



Red Hat Ceph Storage 5.3

Release Notes

Release notes for Red Hat Ceph Storage 5.3

Red Hat Ceph Storage 5.3 Release Notes

Release notes for Red Hat Ceph Storage 5.3

Legal Notice

Copyright © 2024 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 describe the major features, enhancements, known issues, and bug fixes implemented for the Red Hat Ceph Storage 5 product release. This covers the latest updates from the previous release notes up to the current release notes.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	4
PROVIDING FEEDBACK ON RED HAT CEPH STORAGE DOCUMENTATION	5
CHAPTER 1. INTRODUCTION	6
CHAPTER 2. ACKNOWLEDGMENTS	7
CHAPTER 3. NEW FEATURES	8
3.1. THE CEPHADM UTILITY	8
3.2. CEPH DASHBOARD	8
3.3. CEPH FILE SYSTEM	8
3.4. CEPH OBJECT GATEWAY	8
3.5. MULTI-SITE CEPH OBJECT GATEWAY	9
3.6. RADOS	10
3.7. RADOS BLOCK DEVICES (RBD)	10
CHAPTER 4. DEPRECATED FUNCTIONALITY	11
CHAPTER 5. REMOVED FUNCTIONALITY	12
CHAPTER 6. BUG FIXES	13
6.1. THE CEPHADM UTILITY	13
6.2. CEPH FILE SYSTEM	13
6.3. CEPH MANAGER PLUGINS	13
6.4. THE CEPH VOLUME UTILITY	14
6.5. CEPH OBJECT GATEWAY	14
6.6. MULTI-SITE CEPH OBJECT GATEWAY	16
6.7. RADOS	17
6.8. RADOS BLOCK DEVICES (RBD)	20
6.9. RBD MIRRORING	20
6.10. ISCSI GATEWAY	20
6.11. THE CEPH ANSIBLE UTILITY	21
CHAPTER 7. KNOWN ISSUES	22
7.1. THE CEPHADM UTILITY	22
CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES	23
8.1. RED HAT CEPH STORAGE 5.3Z6	23
8.1.1. Enhancements	23
8.1.1.1. Ceph Object Gateway	23
8.1.1.2. Ceph File System	23
8.1.1.3. Ceph Manager plugins	23
8.1.1.4. RADOS	24
8.1.2. Known issues	24
8.1.2.1. Ceph Dashboard	24
8.2. RED HAT CEPH STORAGE 5.3Z5	24
8.3. RED HAT CEPH STORAGE 5.3Z4	24
8.3.1. Known issues	25
8.3.1.1. Multi-site Ceph Object Gateway	25
8.4. RED HAT CEPH STORAGE 5.3Z3	25
8.4.1. Enhancements	25
8.4.1.1. The Cephadm utility	25
8.4.2. Known issues	26

8.5. RED HAT CEPH STORAGE 5.3Z2	26
8.5.1. Enhancements	26
8.5.1.1. Ceph File System	26
8.5.1.2. Ceph Object Gateway	27
8.5.2. Known issues	27
8.5.2.1. The Cephadm utility	27
8.6. RED HAT CEPH STORAGE 5.3Z1	27
8.6.1. Enhancements	27
8.6.1.1. The Cephadm utility	27
8.6.1.2. Ceph Dashboard	28
8.6.1.3. Ceph File System	28
8.6.1.4. Ceph Object Gateway	28
8.6.1.5. Multi-site Ceph Object Gateway	29
8.6.1.6. RADOS	29
8.6.1.7. RADOS Block Devices (RBD)	30
8.6.2. Known issues	30
8.6.2.1. The Cephadm utility	30
8.6.2.2. Ceph Object Gateway	30
CHAPTER 9. SOURCES	31

MAKING OPEN SOURCE MORE INCLUSIVE

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](#).

PROVIDING FEEDBACK ON RED HAT CEPH STORAGE DOCUMENTATION

We appreciate your input on our documentation. Please let us know how we could make it better. To do so, create a Bugzilla ticket:

+ . Go to the [Bugzilla](#) website. . In the Component drop-down, select **Documentation**. . In the Sub-Component drop-down, select the appropriate sub-component. . Select the appropriate version of the document. . Fill in the **Summary** and **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation. . Optional: Add an attachment, if any. . Click **Submit Bug**.

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-us/red_hat_ceph_storage/5.

CHAPTER 2. ACKNOWLEDGMENTS

The Ceph Storage 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 CEPHADM UTILITY

cephadm automatically updates the dashboard Grafana password if it is set in the Grafana service spec

Previously, users would have to manually set the Grafana password after applying the specification.

With this enhancement, if **initial_admin_password** is set in an applied Grafana specification, **cephadm** automatically updates the dashboard Grafana password, which is equivalent to running **ceph dashboard set-grafana-api-password** command, to streamline the process of fully setting up Grafana. Users no longer have to manually set the dashboard Grafana password after applying a specification that includes the password.

OSDs automatically update their Ceph configuration files with the new mon locations

With this enhancement, whenever a **monmap** change is detected, **cephadm** automatically updates the Ceph configuration files for each OSD with the new **mon** locations.



NOTE

This enhancement may take some time to update on all OSDs if you have a lot of OSDs.

3.2. CEPH DASHBOARD

The Block Device images table is paginated

With this enhancement, the Block Device images table is paginated to use with 10000+ image storage clusters as retrieving information for a block device image is expensive.

Newly added cross_origin_url option allows cross origin resource sharing

Previously, IBM developers faced issues with their storage insights product when they tried to ping the REST API using their front-end because of the tight Cross Origin Resource Sharing (CORS) policies set up in Red Hat's REST API.

With this enhancement, CORS is allowed by adding the **cross_origin_url** option that can be set to a particular URL - **ceph config set mgr mgr/dashboard/cross_origin_url localhost** and the REST API allows communication with only that URL.

