

[Subscriptions](#) | [Downloads](#) | [Red Hat Console](#) | [Get Support](#)



Red Hat F

RHSA



6-01-29

Overview

Updated P

Synop

Moderat

Type/Severity

Security Advisory: Moderate

Red Hat Lightspeed patch analysis

Identify and remediate systems affected by this advisory.

[View affected systems](#)

Topic

Red Hat Ceph Storage 9.0 is now available.

Description

Red Hat Ceph Storage is a 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.

Solution

Before applying this update, make sure all previously released errata relevant to your system have been applied.

For details on how to apply this update, refer to:

<https://access.redhat.com/articles/11258> 













For supported configurations, refer to:

<https://access.redhat.com/articles/1548993> 

Affected Products

- Red Hat Enterprise Linux for x86_64 10 x86_64
- Red Hat Enterprise Linux for x86_64 9 x86_64
- Red Hat Enterprise Linux for IBM z Systems 10 s390x
- Red Hat Enterprise Linux for IBM z Systems 9 s390x
- Red Hat Enterprise Linux for Power, little endian 10 ppc64le
- Red Hat Enterprise Linux for Power, little endian 9 ppc64le
- Red Hat Enterprise Linux for ARM 64 10 aarch64
- Red Hat Enterprise Linux for ARM 64 9 aarch64

Fixes

- [BZ - 1696875](#)  - rgw multisite: automate clean up deleted bucket instance and index objects
- [BZ - 1892474](#)  - [RFE] [RADOS - Stretch Cluster] - Allow moving back to normal mode from stretch cluster mode.
- [BZ - 2011600](#)  - On EC reads always read entire stripe
- [BZ - 2022705](#)  - NFS deployment with positional and named arguments is inconsistent.
- [BZ - 2036531](#)  - [RFE][CEE/SD][ceph-metrics] Add per-RGW user metrics to Prometheus module
- [BZ - 2053348](#)  - [RFE] [RGW-MultiSite] continuous consistency reporting and healing
- [BZ - 2154745](#)  - [RFE] : Provide option to see ceph orchestrator ps in tree format with respect to host name
- [BZ - 2161889](#)  - [RFE] [Ceph Dashboard]: add bucket notifications
- [BZ - 2174473](#)  - CVE-2023-25153 containerd: OCI image importer memory exhaustion
- [BZ - 2239586](#)  - [RFE][RHCS 6.1][Please automatically set the correct permissions for the D3N cache directory with no user intervention.]
- [BZ - 2248613](#)  - [ganesha][RFE] subvolume quotas not passed to the nfs client
- [BZ - 2268063](#)  - [7.1]: Deleting an rgw realm, does not clear the 'realm_id' and it is listed in the 'default_info'

- [BZ - 2270403](#) - [Workload-DFG] [ReleaseCriteria 7.0] Regression observed post aging with hybrid workload
- [BZ - 2272916](#) - [Dashboard][RFE]Unable to delete multiple users from Users page
- [BZ - 2273066](#) - Return code for "ceph orch host maintenance enter <hostname>" is 0 on error
- [BZ - 2275081](#) - [RFE][CephFS][Add CephFS performance metrics per volume / subvolume]
- [BZ - 2276931](#) - [rgw][sts]: assume_role_with_web_identity call is failing as validation of signature is failing with invalid padding
- [BZ - 2279848](#) - There must be some way to handle scenario when ceph rgw realm bootstrap failed but real/zonegroup/zone is created
- [BZ - 2282699](#) - [rgw][notif]: object size 0 seen in event record upon lifecycle transition event
- [BZ - 2283916](#) - Observing slow_ops on a mon daemon post site down tests in a 3 AZ cluster
- [BZ - 2293233](#) - [rgw][sts][sns]: without permission for sns:CreateTopic in role or session policy, create topic operation succeeded with sts creds
- [BZ - 2294700](#) - [rgw][ssl][s3-tests]: sts tests failed with access denied for list buckets of alt user in teardown phase
- [BZ - 2298638](#) - [RFE] ceph-osd should offer partial (chunk instead of stripe) reads over EC pools
- [BZ - 2302464](#) - [CEE/SD][RFE][cephadm] "ceph orch" commands should report the exit code after checking the command results.
- [BZ - 2302541](#) - [rgw][sts]: Deny AbortMultipartUpload is not respected if Allow PutObject permission is also present in the session policy
- [BZ - 2303488](#) - rgw: introduce minimally pausing resharding enhancement 8.0
- [BZ - 2305185](#) - [rgw][acl] put-bucket-acl with acl set as public-read-write is allowing only public-write but public-read of objects is denied
- [BZ - 2305659](#) - [8.0][minimal-pausing-reshard]: Error during resharding bucket squid-1:(16) Device or resource busy while an incremental sync was in progress.
- [BZ - 2309383](#) - [8.0][RGW Read-affinity]: Localized reads not happening on a 2 DC stretch cluster
- [BZ - 2309470](#) - [3AZ Stretch pool] [RFE] stretch unset command does not revert all the values set on the pool
- [BZ - 2312579](#) - CVE-2024-11831 npm-serialize-javascript: Cross-site Scripting (XSS) in serialize-javascript
- [BZ - 2312933](#) - [s3-cloud-restore]: Resume InProgress restore operation post RGW service restart
- [BZ - 2312937](#) - [s3-cloud-restore]: Update expiry date if the `restore-object` API request is retried on temporary copies
- [BZ - 2313740](#) - [CephFS - Subvolume] Not all subvolumes are listed when Cluster Storage is in almost full state






























