



Red Hat Ceph Storage 4.0

Release Notes

Release notes for Red Hat Ceph Storage 4.0

Red Hat Ceph Storage 4.0 Release Notes

Release notes for Red Hat Ceph Storage 4.0

Legal Notice

Copyright © 2020 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 document describes the major features and enhancements implemented in Red Hat Ceph Storage in a particular release. The document also includes known issues and bug fixes.

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	7
3.4. CEPH MEDIC	7
3.5. ISCSI GATEWAY	7
3.6. OBJECT GATEWAY	8
3.7. PACKAGES	8
3.8. RADOS	8
3.9. BLOCK DEVICES (RBD)	10
3.10. RBD MIRRORING	11
CHAPTER 4. BUG FIXES	12
4.1. THE CEPH-ANSIBLE UTILITY	12
4.2. OBJECT GATEWAY	12
4.3. RADOS	12
CHAPTER 5. TECHNOLOGY PREVIEWS	14
5.1. CEPH FILE SYSTEM	14
5.2. BLOCK DEVICES (RBD)	14
5.3. OBJECT GATEWAY	14
CHAPTER 6. KNOWN ISSUES	16
6.1. THE CEPH-ANSIBLE UTILITY	16
6.2. CEPH MANAGEMENT DASHBOARD	16
6.3. PACKAGES	18
6.4. RED HAT ENTERPRISE LINUX	18
CHAPTER 7. DEPRECATED FUNCTIONALITY	19
7.1. THE CEPH-ANSIBLE UTILITY	19
7.2. THE CEPH-DISK UTILITY	19
7.3. RADOS	19
CHAPTER 8. SOURCES	20

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 4.0 contains many contributions from the Red Hat Ceph Storage team. Additionally, 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
- UbuntuKylin
- 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.

The main features added by this release are:

- [General availability of the BlueStore OSD back end](#)
- [Ceph Management Dashboard](#)
- [Ceph Object Gateway Beast Asio web server](#)
- [Erasure coding for Ceph Block Device](#)
- [Auto-scaling placement groups](#)
- [Web-based installation interface](#)
- [On-wire encryption](#)
- [RBD performance monitoring and metrics gathering tools](#)
- [Red Hat Enterprise Linux in FIPS mode](#)

3.1. THE CEPH-ANSIBLE UTILITY

Ceph OSDs created with `ceph-disk` are migrated to `ceph-volume` during upgrade

When upgrading to Red Hat Ceph Storage 4, all running Ceph OSDs previously created by the `ceph-disk` utility will be migrated to the `ceph-volume` utility because `ceph-disk` has been deprecated in this release.

For bare-metal and container deployments of Red Hat Ceph Storage, the `ceph-volume` utility does a simple scan and takes over the existing Ceph OSDs deployed by the `ceph-disk` utility. Also, do not use these migrated devices in configurations for subsequent deployments. Note that you cannot create any new Ceph OSDs during the upgrade process.

After the upgrade, all Ceph OSDs created by `ceph-disk` will start and operate like any Ceph OSDs created by `ceph-volume`.

Ansible playbooks for scaling of all Ceph services

Previously, `ceph-ansible` playbooks offered limited scale up and scale down features only to core Ceph products, such as Monitors and OSDs. With this update, additional Ansible playbooks allow for scaling of all Ceph services.

Ceph iSCSI packages merged into single package

The `ceph-iscsi-cli` and `ceph-iscsi-config` packages have been merged to one package named `ceph-iscsi`.

The `nfs-ganesha` service is now supported as a standalone deployment

Red Hat Openstack Directory (OSPd) requires the `ceph-ansible` utility to be able to deploy the `nfs-ganesha` service and configure it so that it points to an external, unmanaged, pre-existing, Ceph cluster. As of Red Hat Ceph Storage 4, `ceph-ansible` allows the deployment of an internal `nfs-ganesha` service

with an external Ceph cluster.