3.3. CEPH FILE SYSTEM

Users can store arbitrary metadata of CephFS subvolume snapshots

With this enhancement, Ceph File System (CephFS) volume users can store arbitrary metadata in the form of key-value pairs for CephFS subvolume snapshots with a set of command-line interface (CLI) commands.

3.4. CEPH OBJECT GATEWAY

STS `max_session_duration` for a role can now be updated

With this enhancement, the STS `max_session_duration` for a role can be updated using the `radosgw-admin` command-line interface.

ListBucket S3 operation now generates JSON output

With this enhancement, on customers' request to facilitate integrations, the `ListBucket` S3 operation generates JSON-formatted output, instead of the default XML, if the request contains an `Accept: application/json` header.

The option to enable TCP keepalive managed by `libcurl` is added

With this enhancement, the option to enable TCP keepalive on the HTTP client sockets managed by `libcurl` is added to make sync and other operations initiated by Ceph Object Gateway more resilient to network instability. This does not apply to connections received by the HTTP frontend, but only to HTTP requests sent by the Ceph Object Gateway, such as Keystone for authentication, sync requests from multi-site, and requests to key management servers for SSE.

Result code 2002 of `radosgw-admin` commands is explicitly translated to 2

Previously, a change in the S3 error translation of internal `NoSuchBucket` result inadvertently changed the error code from the `radosgw-admin bucket stats` command, causing the programs checking the shell result code of those `radosgw-admin` commands to see a different result code.

With this enhancement, the result code 2002 is explicitly translated to 2 and users can see the original behaviour.

You can now use bucket policies with useful errors

Bucket policies were difficult to use since the error indication was wrong. Additionally, silently dropping principals would cause problems during the upgrade. With this update, useful errors from policy parser and a flag to reject invalid principals with `rgw policy reject invalid principals=true` parameter is introduced.

3.5. MULTI-SITE CEPH OBJECT GATEWAY

The `bucket sync run` command provides more details

With this enhancement, user-friendly progress reports on the `bucket sync run` command are added to provide users easier visibility into the progress of the operation. When the user runs the `radosgw-admin bucket sync run` command with `--extra-info` flag, users get a message for the start of generation sync and also for each object that is synced.



WARNING

It is not recommended to use the `bucket sync run` command without contacting Red Hat support.

Multi-site configuration supports dynamic bucket index resharding

Previously, only manual resharding of the buckets for multi-site configurations was supported.

With this enhancement, dynamic bucket resharding is supported in multi-site configurations. Once the storage clusters are upgraded, enable the **resharding** feature, zone level, and zone group. You can either manually reshard the buckets with **radogw-admin bucket reshard command** or automatically reshard them with dynamic resharding, independently of other zones in the storage cluster.

Users can now reshard bucket index dynamically with multi-site archive zones

With this enhancement, multi-site archive zone bucket index can be resharded dynamically when dynamic resharding is enabled for that zone.

3.6. RADOS

Low-level log messages are introduced to warn user about hitting throttle limits

Previously, there was a lack of low-level logging indication that throttle limits were hit, causing these occurrences to incorrectly have the appearance of a networking issue.

With this enhancement, the introduction of low-level log messages makes it much clearer that the throttle limits are hit.

3.7. RADOS BLOCK DEVICES (RBD)

Cloned images can now be encrypted with their own encryption format and passphrase

With this enhancement, layered client-side encryption is now supported that enables each cloned image to be encrypted with its own encryption format and passphrase, potentially different from that of the parent image. The efficient copy-on-write semantics intrinsic to unformatted regular cloned images are retained.

CHAPTER 4. 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.



IMPORTANT

Deprecated functionality continues to be supported until the end of life of Red Hat Ceph Storage 5. Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments. For the most recent list of deprecated functionality within a particular major release, refer to the latest version of release documentation.

NFS support for CephFS is now deprecated

NFS support for CephFS is now deprecated in favor of upcoming NFS availability in OpenShift Data Foundation. Red Hat Ceph Storage support for NFS in OpenStack Manila is not affected. Deprecated functionality will receive only bug fixes for the lifetime of the current release, and may be removed in future releases. Relevant documentation around this technology is identified as "Limited Availability".

iSCSi support is now deprecated

iSCSi support is now deprecated in favor of future NVMeoF support. Deprecated functionality will receive only bug fixes for the lifetime of the current release, and may be removed in future releases. Relevant documentation around this technology is identified as "Limited Availability".

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*.

The `min_compat_client` parameter for Ceph File System (CephFS) is now deprecated

The **min_compat_client** parameter is deprecated for Red Hat Ceph Storage 5.0 and new client features are added for setting-up the Ceph File Systems (CephFS). For details, see the [Client features](#) section in the *Red Hat Ceph Storage File System Guide*.

The snapshot of Ceph File System subvolume group is now deprecated

The snapshot feature of Ceph File System (CephFS) subvolume group is deprecated for Red Hat Ceph Storage 5.0. The existing snapshots can be listed and deleted, whenever needed. For details, see the [Listing snapshots of a file system subvolume group](#) and [Removing snapshots of a file system subvolume group](#) sections in the *Red Hat Ceph Storage Ceph File System guide*.

The Cockpit Ceph Installer is now deprecated

Installing a Red Hat Ceph Storage cluster 5 using Cockpit Ceph Installer is not supported. Use Cephadm to install a Red Hat Ceph Storage cluster. For details, see the [Red Hat Ceph Storage Installation guide](#).

CHAPTER 5. REMOVED FUNCTIONALITY

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

s3cmd RPM is unavailable in Ceph's Tools repository

The **s3cmd** RPM is no longer available in Ceph's Tools repository. Users can download the unsupported community packages from **PyPI** or **EPEL**.

CHAPTER 6. BUG FIXES

This section describes bugs with significant user impact, which 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 CEPHADM UTILITY

Users can upgrade to a local repo image without any issues

Previously, in `cephadm`, **docker.io** would be added to the start of the image name by default, if the image name was not a qualified domain name. Due to this, users were unable to upgrade to images on local repositories.

