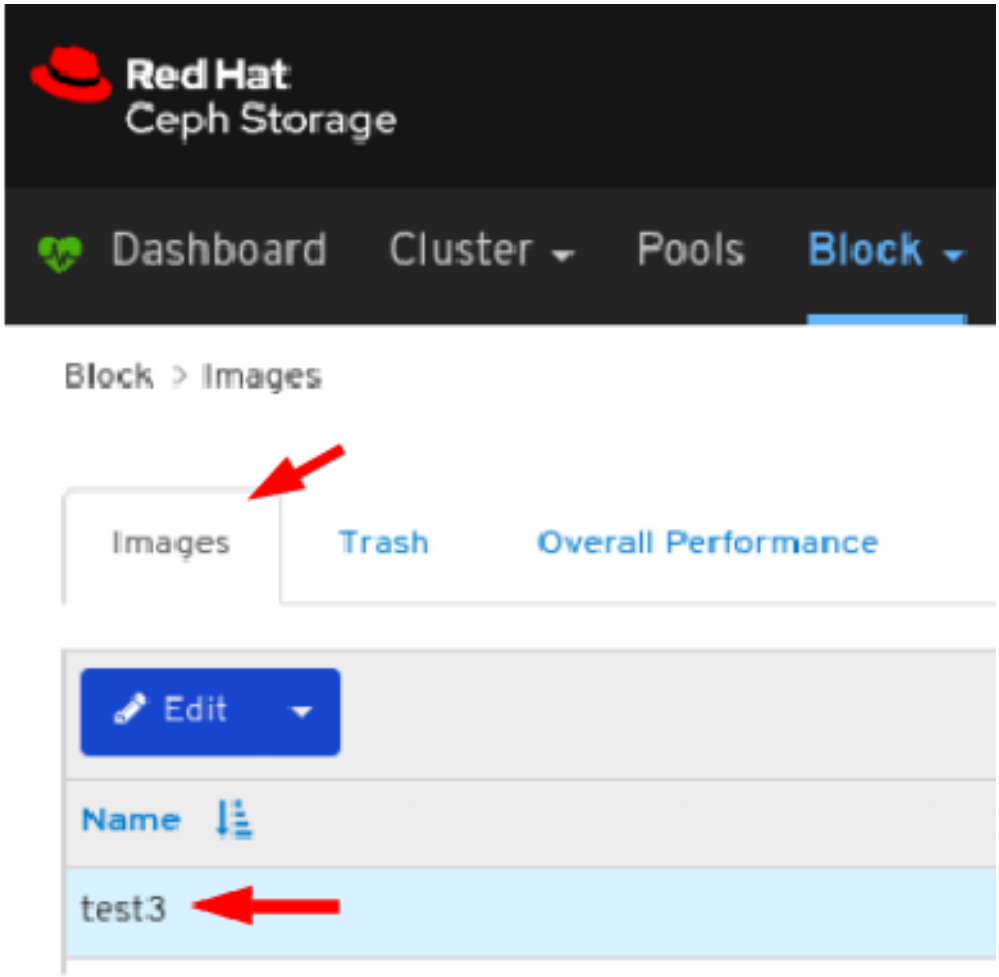

5. Select *Delete* In the *Edit* drop-down:

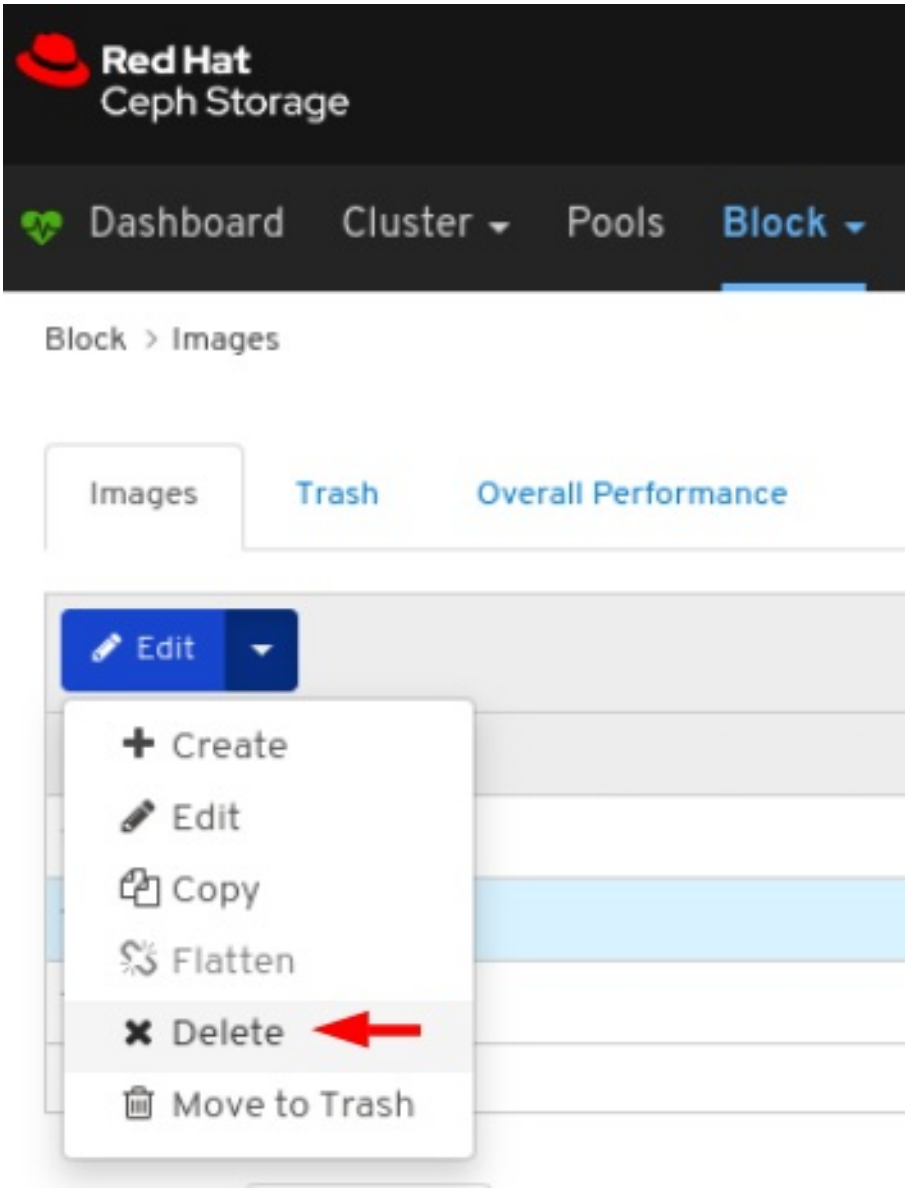

6. In the *Delete RBD* dialog window, Click the *Yes, I am sure* box and then Click *Delete RBD* to save the settings:

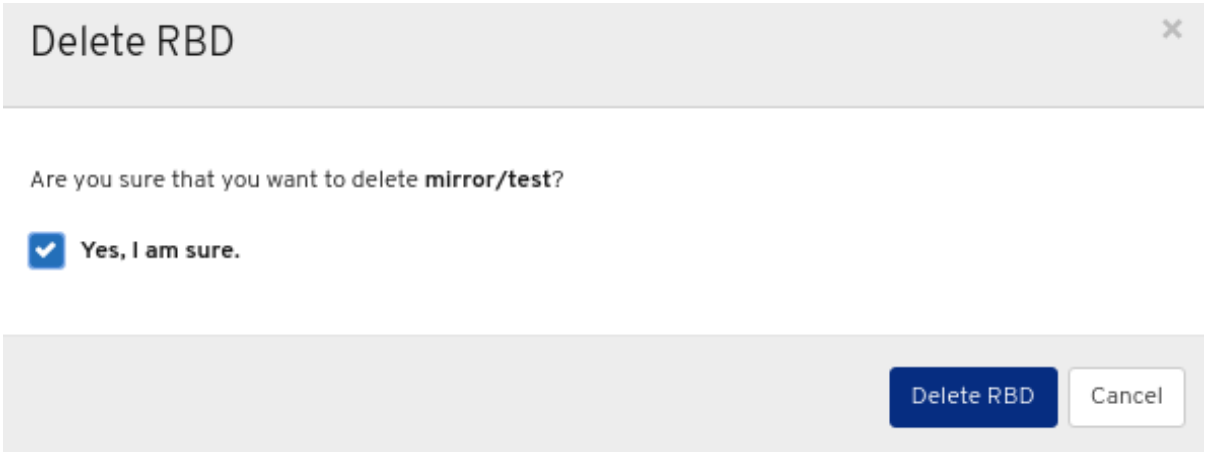

7. A notification towards the top right corner of the page indicates the image was moved to trash successfully.

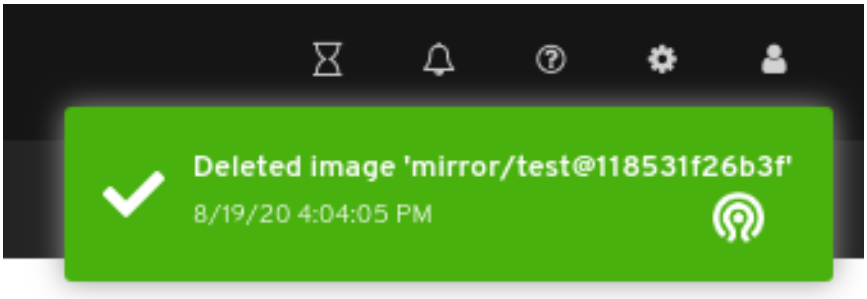

9.2.10. Creating snapshots of images

The dashboard allows you to take snapshots of Ceph block device images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

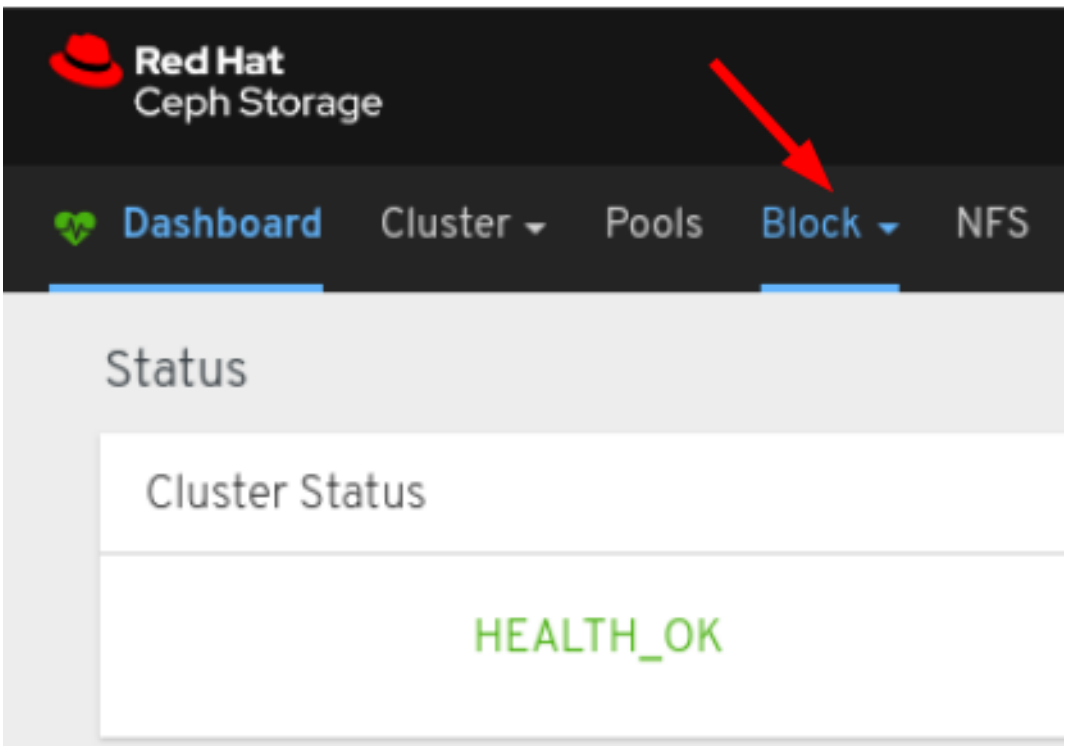

3. Select *Images* from the drop-down:

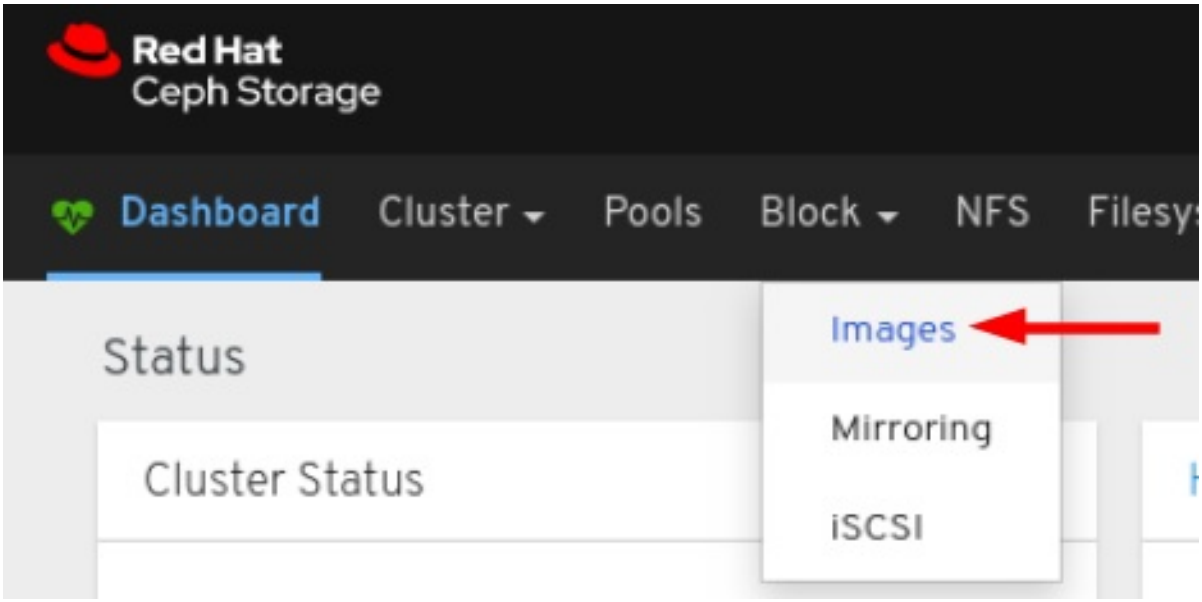

4. To take the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

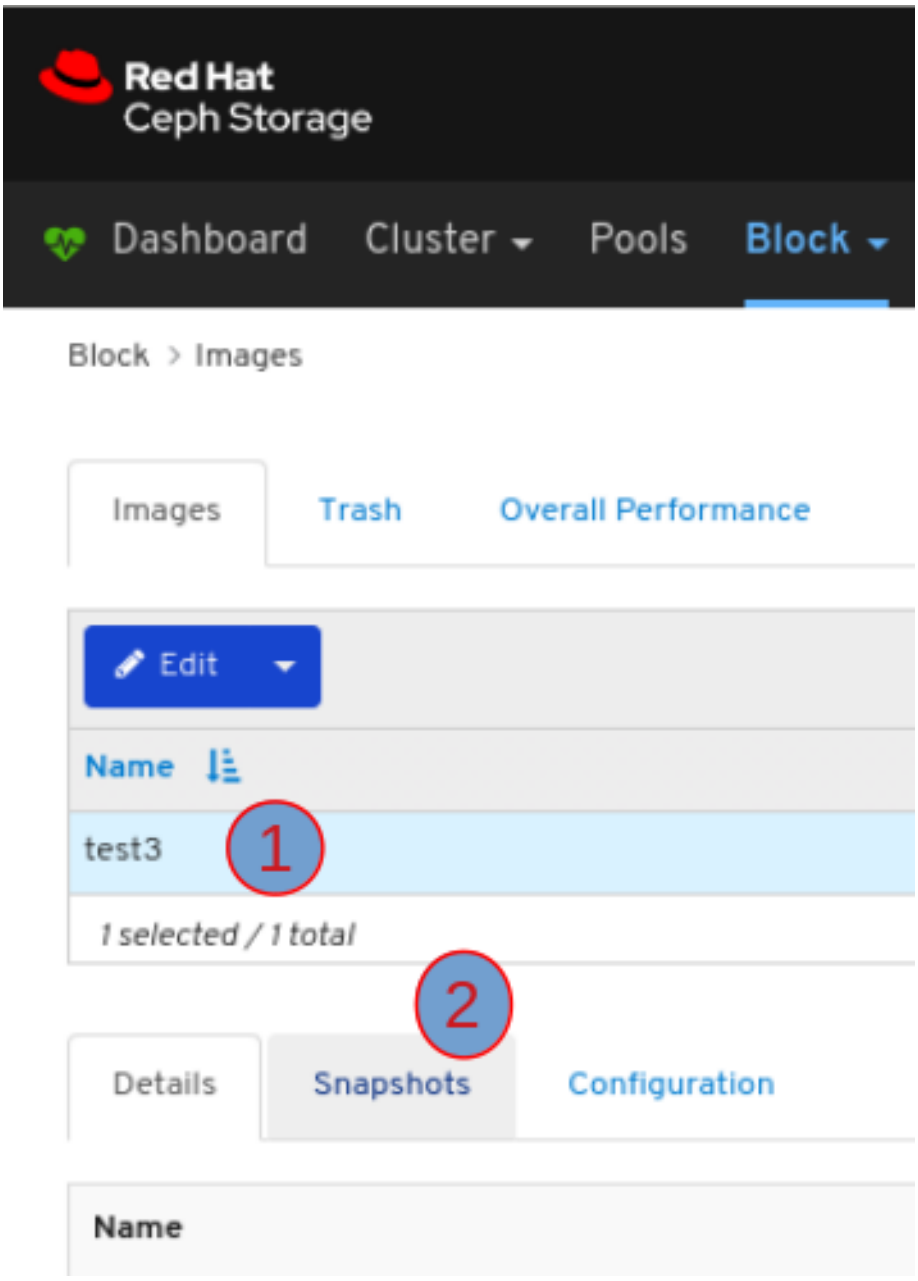

5. Select *Create* in the *Create* drop-down:

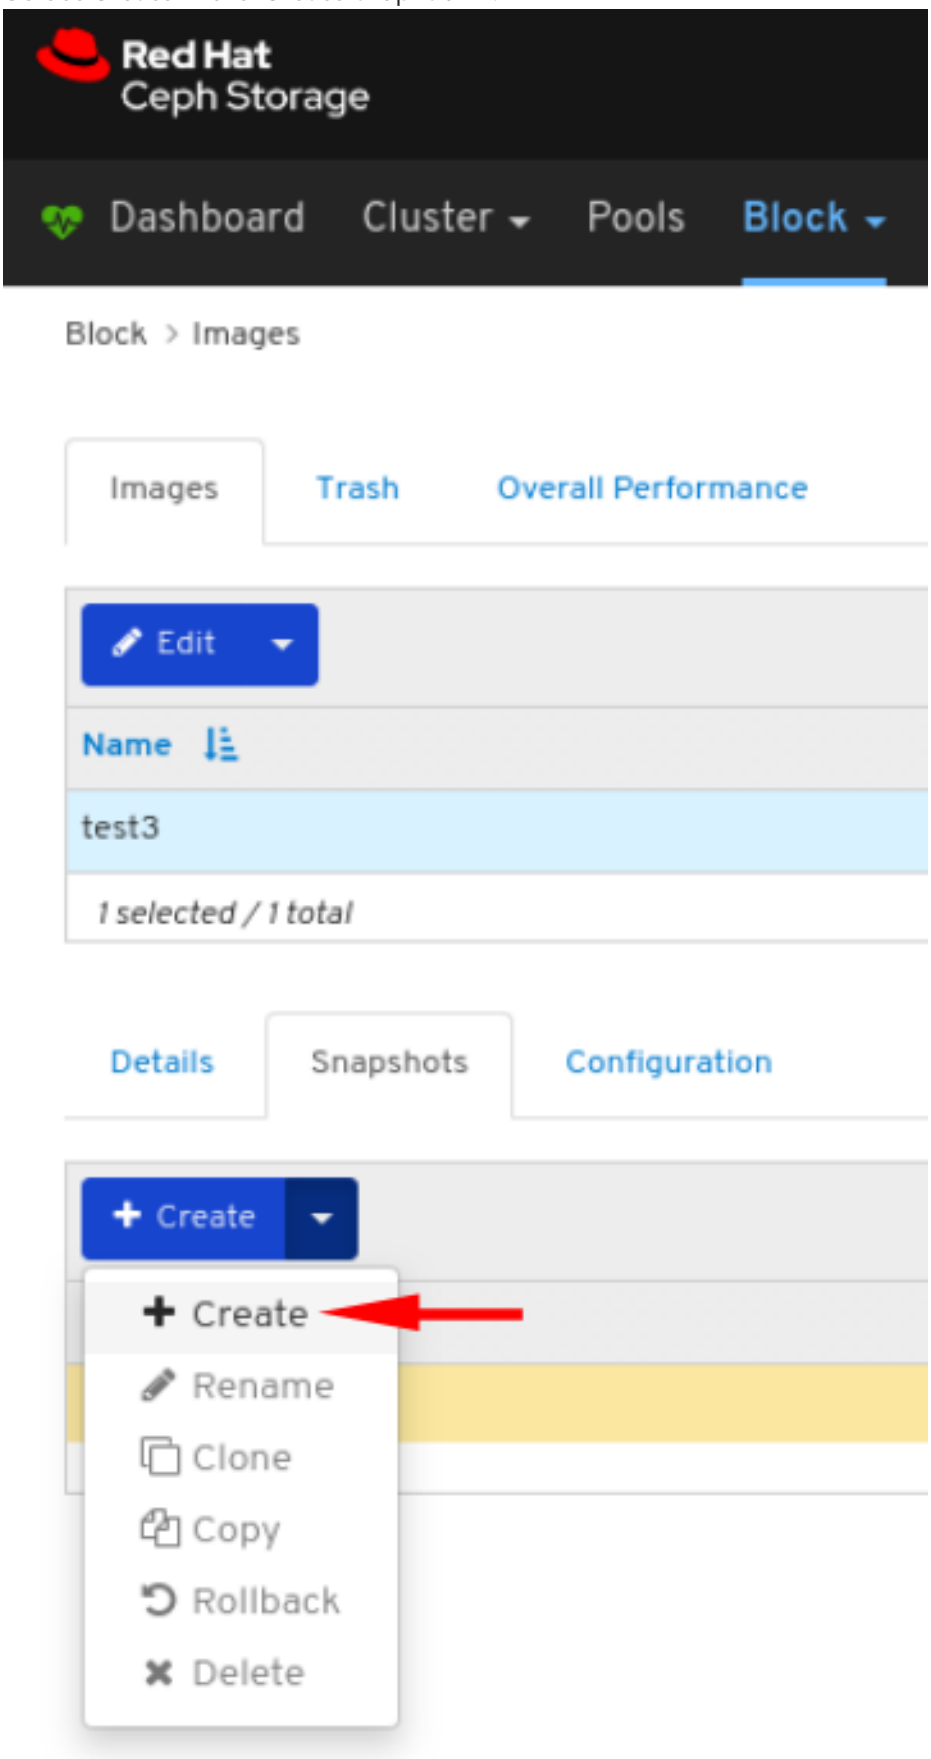

6. In the *CreateRBD Snapshot* dialog, enter the parameters and click the *CreateRBD Snapshot* button:

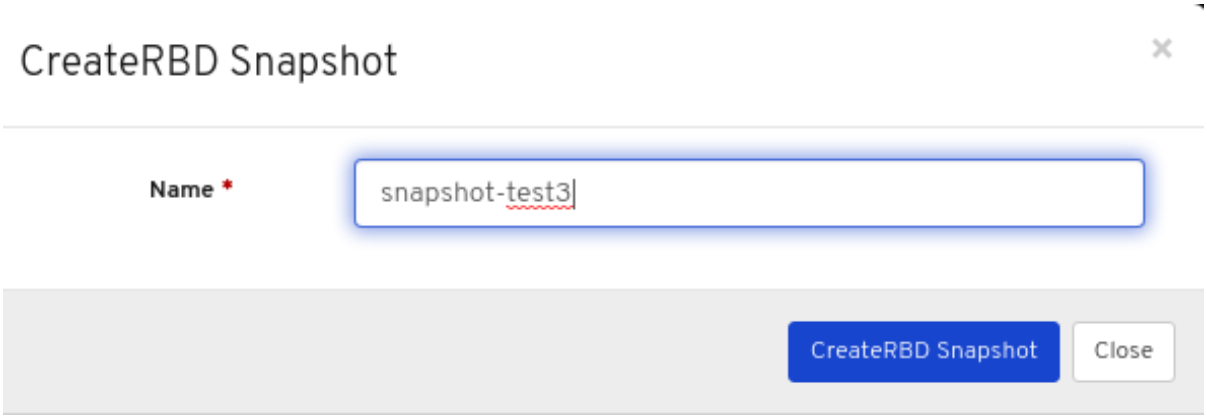

7. A notification towards the top right corner of the page indicates the snapshot of the image was created successfully.

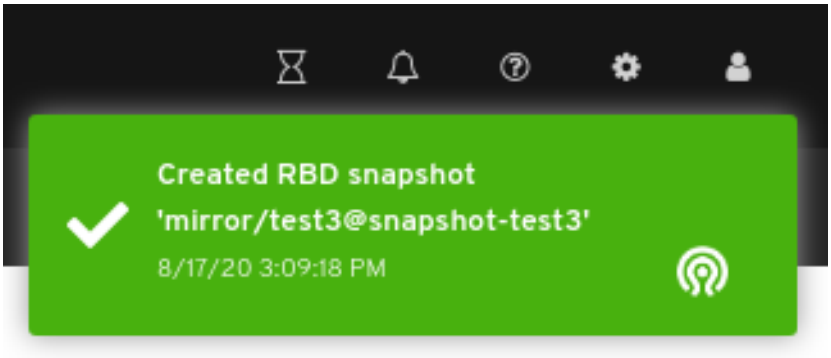

Additional Resources

- See the *Creating a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide*] for more information on creating snapshots.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

9.2.11. Renaming snapshots of images

The dashboard allows you to rename snapshots of Ceph block device images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.

2. On the navigation bar, click *Block*:

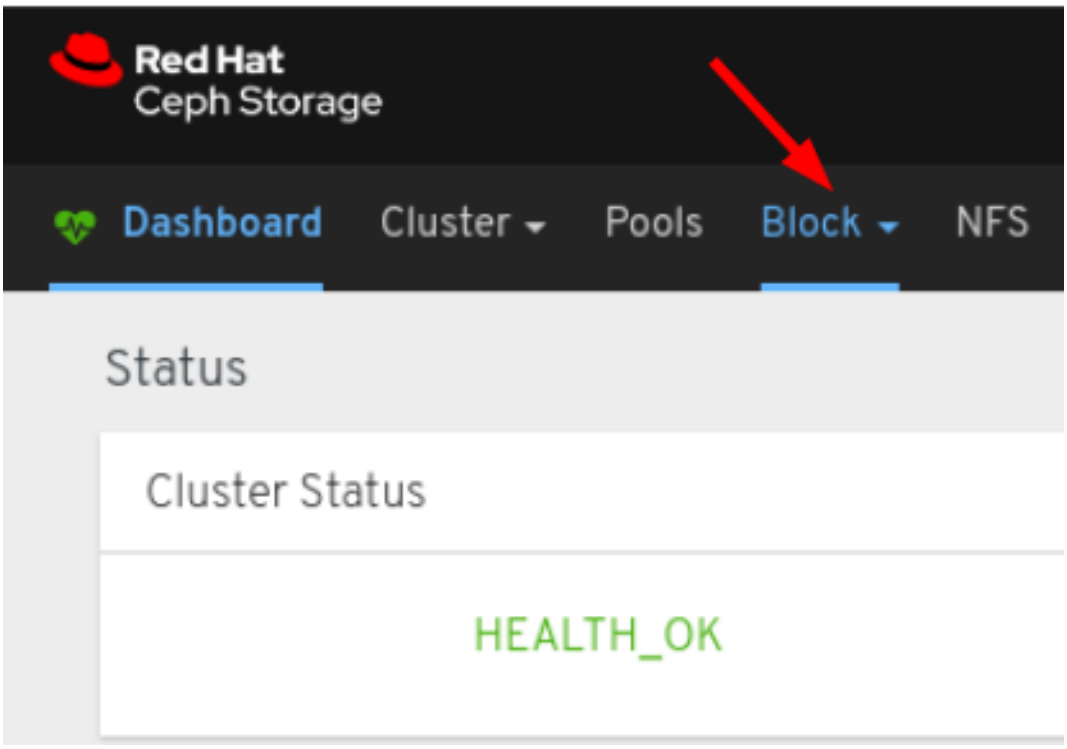

3. Select *Images* from the drop-down:

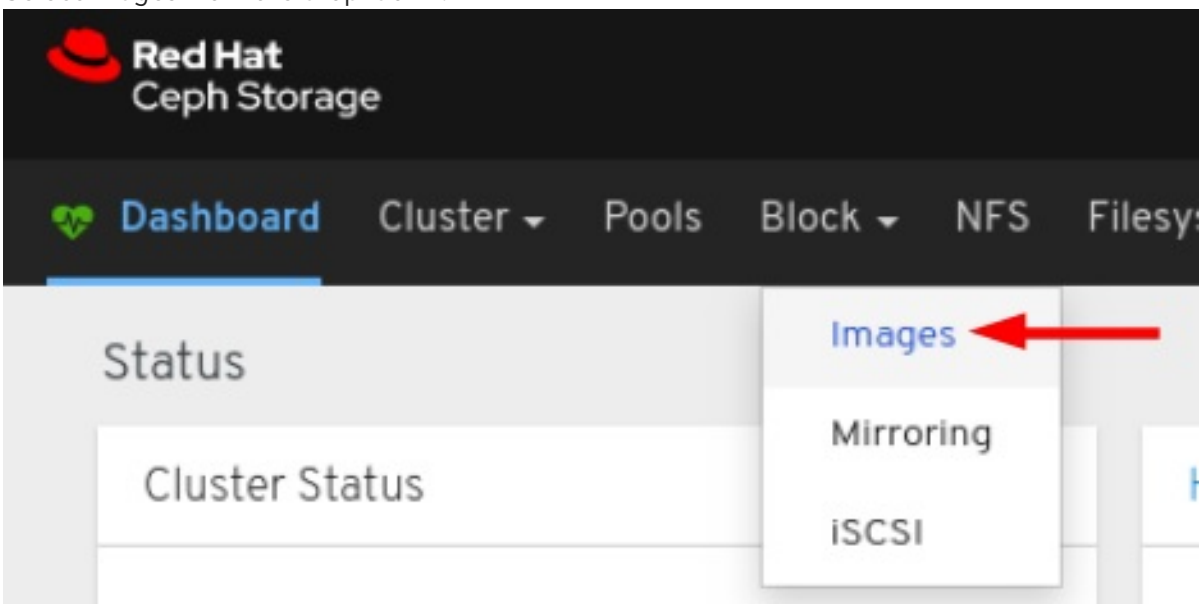

4. To rename the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

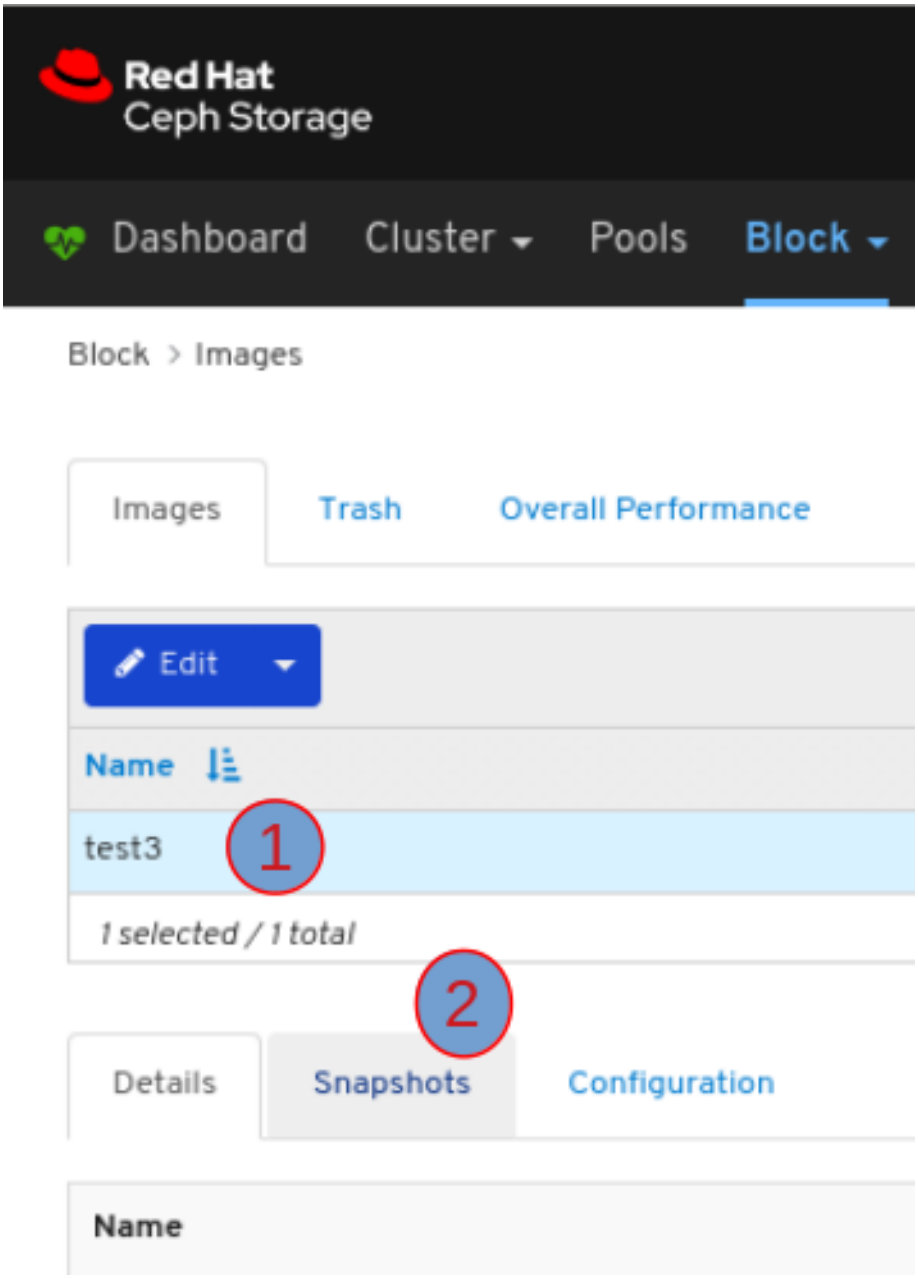

5. Select *Rename* in the the *Rename* drop-down:

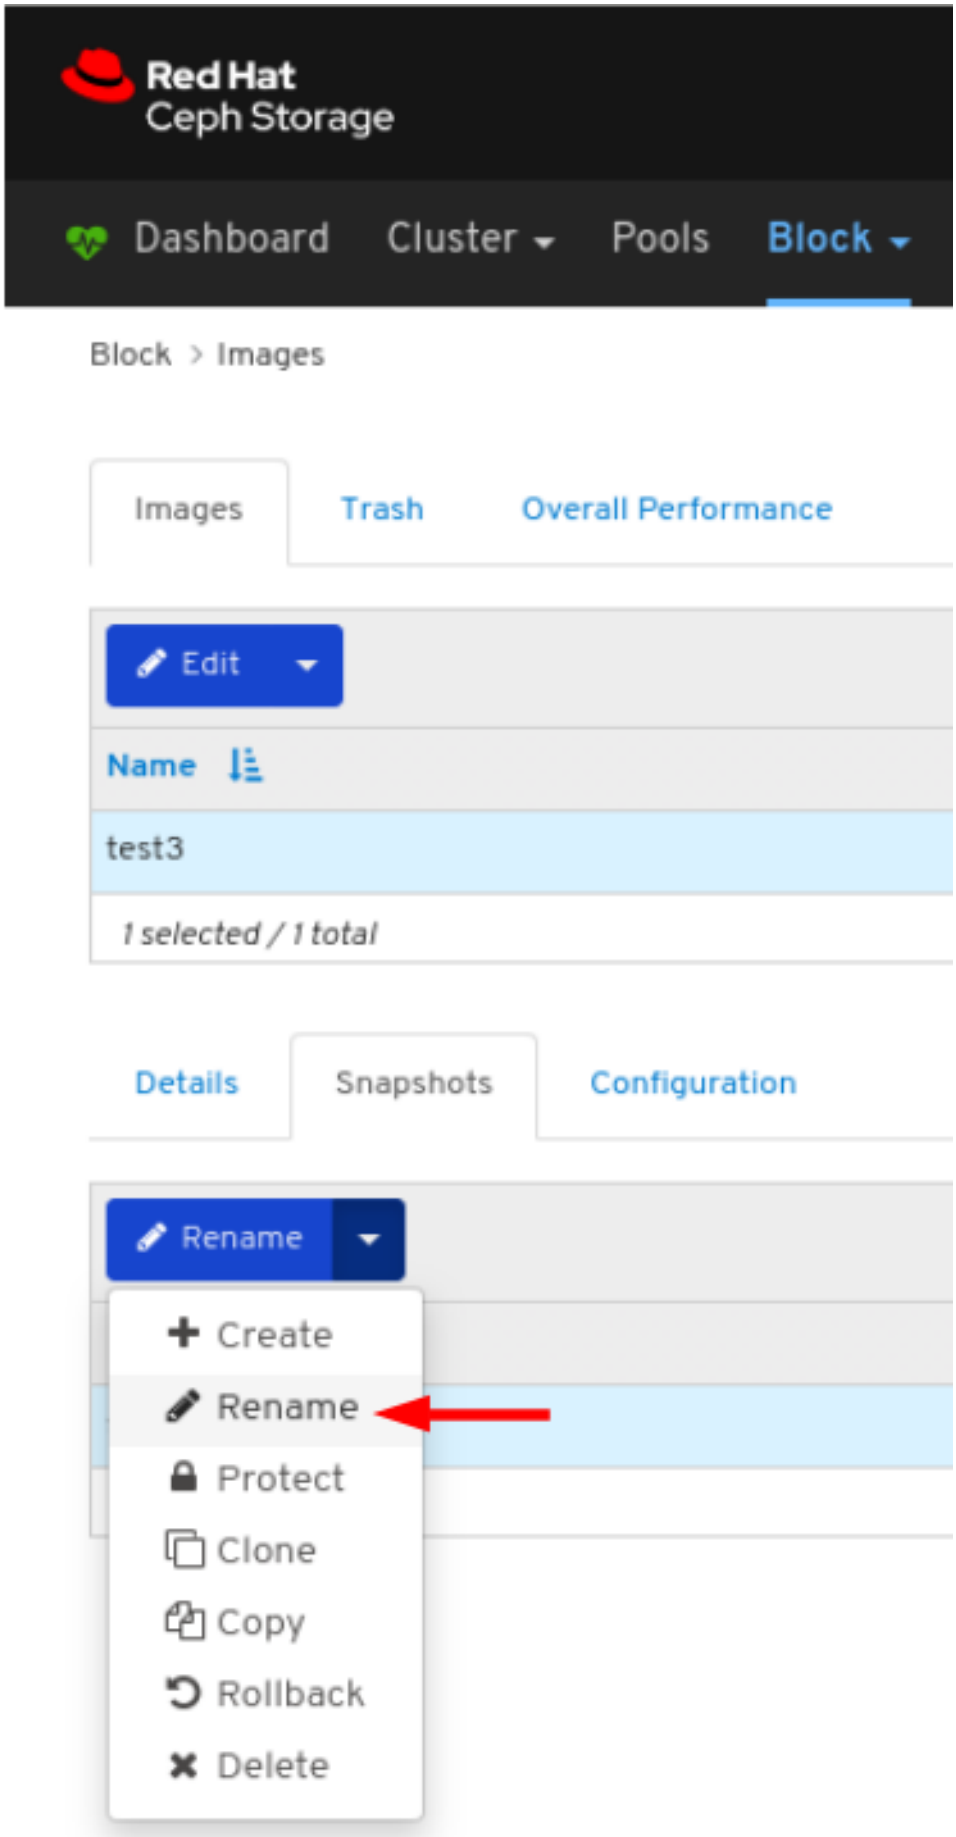

6. In the *RenameRBD Snapshot* dialog, enter the parameters and click the *RenameRBD Snapshot* button:

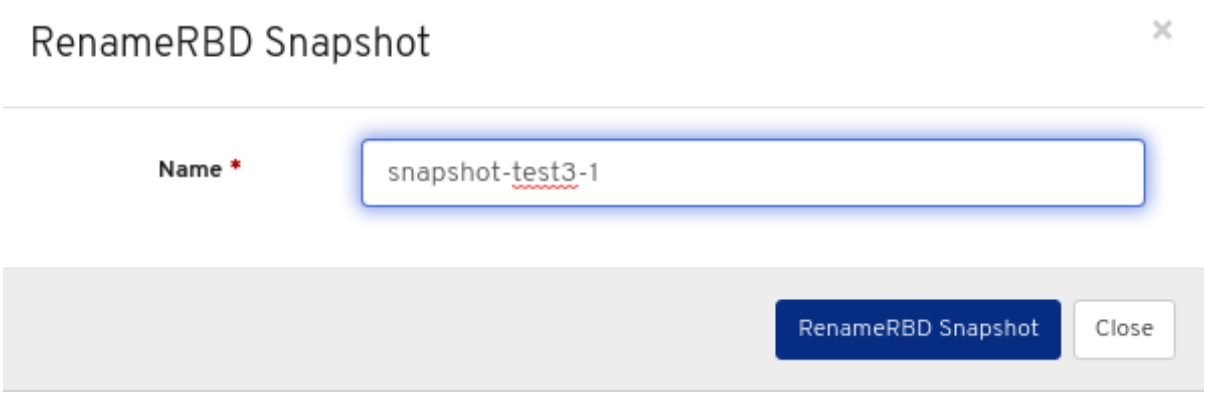

7. A notification towards the top right corner of the page indicates the snapshot of the image was updated successfully.

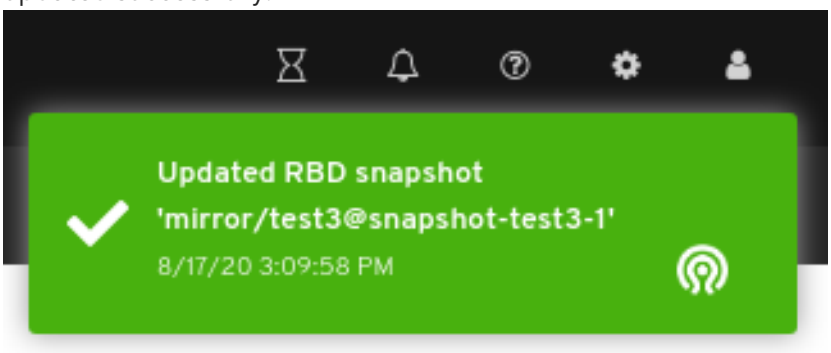

Additional Resources

- See the *Renaming a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide*] for more information.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.

9.2.12. Protecting snapshots of images

The dashboard allows you to protect snapshots of Ceph block device images. This is required when you need to clone the snapshots.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

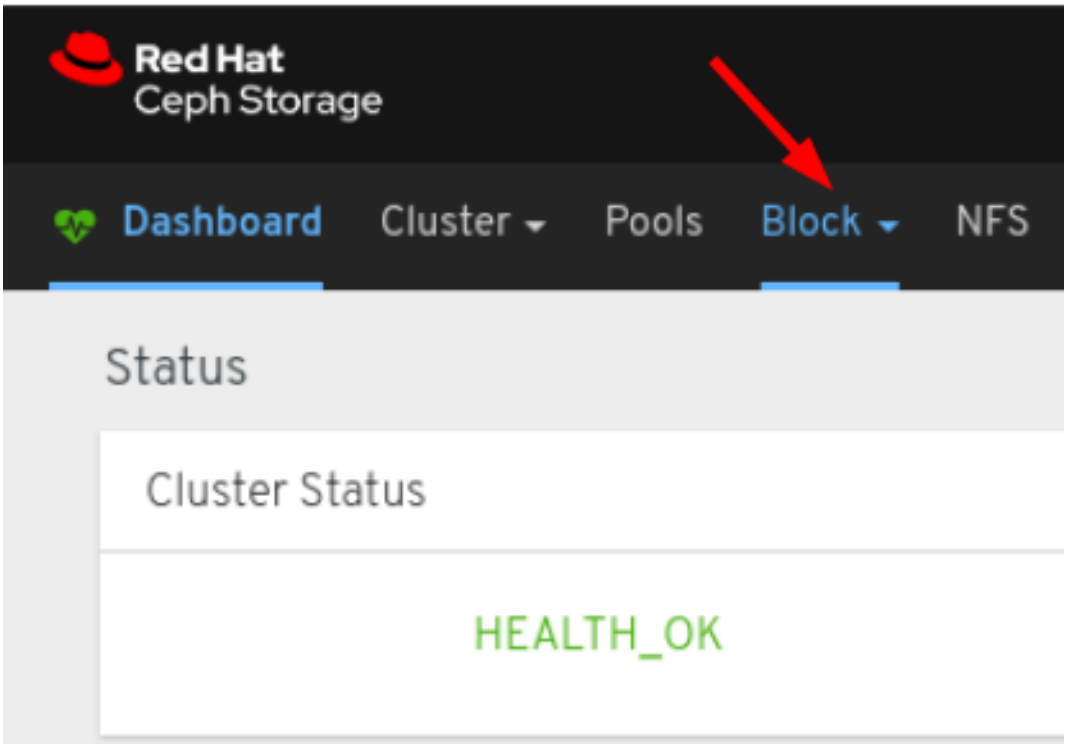

3. Select *Images* from the drop-down:

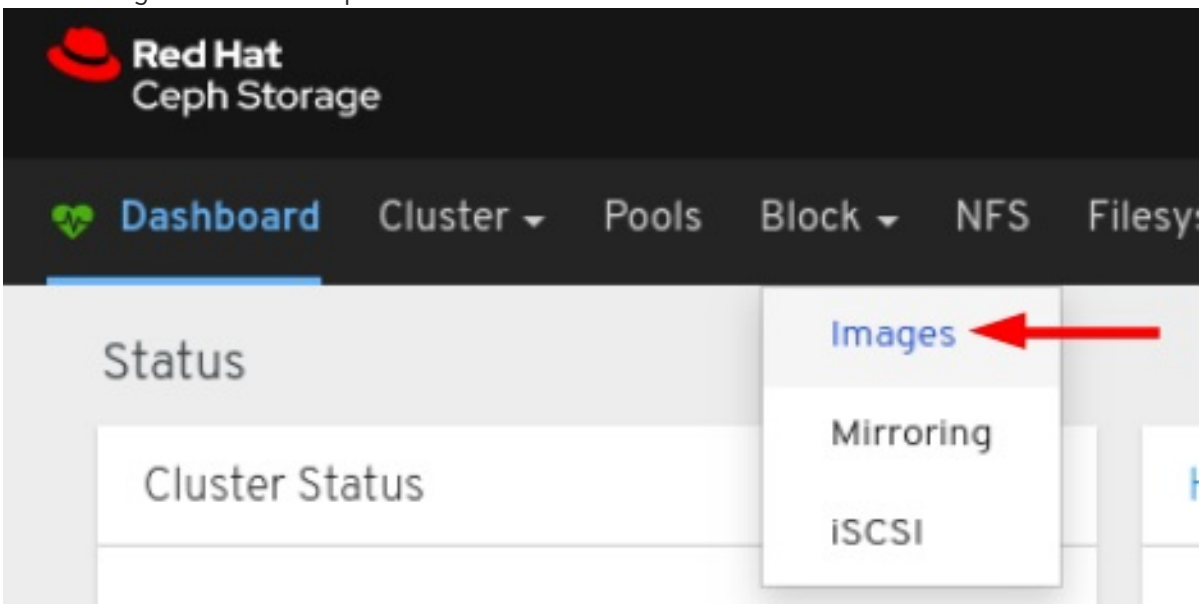

4. To protect the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

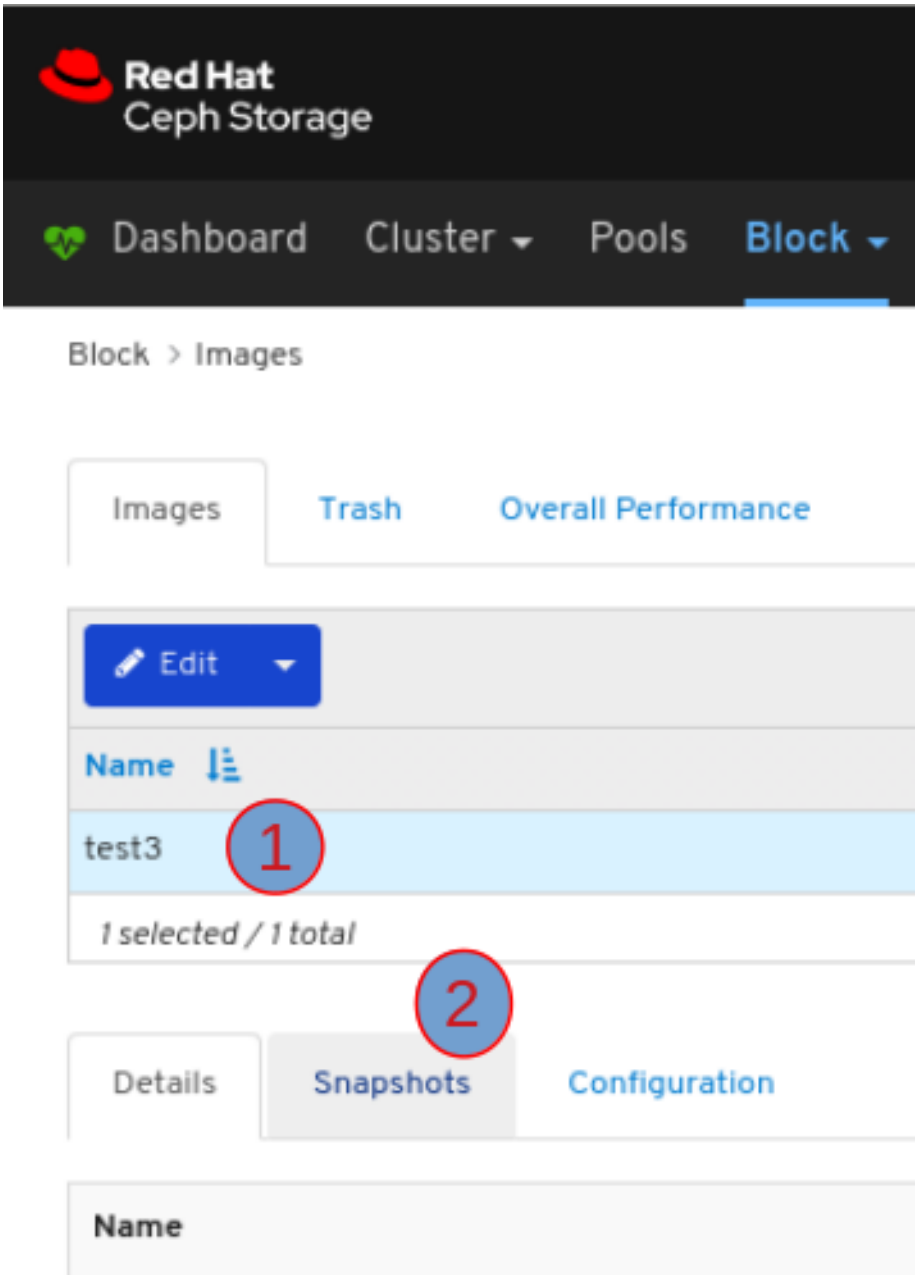

5. Select *Protect* in the the *Rename* drop-down:

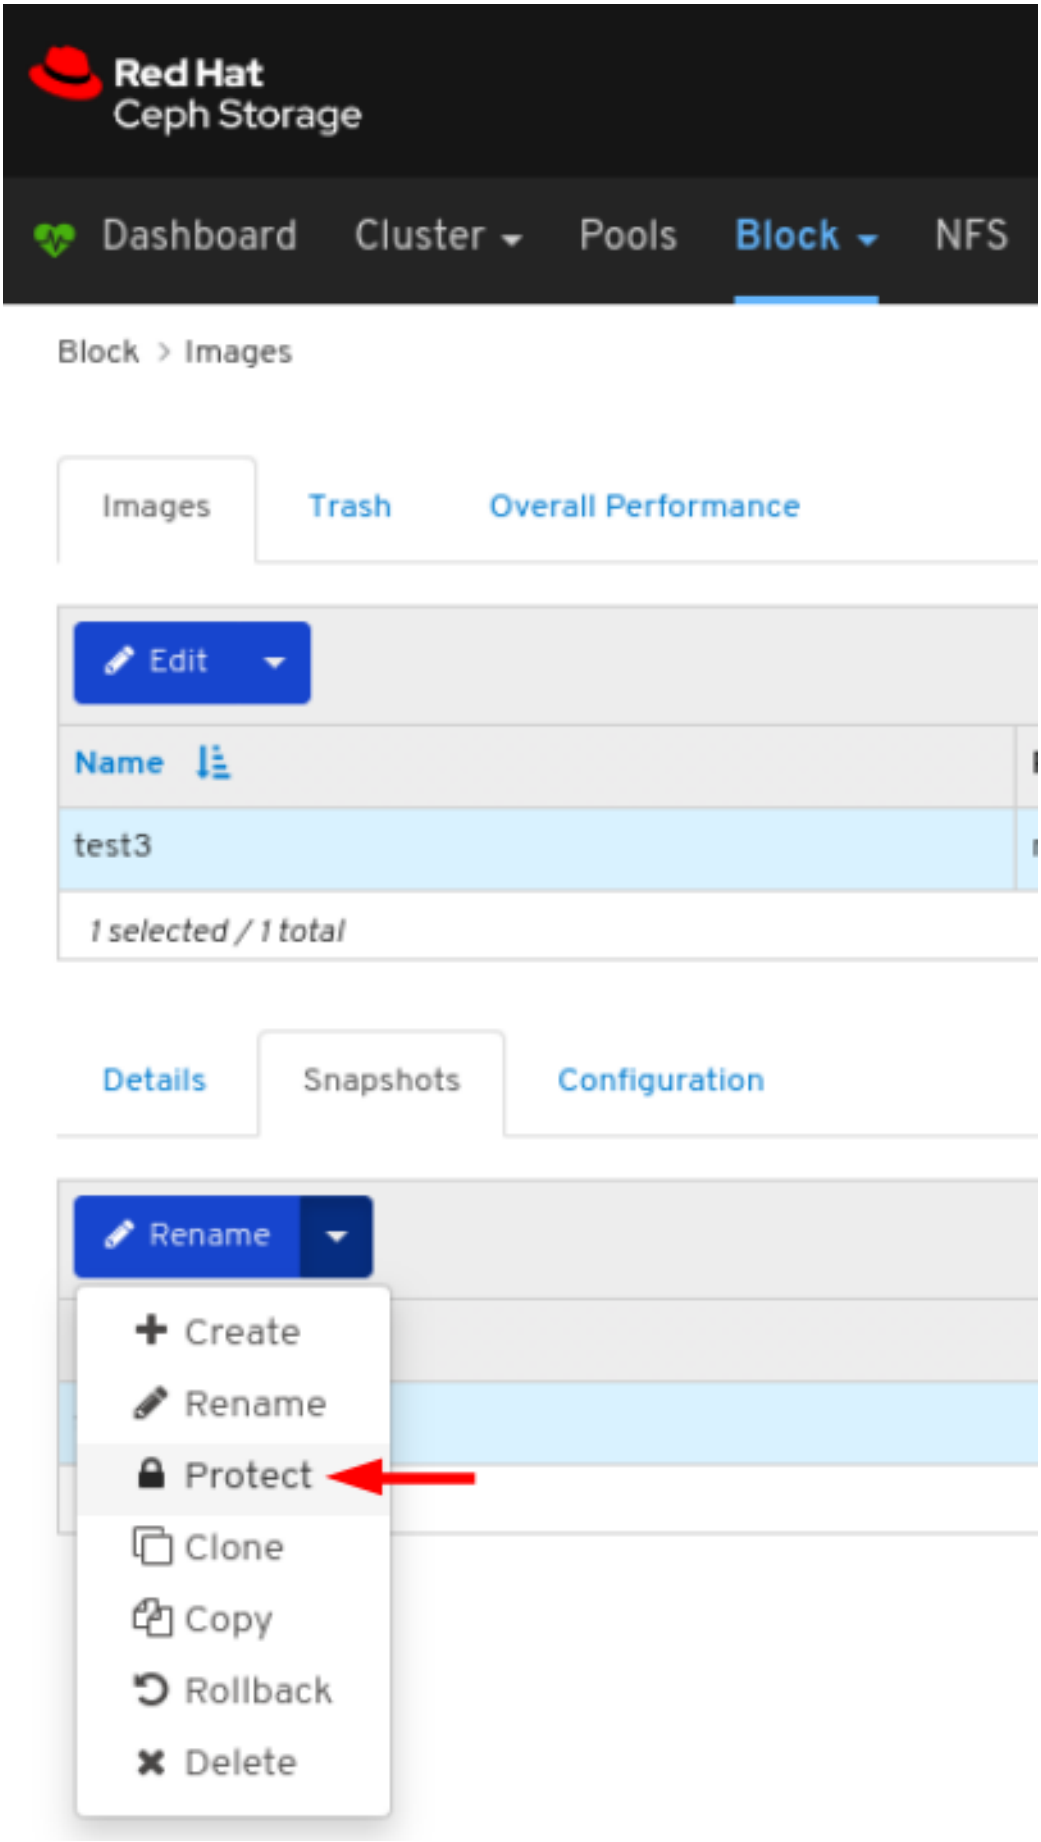

6. A notification towards the top right corner of the page indicates the snapshot of the image was updated successfully.

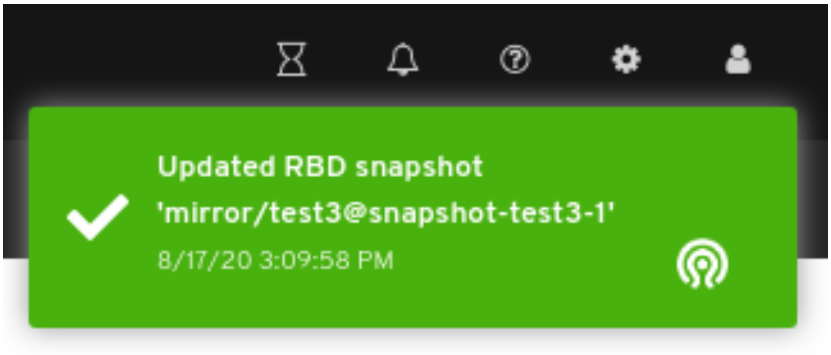

Additional Resources

- See the *Protecting a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide*] for more information.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.