Ceph container can now write logs to a respective daemon file

The previous way of logging for containerized Ceph environments did not allow limiting the `journalctl` output when looking at log data by using the `sosreport` collection. With this release, logging can be enabled or disabled for a particular Ceph daemon with the following command:

```
ceph config set daemon.id log_to_file true
```

Where `daemon` is the type of the daemon and `id` is its ID. For example, to enable logging for the Monitor daemon with ID `mon0`:

```
ceph config set mon.mon0 log_to_file true
```

This new feature makes debugging easier.

Ability to configure Ceph Object Gateway to use TLS encryption

This release of Red Hat Ceph Storage provides the ability to configure the Ceph Object Gateway listener with an SSL certificate for TLS encryption by using the `radosgw_frontend_ssl_certificate` variable to secure the Transmission Control Protocol (TCP) traffic.

Ansible playbook for migrating OSDs from FileStore to BlueStore

A new Ansible playbook has been added to migrate OSDs from FileStore to BlueStore. The object store migration is not done as part of the upgrade process to Red Hat Ceph Storage 4. Do the migration after the upgrade completes. For details, see the [How to migrate the object store from FileStore to BlueStore](#) section in the *Red Hat Ceph Storage Administration Guide*.

3.2. CEPH MANAGEMENT DASHBOARD

Improvements to information for pool usage

With this enhancement, valuable information was added to the pools table. The following columns were added: usage, read bytes, write bytes, read operations, and write operations. Also, the *Placement Groups* column was renamed to *Pg Status*.

Red Hat Ceph Storage Dashboard alerts

Red Hat Ceph Storage Dashboard supports alerts based on Ceph metrics and configured thresholds. The Prometheus AlertManager configures, gathers, and triggers the alerts. The alerts are displayed in the Dashboard as a pop-up notifications in the upper-right corner. You can view details of recent alerts in **Cluster > Alerts**. You can configure the alerts only in Prometheus, but you can temporarily mute them from the Dashboard by creating "Alert Silences" in **Cluster > Silences**.

Displaying and hiding Ceph components in Dashboard

In the Red Hat Ceph Storage Dashboard you can display or hide Ceph components, such as Ceph iSCSI, RBD mirroring, Ceph Block Devices, Ceph File System, or Ceph Object Gateway. This feature allows you to hide components that are not configured.

Ceph dashboard has been added to the Ceph Ansible playbooks

With this release, the Ceph dashboard installation code was merged into the Ceph Ansible playbooks. Ceph Ansible does a containerized deployment of the Ceph dashboard regardless of the Red Hat Ceph Storage deployment type, bare metal or containers. These four new roles were added: **ceph-**

grafana, **ceph-dashboard**, **ceph-prometheus** and **ceph-node-exporter**.

Red Hat Ceph Storage Dashboard alerts

Red Hat Ceph Storage Dashboard supports alerts based on Ceph metrics and configured thresholds. The Prometheus AlertManager configures, gathers, and triggers the alerts.

Viewing the cluster hierarchy from the Dashboard

Red Hat Ceph Storage Dashboard provides the ability to view the cluster hierarchy. For details see the [Viewing the CRUSH map](#) section in the *Dashboard Guide* for Red Hat Ceph Storage 4.

3.3. CEPH FILE SYSTEM

ceph -w now shows information about CephFS scrubs

Previously, it was not possible to check the ongoing Ceph File System (CephFS) scrubs status aside from checking the Metadata server (MDS) logs. With this update, the **ceph -w** command, shows information about active CephFS scrubs to better understand the status.

ceph-mgr volumes module for managing CephFS exports

This release provides Ceph Manager (**ceph-mgr**) volumes module to manage Ceph File System (CephFS) exports. The volumes module implements the following file system export abstractions:

- FS volumes, an abstraction for CephFS file systems
- FS subvolumes, an abstraction for independent CephFS directory trees
- FS subvolume groups, an abstraction for a directory higher than FS subvolumes to effect policies, such as file layouts, across a set of subvolumes.

