



Red Hat Ceph Storage 2.2 Release Notes

Release notes for Red Hat Ceph Storage 2.2

Red Hat Ceph Storage Documentation
Team

Red Hat Ceph Storage 2.2 Release Notes

Release notes for Red Hat Ceph Storage 2.2

Legal Notice

Copyright © 2017 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, 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 Software Collections 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. MAJOR UPDATES	5
CHAPTER 4. TECHNOLOGY PREVIEWS	8
CHAPTER 5. KNOWN ISSUES	9
CHAPTER 6. NOTABLE BUG FIXES	15
CHAPTER 7. SOURCES	18

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.

CHAPTER 2. ACKNOWLEDGMENTS

Red Hat Ceph Storage version 2.2 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

CHAPTER 3. MAJOR UPDATES

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

The rados utility now supports the --omap-key-file option

With this update, the **rados** command-line utility supports the **--omap-key-file** option. You can use this option to specify the path to a file containing the binary key for **omap** key-values pairs. The following commands take **--omap-key-file**:

- » **getomapval**
- » **setomapval**
- » **rmomapkey**

ceph-ansible rebased to 2.1.9

The **ceph-ansible** package has been updated to the upstream version 2.1.9, which provides several important bug fixes to the installation process. In addition, the new version provides compatibility with the Ansible automation application 2.2.1.0.

Ansible now supports purging clusters

With this release, the **ceph-ansible** utility supports purging clusters. See the [Purging a Ceph Cluster](#) section in the Red Hat Ceph Storage 2 Installation Guide for Red Hat Enterprise Linux.

osd_scrub_chunk_max is honored also with objects that have many clones

Previously, deep scrubbing of objects that had a number of clones could impact the client performance. With this enhancement, deep scrubbing honors the limit specified by the **osd_scrub_chunk_max** parameter even when an object has many clones. As a result, the impact on client performance is limited.

The Ceph Object Gateway now supports custom HTTP header logging

Sites that use the Civetweb HTTP web server previously lacked the ability to log custom HTTP headers, as they could when using the Apache web server and the FastCGI protocol. With this update, the Ceph Object Gateway supports custom HTTP header logging.

To log custom HTTP headers, enable the operations log socket on the Ceph Object Gateway instance and list the HTTP headers. Add the following parameters to the Ceph configuration file:

```
rgw enable ops log = true
rgw ops log socket path = <path>
rgw log http headers = "<headers>"
```

Replace **<path>** with the path to the operations log socket and **<headers>** with a comma-separated list of custom HTTP headers, for example:

```
rgw enable ops log = true
rgw ops log socket path = /tmp/opslog
rgw log http headers = "http_x_forwarded_for, http_expect,
http_content_md5"
```

The operations log stream then lists the headers as a JSON-formatted key-value list with the "http_x_headers" key.

The Ceph Object Gateway now supports the S3 multipart copy operation

The Ceph Object Gateway now supports the S3 API for multipart copy, including use of the **x-amz-copy-source** header.

The multipart copy operation provides an optimized mechanism for copying existing objects larger than the 5G upload limit of the Amazon Simple Storage Service (S3). For details, see the [Copy Multipart Upload](#) section in the Developer Guide for Red Hat Ceph Storage 2.

OSD heartbeat_check log messages now include IP addresses

The OSD **heartbeat_check** log messages now include IP addresses of the OSD nodes. This enhancement improves identification of the OSD nodes in the Ceph logs. For example, it is no longer necessary to look up which IP correlates to which OSD node (OSD.<number>) for the **heartbeat_check** message in the log.

ceph rebased to 10.2.5

The **ceph** packages have been updated to the upstream version 10.2.5, which provides a number of bug fixes and enhancements over the previous version.

The Ceph Object Gateway now supports the Swift object versioning

The Ceph Object Gateway now supports the Swift object versioning APIs, including correct handling of the **X-Versions-Location** header. The **X-History-Location** header is not supported.

The object versioning is an object-native version control mechanism of Swift object storage, and is a required capability for RefStack conformance. For details, see the [Object Versioning](#) section of the OpenStack documentation.