9.2.13. Cloning snapshots of images

The dashboard allows you to clone snapshots of Ceph block device images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

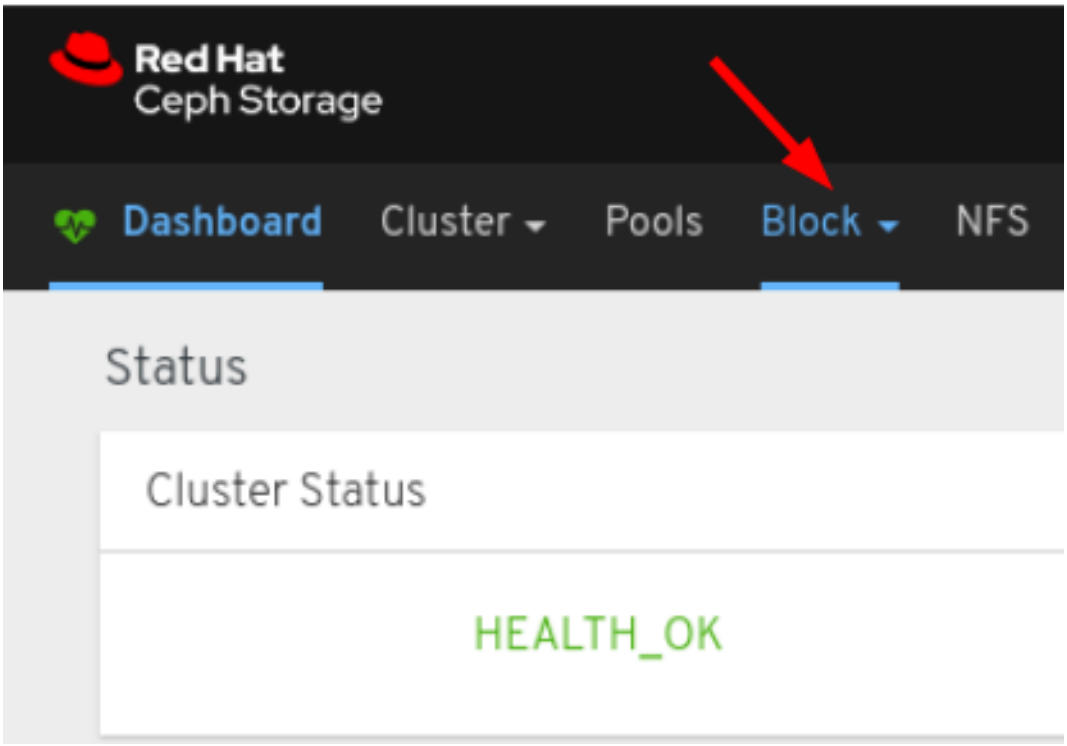

3. Select *Images* from the drop-down:

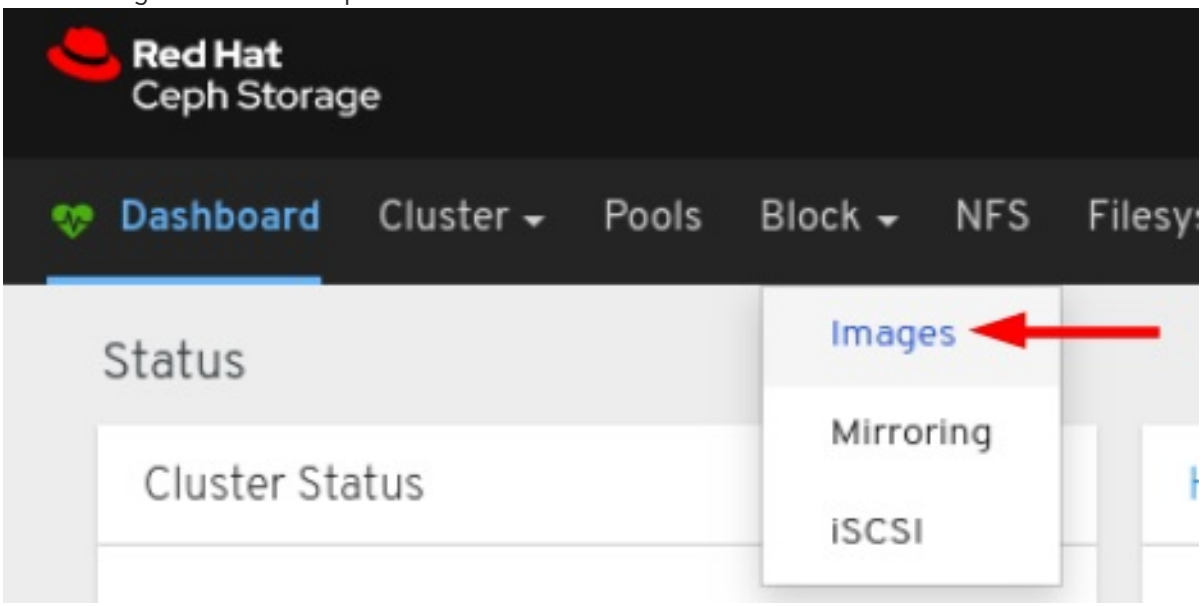

4. To clone the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

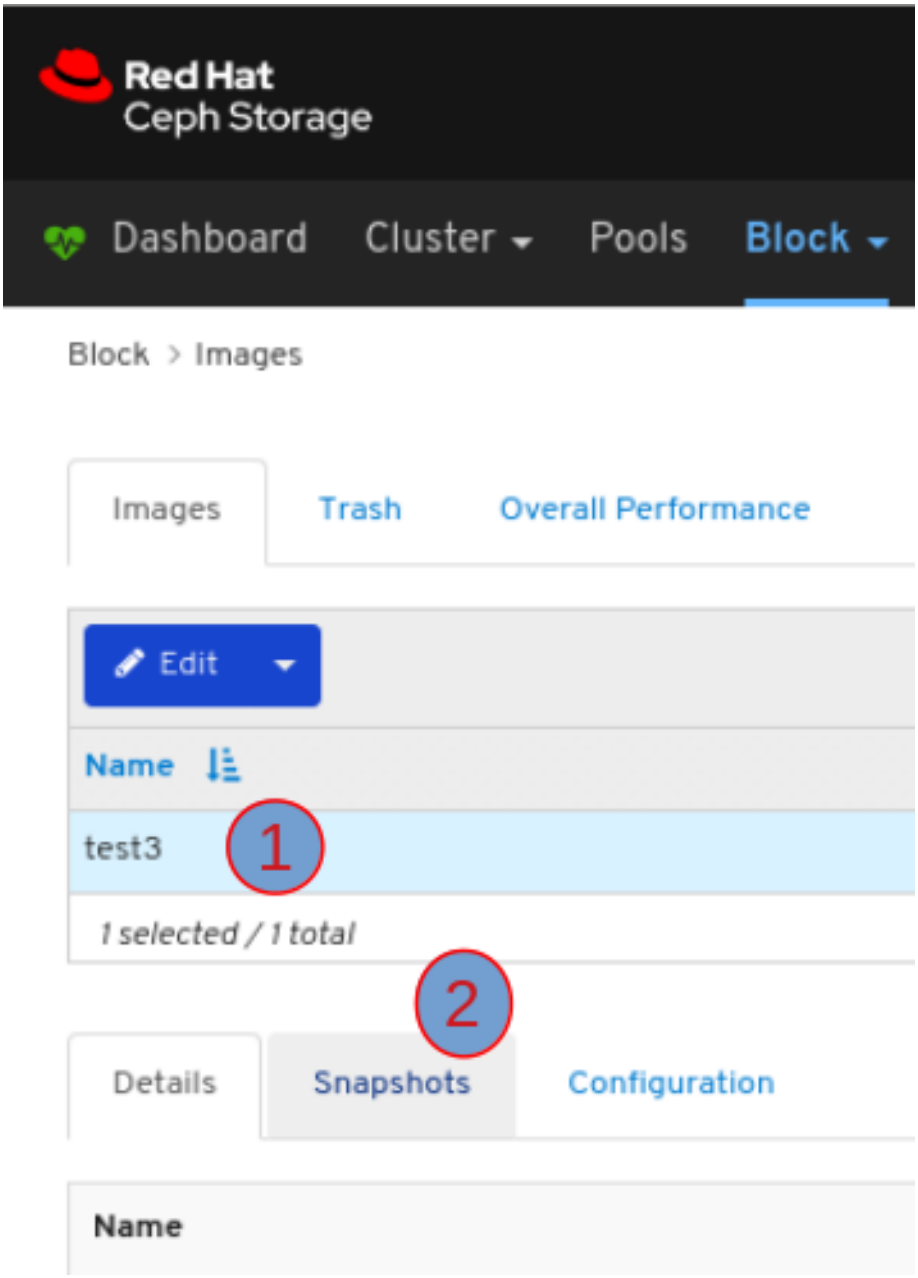

5. Select *Clone* in the the *Rename* drop-down:

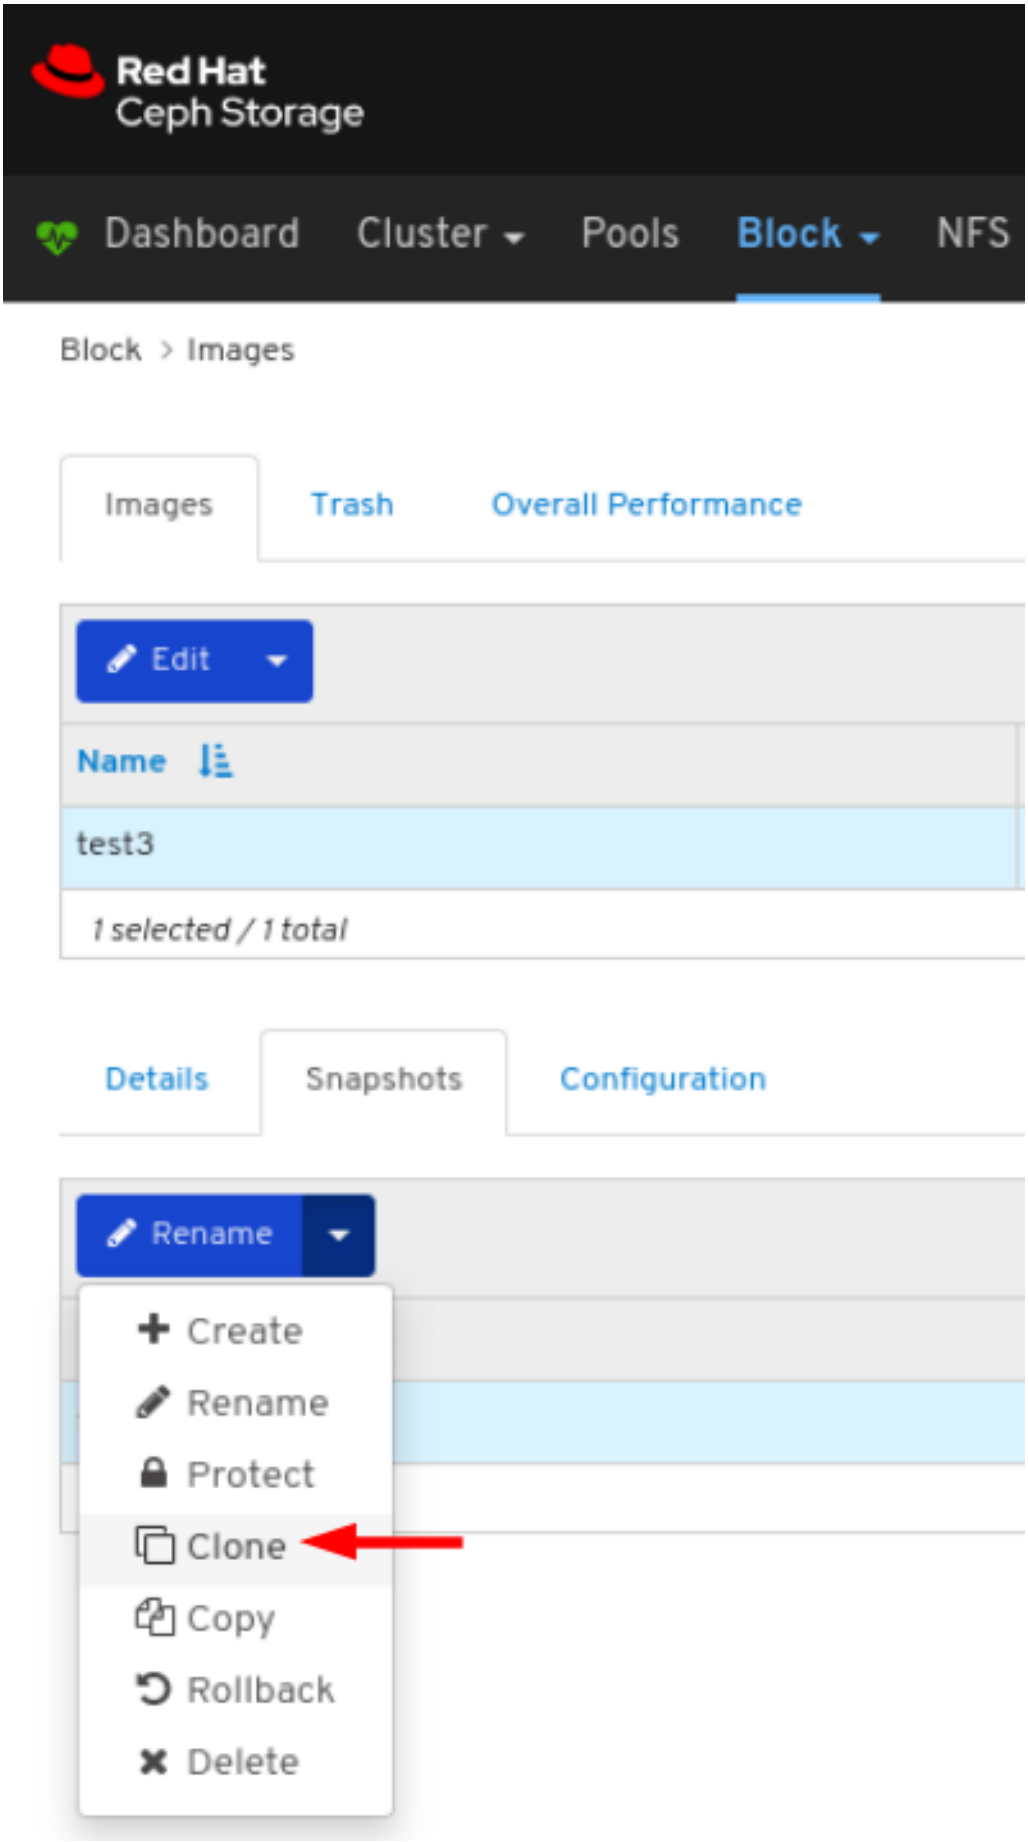

6. In the *CloneRBD* dialog, enter the parameters and click the *CloneRBD* button:

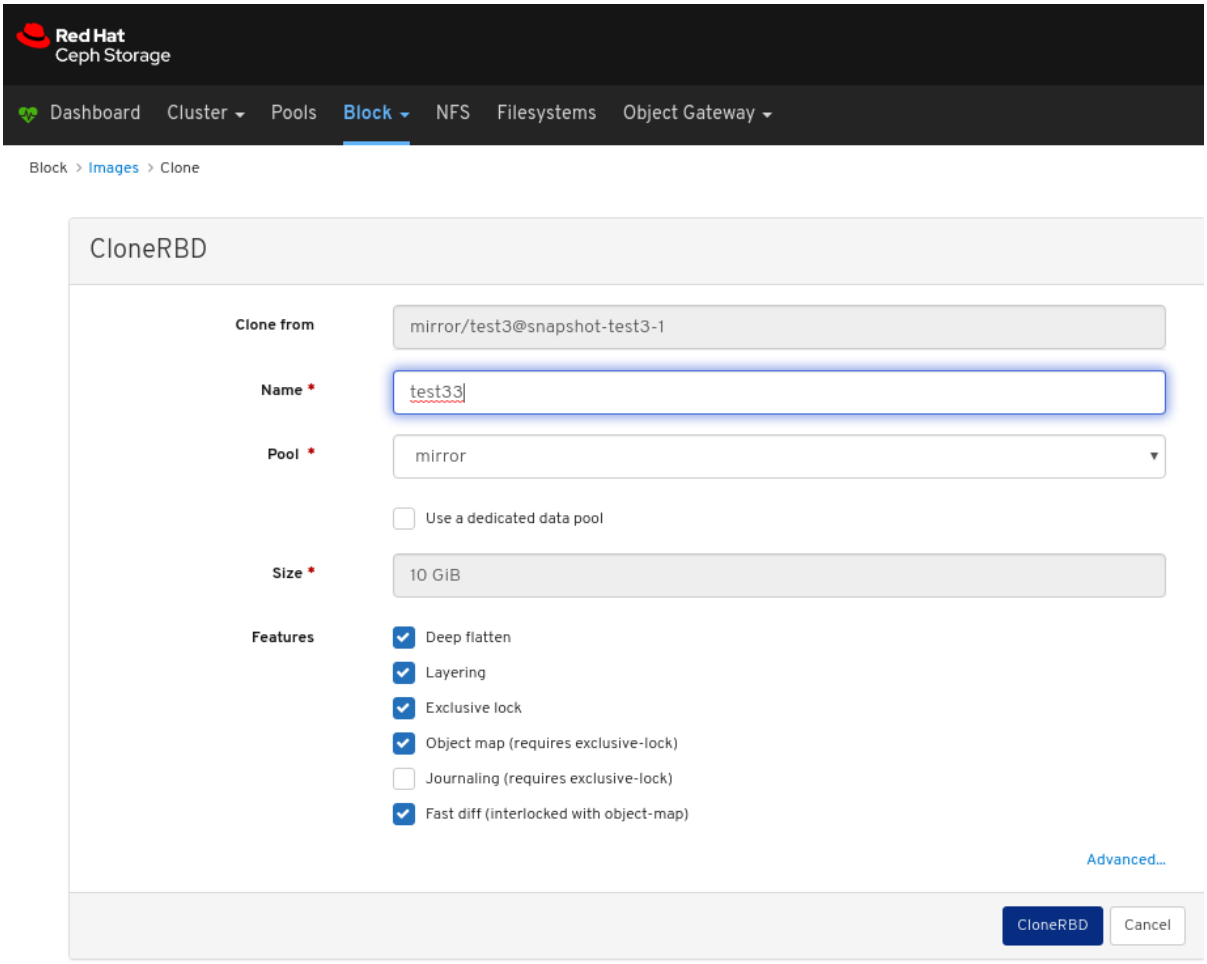

7. A notification towards the top right corner of the page indicates the snapshot of the image was cloned successfully.

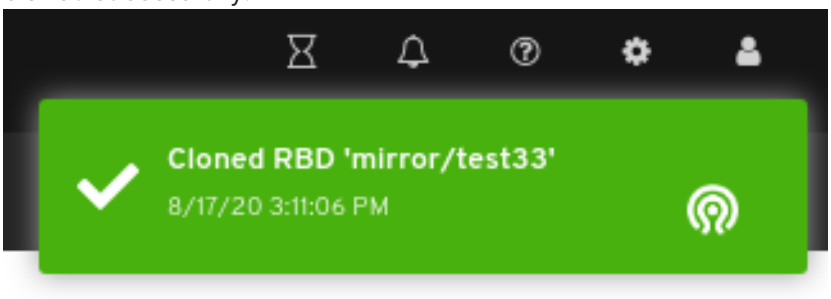

Additional Resources

- See the *Renaming a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide*] for more information.
- See the *Protecting a Block device Snapshot* section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.
- See the *Protecting snapshots of images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

9.2.14. Copying snapshots of images

The dashboard allows you to copy snapshots of Ceph block device images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

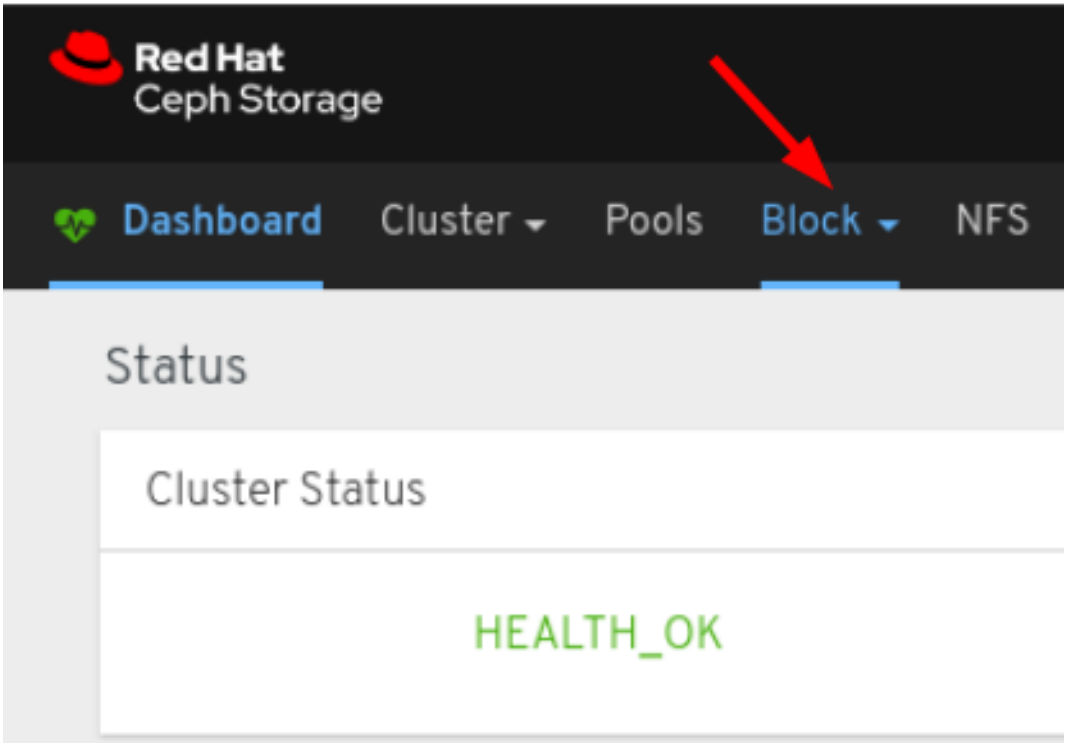

3. Select *Images* from the drop-down:

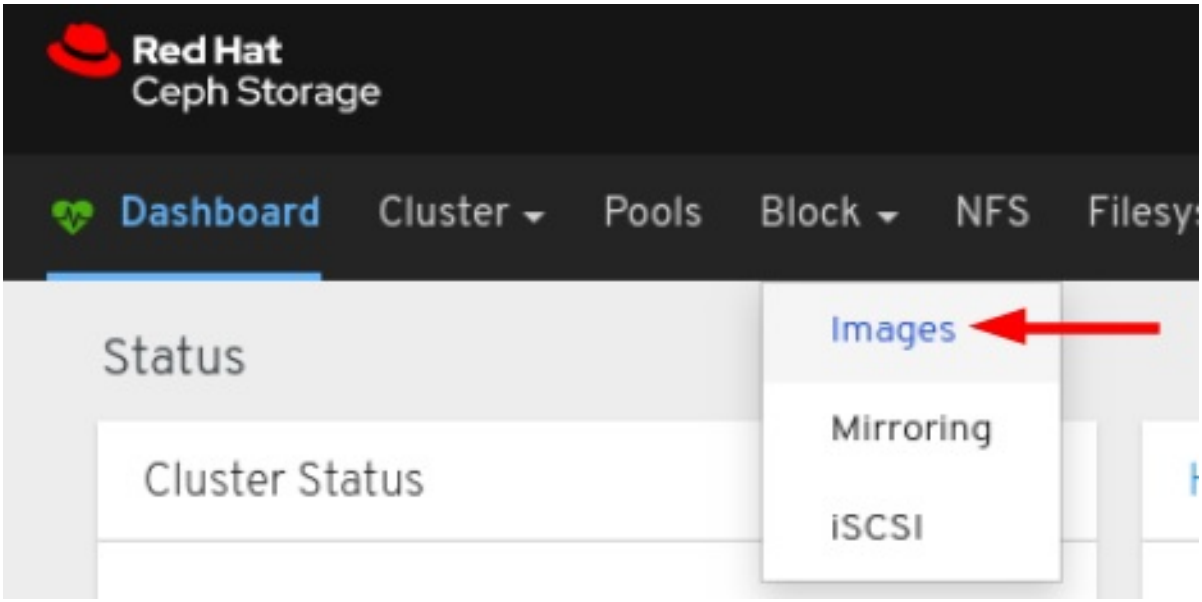

4. To copy the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

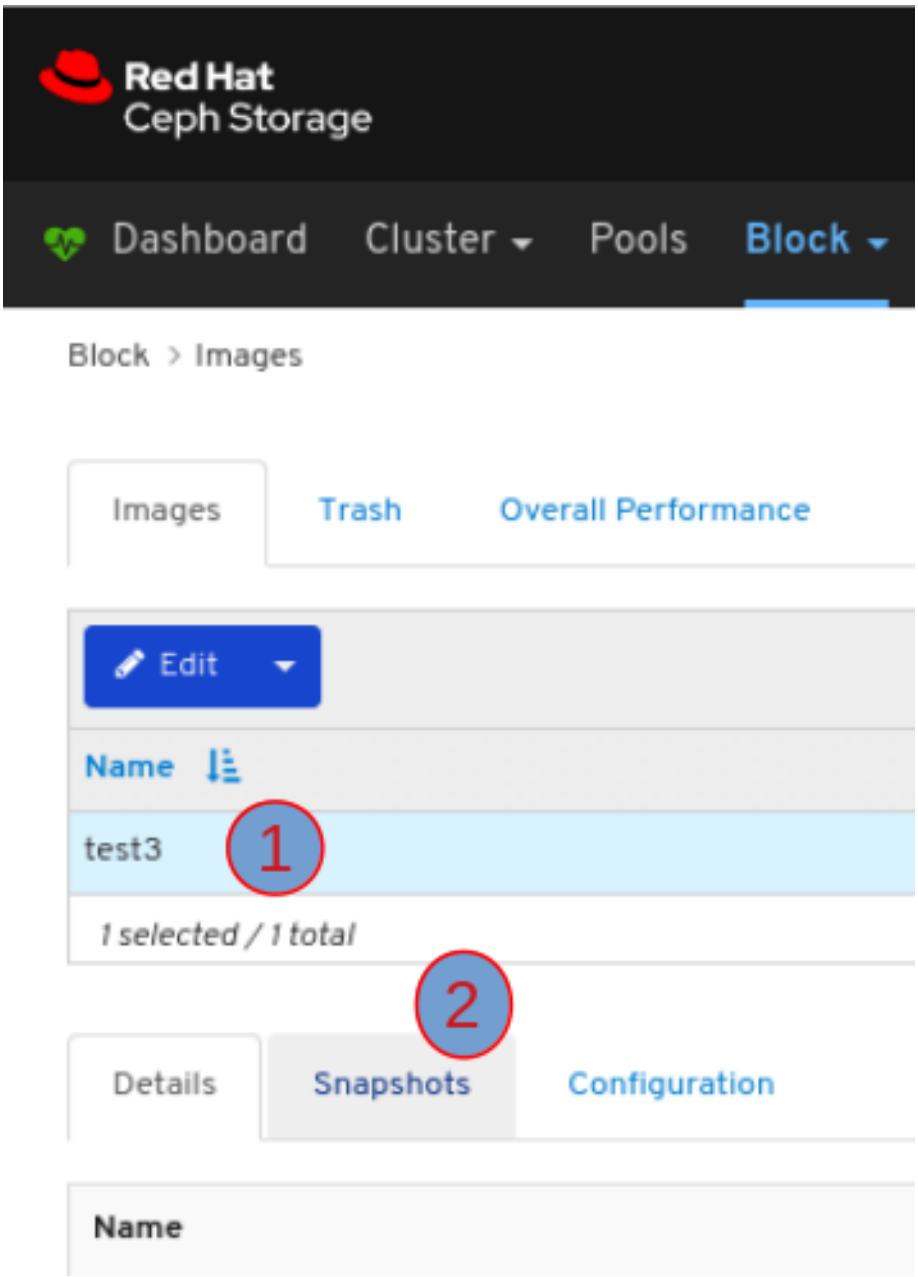

5. Select *Copy* in the the *Rename* drop-down:

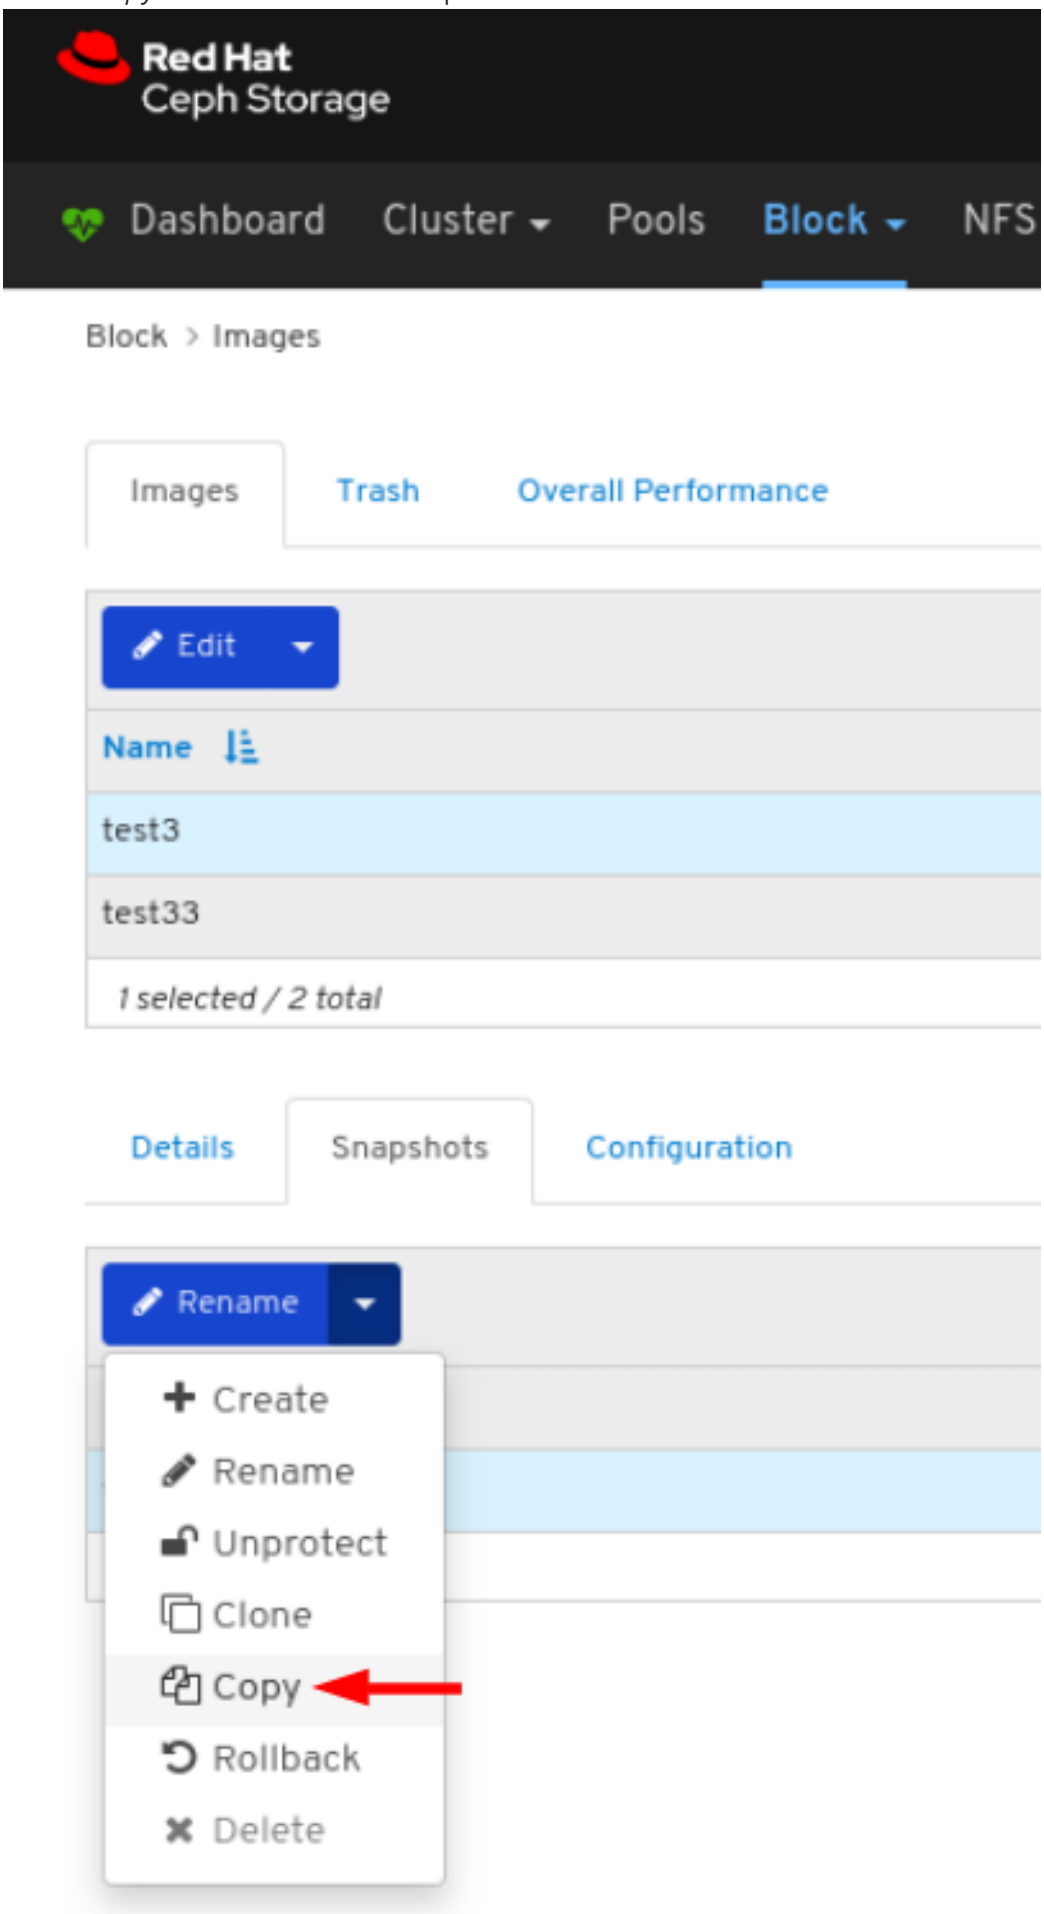

6. In the *CopyRBD* dialog, enter the parameters and click the *CopyRBD* button:

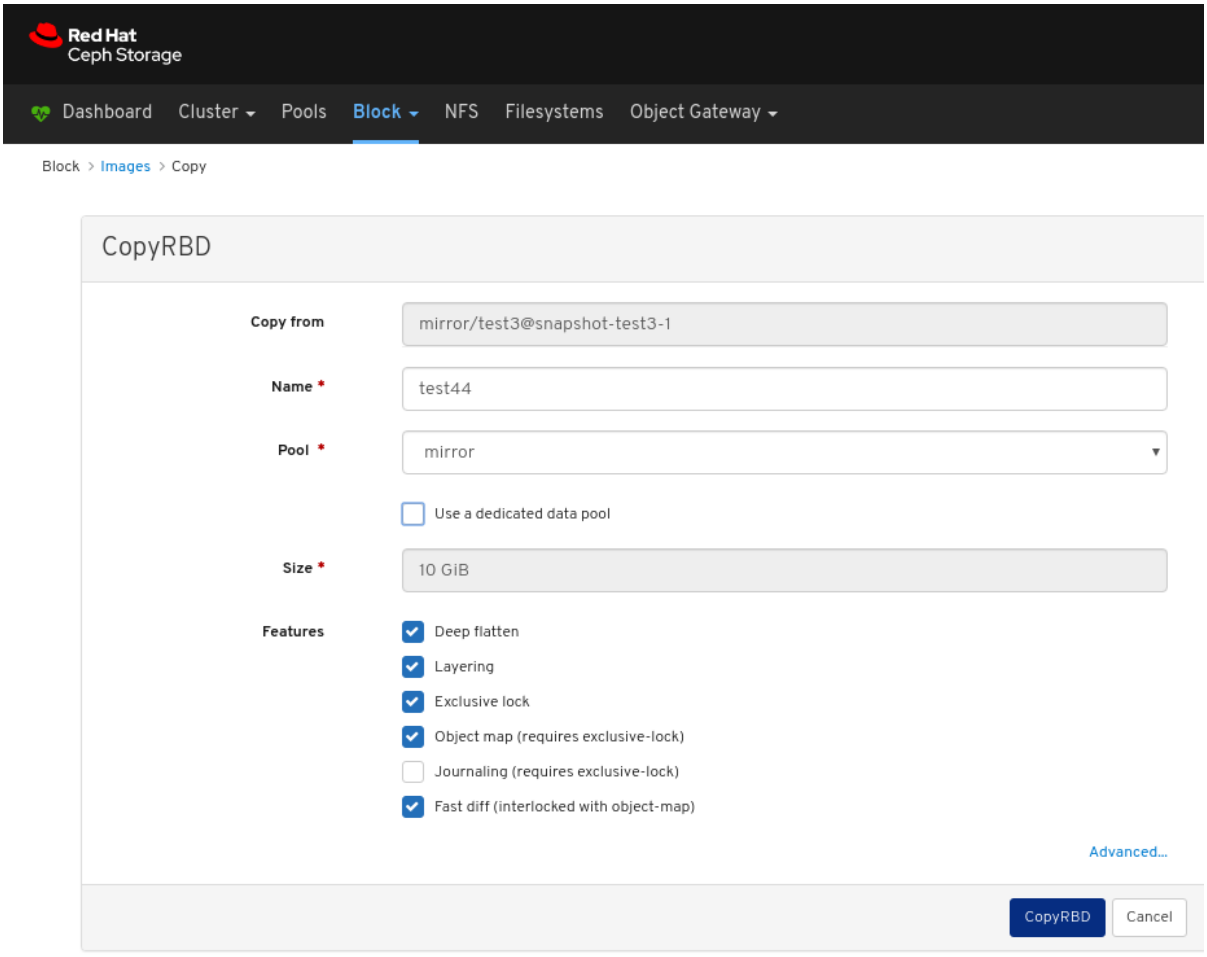

7. A notification towards the top right corner of the page indicates the snapshot of the image was updated successfully.

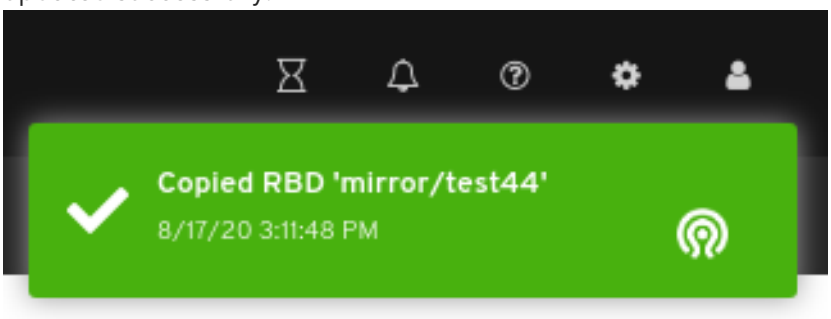

Additional Resources

- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.

9.2.15. Rolling back snapshots of images

The dashboard allows you to rollback snapshots of Ceph block device images.

Prerequisites

- A running Red Hat Ceph Storage cluster.

- Dashboard is installed.
- A pool with the *rbd* application enabled is created.
- An image is created.
- A snapshot of the image is created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

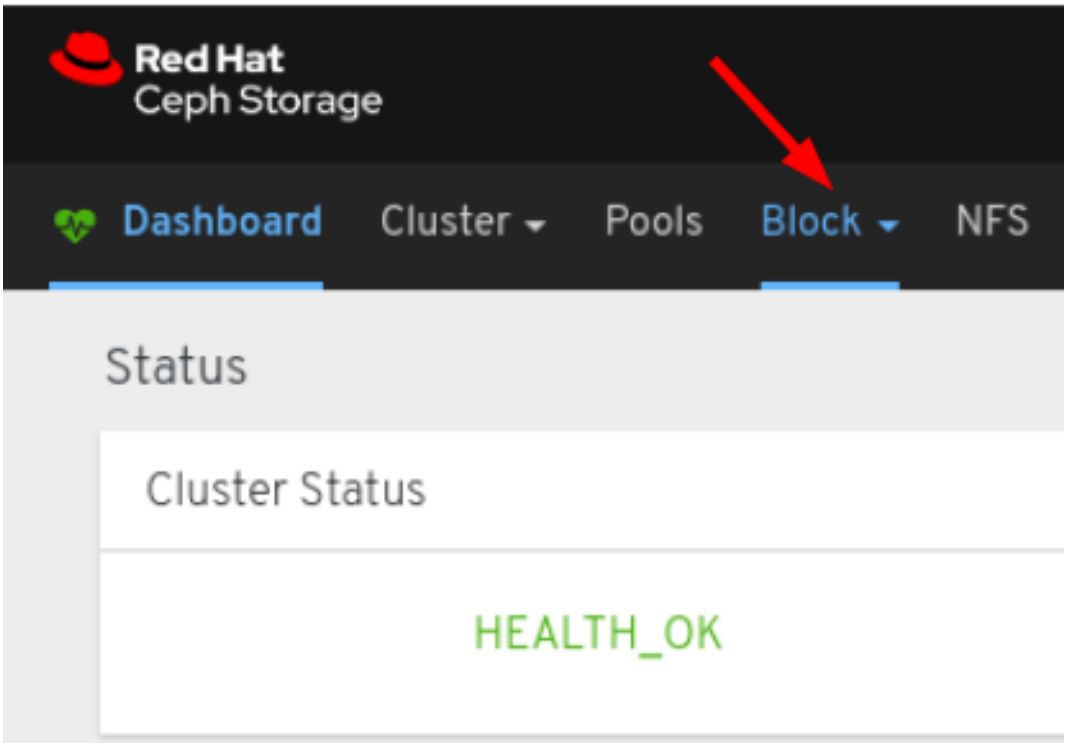

3. Select *Images* from the drop-down:

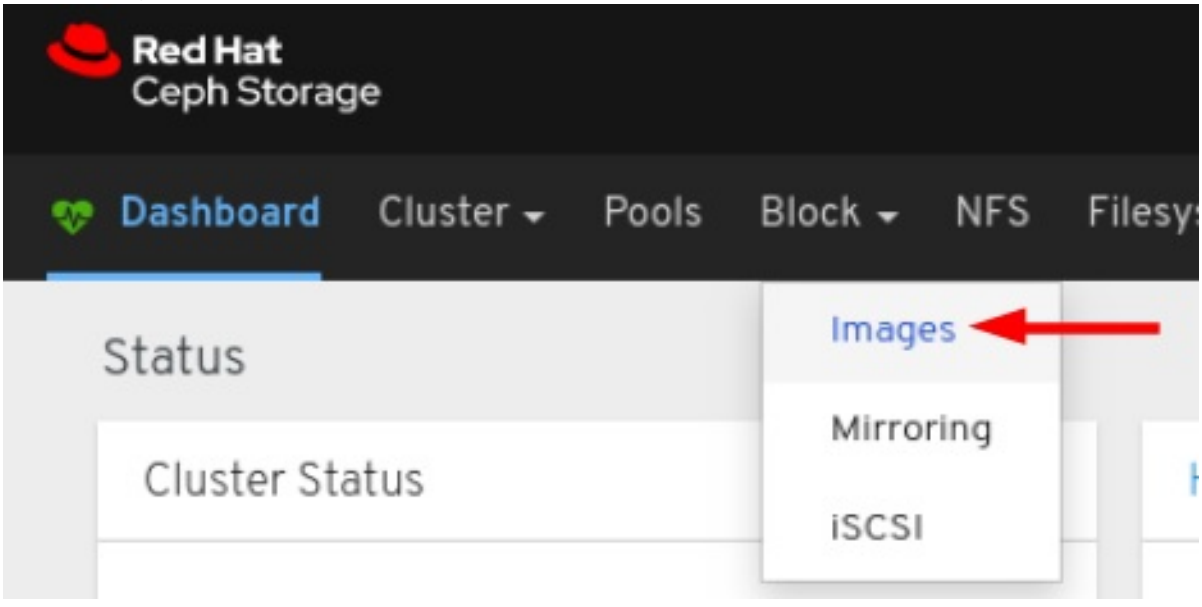

4. To rollback the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

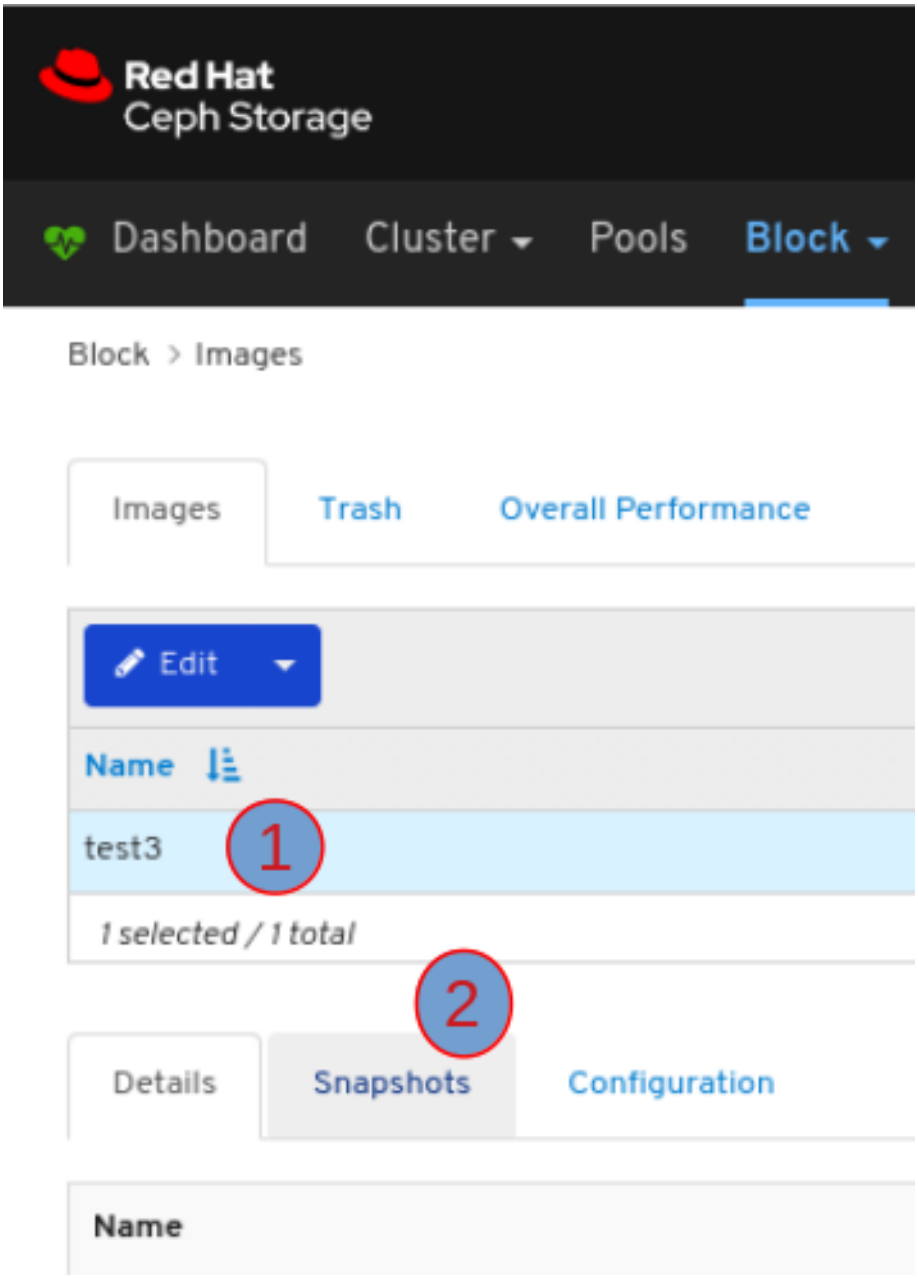

5. Select *Rollback* in the the *Rename* drop-down:

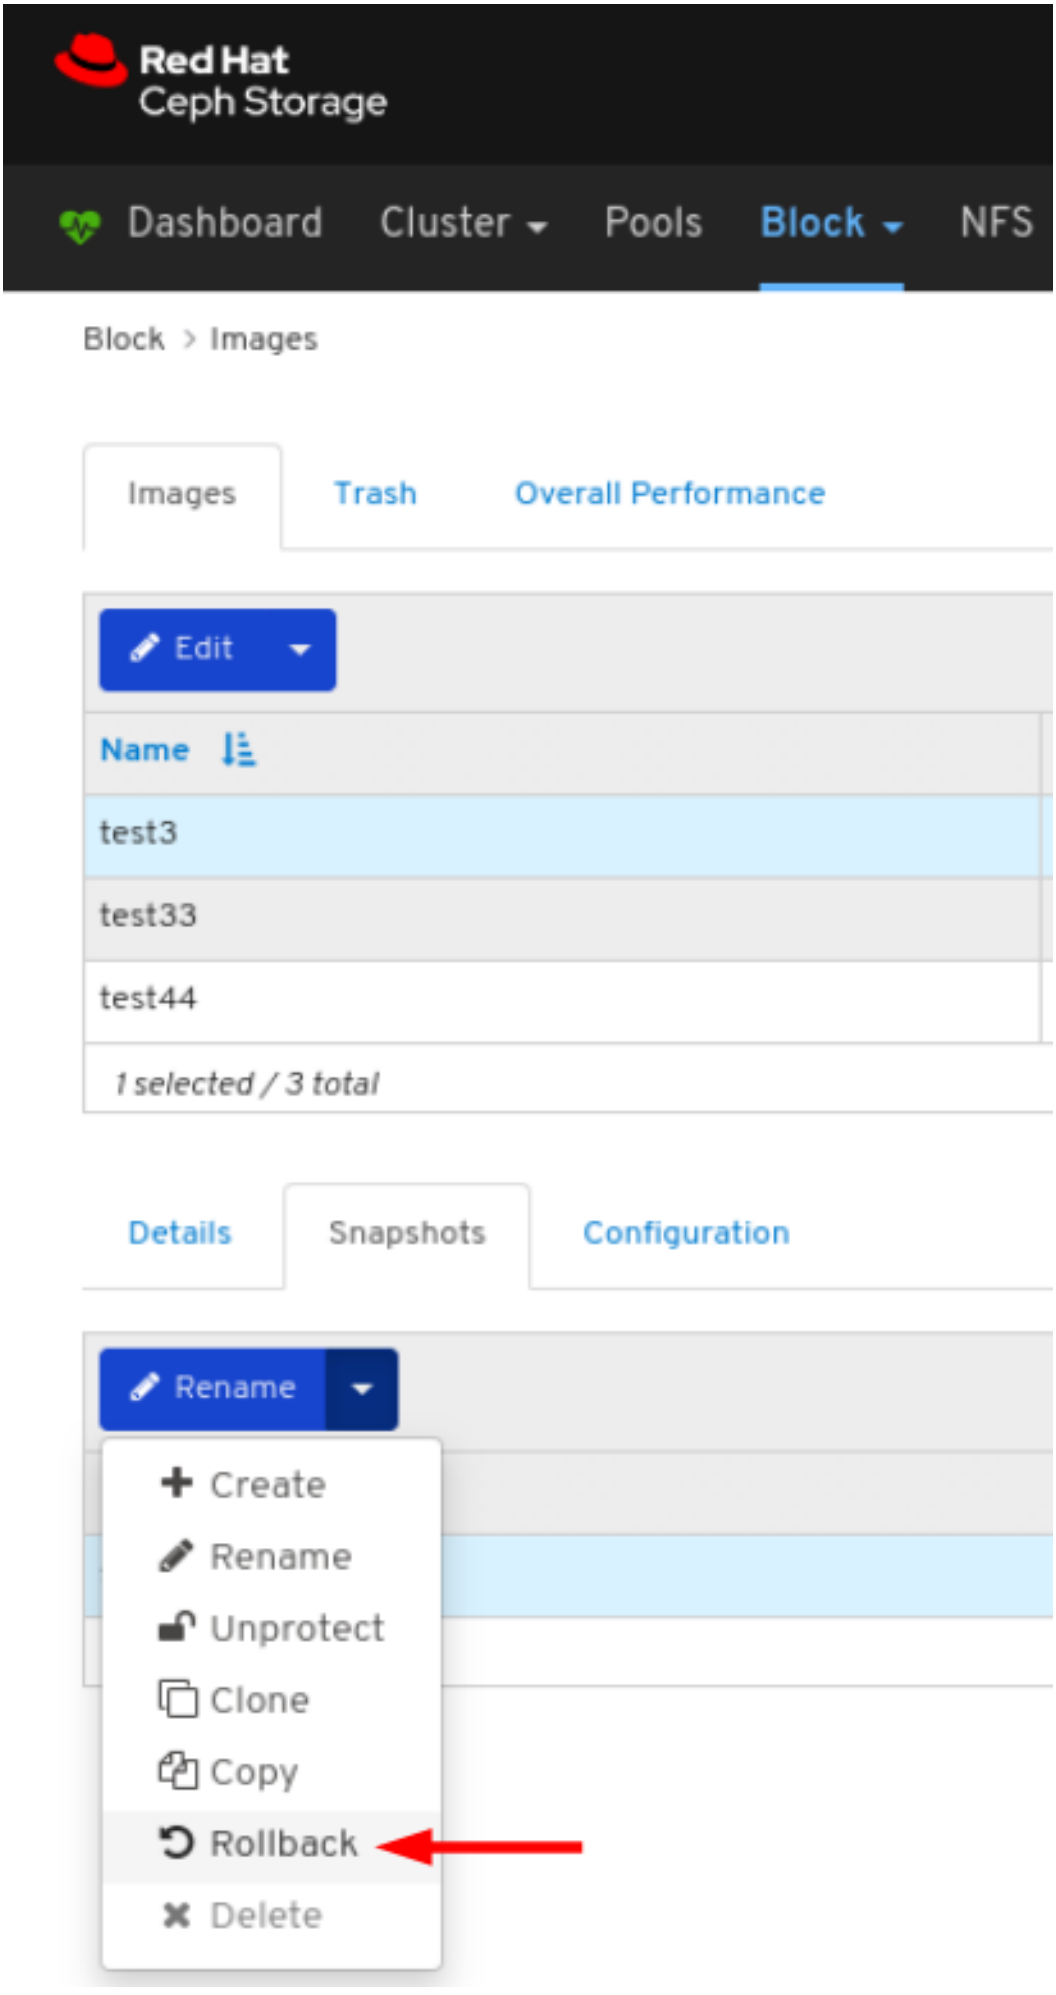
Red Hat
Ceph Storage
Dashboard
Cluster
Pools
Block
NFS
Block > Images
Images
Trash
Overall Performance
Edit
Name
test3
test33
test44
1 selected / 3 total
Details
Snapshots
Configuration
Rename
Create
Rename
Unprotect
Clone
Copy
Rollback
Delete

6. In the *RBD snapshot rollback* window, click the *Rollback* button:

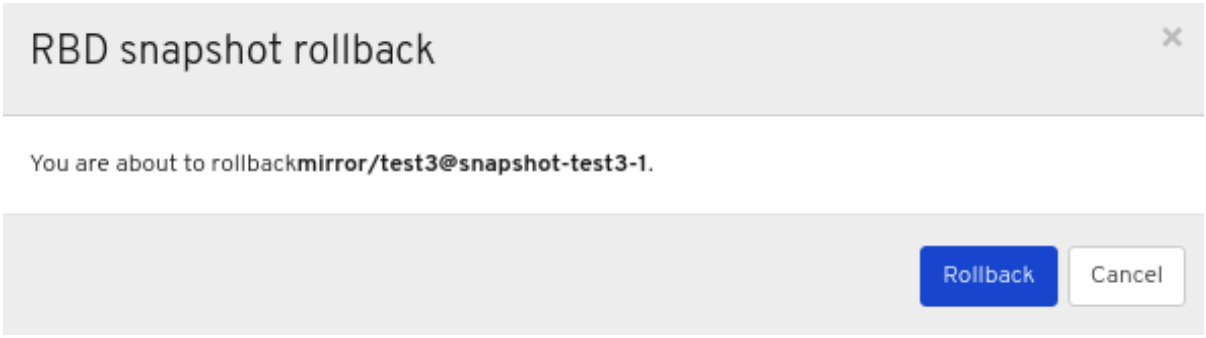

7. A notification towards the top right corner of the page indicates the snapshot of the image was rolled back successfully.

