Red Hat Ceph Storage

Release Notes

9.1

IBM

Contents

Release notes for 9.1

Red Hat Ceph Storage is a hardened, qualified, secure, and supported enterprise software curated from the Ceph open-source project and delivered by Red Hat.

New features and enhancements

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

`cephadm` utility

Improved `cephadm` upgrade workflow with failure-domain scoped OSD upgrades

Previously, OSD upgrades were not aligned with physical failure domains, leading to less controlled upgrade operations.

With this enhancement, `cephadm` supports failure-domain scoped upgrades using CRUSH bucket parameters, enabling controlled, staged upgrades within racks, chassis, or hosts. This improves operational safety and upgrade efficiency.

(IBMCEPH-11980)

Added automated deployment of Ceph Object Gateway D3N/D4N using `cephadm`

Previously, deploying Ceph Object Gateway D3N/D4N configurations required manual setup and orchestration, increasing complexity and risk of misconfiguration.

With this enhancement, `cephadm` now automates the deployment of Ceph Object Gateway D3N/D4N configurations, simplifying setup and ensuring consistent deployment across clusters. This reduces operational effort and improves reliability for object storage workloads.

(ISCE-773)

Image vendor detection enhancement in cephadm inspect-image

Previously, there was no reliable way to determine whether a Ceph container image was built by IBM or Red Hat when inspecting images at runtime.

With this enhancement, `cephadm inspect-image` now detects and reports the image vendor as either `ibm` or `redhat`. This detection is based on the `/etc/redhat-storage-release` file inside the container image, enabling `cephadm` to identify the vendor from the release string. As a result, this enhancement provides clearer visibility into container image origin, improving troubleshooting, validation, and operational awareness when working with Ceph images.

(IBMCEPH-13854)

SSH hardening and command execution enhancements for `cephadm`

Previously, configuring a `cephadm` user required manual steps, including copying SSH public keys and updating the sudoers file. Additionally, `cephadm` executed commands directly over SSH without a centralized control mechanism. Arbitrary command execution paths were not standardized, and SSH hardening was not available.

With this enhancement, new SSH hardening capabilities and supporting commands are introduced:

- Automated SSH user setup.

```
cephadm setup-ssh-user --ssh-user SSH_USER --ssh-pub-key SSH_PUB_KEY
```

  The `cephadm setup-ssh-user` command copies the SSH public key and updates the sudoers configuration. The existing `cephadm set-user` command now invokes this workflow automatically to ensure the user is correctly configured.

- Standardized command execution interface.

```
cephadm exec --command COMMAND
```

All bash command execution in `cephadm` now uses this interface, providing consistent handling across the `mgr/cephadm` code.

- SSH hardening enablement.

```
ceph cephadm prepare-host-and-enable-ssh-hardening USER [HOST_LABEL]
```

This command prepares the host, configures the `cephadm` user, installs the required `cephadm` version if needed, and enforces command execution through the `cephadm_invoker.py` script. As a result, `cephadm` provides a more secure and consistent SSH execution model, reduces manual configuration steps, and enables hardened command execution paths across managed hosts.

(ISCE-404, IBMCEPH-13315)

Ceph Dashboard

New support for `cephadm`-managed SSL certificates for Ceph Object Gateway services

Previously, SSL certificates required manual configuration and were not integrated with `cephadm`.

With this enhancement, `cephadm`-managed SSL certificates are now supported for Ceph Object Gateway deployment through the dashboard, enabling automated certificate handling and simplifying secure service configuration.

(IBMCEPH-13890)

Enhanced MDS status monitoring

Previously, monitoring the health and state of Metadata Servers (MDS) provided limited visibility into detailed status and operational metrics.

With this enhancement, the dashboard provides improved MDS status monitoring, enabling better visibility into performance, health conditions, and operational state. This enhances troubleshooting and cluster observability.

(ISCE-3031)

Added CephFS subvolume metrics to dashboard

Previously, subvolume-level metrics were not readily available in the dashboard, limiting visibility into usage and performance.

With this enhancement, CephFS subvolume metrics are exposed in the dashboard, enabling administrators to monitor usage patterns and resource consumption more effectively.

(ISCE-3111)

Certificate Manager (`CertMgr`) support for Ceph Dashboard

The dashboard now supports `CertMgr` integration with `cephadm`-signed certificates, allowing certificates to be automatically generated and managed by `cephadm`. This simplifies certificate lifecycle management and reduces manual administration.

The dashboard also monitors certificate health and displays warnings for certificates nearing expiration and alerts for certificate-related failures or misconfigurations. This helps administrators proactively identify and resolve certificate issues, ensuring secure and uninterrupted cluster operations.