With this fix, care has been taken to identify the images to which **docker.io** is added by default. Users using a local repo image can upgrade to that image without encountering issues.

([BZ#2100553](#))

6.2. CEPH FILE SYSTEM

snap-schedules are no longer lost on restarts of Ceph Manager services

Previously, in-memory databases were not written to persistent storage on every change to the schedule. This caused **snap-schedules** to get lost on restart of Ceph Manager services.

With this fix, the in-memory databases are dumped into persistent storage on every change or addition to the **snap-schedules**. Retention now continues to work across restarts of Ceph Manager services.

([BZ#2102934](#))

The standby-replay Metadata Server daemon is no longer unexpectedly removed

Previously, the Ceph Monitor would remove a standby-replay Metadata Server (MDS) daemon from the MDS map under certain conditions. This caused the standby-replay MDS daemon to be removed from the Metadata Server cluster, which generated cluster warnings.

With this fix, the logic used in Ceph Monitors during the consideration of removal of an MDS daemon from the MDS map, now includes information about the standby-replay MDS daemons holding a rank. This ensures that the standby-replay MDS daemons are no longer unexpectedly removed from the MDS cluster.

([BZ#2130116](#))

6.3. CEPH MANAGER PLUGINS

Ceph Manager Alert emails are not tagged as spam anymore

Previously, emails sent by the Ceph Manager Alerts module did not have the "Message-Id" and "Date:headers". This increased the chances of flagging the emails as spam.

With this fix, both the headers are added to the emails sent by Ceph Manager Alerts module and the messages are not flagged as spam.

([BZ#2064481](#))

6.4. THE CEPH VOLUME UTILITY

The volume list remains empty when no **ceph-osd** container is found and **cephvolumescan** actor no longer fails

Previously, if Ceph containers ran collocated with other containers without a **ceph-osd** container present among them, the process would try to retrieve the volume list from one non-Ceph container which would not work. Due to this, **cephvolumescan** actor would fail and the upgrade would not complete.

With this fix, if no **ceph-osd** container is found, the volume list will remain empty and the **cephvolumescan** actor does not fail.

([BZ#2141393](#))

Ceph OSD deployment no longer fails when **ceph-volume** treats multiple devices.

Previously, **ceph-volume** computed wrong sizes when there were multiple devices to treat, resulting in failure to deploy OSDs.

With this fix, **ceph-volume** computes the correct size when multiple devices are to be treated and deployment of OSDs work as expected.

([BZ#2119774](#))

6.5. CEPH OBJECT GATEWAY

Users can now set up Kafka connectivity with SASL in a non-TLS environment

Previously, due to a failure in configuring the TLS certificate for Ceph Object Gateway, it was not possible to configure Kafka topic with SASL (user and password).

With this fix, a new configuration parameter, **rgw_allow_notification_secrets_in_cleartext**, is added. Users can now set up Kafka connectivity with SASL in a non-TLS environment.

([BZ#2014330](#))

Internal handling of tokens is fixed

Previously, internal handling of tokens in the refresh path of Java-based client authentication provider jar for AWS SDK for Java and Hadoop S3A Connector, would not deal correctly with the large tokens, resulting in improper processing of some tokens and preventing the renewal of client tokens.

With this fix, the internal token handling is fixed and it works as expected.

([BZ#2055137](#))

The object version access is corrected preventing object lock violation

Previously, inadvertent slicing of version information would occur in some call paths, causing any object version protected by object lock to be deleted contrary to policy.

With this fix, the object version access is corrected, thereby preventing object lock violation.

([BZ#2108394](#))

Ceph Object Gateway no longer crashes with malformed URLs

Previously, a refactoring abstraction replaced a bucket value with a pointer to a bucket value that was not always initialized. This caused malformed URLs corresponding to bucket operations on no buckets resulting in Ceph Object Gateway crashing.

With this fix, a check on the pointer has been implemented into the call path and Ceph Object Gateway returns a permission error, rather than crashing, if it is uninitialized.

([BZ#2109256](#))

The code that parses dates **z-amz-date** format is changed

Previously, the standard format for **x-amz-date** was changed which caused issues, since the new software uses the new date format. The new software built with the latest **go** libraries would not talk to the Ceph Object Gateway.

With this fix, the code in the Ceph Object Gateway that parses dates in **x-amz-date** format is changed to also accept the new date format.

([BZ#2109675](#))

New logic in processing of lifecycle shards prevents stalling due to deleted buckets

Previously, changes were made to cause lifecycle processing to continuously cycle across days, that is, to not restart from the beginning of the list of eligible buckets each day. However, the changes contained a bug which could stall processing of lifecycle shards that contained deleted buckets, causing the processing of lifecycle shards to stall.

With this fix, a logic is introduced to skip over the deleted buckets, due to which the processing no longer stalls.

([BZ#2118295](#))

Header processing no longer causes sporadic swift-protocol authentication failures

Previously, a combination of incorrect HTTP header processing and timestamp handling logic would either cause an invalid Keystone admin token to be used for operations, or non-renewal of Keystone's admin token as required. Due to this, sporadic swift-protocol authentication failures would occur.

With this fix, header processing is corrected and new diagnostics are added. The logic now works as expected.

([BZ#2123335](#))

Warnings are no longer logged in inappropriate circumstances

Previously, an inverted logic would occasionally report an incorrect warning - *unable to find head object*, causing the warning to be logged when it was not applicable in a Ceph Object Gateway configuration.

With this fix, the corrected logic no longer logs the warning in inappropriate circumstances.

([BZ#2126787](#))

PUT object operation writes to the correct bucket index shards

Previously, due to a race condition, a PUT object operation would rarely write to a former bucket index shard. This caused the former bucket index shard to be recreated, and the object would not appear in the proper bucket index. Therefore, the object would not be listed when the bucket was listed.

