



Red Hat Ceph Storage 4.3

Release Notes

Release notes for Red Hat Ceph Storage 4.3

Red Hat Ceph Storage 4.3 Release Notes

Release notes for Red Hat Ceph Storage 4.3

Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

The release notes describes the major features, enhancements, known issues, and bug fixes implemented for the Red Hat Ceph Storage 4.3 product release. These release notes include the previous release notes of the previous Red Hat Ceph Storage 4.2 releases up to the current release. 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. INTRODUCTION	3
CHAPTER 2. ACKNOWLEDGMENTS	4
CHAPTER 3. NEW FEATURES	5
3.1. THE CEPH ANSIBLE UTILITY	5
3.2. CEPH MANAGEMENT DASHBOARD	6
3.3. CEPH FILE SYSTEM	6
3.4. CEPH OBJECT GATEWAY	6
3.5. MULTI-SITE CEPH OBJECT GATEWAY	6
CHAPTER 4. TECHNOLOGY PREVIEWS	7
4.1. BLOCK DEVICES (RBD)	7
4.2. OBJECT GATEWAY	7
CHAPTER 5. DEPRECATED FUNCTIONALITY	8
CHAPTER 6. BUG FIXES	9
6.1. THE CEPH ANSIBLE UTILITY	9
6.2. CEPH MANAGEMENT DASHBOARD	10
6.3. THE CEPH VOLUME UTILITY	10
6.4. CEPH OBJECT GATEWAY	11
6.5. CEPH MANAGER PLUGINS	11
CHAPTER 7. KNOWN ISSUES	12
7.1. CEPH MANAGEMENT DASHBOARD	12
7.2. CEPH OBJECT GATEWAY	12
7.3. THE CEPH ANSIBLE UTILITY	12
CHAPTER 8. SOURCES	13

CHAPTER 1. INTRODUCTION

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

The Red Hat Ceph Storage documentation is available at <https://access.redhat.com/documentation/en/red-hat-ceph-storage/>.

CHAPTER 2. ACKNOWLEDGMENTS

Red Hat Ceph Storage version 4.3 contains many contributions from the Red Hat Ceph Storage team. In addition, the Ceph project is seeing amazing growth in the quality and quantity of contributions from individuals and organizations in the Ceph community. We would like to thank all members of the Red Hat Ceph Storage team, all of the individual contributors in the Ceph community, and additionally, but not limited to, the contributions from organizations such as:

- Intel[®]
- Fujitsu[®]
- UnitedStack
- Yahoo[™]
- Ubuntu Kylin
- Mellanox[®]
- CERN[™]
- Deutsche Telekom
- Mirantis[®]
- SanDisk[™]
- SUSE

CHAPTER 3. NEW FEATURES

This section lists all major updates, enhancements, and new features introduced in this release of Red Hat Ceph Storage.

3.1. THE CEPH ANSIBLE UTILITY

Users can now purge the dashboard and monitoring stack only

Previously, users could not purge only the Ceph Manager Dashboard and Monitoring stack components such as Alertmanager, Prometheus, Grafana, and node-exporter separately.

With the `purge-dashboard.yml` playbook, users can remove only the dashboard and the monitoring stack components.

Purging the storage cluster with `osd_auto_discovery: true` now purges the cluster and removes the Ceph OSDs

Previously, purging the storage cluster deployed with `osd_auto_discovery: true` would not purge the Ceph OSDs. With this release, the purge playbook works as expected and removes the Ceph OSDs when the storage cluster is deployed with `osd_auto_discovery: true` scenario.

The Alertmanager configuration is customizable

With this release, you can customize the Alertmanager configuration using the `alertmanager_conf_overrides` parameter in the `/group_vars/all.yml` file.

The Red Hat Ceph Storage Dashboard deployment is supported on a dedicated network

Previously, `ceph-ansible` stated the address that should be used for deploying the dashboard was on the same subnet as the `public_network`.

With this release, you can override the default dedicated subnet for the dashboard by setting the `dashboard_network` parameter in the `/group_vars/all.yml` file with the CIDR subnet address.

Setting the global NFS options in the configuration file is supported

Previously, `ceph-ansible` would not allow overriding any parameter in the configuration file.

With this release, you can override any parameter in the `NFS_CORE_PARAM` block section in the `ganesh.conf` file by setting the variable `ganesh_core_param_overrides` in `group_vars/all.yml` and update client-related configuration.