(ISCE-2442, ISCE-1407)

New Ceph Object Gateway multisite archive zone support

Previously, the Ceph Dashboard supported configuration of standard Ceph Object Gateway multisite replication but did not support creating or managing archive zones. As a result, users had to rely on CLI commands to configure archive zones, leading to an inconsistent user experience and more complex day-two operations.

With this enhancement, the Ceph Dashboard now supports creating and modifying Ceph Object Gateway zones as archive zones by incorporating the archive tier type into the existing multisite configuration workflow. The Dashboard experience remains consistent with the standard multisite UI, requiring minimal changes to existing workflows. In addition, the Dashboard surfaces archive zone information and operational details, enabling monitoring and management similar to standard zones. As a result, users can configure and manage archive zones directly from the Dashboard, improving usability, consistency, and operational efficiency.

(ISCE-1970)

Ceph File System (CephFS)

Added new MDS metrics for improved observability

Previously, key metrics required for monitoring MDS behavior and performance were not available.

With this enhancement, additional MDS metrics are introduced, improving visibility into metadata operations and helping administrators proactively manage workloads.

(ISCE-3059)

Optimized CephFS default configurations and recommendations

Previously, default configurations for MDS and CephFS required manual tuning for optimal performance in different environments.

With this enhancement, improved defaults and configuration recommendations are introduced, helping administrators achieve better performance and stability with minimal tuning.

(ISCE-2701)

Improved CephFS mirroring performance and reliability

Previously, CephFS mirroring required performance tuning and lacked optimizations for large-scale replication workloads.

With this enhancement, improvements to mirroring functionality enhance performance evaluation and overall reliability of replication across clusters.

(ISCE-2699)

Enhanced CephFS subvolume metrics visibility

Previously, administrators had limited insight into subvolume-level usage and performance characteristics.

With this enhancement, detailed subvolume metrics are introduced for CephFS, enabling improved monitoring, capacity planning, and performance analysis.

(ISCE-3110)

Ceph Block Device

Images can now be added to mirrored Ceph Block Device groups without disabling mirroring

Previously, adding an image to an Ceph Block Device group with mirroring enabled required disabling mirroring, adding the image, and then re-enabling mirroring. As a result, the entire group was resynchronized to the secondary cluster, causing significant network overhead and increased synchronization time.

With this enhancement, images can be added to a mirrored Ceph Block Device group while mirroring remains enabled by using the `rbd group image add` command. Only the newly added image is synchronized to the secondary cluster, while existing images are not resynchronized.

As a result, this reduces unnecessary data transfer, lowers synchronization overhead, and improves operational efficiency.

```
rbd group image add POOL/GROUP POOL/IMAGE
```

(IBMCEPH-13503)

Ceph Object Gateway

Enhanced Lua performance using bytecode caching

Previously, Lua scripts were compiled on every execution, leading to unnecessary overhead and reduced performance.

With this enhancement, Lua bytecode is cached in memory, eliminating repeated compilation and improving execution efficiency for both small and large object workloads.

(IBMCEPH-12115)

Enhanced Ceph Object Gateway policy evaluation logging and testing

Previously, visibility into policy evaluation behavior was limited, making it difficult to trace decisions and troubleshoot access issues.

With this enhancement, setting the **`rgw_copious_policy_logging`** parameter to true enables detailed tracing of policy evaluation in logs. In addition, the `rgw-policy-test` command provides verbose policy evaluation from the command line, helping administrators analyze and validate policy behavior more effectively.

(IBMCEPH-8459)

Enhanced lifecycle processing efficiency for multi-rule configurations

Previously, lifecycle operations involving multiple overlapping rules required multiple passes over objects, impacting performance in large-scale deployments.

With this enhancement, rules are grouped and evaluated in a single pass, improving performance and scalability for multi-site environments.

(IBMCEPH-12820)

Enhanced support for `RestrictPublicBuckets` in PublicAccessBlock configuration

Previously, setting the `RestrictPublicBuckets` option in the PublicAccessBlock configuration had no effect on bucket access behavior because the setting was not implemented.

With this enhancement, support for `RestrictPublicBuckets` is implemented, ensuring that the setting is enforced as expected and improving control over public access restrictions for buckets.

(IBMCEPH-10584)

Enhanced support for Lua-based abort requests

Previously, there was no way to abort a Ceph Object Gateway (`rgw`) request based on the result of a Lua script. If a Lua script encountered a syntax or runtime error, Ceph Object Gateway logged the error but continued processing the request.