With this fix, care is taken to prevent various operations from creating bucket index shards and recover when the race condition is encountered. PUT object operations now always write to the correct bucket index shards.

([BZ#2145022](#))

6.6. MULTI-SITE CEPH OBJECT GATEWAY

Suspending bucket versioning in the primary zone no longer suspends bucket versioning in the archive zone

Previously, if bucket versioning was suspended in the primary zone, bucket versioning in the archive zone would also be suspended.

With this fix, archive zone versioning is always enabled irrespective of bucket versioning changes on other zones. Bucket versioning in the archive zone no longer gets suspended.

([BZ#1957088](#))

The `radosgw-admin sync status` command in multi-site replication works as expected

Previously, in a multisite replication, if one or more participating Ceph Object Gateway nodes are down, you would (5) *Input/output error* output when running the `radosgw-admin sync status` command. This status should get resolved after all the Ceph Object Gateway nodes are back online.

With this update, the `radosgw-admin sync status` command does not get stuck and works as expected.

([BZ#1749627](#))

Processes trimming retired bucket index entries no longer cause `radosgw` instance to crash

Previously, under some circumstances, processes trimming retired bucket index entries could access an uninitialized pointer variable resulting in the `radosgw` instance to crash.

With this fix, code is initialized immediately before use and the `radosgw` instance no longer crashes.

([BZ#2139258](#))

Bucket sync run is given control logic to sync all objects

Previously, to support dynamic bucket resharding on multisite clusters, a singular bucket index log was replaced with multiple bucket index log generations. But, due to how bucket sync run was implemented, only the oldest outstanding generation would be sync run.

With this fix, bucket sync run is given control logic which enables it to run the sync from oldest outstanding to current and all objects are now synced as expected.

([BZ#2066453](#))

Per-bucket replication logical error fix executes policies correctly

Previously, an internal logic error caused failures in per-bucket replication, due to which per-bucket replication policies did not work in some circumstances.

With this fix, the logic error responsible for confusing the source and destination bucket information is corrected and the policies execute correctly.

([BZ#2108886](#))

Variable access no longer causes undefined program behavior

Previously, a coverity scan would identify two cases, where variables could be used after a move, potentially causing an undefined program behavior to occur.

With this fix, variable access is fixed and the potential fault can no longer occur.

([BZ#2123423](#))

Requests with a tenant but no bucket no longer cause a crash

Previously, an upstream refactoring replaced uninitialized bucket data fields with uninitialized pointers. Due to this, any bucket request containing a URL referencing no valid bucket caused crashes.

With this fix, requests that access the bucket but do not specify a valid bucket are denied, resulting in an error instead of a crash.

([BZ#2139422](#))

6.7. RADOS

Performing a DR test with two sites stretch cluster no longer causes Ceph to become unresponsive

Previously, when performing a DR test with two sites stretch-cluster, removing and adding new monitors to the cluster would cause an incorrect rank in **ConnectionTracker** class. Due to this, the monitor would fail to identify itself in the **peer_tracker** copy and would never update its correct field, causing a deadlock in the election process which would lead to Ceph becoming unresponsive.

With this fix, the following corrections are made:

- Added an assert in the function **notify_rank_removed()**, to compare the expected rank provided by the **Monmap** against the rank that is manually adjusted as a sanity check.
- Clear the variable **removed_ranks** from every **Monmap** update.
- Added an action to manually reset **peer_tracker.rank** when executing the command - **ceph connection scores reset** for each monitor. The **peer_tracker.rank** matches the current rank of the monitor.
- Added functions in the **Elector** and **ConnectionTracker** classes to check for clean **peer_tracker** when upgrading the monitors, including booting up. If found unclean, **peer_tracker** is cleared.
- In Red Hat Ceph Storage, the user can choose to manually remove a monitor rank before shutting down the monitor, causing inconsistency in **Monmap**. Therefore, in **Monitor::notify_new_monmap()** we prevent the function from removing our rank or ranks that don't exist in **Monmap**.

The cluster now works as expected and there is no unwarranted downtime. The cluster no longer becomes unresponsive when performing a DR test with two sites stretch-cluster.

([BZ#2142674](#))

Rank is removed from the **live_pinging** and **dead_pinging** set to mitigate the inconsistent connectivity score issue

Previously, when removing two monitors consecutively, if the rank size is equal to Paxos's size, the

monitor would face a condition and would not remove rank from the **dead_pinging** set. Due to this, the rank remained in the **dead_pinging** set which would cause problems, such as inconsistent connectivity score when the stretch-cluster mode was enabled.

With this fix, a case is added where the highest ranked monitor is removed, that is, when the rank is equal to Paxos's size, remove the rank from the **live_pinging** and **dead_pinging** set. The monitor stays healthy with a clean **live_pinging** and **dead_pinging** set.

([BZ#2142174](#))

The Prometheus metrics now reflect the correct Ceph version for all Ceph Monitors whenever requested

Previously, the Prometheus metrics reported mismatched Ceph versions for Ceph Monitors when the monitor was upgraded. As a result, the active Ceph Manager daemon needed to be restarted to resolve this inconsistency.

With this fix, the Ceph Monitors explicitly send metadata update requests with **mon** metadata to **mgr** when MON election is over.

([BZ#2008524](#))

The ceph daemon heap status command shows the heap status

Previously, due to a failure to get heap information through the **ceph daemon** command, the **ceph daemon heap stats** command would return empty output instead of returning current heap usage for a Ceph daemon. This was because **ceph::osd_cmds::heap()** was confusing the **stderr** and **stdout** concept which caused the difference in output.

With this fix, the **ceph daemon heap stats** command returns heap usage information for a Ceph daemon similar to what we get using the **ceph tell** command.