- [BZ - 2313801](#) - [rgw] support for s3:GetObjectVersionAttributes as bucket policy action
- [BZ - 2313820](#) - [rgw] BucketPolicy with s3:GetObjectAttributes permission alone is allowing other users to perform the operation without depending on s3:GetObject permission
- [BZ - 2314311](#) - [NFS-Ganesha] Could not send metrics, aborting: Broken pipe - ganesha.log
- [BZ - 2314624](#) - [8.0] Blank entry for Storage Capacity in dashboard under Cluster > Expand Cluster > Review
- [BZ - 2314829](#) - [RFE]: When we open an about panel, other options at right top corner is still enabled and if we click on any of them, that page will open in the background, and about panel will not shrink automatically
- [BZ - 2315832](#) - [dashboard][rgw][rfe]: Support creation/management of RGW Accounts via dashboard along with root user creation
- [BZ - 2315856](#) - s3 GetObject with partNumber doesn't decrypt SSE-C or SSE-S3
- [BZ - 2315922](#) - [ceph-dashboard] unable to edit pipe config for bucket level policy of a bucket
- [BZ - 2316019](#) - [crimson] OSD device class is not being displayed in outputs
- [BZ - 2316597](#) - [RFE] RGW User Keys creation date info
- [BZ - 2317768](#) - [RFE][RGW-MS]: Object lock configuration rule is not synchronized to the secondary zone
- [BZ - 2317778](#) - Clicking on info section of Lifecycle section navigates page control to the beginning of the page
- [BZ - 2318936](#) - [RFE][MON] Allow monitors to detect netsplit when using connectivity election strategy
- [BZ - 2319055](#) - [RFE] Do not show Silenced health warnings in dashboard landing page
- [BZ - 2319581](#) - [crimson] OSD deployment fails due to low value being set for fs.aio-max-nr
- [BZ - 2322389](#) - Clicking on Bucket policy section navigates page control to the beginning of the page
- [BZ - 2324147](#) - [RGW][Checksum]: Checksum algorithm is not displayed in list-parts
- [BZ - 2329994](#) - [rgw/listing][rfe]: Timeout in list-objects-v2 operation for large bucket listing
- [BZ - 2338350](#) - [RFE][10k osd] Allow registry credentials to define multiple container registries
- [BZ - 2341758](#) - [RFE] Provide user-friendly error message when OSD daemon is deployed on non-existing device
- [BZ - 2343728](#) - [RGW] Delete bucket policy should return 204 instead of 200
- [BZ - 2344639](#) - [rgw] public-access-block with BlockPublicAcls set on a Bucket is denying normal object upload without any public-acl in the request
- [BZ - 2345474](#) - [10k osd][RFE] ceph orchestrator lost the hosts and marked them as stray host
- [BZ - 2345487](#) - [RFE][s3-cloud-restore]: Admin CLIs and OPs to monitor/debug restore status/ops