The radosgw-admin utility now supports new options

The **radosgw-admin** utility now supports the new **--bypass-gc** and **--inconsistent-index** options. Use these options when deleting indexed buckets to bypass the garbage collector and to ignore bucket index consistency, which improves the speed of the deletion.

Ansible now supports adding encrypted OSDs

You can now use the **ceph-ansible** utility to add encrypted OSD nodes.

For details on how to do it, see the [Configuring Ceph OSD Settings](#) section in the Red Hat Ceph Storage 2 Installation Guide for Red Hat Enterprise Linux. For details on how this feature works, see the [Encryption](#) chapter in the Architecture Guide for Red Hat Ceph Storage 2.

Support for the SSL protocol has been added

The Ceph Object Gateway now supports the SSL protocol. Previously, a reverse proxy server with SSL had to be set up to dispatch HTTPS requests. For details, see the [Using SSL with Civetweb](#) chapter in the Ceph Object Gateway Guide.

nfs-ganesha rebased to 2.4.2

The **nfs-ganesha** packages have been updated to the upstream version 2.4.2, which provides a number of bug fixes and enhancements over the previous version.

ceph-client role is now supported

The **ceph-ansible** utility now supports the **ceph-client** role. This new role enables you to copy the Ceph configuration file and administration keyring to node. In addition, you can use this role to create custom pools and clients.

For details, see the [Installing the ceph-client role](#) section in the Red Hat Ceph Storage 2 Installation Guide for Red Hat Enterprise Linux.

The Ceph Object Gateway now supports three zones in a multi-site configuration

You can now configure a third zone in a multi-site configuration of the Ceph Object Gateway. To do so, follow the same steps as when configuring a secondary zone but use different name for the third zone. For details, see the *Multi-site* chapter in the [Object Gateway Guide for Red Hat Enterprise Linux](#) or [Object Gateway Guide for Ubuntu](#).

Red Hat Ceph Storage Developer Guide is now available

The Red Hat Ceph Storage documentation suite now includes a new [Developer Guide](#). This new guide contains the Ceph Object Gateway API reference that was previously included in the Ceph Object Gateway for Red Hat Enterprise Linux or Ubuntu.

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

Red Hat Ceph Storage 2.2 Container Image

Red Hat Ceph Storage 2.2 is now available as a container image. The **rhceph-2-rhel-7** image is now ready to use on the Red Hat registry.

This update also introduces the following changes:

- ✦ The Red Hat Ceph Storage 2.2 image supports the **ceph** user and group. As a result, the Ceph daemons no longer run as **root** and Ceph binaries no longer belongs to **root**.
- ✦ The **OSD_CEPH_DISK_PREPARE** and **OSD_CEPH_DISK_ACTIVATE** commands are now supported. These commands are equivalents of the **ceph-disk prepare** and **ceph-disk activate** commands.

For detailed information how to bootstrap a containerized Ceph cluster see the [Deploying Red Hat Ceph Storage 2 as a Container Image \(Technology Preview\)](#) article on the [Red Hat Customer Portal](#).

To learn more about Linux Containers, see the [Red Hat Enterprise Linux Atomic Host 7 Getting Started with Containers](#) guide.

CHAPTER 5. KNOWN ISSUES

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

Realm names must be updated on each cluster separately

In a multi-site configuration, the name of a realm is only stored locally and is not shared as part of the period. As a consequence, when the name is changed on one cluster, the name is not updated on the other cluster. To rename the realm, execute the `radosgw-admin realm rename` command separately on each cluster. ([BZ#1423886](#))

Calamari sometimes fails to discover some cluster nodes

The Calamari API sometimes fails to discover some cluster nodes. To work around this problem, restart the Calamari service:

```
# systemctl restart calamari.service
```

([BZ#1420537](#))

Multi-site configuration of the Ceph Object Gateway sometimes fails when options are changed at runtime

When the `rgw md log max shards` and `rgw data log num shards` options are changed at runtime in multi-site configuration of the Ceph Object Gateway, the `radosgw` process terminates unexpectedly with a segmentation fault.

