



Red Hat Ceph Storage 1.3.1 Release Notes

Red Hat Ceph Storage 1.3.1 release notes.

Red Hat Customer Content
Services

Red Hat Ceph Storage 1.3.1 Release Notes

Red Hat Ceph Storage 1.3.1 release notes.

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 a list of fixed bugs.

Table of Contents

CHAPTER 1. ACKNOWLEDGMENTS	3
CHAPTER 2. OVERVIEW	4
2.1. PACKAGING	4
2.2. CEPH STORAGE CLUSTER	4
2.3. CEPH BLOCK DEVICE	4
2.4. CEPH OBJECT GATEWAY	5
CHAPTER 3. KNOWN ISSUES	6
CHAPTER 4. FIXED BUGS	9

CHAPTER 1. ACKNOWLEDGMENTS

Red Hat Ceph Storage version 1.3.1 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 2. OVERVIEW

Red Hat Ceph Storage version 1.3.1 is the third release of Red Hat Ceph Storage. New features for Ceph Storage include:

2.1. PACKAGING

Satellite integration

The Red Hat Ceph Storage version 1.3.1 allows users to host package repositories in the Red Hat Satellite 6 server and also manage entitlements in Satellite. Once you registered your Ceph nodes with Satellite, you can deliver upgrades to the cluster without allowing a direct connection to the Internet, as well as search and view errata applicable to the cluster nodes. For detailed information see the [How to Register Ceph with Satellite](#) article on the [Customer Portal](#).

2.2. CEPH STORAGE CLUSTER

Configurable "suicide" option

The suicide timeout option is now configurable. The option ensures that poorly behaving OSDs self-terminate instead of running in degraded states and slowing traffic.

New Foreman-based installer added as a Technology Preview

The **rhcs-installer** package provides a new Foreman-based installer. This update adds the new rhcs-installer package to Red Hat Ceph Storage as a Technology Preview.

New "osd crush tree" command

A new command, **osd crush tree**, has been added. The command lists OSDs in a particular bucket.

Use "recovering" instead of "recovery" when listing PGs

The **pg ls**, **pg ls-by-pool**, **pg ls-by-primary**, and **pg ls-by-osd** commands no longer take the **recovery** argument. Use the **recovering** argument instead to include also the recovering placement groups (PG).

Upstart respawn limit changes

On Ubuntu, the Upstart respawn limit has been changed from 5 restarts in 30 seconds to 3 restarts in 30 minutes for the Ceph OSD and monitor daemons.

2.3. CEPH BLOCK DEVICE

Tracing with the RBD replay feature

The **RBD Replay** toolkit provides utilities for capturing and replaying RADOS Block Device (RBD) workloads. It uses the Linux Trace Toolkit: next generation (LLTng) tracing framework. For detailed information, see the [Tracing RADOS Block Device \(RBD\) Workloads with the RBD Replay Feature](#) article on the [Customer Portal](#).

2.4. CEPH OBJECT GATEWAY

No need to disable "requiretty" for root

The Ceph Object Gateway no longer requires the **requiretty** setting to be disabled in the sudoers configuration for the root user.

New "rgw_user_max_buckets" option

Administrators of the Ceph Object Gateway can now configure the maximum number of buckets for users by using the new **rgw_user_max_buckets** option in the Ceph configuration file.

CHAPTER 3. KNOWN ISSUES

Graphs for monitor hosts are not displayed

Graphs for monitor hosts are not displayed in the Calamari server GUI when selecting them from the **Graphs** drop-down menu. ([BZ#1223335](#))

The "Update" button is not disabled when a check box is cleared

In the Calamari server GUI on the **Manage > Cluster Settings** page, the **Update** button is not disabled when a check box is cleared. Moreover, further clicking on the **Update** button displays an error dialog box, which leaves the button unusable. To work around this issue, reload the page as suggested in the error dialog. ([BZ#1223656](#))

Ceph init script calls a non-present utility

The Ceph init script calls the **ceph-disk** utility which is not present on monitor nodes. Consequently, when running the **service ceph start** command with no arguments, the init script returns an error. This error does not effect functionality of starting and stopping daemons. ([BZ#1225183](#))

Yum upgrade failures after a system upgrade

The **Yum** utility can fail with transaction errors after upgrading Red Hat Enterprise Linux 6 with Red Hat Ceph Storage 1.2 to Red Hat Enterprise Linux 7 with Red Hat Ceph Storage 1.3. This is because certain packages included in the previous versions of Red Hat Enterprise Linux and Red Hat Ceph Storage are not included in the newer versions of these products. Follow the steps below after upgrading from Red Hat Enterprise Linux 6 to 7 to work around this issue. Note that you have to enable all relevant repositories first.

