Rillor

EXPLAINER

Ceph S3 compatibility: what the RADOS Gateway supports

Ceph S3 API compatibility, read from the Umbrella 21.2.0 documentation: which Amazon S3 features the RADOS Gateway lists, what needs configuring, where checksums and identity differ, and how to test.



Rillor

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

Ceph's Object Gateway, also called the RADOS Gateway (RGW), serves a large subset of the Amazon S3 API, and what one gateway answers depends on its release line, its settings and the client's SDK.

- Ceph's Object Gateway (RGW) is an HTTP server built on librados that serves an interface compatible with "a large subset of the Amazon S3 RESTful API", plus a Swift interface in the same namespace. The feature table in Ceph's Umbrella 21.2.0 documentation has 28 rows: 27 marked Supported and Bucket Replication marked Partial, permitted only across zones [1, 2].
- Ceph's S3 feature table and Amazon's API reference count different things. Amazon lists 116 S3 actions, and one Ceph row, such as Multipart Uploads, covers several of them, so no percentage follows. The table also trails Ceph's release notes, which added S3 GetObjectAttributes in Tentacle 20.2.0 with no table row [2, 3, 19, 28].
- Several S3 features of Ceph's Object Gateway (RGW) depend on settings. Virtual-hosted-style bucket addressing needs `rgw_dns_name` or the zonegroup's host names and, with TLS, a wildcard certificate [5]. The Security Token Service (STS) is off by default and needs one key shared by every gateway in a zone [11, 21]. The static website APIs are off by default, and website endpoints need their own domain [5, 21].
- Ceph added user accounts in the style of AWS Identity and Access Management (IAM) in Squid 19.2, with an IAM API whose Umbrella tables list 52 actions. Tentacle 20.2.0 deprecated the older tenant-level IAM APIs in favor of accounts and listed them, with STS Lite (sts:GetSessionToken in Ceph's Security Token Service), for removal "no sooner than the V release" in Ceph's words [9, 10, 19].
- Ceph's Object Gateway (RGW) implements Amazon's three server-side encryption specifications, SSE-C, SSE-KMS and SSE-S3, the last with HashiCorp Vault only, and Umbrella adds authenticated AES-256-GCM. Amazon S3 has encrypted every new object upload automatically since 5 January 2023, applying SSE-S3 wherever a bucket has no other default encryption [12, 35].
- Current AWS SDKs send CRC checksums on uploads by default. Ceph 21.2.0's source marks five checksum types as AWS types, which cover every SDK default in AWS's table, out of the ten algorithms Amazon S3 lists. Ceph's tracker records two checksum failures: the Go SDK v2's checksum trailer, fixed in 21.2.0, and multipart uploads that name only CRC32C, with a fix under review on 8 November 2026 [22, 26, 27, 30, 31].
- The multipart defaults of Ceph's Object Gateway (RGW) at 21.2.0 match Amazon S3's limits: parts of 5 MiB to 5 GiB and up to 10,000 parts, which works out to a maximum object of about 48.8 TiB, the same as Amazon S3's (computed for this report). Its listing cap defaults to 5,000 entries, where Amazon returns at most 1,000 parts or uploads per request [21, 29].
- s3-tests, an S3 compatibility test suite that its own README calls "unofficial", defines 770 test functions at Ceph 21.2.0, with 15 decorated as failing on Ceph's Object Gateway (RGW) and 200 as failing on AWS. The markers record known differences and give no pass rate, and Squid point releases have changed S3 behavior, so a client is tested against the exact release it uses [20, 24, 25].

Background

What the Object Gateway is

The Ceph Object Gateway is "an object storage interface built on top of librados" [1]. It runs as an HTTP server, the radosgw daemon, between applications and the Ceph cluster, and it manages its own users [1]. It offers an S3-compatible interface, for "a large subset of the Amazon S3 RESTful API", and a Swift-compatible interface [1]. The two share one namespace, so data written through one API can be read through the other, and the gateway does not use the Ceph Metadata Server [1]. The S3 API page frames the scope as "compatible with the basic data access model of the Amazon S3 API" [2].