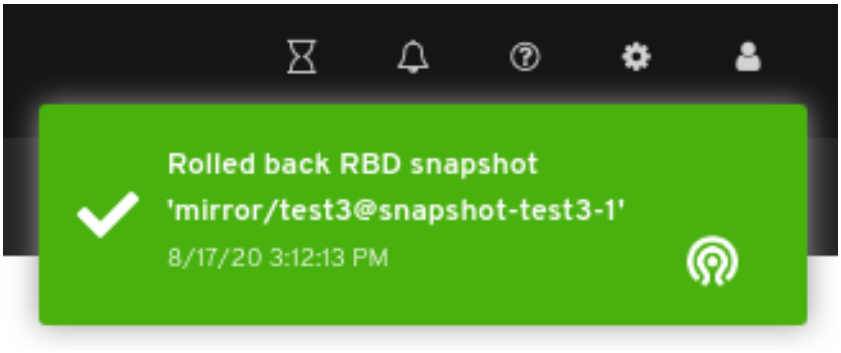

Additional Resources

- See the *Rolling a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide*] for more information.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

9.2.16. Unprotecting snapshots of images

The dashboard allows you to unprotect snapshots of Ceph block device images. This is required when you need to delete the snapshots.

Prerequisites

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

Procedure

1. Log in to the Dashboard.

2. On the navigation bar, click *Block*:

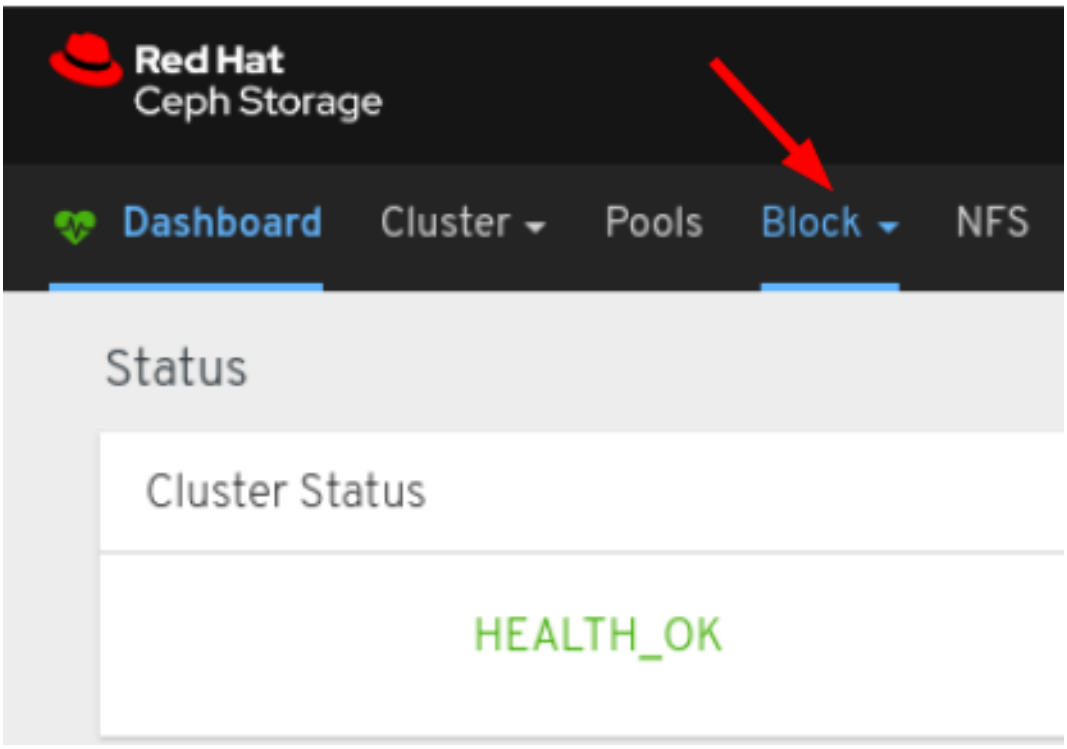

3. Select *Images* from the drop-down:

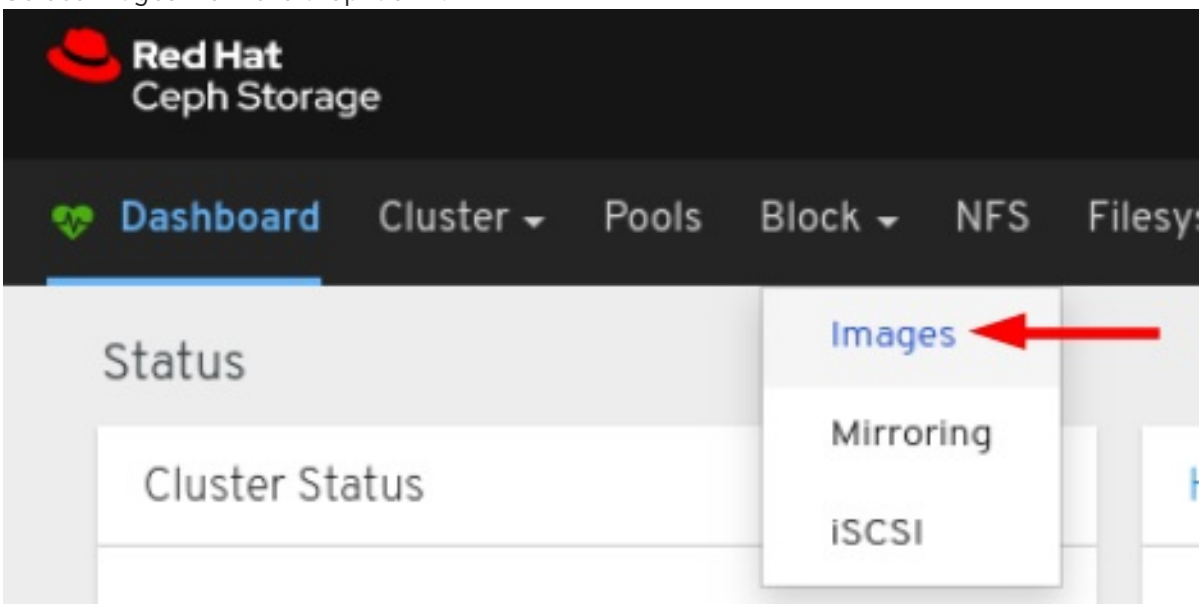

4. To unprotect the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

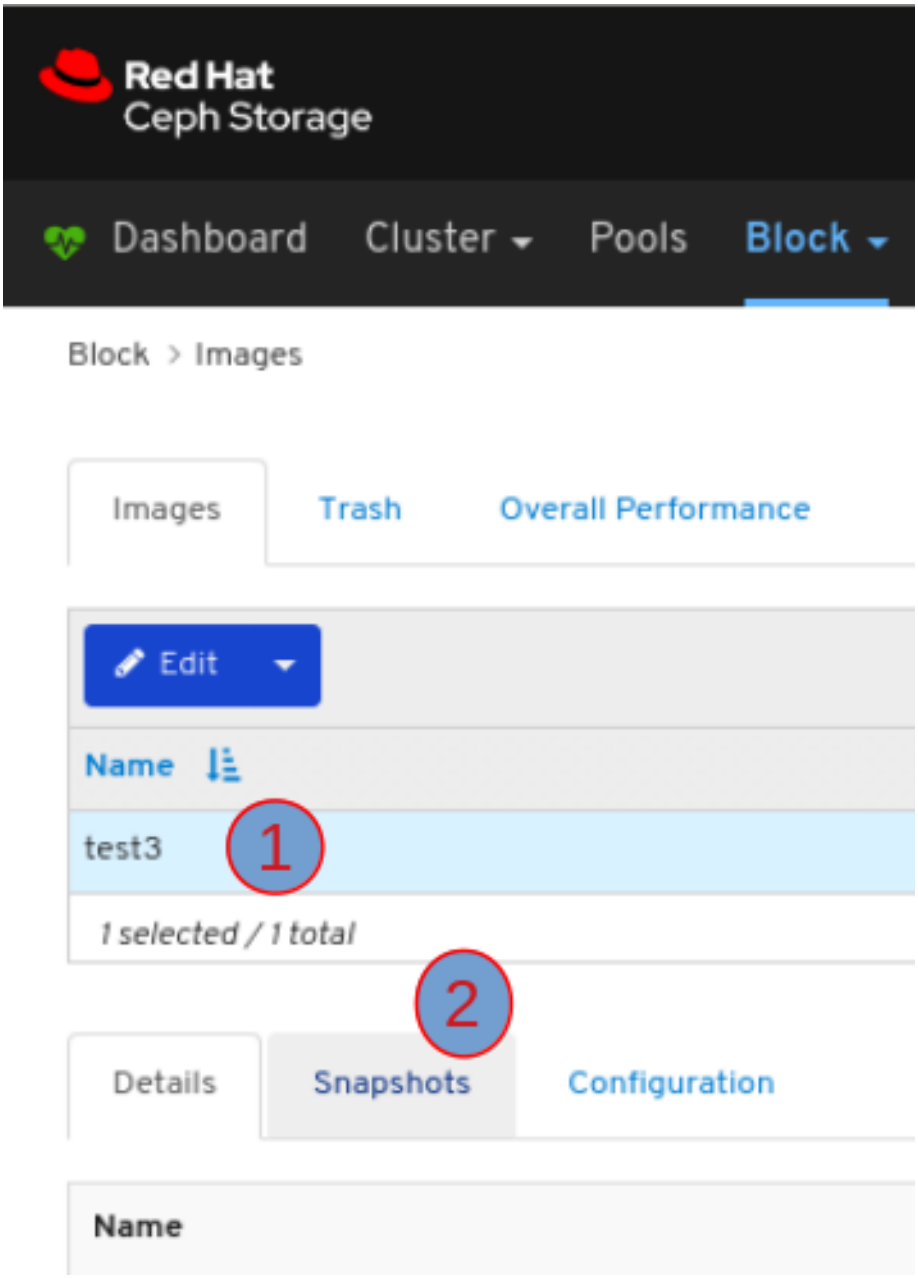

5. Select *UnProtect* in the the *Rename* drop-down:

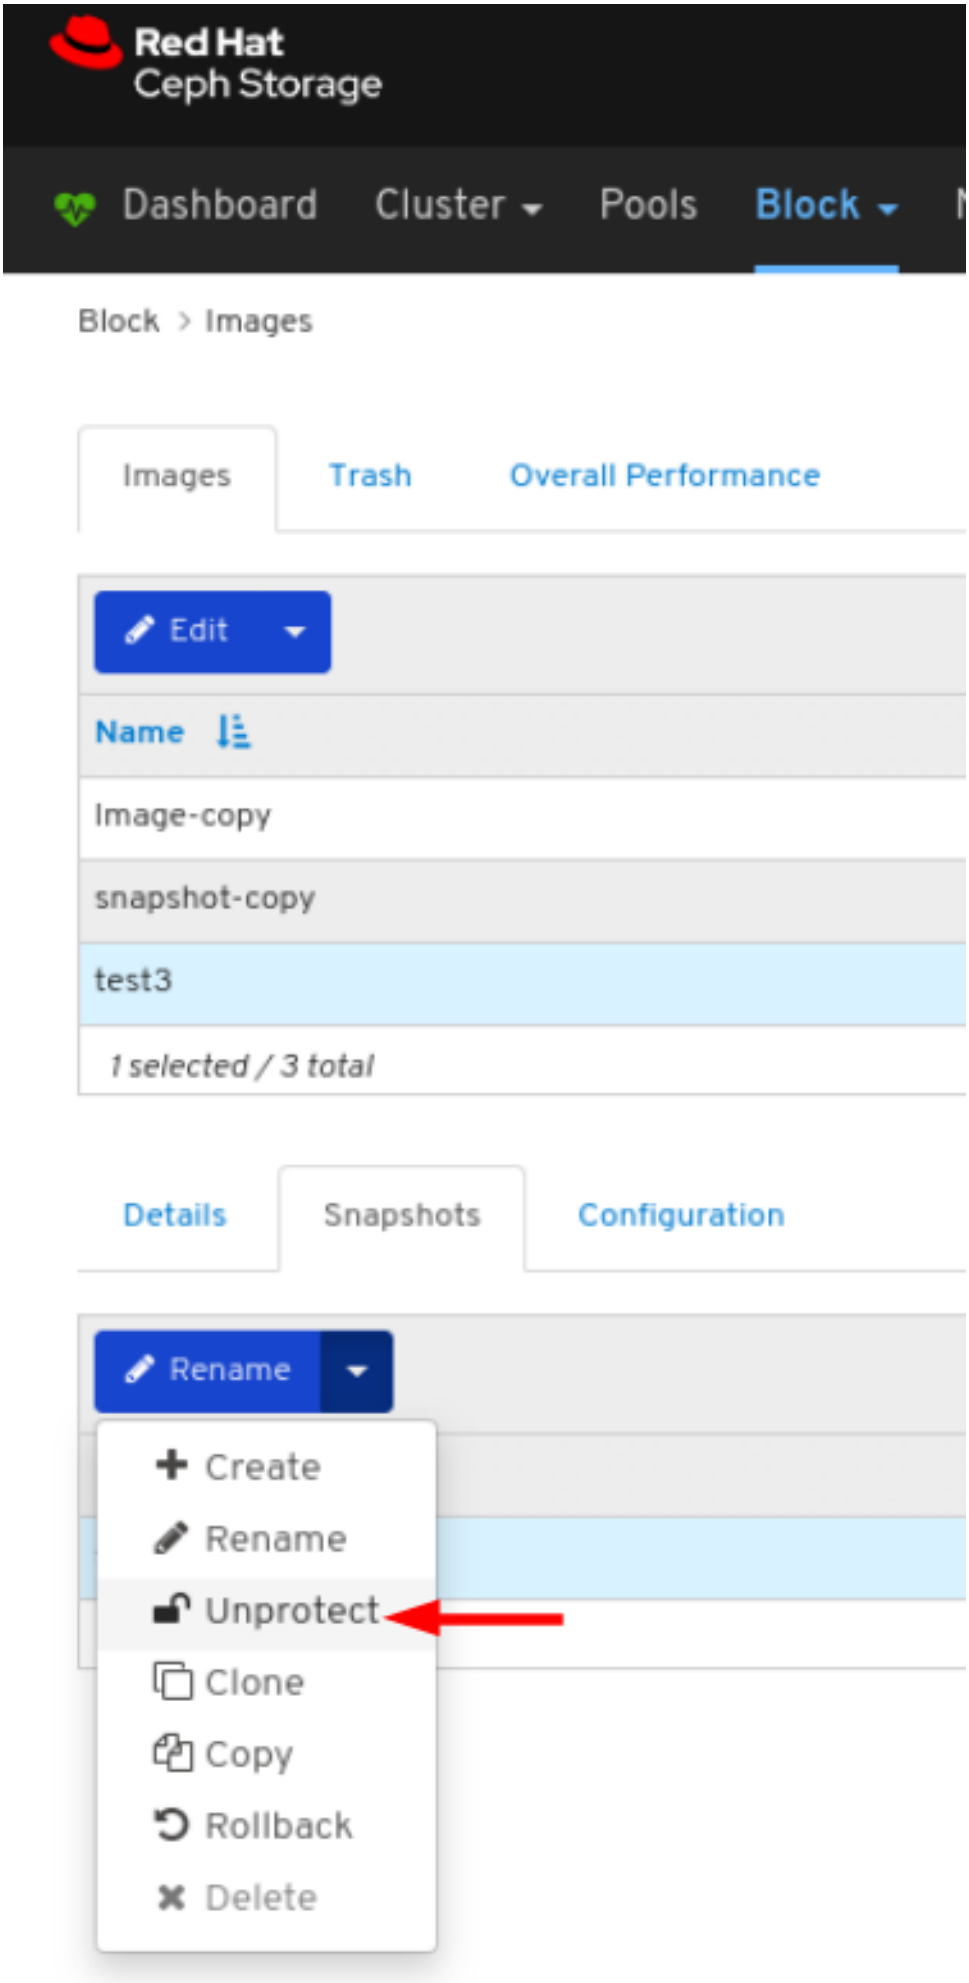
Red Hat
Ceph Storage
Dashboard
Cluster
Pools
Block
Block > Images
Images
Trash
Overall Performance
Edit
Name
Image-copy
snapshot-copy
test3
1 selected / 3 total
Details
Snapshots
Configuration
Rename
Create
Rename
Unprotect
Clone
Copy
Rollback
Delete

6. A notification towards the top right corner of the page indicates the snapshot of the image was updated successfully.

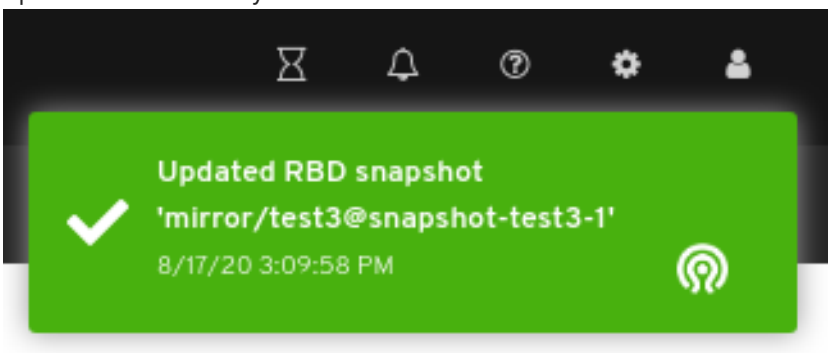

Additional Resources

- See the *Unprotecting a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide* for more information.
- See the *Creating pools* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating RBD pools.
- See the *Creating images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details on creating images in an RBD pool.
- See the *Protecting snapshots of Images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

9.2.17. Deleting snapshots of images

The dashboard allows you to delete snapshots of Ceph block device images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

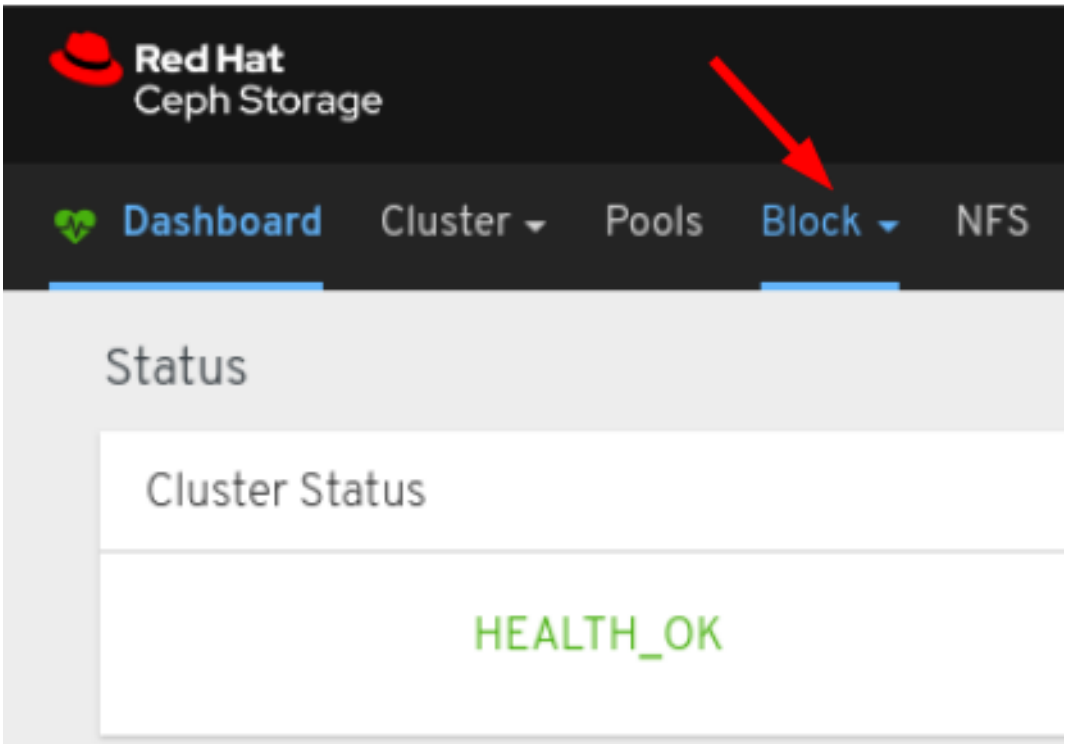

3. Select *Images* from the drop-down:

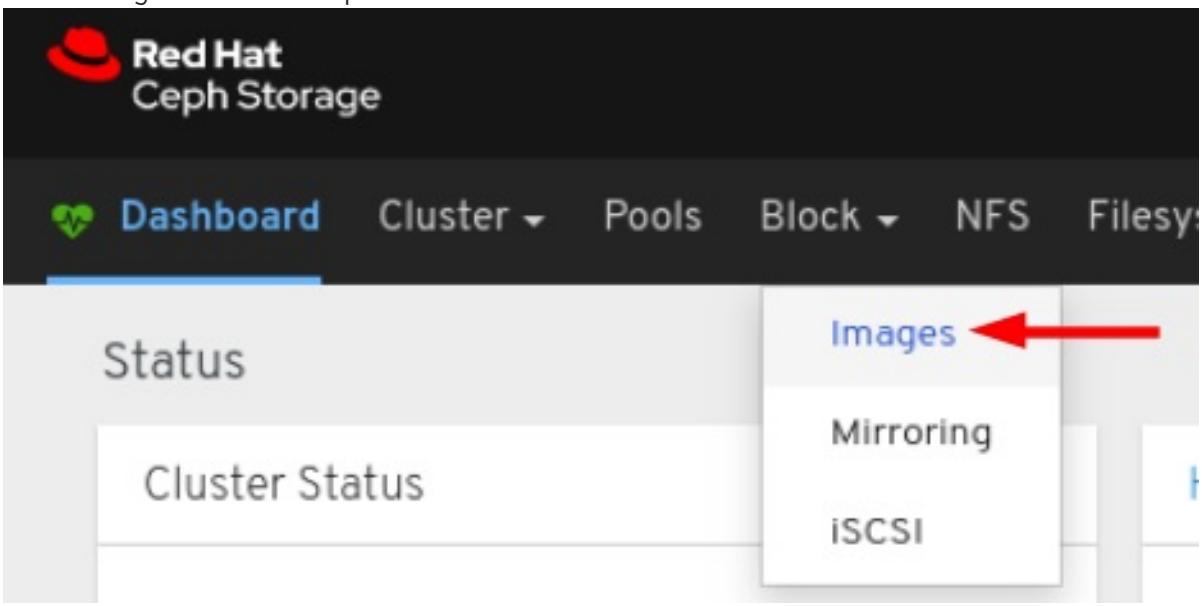

4. To delete the snapshot of the image, in the *Images* tab, click its row, and then click the *Snapshots* tab:

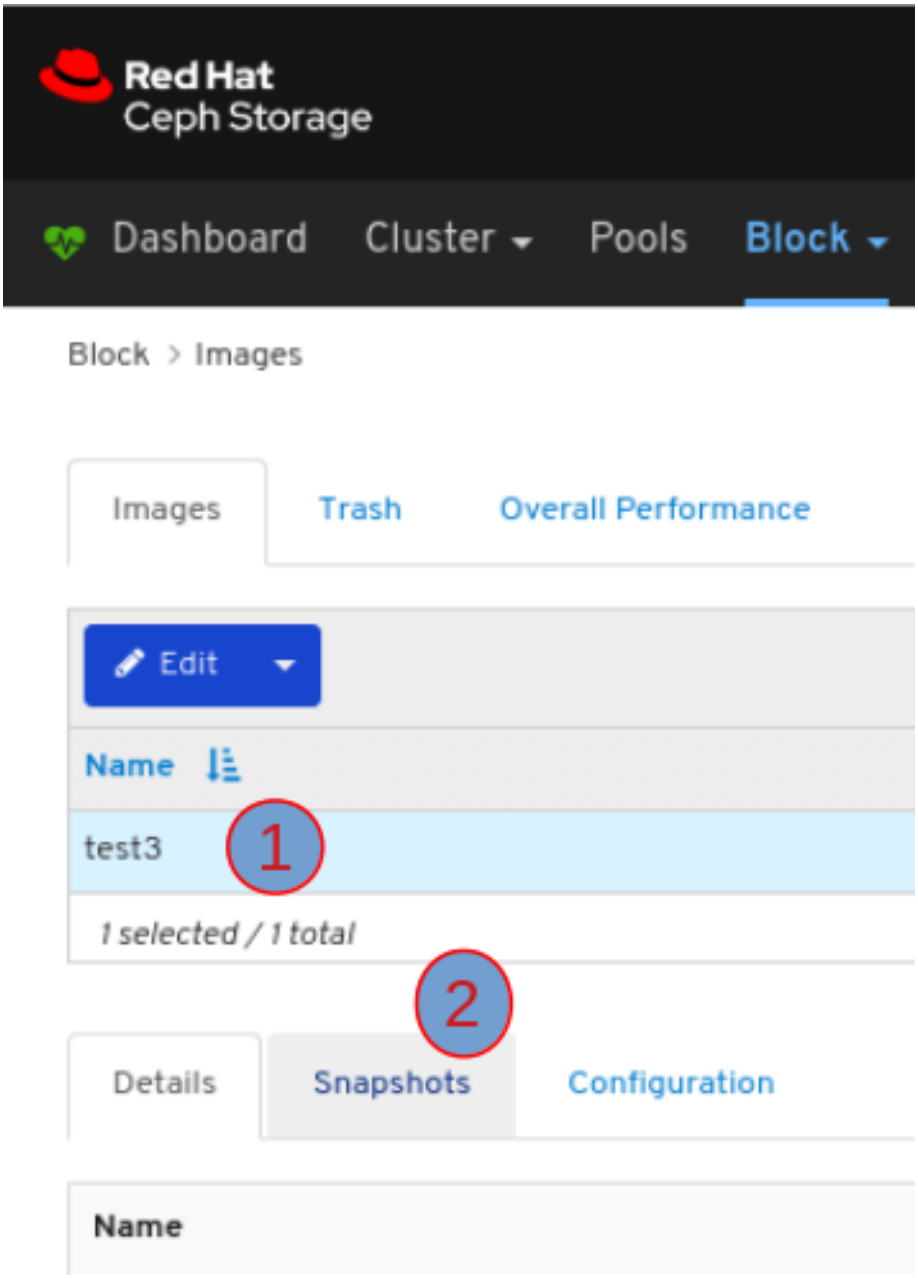

5. Select *Delete* in the the *Rename* drop-down:

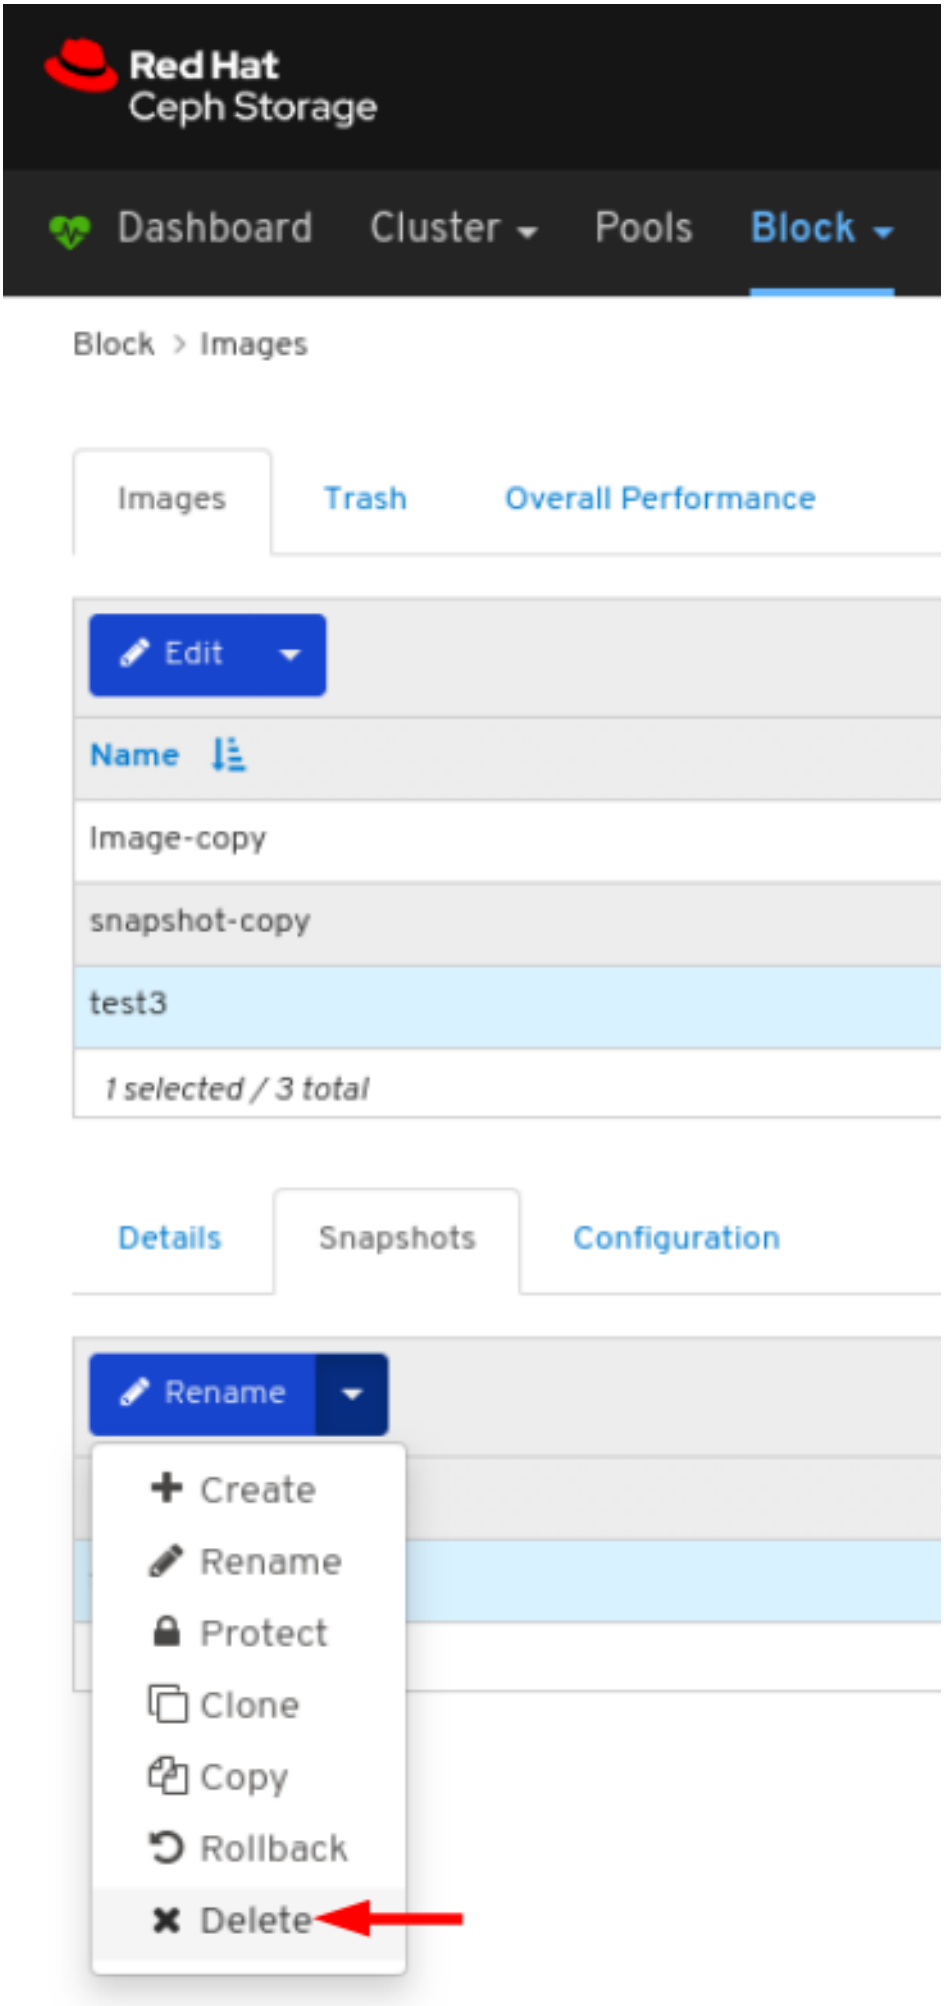
Red Hat
Ceph Storage
Dashboard
Cluster
Pools
Block
Block > Images
Images
Trash
Overall Performance
Edit
Name
Image-copy
snapshot-copy
test3
1 selected / 3 total
Details
Snapshots
Configuration
Rename
Create
Rename
Protect
Clone
Copy
Rollback
Delete

6. A notification towards the top right corner of the page indicates the snapshot of the image was updated successfully.

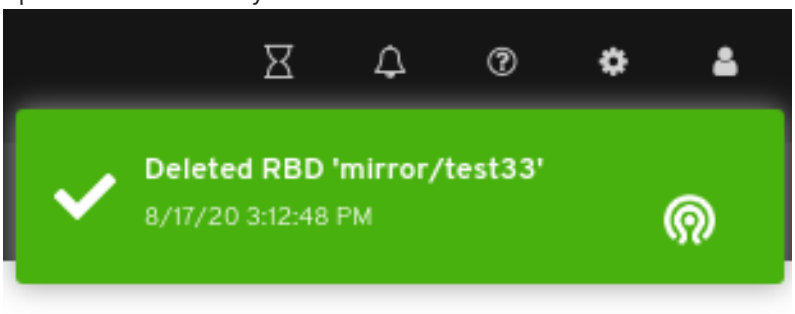

Additional Resources

- See the *Deleting a block device snapshot* section in the *Red Hat Ceph Storage Block Device Guide*] for more information.
- See the *Unprotecting snapshots of Images* section in the *Red Hat Ceph Storage Dashboard Guide* for more details.

9.3. MIRRORING FUNCTIONS

The dashboard allows you to manage and monitor mirroring functions.

9.3.1. Prerequisites

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

9.3.2. Mirroring view

The dashboard allows you to view the overall state of mirroring functions.

Prerequisites

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

Procedure

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

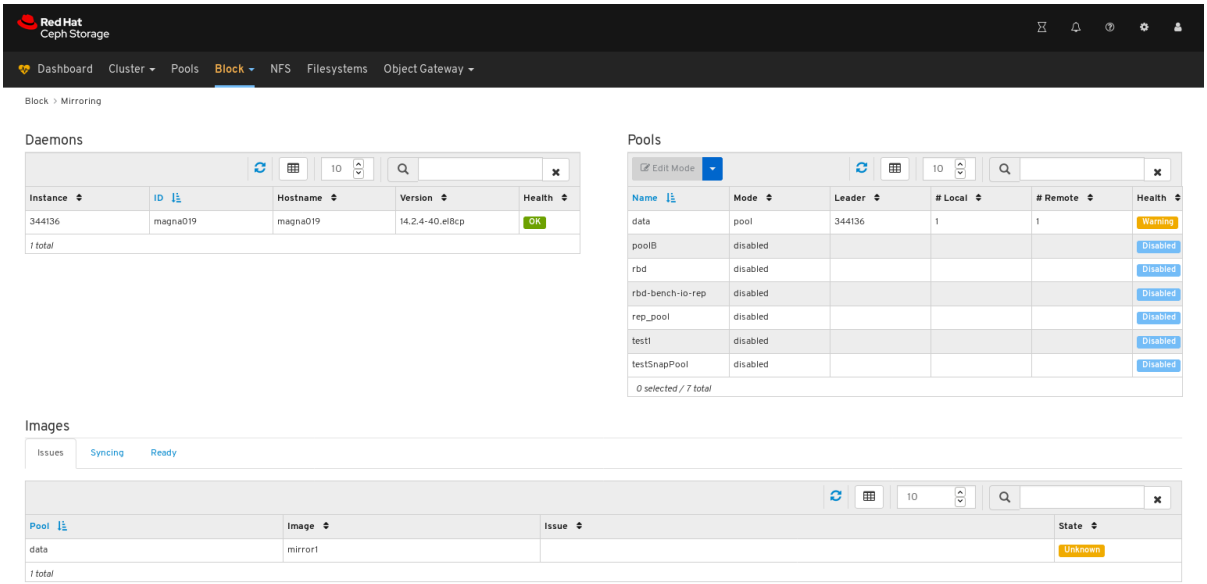

In the above example, you can see mirroring information categorized into tables labeled *Daemons*, *Pools*, and *Images*.

Additional Resources

- For more information on mirroring, see Block Device Mirroring in the Block Device Guide.

9.3.3. Editing mode

The dashboard allows you to edit mode of the overall state of mirroring functions, which includes pools and images.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

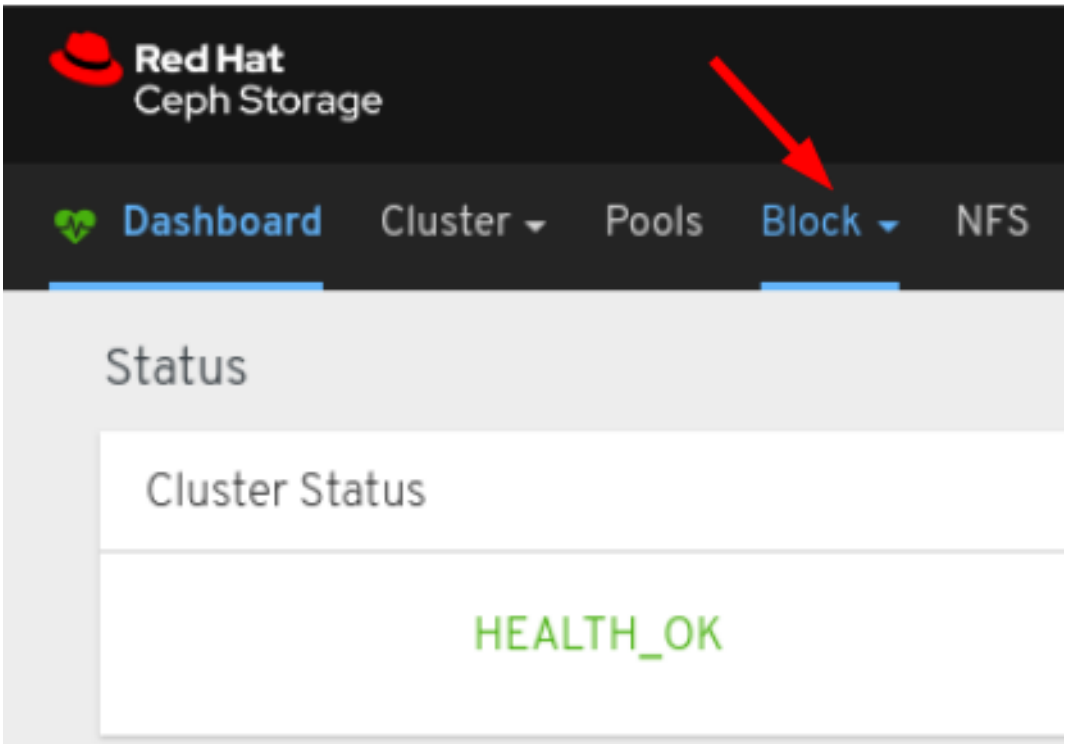

3. Click *Mirroring* from the drop-down:

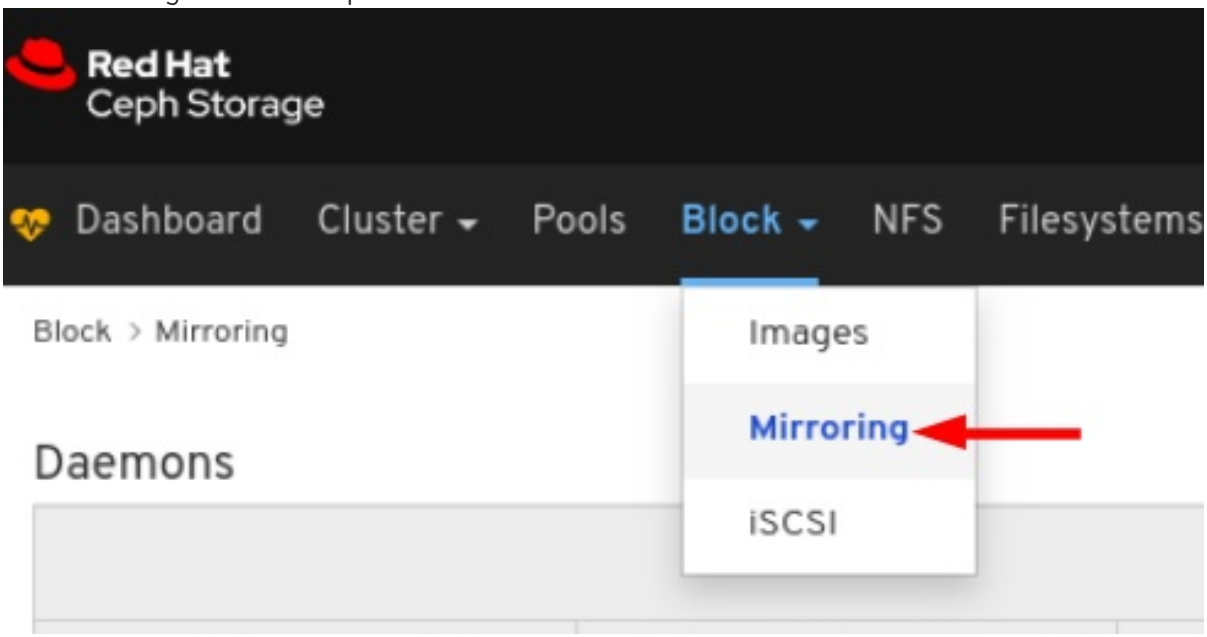

4. In the *Pools* tab, click the row:

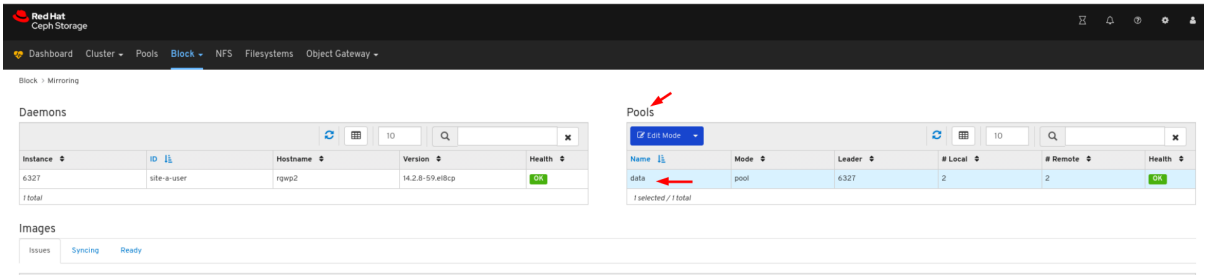

5. In the *Edit Mode* drop-down, select *Edit Mode*:

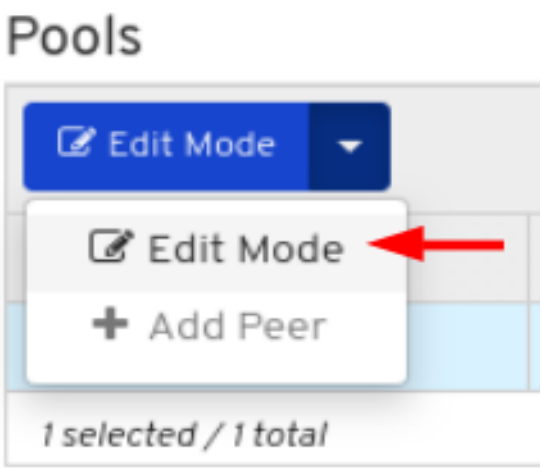

6. In the *Edit Pool mirror mode* window, select the mode from the drop-down, and then click the *Update* button:

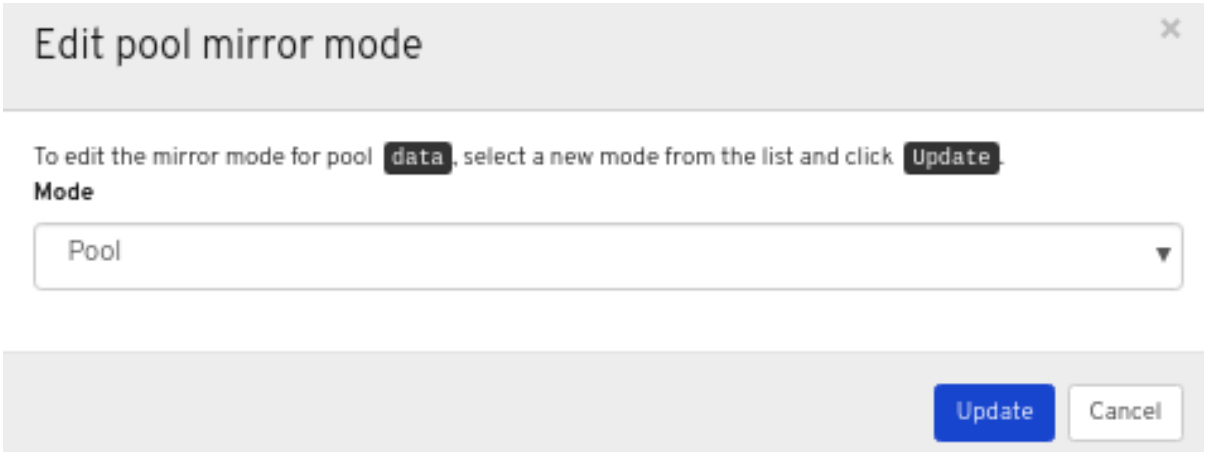

7. A notification towards the top right corner of the page indicates the mirror mode was updated successfully.

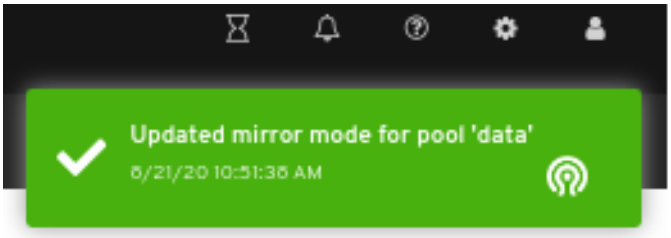

Additional Resources

- See the *Ceph Block Device Mirroring* section in the *Red Hat Ceph Storage Block Device Guide* for more information.