With this enhancement, Lua scripts can now return the RGW_ABORT_REQUEST error code to explicitly signal that the request should be aborted. Ceph Object Gateway interprets this return value as `-EPERM` and stops processing the request. This return value is evaluated only in the prerequest context and is ignored in other RGW request-processing contexts. As a result, users can implement more effective request control using Lua scripts, enabling conditional request rejection and improved policy enforcement.

(IBMCEPH-13203)

RADOS

Added `ok-to-upgrade` command for safer OSD upgrades

Previously, upgrade orchestration relied on less efficient mechanisms that prolonged upgrade times.

With this feature, the **`ceph osd ok-to-upgrade`** command identifies safe sets of OSDs that can be upgraded simultaneously without risking data availability, improving upgrade efficiency and operational safety.

(IBMCEPH-12147)

Optimized EC stripe unit alignment for performance

Previously, stripe unit values could be configured without alignment constraints, leading to suboptimal erasure coding performance.

With this enhancement, stripe unit values must be 4K aligned, ensuring improved performance and efficiency for EC workloads.

(IBMCEPH-13263)

Added RocksDB performance counters

Enhanced visibility into BlueStore RocksDB cache behavior Previously, administrators had limited visibility into RocksDB cache usage and shard-level behavior in BlueStore, making it difficult to analyze cache efficiency and diagnose performance issues.

With this enhancement, Red Hat Ceph Storage provides additional performance counters and administrative commands that enable detailed monitoring of RocksDB cache usage, including shard-level statistics. This

improvement helps administrators identify cache inefficiencies and perform more effective performance analysis and tuning.

(ISCE-2956, IBMCEPH-15595)

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 9.x. 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.

Deprecated method of configuring OIDC federation and IAM roles at the tenant level

All OIDC resources are now managed as resources within a Ceph Object Gateway account. These OIDC resources include providers, roles, and polices, As a result, all OIDC operations that target a tenant, including the global or empty tenant, are considered deprecated. The deprecated operations incldue creating providers, creating roles, and assuming roles.

With the newer per-account model, federated users are directly associated with the account and Ceph Object Gateway no longer creates *shadow users* (for example, *`TENANT$USER_NAMESPACE`*) upon role assumption. The account itself tracks all resources and identities.

Tenant-based OIDC federation users should migrate their configurations to the new Ceph Object Gateway per-account model, before feature removal.

For more information, see Secure Token Service.

(ISCE-3162)

Technology Preview

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 production service level agreements (SLAs), might not be functionally complete, and does not recommend using them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

Ceph File System (CephFS)

Support for CephFS snapshot mirroring

With this feature you can replicate a CephFS to a remote CephFS on another Ceph storage cluster. Snapshot synchronization copies snapshot data to a remote Ceph File System, and creates a new snapshot on the remote target with the same name. You can configure specific directories for snapshot synchronization.

Note: To use CephFS snapshot mirroring, both the source and the target storage clusters must be running the same Red Hat Ceph Storage version.

For more information, see CephFS snapshot mirroring (Technology Preview).

Ceph Object Gateway

New per-user and per-bucket usage counters in Prometheus

Ceph Object Gateway now exports per-user and per-bucket usage counters via performance counters automatically collected by the ceph-exporter and made available in Prometheus. This provides low-overhead, real-time visibility into the following:

- Per-bucket metrics: used bytes, utilized bytes, and number of objects
- Per-user metrics: used bytes and number of objects
- Cache performance metrics: cache hits, misses, updates, and evictions

> **Note:** These metrics are disabled by default. To enable them, configure the appropriate settings in your Ceph Object Gateway configuration.

For more information, see Viewing Ceph Object Gateway per-user and per-bucket performance counters.

(BZ#2036531)

RADOS

Balanced primary placement groups can now be observed in a cluster

Previously, users could only balance primaries with the offline `osdmaptool`.
With this enhancement, autobalancing is available with the `upmap` balancer. Users can now choose between either the `upmap-read` or `read` mode. The `upmap-read` mode offers simultaneous `upmap` and `read` optimization. The `read` mode can only be used to optimize reads.

For more information, see Using the Ceph Manager balancer module.

(BZ#1870804)

Now supports tracking data availability score of a cluster

This release introduces a feature that tracks the data availability score of a Ceph cluster over time. The score represents how accessible your data is at any given moment, based on factors such as OSD health, placement group states, and redundancy policies.

By monitoring this metric, administrators gain a fact-based view of cluster reliability and can validate availability percentages (for example, 99.99%) against service-level objectives. This capability provides actionable insight into operational resilience and helps ensure confidence in Ceph as a storage platform for critical workloads.