In addition, these new commands are now supported:

- **fs subvolume ls** for listing subvolumes
- **fs subvolumegroup ls** for listing subvolumes groups
- **fs subvolume snapshot ls** for listing subvolume snapshots
- **fs subvolumegroup snapshot ls** for listing subvolume group snapshots
- **fs subvolume rm** for removing snapshots

3.4. CEPH MEDIC

ceph-medic can check the health of Ceph running in containers

With this release, the **ceph-medic** utility can now check the health of a Red Hat Ceph Storage cluster running within a container.

3.5. ISCSI GATEWAY

Non-administrators can now be used for the **ceph-iscsi** service

As of Red Hat Ceph Storage 4, non-administrative Ceph users may be used for the **ceph-iscsi** service by setting the **cluster_client_name** in the **/etc/ceph/iscsi-gateway.cfg** on all iSCSI Gateways. This allows resources to be restricted based on users.

Running Ceph iSCSI Gateways can now be removed

As of Red Hat Ceph Storage 4, running iSCSI Gateways can now be removed from a Ceph iSCSI cluster for maintenance or to reallocate resources. The gateway's iSCSI target and its portals will be stopped, and all iSCSI target objects for that gateway will be removed from the kernel and the gateway configuration. Removing gateways that are down is not yet supported.

3.6. OBJECT GATEWAY

The Beast HTTP front end

In Red Hat Ceph Storage 4, the default HTTP front end for the Ceph Object Gateway is Beast. The Beast front end uses the **Boost.Beast** library for HTTP parsing and the **Boost.Asio** library for asynchronous I/O. For details, see the [Using the Beast front end](#) in the *Object Gateway Configuration and Administration Guide* for Red Hat Ceph Storage 4.

Support for S3 MFA-Delete

With this release, the Ceph Object Gateway supports S3 MFA-Delete using Time-Based One-Time Password (TOTP) one-time passwords as an authentication factor. This feature adds security against inappropriate data removal. You can configure buckets to require a TOTP one-time token in addition to standard S3 authentication to delete data.

Addition of storage classes and lifecycle transitions to the Ceph Object Gateway

The Ceph Object Gateway now provides support for storage classes, an S3 compatible representation of the Ceph Object Gateway's underlying placement targets, and for lifecycle transitions, a mechanism to migrate objects between classes. Storage classes and lifecycle transitions provide a higher level of control over data placement for the applications that need it. They have many potential uses, including performance tuning for specific workloads, as well as automatic migration of inactive data to cold storage.

Users can now create new IAM policies and roles using REST APIs

With the release of Red Hat Ceph Storage 4, REST APIs for IAM roles and user policies are now available in the same namespace as S3 APIs and can be accessed using the same endpoint as S3 APIs in the Ceph Object Gateway. This allows end users to create new IAM policies and roles using REST APIs.

3.7. PACKAGES

Ability to install a Ceph cluster using a web-based interface

With this release, the Cockpit web-based interface is supported. Cockpit allows you to install a Red Hat Ceph Storage 4 cluster and other components, such as Metadata Servers, Ceph clients, or the Ceph Object Gateway on base-metal or in containers. For details see the [Installing Red Hat Ceph Storage using the Cockpit Web User Interface](#) chapter in the *Red Hat Ceph Storage 4 Installation Guide*. Note that minimal experience with Red Hat Ceph Storage is required.

3.8. RADOS

Ceph on-wire encryption

Starting with Red Hat Ceph Storage 4, you can enable encryption for all Ceph traffic over the network

with the introduction of the messenger version 2 protocol. For details see the [Ceph on-wire encryption](#) chapter in the *Architecture Guide* and [Encryption in transit](#) section in the *Data Security and Hardening Guide* for Red Hat Ceph Storage 4.

OSD BlueStore is now fully supported

BlueStore is a new back end for the OSD daemons that allows for storing objects directly on the block devices. Because BlueStore does not need any file system interface, it improves performance of Ceph storage clusters. To learn more about the BlueStore OSD back end, see the [OSD BlueStore](#) chapter in the *Administration Guide* for Red Hat Ceph Storage 4.