- [BZ - 2346030](#) - [10k osd]cephadm timeout error during refresh
- [BZ - 2346107](#) - grafana.update.checker trying to reach Internet in a disconnected environment
- [BZ - 2346779](#) - [10k osd][RFE] Dashboard sending `get-credentials` calls without having secure_monitoring_stack enabled
- [BZ - 2348664](#) - [10k osd] segmentation fault osd : tick_without_osd_lock()
- [BZ - 2349154](#) - [RFE] Source information (subvol/snapshot) of clone is removed from .meta file after cloning has been completed.
- [BZ - 2350732](#) - [IBM Support] [RGW] Conditional write doesn't work in certain scenarios
- [BZ - 2354195](#) - CVE-2025-30204 golang-jwt/jwt: jwt-go allows excessive memory allocation during header parsing
- [BZ - 2354779](#) - [RGW] S3 Credentials created at secondary site are erased when new credentials are created at primary using the AdminOps API
- [BZ - 2354783](#) - RGW Multisite. Object ACLs aren't replicating from primary site to secondary site
- [BZ - 2356786](#) - [RFE][Ceph-Object][Tiering]:Tiering tab lacks Local Tiering support and needs Label update
- [BZ - 2356814](#) - [IBM SUPPORT] [cephFs][Mirroring] I/O performance is varying between mirroring ceph clusters
- [BZ - 2358809](#) - In the helper text of retain head object, need a correction as default: false
- [BZ - 2358811](#) - [RFE]: Display value of retain head object in the tiering details page
- [BZ - 2359400](#) - [CephFS - FScrypt] Enctag get for tag with 256 character count returns 'Unexpected error'
- [BZ - 2359626](#) - [GSS] MGRs constantly failed over and interfering with upgrade
- [BZ - 2360425](#) - [8.1][rgw] radosgw-admin topic list output is not wrapped with "topics" keyword
- [BZ - 2360695](#) - [RFE][rgw][accounts]: Include a field 'account_root_user' in account get command output to mention the account_root_user
- [BZ - 2361474](#) - Renamed mirror group on secondary cluster is not reflected on primary side and a subsequent planned failover results in up+error state
- [BZ - 2361672](#) - [RFE] Request Ganesha to support binding to specific address for monitoring
- [BZ - 2362686](#) - [CephFS - FScrypt] Prompt a error/warning message during encrypt if unsupported encryption algorithm specified in fscrypt.conf file
- [BZ - 2362867](#) - OSDs unreachable with IPv6 subnet and in subnet all public address
- [BZ - 2363597](#) - [NFS-Ganesha] Incorrect IP shown in NFS Ganesha info when bound to non-primary interface on multi-NIC server
- [BZ - 2363934](#) - URL updates needed
- [BZ - 2364323](#) - RGW bucket index downsharding does not shrink shards to default 11 shards after LC deletion of large number of objects.

- [BZ - 2364399](#) - [RFE][rgw][server-access-logging] send flushed object name in API reply
- [BZ - 2364414](#) - Optimize daemon orchestration with concurrency
- [BZ - 2365095](#) - [RFE] Support Ceph RGW S3 Restore for IBM COS ACCELERATED Storage Class (Equivalent to AWS Glacier)
- [BZ - 2365225](#) - [RFE][RGW-Dashboard]: Provide value of tier_type in the details section of storage class tiering to differentiate between different storage class
- [BZ - 2365230](#) - [RFE][RGW-Dashboard]: Provide an option to select cloud type while creating storage class tiering in dashboard
- [BZ - 2365291](#) - [RGW-Dashboard]: Provide an option to set "acl_mappings" while creating storage class tiering
- [BZ - 2365292](#) - [RFE][RGW-Dashboard]: Include value of acl_mappings in the details section of storage class tiering
- [BZ - 2365697](#) - [rgw][server-access-logging]: REST.POST.UPLOAD (complete multipart) log record is not populated with object size even in the Standard mode
- [BZ - 2365931](#) - [rgw][server-access-logging]: please modify the operation name for multipart upload_part to REST.PUT.PART in the Standard mode
- [BZ - 2366182](#) - [rgw][rfe] The zonegroup rename should also update the api_name field in the zonegroup map
- [BZ - 2367347](#) - [GSS] rook-ceph-nfs-ocs-storagecluster-cephnfs-a pod in CrashLoopBackOff after ODF upgraded to 4.18.3
- [BZ - 2367725](#) - upload-part-copy fails when object name contains % in it
- [BZ - 2368270](#) - [IBM Support][RGW] Deleted bucket appears in Dashboard and Edit button grey out until the page is not refreshed
- [BZ - 2369130](#) - [EC 8+6 MSR] Unable to create 8+6 MSR EC profile with ISA-L plugin
- [BZ - 2370197](#) - [NFS-Ganesha][HAProxy-Protocol] VIP failback to original node causes I/O failures with "Remote I/O error" on clients
- [BZ - 2370353](#) - AttributeError encountered when attempting to log in to Cephadm Shell of a stopped OSD
- [BZ - 2370358](#) - The osd_scrub_extended_sleep parameter exhibits no effect on scrub time when its value is increased.
- [BZ - 2370450](#) - [Test Build] Not all NFS Daemons are coming up when deploying NFS cluster in parallel at scale
- [BZ - 2370541](#) - [Test Build] NFS service deployment failing on freshly installed cluster - grace tool failed: rados_pool_create: -1
- [BZ - 2371109](#) - [rgw][server-access-logging]: AuthenticationType in standard log record is populated with QueryString for unauthenticated requests as well
- [BZ - 2371110](#) - [rgw][server-access-logging]: aclRequired is not populated with yes for put/get requests involved with ACL
- [BZ - 2371418](#) - [Hotfix Build] Ceph Manager Daemon Crashes During Parallel NFS Deployment and Node Reboot in Hotfix Build