To avoid this issue, do not change the aforementioned options at runtime, but set them during the initial configuration of the Ceph Object Gateway. ([BZ#1330952](#))

Dynamic feature updates are not replicated

When a feature is disabled or enabled on an already existing image and the image is mirrored to a peer cluster, the feature is not disabled or enabled on the replicated image. ([BZ#1344262](#))

Unable to write data on a promoted image after a non orderly shutdown

In RBD mirroring configuration, after a non orderly shutdown of the local cluster, images are demoted to non-primary on the local cluster and promoted to primary on the remote cluster. If this happens and the `rbd-mirror` daemon is not restarted on the remote cluster, it is not possible to write data on the promoted image because `rbd-daemon` considers the demoted image on the local cluster to be the primary one. To avoid this issue, restart the `rbd-mirror` daemon to gain the read/write access to the promoted image. ([BZ#1365648](#))

Mirroring image metadata is not supported

Image metadata are not currently replicated to a peer cluster. ([BZ#1344212](#))

Disabling image features is incorrectly allowed on non-primary images

With RADOS Block Device (RBD) mirroring enabled, non-primary images are expected to be read-only. An attempt to disable image features on non-primary images could cause an indefinite wait. This operation should be disallowed on non-primary images.

To avoid this issue, make sure to disable image features only on the primary image. ([BZ#1353877](#))

Users created by using the Calamari API do not have permissions to run the API commands

When a user is created by using the Calamari REST API (`api/v2/user`), the user does not have permissions to run most of the Calamari API commands. Consequently, an attempt to run the commands fails with the following error message:

```
"You do not have permission to perform this action"
```

To work around this issue, use the `calamari-ctl add_user` command from the command line when creating new users. ([BZ#1356872](#))

The GNU tar utility currently cannot extract archives directly into the Ceph Object Gateway NFS mounted file systems

The current version of the GNU tar utility makes overlapping write operations when extracting files. This behavior breaks the strict sequential write restriction in the current version of the Ceph Object Gateway NFS. In addition, GNU tar reports these errors in the usual way, but it also by default continues extracting the files after reporting the errors. As a result, the extracted files can contain incorrect data.

To work around this problem, use alternate programs to copy file hierarchies into the Ceph Object Gateway NFS. Recursive copying by using the `cp -r` command works correctly. Non-GNU archive utilities might be able to correctly extract the tar archives, but none have been verified. ([BZ#1418606](#))

Ansible fails to install OSDs if they point to directories

Ansible does not support installation of OSDs that point to directories and not to partitions. As a consequence, an attempt to install such OSDs fails. ([BZ#1361228](#))

Results from deep scrubbing are overwritten by shallow scrubbing

When performing shallow scrubbing after deep scrubbing, results from deep scrubbing are overwritten by results from shallow scrubbing. As a consequence, the deep scrubbing results are lost. ([BZ#1330023](#))

The NFS interface for the Ceph Object Gateway does not show bucket size or number of blocks

The NFS interface of the Ceph Object Gateway lists buckets as directories. However, the interface always shows that the directory size and the number of blocks is `0`, even if some data is written to the buckets. ([BZ#1359408](#))

Certain image features are not supported with the RBD kernel module

The following image features are not supported with the current version of the RADOS Block Device (RBD) kernel module (**krbd**) that is included in Red Hat Enterprise Linux 7.3:

- ✦ **object-map**
- ✦ **deep-flatten**
- ✦ **journaling**
- ✦ **fast-diff**

However, by default the **ceph-installer** utility creates RBDs with the aforementioned features enabled. As a consequence, an attempt to map the kernel RBDs by running the **rbd map** command fails.

To work around this issue, disable the unsupported features by setting the **rbd_default_features = 1** option in the Ceph configuration file for kernel RBDs or dynamically disable them by running the following command:

```
rbd feature disable <image> <feature>
```

This issue is a limitation only in kernel RBDs, and the features work as expected with user-space RBDs. ([BZ#1340080](#))

Swift SLOs cannot be read from any other zones

The Ceph Object Gateway fails to fetch manifest files of Swift Static Large Objects (SLO). As a consequence, an attempt to read these objects from any other zone than the zone where the object was originally uploaded fails. ([BZ#1423858](#))

The Calamari REST-based API fails to edit user details

An attempt to use the Calamari REST-based API to edit user details fails with an error. To change user details, use the **calamari-ctl** command-line utility. ([BZ#1338649](#))

The **rbd bench write** command fails when **--io-size** is equal to the image size

The **rbd bench-write --io-size <size> <image>** command fails with a segmentation fault if the size specified by the **--io-size** option is equal to the image size.