([BZ#2119100](#))

Ceph Monitors no longer crash when using **ceph orch apply mon <num>** command

Previously, when the command **ceph orch apply mon <num>** was used to decrease monitors in a cluster, the monitors were removed before shutting down in **ceph-adm** causing the monitors to crash.

With this fix, a sanity check is added to all code paths that check whether the peer rank is more than or equal to the size of the ranks from the monitor map. If the condition is satisfied, then skip certain operations that lead to the monitor crashing. The peer rank eventually resolves itself in the next version of the monitor map. The monitors no longer crash when removed from the monitor map before shutting down.

([BZ#2142141](#))

End-user can now see the scrub or deep-scrub starts message from the Ceph cluster log

Previously, due to the scrub or deep-scrub starts message missing in the Ceph cluster log, the end-user would fail to know if the PG scrubbing had started for a PG from the Ceph cluster log.

With this fix, the scrub or deep-scrub **starts** message is reintroduced. The Ceph cluster log now shows the message for a PG, whenever it goes for a scrubbing or deep-scrubbing process.

([BZ#2091773](#))

No assertion during the Ceph Manager failover

Previously, when activating the Ceph Manager, it would receive several **service_map** versions sent by the previously active manager. This incorrect check in code would cause assertion failure when the newly activated manager received a map with a higher version sent by the previously active manager.

With this fix, the check in the manager that deals with the initial service map is relaxed and there is no assertion during the Ceph Manager failover.

([BZ#2095062](#))

Users can remove cloned objects after upgrading a cluster

Previously, after upgrading a cluster from Red Hat Ceph Storage 4 to Red Hat Ceph Storage 5, removing snapshots of objects created in earlier versions would leave clones, which could not be removed. This was because the **SnapMapper** keys were wrongly converted.

With this fix, **SnapMapper**'s legacy conversation is updated to match the new key format. The cloned objects in earlier versions of Ceph can now be easily removed after an upgrade.

([BZ#2107405](#))

RocksDB error does not occur for small writes

BlueStore employs a strategy of deferring small writes for HDDs and stores data in RocksDB. Cleaning deferred data from RocksDB is a background process which is not synchronized with BlueFS.

With this fix, deferred replay no longer overwrites BlueFS data and some RocksDB errors do not occur, such as:

- **osd_superblock** corruption.
- CURRENT does not end with newline.
- **.sst** files checksum error.



NOTE

Do not write deferred data as the write location might either contain a proper object or be empty. It is not possible to corrupt object data this way. BlueFS is the only entity that can allocate this space.

([BZ#2109886](#))

Corrupted dups entries of a PG Log can be removed by off-line and on-line trimming

Previously, trimming of PG log dups entries could be prevented during the low-level PG split operation, which is used by the PG autoscaler with far higher frequency than by a human operator. Stalling the trimming of dups resulted in significant memory growth of PG log, leading to OSD crashes as it ran out of memory. Restarting an OSD did not solve the problem as the PG log is stored on disk and reloaded to RAM on startup.

With this fix, both off-line, using the **ceph-objectstore-tool** command, and on-line, within OSD, trimming can remove corrupted dups entries of a PG log that jammed the on-line trimming machinery and were responsible for the memory growth. A debug improvement is implemented that prints the number of dups entries to the OSD's log to help future investigations.

([BZ#2119853](#))

6.8. RADOS BLOCK DEVICES (RBD)

rbd info command no longer fails if executed when the image is being flattened

Previously, due to an implementation defect, **rbd info** command would fail, although rarely, if run when the image was being flattened. This caused a transient *No such file or directory* error to occur, although, upon rerun, the command always succeeded.

With this fix, the implementation defect is fixed and **rbd info** command no longer fails even if executed when the image is being flattened.

([BZ#1989527](#))

Removing a pool with pending Block Device tasks no longer causes all the tasks to hang

Previously, due to an implementation defect, removing a pool with pending Block Device tasks caused all Block Device tasks, including other pools, to hang. To resume hung Block Device tasks, the administrator had to restart the **ceph-mgr** daemon.

With this fix, the implementation defect is fixed and removing a pool with pending RBD tasks no longer causes any hangs. Block Device tasks for the removed pool are cleaned up. Block Device tasks for other pools continue executing uninterrupted.

([BZ#2150968](#))

6.9. RBD MIRRORING

The image replayer shuts down as expected

Previously, due to an implementation defect, a request to shut down a particular image replayer would cause the **rbd-mirror** daemon to hang indefinitely, especially in cases where the daemon was blocklisted on the remote storage cluster.

With this fix, the implementation defect is fixed and a request to shut down a particular image replayer no longer causes the **rbd-mirror** daemon to hang and the image replayer shuts down as expected.

([BZ#2086471](#))

The **rbd mirror pool peer bootstrap create** command guarantees correct monitor addresses in the bootstrap token

Previously, a bootstrap token generated with the **rbd mirror pool peer bootstrap create** command contained monitor addresses as specified by the **mon_host** option in the **ceph.conf** file. This was fragile and caused issues to users, such as causing confusion between V1 and V2 endpoints, specifying only one of them, grouping them incorrectly, and the like.

With this fix, the **rbd mirror pool peer bootstrap create** command is changed to extract monitor address from the cluster itself, guaranteeing the monitor addresses contained in a bootstrap token to be correct.

([BZ#2122130](#))

6.10. ISCSI GATEWAY

Upgrade from Red Hat Ceph Storage 4.x to 5.x with iSCSI works as expected

Previously, due to version conflict between some of the **ceph-iscsi** dependent libraries, upgrades from Red Hat Ceph Storage 4.x to 5.x would lead to a persistent HTTP 500 error.

With this fix, the versioning conflict is resolved and the upgrade works as expected. However, as a result of this fix, iSCSI REST API responses aren't pretty-printed.