`ceph-ansible` checks for the Ceph Monitor quorum before starting the upgrade

Previously, when the storage cluster was in a `HEALTH_ERR` or `HEALTH_WARN` state due to one of the Ceph monitors being down, the `rolling_upgrade.yml` playbook would run. However, the upgrade would fail and the quorum was lost resulting in I/O down or a cluster failure.

With this release, an additional condition occurs where `ceph-ansible` checks the Ceph Monitor quorum before starting the upgrade.

The systemd target units for containerized deployments are now supported

Previously, there was no way to stop all Ceph daemons on a node in a containerized deployment.

With this release, systemd target units for containerized deployments are supported and you can stop all the Ceph daemons on a host or specific Ceph daemons similar to bare-metal deployments.

ceph-ansible now checks the relevant release version during an upgrade before executing the playbook

With this release, during a storage cluster upgrade, **ceph-ansible** first checks for the relevant release version and the playbook fails with an error message if a wrong Ceph version is provided.

3.2. CEPH MANAGEMENT DASHBOARD

A new Grafana Dashboard to display graphs for Ceph Object Gateway multi-site setup

With this release, a new Grafana dashboard is now available and displays graphs for Ceph Object Gateway multisite sync performance including two-way replication throughput, polling latency, and unsuccessful replications.

See the [Monitoring Ceph object gateway daemons on the dashboard](#) section in the *Red Hat Ceph Storage Dashboard Guide* for more information.

3.3. CEPH FILE SYSTEM

Use `max_concurrent_clones` option to configure the number of clone threads

Previously, the number of concurrent clones was not configurable and the default was 4.

With this release, the maximum number of concurrent clones is configurable using the manager configuration option:

Syntax

```
ceph config set mgr mgr/volumes/max_concurrent_clones VALUE
```

Increasing the maximum number of concurrent clones could improve the performance of the storage cluster.

3.4. CEPH OBJECT GATEWAY

The role name and the role session information is displayed in the ops log for S3 operations

With this release, you get information such as the role name and the role session in the ops log for all the S3 operations that use temporary credentials returned by AssumeRole* operations for debugging and auditing purposes.

3.5. MULTI-SITE CEPH OBJECT GATEWAY

Data sync logging experienced delays in processing

Previously, data sync logging could be subject to delays in processing large backlogs of log entries.

With this release, data sync includes caching for bucket sync status. The addition of the cache speeds the processing of duplicate datalog entries when a backlog exists.

CHAPTER 4. TECHNOLOGY PREVIEWS

This section provides an overview of Technology Preview features introduced or updated in this release of Red Hat Ceph Storage.



IMPORTANT

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

For more information on Red Hat Technology Preview features support scope, see <https://>

4.1. BLOCK DEVICES (RBD)

Mapping RBD images to NBD images

The **rbd-nbd** utility maps RADOS Block Device (RBD) images to Network Block Devices (NBD) and enables Ceph clients to access volumes and images in Kubernetes environments. To use **rbd-nbd**, install the **rbd-nbd** package. For details, see the **rbd-nbd(7)** manual page.

4.2. OBJECT GATEWAY

Object Gateway archive site

With this release an archive site is supported as a Technology Preview. The archive site allows you to have a history of versions of S3 objects that can only be eliminated through the gateways associated with the archive zone. Including an archive zone in a multizone configuration allows you to have the flexibility of an S3 object history in only one zone while saving the space that the replicas of the versions S3 objects would consume in the rest of the zones.

CHAPTER 5. DEPRECATED FUNCTIONALITY

This section provides an overview of functionality that has been deprecated in all minor releases up to this release of Red Hat Ceph Storage.

Ubuntu is no longer supported

Installing a Red Hat Ceph Storage 4 cluster on Ubuntu is no longer supported. Use Red Hat Enterprise Linux as the underlying operating system.

Configuring iSCSI gateway using `ceph-ansible` is no longer supported

Configuring the Ceph iSCSI gateway by using the `ceph-ansible` utility is no longer supported. Use `ceph-ansible` to install the gateway and then use the `gwcli` utility of the to configure the Ceph iSCSI gateway. For details, see the [The Ceph iSCSI Gateway](#) chapter in the *Red Hat Ceph Storage Block Device Guide*.