9.3.4. Adding peer in mirroring

The dashboard allows you to add storage cluster peer for the `rbd-daemon` mirror to discover its peer storage cluster.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*:

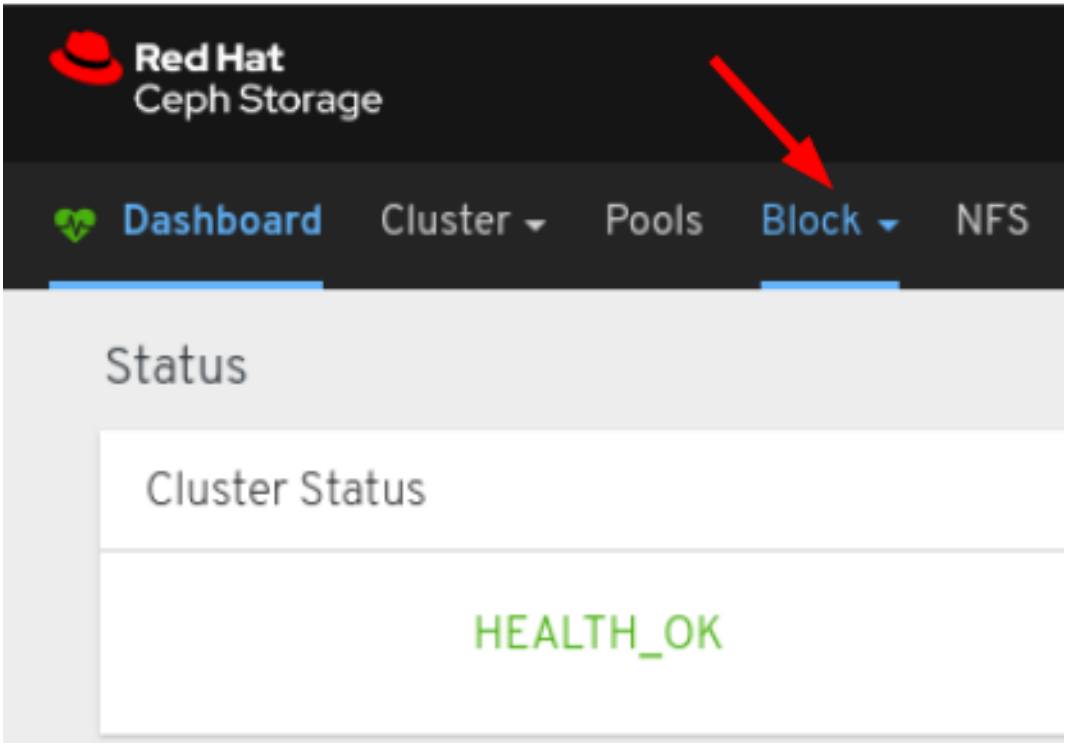

3. Click *Mirroring* from the drop-down:

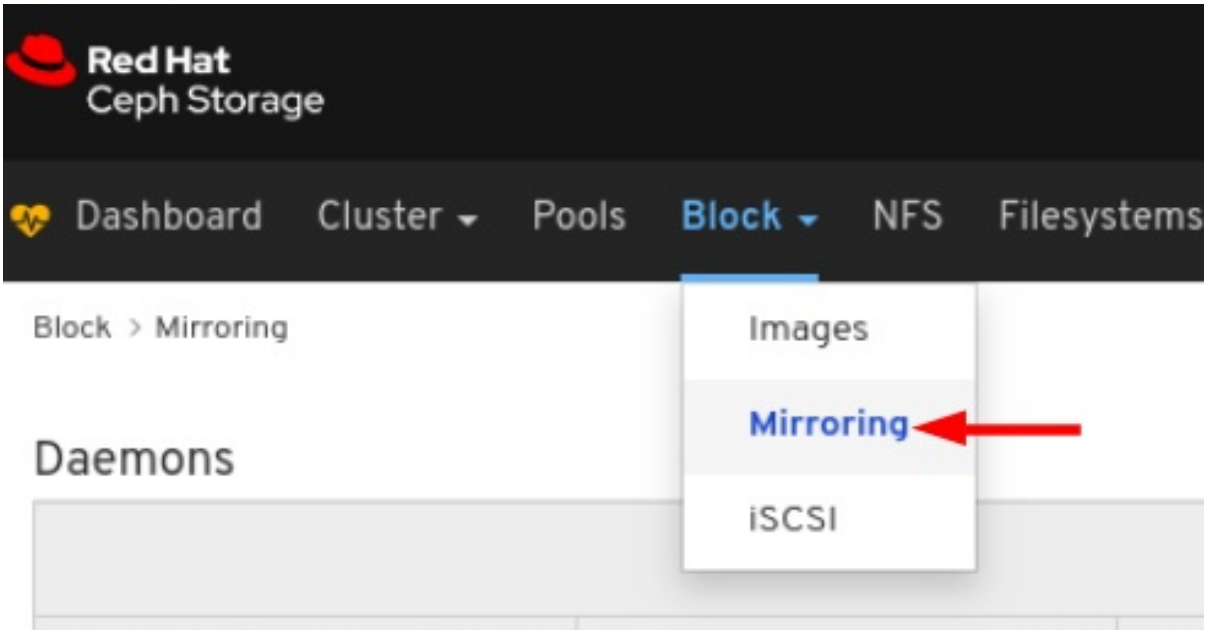

4. In the *Pools* tab, click the row

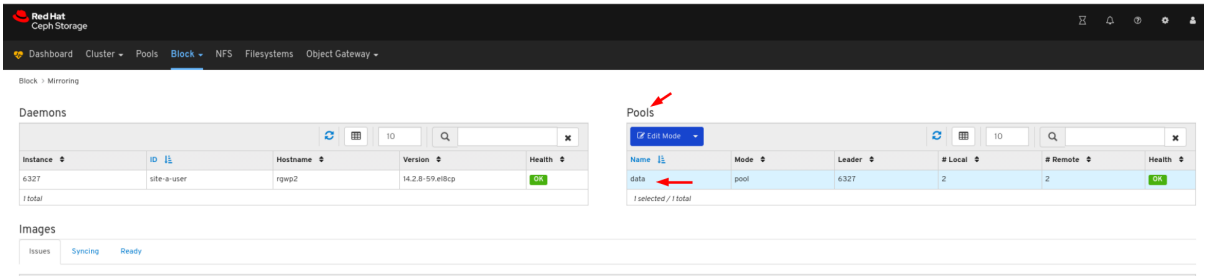

5. In the *Edit Mode* drop-down, select *Add peer*:

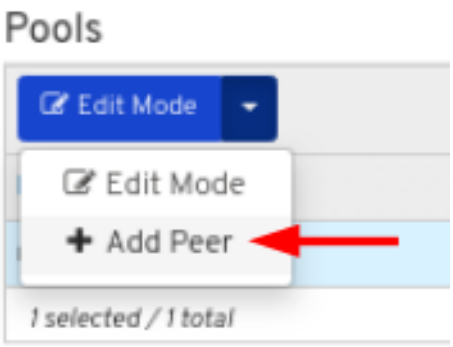

6. In the *Add pool mirror peer* window, enter the parameters, and then click the *Submit* button:

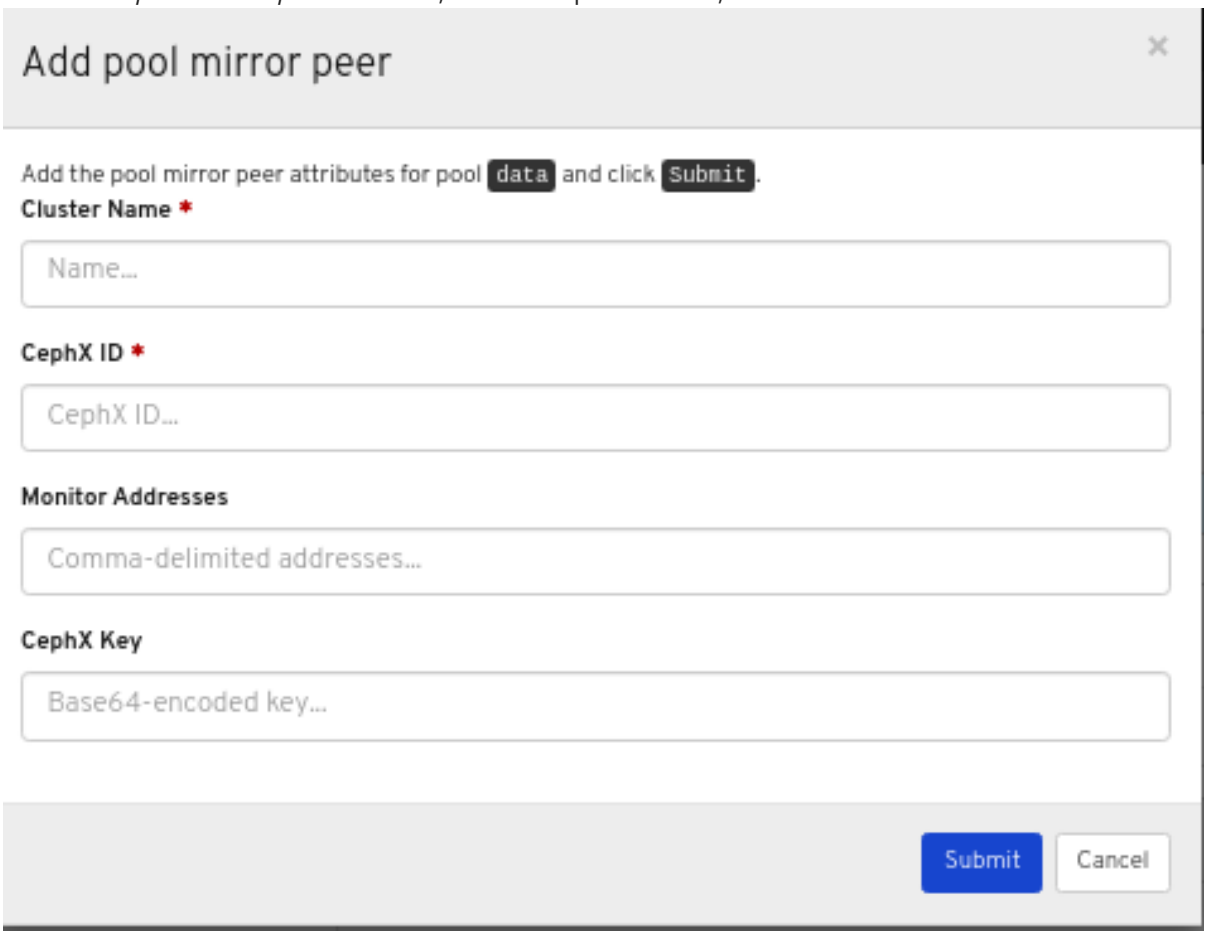

7. A notification towards the top right corner of the page indicates the mirror peer was created successfully.

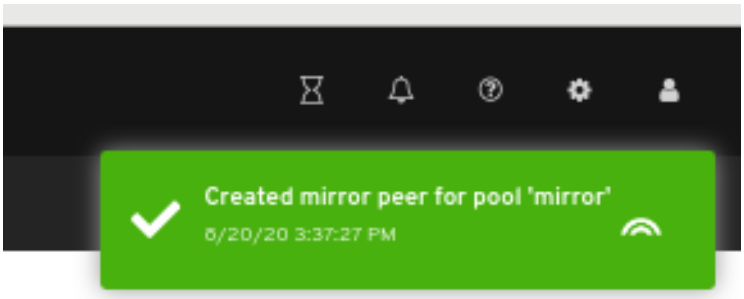

Additional Resources

- See the *Adding a storage cluster peer* section in the *Red Hat Ceph Storage Block Device Guide* for more information.

9.3.5. Editing peer in mirroring

The dashboard allows you to edit storage cluster peer for the`rbd-daemon` mirror to discover its peer storage cluster.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.

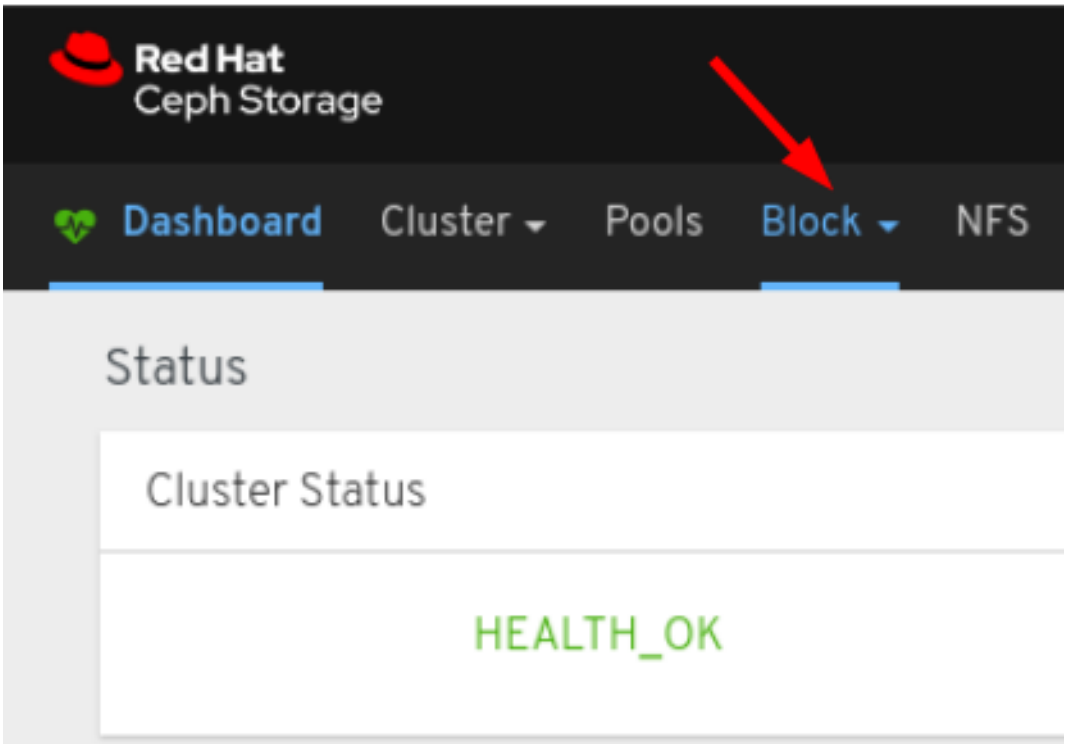

3. Click *Mirroring* from the drop-down:

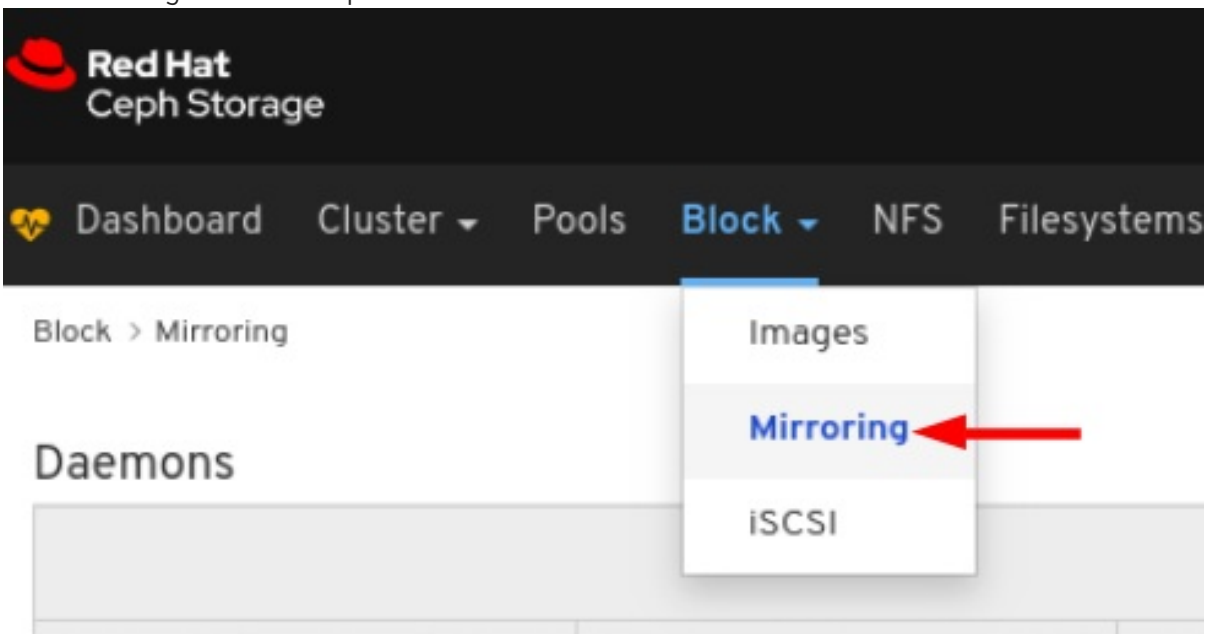

4. In the *Pools* tab, click the row

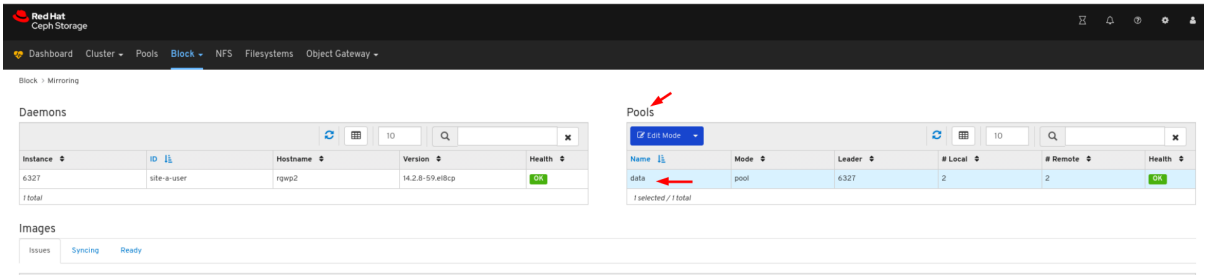

5. In the *Edit Mode* drop-down, select *Edit peer*:

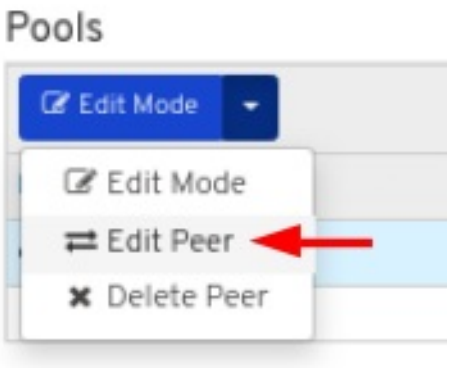

6. In the *Edit pool mirror peer* window, edit the parameters, and then click the *Submit* button:

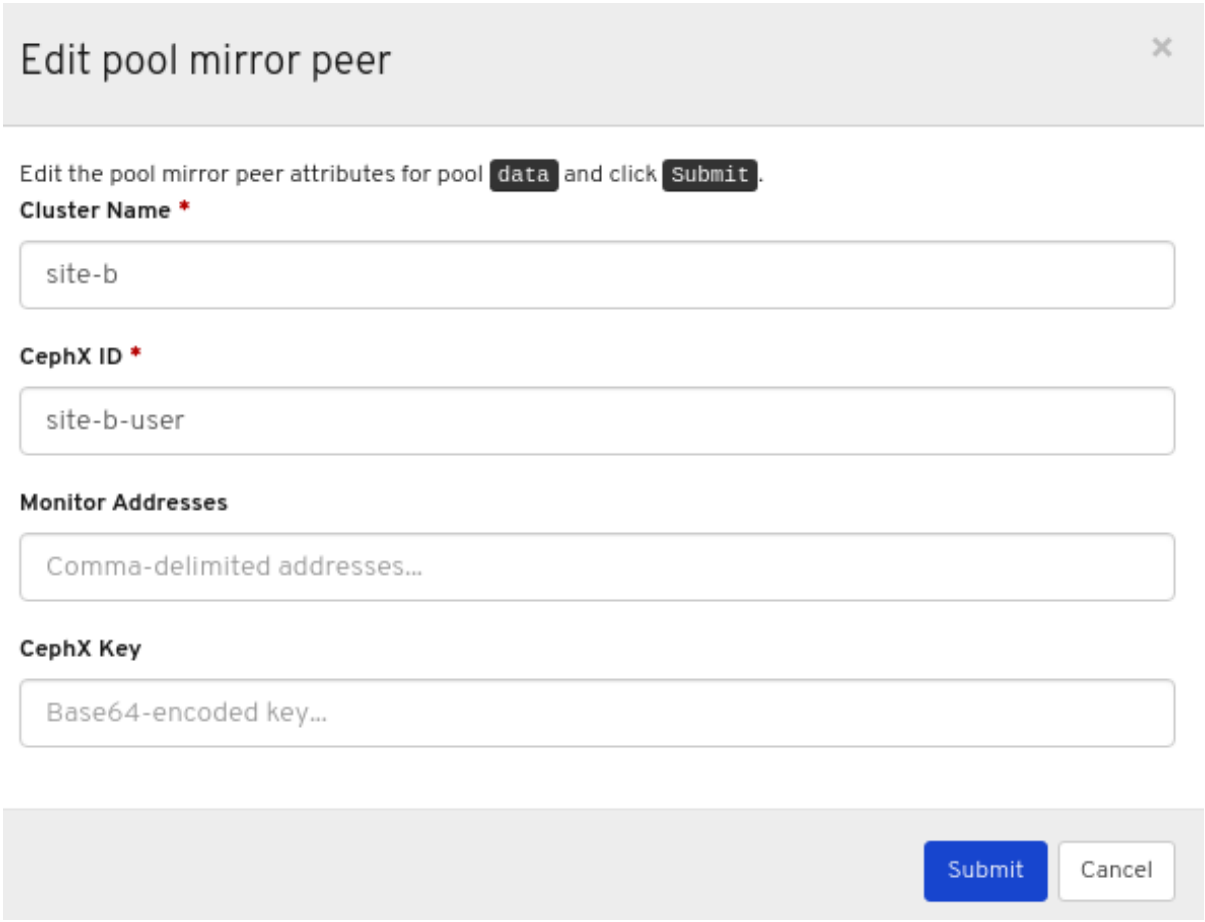

7. A notification towards the top right corner of the page indicates the mirror peer was updated successfully.

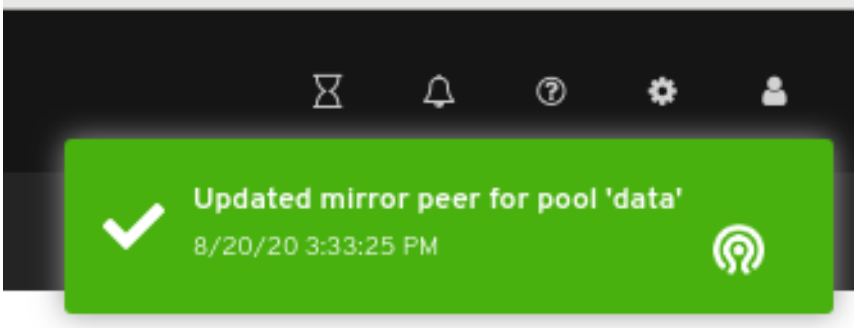

Additional Resources

- See the *Adding peer in mirroring* section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

9.3.6. Deleting peer in mirroring

The dashboard allows you to edit storage cluster peer for the`rbd-daemon` mirror to discover its peer storage cluster.

Prerequisites

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

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.

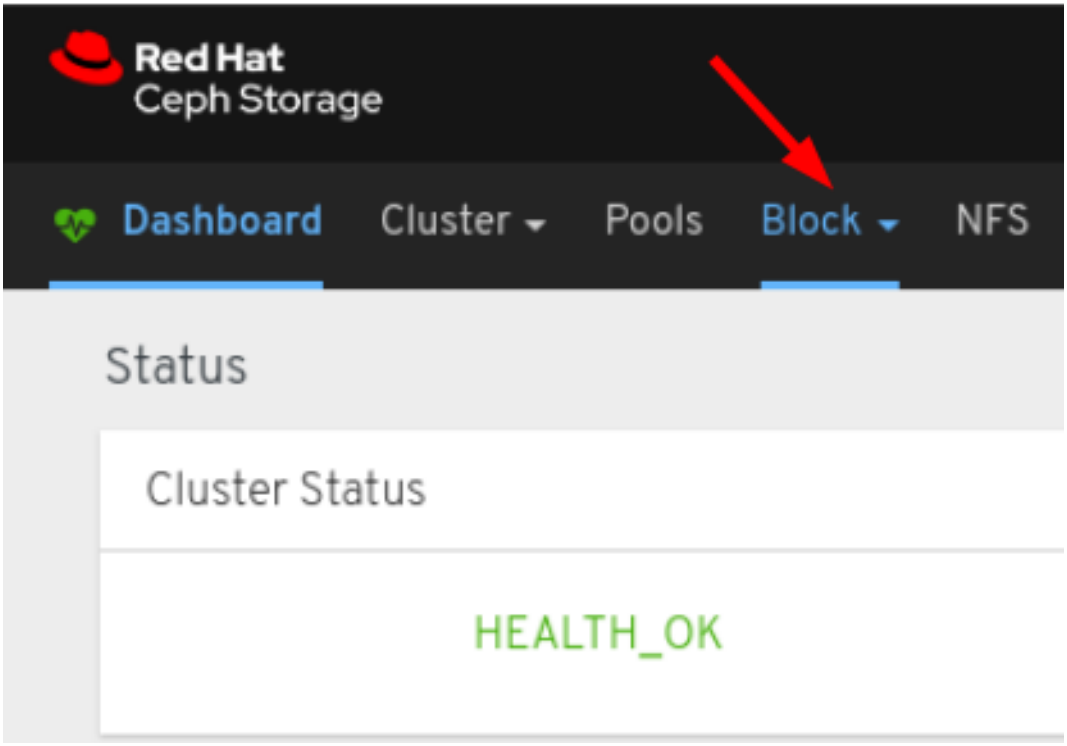

3. Click *Mirroring* from the drop-down:

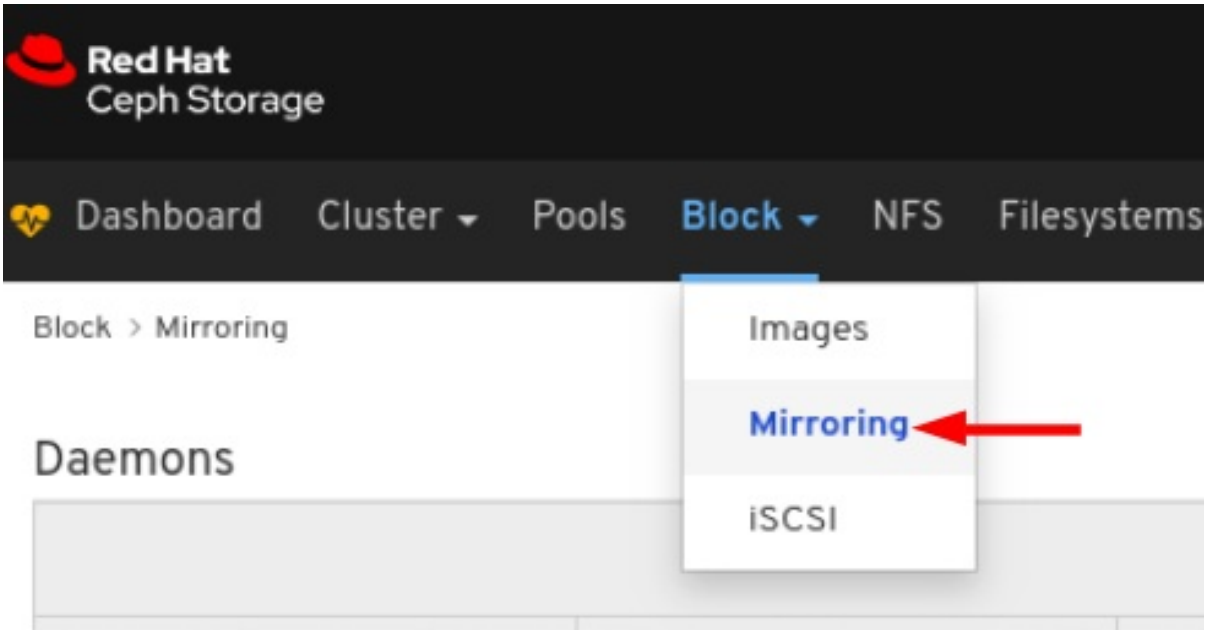

4. In the *Pools* tab, click the row:

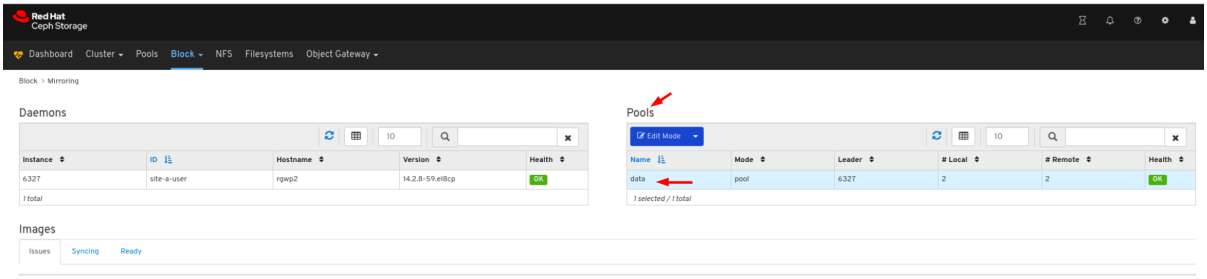

5. In the *Edit Mode* drop-down, select *Delete peer*:

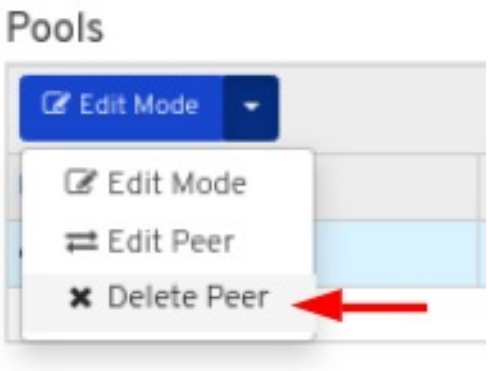

6. In the *Delete mirror peer* dialog window, Click the *Yes, I am sure* box and then Click *Delete mirror peer* to save the settings:

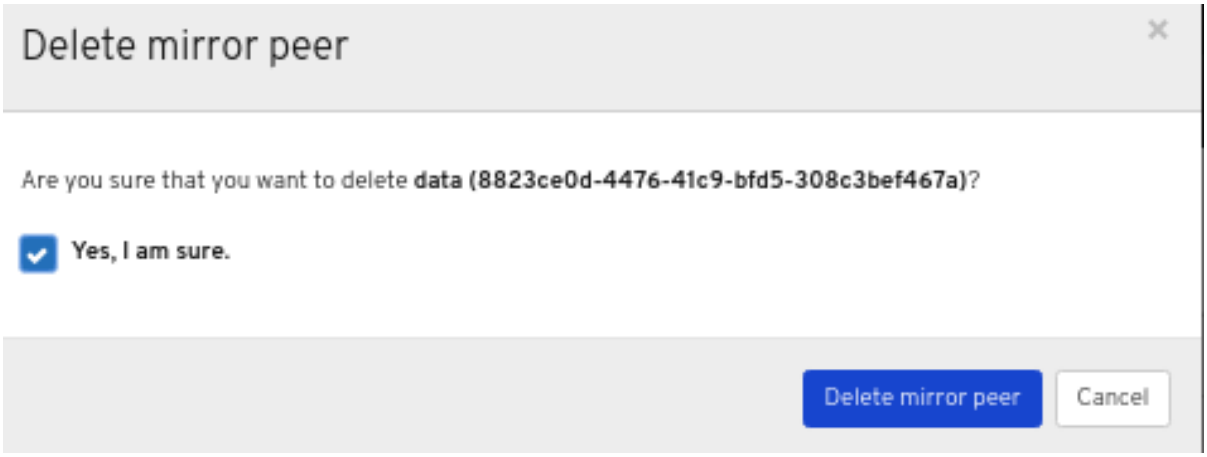

7. A notification towards the top right corner of the page indicates the image was moved to trash successfully.

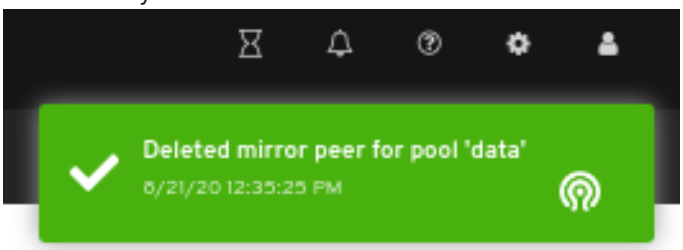

Additional Resources

- See the *Adding peer in mirroring* section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

9.4. ISCSI FUNCTIONS (LIMITED AVAILABILITY)

The dashboard allows you to manage and monitor iSCSI images and targets. Before you can use the dashboard to manage and monitor iSCSI images and targets, you must add gateways to it and enable the dashboard iSCSI feature.

NOTE

This technology is Limited Availability. See the *Deprecated functionality* chapter for additional information.

9.4.1. Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- iSCSI gateways are added to the dashboard.

9.4.2. Manually adding iSCSI gateways to the dashboard

The Red Hat Ceph Storage Dashboard can manage iSCSI targets using the REST API provided by the **rbd-target-api** service of the Ceph iSCSI Gateway. You must add the API address to the dashboard before the dashboard can access it.

Prerequisites

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

IMPORTANT

The Ceph iSCSI gateway requires at lest two gateways to provide high availability.

Procedure

1. Log in to a Monitor node.
2. Optional: If the REST API for the iSCSI gateway is configured in HTTPS mode using a self-signed certificate, you must configure the dashboard to avoid SSL certificate verification when accessing the API. Run the following command to disable SSL verification.

```
# ceph dashboard set-iscsi-api-ssl-verification false
```

   Example:

```
[root@mon ~]# ceph dashboard set-iscsi-api-ssl-verification false
Option ISCSI_API_SSL_VERIFICATION updated
```

3. Create two files for the gateways:

```
# echo "SCHEME :// USERNAME : PASSWORD @ HOST [: PORT ]" >
FILE_CONTAINING_GATEWAY_URL
```

   Example:

```
[root@mon ~]# echo "http://admin:admin@192.168.122.157:5000" > /tmp/first_gateway
[root@mon ~]# echo "http://admin:admin@192.168.122.193:5000" > /tmp/second_gateway
```

NOTE

The *USERNAME* and *PASSWORD* were set when you configured the iSCSI target. The credentials can be retrieved from the **iscsi-gateway.cfg** file on the iSCSI Gateway node.

4. Add the two gateways to the dashboard:

```
# ceph dashboard iscsi-gateway-add -i FILE_CONTAINING_GATEWAY_URL
```

   Example:

```
[root@mon ~]# ceph dashboard iscsi-gateway-add -i /tmp/first_gateway
Success
[root@mon ~]# ceph dashboard iscsi-gateway-add -i /tmp/second_gateway
Success
```

5. Verify the gateways were added correctly:

```
# ceph dashboard iscsi-gateway-list
```

Example:

```
[root@mon ~]# ceph dashboard iscsi-gateway-list
{"gateways": {"ceph4": {"service_url": "http://admin:admin@192.168.122.193:5000"}, "ceph4":
{"service_url": "http://admin:admin@192.168.122.193:5000"}}}
```

6. Optional: If you make a mistake adding a gateway you can remove it by specifying its hostname as mentioned in the command **iscsi-gateway-list**:

```
# ceph dashboard iscsi-gateway-rm GATEWAY_NAME
```

Example:

```
[root@mon ~]# ceph dashboard iscsi-gateway-rm ceph4
Success
```

Additional Resources

- For information on how to install the Ceph iSCSI Gateway, see Installing the iSCSI gateway in the Block Device Guide.

9.4.3. iSCSI overview

The dashboard provides an overview that displays iSCSI gateway hosts and images exported over iSCSI.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph iSCSI gateway is installed.
- iSCSI gateways are added to the dashboard.
- The dashboard iSCSI feature is enabled.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.
3. Click *iSCSI*:

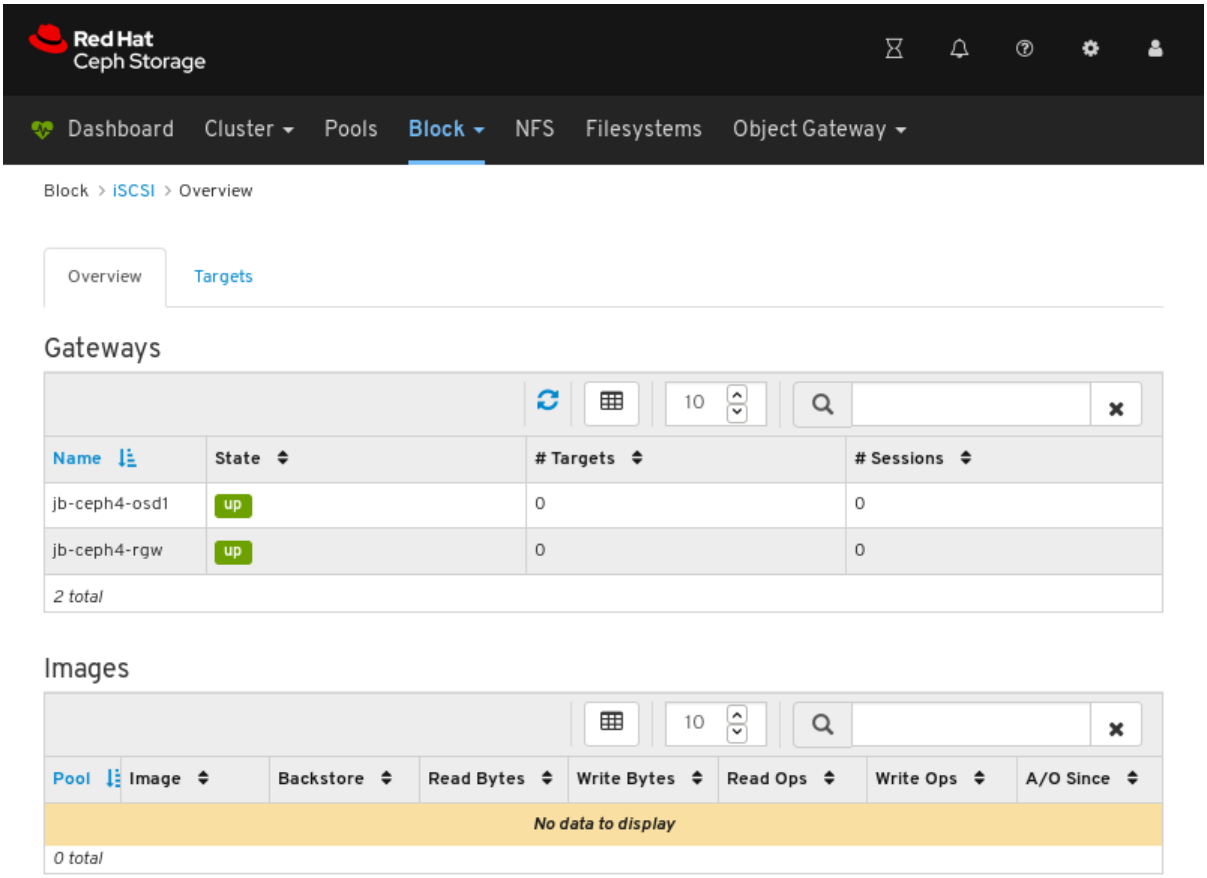

Additional Resources

- For information on how to install the Ceph iSCSI Gateway, see Installing the iSCSI gateway in the Block Device Guide.
- For information on how to add iSCSI gateways to the dashboard, see Manually adding iSCSI gateways to the dashboard in the Dashboard Guide.
- For information on how to enable the dashboard iSCSI feature see Enabling the dashboard iSCSI feature in the dashboard in the Dashboard Guide.

9.4.4. Creating iSCSI targets

The dashboard allows you to create iSCSI targets.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph iSCSI gateway is installed with at least two gateways.
- iSCSI gateways are added to the dashboard.
- The dashboard iSCSI feature is enabled.
- A replicated pool with the RBD application enabled.

- An erasure coded pool with the RBD application enabled.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.
3. Click *iSCSI*:

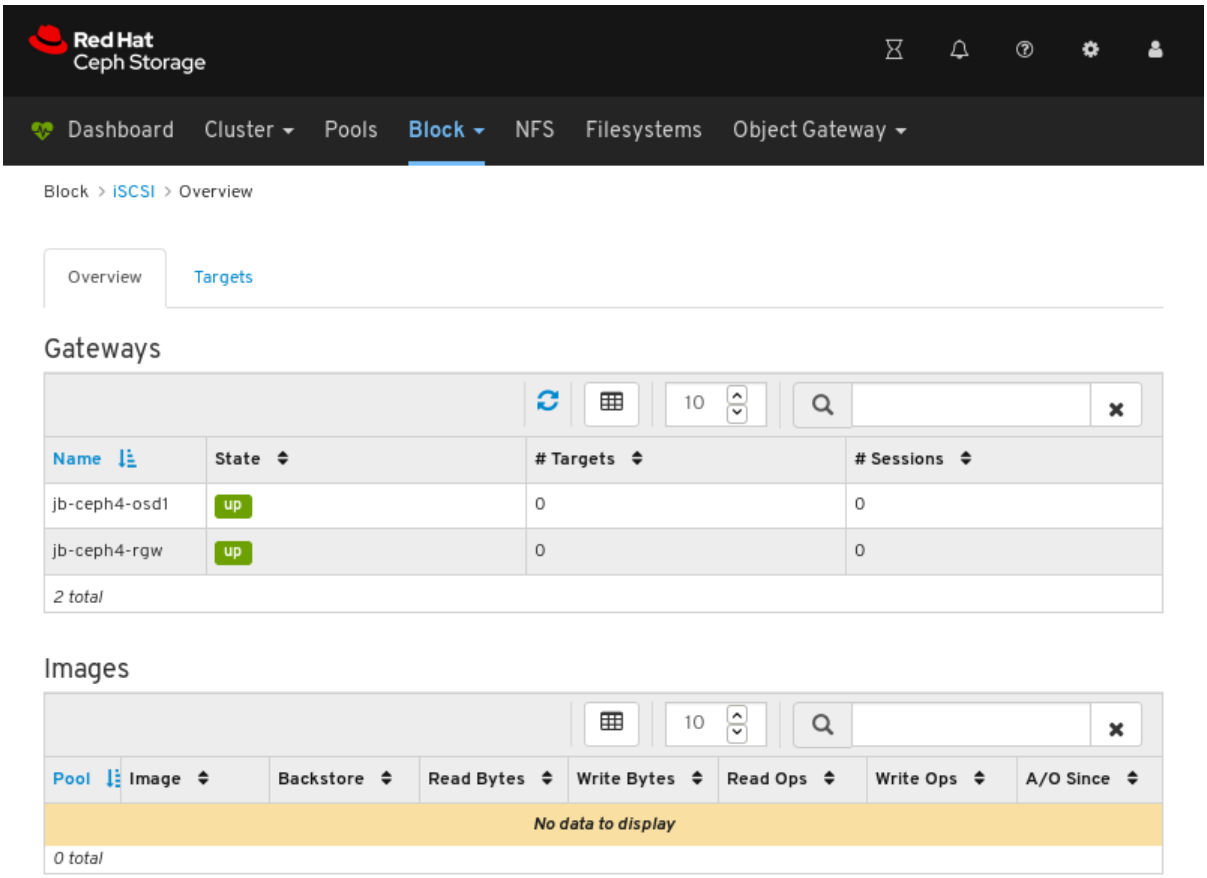

4. Towards the upper left corner of the page, click the *Targets* tab:

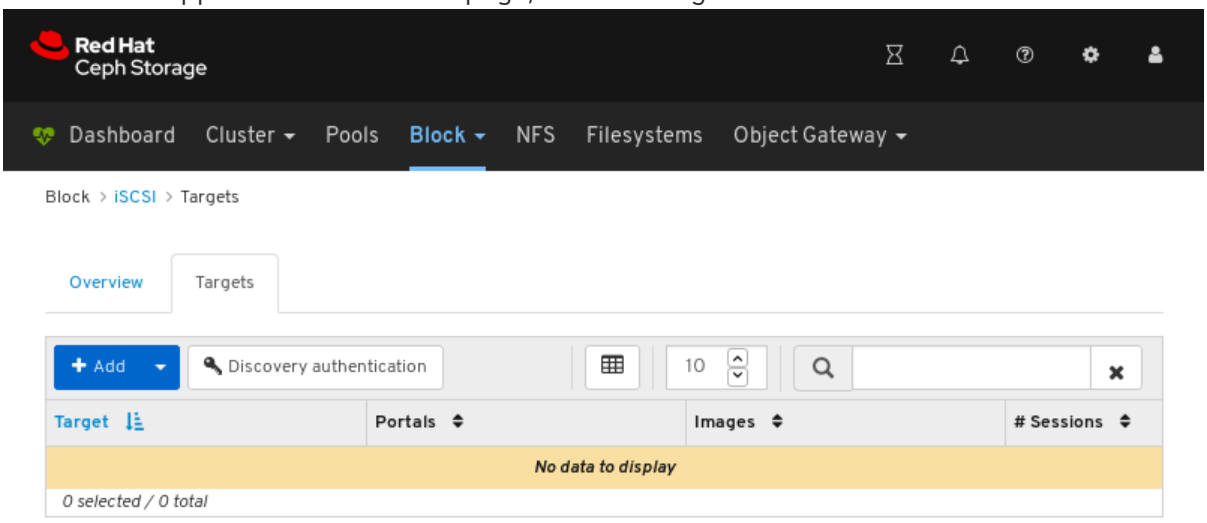

5. Towards the upper left corner of the page, click the *Add* button:

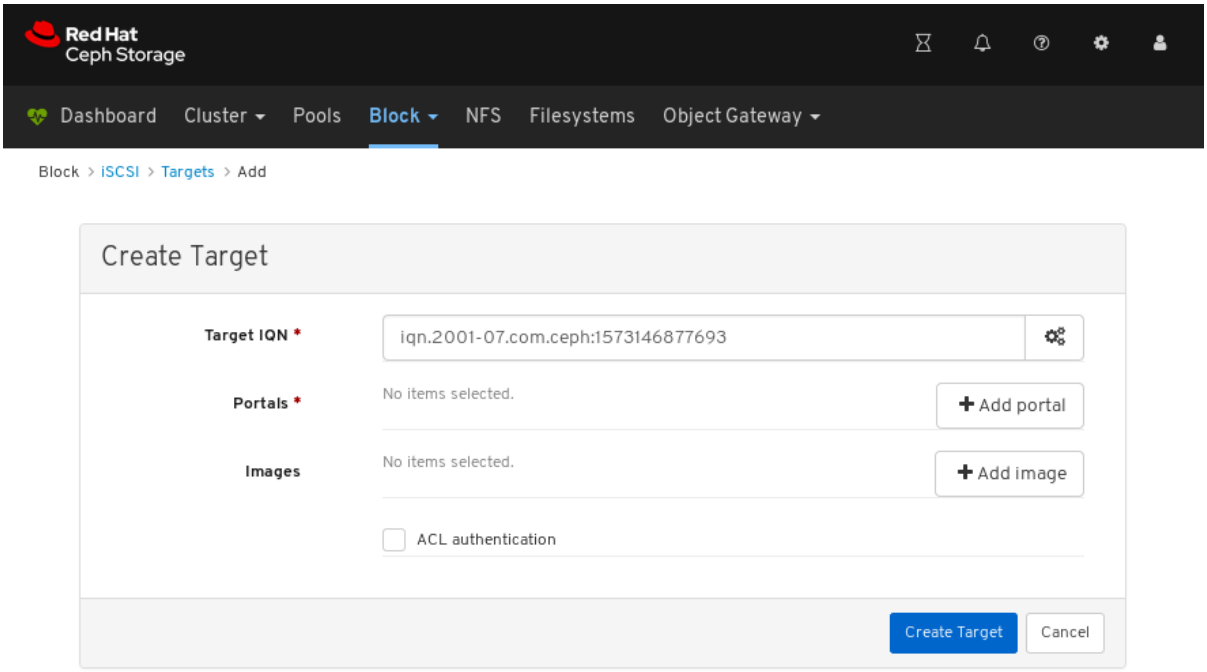

6. Optional: Modify the Target IQN.

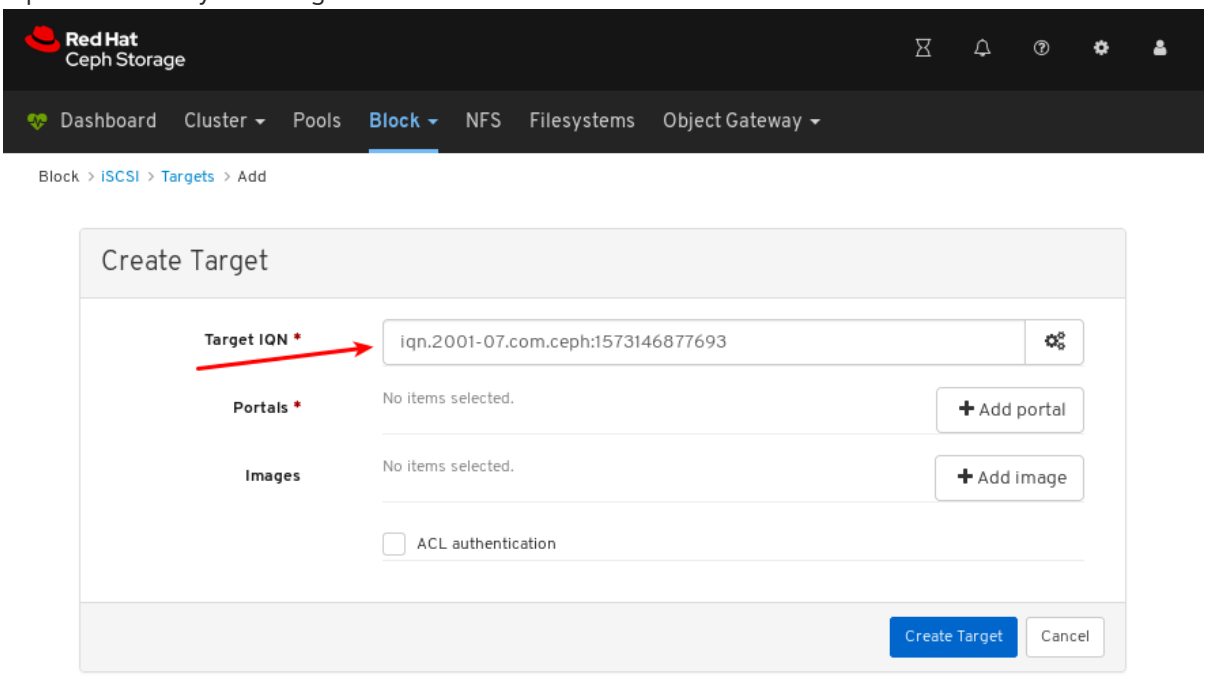

7. Optional: Set advanced settings for the target.

   a. Click the gear to set advanced settings for the target.

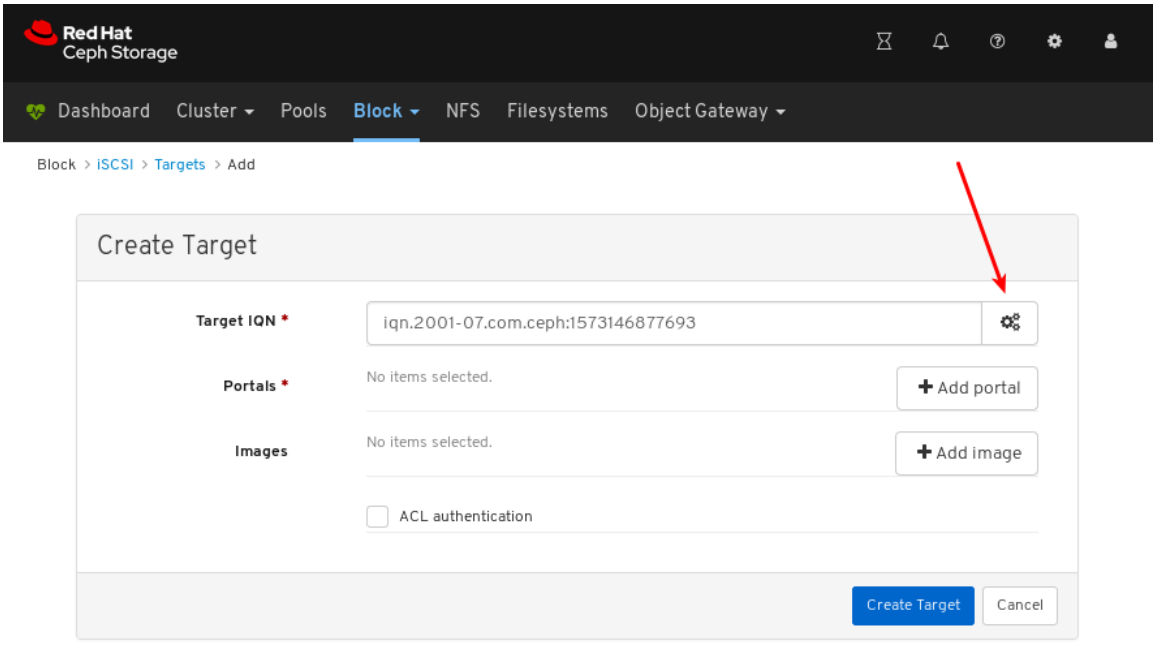

b. Set advanced settings in the *Advanced Settings* dialog window.

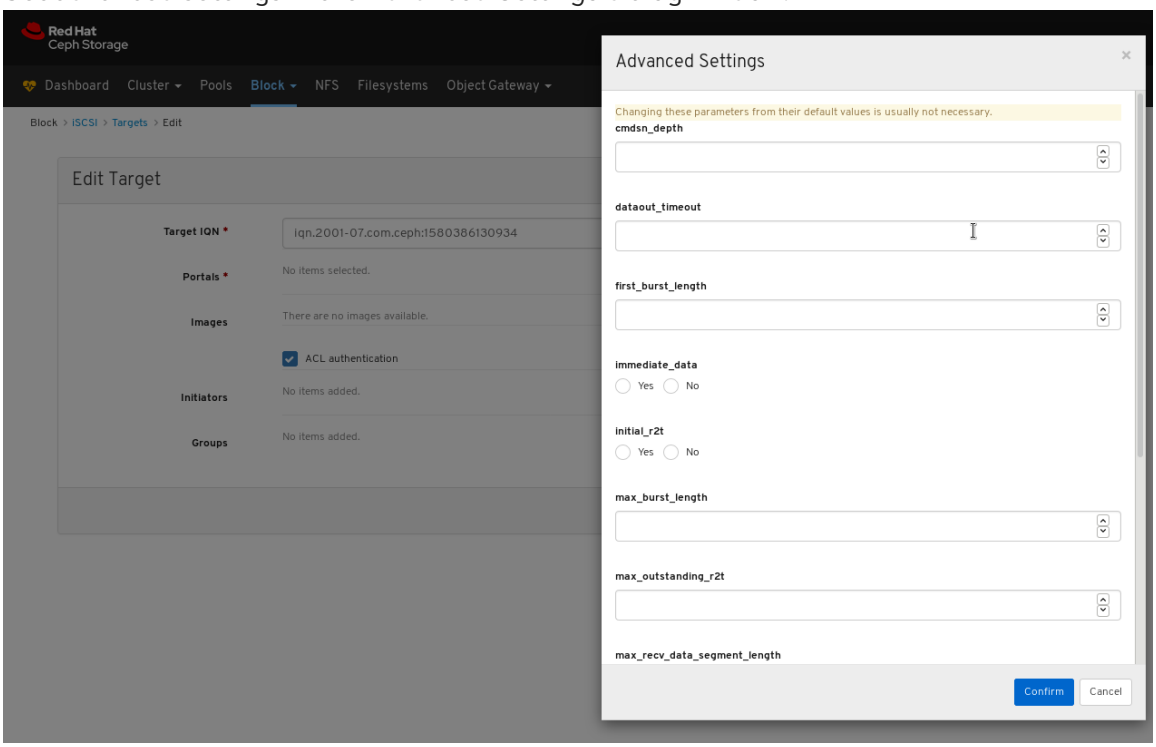

c. Click *Confirm* to save the settings.