Red Hat Enterprise Linux in FIPS mode

With this release, you can install Red Hat Ceph Storage on Red Hat Enterprise Linux where the FIPS mode is enabled.

Changes the `ceph df` output and a new `ceph osd df` command

The output of the `ceph df` command has been improved. Notably, the `RAW USED` and `%RAW USED` values now show the preallocated space for the `db` and `wal` BlueStore partitions. The `ceph osd df` command shows the OSD utilization stats, such as the amount of written data.

Asynchronous recovery for non-acting OSD sets

Previously, recovery with Ceph was a synchronous process by blocking write operations to objects until those objects were recovered. In this release, the recovery process is now asynchronous by not blocking write operations to objects only in the non-acting set of OSDs. This new feature requires having more than the minimum number of replicas, as to have enough OSDs in the non-acting set.

The new configuration option, `osd_async_recovery_min_cost`, controls how much asynchronous recovery to do. The default value for this option is `100`. A higher value means asynchronous recovery will be less, whereas a lower value means asynchronous recovery will be more.

Configuration is now stored in Monitors accessible by using `ceph config`

In this release, Red Hat Ceph Storage centralizes configuration in Monitors instead of using the Ceph configuration file (`ceph.conf`). Previously, changing configuration included manually updating `ceph.conf`, distributing it to appropriate nodes, and restarting all affected daemons. Now, the Monitors manage a configuration database that has the same semantic structure as `ceph.conf`. The database is accessible by the `ceph config` command. Any changes to the configuration are applied to daemons or clients in the system immediately and restarting them is no longer needed. Use the `ceph config -h` command for details on the available set of commands. Note that a Ceph configuration file is still required to identify the Monitor nodes.

Placement groups can now be auto-scaled

Red Hat Ceph Storage 4 introduces the ability to auto-scale placement groups (PGs). The number of placement groups (PGs) in a pool plays a significant role in how a cluster peers, distributes data, and rebalances. Auto-scaling the number of PGs can make managing the cluster easier. The new `pg-autoscaling` command provides recommendations for scaling PGs, or automatically scales PGs based on how the cluster is being used. For more details about auto-scaling PGs, see the [Auto-scaling placement groups](#) section in the *Storage Strategies Guide* for Red Hat Ceph Storage 4.

Introduction of `diskprediction` module

The Red Hat Ceph Storage `diskprediction` module gathers metrics to predict disk failures before they happen. The module has two modes, cloud and local. With this release, only the local mode is supported. The local mode does not require any external server for data analysis. It uses an internal predictor module for disk prediction service, and then returns the disk prediction result to the Ceph system.

To enable the **diskprediction** module:

```
ceph mgr module enable diskprediction_local
```

To set the prediction mode:

```
ceph config set global device_failure_prediction_mode local
```

To disable the **diskprediction** module:

```
ceph config set global device_failure_prediction_mode none
```

New configurable option: **mon_memory_target**

Red Hat Ceph Storage 4 introduces a new configurable option, **mon_memory_target**, used to set the target amount of bytes for Monitor memory usage. It specifies the amount of memory to allocate and manage using the priority cache tuner for the associated Monitor daemon caches. The default value of **mon_memory_target** is set to 2 GiB and you can change it during runtime with:

```
# ceph config set global mon_memory_target size
```

Prior to this release, as a cluster scaled, the Monitor specific RSS usage exceeded the limits that were set using the **mon_osd_cache_size** option, which led to issues. This enhancement allows for improved management of memory allocated to the monitor caches and keeps the usage within specified limits.

3.9. BLOCK DEVICES (RBD)

Erasure coding for Ceph Block Device

Erasure coding for Ceph Block Device (RBD) is now fully supported. This feature allows RBD to store their data in an erasure coded pool. For details see the [Erasure Coding with Overwrites](#) section in the *Storage Strategies* for Red Hat Ceph Storage 4.

RBD performance monitoring and metrics gathering tools

Red Hat Ceph Storage 4 now incorporates new Ceph Block Device performance monitoring utilities for aggregated RBD image metrics for IOPS, throughput, and latency. Per-image RBD metrics are now available using the Ceph Manager Prometheus module, the Ceph Dashboard, and the **rbd** CLI using the **rbd perf image iostat** or **rbd perf image iotop** commands.