How Ceph documents compatibility

Ceph spreads the answer over several kinds of document. The S3 API page holds the feature table and the unsupported headers [2]. Operation pages cover authentication, common entities such as host names, bucket calls and object calls [4, 5, 6, 7]. Feature pages cover bucket policies, user accounts, the IAM API, STS, encryption, notifications, S3 Select, placement and S3 Control [8, 9, 10, 11, 12, 13, 14, 15, 16]. Configuration defaults sit in the source file `rgw.yaml.in` [21]. Release notes record changes per point release [18, 19, 20], the tracker records defects and backports [26, 27], and the s3-tests suite exercises the API [24, 25].

Release lines

Release line	First release	Point release read	Estimated end of life
Umbrella (21.2)	28 October 2026	21.2.0	1 June 2028
Tentacle (20.2)	18 November 2025	20.2.5, 2 November 2026	1 June 2027
Squid (19.2)	26 September 2024	19.2.7, 12 October 2026	2 December 2026

Table 1. Ceph's active release lines as listed on 8 November 2026, with end-of-life dates as Ceph's estimates [17, 19, 20]. Reef 18.2.8 reached end of life on 20 March 2026 [17].

This report is for engineers and architects who point S3 clients, SDKs or tools at a Ceph cluster and need to know what the gateway answers and where it has differed from Amazon S3.

Method

What was read. Ceph's Umbrella documentation for the Object Gateway and its S3 interface [1, 2, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16], Tentacle's S3 API page for comparison [3], the releases index and the Umbrella, Tentacle and Squid release notes [17, 18, 19, 20], the gateway's configuration options and checksum table in the source at tag v21.2.0 [21, 22], one merged pull request [23], the s3-tests README and functional tests vendored at v21.2.0 [24, 25], two tracker issues [26, 27] and eight Amazon Web Services pages on the API being matched [28, 29, 30, 31, 32, 33, 34, 35].

Version pinned. Every Ceph documentation page read is the Umbrella version and every source file is at tag v21.2.0, unless a sentence names another release, and every release-note statement names its release line or version. Where an Umbrella page and a release note disagree, both are given with their sources.

Counts. Feature table rows were counted on the Umbrella and Tentacle pages [2, 3]. Amazon's actions were counted from the action links on its API reference [28]. IAM actions, managed policies and bucket policy entries were counted from Ceph's tables [8, 10]. Test functions and markers were counted in `test_s3.py` at tag v21.2.0 [25]. Figures computed for this report are labelled.

Search check. Google autocomplete, checked on 8 November 2026, completes "ceph s3" and "ceph s3 compatibility" with "ceph s3 api compatibility", and offers "ceph rgw s3 compatibility", "ceph object gateway s3 api" and "ceph s3 tests" beside it.

All pages were viewed on 8 November 2026, and tracker statuses are as viewed that day.

Findings

1. The feature table: 28 rows, one partial

Umbrella's Features Support table has 28 rows, counted for this report [2]. 27 are marked Supported [2]. Bucket Replication is marked Partial, with the remark "Permitted only across zones" [2]. The remarks on Create Bucket and Bucket ACLs note a different set of canned ACLs, Policy (Buckets, Objects) notes ACLs and bucket policies, Bucket Notification points to the notification compatibility page, Object Tagging points to the object operations for policy verbs, Storage Class points to the storage classes page, and Account PublicAccessBlock points to the S3 Control page [2].

Group	Rows in the Umbrella table
Buckets (15)	List Buckets, Create Bucket, Delete Bucket, Get Bucket Info (HEAD), Bucket Location, Bucket Lifecycle, Bucket Replication (Partial), Bucket Website, Bucket ACLs, Policy (Buckets, Objects), Bucket Notification, Bucket Object Versions, Bucket Request Payment, Bucket Tagging, Bucket Logging
Objects (10)	Put Object, Get Object, Delete Object, Get Object Info (HEAD), POST Object, Copy Object, Object ACLs, Object Tagging, Multipart Uploads, Storage Class
Ownership and public access (3)	Object Ownership, Bucket PublicAccessBlock, Account PublicAccessBlock

Table 2. The rows are Ceph's [2]. The grouping is this report's.

Tentacle's table has 25 rows, 24 Supported and Bucket Replication Partial, and lacks the three ownership and public-access rows [3]. The Umbrella page lists one unsupported header field, `x-amz-id-2`, a response header [2].