8. Click the *Add portal* button and select the first of at least two gateways:

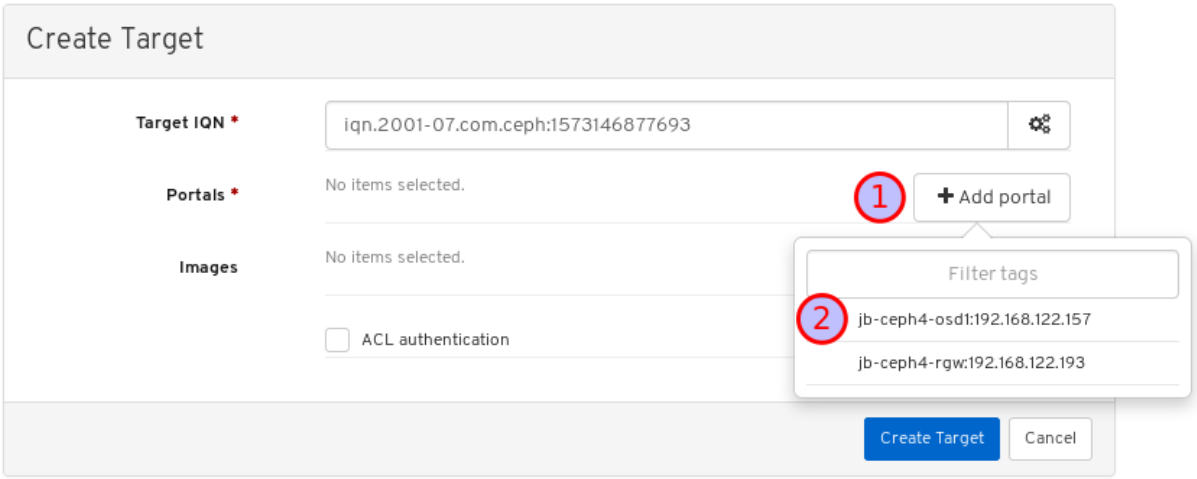

9. Click the *Add portal* button and select the second of at least two gateways:

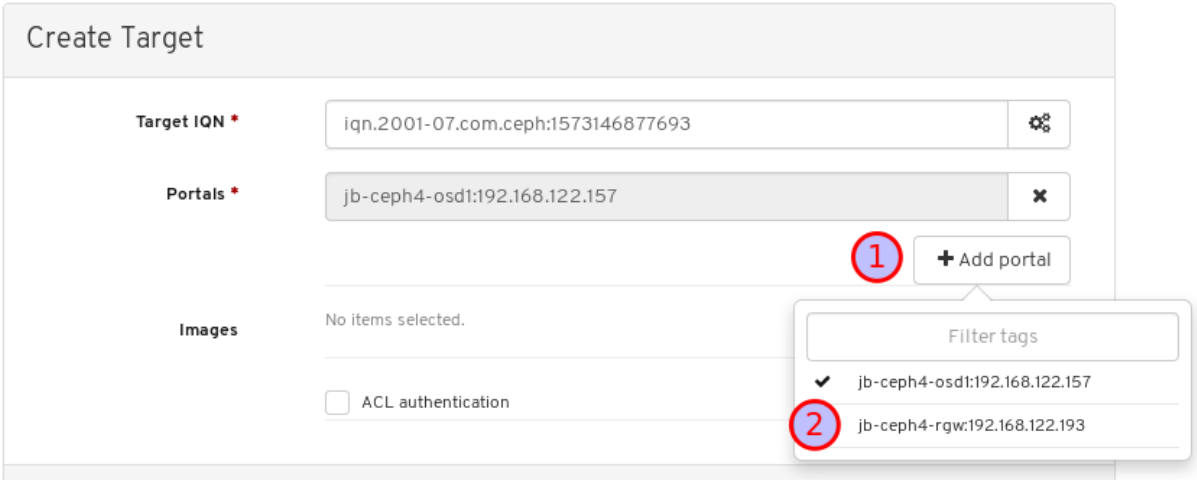

Repeat this step for any additional gateways.

10. Click the *Add image* button and select an image to be exported by the target:

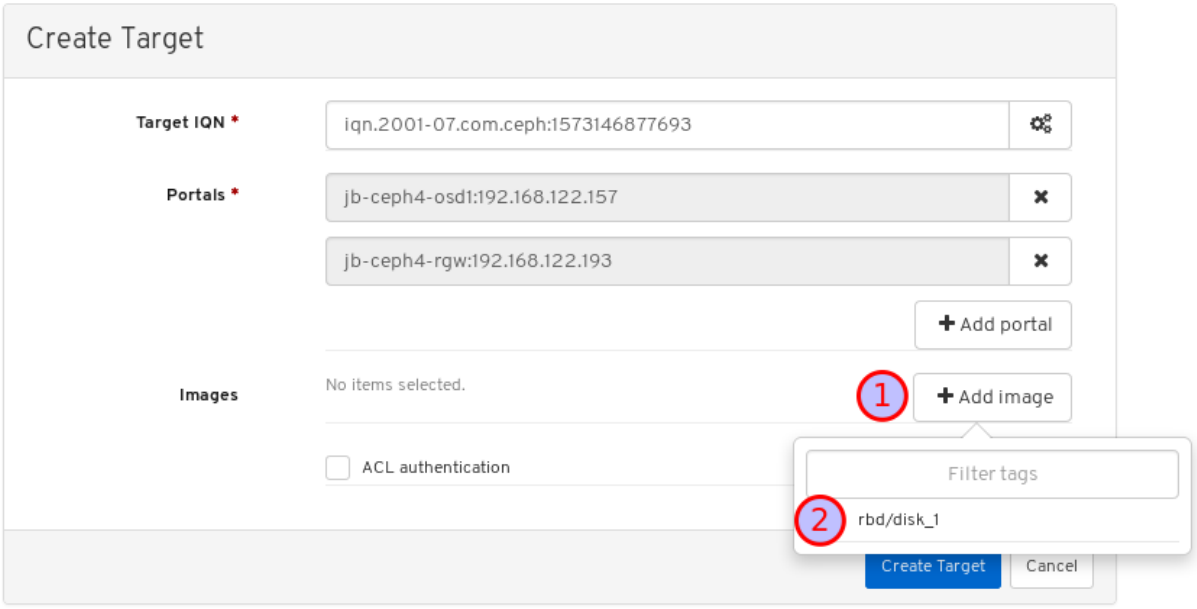

Repeat this step for any additional images.

11. Optional: Modify the Images.

    a. Click the gear to the right of the image

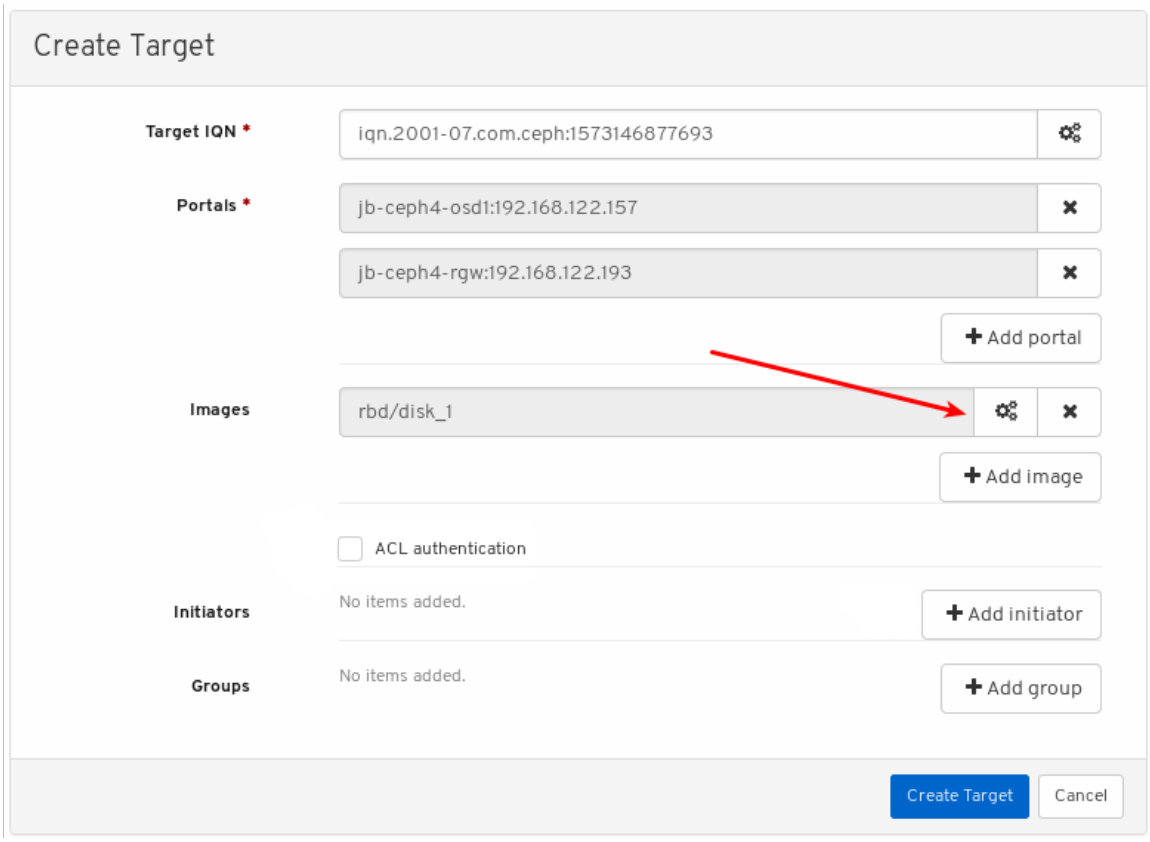

b. Modify image settings in the *Configure* dialog window:

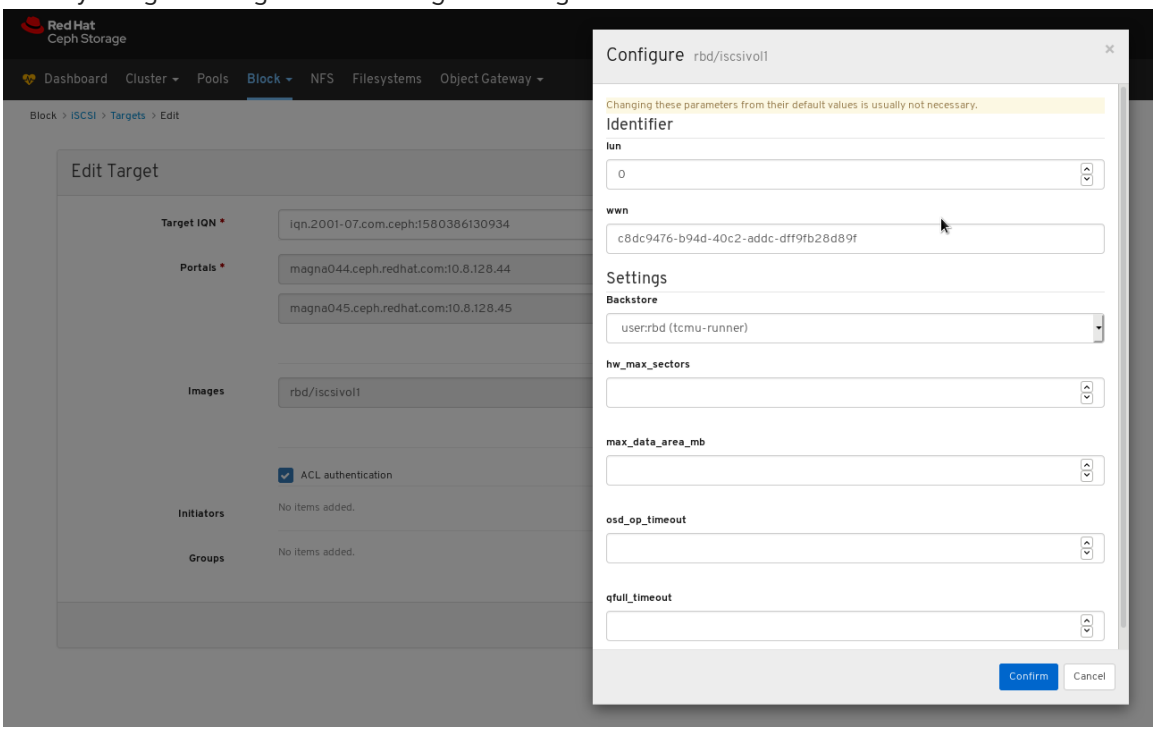

c. Click *Confirm* to save the settings.

12. Click the *ACL authentication* box and then click the *Add initiator* button:

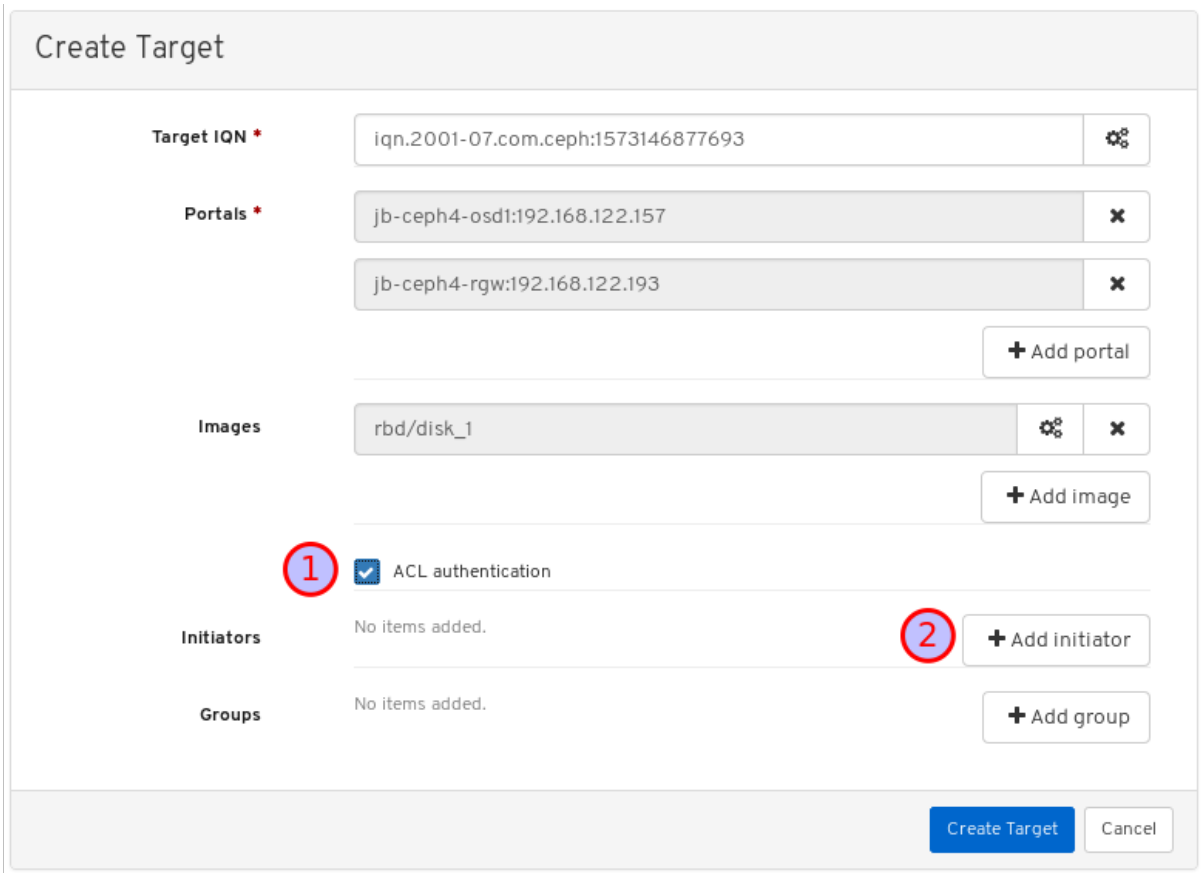

13. Enter the IQN from your client in the first text box:

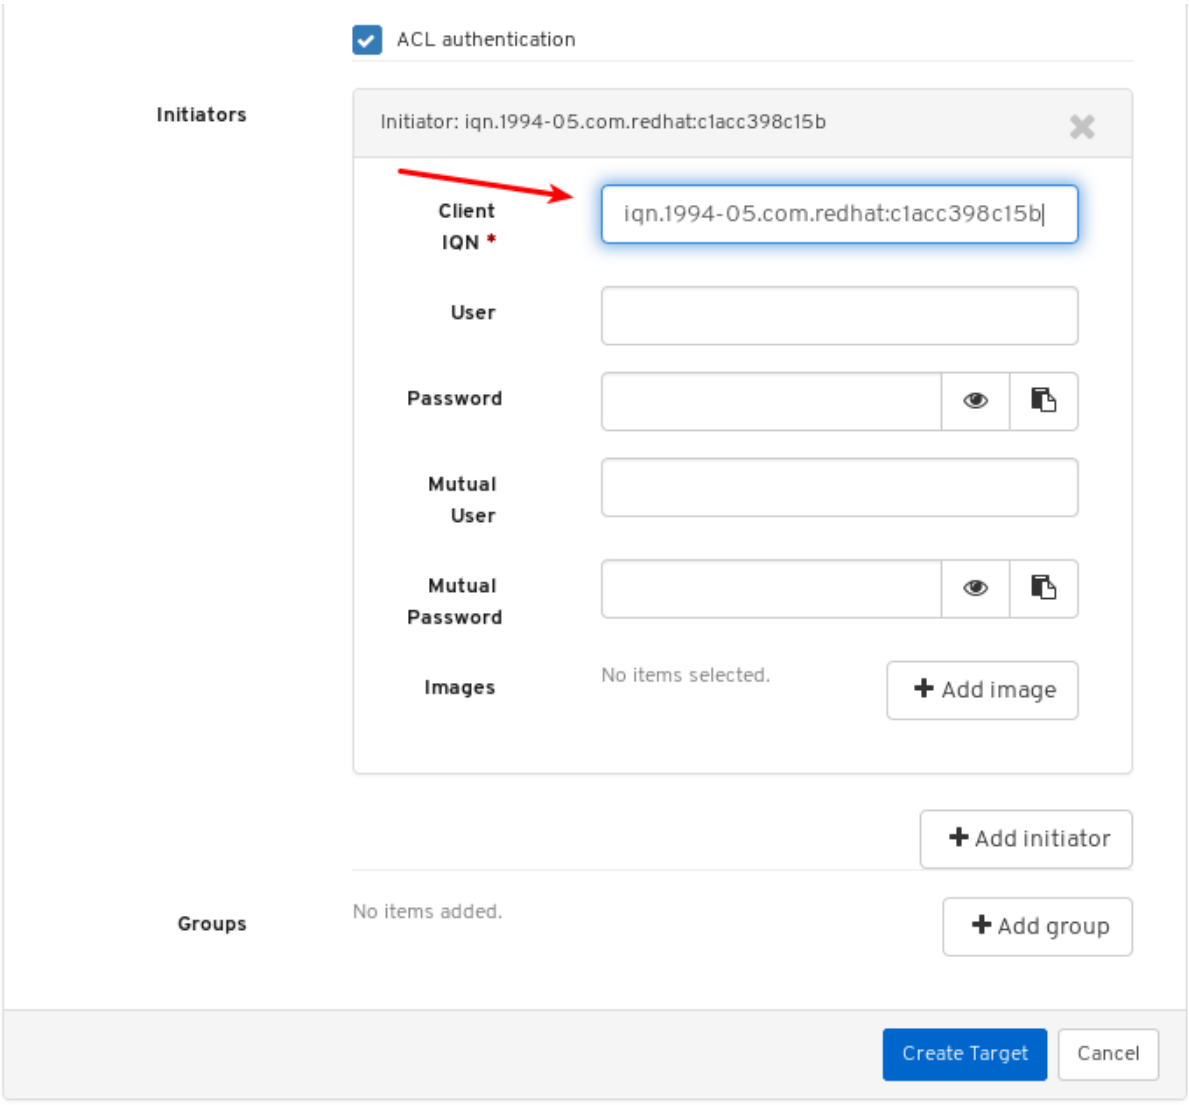

Retrieve the client IQN from the system where the initiator software runs. See Configuring the iSCSI initiator in the Block Device Guide for more information.

14. Enter a user name and password details for the target:

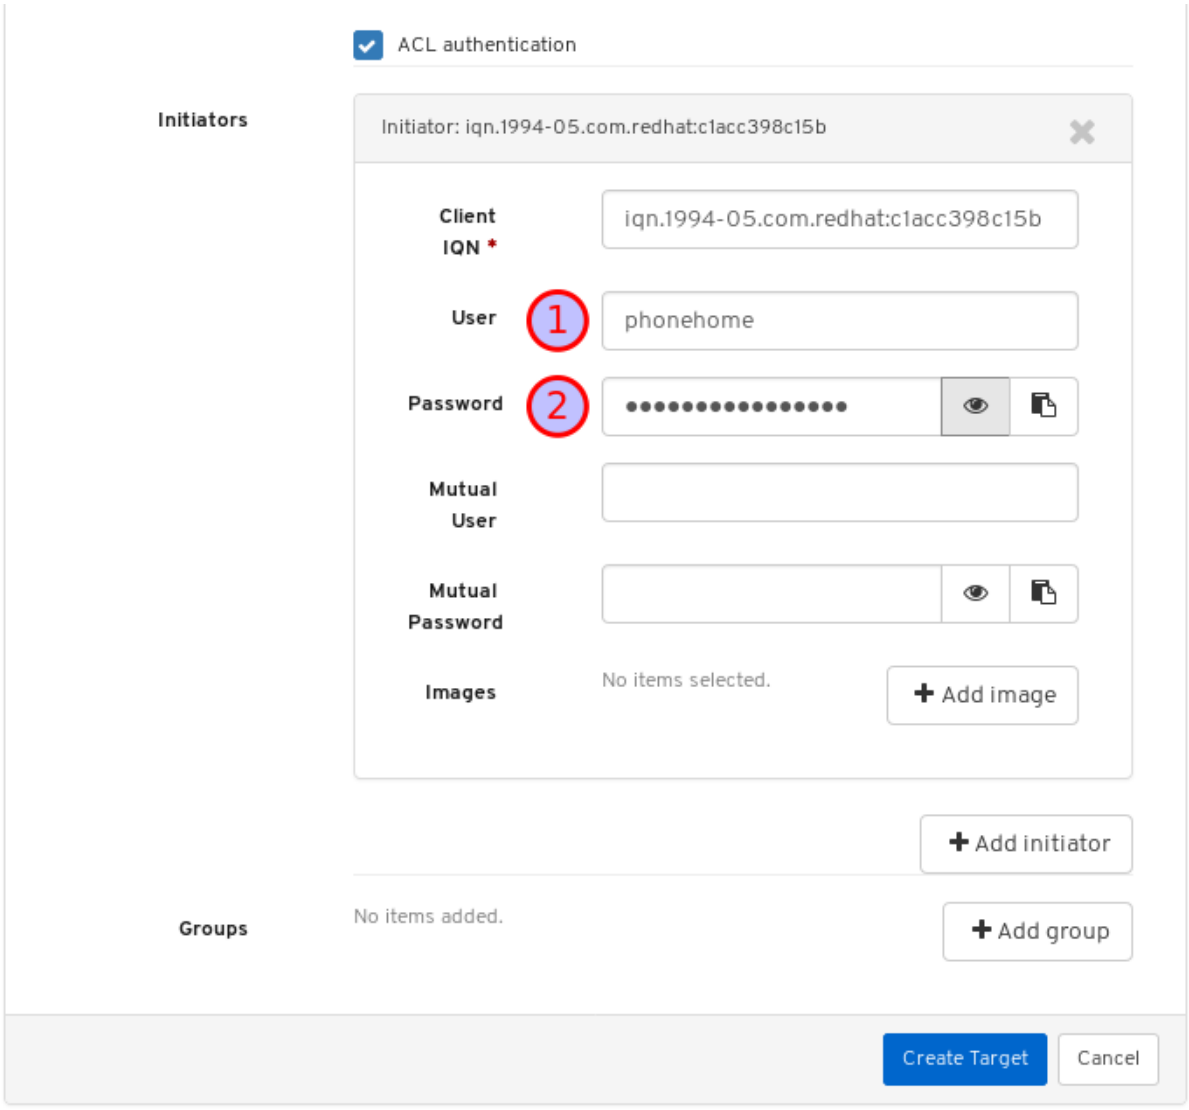

15. Click *Add image* and select an image:

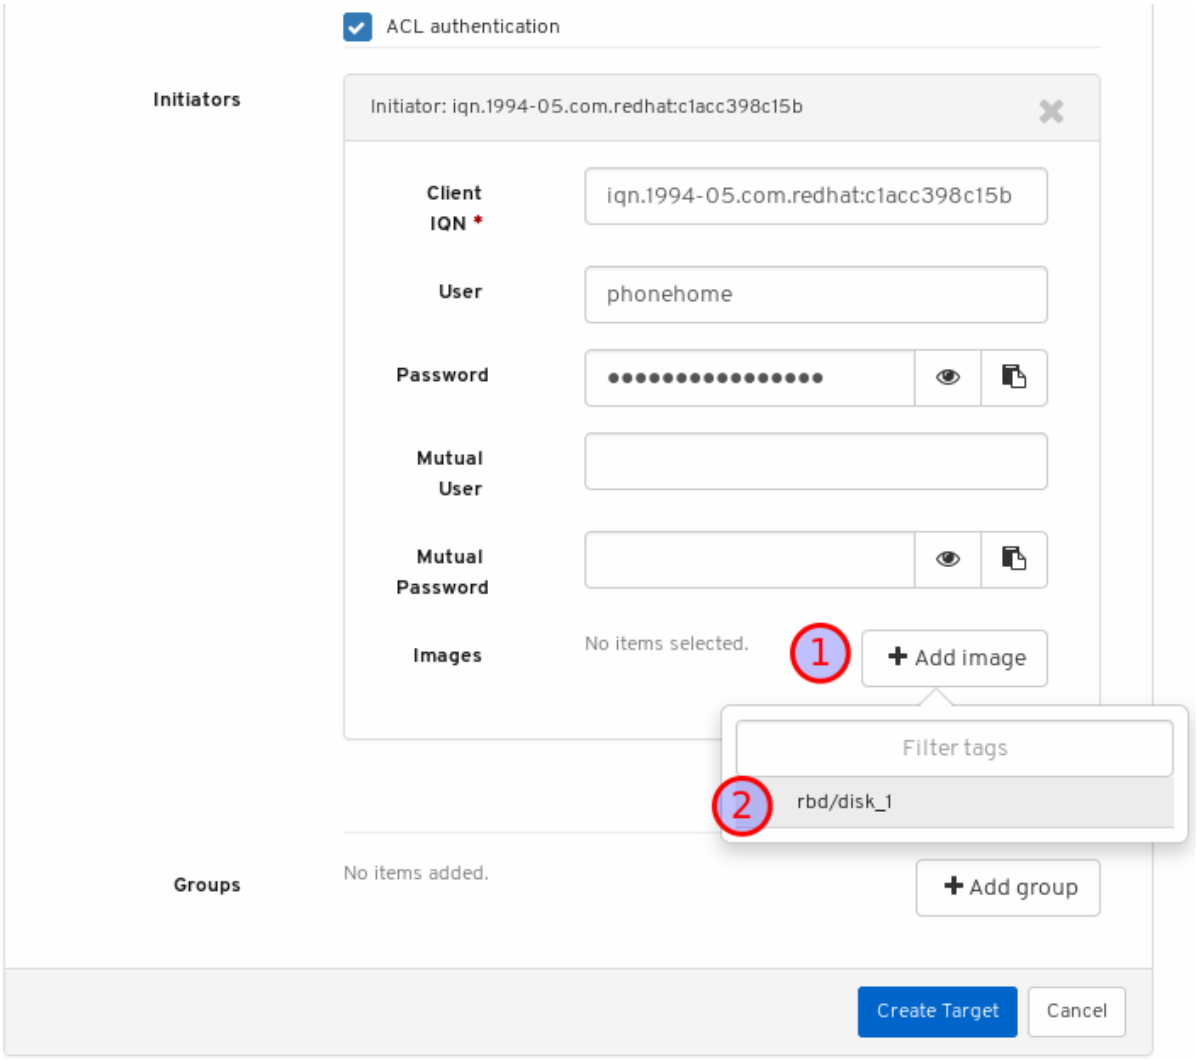

Repeat this step for any additional images.

16. Finish the procedure by clicking the *Create Target* button:

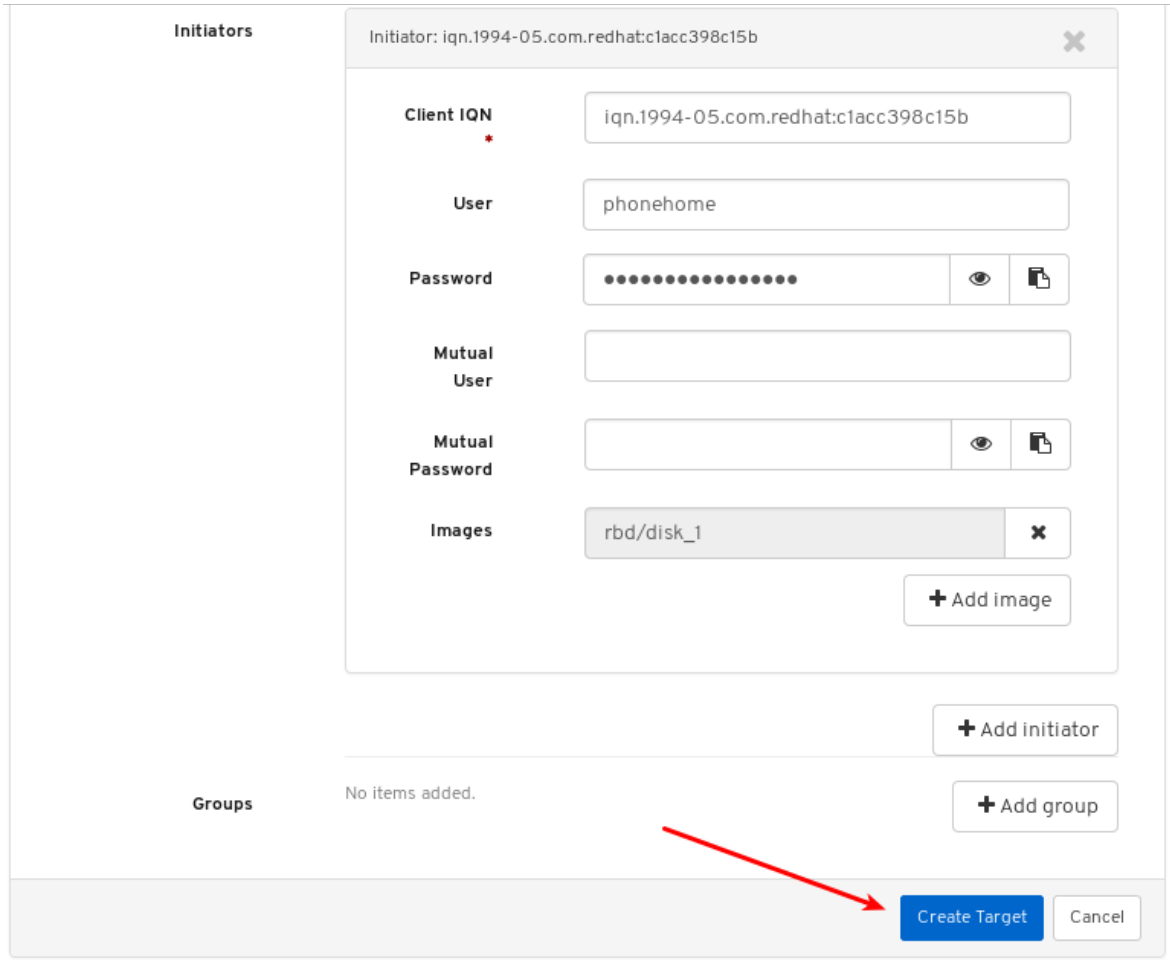

17. Verify the target was added by looking for it on the *Targets* page.
To locate the *Targets* page, follow the procedure Viewing iSCSI targets in the Dashboard guide.

Additional Resources

- For information on how to install the Ceph iSCSI Gateway, see Installing the iSCSI gateway in the Block Device Guide.
- For information on how to add iSCSI gateways to the dashboard, see Manually adding iSCSI gateways to the dashboard in the Dashboard Guide.
- For information on how to enable the dashboard iSCSI feature see Enabling the dashboard iSCSI feature in the dashboard in the Dashboard Guide.
- For information on how to create a pool with the RBD application enabled, see Creating Block Device Pools in the Block Device Guide
- For information on how to create images see Creating block device images in the Block Device Guide.

9.4.5. Viewing iSCSI targets

The dashboard allows you to view iSCSI targets.

Prerequisites

- A running Red Hat Ceph Storage cluster.

- Dashboard is installed.
- The Ceph iSCSI gateway is installed.
- An iSCSI target is created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block* and then click *iSCSI*:

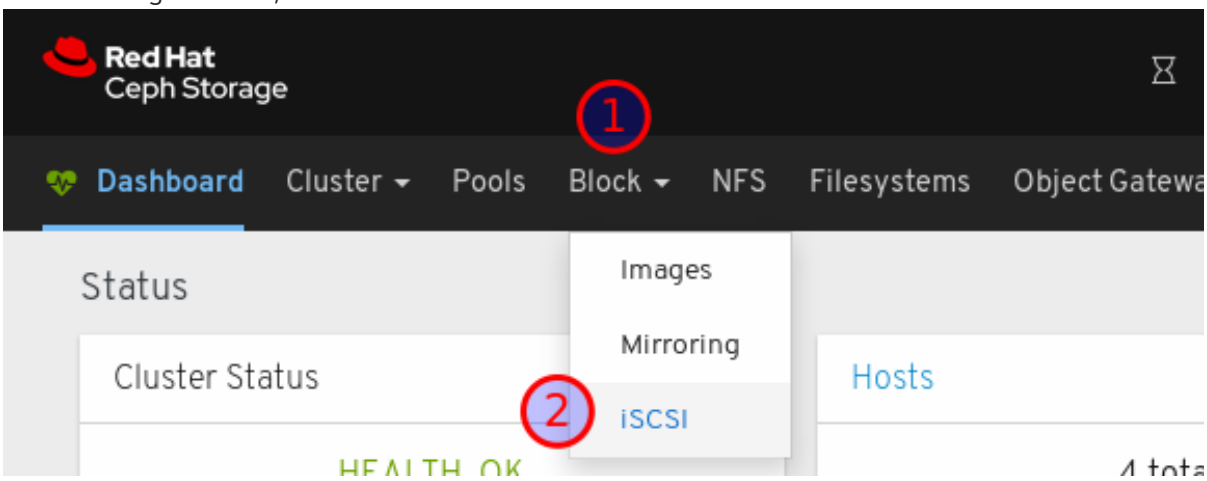

3. Towards the upper left corner of the page, click the *Targets* tab:

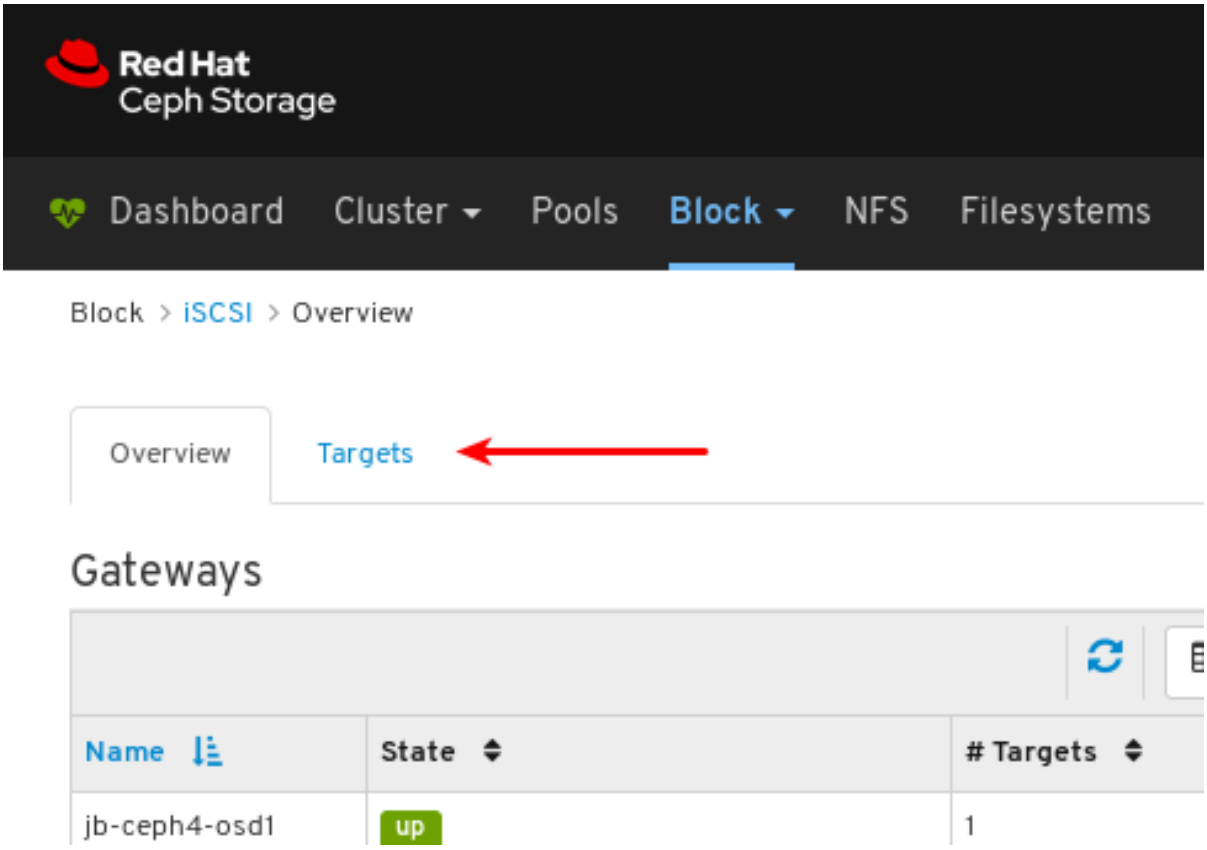

4. To view details about a target, click on its row:

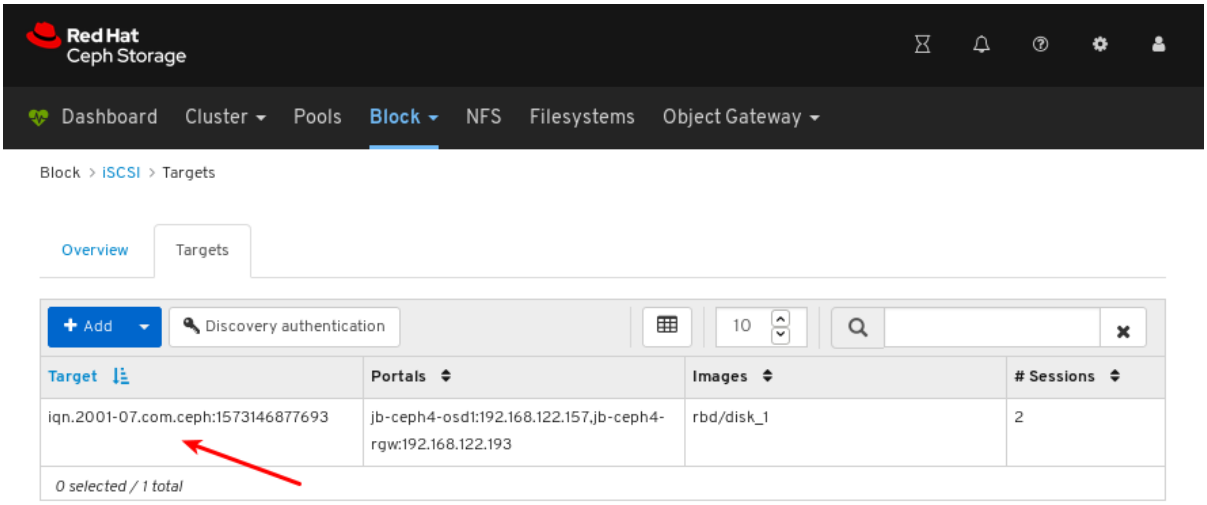

5. You can see the iSCSI topology, including whether an initiator is logged in:

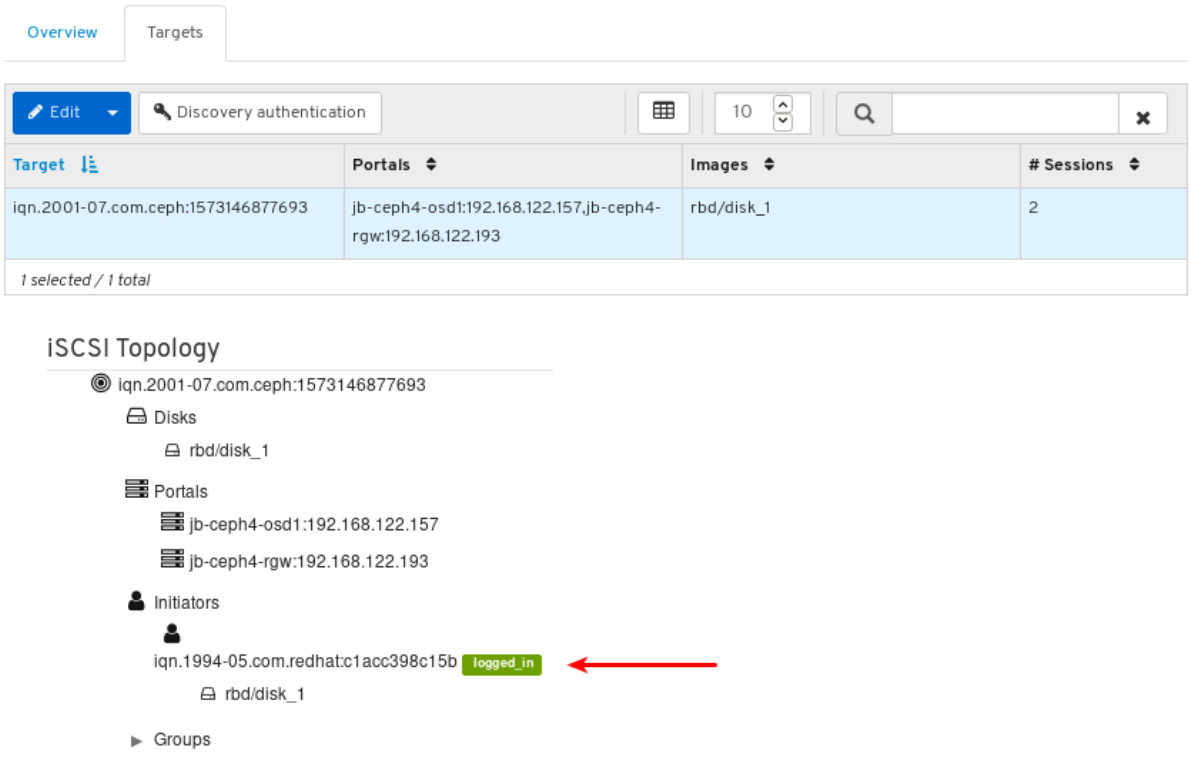

6. Click an object to view detailed information about it:

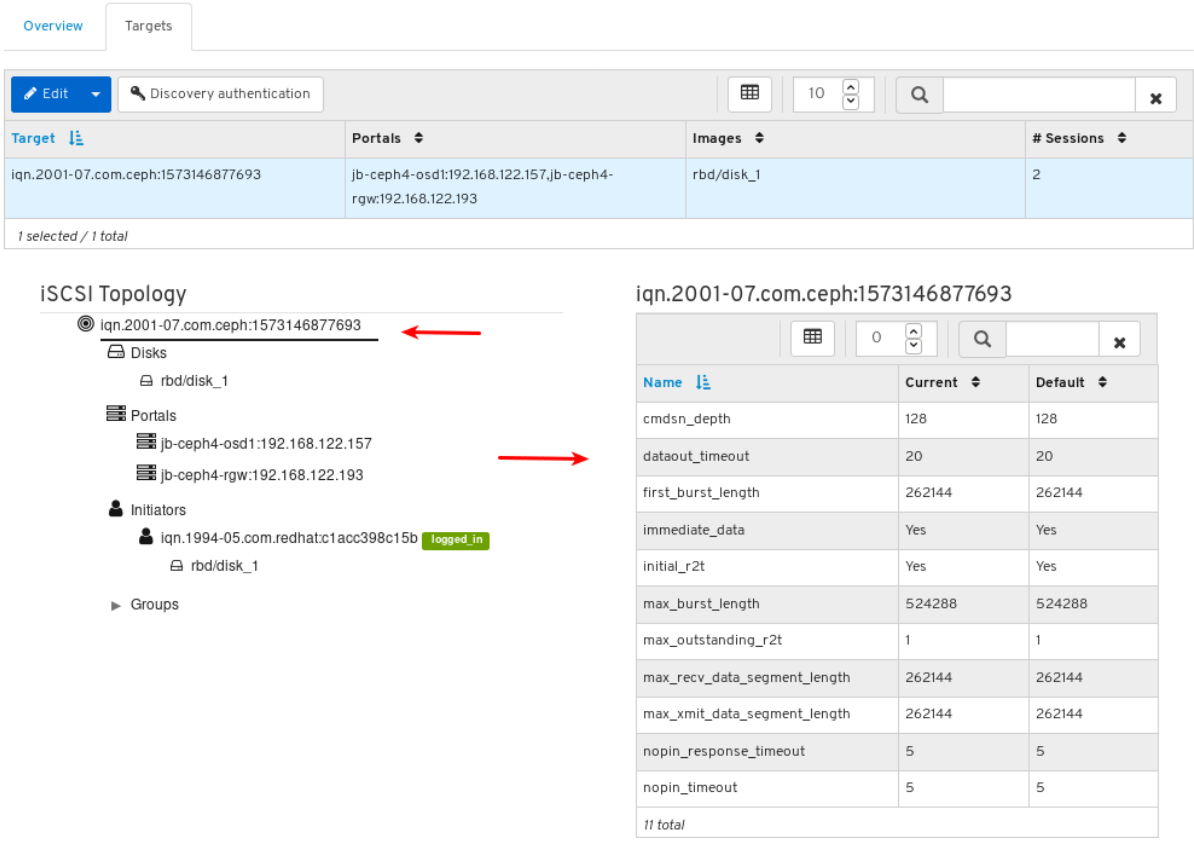

Note: Only some objects display detailed information when clicked.

Additional Resources

- For information on how to install the Ceph iSCSI gateway, see Installing the iSCSI gateway in the Block Device Guide.
- For information on how to create iSCSI targets in dashboard, see Creating iSCSI targets in the Dashboard guide.