For more information, see Track the data availability score of a cluster (Technology Preview).

(ISCE-2144)

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.

`cephadm` utility

Improved error reporting for OSD redeployment when DB device slots are unavailable

Previously, the `ceph-volume batch` command did not correctly count devices with existing LVM data. After zapping an OSD, redeployment could fail silently because no DB device slots were available.

With this fix, the condition is treated as an error and a clear message is logged. As a result, when deployment fails due to unavailable DB device slots, users are informed through explicit log errors.

(IBMCEPH-12884)

Invalid HA cluster port values are now prevented

Previously, when creating a high availability (HA) cluster with a custom an internal offset caused the assigned port to exceed the valid range.

With this fix, validation has been added to ensure that the resulting port value does not exceed the maximum allowed limit of 65535. As a result, if the provided custom port plus the offset exceeds the valid range, the HA cluster creation command fails with a validation error instead of assigning an invalid port.

(IBMCEPH-11819)

Ceph Manager

Manager modules now load reliably after failover

Previously, a race condition in the manager (MGR) component could cause modules to fail to load in time after a failover. This could happen during an upgrade or when running the **`ceph mgr fail`** command. As a result, MGR module commands could fail with misleading error messages indicating that a module was not enabled or loaded, even though the issue was related to delayed loading.

With this fix, the MGR declares itself active only after all modules have been successfully loaded and are ready to receive commands. In addition, error messages have been improved to clearly distinguish between modules that failed to load in time and those that are not enabled. As a result, MGR module commands now execute reliably after failover, and clearer health warnings are provided if any modules fail to load in time.

(IBMCEPH-9881)

Ceph Dashboard

Improved clarity of performance card empty states

Previously, a generic empty state message was displayed when storage or Prometheus was not configured, providing no clear indication of the issue.

With this fix, empty state messages are displayed dynamically based on the configuration status. As a result, users can clearly identify whether Prometheus is disabled, not configured, storage is not configured, or performance charts are displayed as expected.

(IBMCEPH-13420)

Duplicate storage class name validation in Ceph Dashboard

Previously, the Ceph Dashboard did not validate duplicate storage class names during creation. As a result, creating a storage class with a name that already existed unintentionally modified the existing storage class instead of creating a new one.

With this fix, validation is added to prevent duplicate storage class names. As a result, users can no longer create a storage class with a name that already exists, preventing unintended modifications to existing storage classes.

(IBMCEPH-13044)

Ceph File System (CephFS)

Quota display is now accurate in mixed quota modes

Previously, if a quota was set on a parent directory, it was used for displaying quota information. This caused incorrect values when different quota types were defined at different directory levels.

With this fix, the system identifies the correct quota root separately for `max_files` and `max_bytes`. As a result, quotas are displayed correctly in mixed quota configurations.

(IBMCEPH-12286)

Fixed stability issues in `LogSegment` lifecycle management

Previously, `LogSegment` objects were managed using raw pointers, allowing multiple MDS components to reference them without a clear ownership model. This could leave dangling pointers when segments were trimmed or expired.

With this fix, reference counting has been introduced for `LogSegment`. As a result, segments are released only after all references are cleared, eliminating dangling pointers and improving MDS stability during operations such as replay and recovery.

(IBMCEPH-11811)

Snapshot info now correctly includes enctag

Previously, the `enctag` field was not populated when querying filesystem snapshot information. As a result, an error occurred and the snapshot info output was not fully displayed.

With this fix, the `enctag` field is populated during snapshot queries and snapshot information now displays correctly, including all attributes such as `enctag`.

(IBMCEPH-13575)

MDS no longer crashes due to improper `mds_lock` usage

Previously, a bug in the MDS allowed certain code paths to execute without holding the required `mds_lock`, leading to unsafe behavior and unpredictable crashes.

With this fix, the `mds_lock` is consistently held wherever required. As a result, the MDS operates reliably without unexpected crashes.

(IBMCEPH-13700, IBMCEPH-13705)

Ceph Block Device

Clone image sparseness is preserved during live migration

Previously, due to an implementation defect, sparseness was lost when live migrating a clone image. As a result, the destination image could consume significantly more space than the source clone image, including its entire parent chain.

With this fix, the implementation has been corrected to preserve sparseness during live migration. As a result, the destination image now consumes the same amount of space as the source clone image, including its parent chain.

(IBMCEPH-13504)

Ceph Object Gateway

Replayed multipart upload requests now return correct `etag`

Previously, the `CompleteMultipartUpload::etag` member was not assigned when replaying an already completed multipart upload request. As a result, an empty `etag` header was returned for replayed requests.