- [BZ - 2372311](#) - [rgw][server-access-logging]: please modify operation name for bucket operations like REST.GET.get_bucket_encryption to REST.GET.ENCRYPTION in the standard mode
- [BZ - 2372438](#) - [9.0] Parallel NFS daemon deploy processing does synchronous deletes during reschedules
- [BZ - 2372526](#) - UI does not accept plain ip:port for rgw_crypt_kmip_addr in SSE-KMS ? GKLM (KMIP) configuration , even though it is a valid format accepted via CLI
- [BZ - 2372675](#) - [10k osd] Prometheus configuration file is re-generated every time the mgr fails over
- [BZ - 2374255](#) - [Scale transition/Multiple clouds]: Incorrect cloud transition for a few objects in multi-cloud Ceph RGW
- [BZ - 2374279](#) - Ceph Dashboard raises a MTU mismatch alert when not expected
- [BZ - 2374369](#) - [NFS-Ganesha] NFS export created with same CephFS subvolume path on two different NFS clusters leads to inconsistent visibility and synchronization issues across clients
- [BZ - 2374412](#) - CVE-2025-52555 ceph: privilege escalation by unprivileged users in a ceph-fuse mounted CephFS
- [BZ - 2375000](#) - [RFE] Add Conditional Delete for DeleteObject and DeleteObjects
- [BZ - 2376165](#) - [RGW]: RGW Crash in thread 7f542bfd2640 thread_name:io_context_pool found with delete_obj operation
- [BZ - 2376756](#) - [9.0][NFS-Ganesha][BYOK] BYOK export allows enabling encryption on non-empty directory; junk entries shown on NFS mount for existing files and dirs
- [BZ - 2377090](#) - [NFS-Ganesha] Ganesha attempts to start twice on standby node after active node reboot in HA deployment, causing IO disruption
- [BZ - 2379131](#) - [rgw-multisite] [9.0] [IBM_Support] multisite replication is broken for a customer
- [BZ - 2379222](#) - [Versioned-Bucket]: Conditional delete with delete-object, deletes latest version of object irrespective of the condition provided
- [BZ - 2379275](#) - [Versioned-Bucket]: Crash in thread 7fbbf7813640 thread_name:io_context_pool seen while performing delete-object using version-id
- [BZ - 2379788](#) - [9.0] [NFS-Ganesha] [BYOK]After Client node reboot during IO, remount succeeds, IO resumed, but mount path becomes inaccessible
- [BZ - 2379790](#) - [9.0][NFS-Ganesha][BYOK] Git clone hangs indefinitely on BYOK-enabled NFS exports
- [BZ - 2379861](#) - [Versioned-Bucket]: Conditional delete using delete-objects removes the version of the object for which condition not matched
- [BZ - 2379872](#) - [Versioned-Bucket]: With delete-objects method, if we provide multiple condition along with version-id in the json, other condition is not respected
- [BZ - 2379914](#) - [Conditional-delete]: RGW debug log even with debug20 doesn't provide detailed error during failure of conditional delete with PreconditionFailed

- [BZ - 2380071](#) - [9.0][ceph-dashboard] grafana graphs displaying parse error: unexpected ", " in label matching
- [BZ - 2380400](#) - [CEE] mon_memory_target is ignored at startup when set without mon_memory_autotune in the config database
- [BZ - 2380412](#) - [IBM 8.1][SMB] .smb pool gets created automatically during initial bootstrap of ceph cluster
- [BZ - 2380443](#) - [Versioned-Bucket]: condition --if-none-match is not respected in case of version bucket conditional put
- [BZ - 2380738](#) - [Versioned-Bucket]: condition --if-match <ETAG> is not respected in case of version bucket conditional put
- [BZ - 2381576](#) - [RFE] Add the following API calls: iam:GetAccountSummary
- [BZ - 2381577](#) - [RFE] Add the following API calls: sts:GetCallerIdentity
- [BZ - 2381960](#) - [BYOK] [RFE] Cephadm NFS service spec does not support SNI (Server Name Indication) for mTLS
- [BZ - 2383180](#) - [9.0][MGR] Ceph MGR crashes with the assert: virtual ceph::common::RefCountedObject::~~RefCountedObject(): ceph_assert(nref == 0)
- [BZ - 2383217](#) - Mirrored images have incorrect info in the GUI
- [BZ - 2383548](#) - Ceph osd aborted in thread 7f71d41a9640 thread_name:safe_timer
- [BZ - 2385921](#) - [RFE][Cephadm][NFS-Ganesha] Cephadm support for SNI (Server Name Indication) for mTLS
- [BZ - 2385946](#) - [NFS-Ganesha][BYOK] NFS-Ganesha crashing while restarting the ganeshas service with 2 NFS exports having diff --kmip_key_id set
- [BZ - 2385951](#) - [NFS-Ganesha][BYOK] NFS-Ganesha BYOK crashed when Cthon04 Large Compile tests are executed
- [BZ - 2385952](#) - [NFS-Ganesha] - NFS Ganesha daemon crashes after upgrade
- [BZ - 2385959](#) - NFS daemons continuously provisioning and deprovisioning to avoid a port or conflict
- [BZ - 2385960](#) - QOS config not supported on nfs orch create
- [BZ - 2385962](#) - When all Ganesha instance fails to get KMIP key from all configured hosts, instead of failing to mount, File share is able to mount in plain text
- [BZ - 2386473](#) - MDS crashes when case sensitivity is disabled
- [BZ - 2386743](#) - [GSS] MGRs constantly failed over and interfering with upgrade due to frequent list_servers from cephadm
- [BZ - 2387075](#) - [rgw][server-access-logging]: some of the old log objects download failing with 404 NoSuchKey error, but head-object works.
- [BZ - 2387592](#) - [Ceph-Dashboard] improve the /mirroring/summary API
- [BZ - 2387764](#) - [RGW] Bulk delete on versioned bucket causes excessive OLH updates and high latency
- [BZ - 2388210](#) - cephadm fails to build on RHEL 10 in Brew
- [BZ - 2388477](#) - haproxy.cfg changes for active-active deployment of NFS Ganesha