9.4.6. Editing iSCSI targets

The dashboard allows you to edit iSCSI targets.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph iSCSI gateway is installed with at least two gateways.
- iSCSI gateways are added to the dashboard.
- The dashboard iSCSI feature is enabled.
- A replicated pool with the RBD application enabled.
- An erasure coded pool with the RBD application enabled.
- An iSCSI target is created.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.
3. Click *iSCSI*:

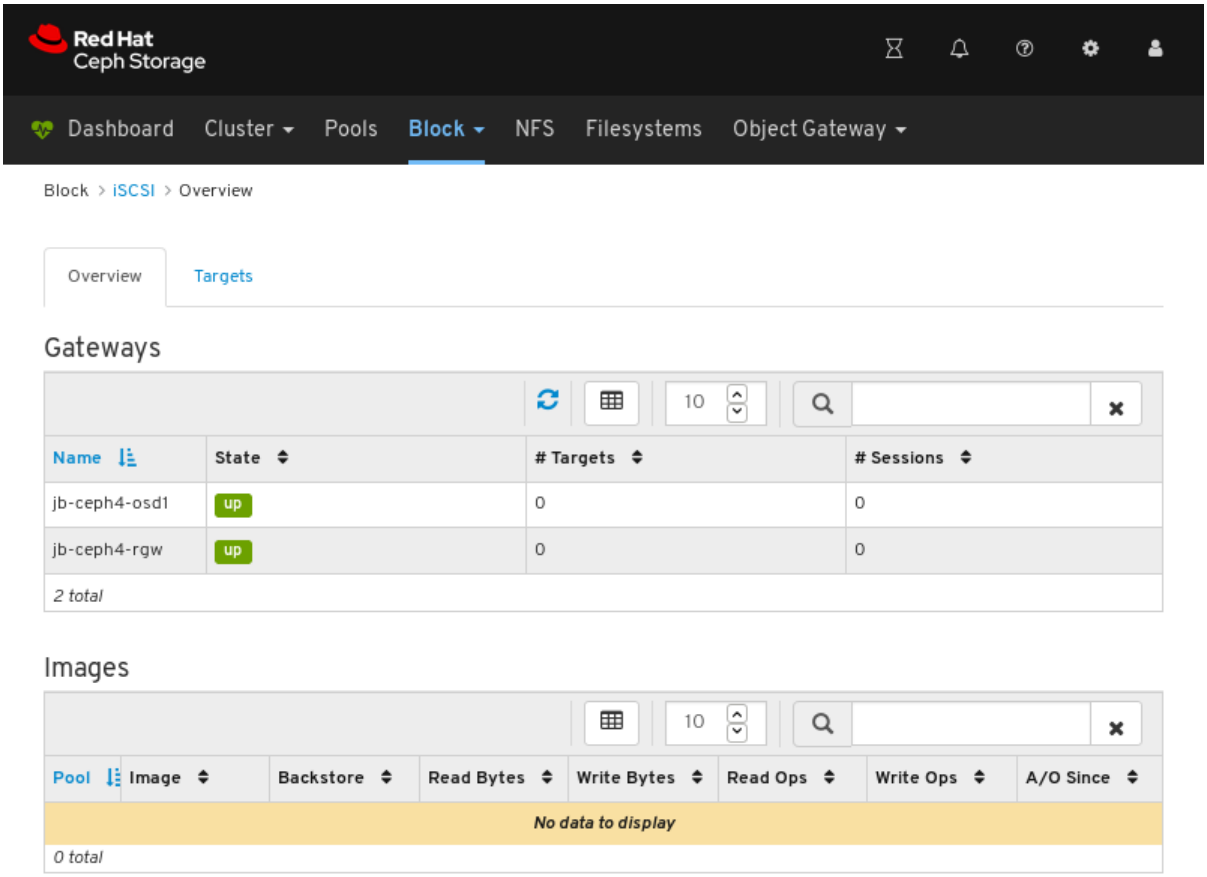

4. Towards the upper left corner of the page, click the *Targets* tab:

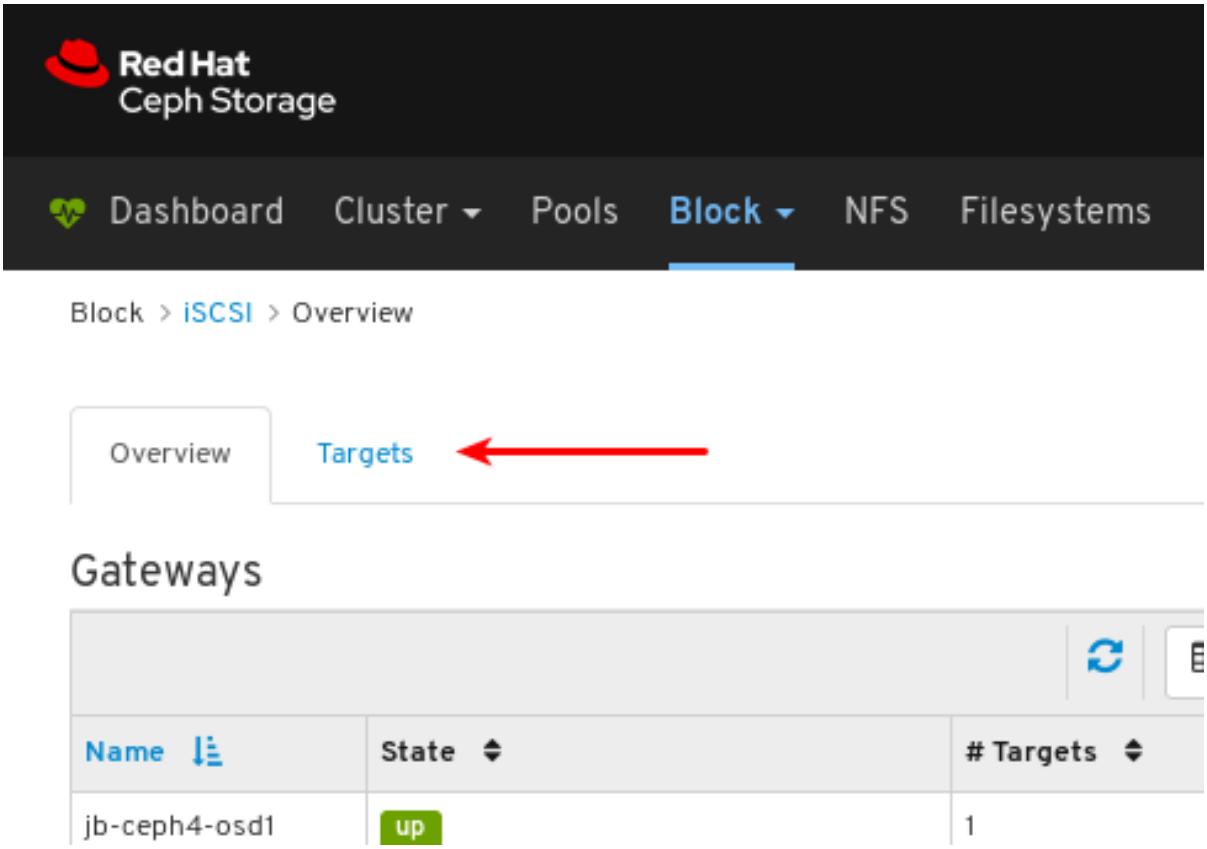

5. To edit a target, click on its row:

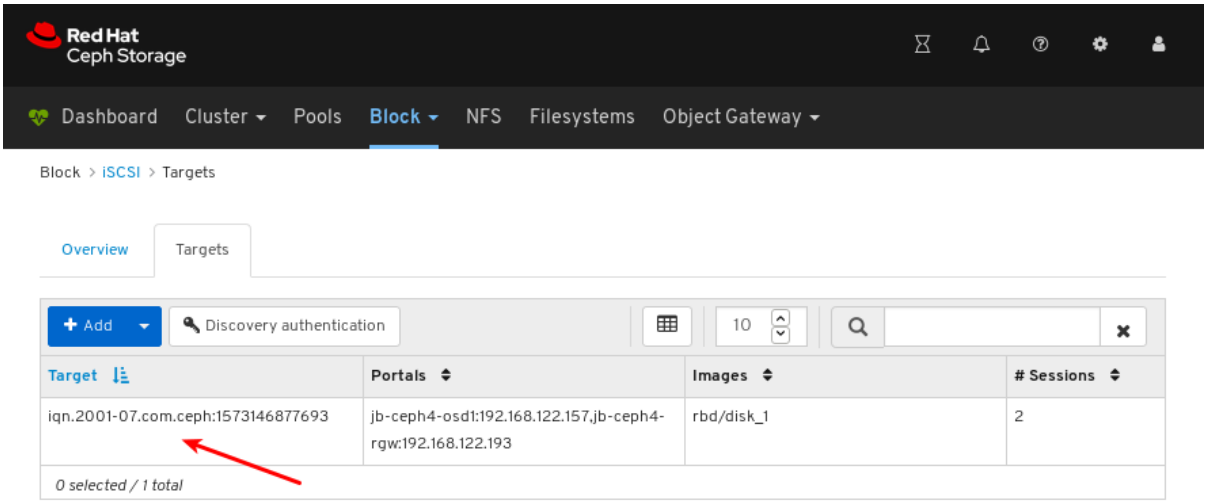

6. Towards the upper left corner of the page, click the *Edit* button.

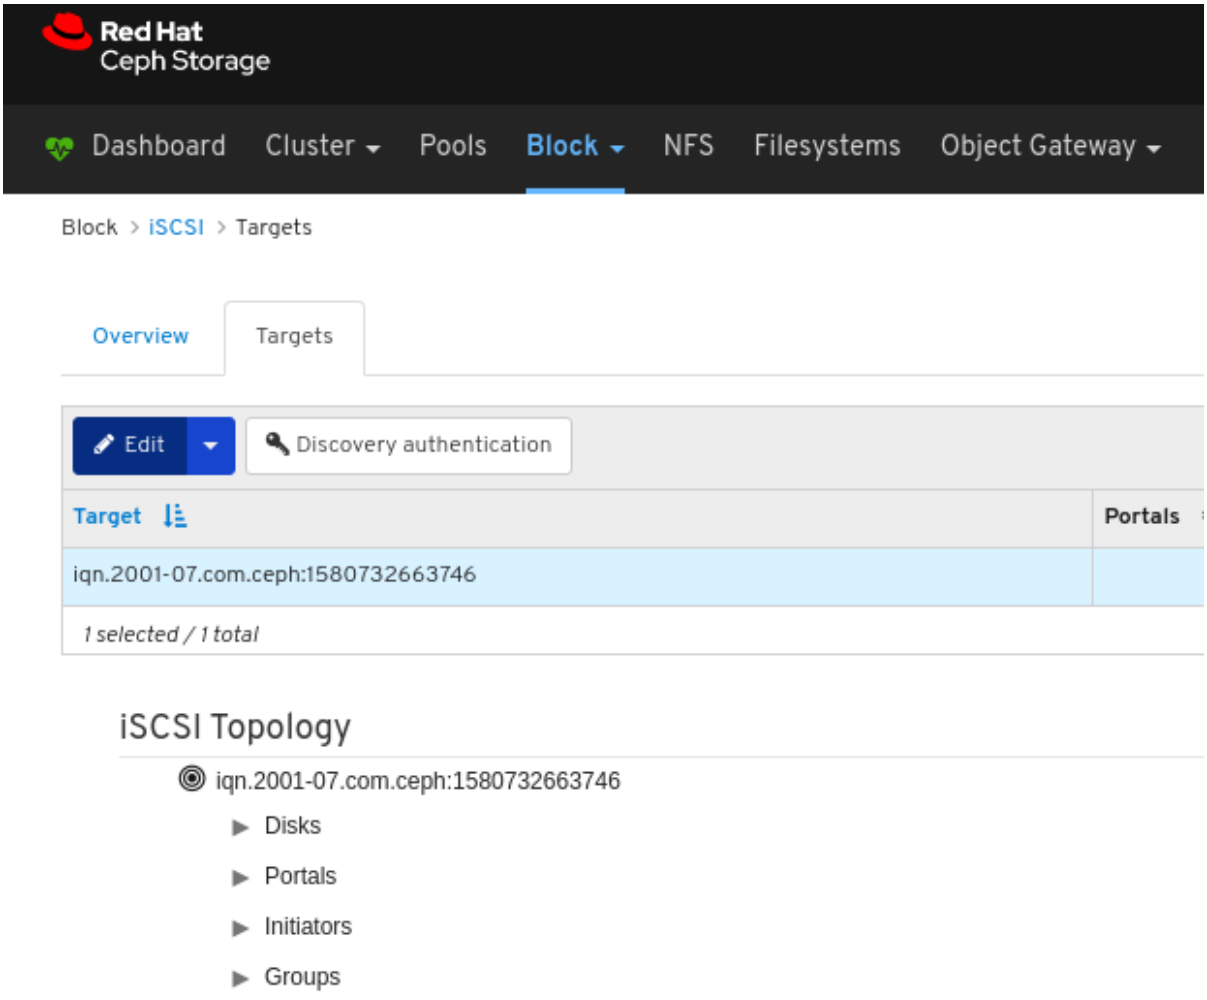

7. Edit the parameters and click the *Edit Target* button.

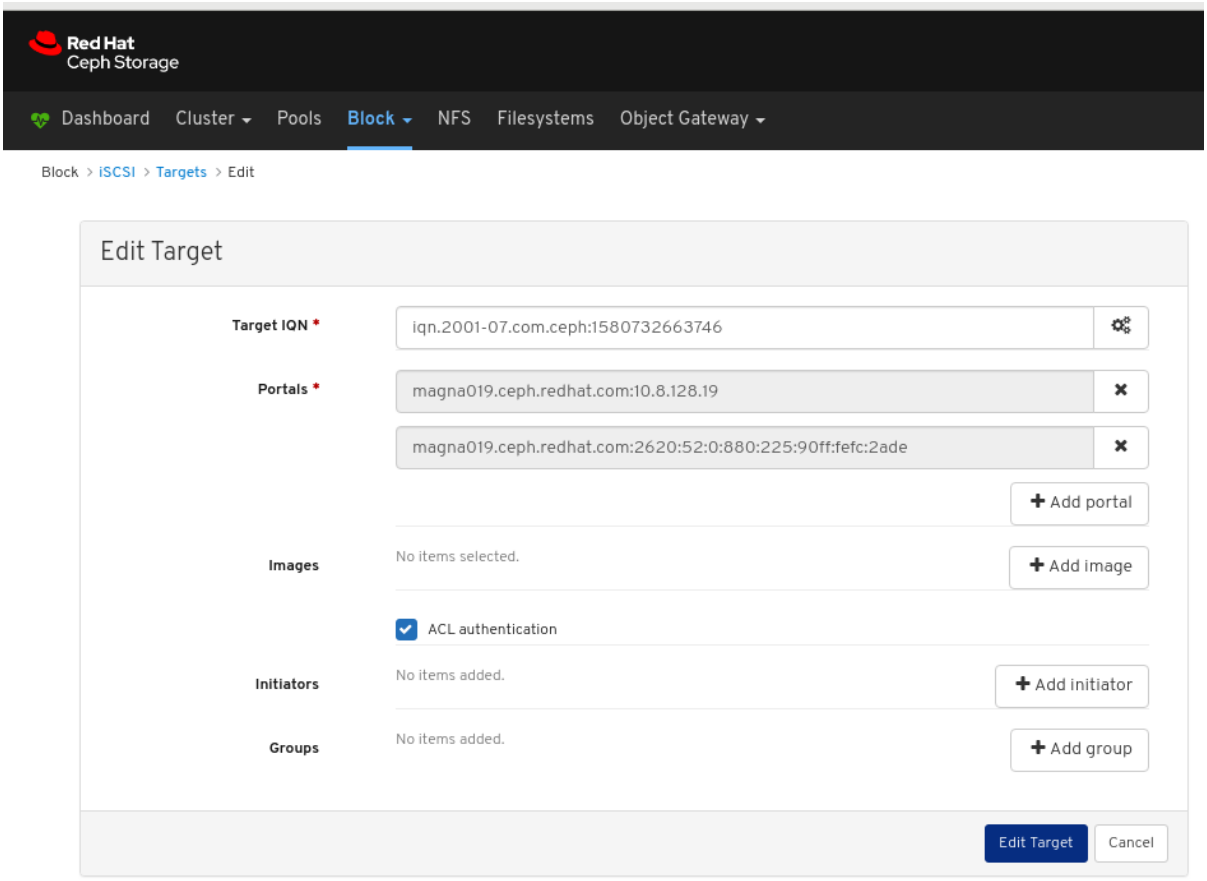

8. Verify the target was edited by looking for it on the *Targets* page.
   To locate the *Targets* page, follow the procedure Viewing iSCSI targets in the Dashboard guide.

Additional Resources

- For information on how to install the Ceph iSCSI Gateway, see Installing the iSCSI gateway in the Block Device Guide.
- For information on how to add iSCSI gateways to the dashboard, see Manually adding iSCSI gateways to the dashboard in the Dashboard Guide.
- For information on how to enable the dashboard iSCSI feature see Enabling the dashboard iSCSI feature in the dashboard in the Dashboard Guide.
- For information on how to create a pool with the RBD application enabled, see Creating Block Device Pools in the Block Device Guide
- For information on how to create images see Creating block device images in the Block Device Guide.

9.4.7. Deleting iSCSI targets

The dashboard allows you to delete iSCSI targets.

Prerequisites

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

- The Ceph iSCSI gateway is installed with at least two gateways.
- iSCSI gateways are added to the dashboard.
- The dashboard iSCSI feature is enabled.
- An iSCSI target is created.
- Disconnect all iSCSI Initiators. See Disconnecting iSCSI initiators.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.
3. Click *iSCSI*:

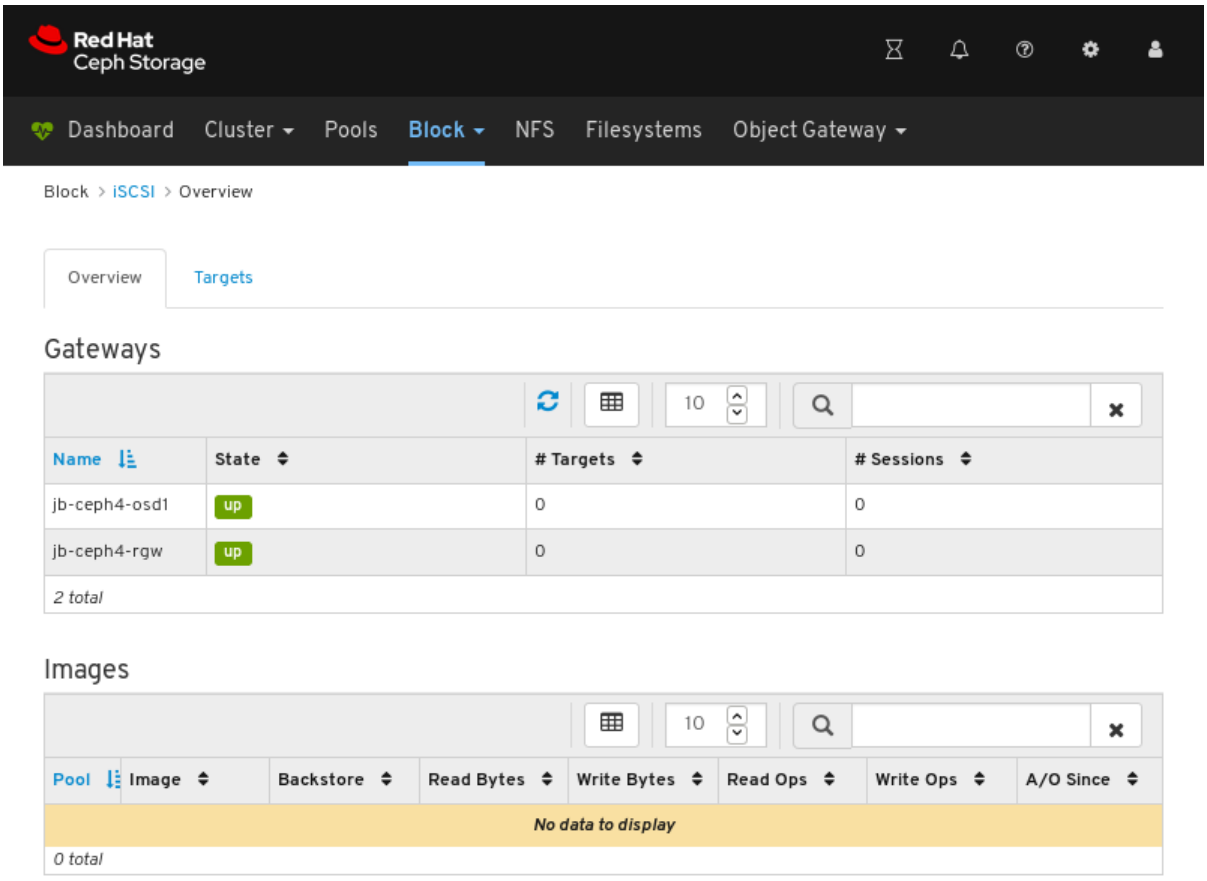

4. Towards the upper left corner of the page, click the *Targets* tab:

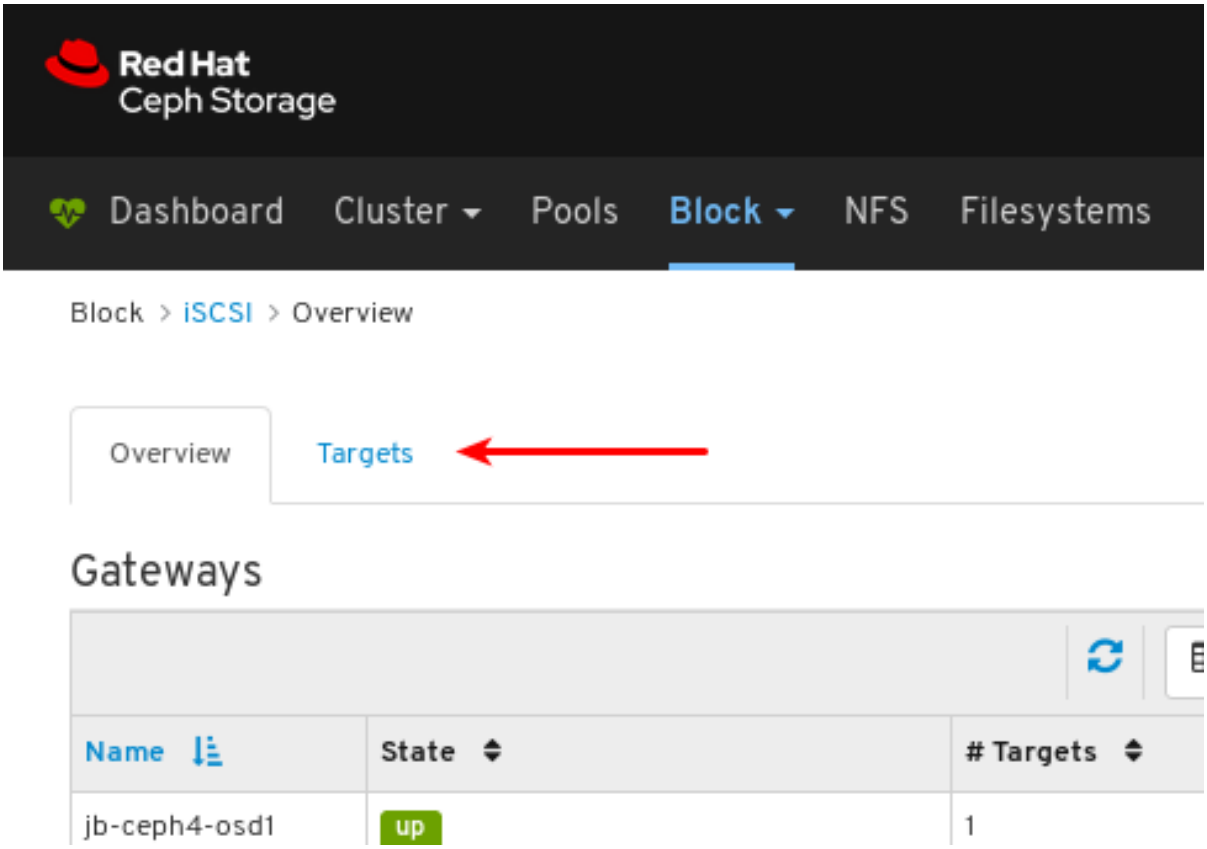

5. To delete a target, click on its row:

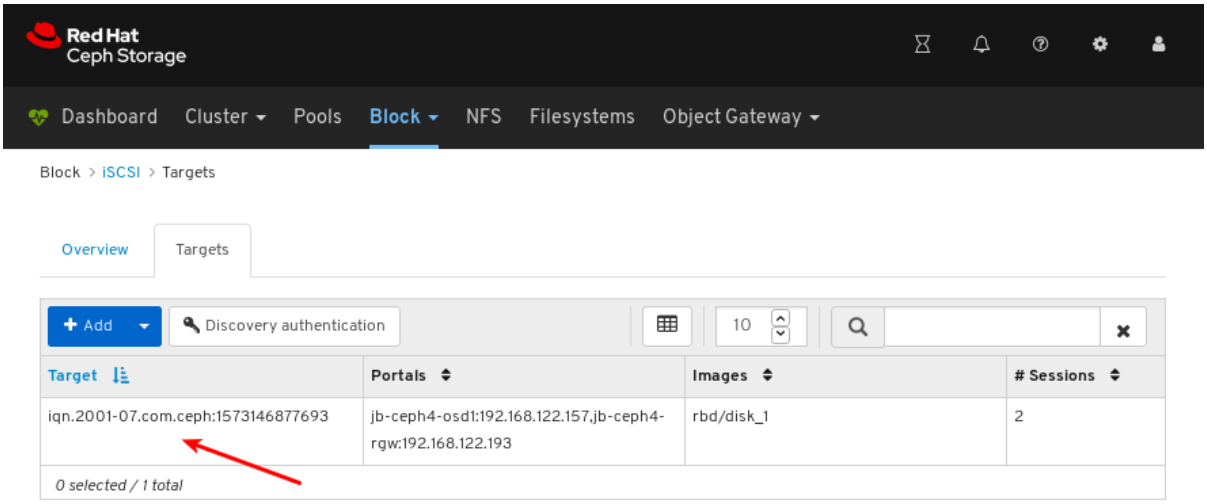

6. Towards the upper left corner of the page, click the *Edit* drop down.

7. From the drop-down, select *Delete*:

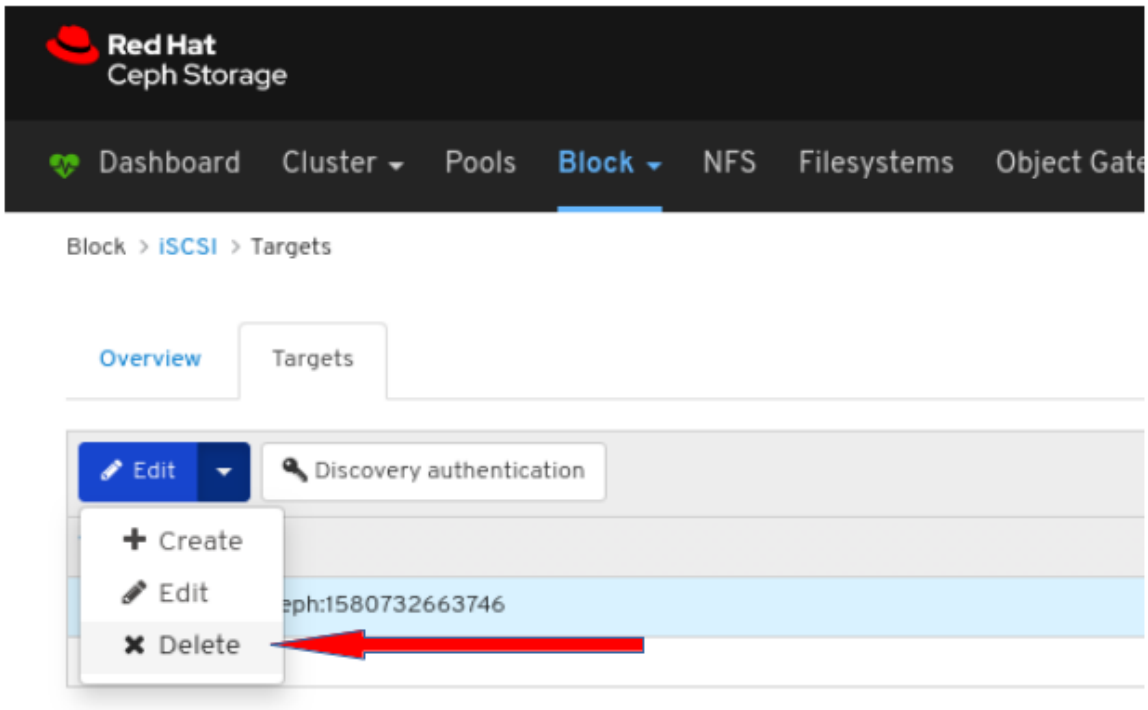

8. Click the *Yes, I am sure* box and then Click *Delete iSCSI* to save the settings.

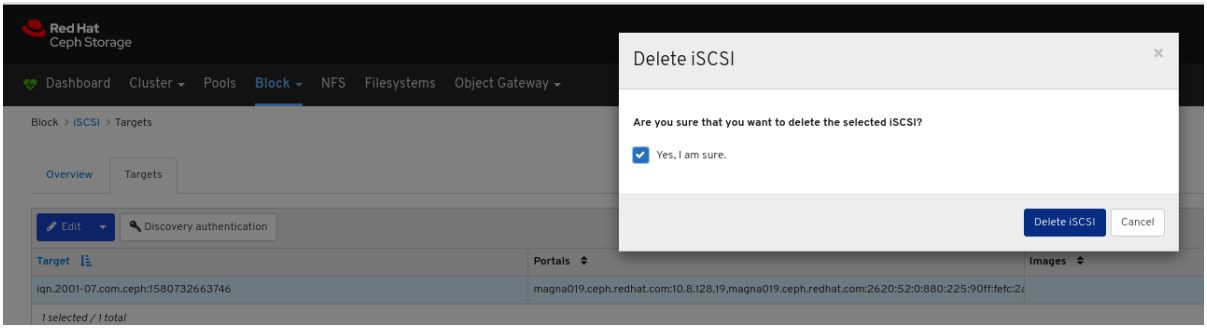

Additional Resources

- For information on how to install the Ceph iSCSI Gateway, see Installing the iSCSI gateway in the Block Device Guide.
- For information on how to add iSCSI gateways to the dashboard, see Manually adding iSCSI gateways to the dashboard in the Dashboard Guide.
- For information on how to disconnect iSCSI initiators see Disconnecting iSCSI initiators in the Block Device Guide.

9.4.8. Setting Discovery Authentication

The dashboard allows Discovery Authentication by using CHAP/CHAP_MUTUAL.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- The Ceph iSCSI gateway is installed with at least two gateways.

- iSCSI gateways are added to the dashboard.
- The dashboard iSCSI feature is enabled.
- A replicated pool with the RBD application enabled.
- An erasure coded pool with the RBD application enabled.

Procedure

1. Log in to the Dashboard.
2. On the navigation bar, click *Block*.
3. Click *iSCSI*:

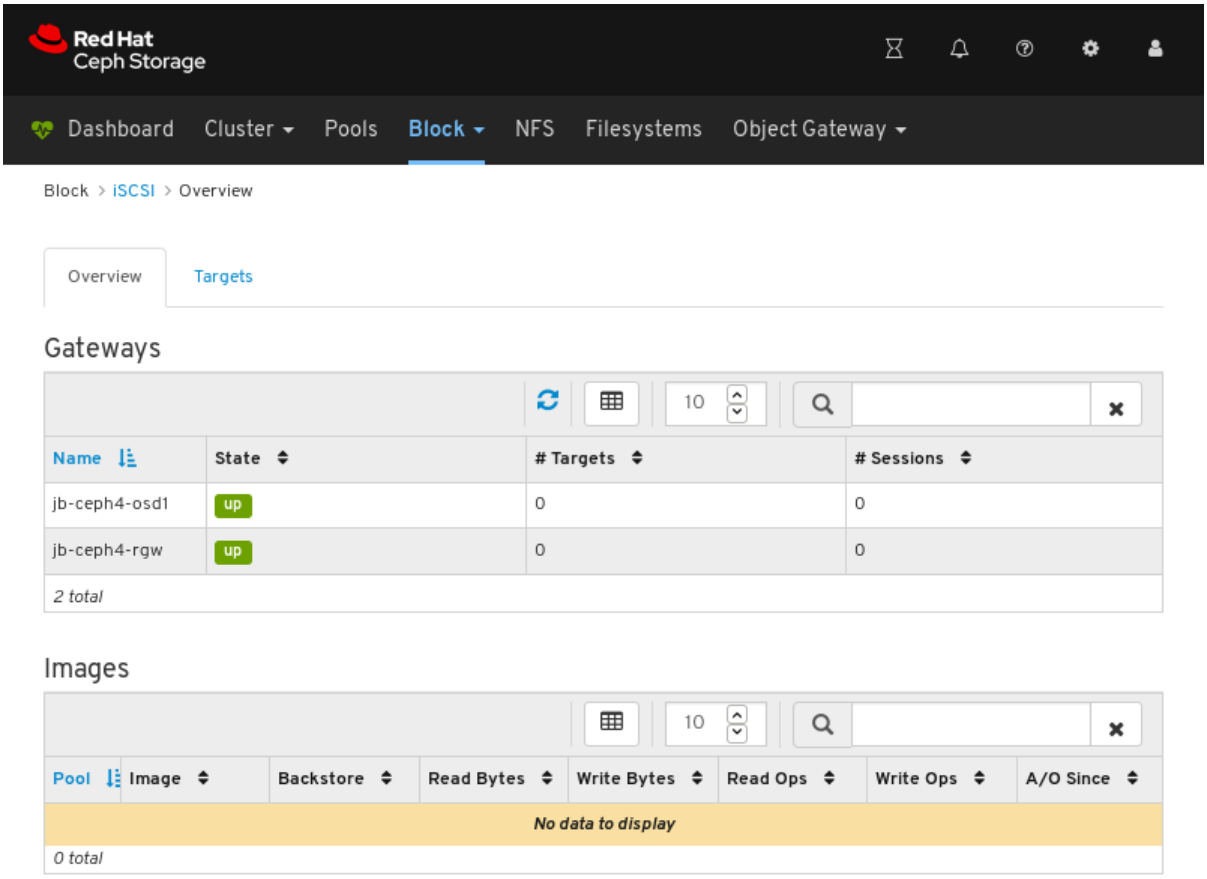

4. Towards the upper left corner of the page, click the *Targets* tab.

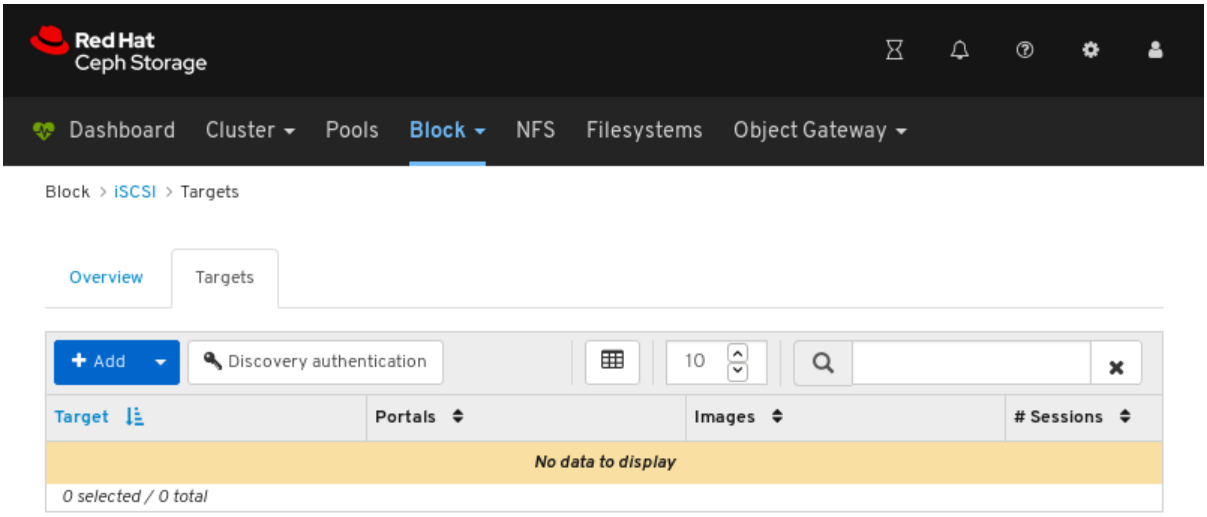

5. Towards the upper left corner of the page, click the *Discovery authentication* button.

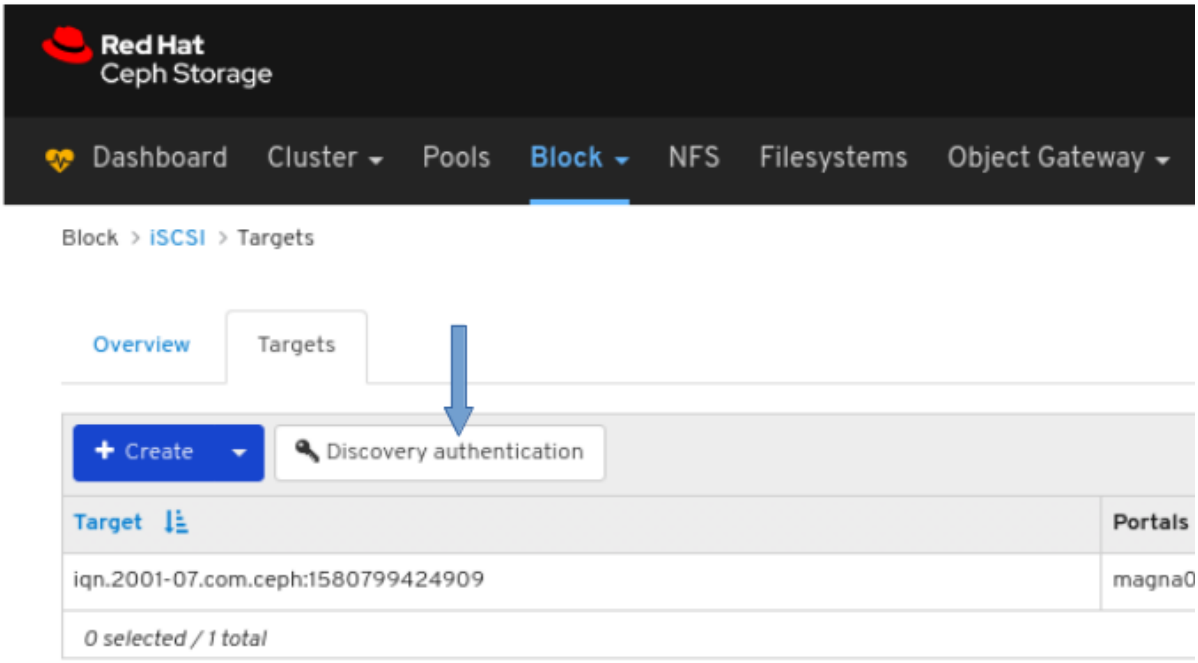

6. in *Discovery Authentication* window, provide the details and then Click the *Submit* button.

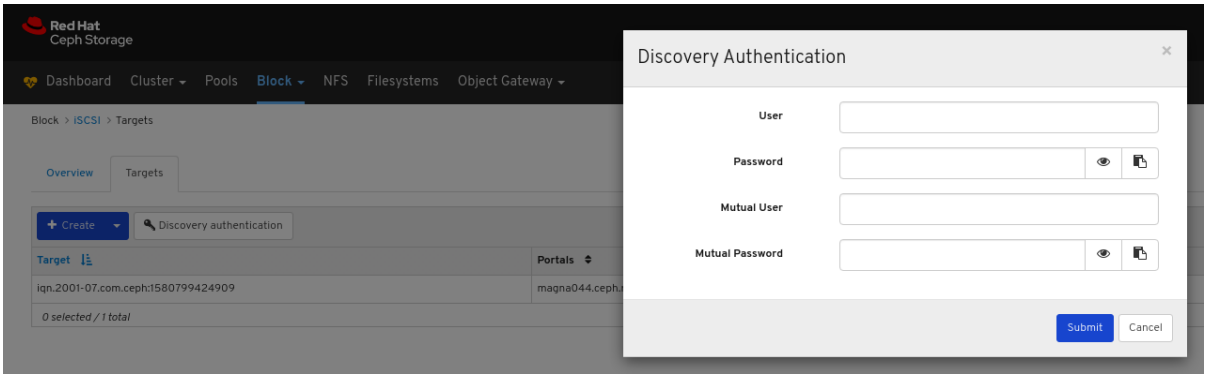

9.5. QUALITY OF SERVICE CONFIGURATION

As a storage administrator, you can use Quality of Service (QoS) limits to prioritize or deprioritize the performance of pools or images so all images get the resources they need to meet specific business needs.

9.5.1. Prerequisites

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

9.5.2. Configuring Quality of Service on an existing image

As a storage administrator, you can use Quality of Service (QoS) limits to priortize or deprioritize the performance of an existing image.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A replicated pool with the RBD application enabled.
- An erasure coded pool with the RBD application enabled.
- An existing image.

Procedure

1. Log in to the dashboard.
2. On the navigation bar, click *Block* and then click *Images*:

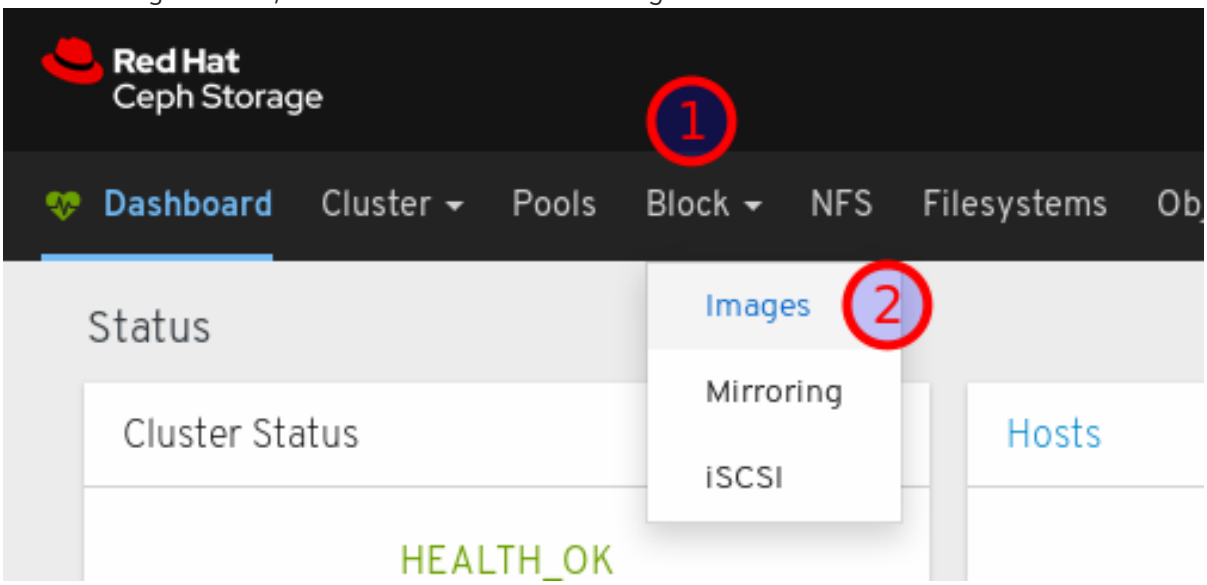

3. Click the row of an image to select it for editing:

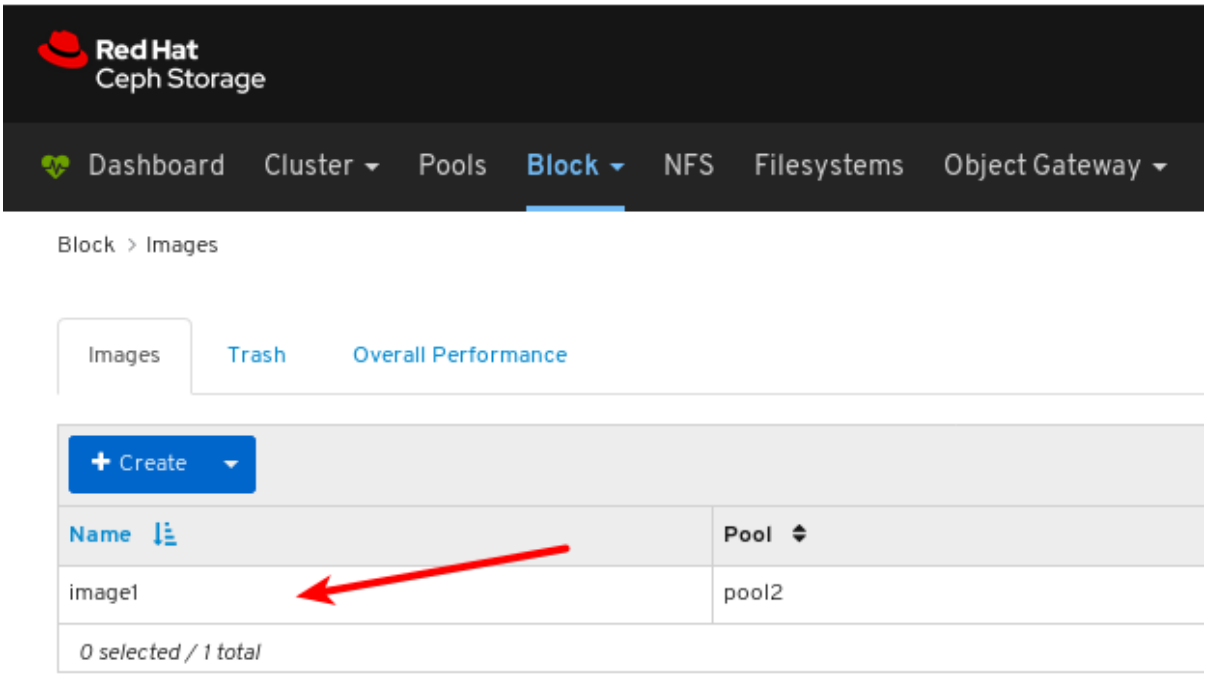

4. Click the *Edit* button:

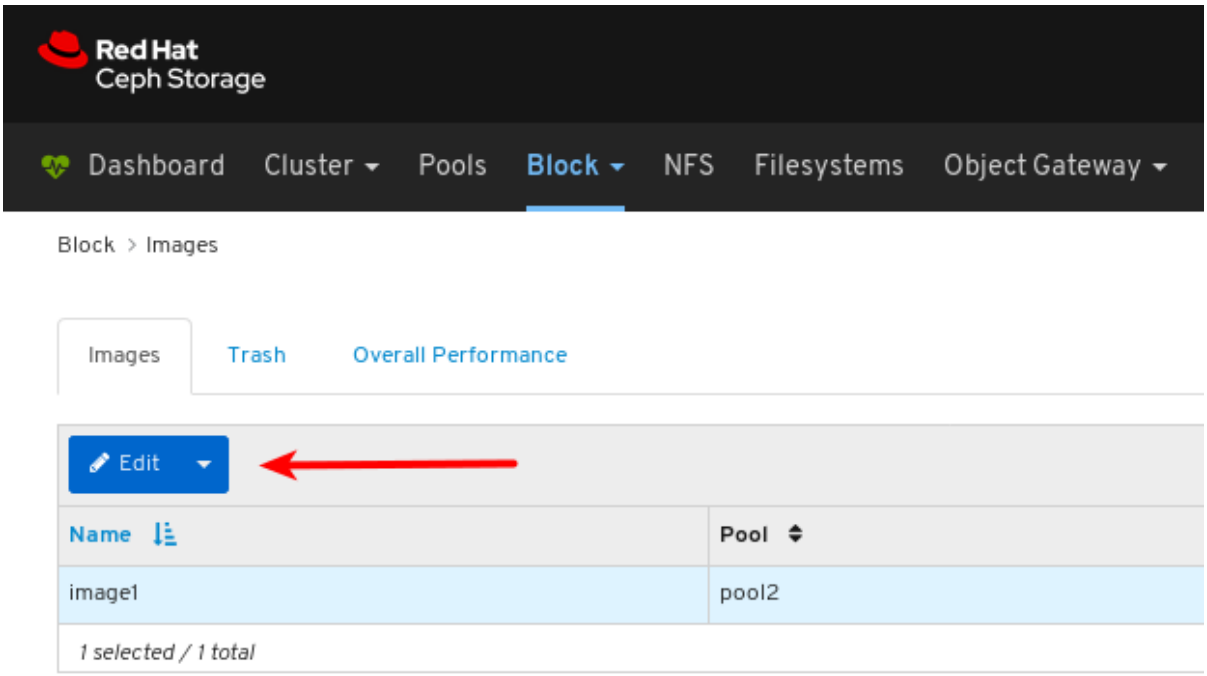

5. Click *Advanced* towards the bottom right corner of the dialog:

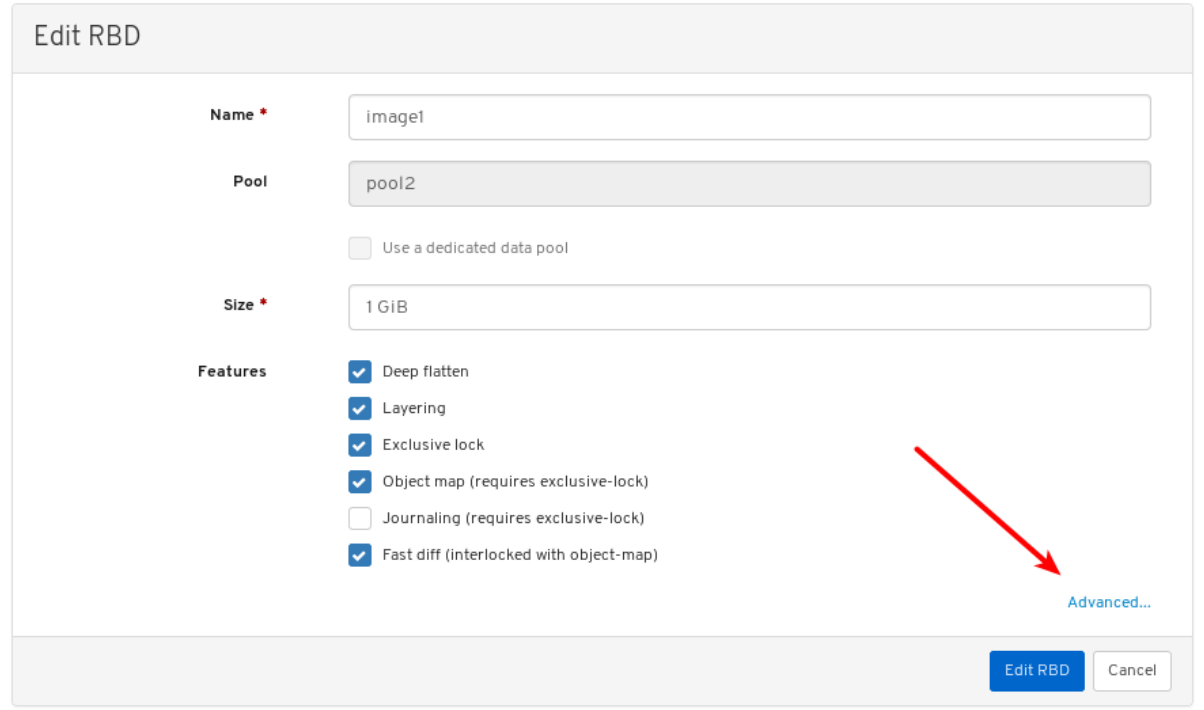

6. Click the plus symbol next to *Quality of Service* to open the QoS settings:

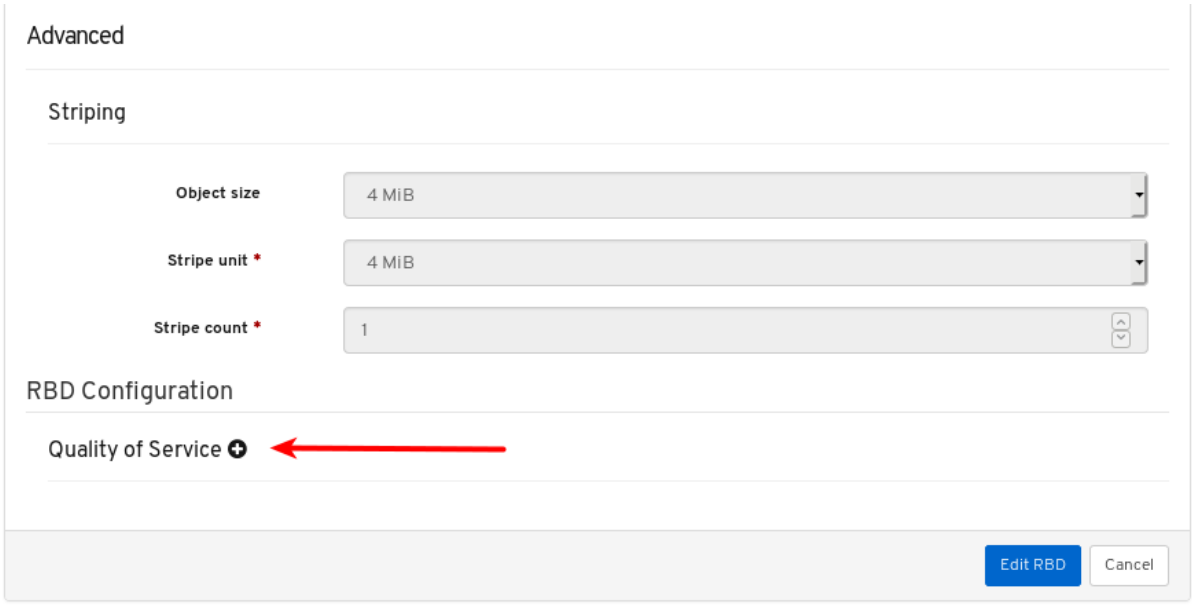

7. Optional: Click the question mark symbol next to an individual setting to find out more about it.

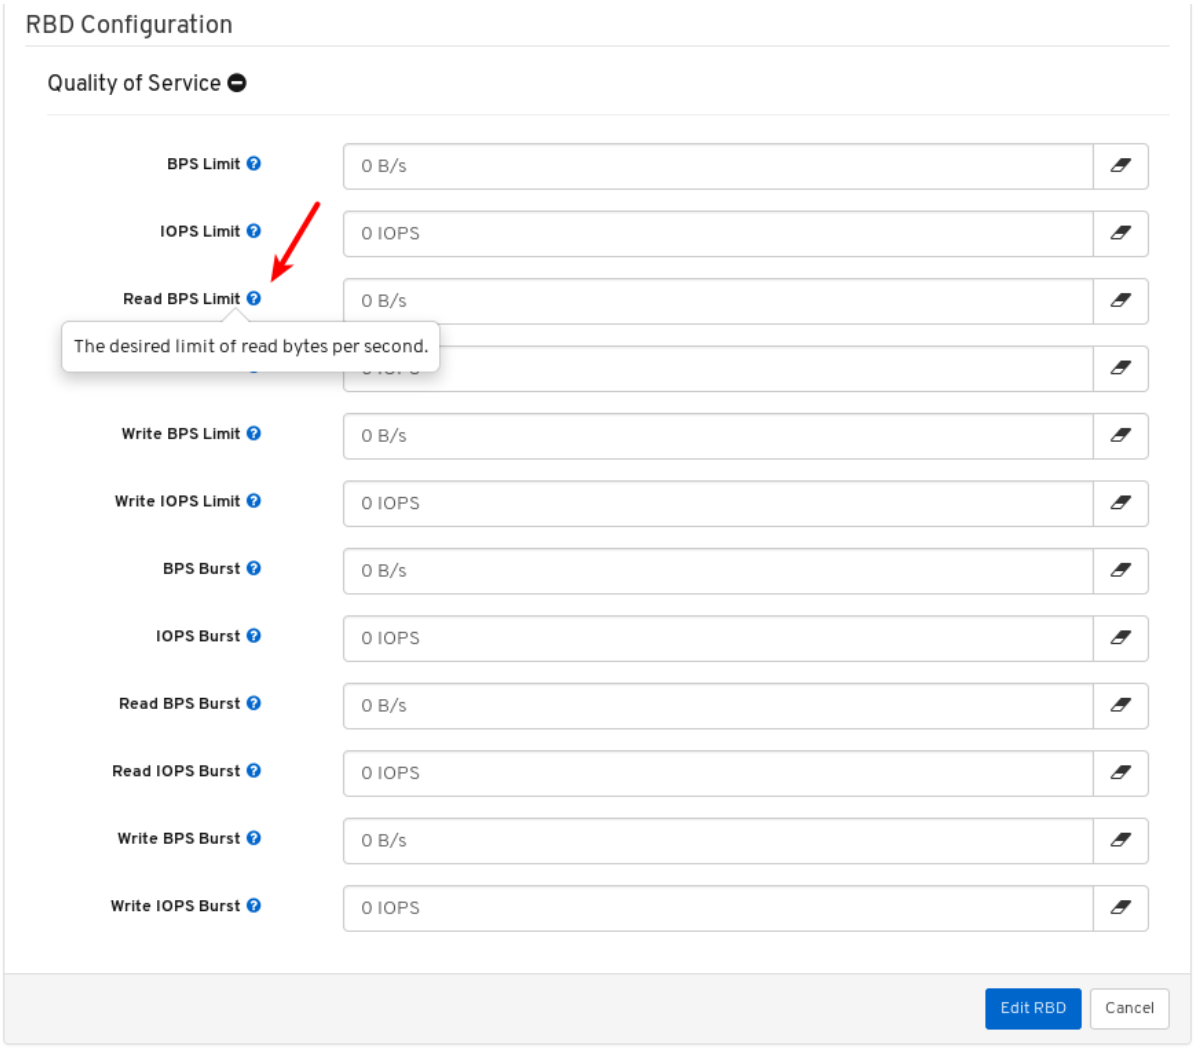

8. Enter or edit values for the QoS settings you want to change:

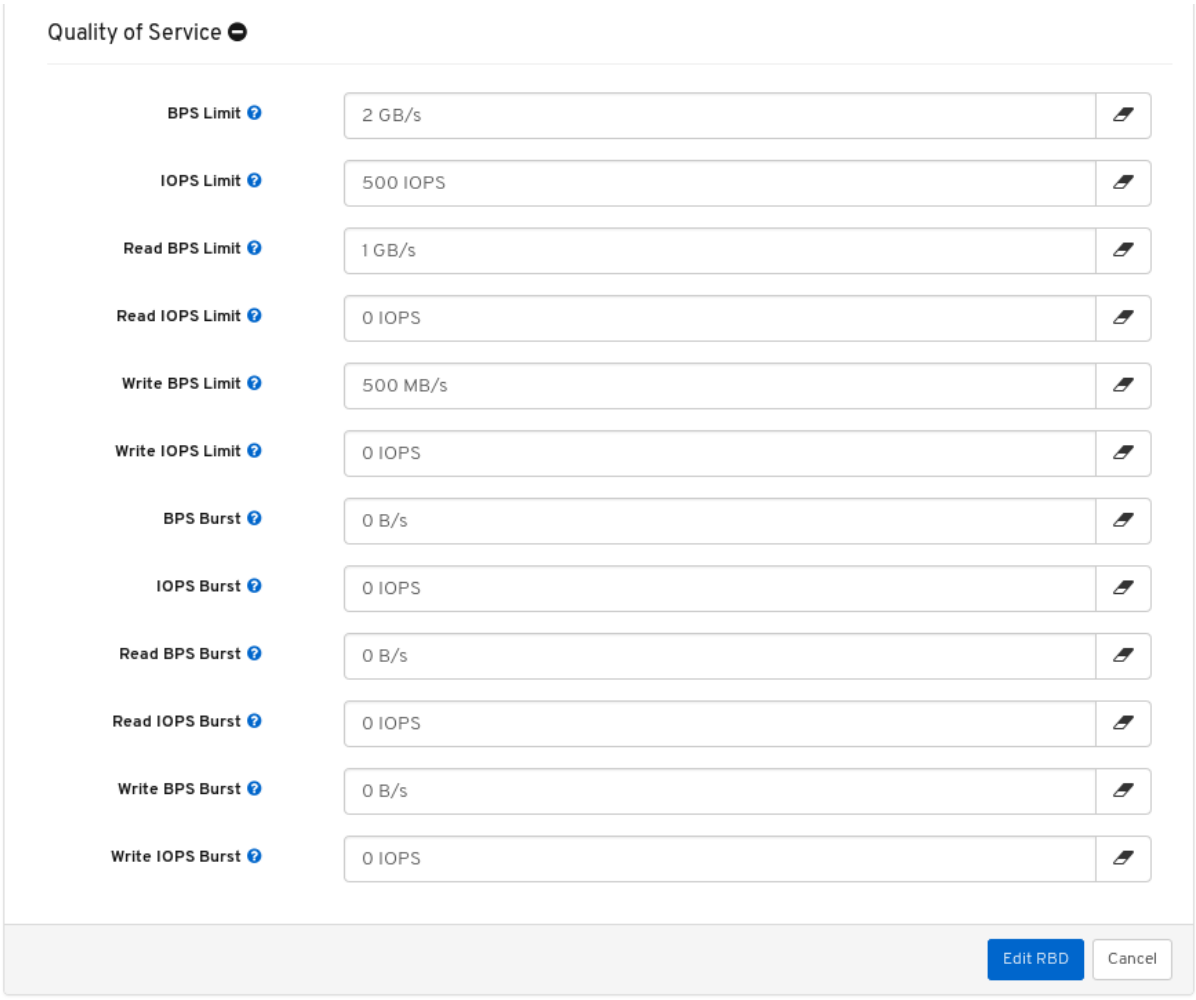

9. Optional: Click the eraser symbol for any setting to remove the local value and inherit the value from the parent pool or global configuration.