With this fix, the `etag` member is assigned the validated value in the replay response path. As a result, replayed requests now return the correct `etag` as expected.

(IBMCEPH-13821)

Conditional multipart upload operations now function as expected

Previously, while S3 operations were supported, HTTP preconditions for conditional multipart upload were not.

With this fix, HTTP preconditions for conditional multipart upload are fully implemented. As a result, conditional multipart upload operations are now supported as expected.

(IBMCEPH-13557)

Unordered bucket listings no longer repeat entries or loop indefinitely

Previously, unordered bucket listing had a bug when traversing regions of the entry space, which could cause entries to repeat or the operation to loop indefinitely.

With this fix, the traversal logic has been corrected. As a result, unordered bucket listing completes as expected without duplicate entries or infinite loops.

(IBMCEPH-12291)

Ceph Object Gateway no longer crashes during lifecycle processing of versioned buckets

Previously, the `rgw` daemon could crash during lifecycle (LC) processing on versioned buckets with a large number of object versions, particularly when expiration and noncurrent version expiration rules were applied concurrently.

With this fix, lifecycle processing handles object listing boundaries correctly, preventing invalid memory access during concurrent or batched operations. As a result, Ceph Object Gateway (`rgw`) no longer crashes during lifecycle deletion, and lifecycle processing completes reliably on large versioned buckets.

(IBMCEPH-14308)

Log buckets are now supported in erasure-coded pools

Previously, log buckets could not be created in erasure-coded (EC) pools due to append operation limitations.

With this fix, log records are first written to a temporary object in the replicated `default.rgw.log` pool and then copied to the EC pool upon commit. Implicit commit operations are handled asynchronously by a new `BucketLoggingManager`, while explicit commits continue to run synchronously. As a result, log buckets can now be used with EC pools, enabling more flexible storage configurations.

(IBMCEPH-11747)

Restore operations now validate storage class and return correct results

Previously, restore operations targeting a non-existent storage class succeeded silently and defaulted to the `STANDARD` class. However, the head-object API still reported the non-existent storage class, creating a mismatch between actual data placement and reported metadata.

With this fix, the restore operation validates whether the requested storage class exists before proceeding and fails with an appropriate error if it does not. As a result, restore operations no longer succeed with invalid storage classes, and object metadata accurately reflects the actual storage class.

(IBMCEPH-11517)

Cloud restore and transition operations are more robust and reliable

Previously, multiple issues affected cloud restore and transition workflows. Setting `Days=0` in restore requests could leave objects in a broken state, multipart upload resumption could fail, restored objects could include incorrect ETag formatting, and the restore process was not shut down in the correct order.

With this fix, these issues are addressed across the restore and transition modules to ensure correct request handling, proper ETag formatting, and orderly process management. As a result, cloud restore and transition operations are more reliable, improving overall stability and service availability.

(IBMCEPH-9849)

Ceph Object Gateway multi-site

Bucket index metadata remains consistent after lifecycle expiration

Previously, in rare cases, lifecycle expiration in versioned buckets could leave behind stale bucket index metadata. This occurred due to incorrect shutdown and destruction ordering, which could trigger a segmentation violation and prevent proper cleanup. As a result, stale index entries remained even after objects were removed, potentially impacting operations that rely on accurate bucket index listings and causing errors such as `(27) File too large`.

With this fix, the Asio run queue is drained before destroying the storage backend, ensuring proper shutdown sequencing and cleanup. As a result, bucket index metadata is now correctly maintained after lifecycle expiration, preventing stale entries and improving stability.

(IBMCEPH-12980)

Ceph Object Gateway no longer crashes on client disconnect and bulk delete operations

Previously, two issues could cause Ceph Object Gateway to crash during request handling. When a client disconnected while Ceph Object Gateway was sending an error response, an unhandled exception could trigger an abort. In addition, bulk delete requests could cause an assertion failure due to attempts to close an already closed XML section.

With this fix, broken pipe errors from client disconnects are handled gracefully at write time, and redundant XML section close calls are removed. As a result, Ceph Object Gateway no longer crashes during client disconnects or bulk delete operations, improving overall stability.

(IBMCEPH-14447)

RADOS

Idle PGs are no longer incorrectly flagged as `stuck peering`

Previously, when OSDs restarted or were marked down, the cluster could report placement groups (PGs) as `stuck peering` even when they were active. This led to misleading `HEALTH_WARN` messages.