Cloned images can be created from non-primary images

Creating cloned child RBD images from mirrored non-primary parent image is now supported. Previously, cloning of mirrored images was only supported for primary images. When cloning golden images for virtual machines, this restriction prevented the creation of new cloned images from the golden non-primary image. This update removes this restriction, and cloned images can be created from non-primary mirrored images.

Segregating RBD images within isolated namespaces within the same pool

RBD images can now be segregated within isolated namespaces within the same pool. When using Ceph Block Devices directly without a higher-level system, such as OpenStack or OpenShift Container Storage, it was not possible to restrict user access to specific RBD images. When combined with CephX capabilities, users can be restricted to specific pool namespaces to restrict access to RBD images.

Moving RBD images between different pools within the same cluster

This version of Red Hat Ceph Storage adds the ability to move RBD images between different pools within the same cluster. For details, see the [Moving images between pools](#) section in the *Block Device Guide* for Red Hat Ceph Storage 4.

Long-running RBD operations can run in the background

Long-running RBD operations, such as image removal or cloned image flattening, can now be scheduled to run in the background. RBD operations that involve iterating over every backing RADOS object for the image can take a long time depending on the size of the image. When using the CLI to perform one of these operations, the **rbid** CLI is blocked until the operation is complete. These operations can now be scheduled to run by the Ceph Manager as a background task by using the **ceph rbd task add** commands. The progress of these tasks is visible on the Ceph dashboard as well as by using the CLI.

3.10. RBD MIRRORING

Support for multiple active instances of RBD mirror daemon in a single storage cluster

Red Hat Ceph Storage 4 now supports deploying multiple active instances of the RBD mirror daemon in a single storage cluster. This enables multiple RBD mirror daemons to perform replication for the RBD images or pools using an algorithm for chunking the images across the number of active mirroring daemons.

CHAPTER 4. 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.

4.1. THE CEPH-ANSIBLE UTILITY

The Ansible playbook no longer takes hours to complete the `fuser` command

Previously, on a system running thousands of processes, the `fuser` command in the Ansible playbook could take several minutes or hours to complete because it iterates over all PID present in the `/proc` directory. Because of this, the handler tasks take a long time to complete checking if a Ceph process is already running. With this update, instead of using the `fuser` command, the Ansible playbook now checks the socket file in the `/proc/net/unix` directory, and the handler tasks that check the Ceph socket completes almost instantly.

([BZ#1717011](#))

The `purge-docker-cluster.yml` Ansible playbook no longer fails

Previously, the `purge-docker-cluster.yml` Ansible playbook could fail when trying to unmap RADOS Block Devices (RBDs) because either the binary was absent, or because the Atomic host version provided was too old. With this update, Ansible now uses the `sysfs` method to unmap devices if there are any, and the `purge-docker-cluster.yml` playbook no longer fails.

([BZ#1766064](#))

4.2. OBJECT GATEWAY

The `Expiration, Days S3 Lifecycle` parameter can now be set to `0`

The Ceph Object Gateway did not accept the value of `0` for the `Expiration, Days Lifecycle` configuration parameter. Consequently, setting the expiration to `0` could not be used to trigger background delete operation of objects. With this update, `Expiration, Days` can be set to `0` as expected.

([BZ#1493476](#))

Archive zone fetches the current version of the source objects

An object is synced from multiple source zones to archive zone. This behavior can lead to different versions of the same object on the archive zone. This update eliminates duplicate versions from getting created by ensuring that the archive zone fetches the current version of the source object.

([BZ#1760862](#))

4.3. RADOS

A message to set `expected_num_objects` is no longer shown when using BlueStore

With this update, a message that recommends setting the `expected_num_obejcts` parameter during the creation of the BlueStore pools has been removed because this message does not apply when using the BlueStore OSD back end.