To avoid this problem, make sure that the value of **--io-size** is smaller than the image size. ([BZ#1362014](#))

Calamari sometimes does not respond when sending a PATCH Request

The Calamari API does not respond when making PATCH requests to **/api/v2/cluster/FSID/osd/OSD_ID** if the requests does not change any fields on the OSD from their present values. ([BZ#1338688](#))

The **rados list-inconsistent-obj** command does not highlight inconsistent shards when it could have

The output of the **rados list-inconsistent-obj** command does not explicitly show which shard is inconsistent when it could have. ([BZ#1363949](#))

An LDAP user can access buckets created by a local RGW user with the same name

The RADOS Object Gateway (RGW) does not differentiate between a local RGW user and an LDAP user with the same name. As a consequence, the LDAP user can access the buckets created by the local RGW user.

To work around this issue, use different names for RGW and LDAP users. ([BZ#1361754](#))

Simultaneous upload operations to the same file cause I/O errors

Simultaneous upload operations to the same file location by different NFS clients cause I/O errors on both clients. Consequently, no data is updated in the Ceph Object Gateway cluster; if an object already existed in the cluster in the same location, it is unchanged.

To work around this problem, do not simultaneously upload to the same file location. ([BZ#1420328](#))

Ansible and "ceph-disk" fail to create encrypted OSDs if the cluster name is different than "ceph"

The **ceph-disk** utility does not support configuring the **dmccrypt** utility if the cluster name is different than "ceph". Consequently, it is not possible to use the **ceph-ansible** utility to create encrypted OSDs if you use a custom cluster name.

To avoid this problem, use the default cluster name, which is "ceph". ([BZ#1391920](#))

Ansible fails to add a monitor to an upgraded cluster

An attempt to add a monitor to a cluster by using the Ansible automation application after upgrading the cluster from Red Hat Ceph Storage 1.3 to 2 fails on the following task:

```
TASK: [ceph-mon | collect admin and bootstrap keys]
```

This happens because the original monitor keyring was created with the **mds "allow"** capability while the newly added monitor requires a keyring with the **mds "allow *"** capability.

To work around this issue, after installing the **ceph-mon** package, manually copy the administration keyring from an already existing monitor node to the new monitor node:

```
scp /etc/ceph/<cluster_name>.client.admin.keyring  
<target_host_name>:/etc/ceph
```

For example:

```
# scp /etc/ceph/ceph.client.admin.keyring node4:/etc/ceph
```

Then use Ansible to add the monitor as described in the [Adding a Monitor with Ansible](#) section of the [Administration Guide](#) for Red Hat Ceph Storage 2. ([BZ#1357292](#))

Ansible does not support removing monitor or OSD nodes

The current version of the **ceph-ansible** utility does not support removing monitor or OSD nodes. To remove monitor or OSD nodes from a cluster, use the manual procedure. For more information, see the [Administration Guide](#) for Red Hat Ceph Storage 2. ([BZ#1366807](#))

ceph-radosgw does not start after upgrading from 1.3 to 2 if a non-default value is used for rgw_region_root_pool and rgw_zone_root_pool

The **ceph-radosgw** service does not start after upgrading the Ceph Object Gateway from 1.3 to 2, if the Gateway uses non-default values for the **rgw_region_root_pool** and **rgw_zone_root_pool** parameters.

See the [Inconsistent zonegroup/zone state in Rados GW after upgrade of multizone site to Ceph 2](#) solution on the Red Hat Customer Portal for details on how to work around this issue. (BZ#1396956)

Old zone group name is sometimes displayed alongside with the new one

In a multi-site configuration when a zone group is renamed, other zones can in some cases continue to display the old zone group name in the output of the **radosgw-admin zonegroup list** command.