With this fix, PG statistics are updated more frequently, including for idle PGs, and the `last_active` and `last_peered` timestamps and reporting logic are adjusted. As a result, idle PGs are no longer incorrectly flagged as stuck peering, and `HEALTH_WARN` messages are triggered only after genuine timeouts (greater than 60 seconds).

(IBMCEPH-6744)

All bug fixes

This section lists a complete listing of all bug fixes in this release of Red Hat Ceph Storage 9.1.

Issue key	Severity	Summary
IBMCEPH-12414	Critical	OSD journal replay failures during startup
IBMCEPH-13341	Critical	Get request fails with 404 after successful HEAD operation
IBMCEPH-13575	Critical	Subvolume info CLI command fails due to undefined variable
IBMCEPH-13576	Critical	Cephadm bootstrap fails with SSH user and custom keys
IBMCEPH-13592	Critical	RGW dashboard delete confirmation text not clearly visible
IBMCEPH-13611	Critical	Dashboard missing archive option during multisite import
IBMCEPH-13617	Critical	Dashboard missing option to edit multisite sync configuration
IBMCEPH-13672	Critical	Object browser page remains loading indefinitely
IBMCEPH-13680	Critical	Bucket creation page stuck without response in dashboard
IBMCEPH-13718	Critical	Clean installation of object browser fails
IBMCEPH-13749	Critical	Per user and per bucket counters not available
IBMCEPH-13782	Critical	Uploaded objects show incorrect file size in dashboard
IBMCEPH-13790	Critical	Dashboard user export creation fails with internal server error
IBMCEPH-13799	Critical	S3 browser not showing buckets with SSL endpoint
IBMCEPH-13861	Critical	Dashboard should prevent login with invalid credentials and remain on login page
IBMCEPH-13998	Critical	Cluster deployment fails when using sudo user
IBMCEPH-14005	Critical	Dashboard sync from option retains old zone after configuration change
IBMCEPH-14031	Critical	Incorrect archive option behavior during multisite setup
IBMCEPH-14082	Critical	Dashboard support diagnostics section shows incorrect entries
IBMCEPH-14088	Critical	Dashboard documentation link returns 404 page
IBMCEPH-14142	Critical	Uploaded files not shown immediately in dashboard object view
IBMCEPH-14191	Critical	Dashboard bucket creation stuck after service creation

Issue key	Severity	Summary
IBMCEPH-14193	Critical	EC optimization default incorrectly set to false
IBMCEPH-14223	Critical	ceph orch accept call home enabled command not working correctly
IBMCEPH-14228	Critical	Dashboard API missing functionality to set and clear MOTD
IBMCEPH-14249	Critical	Dashboard multisite sync not working as expected
IBMCEPH-14255	Critical	Orchestrator cannot redeploy DB devices
IBMCEPH-14267	Critical	Keepalived service default configuration causes service issues
IBMCEPH-14273	Critical	RGW segfault during AWS cloud transition with HTTPS endpoint
IBMCEPH-14278	Critical	RGW multisite sync daemon crashes during large bucket full sync
IBMCEPH-14308	Critical	RGW crash during lifecycle deletion on versioned buckets
IBMCEPH-14314	Critical	Ceph orch upgrade violates enforced upgrade order
IBMCEPH-14377	Critical	s3 browser refresh shows duplicate entries
IBMCEPH-14378	Critical	RGW not able to configure TLS 1.3 cipher suites using ssl_ciphers parameter
IBMCEPH-14441	Critical	Cephadm bootstrap fails on RHCS 9.1 while configuring dashboard SSL port
IBMCEPH-15007	Critical	cephadm bootstrap fails with crimson image if automatically accept license flag not provided
IBMCEPH-15196	Critical	RGW Cloud Tier Restore fails for AWS S3 with SSE-S3 encryption error
IBMCEPH-15708	Critical	OSDs going down after monitor remove add operations
IBMCEPH-12080	Important	Bootstrap fails due to incorrect IP address validation
IBMCEPH-12115	Important	Performance degradation observed for certain object sizes
IBMCEPH-12147	Important	ok-to-upgrade command behavior needs improvement
IBMCEPH-12286	Important	Quota reset not reflected until quota files refreshed
IBMCEPH-12701	Important	RGW segfault occurs during object upload
IBMCEPH-12719	Important	OSD service count misreported in orchestrator output