([BZ#2121462](#))

6.11. THE CEPH ANSIBLE UTILITY

Upgrade workflow with Ceph Object Gateway configuration is fixed

Previously, whenever **set_radosgw_address.yml** was called from the dashboard playbook execution, the fact **is_rgw_instances_defined** was expected to be set if **rgw_instances** was defined in **group_vars/host_vars** by the user. Otherwise, the next task that sets the fact **rgw_instances** will be executed under the assumption that it wasn't user defined. This caused the upgrade workflow to break when deploying the Ceph Object Gateway multisite and Ceph Dashboard.

With this fix, **ceph-ansible** sets the parameter when **set_radosgw_address.yml** playbook is called from the dashboard playbook and the upgrade workflow works as expected.

([BZ#2117672](#))

The fact condition is updated to execute only on the Ceph Object Gateway nodes

Previously, due to **set_fact_radosgw_address_to_radosgw_address_block_ipv4** being executed on all nodes, including the ones where no Ceph Object Gateway network range was present, playbooks failed to work.

With this fix, the **when** condition is updated to execute the fact setting only on the Ceph Object Gateway nodes and now works as expected.

([BZ#2136551](#))

CHAPTER 7. KNOWN ISSUES

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

7.1. THE CEPHADM UTILITY

Adding or expanding iSCSI gateways in `gwcli` across the iSCSI daemons works as expected

Previously, due to iSCSI daemons not being reconfigured automatically when a trusted IP list was updated in the specification file, adding or expanding iSCSI gateways in `gwcli` would fail due to the `iscsi-gateway.cfg` not matching across the iSCSI daemons.

With this fix, you can expand the gateways and add it to the existing gateways with `gwcli` command. ([BZ#2099470](#))

`ceph orch ps` does not display a version for monitoring stack daemons

In `cephadm`, due to the version grabbing code currently being incompatible with the downstream monitoring stack containers, version grabbing fails for monitoring stack daemons, such as `node-exporter`, `prometheus`, and `alertmanager`.

As a workaround, if the user needs to find the version, the daemons' container names include the version.

([BZ#2125382](#))

CHAPTER 8. ASYNCHRONOUS ERRATA UPDATES

This section describes the bug fixes, known issues, and enhancements of the z-stream releases.

8.1. RED HAT CEPH STORAGE 5.3Z6

Red Hat Ceph Storage release 5.3z6 is now available. The bug fixes and security updates that are included in the update are listed in [RHSA-2024:0745](#) advisory.

8.1.1. Enhancements

8.1.1.1. Ceph Object Gateway

The **rgw-restore-bucket-index** experimental tool restores bucket indices for versioned and un-versioned buckets

With this enhancement, you can restore the bucket indices for versioned buckets with the **rgw-restore-bucket-index** experimental tool, in addition to its existing ability to work with un-versioned buckets.

[BZ#2224636](#)

Enhanced ordered bucket listing

Previously, in some cases, buckets with larger number of shards and several pseudo-subdirectories would take an unnecessarily long time to complete.

With this enhancement, such buckets perform an ordered bucket listing more quickly.

[BZ#2239433](#)

The **radosgw-admin bucket stats** command prints bucket versioning

With this enhancement, the **radosgw-admin bucket stats** command prints the versioning status for buckets as one of three values of **enabled**, **off**, or **suspended** since versioning can be enabled or disabled after creation.

[BZ#2240089](#)

8.1.1.2. Ceph File System

The MDS default balancer is now disabled by default

With this release, the MDS default balancer or the automatic dynamic subtree balancer is disabled by default. This prevents accidental subtree migrations, Subtree migrations can be expensive to undo when the operator increases the file system **max_mds** setting without planning subtree delegations, such as, with pinning.

[BZ#2255436](#)

8.1.1.3. Ceph Manager plugins

Each Ceph Manager module has a separate thread to run commands

Previously, there was one thread through which all the **ceph-mgr** module commands were run. If one of the module's commands were stuck, all the other module's commands would hang, waiting on the same thread.

With this update, one finisher thread for each Ceph Manager module is added. Each module has a separate thread for commands run. Even if one of the module's command hangs, the other modules are able to run.

[BZ#2234610](#)

8.1.1.4. RADOS

Improved protection against running BlueStore twice

Previously, advisory locking was used to protect against running BlueStore twice. This works well on baremetal deployments. However, when used on containers it would create unrelated inodes that targeted same **mknod b** block device. As a result, two containers might assume that they can have exclusive access which led to severe errors.

With this release, you can improve protection against running OSDs twice at the same time on one block device. You can reinforce advisory locking with `O_EXCL` open flag dedicated for block devices. It is no longer possible to open one BlueStore instance twice and the overwrite and corruption does not occur.

[BZ#2239455](#)

New reports available for sub-events for delayed operations

Previously, slow operations were marked as delayed but without a detailed description.

With this enhancement, you can view the detailed descriptions of delayed sub-events for operations.

[BZ#2240839](#)

8.1.2. Known issues

8.1.2.1. Ceph Dashboard

Some metrics are displayed as null leading to blank spaces in graphs

Some metrics on the Ceph dashboard are shown as null, which leads to blank space in the graphs since you do not initialize a metric until it has some value.

As a workaround, edit the Grafana panel in which the issue is present. From the *Edit* menu, click *Migrate* and select *Connect Nulls*. Choose *Always* and the issue is resolved.

[Bugzilla:2228128](#)

8.2. RED HAT CEPH STORAGE 5.3Z5

Red Hat Ceph Storage release 5.3z5 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:4760](#) advisory.

8.3. RED HAT CEPH STORAGE 5.3Z4

Red Hat Ceph Storage release 5.3z4 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:4213](#) advisory.

8.3.1. Known issues

8.3.1.1. Multi-site Ceph Object Gateway

md5 mismatch of replicated objects when testing Ceph Object gateway's server-side encryption in multi-site

Presently, a md5 mismatch of replicated objects is observed when testing Ceph Object gateway's server-side encryption in multi-site. The data corruption is specific to S3 multipart uploads with SSE encryption enabled. The corruption only affects the replicated copy. The original object remains intact.