- [BZ - 2389164](#)  - cephadm support for NFS BYOK
- [BZ - 2389280](#)  - [RFE] rgw: rate limit for LIST and DELETE ops
- [BZ - 2389764](#)  - [cephfs][RFE] subvolume quotas not passed to the nfs client
- [BZ - 2389907](#)  - CVE-2024-31884 pybind: Improper use of Pybind
- [BZ - 2390032](#)  - [GSS] MGRs constantly failed over and interfering with upgrade
- [BZ - 2390060](#)  - Fix snapdiff result fragmentation
- [BZ - 2390658](#)  - copy_object from versioning_suspended bucket to versioning-disabled bucket taking olh attributes along
- [BZ - 2390708](#)  - When removing gw service - the gws remain in mon db as zombies
- [BZ - 2391477](#)  - Ceph 9 fails to build on aarch64 with cmake error
- [BZ - 2391529](#)  - [IBM_Support] RGW returns false success [Code 200] when setting "rateLimit" using storage router endpoint
- [BZ - 2391692](#)  - last_purged_snaps_scrub updates can cause high frequency OSDMap updates leading to performance issues
- [BZ - 2392322](#)  - Forward porting commits from the 8.1 branch
- [BZ - 2392374](#)  - Need access to rbd mirror schedule interval through the API
- [BZ - 2392398](#)  - CG sync stopped after secondary is restarted
- [BZ - 2392604](#)  - rgw: checksum backports 9.0
- [BZ - 2393264](#)  - [9.0]mgr/dashboard: Remove the time dropdown from grafana iframe.
- [BZ - 2393415](#)  - rgw: fix enqueueing overwritten objects for gc -- defect introduced by not-shipped conditional write change
- [BZ - 2393440](#)  - [rgw][server-access-logging][missing-log-objects]: modifying the logging conf while the workload is running on the src bucket is causing the committed log object download failure after gc kicks in
- [BZ - 2393477](#)  - [8.1]The Ratelimit configuration from secondary site not getting sync to primary
- [BZ - 2393790](#)  - [RGW][STS] User ratelimits are not enforced when requests use STS assumed-role credentials
- [BZ - 2393857](#)  - Fix mypy issues of 9.0 branch
- [BZ - 2394062](#)  - [RFE][RGW][server-access-logging] allow committing of empty objects
- [BZ - 2394077](#)  - [9.0][add-host] While adding RHEL 8 host to the RHEL 9 admin node it fails with error "No module named 'dataclasses'"
- [BZ - 2394196](#)  - Image uploads to S3-compatible storage (RGW) fail due to boto3 checksum behavior changes
- [BZ - 2394366](#)  - [cephadm] Improve HW QAT setup experience with cephadm
- [BZ - 2394395](#)  - [GSS] ceph configuration option bdev_async_discard is removed
- [BZ - 2394511](#)  - rgw: fix copy object encryption returned attributes
- [BZ - 2394541](#)  - Nfs instance not creating properly
- [BZ - 2395169](#)  - RFE [Cephadm] [NFS-Ganesha]: Cephadm support for NFS-Ganesha TLS configuration