Issue key	Severity	Summary
IBMCEPH-13044	Important	Duplicate storage class creation overwrites existing configuration
IBMCEPH-13315	Important	Cephadm SSH channel hardening incomplete
IBMCEPH-13318	Important	Directory configuration issue impacts namespace behavior
IBMCEPH-13340	Important	Required backport request needed for upstream fix
IBMCEPH-13401	Important	Firewall port not opened for HA proxy peer communication
IBMCEPH-13403	Important	Required backport missing in release build
IBMCEPH-13450	Important	Optional service dependency installation missing during host preparation
IBMCEPH-13491	Important	RGW S3 requests return HTTP 409 error
IBMCEPH-13494	Important	HeadObject request returns incorrect 403 response
IBMCEPH-13540	Important	Bootstrap CLI help message duplication for license options
IBMCEPH-13557	Important	Multipart upload handling requires fix
IBMCEPH-13594	Important	Unable to modify multisite endpoints in dashboard
IBMCEPH-13605	Important	Dashboard landing page contains excessive empty space
IBMCEPH-13621	Important	Duplicate capacity metrics displayed for CephFS in dashboard
IBMCEPH-13639	Important	CephFS mirror directories stuck in syncing state
IBMCEPH-13651	Important	Dashboard Grafana capacity metrics not displayed correctly
IBMCEPH-13662	Important	RGW crashes when admin sends invalid AssumeRole request
IBMCEPH-13669	Important	Object browser login page missing example inputs and guidance
IBMCEPH-13784	Important	Creating folder generates unintended dummy file
IBMCEPH-13788	Important	RGW multisite full sync fails silently while reporting success
IBMCEPH-13794	Important	Dashboard redirects to add storage page after re login
IBMCEPH-13851	Important	RGW segfault occurs under high S3 write workload
IBMCEPH-13864	Important	Edit client config option stuck in loading state
IBMCEPH-13894	Important	Dashboard device availability warnings are misleading

Issue key	Severity	Summary
IBMCEPH-13921	Important	Read only users cannot access dashboard overview page
IBMCEPH-13931	Important	Dashboard filesystem overview missing MDS details
IBMCEPH-13973	Important	Dashboard shows RAM disks incorrectly in device inventory
IBMCEPH-14053	Important	Bucket creation remains stuck when bucket already exists
IBMCEPH-14054	Important	Object browser image should use ICR image instead of current image
IBMCEPH-14059	Important	Bucket settings show incorrect object count instead of actual objects
IBMCEPH-14135	Important	Dashboard missing MDS observability metrics
IBMCEPH-14140	Important	Dashboard reports incorrect CephFS capacity usage
IBMCEPH-14207	Important	Dashboard pool manager role unable to create pools due to access denied
IBMCEPH-14282	Important	Dashboard icon misalignment in add storage and overview buttons
IBMCEPH-14429	Important	Unable to recreate topic created with invalid characters
IBMCEPH-11602	Moderate	Directory removal implementation inefficient
IBMCEPH-11846	Moderate	Dashboard table layout expands incorrectly when selecting modules
IBMCEPH-11914	Moderate	Dashboard replication policy changes not applied correctly
IBMCEPH-12403	Moderate	Deleting storage class does not remove associated zone entry
IBMCEPH-12520	Moderate	Dashboard metrics not visible after upgrade
IBMCEPH-12659	Moderate	Deep-scrub completion timing inconsistent
IBMCEPH-12689	Moderate	OSD deployment fails on nodes with high OSD count
IBMCEPH-12820	Moderate	Replication for new buckets does not function correctly
IBMCEPH-12873	Moderate	Linking bucket to account fails with internal server error
IBMCEPH-12987	Moderate	Dashboard unable to restart or stop RGW service
IBMCEPH-13263	Moderate	Stripe unit alignment restriction missing for EC optimization
IBMCEPH-13555	Moderate	OSD delete popup missing option to preserve OSD ID

Issue key	Severity	Summary
IBMCEPH-13750	Moderate	Dashboard empty state illustration missing image asset
IBMCEPH-13890	Moderate	Dashboard missing support for `cephadm` managed SSL certificates
IBMCEPH-14089	Moderate	Dashboard overview graphs display incorrect decimal values
IBMCEPH-14144	Moderate	Dashboard bucket content not visible without manual refresh
IBMCEPH-13731	Low	Dashboard host selection not populated during namespace creation
IBMCEPH-13785	Low	Dashboard shows duplicate objects intermittently in object browser

Known issues

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

`cephadm` utility

Promtail image continues to appear in `cephadm` list-images output

The Promtail container image might still appear in `cephadm` list-images output after upgrading to Alloy.

Currently, during the transition from Promtail to Alloy, `cephadm` continues to register the Promtail container image to maintain backward compatibility. As a result, the Promtail image remains visible even though Alloy is the default logging solution.