([BZ#1650922](#))

Deprecated JSON fields have been removed

This update removes deprecated fields from the JSON output of the **ceph status** command.

([BZ#1739233](#))

CHAPTER 5. 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://access.redhat.com/support/offerings/techpreview/>.

5.1. CEPH FILE SYSTEM

CephFS snapshots

Ceph File System (CephFS) supports taking snapshots as a Technology Preview. Snapshots create an immutable view of the file system at the point in time they are taken.

5.2. 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.

5.3. 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.

Tiering within a cluster by disk type

This release adds the ability to tier within a cluster by disk type as a Technology Preview by using the capabilities for mapping for example pools to placement targets and storage classes. Using lifecycle transition rules, it is possible to cause objects to migrate between storage pools based on policy.

S3 bucket notifications

S3 bucket notifications are now supported as a Technology Preview. When certain events are triggered on an S3 bucket, the notifications can be sent from the Ceph Object Gateway to HTTP, Advanced Message Queuing Protocol (AMQP) 9.1, and Kafka endpoints. Additionally, the notifications can be stored in a "PubSub" zone instead of, or in addition to sending them to the endpoints. "PubSub" is a publish-subscribe model that enables recipients to pull notifications from Ceph.

To use the S3 notifications, install the **librabbitmq** and **librdkafka** packages.

CHAPTER 6. KNOWN ISSUES

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

6.1. THE CEPH-ANSIBLE UTILITY

Red Hat Ceph Storage installation on Red Hat OpenStack Platform fails

The **ceph-ansible** utility becomes unresponsive when attempting to install Red Hat Ceph Storage together with the Red Hat OpenStack Platform 16, and it returns an error similar to the following one:

```
'Error: unable to exec into ceph-mon-dcn1-computehci1-2: no container with name or ID ceph-mon-dcn1-computehci1-2 found: no such container'
```

To work around this issue, update the following part of the **handler_osds.yml** file, located in the **ceph-ansible/roles/ceph-handler/tasks/** directory:

```
- name: unset noup flag
  command: "{{ container_exec_cmd | default('') }}" ceph --cluster {{ cluster }} osd unset noup"
  delegate_to: "{{ groups[mon_group_name][0] }}"
  changed_when: False
```

To:

```
- name: unset noup flag
  command: "{{ hostvars[groups[mon_group_name][0]]['container_exec_cmd'] | default('') }}" ceph --
cluster {{ cluster }} osd unset noup"
  delegate_to: "{{ groups[mon_group_name][0] }}"
  changed_when: False
```

And start the installation process again.

([BZ#1792320](#))

Ansible does not unset the **norebalance** flag after it completes

The **rolling-update.yml** Ansible playbook does not unset the **norebalance** flag after it completes. To work around this issue, unset the flag manually.

([BZ#1793564](#))

Ansible fails to upgrade a multisite Ceph Object Gateway when the Dashboard is enabled

When the Red Hat Ceph Storage Dashboard is enabled, an attempt to use Ansible to upgrade to a further version of Red Hat Ceph Storage fails when trying to upgrade the secondary Ceph Object Gateway site in a multisite setup. This bug does not occur on the primary site or if the Dashboard is not enabled.

([BZ#1794351](#))

6.2. CEPH MANAGEMENT DASHBOARD

The Dashboard does not show correct values for certain configuration options

Both, the Red Hat Ceph Storage Dashboard and the underlying **ceph config show** command do not

return the current values for certain configuration options, such as **fsid**. This is probably because certain core options that are not meant for further modification after deploying a cluster are not updated and a default value is used. As a result, the Dashboard does not show correct values for certain configuration options.

([BZ#1765493](#), [BZ#1772310](#))

NFS Ganesha in Dashboard

The Red Hat Ceph Storage Dashboard currently does not support managing NFS Ganesha.

([BZ#1772316](#))

Dashboard does not support email verification

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.

([BZ#1778608](#))

The OSD histogram graph for read and write operations is not clear

The Red Hat Ceph Storage Dashboard does not display any numbers or description in the OSD histogram graph for read and write operations, and therefore the graph is not clear.