- [BZ - 2395210](#) - [server-access-logging] improve error message when log bucket fails in journal mode
- [BZ - 2396437](#) - Enable EC optimizations by default
- [BZ - 2396582](#) - [RFE] make mirror group enable asynchronous
- [BZ - 2396583](#) - [RFE] integrate the new GroupSnapshotNamespaceMirror::complete field
- [BZ - 2396637](#) - [RFE][rgw][Ratelimit-interval]: Prevent setting ratelimit interval to 0 due to non-recoverable IO blocking
- [BZ - 2397780](#) - [NFS-Ganesha] Unexpected restart of Ganesha services on active nodes during VIP failback after node reboot
- [BZ - 2397793](#) - [RFE] Cephadm. CertMGR promote to GA
- [BZ - 2398027](#) - monitoring: Migrate from promtail to grafana alloy in centralized logging
- [BZ - 2398033](#) - [ceph-dashboard]: mgr/dashboard: Tiering form - Placement Target in Advanced Section
- [BZ - 2398040](#) - when prometheus is not deployed in the cluster: getting 404 - Not Found Could not reach Prometheus's API error in dashboard
- [BZ - 2398042](#) - [dashboard]: Maintain same indentation in about panel
- [BZ - 2398047](#) - Provisioner UI: Expecting error on mismatched strings for "Enter new password" & "Confirm new password" fields on Reset password screen
- [BZ - 2398048](#) - [ceph-dashboard] mgr/dashboard: FS - Attach Command showing undefined for MountData
- [BZ - 2398058](#) - Local storage class creation via dashboard doesn't handle creation of pool and pool updation in zone configuration
- [BZ - 2398061](#) - Mismatch in the object size observed in Advance section of Tiering Storage Class and In the details section
- [BZ - 2398063](#) - [RGW-Dashboard]: ACLs Mapping should not be a mandatory setting while creating s3 storage class
- [BZ - 2398150](#) - [Ceph Common]: "ceph-diff-sorted" missing from package
- [BZ - 2399618](#) - [rbd-mirror] rbd command crashes if trying to remove a non-existing group
- [BZ - 2399622](#) - [RGW] [20.1.0] Segfault on one of 3 RGWs during ~20M-object workload; IO continued and RGW auto-restarted
- [BZ - 2399636](#) - Provisioner UI: Expecting some sort of progress information after clicking "Deploy ceph cluster"
- [BZ - 2399751](#) - [RGW][Lua] Multipart object uploaded using s3cmd having storage class "STANDARD"
- [BZ - 2400114](#) - [cee/sd][RGW-multisite] after manual bucket index reshard old bucket index objects are still present
- [BZ - 2400121](#) - [NFS-Ganesha][Active-Active HA] Post node reboot, I/O and basic commands on mount point remain stuck
- [BZ - 2400177](#) - [rgw][9.0]: rgw crashing while performing get-caller-identity on sts temporary user

- [BZ - 2400414](#) - [9.0][ceph-dashboard] Group similar alerts in UI
- [BZ - 2400427](#) - [Fast EC] Scrub errors on pools with EC enhancements enabled
- [BZ - 2400565](#) - [Tentacle] [Upgrade][RHEL10] Upgrade to latest build of Tentacle failing with mgr getting crash in SnapRealmInfoNew::decode()
- [BZ - 2400686](#) - Add mgmt-gateway and oauth2-proxy to the upgrade list + RGW SSL migration
- [BZ - 2400920](#) - [IBM_Support] Dashboard is grayed out post upgrade to 8.1
- [BZ - 2401174](#) - [GSS] [RGW]: segfault:
(RGWRados::remove_objs_from_index(DoutPrefixProvider const*, RGWBucketInfo&, std::__cxx11::list<rgw_obj_index_key, std::allocator<rgw_obj_index_key> > const&)+0x10bd) [0x5617746101bd]
- [BZ - 2401206](#) - [Tentacle] Ceph cluster bootstrap with RHEL 10 based container images on RHEL 9 host fails with error: ?PyO3 modules do not yet support subinterpreters?
- [BZ - 2401399](#) - [RGW-Dedup][20.1.0-30] radosgw-admin dedup stats aborts making the utility unusable
- [BZ - 2401685](#) - [RGW][Dedup-estimate][20.1.0-30]: md5_stats remain not populated for duplicate objects > 4 MB, dedup ratio stays 0
- [BZ - 2401750](#) - [RFE][RGW Dedup]: The cli dedup restart is misleading, instead use dedup process
- [BZ - 2401776](#) - HAProxy daemon should not be restarted on haproxy.cfg update for NFS Ganesha
- [BZ - 2401806](#) - logical volume DB devices specified in the OSD specification are not recognized or assigned during OSD deployment
- [BZ - 2401876](#) - Disable the availability tracking option by default, since it's a Tech Preview feature
- [BZ - 2402035](#) - [NFS-Ganesha] NFS Ganesha daemons are marked as ?stray? by cephadm after deployment using a spec file, changing cluster health from HEALTH_OK to HEALTH_WARN.
- [BZ - 2402054](#) - [9.0][rgw][server-access-logging]: almost all of the log-objects download failing with 404 after gc process them
- [BZ - 2402146](#) - [rgw] more than 100% performance dip observed with 1M small objects PUT workloads after an upgrade to 9.0
- [BZ - 2402468](#) - Communication between Internal services (Dashboard and Prometheus i.e) is failing when running mgmt-gateway in HA mode
- [BZ - 2402526](#) - RGW returns wrong HTTP status code for expired STS token
- [BZ - 2402528](#) - [GSS] The ceph RHCS Image does not contain tar which is complicating some scripting operations.
- [BZ - 2402656](#) - [RFE][ceph-dashboard] Rename RGW left navigation items aligned with S3 standards