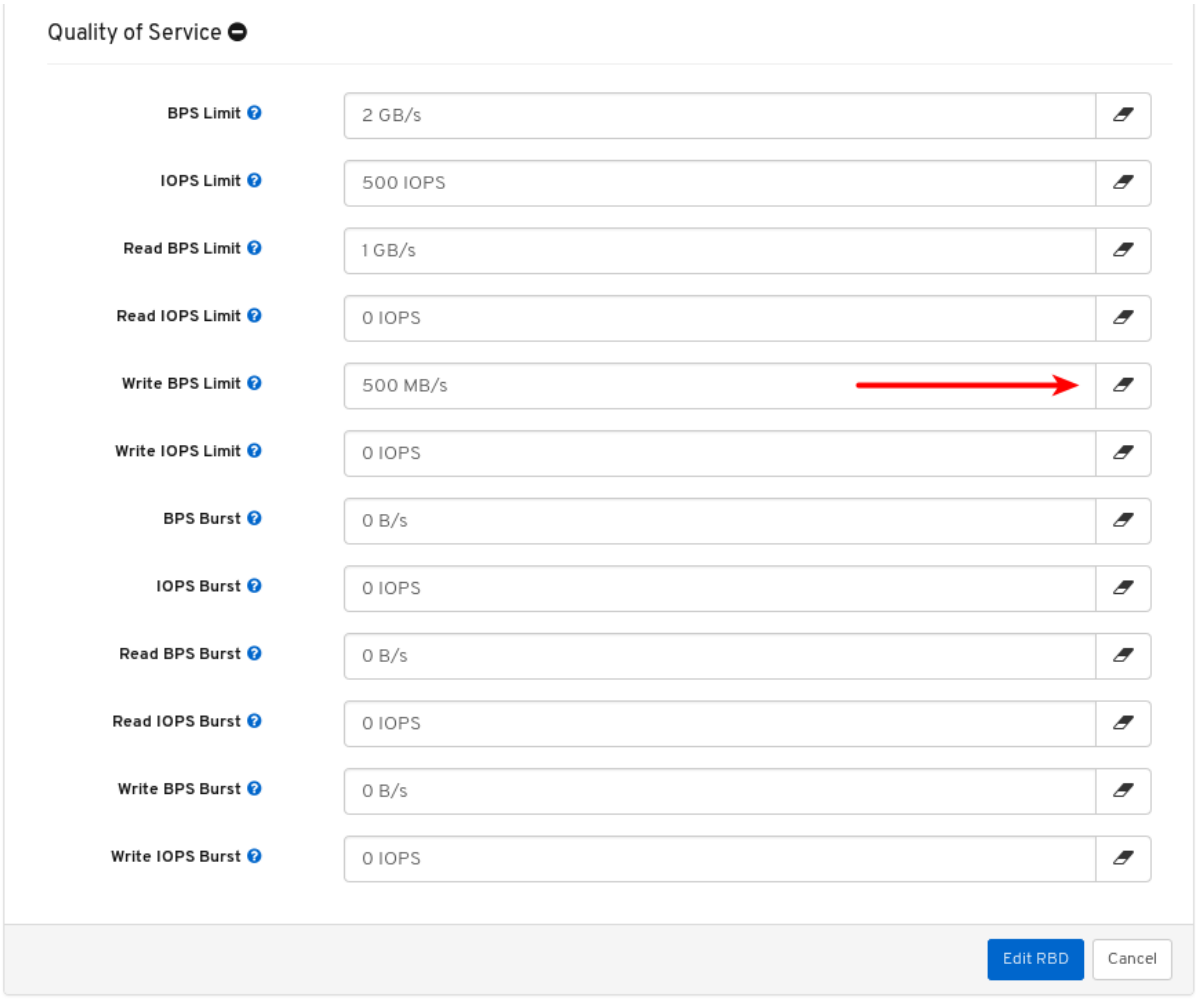

The field for the setting is disabled to indicate it is inheriting the value from the parent.

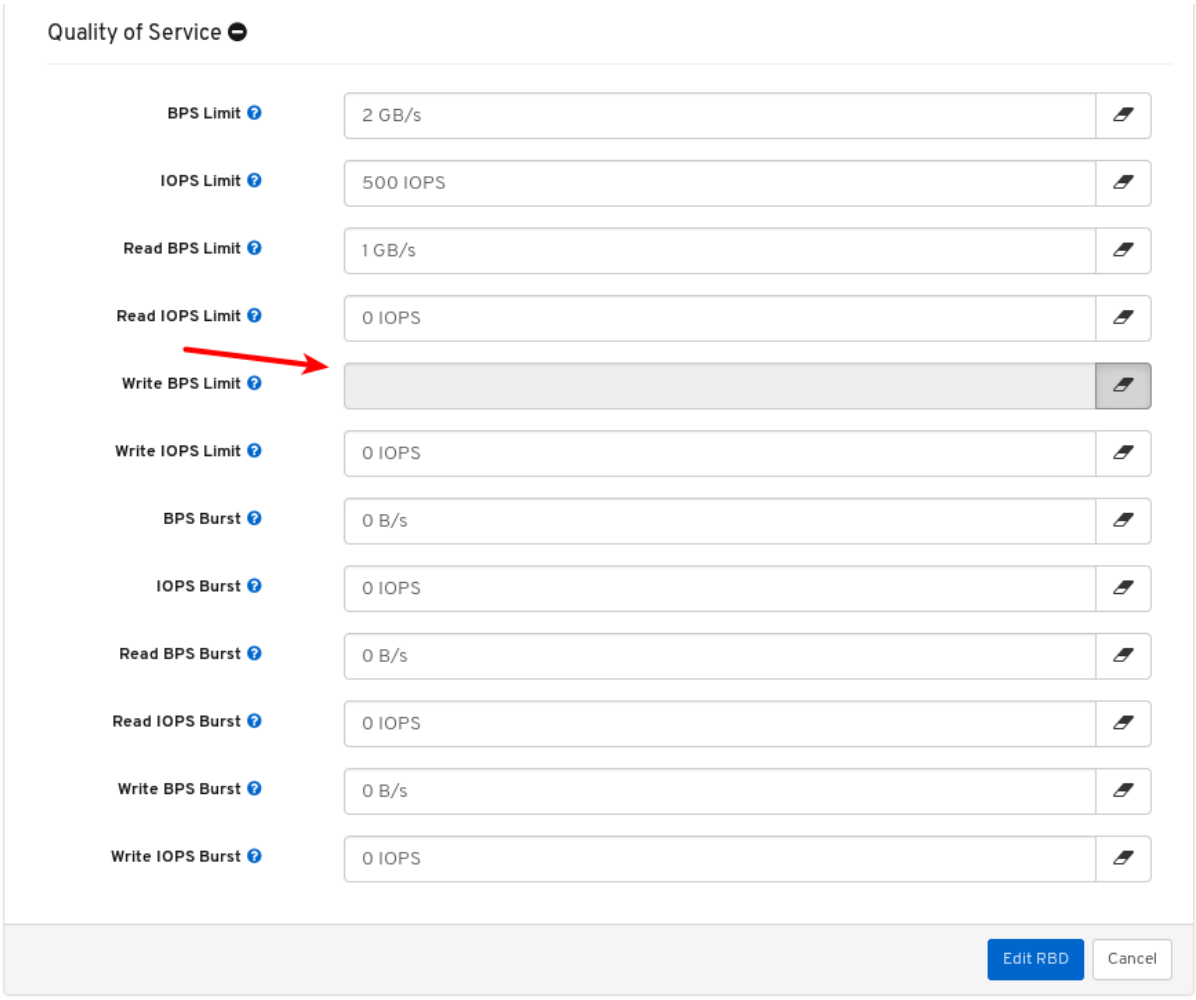

10. Click the *Edit* button to save the changes.

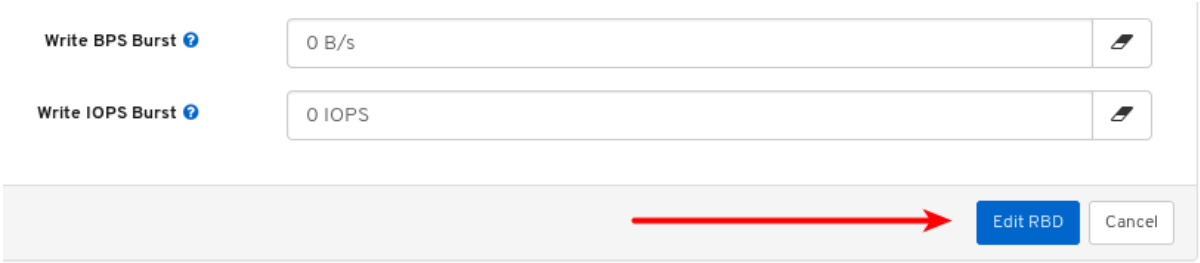

9.5.3. Configuring Quality of Service on an existing pool

As a storage administrator, you can configure Quality of Service (QoS) on an existing pool.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A replicated pool with the RBD application enabled.
- An erasure coded pool with the RBD application enabled.

Procedure

1. Log in to the dashboard.

2. On the navigation bar, click *Pools*.

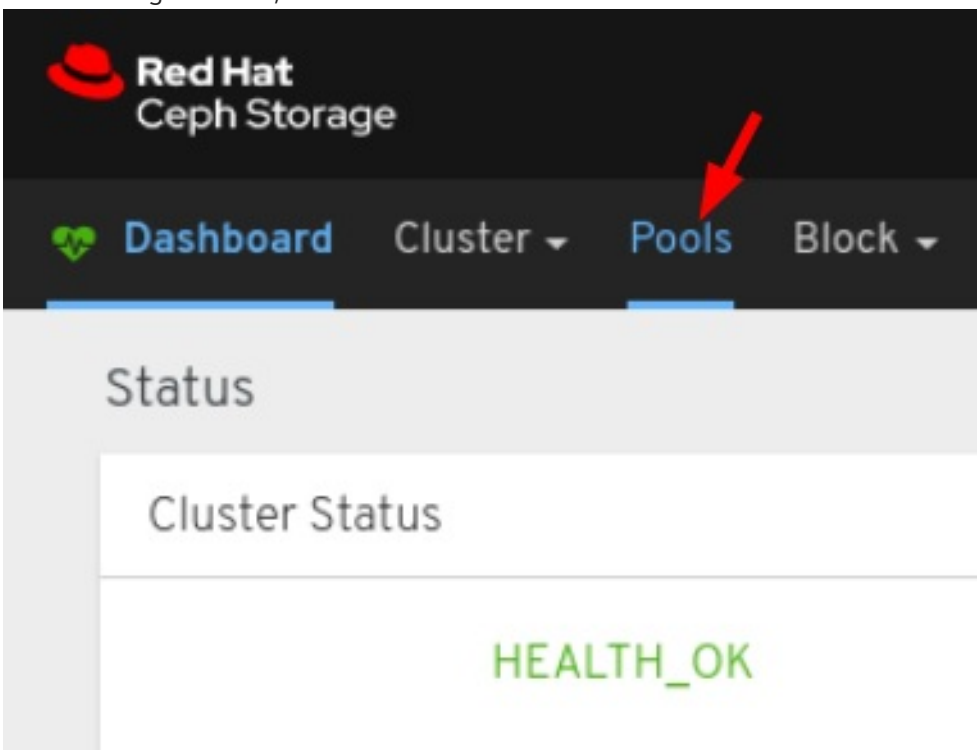

3. Click the row of a replicated pool with the RBD application to select it for editing:

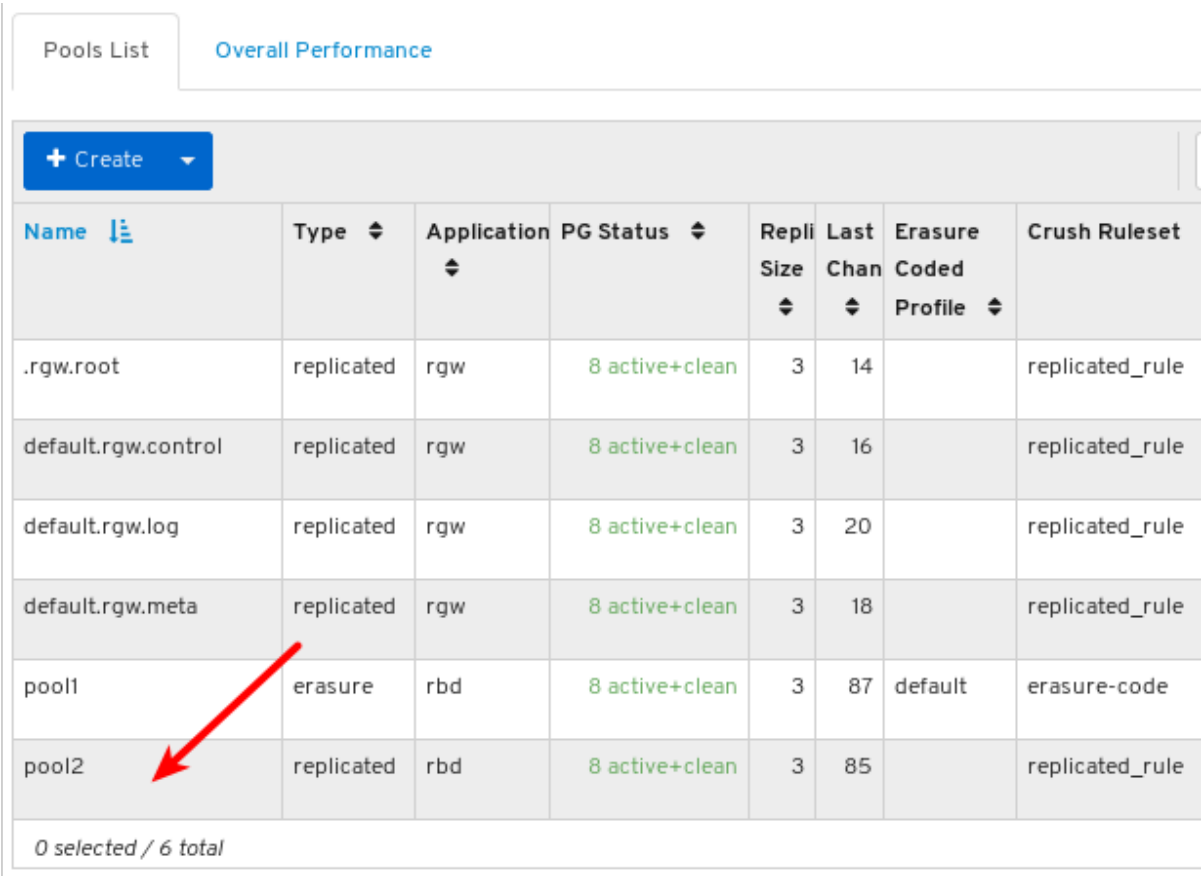

4. Click the *Edit* button:

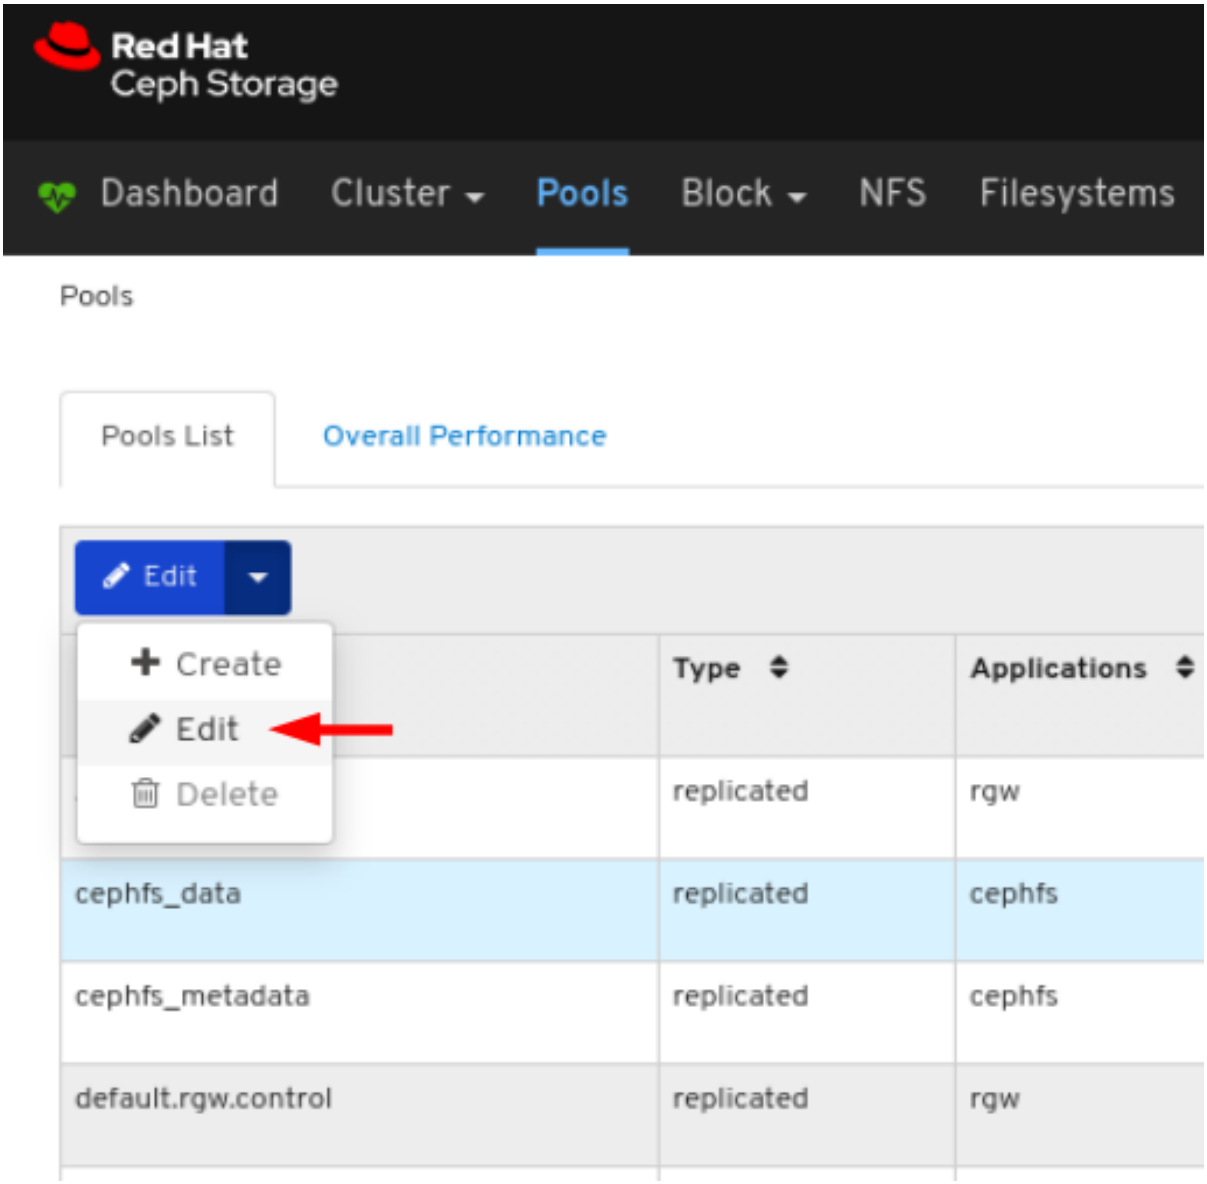

5. Click the plus symbol next to *Quality of Service* to open the QoS settings:

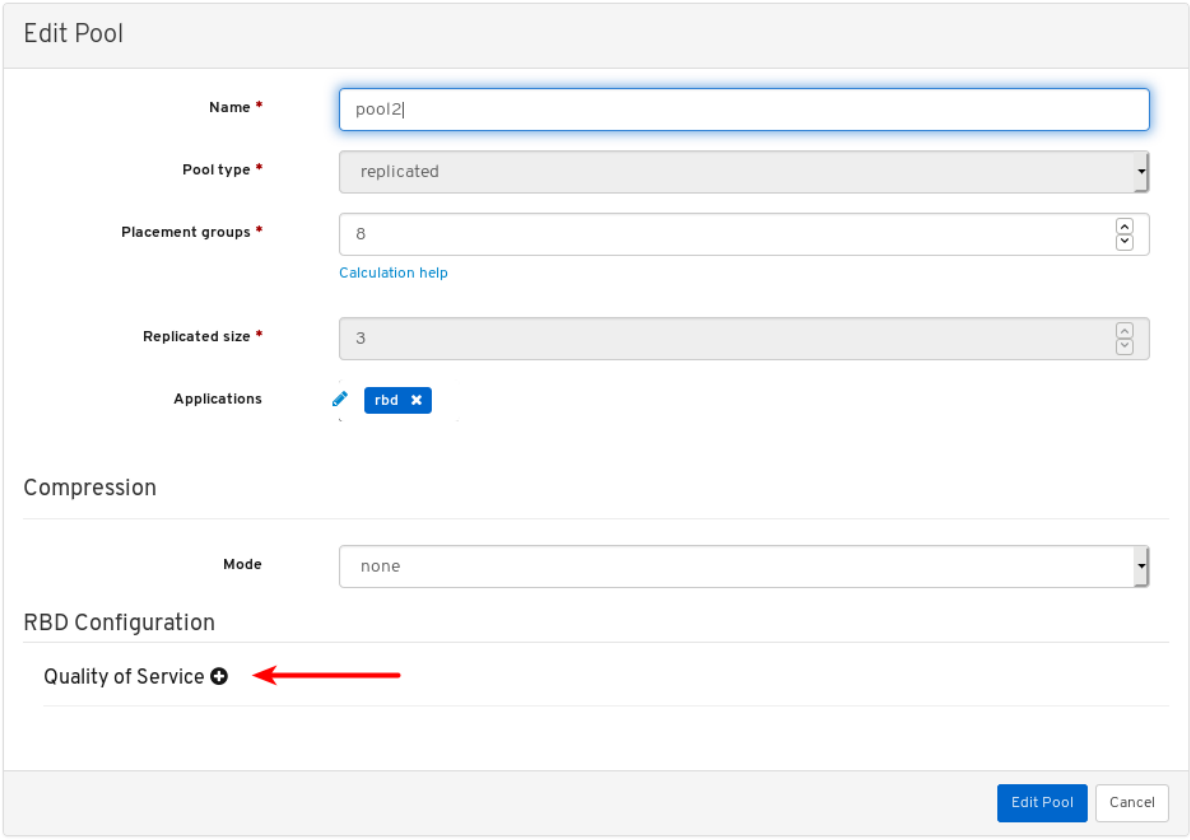

6. Optional: Click the question mark symbol next to an individual setting to find out more about it.

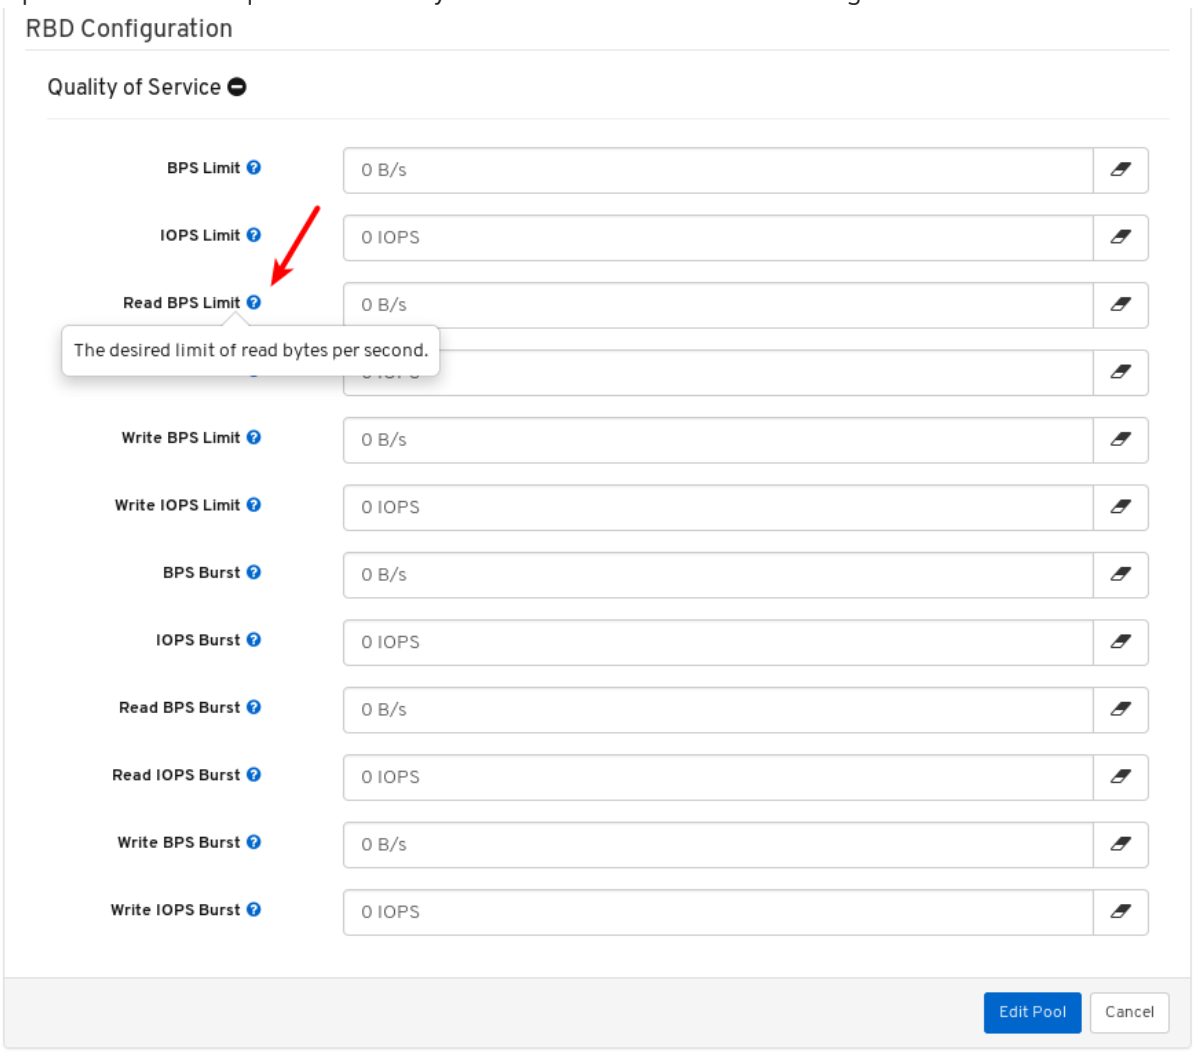

7. Enter or edit values for the QoS settings you want to change:

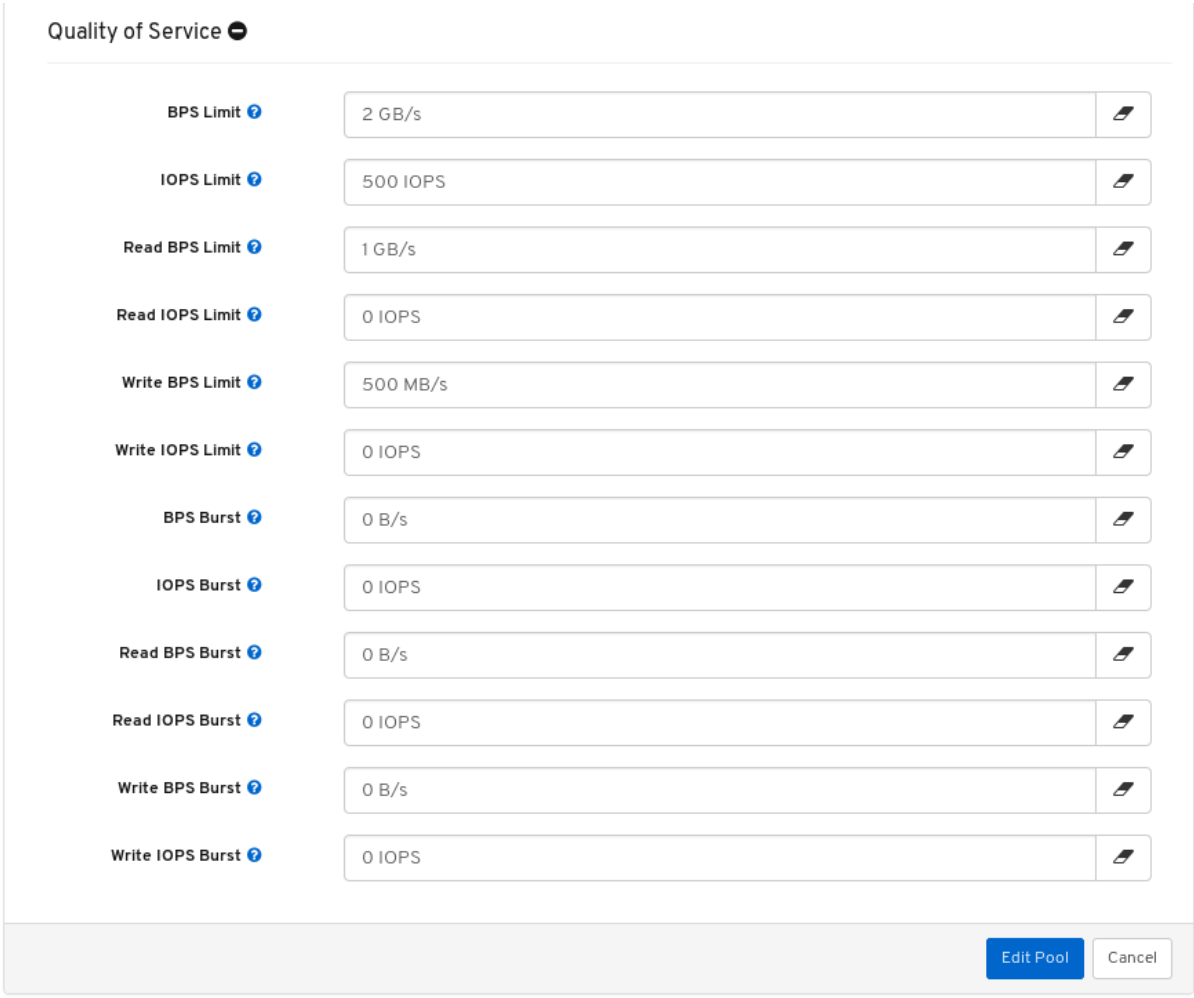

8. Optional: Click the eraser symbol for any setting to remove the local value and inherit the value from the parent global configuration.

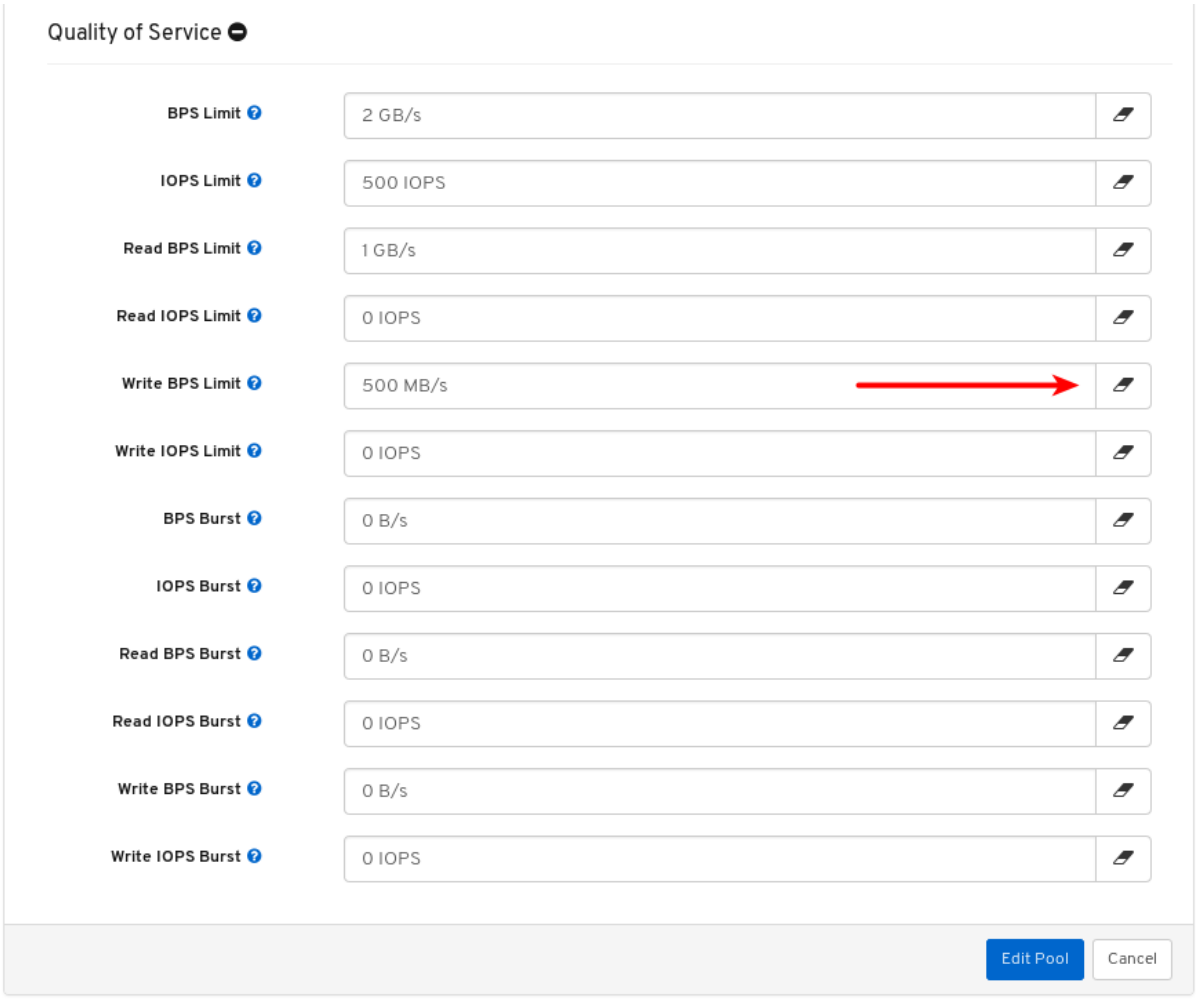

The field for the setting is disabled to indicate it is inheriting the value from the parent.

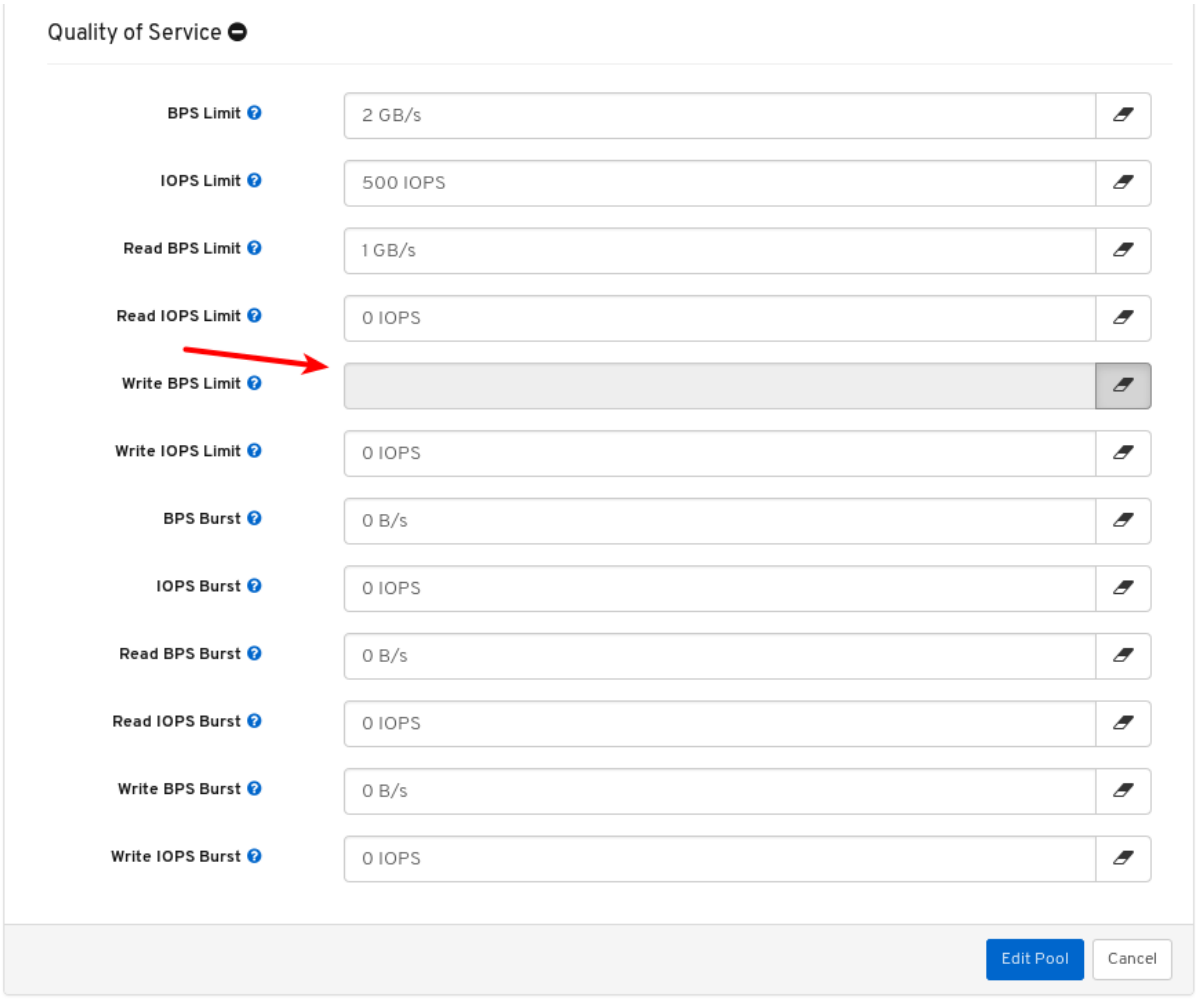

9. Click the *Edit* button to save the changes.

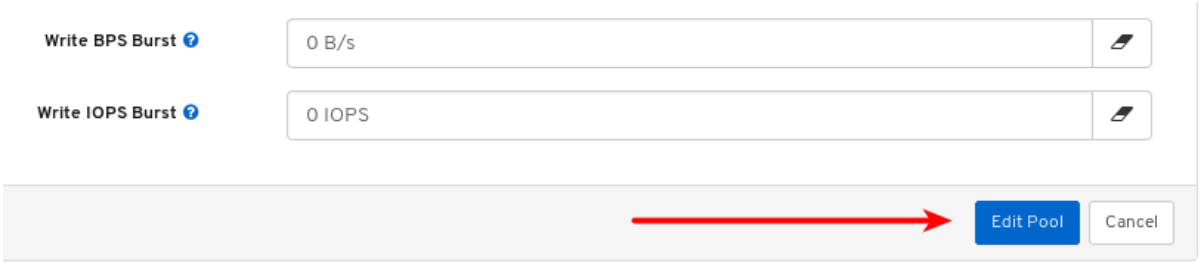

9.5.4. Configuring Quality of Service when creating an image

As a storage administrator, you can configure Quality of Service (QoS) limits when creating an image.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A replicated pool with the RBD application enabled.
- An erasure coded pool with the RBD application enabled.

Procedure

1. Log in to the dashboard.

2. On the navigation bar, click *Block* and then click *Images*:

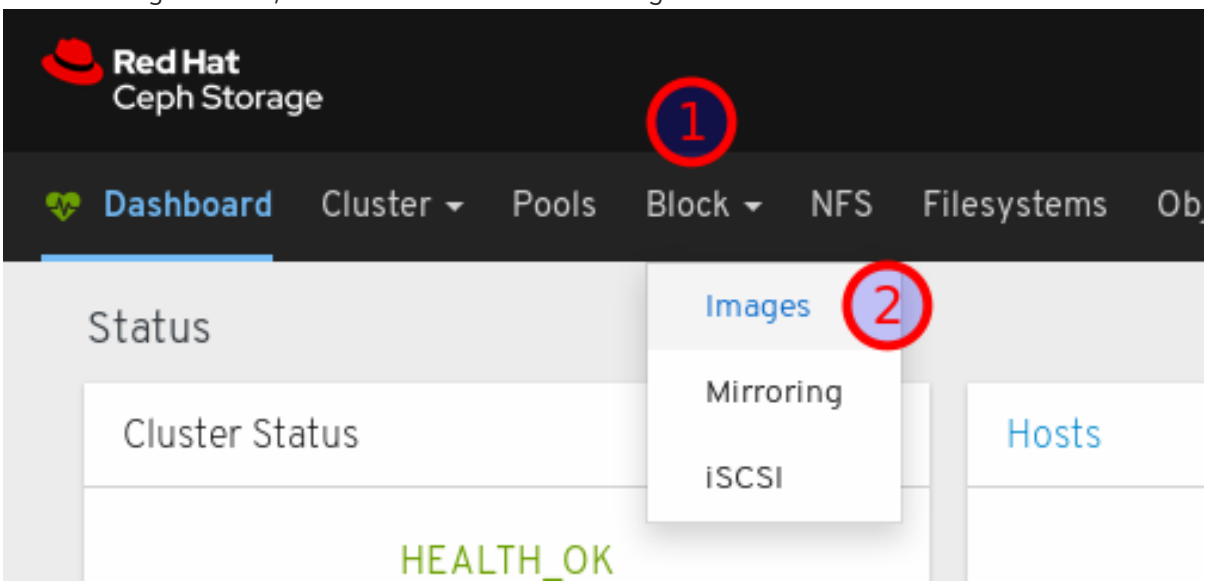

3. Click the *Create* button:

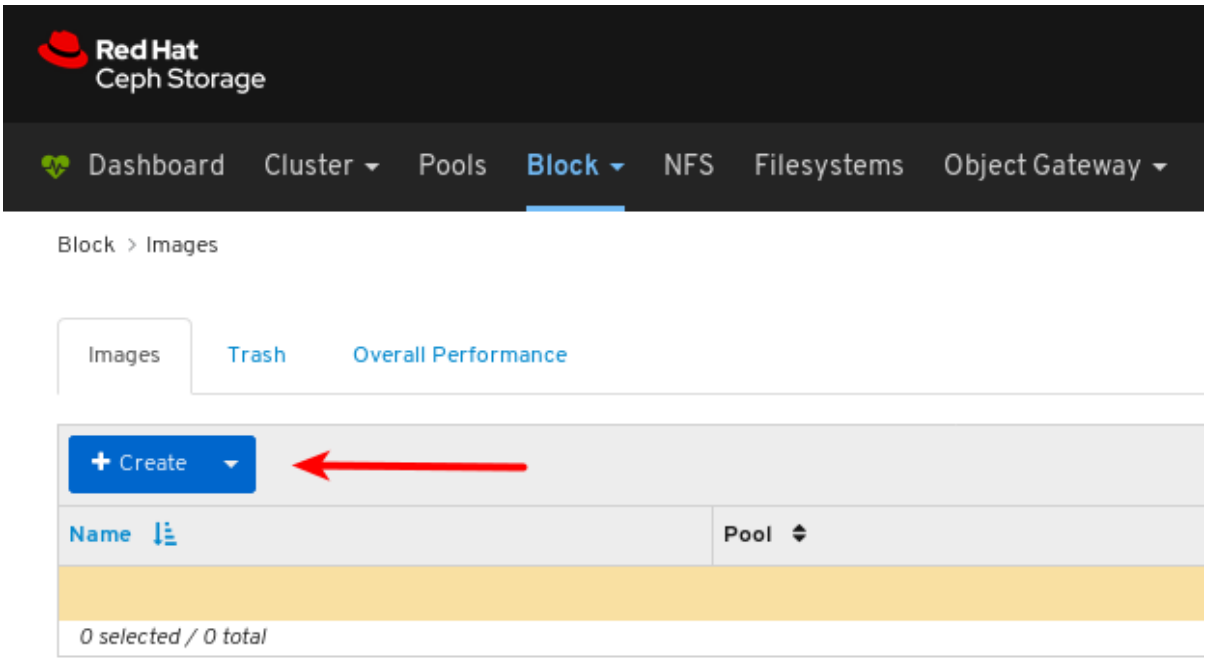

4. At *1*, set the name, at *2*, set the pool, at *3*, set the size, and at *4*, click *Advanced*.

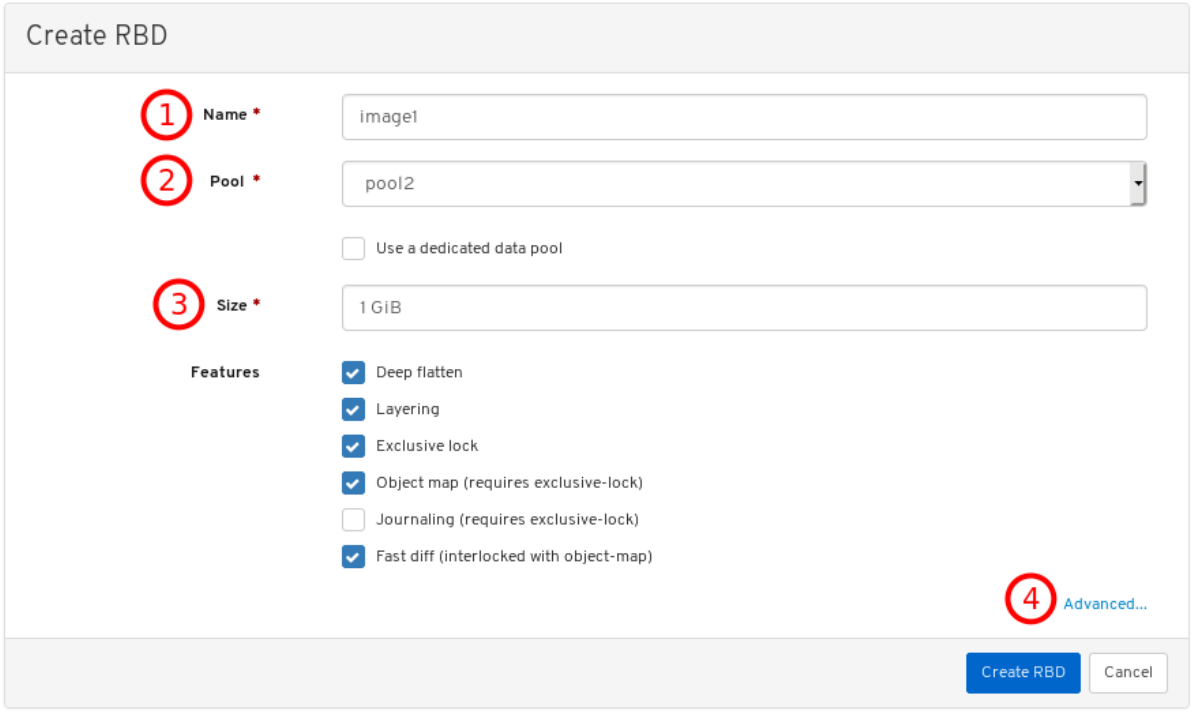

5. Click the plus symbol next to *Quality of Service* to open the QoS settings:

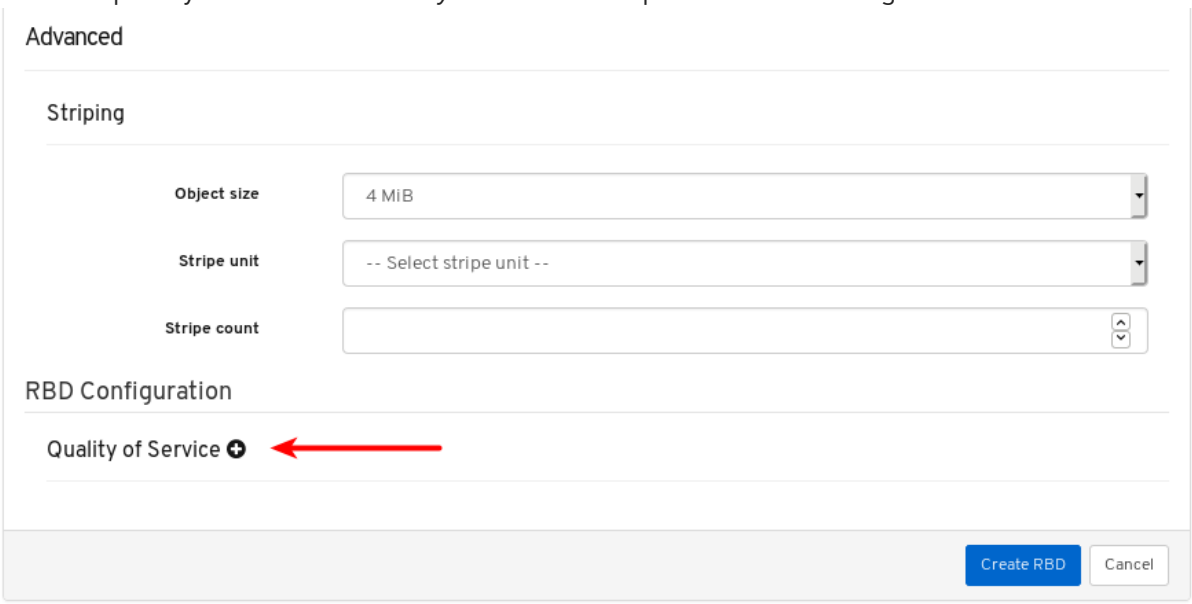

6. Optional: Click the question mark symbol next to an individual setting to find out more about it.

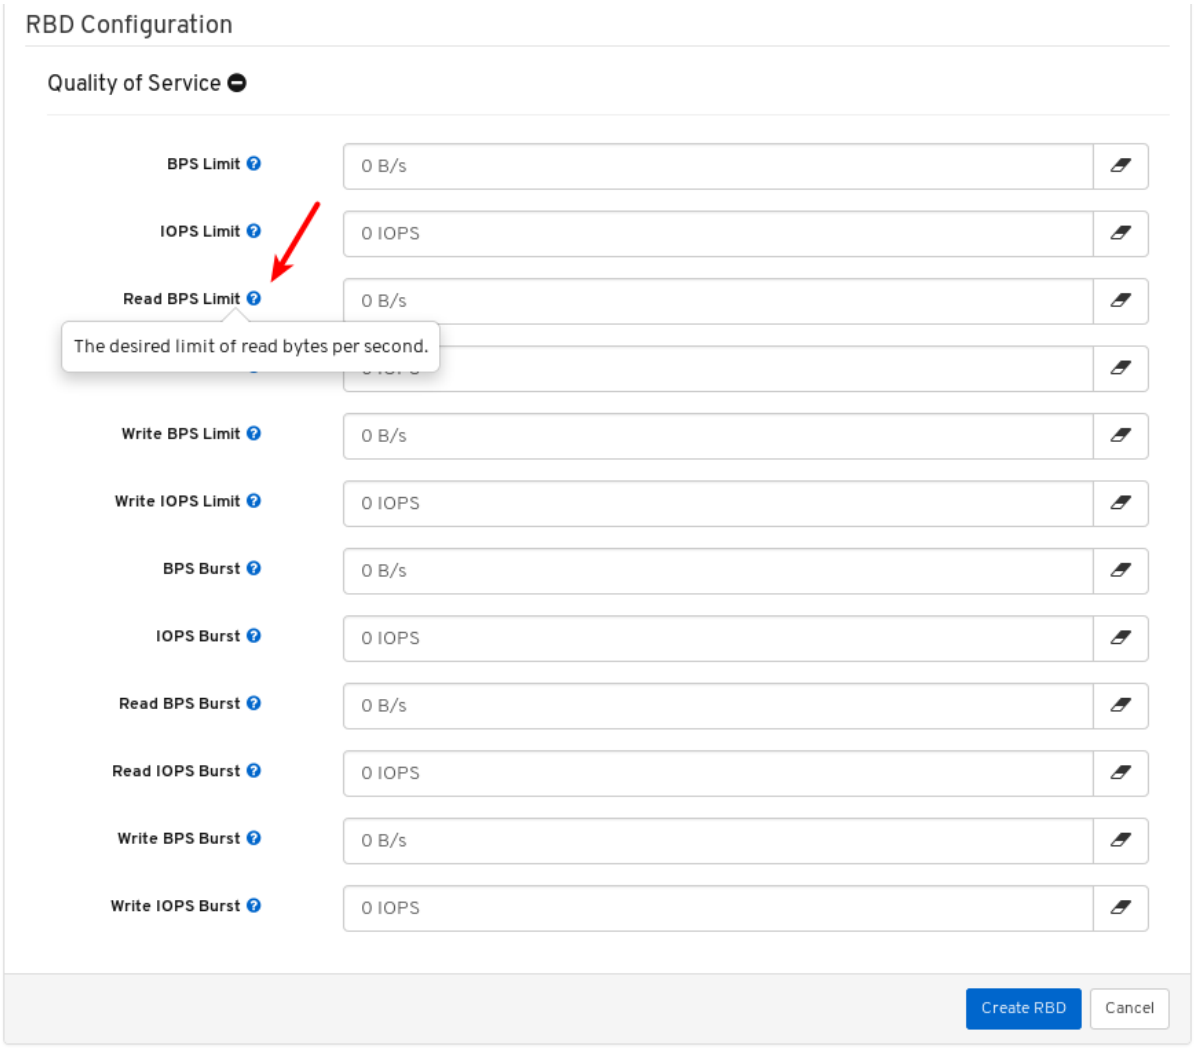

7. Enter or edit values for the QoS settings you want to change:

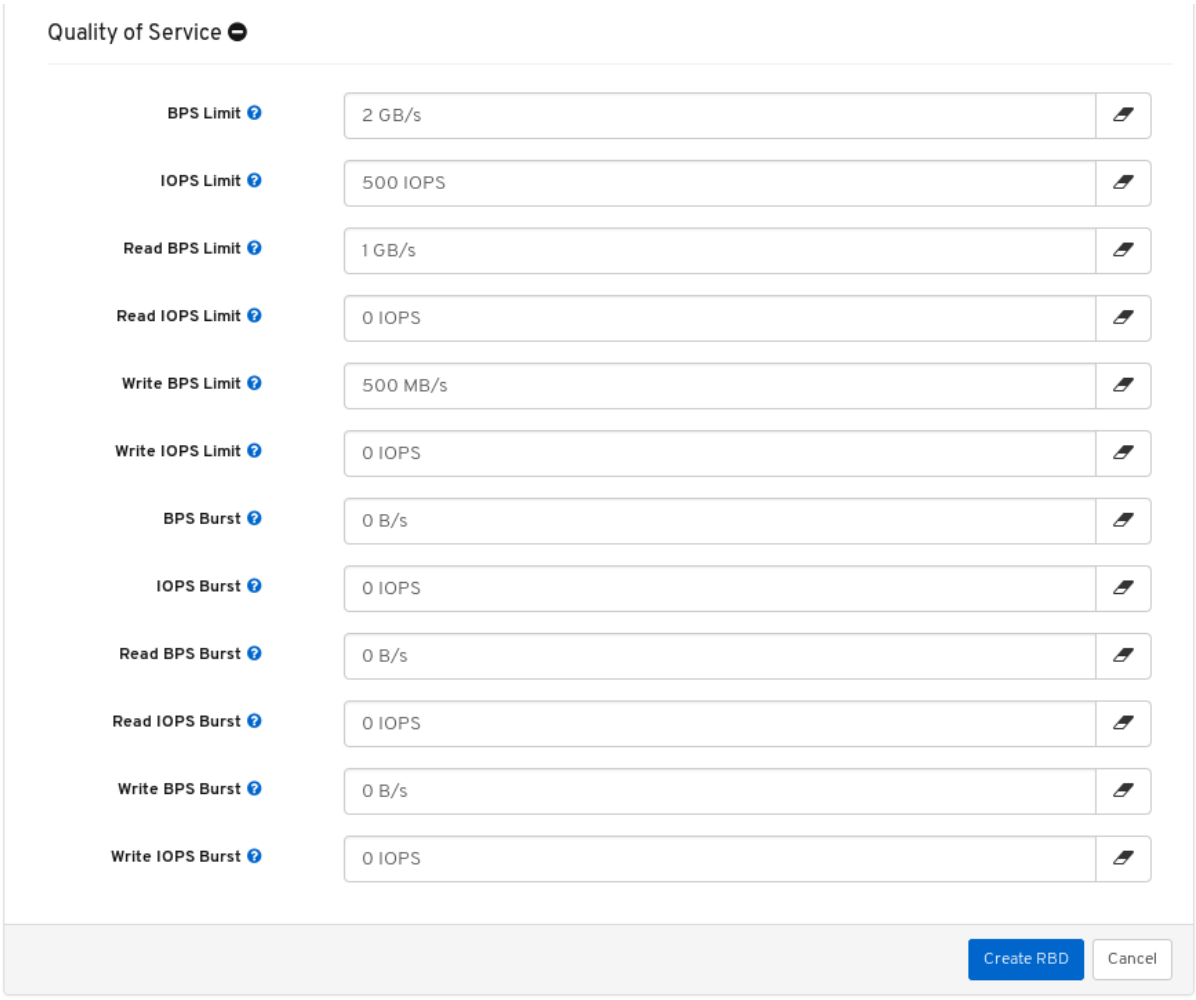

8. Optional: Click the eraser symbol for any setting to remove the local value and inherit the value from the parent pool configuration.