To work around this issue:

1. Verify that the new zone group name is present on each cluster.
2. Remove the old zone group name:

```
$ rados -p .rgw.root rm zonegroups_names.<old-name>
```

(BZ#1423402)

Calamari sometimes incorrectly outputs "null" as a value

When the Calamari REST-based API is used to get details of a CRUSH rule in the Ceph cluster, the output contains "null" as a value for certain fields in the **steps** section of the CRUSH rule. The fields containing null values can be safely ignored for the respective steps in the CRUSH rule. However, do not use "null" as a value for any field when doing a PATCH operation. Using null values in such a case causes the operation to fail. (BZ#1342504)

The Calamari API returns the "server error (500)" error when changing the take step

When changing a CRUSH rule, modifying the **take** step type to any other value than **take** causes the Calamari API to return the "server error (500)" error.

To avoid this issue, do not change the **take** step to any other value. (BZ#1329216)

Ansible does not properly handle unresponsive tasks

Certain tasks, for example adding monitors with the same host name, cause the **ceph-ansible** utility to become unresponsive. Currently, there is no timeout set after which the unresponsive tasks is marked as failed. (BZ#1313935)

Object sync requests are sometimes skipped

In multi-site configurations of the Ceph Object Gateway, a non-master zone can be promoted to the master zone. In most cases, the master zone's gateway or gateways are still running when this happens. However, if the gateways are down, it can take up to 30 seconds after their restart until the gateways notice that another zone was promoted. During this time, the gateways can miss changes

to buckets that occur on other zones. Consequently, object sync requests are skipped.

To work around this issue, pull the new master's period to the old master zone before restarting the old master zone:

```
$ radosgw-admin period pull --remote=<new-master-zone-id>
```

For details on pulling the period, see the [Ceph Object Gateway Guide for Red Hat Enterprise Linux](#) or the [Ceph Object Gateway Guide for Ubuntu](#). (BZ#1362639)

CHAPTER 6. NOTABLE BUG FIXES

This section describes bugs fixed in this release of Red Hat Ceph Storage that have significant impact on users.

"Operation not permitted" errors are no longer incorrectly returned

When using a client whose MDS capabilities are limited by the `path=` parameter, operations in newly created directories in certain cases failed with the "Operation not permitted" errors (EPERM). The underlying source code has been modified, and such errors are no longer returned. ([BZ#1415260](#))

Buckets no longer have incorrect time stamps

Previously, buckets created by the Simple Storage Service (S3) API on the Ceph Object Gateway before mounting the Ganesha NFS interface had incorrect time stamps when viewed from NFS. With this update, the NFS service uses time stamps that are based on the correct times of creation or modification of buckets. As a result, buckets created by the S3 API no longer have incorrect time stamps. ([BZ#1359404](#))

Setting file permissions and ownership attributes no longer fails on existing files and directories

Previously, the NFS Ganesha file system failed to serialize and store UNIX attributes on existing files and directories. Consequently, file permissions and ownership attributes that were set after file or directory creation were not correctly stored. The underlying source code has been modified, and setting file permissions and ownership attributes no longer fails on existing files and directories. ([BZ#1358020](#))

The `radosgw-admin orphan find` command works as expected

When listing objects, a segment marker caused incorrect listing of a subset of the Ceph Object Gateway internal objects. This behavior caused the `radosgw-admin orphan find` command to enter an infinite loop. This bug has been fixed, and the `radosgw-admin orphan find` command now works correctly. ([BZ#1371212](#))

The "ceph df" output no longer includes OSD nodes marked as "out"

The `ceph df` command shows cluster free space. Previously, the OSD node that were marked as `out` were incorrectly included in the output of `ceph df`. Consequently, if the Ceph cluster included an OSD node that was marked as `out`, the output of `ceph df` was incorrect. This bug has been fixed, and `ceph df` now correctly reports cluster free space. ([BZ#1391250](#))

Listing bucket info data no longer causes the OSD daemon to terminate unexpectedly