As a workaround, no action is required. Ignore the Promtail image entry during the transition phase.

(IBMCEPH-13162)

Grafana certificate does not migrate during upgrade

When you upgrade from Red Hat Ceph Storage 8.1 to 9.0, the existing user-signed Grafana certificate is not migrated. Instead, Grafana switches to a `cephadm`-signed certificate. As a result, duplicate certificate entries may appear, and certificate-related health warnings can persist. Manual reconfiguration is required if you want to use custom TLS certificates.

> **Note:** Data services remain unaffected.

To work without custom TLS certificates, you can continue using the `cephadm`-signed certificate.

As a workaround to use custom TLS certificates, complete the following steps:

1. Change the Grafana specification to use `certificate_source: reference`.
2. Use `certmgr` to upload a valid user-signed certificate and key for each host.
3. Run the **`ceph orch reconfig grafana`** command.

(IBMCEPH-13080)

Monitor configuration updates are not applied on restart

When you update monitor configuration settings, such as `public_network`, the changes are not applied when the monitor daemon is restarted. This occurs because monitor daemon configurations are not dynamically refreshed during a restart.

As a result, the monitor continues to run with the previous configuration, and the updated values do not take effect.

As a workaround, redeploy the monitor daemon instead of restarting it after updating the configuration. This ensures that the updated configuration is applied successfully.

(IBMCEPH-12242)

Crash daemon cannot access crash directory due to permission changes

When certain services, such as Grafana, are deployed, the permissions of the crash directory can change. As a result, the crash daemon cannot access the directory, preventing it from functioning correctly.

As a workaround, manually update the permissions of the crash directory to `167`. You must repeat this action each time a daemon deployment changes the directory permissions to ensure proper access.

(IBMCEPH-12678)

Ceph File System (CephFS)

`root_squash` kernel client may cause data inconsistency and triggers `HEALTH_ERR`

A bug in the root_squash implementation can cause changes made by a kernel client restricted with `root_squash` capabilities to be lost. Although the issue is fixed for the FUSE client and the MDS, the kernel client remains affected. As a result, the cluster emits the following error when it detects a client with the broken `root_squash` implementation:

```
HEALTH_ERR: MDS_CLIENTS_BROKEN_ROOTSQUASH
```

This occurs, due to the risk of data inconsistency and lost updates.

To avoid this issue, it is recommended to discontinue using `root_squash` with kernel clients until a fix is available.

To prevent affected clients from connecting, you can evict and permanently block them by setting the required client feature.

```
ceph fs required_client_features add client_mds_auth_caps
```

This helps protect the cluster from inconsistent behavior caused by affected clients.

(IBMCEPH-14902)

Ceph Object Gateway multi-site

Secondary site displays old zonegroup name after rename

After renaming a zonegroup in a multisite configuration, the secondary site might still display the previous zonegroup name.

Currently, when a zonegroup is renamed on the primary site, the old name is not removed from the .rgw.root pool. As a result, both the old and new zonegroup names appear in the `radosgw-admin zonegroup list` output, and sync operations might be impacted.

As a workaround, perform the following steps:

1. Verify that the new zonegroup name exists.

   ```
   radosgw-admin zonegroup list
   ```

2. List entries in the `.rgw.root` pool and locate the old zonegroup name.

   ```
   rados -p .rgw.root ls
   ```

   The old name appears in the format:

   ```
   zonegroups_names.OLD_ZONEGROUP_NAME
   ```

3. Remove the old zonegroup name from the pool:

   ```
   rados -p .rgw.root rm zonegroups_names.OLD_ZONEGROUP_NAME
   ```

   Removing the old zonegroup name restores normal sync operations.

(IBMCEPH-13140)

RADOS

`ceph versions -f xml` command produces non-well-formed XML output

When you run the **`ceph versions -f xml`** command, the generated output is not well-formed XML and cannot be parsed by standard XML parsers. This occurs because the command uses full Ceph version strings (including special characters such as dots, parentheses, spaces, and hashes) as XML tag names, which violates XML syntax rules.

As a result, XML parsing fails with errors such as not well-formed (invalid token), preventing automated processing or validation of the output.

Currently there is no workaround.

(IBMCEPH-13690)

Sources

Use this information to attain Red Hat Ceph Storage source code packages.

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

- For Red Hat Enterprise Linux 9: http://ftp.redhat.com/redhat/linux/enterprise/9Base/en/RHCEPH/SRPMS/
- For Red Hat Enterprise Linux 10: https://ftp.redhat.com/redhat/linux/enterprise/10Base/en/RHCEPH/SRPMS/

Published date: 2026-06-