Encryption of multipart uploads requires special handling around the part boundaries because each part is uploaded and encrypted separately. In multi-site, objects are encrypted, and multipart uploads are replicated as a single part. As a result, the replicated copy loses its knowledge about the original part boundaries required to decrypt the data correctly, which causes this corruption.

As a workaround, multi-site users should not use server-side encryption for multipart uploads. For more detailed information, see the KCS [Sever side encryption with RGW multisite configuration might lead to data corruption of multipart objects](#).

[Bugzilla:2214252](#)

8.4. RED HAT CEPH STORAGE 5.3Z3

Red Hat Ceph Storage release 5.3z3 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:3259](#) advisory.

8.4.1. Enhancements

8.4.1.1. The Cephadm utility

Users can now set `crush_device_class` in OSD specification

Previously, users would manually set the `crush_device_class` after the OSDs were made.

With this release, users can set the `crush_device_class` in an OSD specification, which gets `cephadm` to mark all OSDs created based on that specification as being that crush device class.

Syntax

```
service_type: osd
service_id: SERVICE_ID_OF_OSD
placement:
hosts:
- HOSTNAME_01
- HOSTNAME_01
crush_device_class: CRUSH_DEVICE_CLASS(SSD/HDD)
spec:
data_devices:
paths:
- DATA_DEVICES
```

```
db_devices:  
paths:  
- DB_DEVICES  
wal_devices:  
paths:  
- WAL_DEVICES
```

Users can now set retention time in Prometheus specification

Previously, setting the retention time required manually editing the `unit.run` file, and that would be overwritten whenever the Prometheus daemon was redeployed.

With this release, you can set the retention time in the Prometheus specification file as follows:

Example

```
service_type: prometheus  
placement:  
  count: 1  
spec:  
  retention_time: "1y"
```

In this example, the retention time is set to one year instead of the default 15 days.

8.4.2. Known issues

- **Documentation for users to manage Ceph File system snapshots on the Red Hat Ceph Storage Dashboard**
Details for this feature will be included in the next version of the *Red Hat Ceph Storage Dashboard Guide*.
- **Documentation for users to manage hosts on the Red Hat Ceph Storage Dashboard**
Details for this feature will be included in the next version of the *Red Hat Ceph Storage Dashboard Guide*.
- **Documentation for users to import RBD images instantaneously**
Details for the **rbd import** command will be included in the next version of the *Red Hat Ceph Storage Block Device Guide*.

8.5. RED HAT CEPH STORAGE 5.3Z2

Red Hat Ceph Storage release 5.3z1 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:1732](#) advisory.

8.5.1. Enhancements

8.5.1.1. Ceph File System

Client request counters are converted from `_u8` type to `_u32` type and the limit is set to 256 times

Previously, in multiple active MDSs cases, if a single request failed in the current MDS, the client would forward the request to another MDS. If no MDS could successfully handle the request, it would bounce infinitely between MDSs. The old `num_fwd/num_retry` counters are `_u8` type, which would overflow

after bouncing 256 times.

With this enhancement, the counters are converted from `_u8` type to `_u32` type and the limit for forwarding and retrying is set to 256 times. The client requests stop forwarding and retrying after 256 times and fails directly instead of infinitely forwarding and retrying.

8.5.1.2. Ceph Object Gateway

Administrators can now reuse output from `rados ls` to complete bucket reindexing quickly

Previously, running `rados ls` command for each bucket was very time-consuming and therefore, slowed down the reindexing of buckets.

With this enhancement, the `rgw-restore-bucket-index` tool is enhanced to allow it to reuse a pre-existing output of a `rados ls` command, thereby allowing administrators to reuse the output from one `rados ls` command. This allows bucket index recovery of multiple non-versioned buckets to be completed more quickly.

8.5.2. Known issues

8.5.2.1. The Cephadm utility

Adding or expanding iSCSI gateways in `gwcli` across the iSCSI daemons works as expected

Previously, due to iSCSI daemons not being reconfigured automatically when a trusted IP list was updated in the specification file, adding or expanding iSCSI gateways in `gwcli` would fail due to the `iscsi-gateway.cfg` not matching across the iSCSI daemons.

With this fix, you can expand the gateways and add it to the existing gateways with `gwcli` command.

([BZ#2099470](#))

`ceph orch ps` does not display a version for monitoring stack daemons

In `cephadm`, due to the version grabbing code currently being incompatible with the downstream monitoring stack containers, version grabbing fails for monitoring stack daemons, such as `node-exporter`, `prometheus`, and `alertmanager`.

As a workaround, if the user needs to find the version, the daemons' container names include the version.

([BZ#2125382](#))

8.6. RED HAT CEPH STORAGE 5.3Z1

Red Hat Ceph Storage release 5.3z1 is now available. The bug fixes that are included in the update are listed in the [RHBA-2023:0981](#) advisory.

8.6.1. Enhancements

8.6.1.1. The Cephadm utility

`cephadm` automatically updates the dashboard Grafana password if it is set in the Grafana service spec

Previously, users would have to manually set the Grafana password after applying the specification.

With this enhancement, if **initial_admin_password** is set in an applied Grafana specification, **cephadm** automatically updates the dashboard Grafana password, which is equivalent to running **ceph dashboard set-grafana-api-password** command, to streamline the process of fully setting up Grafana. Users no longer have to manually set the dashboard Grafana password after applying a specification that includes the password.

OSDs automatically update their Ceph configuration files with the new **mon** locations

With this enhancement, whenever a **monmap** change is detected, **cephadm** automatically updates the Ceph configuration files for each OSD with the new **mon** locations.



NOTE

This enhancement may take some time to update on all OSDs if you have a lot of OSDs.

8.6.1.2. Ceph Dashboard

The Block Device images table is paginated

With this enhancement, the Block Device images table is paginated to use with 10000+ image storage clusters as retrieving information for a block device image is expensive.