- [BZ - 2402662](#) - [IBM_Support][CLOUDTIER] Transitioning ceph data to Amazon S3 is failing
- [BZ - 2402769](#) - [Tentacle] [Upgrade][RHEL10] Upgrade to latest build of Tentacle failing with mgr Crashes and cephadm modules unable to start with upgraded mgr daemons
- [BZ - 2403464](#) - [Ceph Dashboard]: can't enable mirroring while editing an RBD image
- [BZ - 2403487](#) - [Regression] The Subuser creation doesn't generating the secret keys automatically
- [BZ - 2403699](#) - [RGW]: Getting 500 - Internal Server Error, while expanding bucket section
- [BZ - 2403700](#) - [RGW]: Font and background color in edit tiering section makes it unreadable
- [BZ - 2403702](#) - [RGW]: Post setting tiering configuration/lifecycle configuration, Policy section in UI does not display lifecycle rule
- [BZ - 2403703](#) - [RGW]: Edit user via UI throwing multiple server errors
- [BZ - 2403705](#) - [RGW]: Edit bucket throws 500 server error and page gets stuck in loading form data
- [BZ - 2404088](#) - [RFE][ceph-dashboard] Add customizations to table and multi-site actions
- [BZ - 2404331](#) - [9.0] Bootstrap fails with error Failed to set orch backend to cephadm: maximum retries reached
- [BZ - 2404347](#) - Removing offline hosts using the command `ceph orch host rm <hostname> --force --offline` fails with the error "Attempted to access cert for unknown TLS object name cephadm-signed_node-exporter_cert"
- [BZ - 2404412](#) - RGW Crashes When Creating 1K Buckets with Dynamic Resharding Disabled
- [BZ - 2405377](#) - rgw-multisite: object deletion is not properly synced on versioning-suspended bucket
- [BZ - 2405397](#) - cephadm crashes and doesn't recover with ganesha-rados-grace tool failed: Failure: -126
- [BZ - 2405692](#) - mgr/dashboard: fix issues in SMB Overview grafana dashboards
- [BZ - 2405986](#) - [RGW][Dedup]: Observing segfault in thread_name:dedup_bg when doing 'dedup estimate and dedup stats'
- [BZ - 2405989](#) - mgr/dashboard: Fix Filesystem grafana dashboard units
- [BZ - 2406035](#) - [ceph-dashboard] unable to create bucket level sync policy from Object >> Multi-site >> Sync-policy
- [BZ - 2406182](#) - mgr/dashboard: Remove Multi-Cluster Tab from Dashboard
- [BZ - 2406510](#) - [ceph-dashboard][rgw] post configuring multisite, secondary zone created does not have endpoint configured for the same
- [BZ - 2406696](#) - [RGW-Module]:Endpoint is not getting added to the zone post realm bootstrap
- [BZ - 2406794](#) - [RDR] A few RBD images report error due to incomplete group snapshots on the secondary cluster after workload deployment


- [BZ - 2406837](#) - [RGW]GetBucketLocation fails in Ceph 8.1 when bucket owner assumes a role (ownership check skipped for non account users)
- [BZ - 2406973](#) - [NFS-Ganesha][TLS-Support]: Unable to enable TLS or mTLS without manually adding sectype as "sys" to export block
- [BZ - 2408795](#) - [9.0][CephAdm] Unable to set cepahdm signed cert for RGW service using certmgr enhancements
- [BZ - 2411930](#) - Per-bucket per-user counters not getting reflected in prometheus
- [BZ - 2411963](#) - Group replayer shutdown can hang in the face of active m_in_flight_op_tracker ops
- [BZ - 2412218](#) - RGW crash observed during copy-object ops with copy-source being empty
- [BZ - 2412220](#) - Object keys with a '+' symbol causes sync in multisite to enter into a loop
- [BZ - 2412238](#) - [9.0][Dashboard] oauth2-proxy service creation is stuck for a long time if improper values are passed for service creation
- [BZ - 2412273](#) - RGW crashes post zonegroup rename and period commit
- [BZ - 2412500](#) - [RGW-MS]: One secondary rgw daemon crash during initial multisite sync from primary with ~22M objects [RGWDataChangesLog::add_entry]
- [BZ - 2413062](#) - [RGW][Dedup]: Don't increment Corrupted_ETAG counters for DeleteMarkers
- [BZ - 2413276](#) - [cephadm] Ceph Prometheus and Dashboard modules crashes with : cherryPy ChannelFailures due to "Port <port> not free"
- [BZ - 2413522](#) - [Ceph 9.0] [NFS-Ganesha 7.0.0.6.5] NFS crashed and dumped core (ceph_fsal_release) while performing failover failback operations with HA - haproxy-protocol
- [BZ - 2413723](#) - RFE [Cephadm] [NFS-Ganesha]: Cephadm support for NFS-Ganesha TLS configuration adding new option xprtsec
- [BZ - 2413802](#) - [RGW-Dedup]: Dedup estimate connection times out while doing for many objects (~ 250K objects)
- [BZ - 2414044](#) - [RGW-MS]: RGW crash in data sync thread in RGWCoroutinesManager/RGWRemoteDataLog after remote data log shard failure.
- [BZ - 2414677](#) - Prometheus module error causing cluster to go to HEALTH_ERR
- [BZ - 2414841](#) - [IBM_Support] Application Pod stays in Init state as the CephFS VolumeAttachment doesn't complete.
- [BZ - 2415313](#) - Issue installing ceph-common package due to libcrypto.so.3 Dependency
- [BZ - 2415656](#) - [RGW][Dedup]: Dedup background thread crashes while processing 500K identical 5MB objects
- [BZ - 2415796](#) - [Fast EC] 2/3 mon daemons crashed with thread 7feeaaccd640 thread_name:safe_timer
- [BZ - 2415823](#) - [rhel-10] ceph 9 ubi-9 ceph-base client package installation fails on rhel-10
- [BZ - 2416554](#) - rbd-mirror daemon restart fails to resume partially synced demote snapshot synchronization on secondary