Due to invalid memory access in an object class operation, the `radosgw-admin bi list --max-entries=1` command in some cases caused the Ceph OSD daemon to terminate unexpectedly with a segmentation fault. This bug has been fixed, and listing bucket info data no longer causes the OSD daemon to crash. ([BZ#1390716](#))

The Ceph Object Gateway now correctly logs when HTTP clients get disconnected

Due to incorrect error translation, unexpected disconnections by HTTP clients were incorrectly logged as **HTTP 403: authorization failed** errors. As a consequence, administrators could believe that an actual authentication failure had occurred, and that this failure was visible to clients. With this update, the Ceph Object Gateway handles the error translation correctly and logs a proper error message when the HTTP clients get disconnected. (BZ#1417178)

OSD nodes no longer fail to start after a reboot

When the **ceph-osd** service was enabled for a given OSD device, a race condition in some cases occurred between the **ceph-osd** and the **ceph-disk** services at boot time. As a consequence, the OSD did not start after a reboot. With this update, the **ceph-disk** utility now calls the **systemctl enable** and **disable** commands with the **--runtime** option so that the **ceph-osd** units are lost after a reboot. As result, OSD nodes start as expected after a reboot. (BZ#1391197)

Restart of the radosgw service on clients is no longer needed after rebooting the cluster

Previously, after rebooting the Ceph cluster, it was necessary to restart the **radosgw** service on the Ceph Object Gateway clients to restore the connection with the cluster. With this update, the restart of **radosgw** is no longer needed. (BZ#1363689)

Upgrading encrypted OSDs is now supported

Previously, the **ceph-ansible** utility did not support adding encrypted OSD nodes. As a consequence, an attempt to upgrade to a newer, minor, or major version failed on encrypted OSD nodes. In addition, Ansible returned the following error message during the disk activation task:

```
mount: unknown filesystem type 'crypto_LUKS'
```

With this update, **ceph-ansible** supports adding encrypted OSD nodes, and upgrading works as expected. (BZ#1366808)

The Ceph Object Gateway now passes all Swift tests in the RefSTack Tempest test suite version 10.0.0-3

Previously, the Ceph Object Gateway failed certain RefSTack Tempest tests, such as the TempURL and object versioning tests. With this update, the underlying source code has been modified, and the Ceph Object Gateway now correctly passes all tests.

In addition, to pass the "(0) content-length header after object deletion present" test, set the **rgw print prohibited content length** setting in the Ceph configuration file to **true**.

If the Ceph Object Gateway is configured for Object Store and not for Swift, perform the following steps to pass the tests:

1. During the Tempest configuration, set the following parameters in the Ceph configuration file:

```
rgw_swift_url_prefix = "/"
rgw_enable_apis=swift, swift_auth, admin
```

2. Once the configuration is complete, comment the parameters out:

■

```
# rgw_swift_url_prefix  
# rgw_enable_apis
```

See the [config_tempest.py breaks if Rados Gateway is configured for object-store solution](#) for details. ([BZ#1252600](#))

Removing a Ceph Monitor no longer returns an error

When removing a Ceph Monitor by using the `ceph mon remove` command, the Monitor was successfully removed, but an error message similar to the following was returned:

```
Error EINVAL: removing mon.magna072 at 10.8.128.72:6789/0, there will  
be 3 monitors
```

The underlying source code has been fixed, and the error is no longer returned when removing Ceph Monitors. ([BZ#1394495](#))

OSD nodes no longer crash when an I/O error occurs

Previously, if an I/O error occurred on one of the objects in an erasure-coded pool during recovery, the primary OSD node of the placement group containing the object hit the runtime check. Consequently, this OSD terminated unexpectedly. With this update, Ceph leaves the object unrecovered without hitting the runtime check. As a result, OSDs no longer crash in such a case. ([BZ#1414613](#))

CHAPTER 7. SOURCES

The updated Red Hat Ceph Storage packages are available at the following locations:

- ✦ for Red Hat Enterprise Linux:
<ftp://ftp.redhat.com/redhat/linux/enterprise/7Server/en/RHCEPH/SRPMS/>
- ✦ for Ubuntu: <https://rhcs.download.redhat.com/ubuntu/>