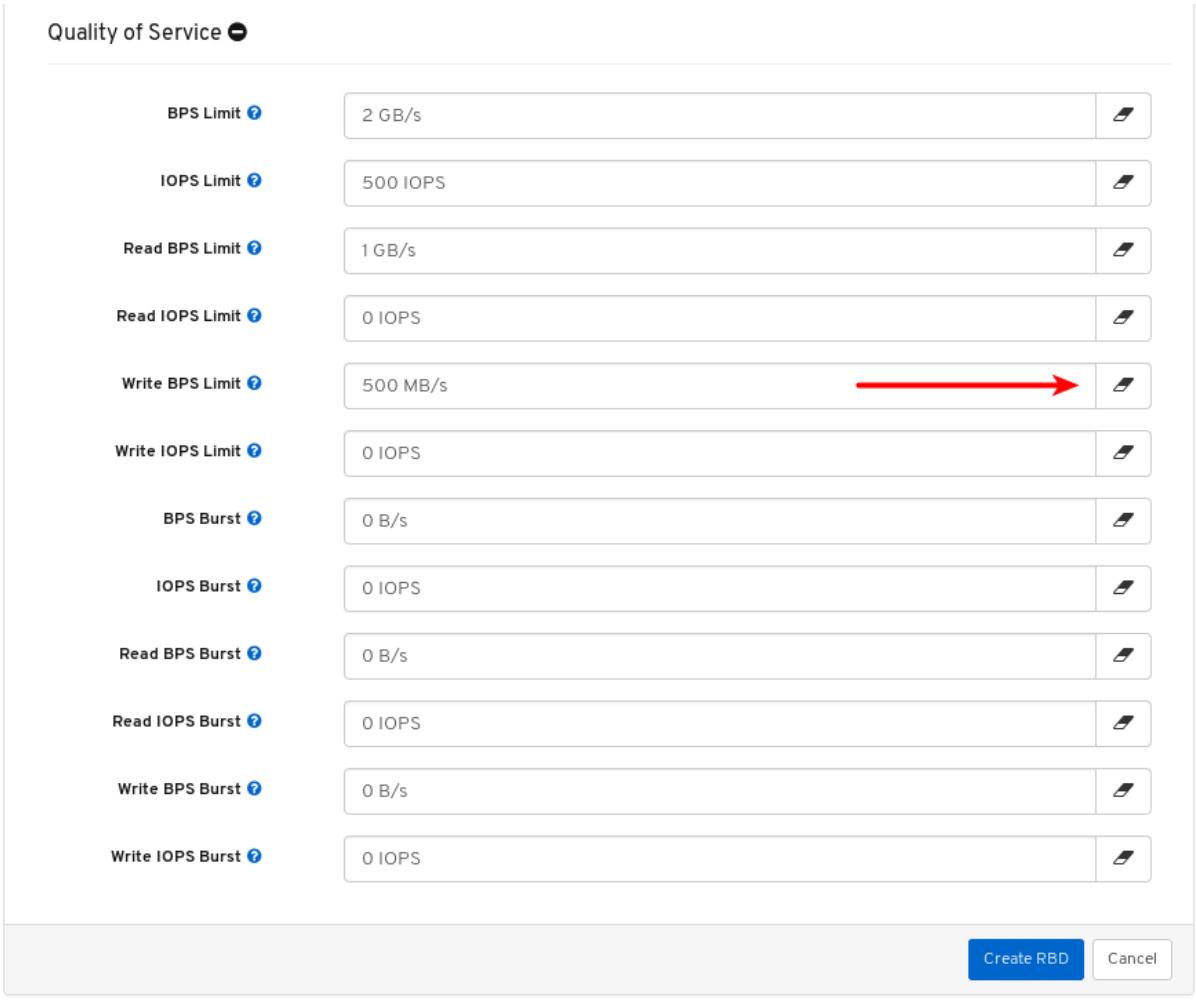

The field for the setting is disabled to indicate it is inheriting the value from the parent.

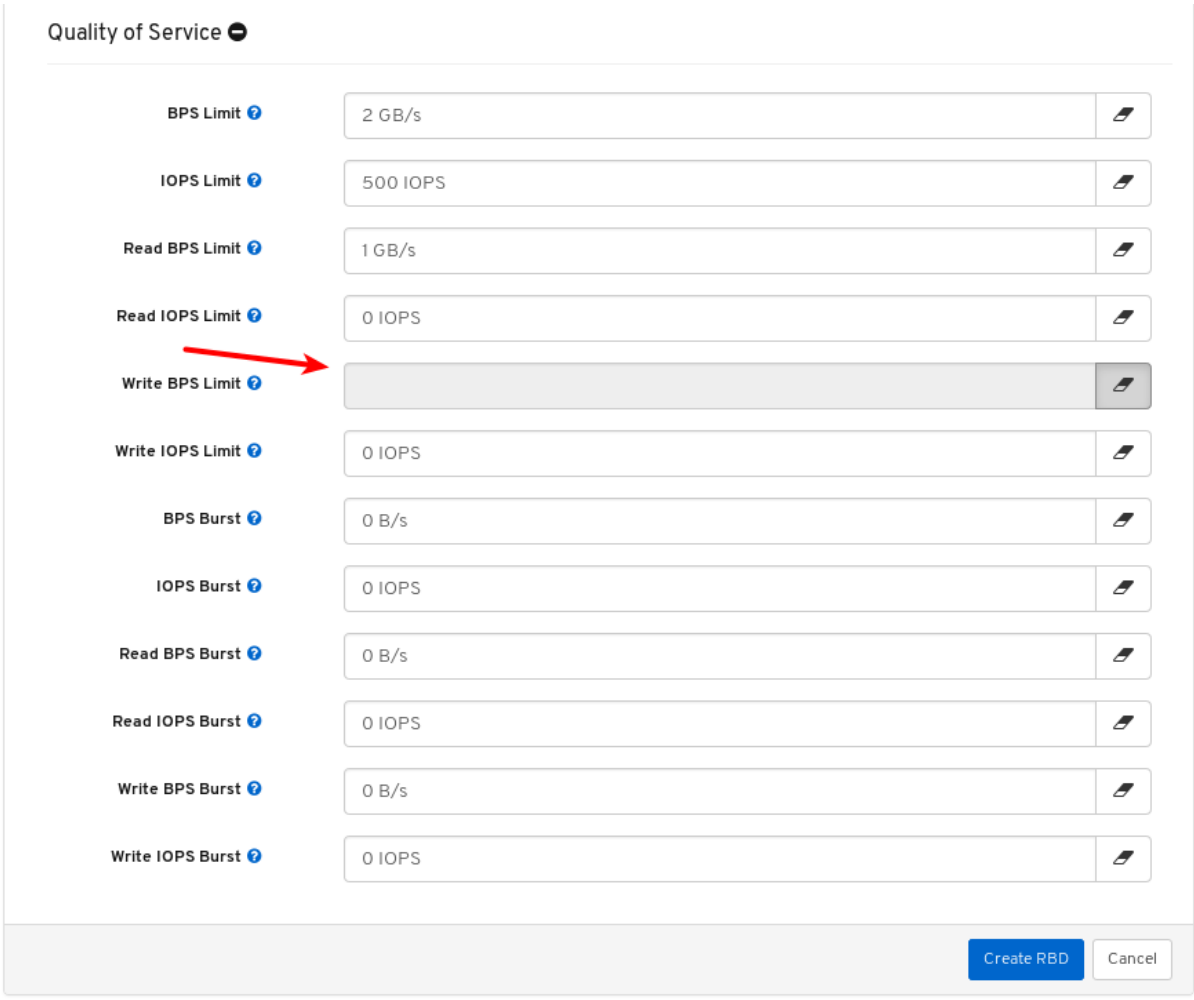

9. Click the *Create RBD* button:

9.5.5. Configuring Quality of Service when creating a pool

As a storage administrator, you can configure Quality of Service (QoS) when creating a pool.

Prerequisites

- A running Red Hat Ceph Storage cluster.
- Dashboard is installed.
- A replicated pool with the RBD application enabled.
- An erasure coded pool with the RBD application enabled.

Procedure

1. Log in to the dashboard.

2. On the navigation bar, click *Pools*.

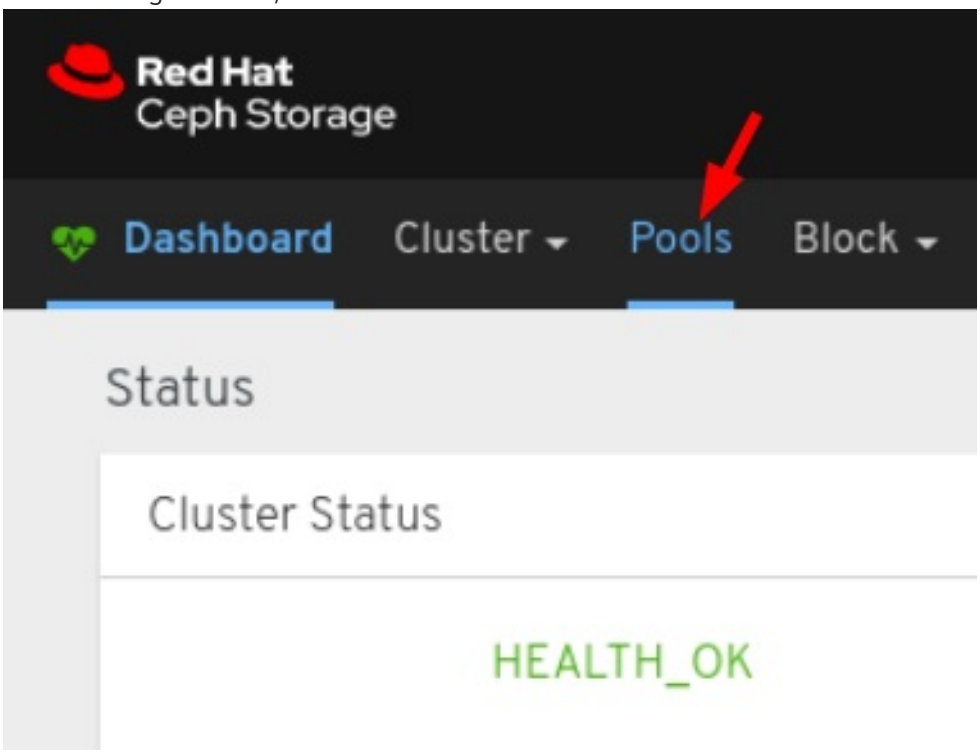

3. Click the *Create* button towards the top left corner of the page:

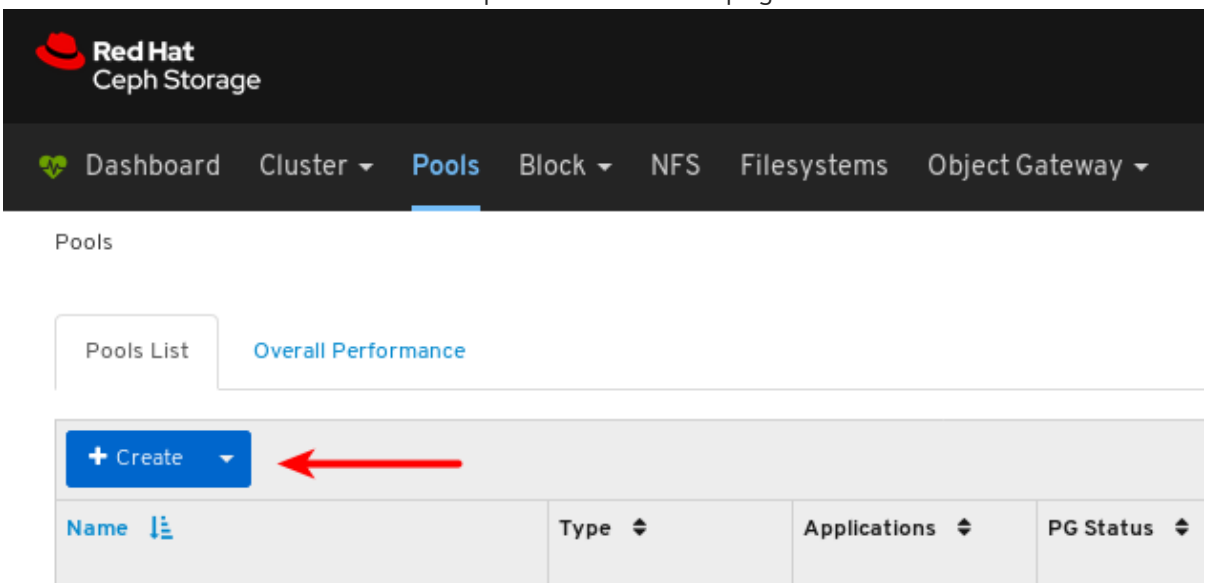

4. In the dialog box, at *1*, set the pool name, at *2*, set the pool type to *erasure*, at *3*, set the number of placement groups, at *4*, enable *EC Overwrites*, at *5*, set the *rbd* application, finally, at *6*, click *Create Pool*.

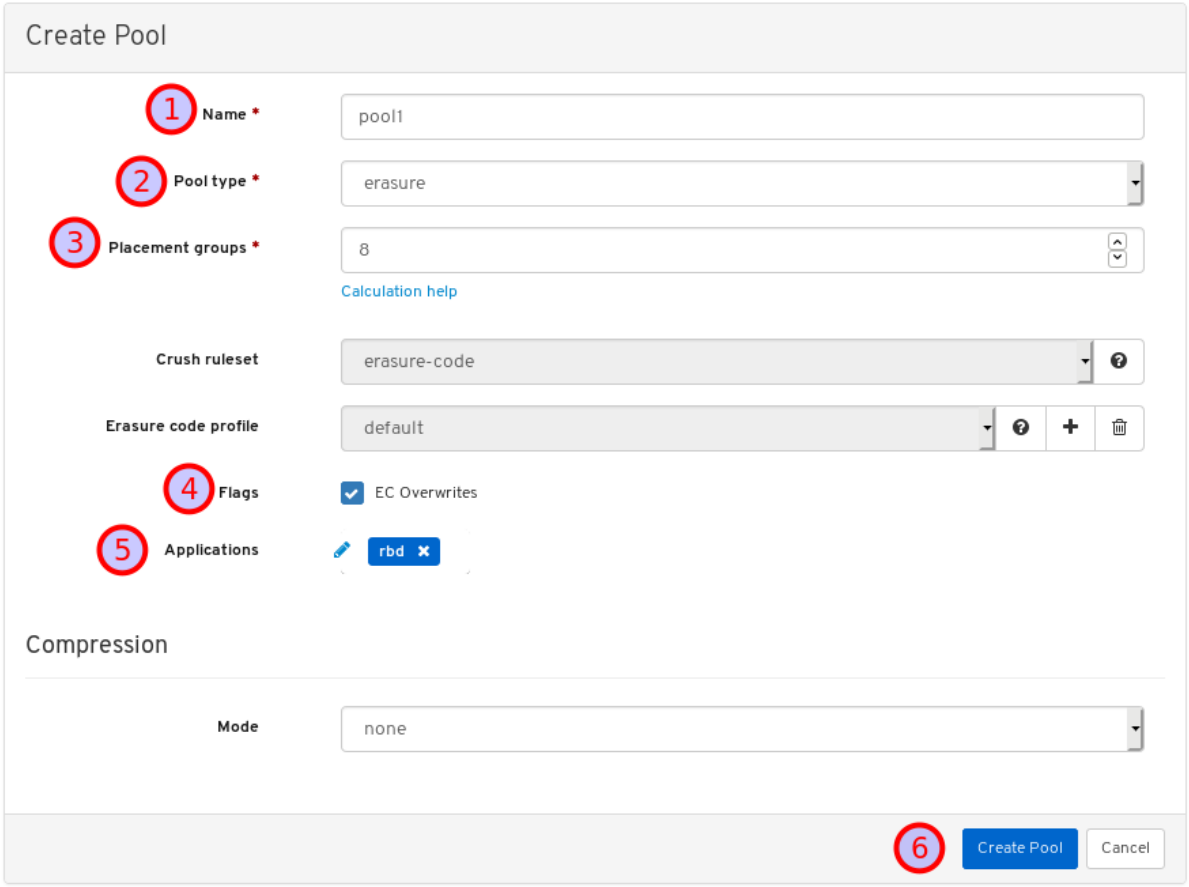

5. Create another pool but this time set its type to *replicated*:

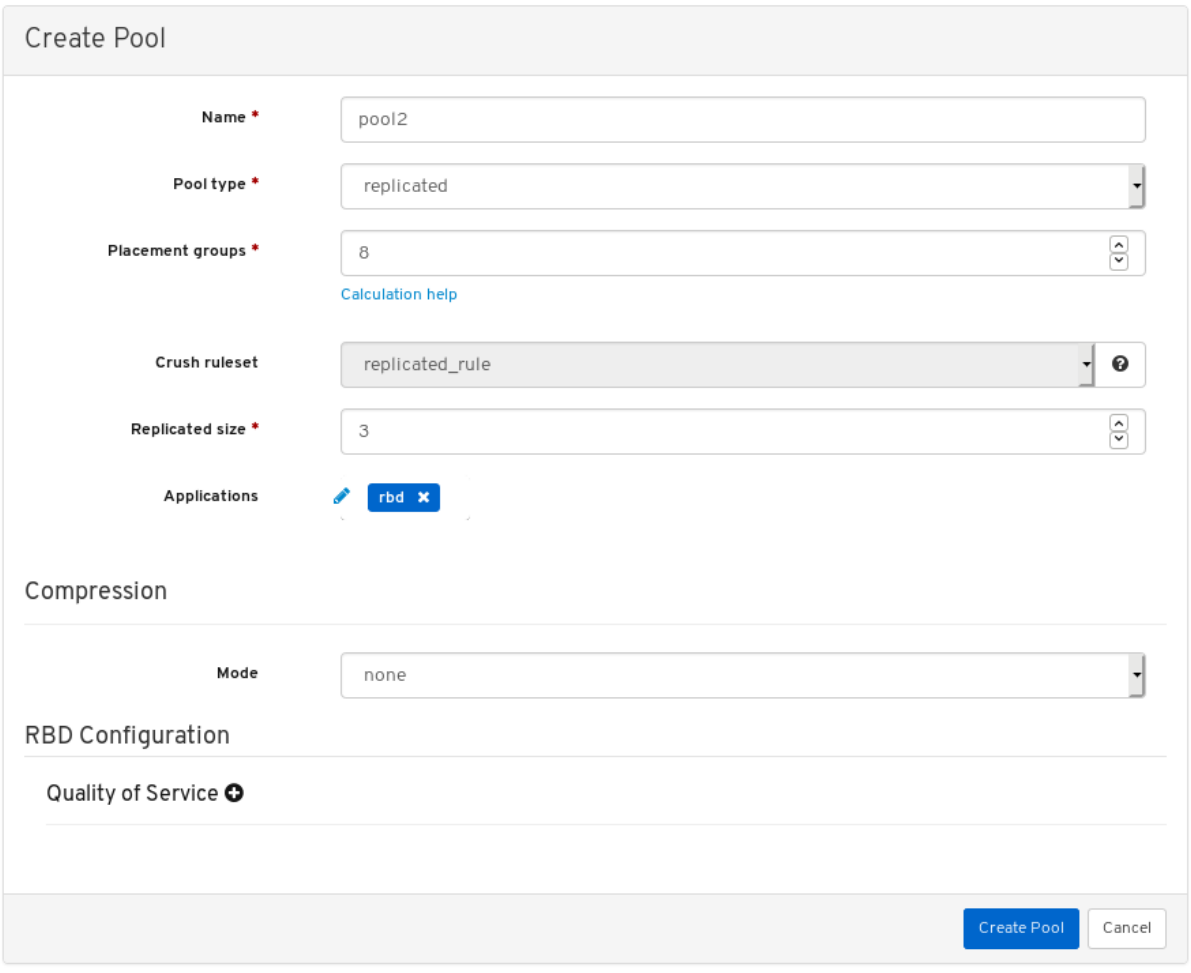

6. Click the plus symbol next to *Quality of Service* to open the QoS settings:

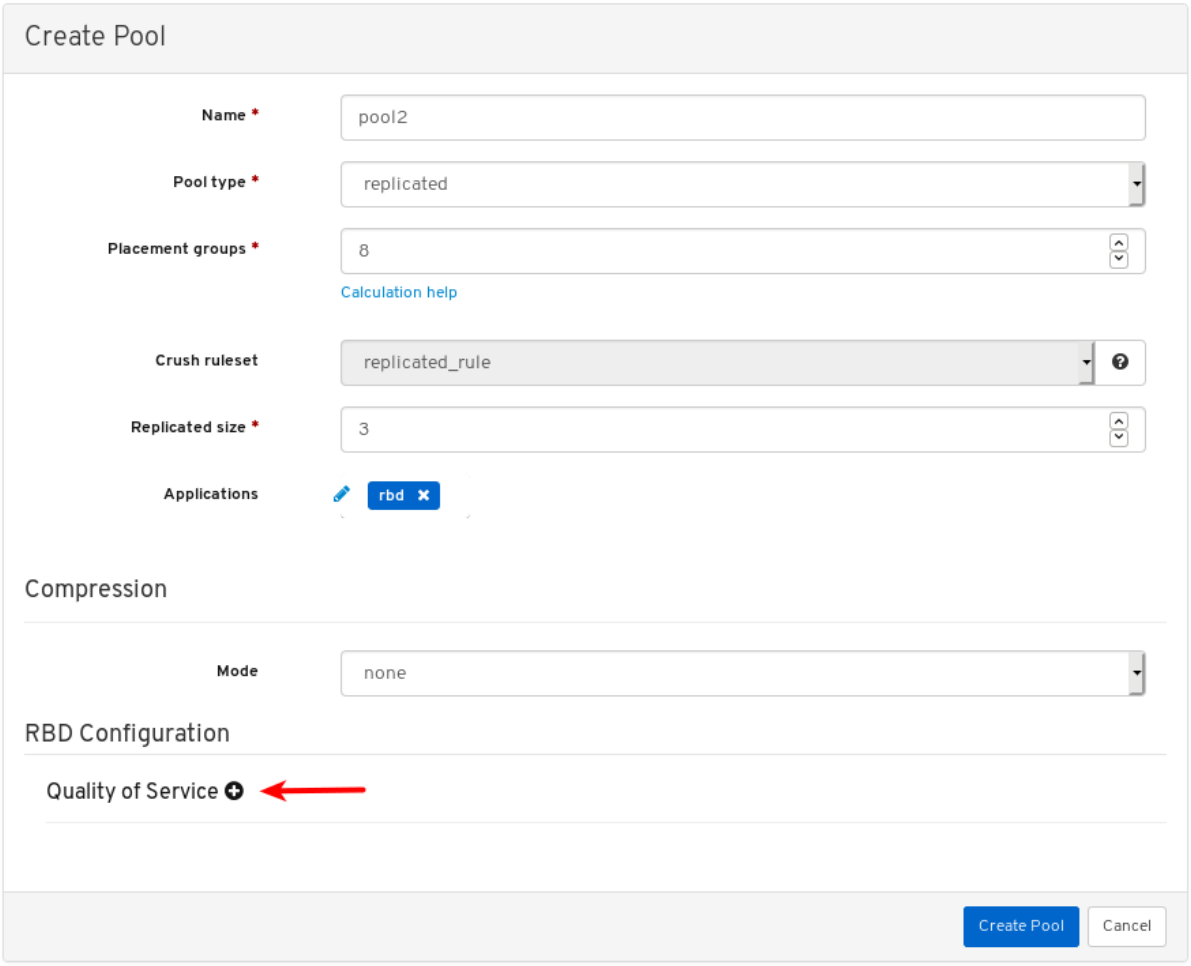

7. Optional: Click the question mark symbol next to an individual setting to find out more about it.

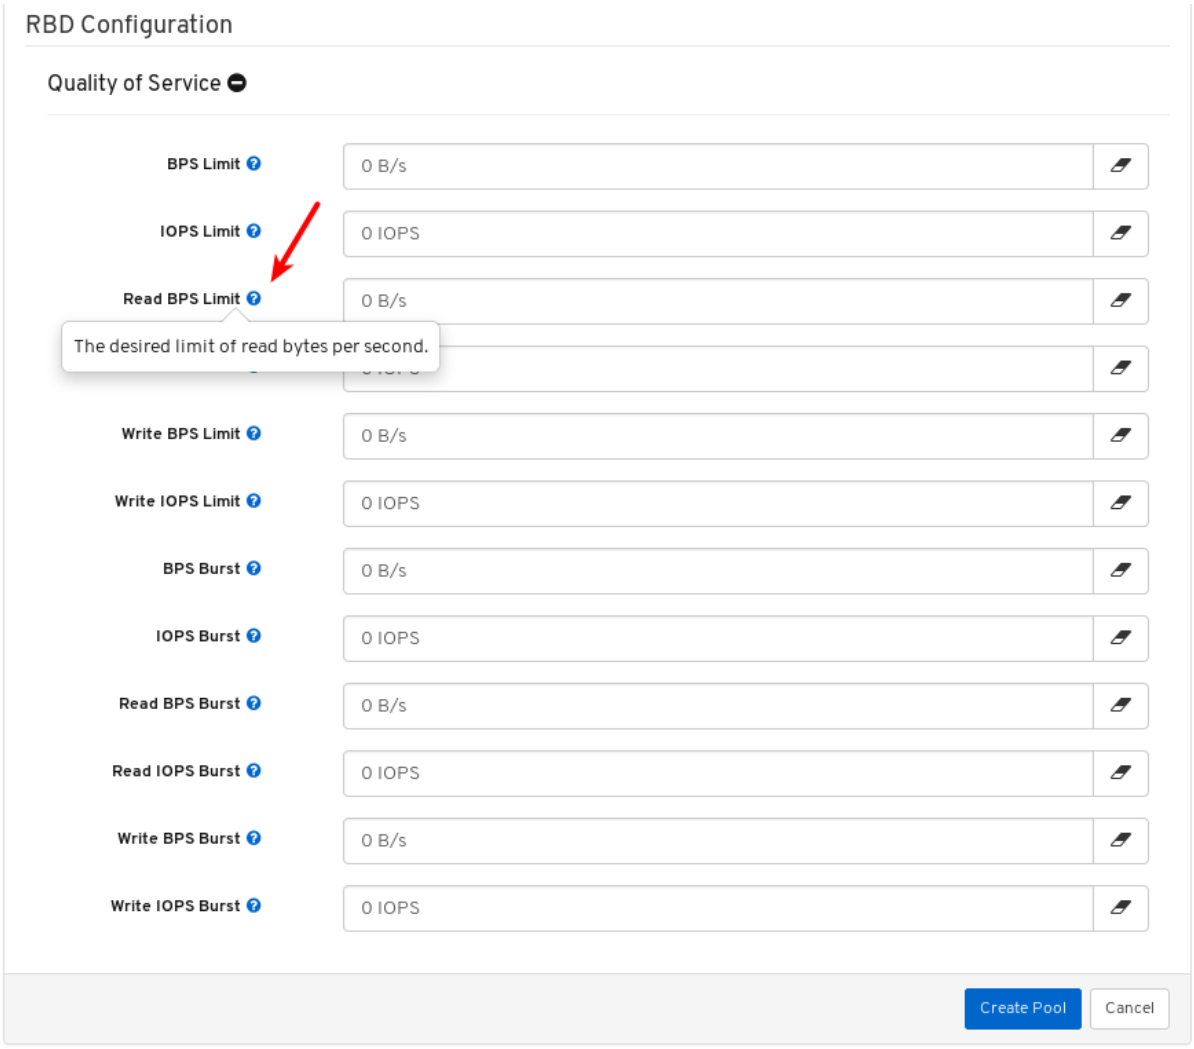

8. Enter or edit values for the QoS settings you want to change:

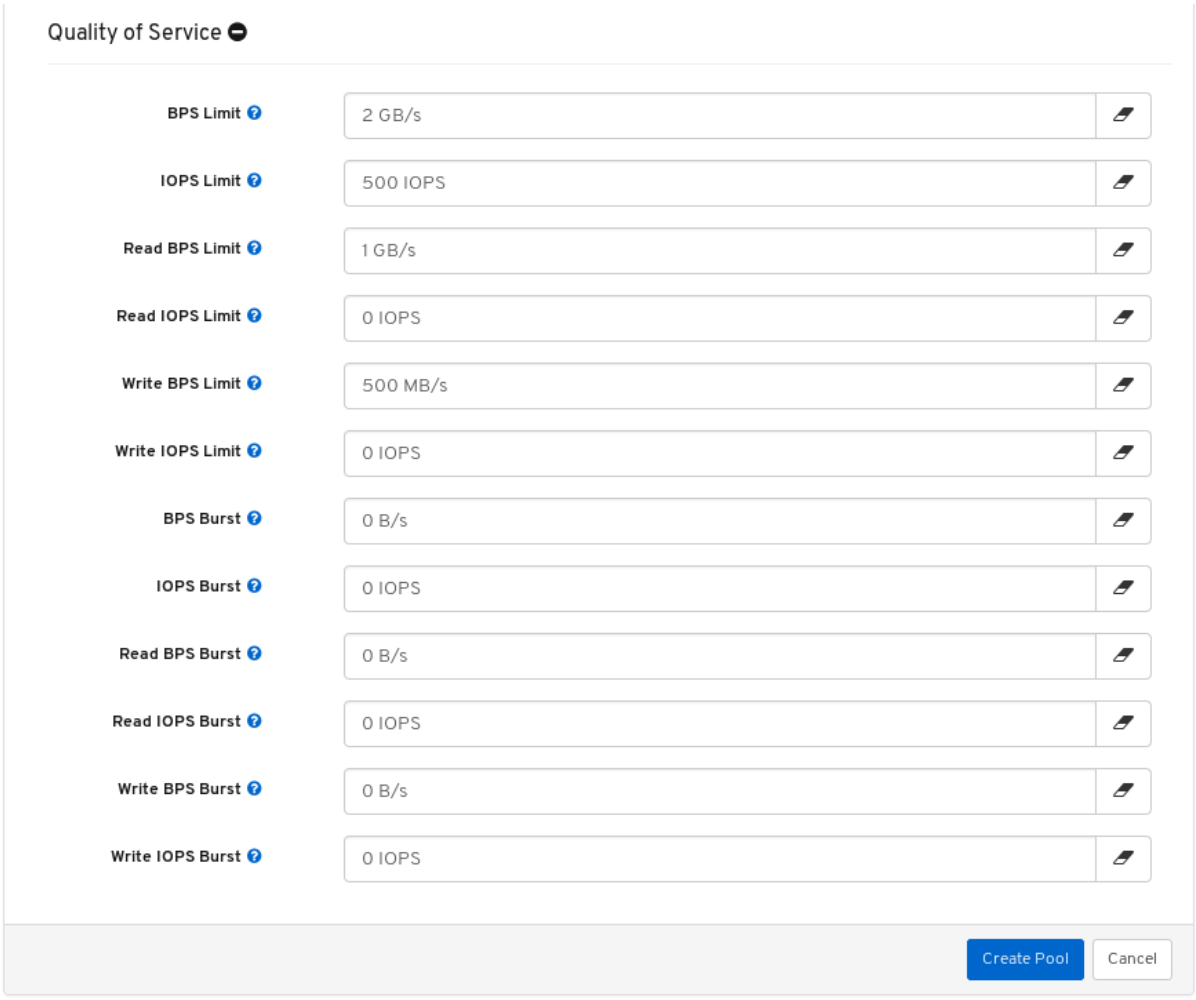

9. Optional: Click the eraser symbol for any setting to remove the local value and inherit the value from the parent global configuration.

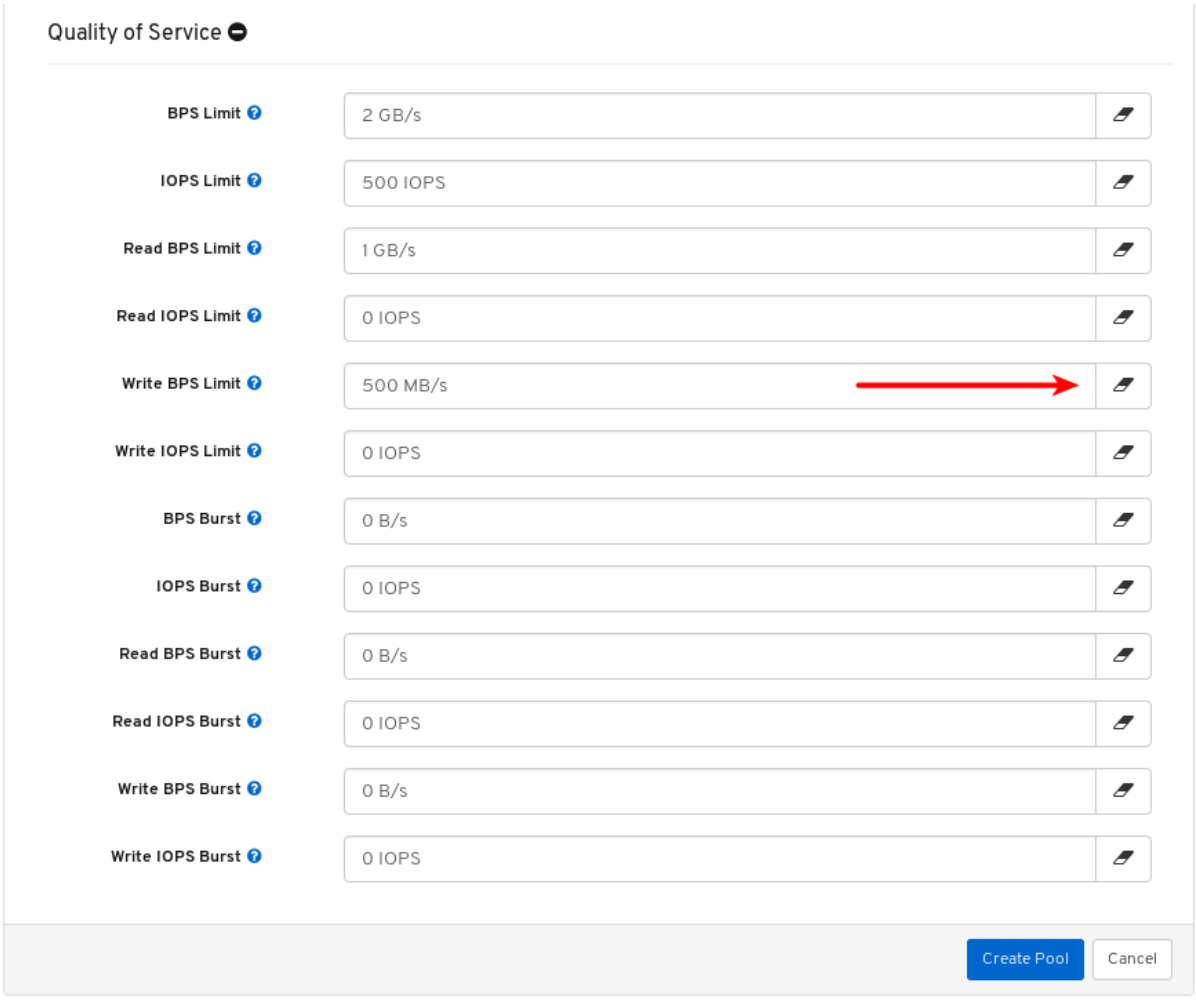

The field for the setting is disabled to indicate it is inheriting the value from the parent.

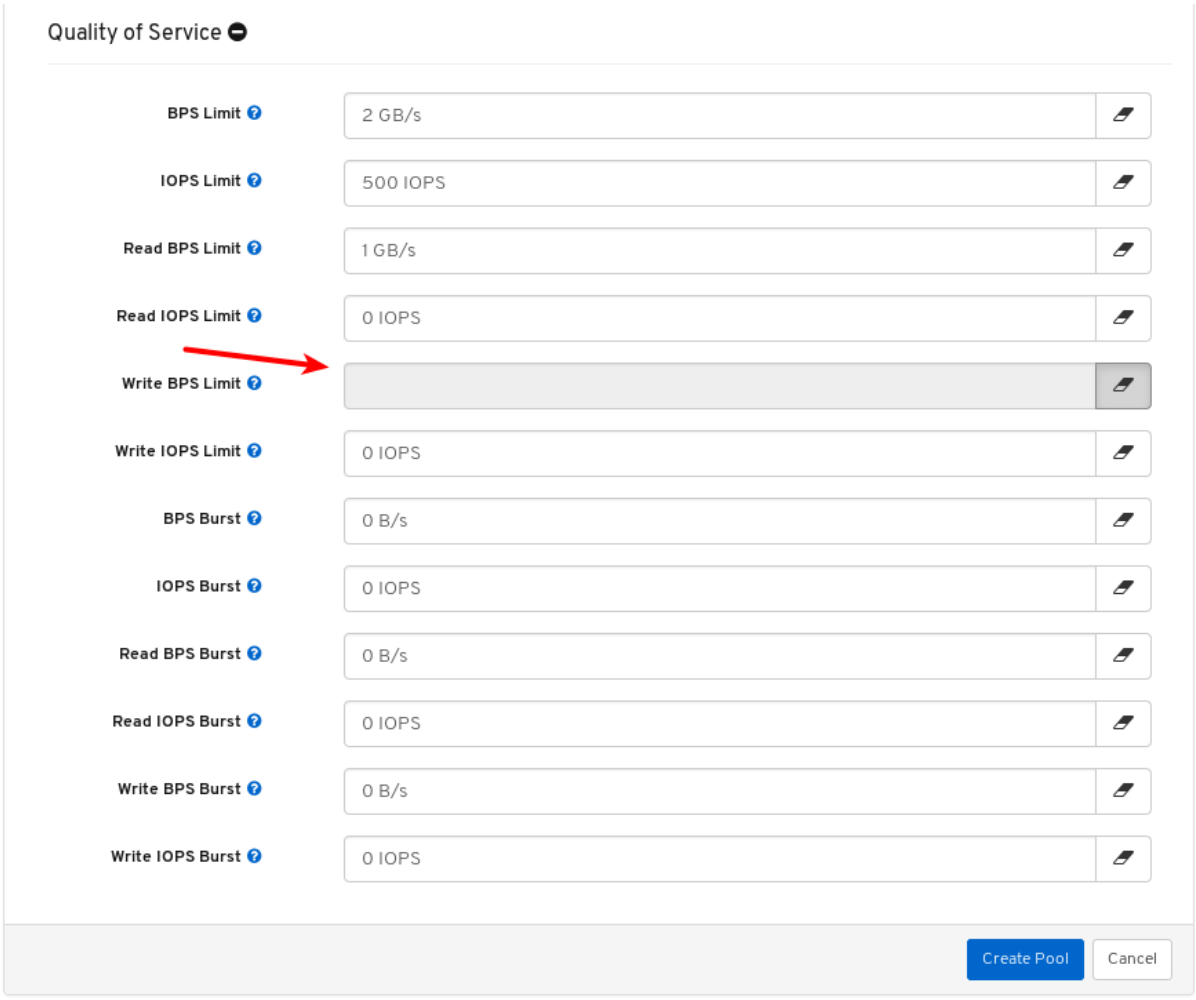

10. Click the *Create Pool* button:

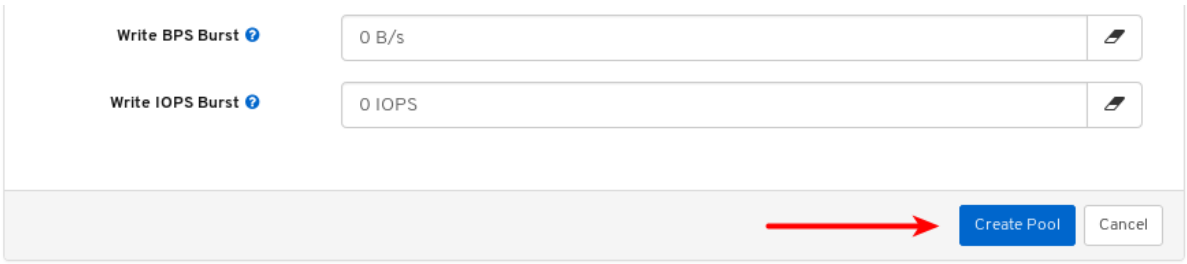

9.5.6. Additional Resources

- For more information on block devices, see the Block Device guide.

APPENDIX A. TROUBLESHOOTING

This section provides the multiple troubleshooting scenarios while using the dashboard.

A.1. DASHBOARD RESPONSE IS SLOW

If the dashboard response is slow, clear the browser cache and reload the dashboard.

A.2. DASHBOARD SHOWS A SERVICE IS DOWN

Dashboard is only a replica of the cluster. If the service is down, check the service status on the node as dashboard displays information collected via node-exporter running on the node. The issue may be in the cluster, configuration, or network.

A.3. TASK FAILURE ON DASHBOARD

While performing any task on the dashboard, if there is any failure, check the respective Ceph daemons. For more information refer to the Troubleshooting Guide

A.4. IMAGES CANNOT BE VIEWED

An image can only be viewed under *Block* > *Images* if the pool it is in has the RBD application enabled on it.

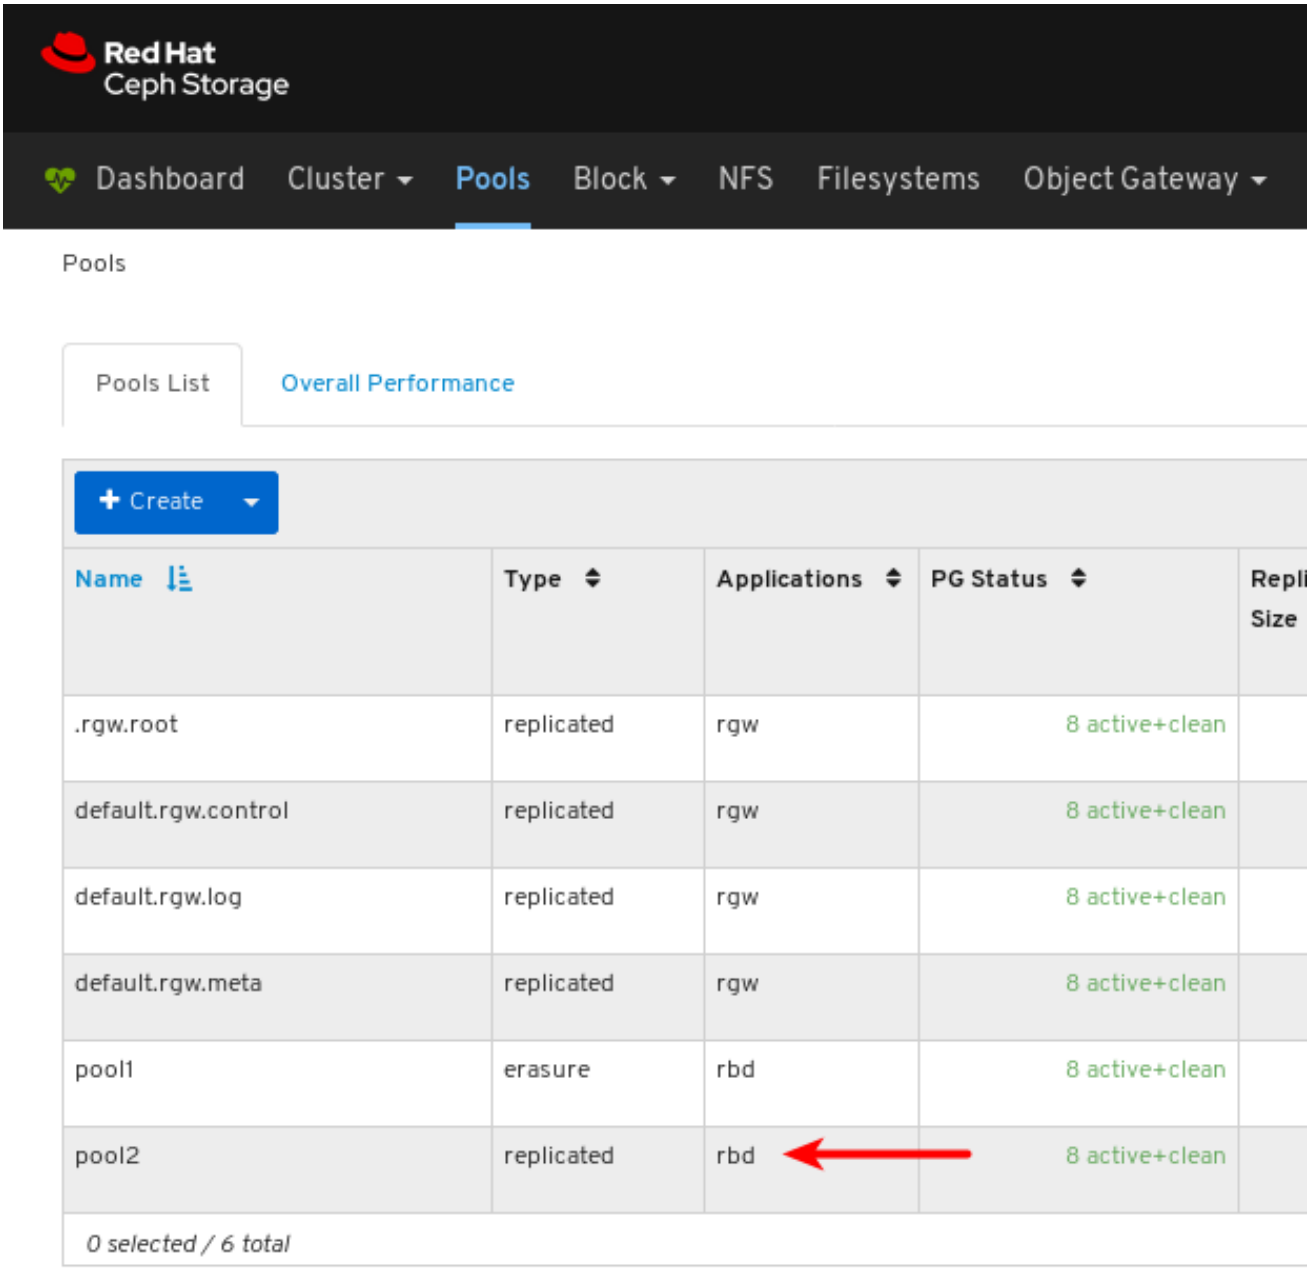
Red Hat
Ceph Storage
Dashboard
Cluster
Pools
Block
NFS
Filesystems
Object Gateway
Pools
Pools List
Overall Performance
+ Create
Name
Type
Applications
PG Status
Repli
Size
.rgw.root
replicated
rgw
8 active+clean
default.rgw.control
replicated
rgw
8 active+clean
default.rgw.log
replicated
rgw
8 active+clean
default.rgw.meta
replicated
rgw
8 active+clean
pool1
erasure
rbd
8 active+clean
pool2
replicated
rbd
8 active+clean
0 selected / 6 total

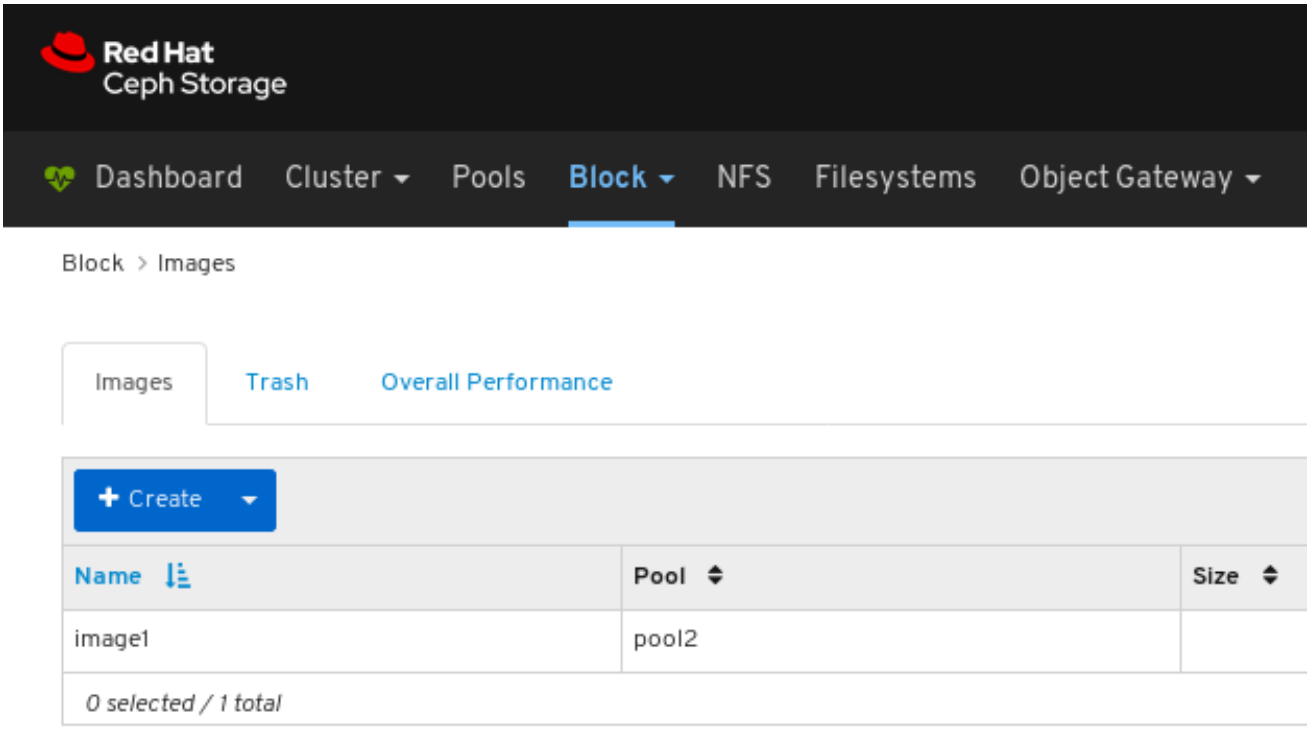

Additional resources

For more information, refer to the Troubleshooting Guide