Newly added **cross_origin_url** option allows cross origin resource sharing

Previously, IBM developers faced issues with their storage insights product when they tried to ping the REST API using their front-end because of the tight Cross Origin Resource Sharing (CORS) policies set up in Red Hat's REST API.

With this enhancement, CORS is allowed by adding the **cross_origin_url** option that can be set to a particular URL - **ceph config set mgr mgr/dashboard/cross_origin_url localhost** and the REST API allows communication with only that URL.

8.6.1.3. Ceph File System

Users can store arbitrary metadata of CephFS subvolume snapshots

With this enhancement, Ceph File System (CephFS) volume users can store arbitrary metadata in the form of key-value pairs for CephFS subvolume snapshots with a set of command-line interface (CLI) commands.

8.6.1.4. Ceph Object Gateway

STS **max_session_duration** for a role can now be updated

With this enhancement, the STS **max_session_duration** for a role can be updated using the **radosgw-admin** command-line interface.

ListBucket S3 operation now generates JSON output

With this enhancement, on customers' request to facilitate integrations, the **ListBucket** S3 operation generates JSON-formatted output, instead of the default XML, if the request contains an **Accept: application/json** header.

The option to enable TCP keepalive managed by **libcurl** is added

With this enhancement, the option to enable TCP keepalive on the HTTP client sockets managed by **libcurl** is added to make sync and other operations initiated by Ceph Object Gateway more resilient to network instability. This does not apply to connections received by the HTTP frontend, but only to HTTP requests sent by the Ceph Object Gateway, such as Keystone for authentication, sync requests from multi-site, and requests to key management servers for SSE.

Result code 2002 of **radosgw-admin** commands is explicitly translated to 2

Previously, a change in the S3 error translation of internal **NoSuchBucket** result inadvertently changed the error code from the **radosgw-admin bucket stats** command, causing the programs checking the shell result code of those **radosgw-admin** commands to see a different result code.

With this enhancement, the result code 2002 is explicitly translated to 2 and users can see the original behaviour.

You can now use bucket policies with useful errors

Bucket policies were difficult to use since the error indication was wrong. Additionally, silently dropping principals would cause problems during the upgrade. With this update, useful errors from policy parser and a flag to reject invalid principals with **rgw policy reject invalid principals=true** parameter is introduced.

8.6.1.5. Multi-site Ceph Object Gateway

The **bucket sync run** command provides more details

With this enhancement, user-friendly progress reports on the **bucket sync run** command are added to provide users easier visibility into the progress of the operation. When the user runs the **radosgw-admin bucket sync run** command with **--extra-info** flag, users get a message for the start of generation sync and also for each object that is synced.



WARNING

It is not recommended to use the **bucket sync run** command without contacting Red Hat support.

Multi-site configuration supports dynamic bucket index resharding

Previously, only manual resharding of the buckets for multi-site configurations was supported.

With this enhancement, dynamic bucket resharding is supported in multi-site configurations. Once the storage clusters are upgraded, enable the **resharding** feature, zone level, and zone group. You can either manually reshard the buckets with **radosgw-admin bucket reshard command** or automatically reshard them with dynamic resharding, independently of other zones in the storage cluster.

Users can now reshard bucket index dynamically with multi-site archive zones

With this enhancement, multi-site archive zone bucket index can be resharded dynamically when dynamic resharding is enabled for that zone.

8.6.1.6. RADOS

Low-level log messages are introduced to warn user about hitting throttle limits

Previously, there was a lack of low-level logging indication that throttle limits were hit, causing these occurrences to incorrectly have the appearance of a networking issue.

With this enhancement, the introduction of low-level log messages makes it much clearer that the throttle limits are hit.

8.6.1.7. RADOS Block Devices (RBD)

Cloned images can now be encrypted with their own encryption format and passphrase

With this enhancement, layered client-side encryption is now supported that enables each cloned image to be encrypted with its own encryption format and passphrase, potentially different from that of the parent image. The efficient copy-on-write semantics intrinsic to unformatted regular cloned images are retained.

8.6.2. Known issues

8.6.2.1. The Cephadm utility

Adding or expanding iSCSI gateways in `gwcli` across the iSCSI daemons works as expected

Previously, due to iSCSI daemons not being reconfigured automatically when a trusted IP list was updated in the specification file, adding or expanding iSCSI gateways in `gwcli` would fail due to the `iscsi-gateway.cfg` not matching across the iSCSI daemons.

With this fix, you can expand the gateways and add it to the existing gateways with `gwcli` command.

([BZ#2099470](#))

`ceph orch ps` does not display a version for monitoring stack daemons

In `cephadm`, due to the version grabbing code currently being incompatible with the downstream monitoring stack containers, version grabbing fails for monitoring stack daemons, such as `node-exporter`, `prometheus`, and `alertmanager`.

As a workaround, if the user needs to find the version, the daemons' container names include the version.

([BZ#2125382](#))

8.6.2.2. Ceph Object Gateway

Resharding a bucket having `num_shards = 0` results in the bucket's metadata being lost

Upgrade to Red Hat Ceph Storage 5.3 from older releases with buckets having `num_shards = 0` can result in bucket's metadata loss leading to the bucket's unavailability while trying to access it. This is a known issue that will be fixed in an upcoming release. The [Upgrade guide](#) contains the workaround to disable the dynamic bucket resharding and setting `num_shards` to a non-zero value before going for the upgrade. For any help on the upgrade or to know more about the issue, contact [Red Hat Support](#).

([BZ#2174235](#))

CHAPTER 9. SOURCES

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

- For Red Hat Enterprise Linux 8:
<http://ftp.redhat.com/redhat/linux/enterprise/8Base/en/RHCEPH/SRPMS/>
- For Red Hat Enterprise Linux 9:
<http://ftp.redhat.com/redhat/linux/enterprise/9Base/en/RHCEPH/SRPMS/>