- [BZ - 2416700](#) - [RGW-LC] LC transition crash in LCOpRule::process after cloud tier region update (eu-west-1 to us-east-1) and period commit
- [BZ - 2416941](#) - Further upload to the bucket, will not increment per_bucket per_user counter, instead it displays only the last uploaded count/bytes
- [BZ - 2417210](#) - Update Images.py with Redhat Registry locations
- [BZ - 2417224](#) - Backport qa/workunits/rbd/rbd_groups.sh test fixup missed in BZ 2396583
- [BZ - 2417952](#) - [9.0] RGW segfault in RGWIndexCompletionManager::create_completion() via LC expiration on large versioned buckets
- [BZ - 2419075](#) - Multisite full sync misses a subset of small objects.
- [BZ - 2419090](#) - [Fast EC] EC pools with optimizations are being created with default 4K stripe_unit instead of expected 16K
- [BZ - 2419138](#) - [NFS-Ganesha] NFS crashed and dumped core (deconstruct_handle) while performing Cluster QOS tests with NFS service restart
- [BZ - 2419151](#) - [RBD Mirror] RBD group mirror replayer logs repeated ?remote group snapshot not found? messages after remote peer UUID is removed
- [BZ - 2419667](#) - Fast EC OSD key-not-found exception on truncate
- [BZ - 2419834](#) - Multiple OSDs errors due to Ceph assertion failures during EC4+2 FIO tests
- [BZ - 2420485](#) - ODF Internal mode cluster deployment fails with MGR HEALTH_ERR Module prometheus crash
- [BZ - 2420499](#) - [9.0][IBM] cephadm list-images contains redhat registry images
- [BZ - 2421370](#) - RBD mirror operation dumping core
- [BZ - 2421457](#) - tentacle: Revert "PrimaryLogPG: don't accept ops with mixed balance_reads and rworded flags
- [BZ - 2423160](#) - [RGW-MS]: Incremental sync of a bucket 'testbucket3' out of 5 buckets, is stuck after a site outage and ceph service restarts
- [BZ - 2423971](#) - osd: Fast EC incorrectly deletes object due to diverged log with partial writes
- [BZ - 2424287](#) - [9.0] [RGW-MS]: RGW LC goes FAILED when expiring delete markers on versioned bucket [testver1] and with error: (27) File too large
- [BZ - 2425002](#) - [Fast EC] In Fast EC pools, osd crashed with osd/PGLog.h: FAILED ceph_assert(missing.is_missing(hoid) && missing.get_items().at(hoid).need >= objiter->second->version)
- [BZ - 2425606](#) - [Fast EC] OSD crashed with error "assert_condition": "shard_size >= tobj_size" in Fast EC pools
- [BZ - 2427585](#) - [9.0 backport]erasure-code: Enable isa for ppc64le target.



CVEs


- [CVE-2023-25153](#)
- [CVE-2024-11831](#)
- [CVE-2024-31884](#)
- [CVE-2025-30204](#)
- [CVE-2025-52555](#)


References


- <https://access.redhat.com/security/updates/classification/#moderate> 


The Red Hat security contact is secalert@redhat.com. More contact details at <https://access.redhat.com/security/team/contact/>.




Quick Links 

Help 

Site Info 

Related Sites 

 Partial system outage



About Red Hat

Jobs

Events

Locations

Contact Red Hat

Red Hat Blog

Inclusion at Red Hat

Cool Stuff Store

Red Hat Summit

© 2026 Red Hat

[Privacy statement](#)

[Terms of use](#)

[All policies and guidelines](#)

[Digital accessibility](#)

[Cookie Preferences and Opt-Out Rights](#)