`ceph-disk` is deprecated

The `ceph-disk` utility is no longer supported. The `ceph-volume` utility is used instead. For details, see the [Why does `ceph-volume` replace `ceph-disk`](#) section in the *Administration Guide* for Red Hat Ceph Storage 4.

FileStore is no longer supported in production

The FileStore OSD back end is now deprecated because the new BlueStore back end is now fully supported in production. For details, see the [How to migrate the object store from FileStore to BlueStore](#) section in the *Red Hat Ceph Storage Installation Guide*.

Ceph configuration file is now deprecated

The Ceph configuration file (`ceph.conf`) is now deprecated in favor of new centralized configuration stored in Ceph Monitors. For details, see the [The Ceph configuration database](#) section in the *Red Hat Ceph Storage Configuration Guide*.

CHAPTER 6. BUG FIXES

This section describes bugs with significant impact on users that were fixed in this release of Red Hat Ceph Storage. In addition, the section includes descriptions of fixed known issues found in previous versions.

6.1. THE CEPH ANSIBLE UTILITY

Alertmanager does not log errors when self-signed or untrusted certificates are used

Previously, when using untrusted CA certificates, Alertmanager generated many errors in the logs.

With this release, the **ceph-ansible** can set the **insecure_skip_verify** parameter to **true** in the **alertmanager.yml** file by setting **alertmanager_dashboard_api_no_ssl_verify: true** in the **group_vars/all.yml** file when using self-signed or untrusted certificates and the Alertmanager does not log those errors anymore and works as expected.

([BZ#1936299](#))

Use a fully-qualified domain name (FQDN) when HTTPS is enabled in a multi-site configuration

Previously, in a multi-site Ceph configuration, **ceph-ansible** would not differentiate between HTTP and HTTPS and set the zone endpoints with the IP address instead of the host name when HTTPS was enabled.

With this release, **ceph-ansible** uses the fully-qualified domain name (FQDN) instead of the IP address when HTTPS is enabled and the zone endpoints are set with the FQDN and match the TLS certificate CN.

([BZ#1965504](#))

Add the `--pid-limits` parameter as `-1` for podman and `0` for docker in the systemd file to start the container

Previously, the number of processes allowed to run in containers, 2048 for podman and 4096 for docker, were not sufficient to start some containers which needed to start more processes than these limits.

With this release, you can remove the limit of maximum processes that can be started by adding the **`--pid-limits`** parameter as **`-1`** for podman and as **`0`** for docker in the systemd unit files. As a result, the containers start even if you customize the internal processes which might need to run more processes than the default limits.

([BZ#1987041](#))

ceph-ansible pulls the monitoring container images in a dedicated task behind the proxy

Previously, **ceph-ansible** would not pull the monitoring container images such as Alertmanager, Prometheus, node-exporter, and Grafana in a dedicated task and would pull images when the systemd service was started.

With this release, **ceph-ansible** supports pulling monitoring container images behind a proxy.

([BZ#1995574](#))

The **ceph-ansible** playbook creates the radosgw system user and works as expected

Previously, the **ceph-ansible** playbook failed to create the radosgw system user and failed to deploy the dashboard when **rgw_instances** was set at the **host_vars** or **group_vars** level in a multi-site deployment. This variable is not set on Ceph Monitor nodes and given that this where the tasks are delegated, it failed.

With this release, **ceph-ansible** checks all the Ceph Object Gateway instances that are defined and sets a boolean fact to check if at least one instance has the **rgw_zonemaster** set to 'True'. The radosgw system user is created and the playbook works as expected.

([BZ#2034595](#))

The Ansible playbook does not fail when used with `--limit` option

Previously, the **dashboard_server_addr** parameter was unset when the Ansible playbook was run with the **--limit** option and the playbook would fail if the play target did not match the Ceph Manager hosts in a non-collocated scenario.

With this release, you have to set the **dashboard_server_addr** parameter on the Ceph Manager nodes and the playbook works as expected.

([BZ#2063029](#))

6.2. CEPH MANAGEMENT DASHBOARD

The "Client Connection" panel is replaced with "MGRs" on the Grafana dashboard

Previously, the "Client Connection" panel displayed the Ceph File System information and was not meaningful.

With this release, "MGRs" replaces the "Client Connection" panel and displays the count of the active and standby Ceph Managers.