1. Firstly, update the **python-flask** package by running the following command as root:

```
# yum update python-flask
```

2. Next, remove the following packages:

- ✦ **python-argparse**
- ✦ **libproxy-python**
- ✦ **python-iwlib**
- ✦ **libreport-compat**
- ✦ **libreport-plugin-kerneloops**
- ✦ **libreport-plugin-logger**
- ✦ **subversion-perl**
- ✦ **python-jinja2-26**

To do so, execute the following command as root:

```
# yum remove python-argparse libproxy-python python-iwlib \
    libreport-compat libreport-plugin-kerneloops \
    libreport-plugin-logger subversion-perl python-jinja2-26
```

3. Finally, update the system:

```
# yum update
```

(BZ#1230679)

PGs creation hangs after creating a pool using incorrect values

An attempt to create a new erasure-coded pool using values that do not align with the OSD crush map causes placement groups (PGs) to remain in the "creating" state indefinitely. As a consequence, the Ceph cluster cannot achieve the **active+clean** state. To fix this problem, delete the erasure-encoded pool and associated crush ruleset, delete the profile that was used to create that pool, and use a new corrected erasure-encoding profile that aligns with the crush map.

(BZ#1231630)

IPv6 addresses are not supported destinations for radosgw-agent

The Ceph Object Gateway Sync Agent (**radosgw-agent**) does not support IPv6 addresses when specifying a destination. To work around this issue, specify a host name with an associated IPv6 address instead of the IPv6 address itself. (BZ#1232036)

Ceph Block Device sometimes fails in VM

When the writeback process is blocked by I/O errors, a Ceph Block Device terminates unexpectedly after a force shutdown in the Virtual Manager (VM). This issue is specific to Ubuntu only.

(BZ#1250042)

Upstart cannot stop or restart the initial "ceph-mon" process on Ubuntu

When adding a new monitor on Ubuntu, either manually or by using the **ceph-deploy** utility, the initial **ceph-mon** process cannot be stopped or restarted using the Upstart init system. To work around this issue, use the **kill** utility or reboot the system to stop the **ceph-mon** process. Then, it is possible to restart the process using Upstart as expected. (BZ#1255497)

Missing Calamari graphs on Ubuntu nodes

After installing a new Ceph cluster on Ubuntu and initializing the Calamari GUI server, Calamari graphs can be missing. The graphs are missing on any node connected to Calamari after the **sudo calamari-ctl initialize** command was run. To work around this issue, run **sudo calamari-ctl initialize** after connecting additional Ceph cluster nodes to Calamari. This command can be run multiple times without any issues. (BZ#1273943)

Removing manually added monitors by using "ceph-deploy" fails

When attempting to remove a manually added monitor host by using the **ceph-deploy mon destroy** command, the command fails with the following error:

```
UnboundLocalError: local variable 'status_args' referenced before
```

assignment"

The monitor is removed despite the error, however, **ceph-deploy** fails to remove the monitor's configuration directory located in the `/var/lib/ceph/mon/` directory. To work around this issue, remove the monitor's directory manually. ([BZ#1278524](#))

CHAPTER 4. FIXED BUGS

Bug ID	Component	Summary
1261606	Vulnerability	CVE-2015-5245 Ceph: RGW returns requested bucket name raw in Bucket response header
1231203	Installer	1.3.0: "mon add <mon2>" failed to complete in 300 sec. After interrupting the command(ctrl-c), the ceph commands time out.
1213026	Installer	Install: Identify and package all foreman dependencies
1213086	Installer	install: package upstream mariner-installer into rhcs-installer
1220961	Distribution	New component: ceph-puppet-modules
1254814	Build	RHCS is compiled --without-libatomic-ops
1224877	Calamari	calamari-crush-location causes OSD start failure when sudo is not present
1277393	Calamari	One of the calamari Rest API for OSD results in URL not found.
1224174	Ceph	Flag of rbd image info is showing as "flags: object map invalid"
1219040	Ceph	[MON]: Monitor crash after doing so many crush map edits
1207362	Ceph	RBD: Ltng tracepoints
1238415	Distribution	rebase ceph to 0.94.3

Bug ID	Component	Summary
1223532	Ceph	PG::find_best_info exclude incomplete for max_last_epoch_started_found
1223941	Ceph	OSD crash in release_op_ctx_locks with rgw and pool snaps
1265973	Ceph	After an upgrade from 1.1 to 1.3 through 1.2.3, OSD process is crashing.
1228842	Installer	ceph-deploy rgw create should echo the port number (e.g., 7480)
1235845	Ceph	rgw: usage reporting with civetweb doesn't reflect partial download
1253766	Ceph	[0.94.1-17.el7cp] AssertionError: 403 != 200 s3tests.functional.test_s3.test_object_copy_canned_acl
1258618	Ceph	rgw: COPYING an old object onto itself produces a truncated object
1258621	Ceph	RGW returns requested bucket name raw in "Bucket" response header
1238521	Ceph	radosgw start-up script relies on requiretty being disabled and silently fails if it is not
1255830	Ceph	FAILED assert(m_seed < old_pg_num) in librbd when increasing placement groups
1210825	Ceph	RFE: OSD Suicide Timeout should be configurable
1200182	Installer	ceph-deploy osd create and zap disk should check for block device type based on cli syntax

Bug ID	Component	Summary
1254343	Ceph	[RFE] RGW : setting max number of buckets for user via ceph.conf option
1245663	Ceph	HTTP return code is not being logged by CivetWeb