



Red Hat Ceph Storage 1.3.2 Release Notes

Release notes for Red Hat Ceph Storage 1.3.2

Red Hat Ceph Storage Documentation
Team

Red Hat Ceph Storage 1.3.2 Release Notes

Release notes for Red Hat Ceph Storage 1.3.2

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.

Table of Contents

CHAPTER 1. INTRODUCTION	3
CHAPTER 2. ACKNOWLEDGMENTS	4
CHAPTER 3. MAJOR UPDATES	5
CHAPTER 4. KNOWN ISSUES	7
CHAPTER 5. TECHNOLOGY PREVIEWS	10
CHAPTER 6. SOURCES	11

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 1.3.2 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.

Support for SELinux has been added

With this release, the SELinux policy for Red Hat Ceph Storage has been added. SELinux provides another security layer by enforcing Mandatory Access Control (MAC) mechanism over all processes. To learn more about SELinux, see the [SELinux User's and Administrator's Guide](#) for Red Hat Enterprise Linux 7.

SELinux support for Ceph is not enabled by default. To use it, install the **ceph-selinux** package. For detailed information about this process, see the [SELinux](#) section in the [Red Hat Ceph Storage Installation Guide for Red Hat Enterprise Linux](#).



Note

All Ceph daemons will be down for the time the **ceph-selinux** package is being installed. Therefore, your cluster node will not be able to serve any data at this point. This operation is necessary in order to update the metadata of the files located on the underlying file system and to make Ceph daemons run with the correct context. This operation may take several minutes depending on the size and speed of the underlying storage.

Package caching for Ubuntu is now supported

With this release, a caching server can be set up to provide Red Hat Ceph Storage repositories for offline Ceph clusters. See the [Package Caching for Red Hat Ceph Storage on Ubuntu](#) article on the Red Hat Customer Portal to learn more.

A new "ceph osd crush tree" command has been added

The CRUSH map contains a list of buckets for aggregating the devices into physical locations. With this update, a new **ceph osd crush tree** command has been added to Red Hat Ceph Storage. The command prints CRUSH buckets and items in a tree view. As a result, it is now easier to analyze the CRUSH map to determine a list of OSD daemons in a particular bucket.

TCMalloc thread cache is now configurable

With Red Hat Ceph Storage 1.3.2, support for modifying the size of the **TCMalloc** thread cache has been added. Increasing the thread cache size significantly improves Ceph cluster performance.

To set the thread cache size, edit the value of the **TCMALLOC_MAX_TOTAL_THREAD_CACHE_BYTES** parameter in the Ceph system configuration file, that is **/etc/sysconfig/ceph** for Red Hat Enterprise Linux and **/etc/default/ceph** for Ubuntu.

In addition, the default value of **TCMALLOC_MAX_TOTAL_THREAD_CACHE_BYTES** has been changed from 32 MB to 128 MB.

Red Hat Satellite 5 and Red Hat Ceph Storage integration

Red Hat Ceph Storage nodes can be connected to the Red Hat Satellite 5 Server. The server then

hosts package repositories and provides system updates.

Once you register your Ceph nodes with the Satellite 5 server, you can deliver upgrades to the Ceph cluster without allowing a direct connection to the Internet, as well as search and view errata applicable to the cluster nodes.

To learn more, see the [How to Register Ceph with Satellite 5](#) article on the Red Hat Customer Portal.

CHAPTER 4. KNOWN ISSUES

This section describes known issues found in this release.

The "--orphan-stale-secs" option is not documented

The `radosgw-admin(8)` manual page does not include the description of the `--orphan-stale-secs` option of the `orphans find` command. The option sets the number of seconds to wait before declaring an object to be an orphan. ([BZ#1301706](#))

Missing password field in Calamari

An attempt to create a new user in the Calamari UI fails with the `password field is required` error because the password field is missing from the Calamari UI. To work around this issue, use the Raw data form and create the new user by using the JSON format, for example:

```
{
  "username": "user",
  "email": "user@example.com",
  "password": "password"
}
```

([BZ#1305478](#))

"ice_setup" returns various errors when installing Calamari on Ubuntu

When running the `ice_setup` utility on Ubuntu, the utility returns various warnings, similar to the ones in the output below:

```
apache2_invoke: Enable configuration javascript-common
invoke-rc.d: initscript apache2, action "reload" failed.
apache2_invoke: Enable module wsgi
Adding user postgres to group ssl-cert
Creating config file /etc/logrotate.d/postgresql-common with new
version
Building PostgreSQL dictionaries from installed myspell/hunspell
packages...
Removing obsolete dictionary files:
 * No PostgreSQL clusters exist; see "man pg_createcluster"
ERROR: Module version does not exist!
salt-master: no process found
```

These warnings do not affect the installation process of Calamari and can be safely ignored. ([BZ#1305133](#))

Upstart cannot stop or restart the initialceph-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

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

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](#))

Various race conditions occur when using "ceph-disk"

The process of preparing and activating disks using the **ceph-disk** utility involves also usage of the **udev** device manager and the **systemd** service manager. Consequently, various race conditions can occur causing the following problems:

- ✦ Mount points for OSDs are duplicated
- ✦ **ceph-disk** fails to activate a device even though the device has been successfully prepared with the **ceph-disk prepare** command
- ✦ Some OSDs are not activated at boot time

To work around these issues, perform the steps below:

1. Manually remove the **udev** rules by running the following command as **root**:

```
# rm /usr/lib/udev/rules.d/95-ceph-osd.rules
```

2. Prepare the disks:

```
$ ceph-disk prepare
```

3. Add the "ceph-disk activate-all" string to the **/etc/rc.local** file. Run the following command as **root**:

```
# echo "ceph-disk activate-all" | tee -a /etc/rc.local
```

4. Reboot the system or activate the disks by running the following command as **root**:

```
# ceph-disk activate-all
```

([BZ#1300703](#), [BZ#1300617](#))

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

In the Calamari web UI 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](#))

The "salt '*' state.highstate" command fails to restart the "diamond" service

The `salt '*' state.highstate` command fails to restart the `diamond` service after installation of Red Hat Ceph Storage because the command cannot load the `diamond.service` unit file. As a consequence, the Calamari web UI does not show any data for the graphs in the **IOPS** and **Usage** sections of the Calamari dashboard. To work around this issue, restart `diamond` on each node by running the following command as **root**:

```
# /etc/init.d/diamond restart
```

Then run `salt '*' state.highstate` again:

```
# salt '*' state.highstate
```

([BZ#1310829](#))

CHAPTER 5. TECHNOLOGY PREVIEWS

This section contains a list of product features that are not fully supported yet.

Ceph Docker image added as a Technology Preview

With this release, Red Hat Ceph Storage is available as a Docker Container Image. The **rhceph-rhe17** image is now ready to use on the Red Hat registry. For detailed information how to bootstrap a containerized Ceph cluster see the [Deploying Red Hat Ceph Storage as a Docker Container \(Technology Preview\)](#) article on the [Customer Portal](#).

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

CHAPTER 6. 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: <ftp://ftp.redhat.com/redhat/rhceph-ubuntu/1.3.2>