([BZ#1992178](#))

The Red Hat Ceph Storage Dashboard displays the values for disk IOPS

Previously, the Red Hat Ceph Storage Dashboard would not display the Ceph OSD disk performance in the *Hosts* tab.

With this release, the Red Hat Ceph Storage Dashboard displays the expected information about the Ceph OSDs, host details, and the Grafana graphs.

([BZ#1992246](#))

6.3. THE CEPH VOLUME UTILITY

The `add-osd.yml` playbook does not fail anymore while creating new OSDs

Previously, the **add-osd.yml** playbook would fail when new OSDs were added using **ceph-ansible**. This was due to the **ceph-volume lvm batch** limitation which does not allow addition of new OSDs in a non-interactive mode.

With this release, the **--yes** and **--report** options are not passed to the command-line interface and the **add-osd.yml** playbook works as expected when creating new OSDs.

([BZ#1896803](#))

6.4. CEPH OBJECT GATEWAY

The `rgw_bucket_quota_soft_threshold` parameter is disabled

Previously, the Ceph Object Gateway fetched utilization information from the bucket index if the cached utilization reached `rgw_bucket_quota_soft_threshold` causing high operations on the bucket index and slower requests.

This release removes the `rgw_bucket_quota_soft_threshold` parameter and uses the cached stats resulting in better performance even if the quota limit is almost reached.

(BZ#1965314)

The `radosgw-admin datalog trim` command does not crash while trimming a marker

Previously, the `radosgw-admin datalog trim` command would crash when trimming a marker in the current generation from `radosgw-admin` due to a logic error.

This release fixes a logic error and log trimming occurs without the `radosgw-admin datalog trim` command crashing.

(BZ#1981860)

6.5. CEPH MANAGER PLUGINS

The cluster health changes are no longer committed to persistent storage

Previously, rapid changes to the health of the storage cluster caused excessive logging to the `ceph.audit.log`.

With this release, the `health_history` is not logged to the `ceph.audit.log` and cluster health changes are no longer committed to persistent storage.

(BZ#2004738)

CHAPTER 7. KNOWN ISSUES

This section documents known issues found in this release of Red Hat Ceph Storage.

7.1. CEPH MANAGEMENT DASHBOARD

Disk AVG utilization panel shows N/A on the Red Hat Ceph Storage Dashboard

The Red Hat Ceph Storage Dashboard displays a value of N/A on the Overall host performance AVG disk utilization panel thereby showing incorrect Grafana queries.

7.2. CEPH OBJECT GATEWAY

Lifecycle processing stuck in “PROCESSING” state for a given bucket

If a Ceph Object Gateway server is unexpectedly restarted when the lifecycle processing is in progress for a given bucket, that bucket does not resume processing lifecycle work for at least two scheduling cycles and is stuck in “PROCESSING” state. This is an expected behavior as it is intended to avoid multiple Ceph Object gateway instances or threads from processing the same bucket simultaneously, especially when the debugging is in progress in production. For future releases, the lifecycle processing shall restart on the following day, when debugging is not enabled.

(BZ#2072681)

7.3. THE CEPH ANSIBLE UTILITY

Ceph containers fails during startup

The new deployment of Red Hat Ceph Storage-4.3.z1 on Red Hat Enterprise Linux-8.7 (or higher) or Upgrade of Red Hat Ceph Storage-4.3.z1 to 5.X with host OS as Red Hat Enterprise Linux-8.7(or higher) fails at **TASK [ceph-mgr : wait for all mgr to be up]**. The behavior of **podman** released with Red Hat Enterprise Linux 8.7 had changed with respect to SELinux relabeling. Due to this, depending on their startup order, some Ceph containers would fail to start as they would not have access to the files they needed.

As a workaround, refer to the knowledge base [RHCS 4.3 installation fails while executing the command ``ceph mgr dump``](#)

(BZ#2235299)

CHAPTER 8. SOURCES

The updated Red Hat Ceph Storage source code packages are available at the following location:

- For Red Hat Enterprise Linux 7:
<http://ftp.redhat.com/redhat/linux/enterprise/7Server/en/RHCEPH/SRPMS/>
- For Red Hat Enterprise Linux 8:
<http://ftp.redhat.com/redhat/linux/enterprise/8Base/en/RHCEPH/SRPMS/>