([BZ#1779170](#))

Dashboard returns an error when a **ceph-mgr** module is enabled from the **ceph CLI**

When enabling a Ceph Manager (**ceph-mgr**) module, such as **telemetry**, from the **ceph CLI**, the Red Hat Ceph Storage Dashboard displays the following error message:

```
0 - Unknown Error
```

In addition, the Dashboard does not mark the module as enabled, until the **Refresh** button is clicked.

([BZ#1785077](#))

The Dashboard allows modifying the LUN ID and WWN, which can lead to data corruption

The Red Hat Ceph Storage Dashboard allows you to modify the LUN ID and its World Wide Name (WWN), which is not required after creating the LUN. Moreover, editing these parameters can be dangerous for certain initiators that do not fully support this feature by default. Consequently, editing these parameters after creating them can lead to data corruption. To avoid this, do not modify the LUN ID and WWN in the Dashboard.

([BZ#1786455](#))

The Dashboard does not provide correct Ceph iSCSI error messages

If the Ceph iSCSI returns an error, for example the HTTP "400" code when trying to delete an iSCSI target while a user is logged in, the Red Hat Ceph Storage Dashboard does not forward that error code and message to the Dashboard user using the pop-up notifications, but displays a generic "500 Internal Server Error". Consequently, the message that the Dashboard provides is not informative and even misleading; an expected behavior ("users cannot delete a busy resource") is perceived as an operational failure ("internal server error"). To work around this issue, see the Dashboard logs.

([BZ#1786457](#))

The Dashboard requires to disable iptables rules

The Red Hat Ceph Storage Dashboard is unable to perform any iSCSI operations, such as creating a gateway, unless all **iptables** rules are manually disabled on the Ceph iSCSI node. To do so, use the following command as **root** or the **sudo** user:

```
# iptables -F
```

Note that after a reboot, the rules are enabled again. Either disable them again or delete them permanently.

([BZ#1792818](#))

6.3. PACKAGES

Current version of Grafana causes certain bugs in the Dashboard

Red Hat Ceph Storage 4 uses the Grafana version 5.2.4. This version causes the following bugs in the Red Hat Ceph Storage Dashboard:

- When navigating to **Pools > Overall Performance**, Grafana returns the following error:

```
TypeError: l.c[t.type] is undefined  
true
```

- When viewing a pool's performance details (**Pools > select a pool from the list > Performance Details**) the Grafana bar is displayed along with other graphs and values, but it should not be there.

These bugs will be fixed after rebasing to a newer Grafana version in a future release of Red Hat Ceph Storage.

([BZ#1786107](#), [BZ#1762197](#), [BZ#1765536](#))

6.4. RED HAT ENTERPRISE LINUX

Ansible cannot start NFS Ganesha if SELinux is in enforcing mode

When using SELinux in enforcing mode on Red Hat Enterprise Linux 8.1, the **ceph-ansible** utility fails to start the NFS Ganesha service because SELinux policy currently does not allow creating a directory required for NFS Ganesha.

([BZ#1794027](#), [BZ#1796160](#))

CHAPTER 7. 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.

7.1. THE CEPH-ANSIBLE UTILITY

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 either the `gwcli` utility of the Red Hat Ceph Storage Dashboard to configure the gateway. For details, see the [Using the Ceph iSCSI Gateway](#) chapter in the *Block Device Guide* for Red Hat Ceph Storage 4.

7.2. THE CEPH-DISK UTILITY

`ceph-disk` is deprecated

With this release, 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.

7.3. RADOS

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 4 *Installation Guide*.

Ceph configuration is now deprecated

The Ceph configuration file (`ceph.conf`) is now deprecated in favor of new centralized configuration stored in Monitors. For details, see the [Configuration is now stored in Monitors accessible by using `ceph config`](#) release note.

CHAPTER 8. SOURCES

The updated Red Hat Ceph Storage source code packages are available at <http://ftp.redhat.com/redhat/linux/enterprise/7Server/en/RHCEPH/SRPMS/>.