2. The table, Amazon's action list and the release notes

Amazon's S3 API reference lists 116 actions, counted on 8 November 2026 [28]. They include CreateSession, ListDirectoryBuckets, RenameObject, the bucket metadata table configurations, inventory, analytics, intelligent-tiering and SelectObjectContent [28]. A Ceph row and an Amazon action are different units. Multipart Uploads alone spans several actions, so 28 rows against 116 actions gives no meaningful percentage (this report's reasoning). A missing row says what the table lists, and the gateway's behavior for that call is a question for the operation pages, the release notes and a test.

The table trails the release notes. Tentacle's v20.2.0 notes say "Added support for S3 GetObjectAttributes", and neither the Tentacle nor the Umbrella table has a GetObjectAttributes row [2, 3, 19]. Umbrella's notes list RGW changes that the table does not name: GetAccountSummary, admin REST APIs for account quota, STS

GetCallerIdentity, global CORS rule support, a rate-limit response with Retry-After and timestamp-based epochs for versioned objects, beside the account-wide PublicAccessBlock and S3 Object Ownership that do have rows [18].

Tentacle v20.2.0 also changed two responses. For compatibility with AWS S3, LastModified timestamps are truncated to the second, and during an upgrade they may appear to move backwards. HeadBucket reports X-RGW-Bytes-Used and X-RGW-Object-Count only when the read-stats query string is set [19].

3. Signatures and access control

RGW authenticates S3 requests with AWS signatures derived from the user's access key and secret key [4]. For Signature Version 4, the page lists five accepted values of `x-amz-content-sha256`: the actual payload checksum, `UNSIGNED-PAYLOAD`, `STREAMING-UNSIGNED-PAYLOAD-TRAILER`, `STREAMING-AWS4-HMAC-SHA256-PAYLOAD` and `STREAMING-AWS4-HMAC-SHA256-PAYLOAD-TRAILER` [4]. The trailer modes matter for the checksum behavior in Finding 9. For Signature Version 2, Ceph's note says v2 signatures are deprecated in AWS and that RGW continues to support them [4]. In OpenStack deployments, `rgw s3 auth use keystone = true` makes Keystone the authority for S3 authentication, with EC2-style credentials [4].

Access control keeps S3's ACL model: canned ACLs and grants of READ, WRITE, READ_ACP, WRITE_ACP and FULL_CONTROL mapped to S3 operations [4]. PUT Bucket accepts the canned values private, public-read, public-read-write and authenticated-read [6].

Signature and token code has carried security fixes. Squid v19.2.6 and Tentacle v20.2.4, both released on 19 August 2026, fixed CVE-2026-39944, a privilege escalation through CBC bit flips in STS session tokens, and CVE-2026-54330, a header injection in the SigV4 verifier [19, 20].

4. Addressing, DNS and bucket names

Ceph's common entities page describes two addressing modes, path-style and virtual-hosted-style. It says "Most S3 clients default to virtual-hosted-style access" and that AWS has deprecated path-style [5]. To answer virtual-hosted-style requests, RGW needs `rgw_dns_name`, which takes several names separated by commas, or the name added to the zonegroup's host names [5]. With SSL, certificates need a wildcard to match bucket subdomains, and behind a proxy the proxy's DNS name is the one to set [5]. The static website endpoint, s3website, must be served under a different domain, set in `rgw_dns_s3website_name` [5], and the static website APIs are off by default, with `rgw_enable_static_website` false [21].

AWS's own schedule moved. Its post of 8 May 2019 set out a plan in which path-style continued for buckets created on or before 30 September 2020 and later buckets used virtual-hosted references. An update of 23 September 2020 delayed the deprecation [33].

Bucket names follow default rules of 3 to 63 characters of lowercase letters, numbers, hyphens and periods, not in the form of an IP address, with labels that start and end with a letter or number, and unique [6]. With `rgw_relaxed_s3_bucket_names` set to true, names may use letters, numbers, periods, dashes and

underscores up to 255 characters, still unique and not an IP address [6]. The option defaults to false [21].

5. Identity: users, tenants, accounts, IAM and STS

Squid introduced User Accounts, an optional feature for self-service management of users, groups and roles "similar to those in AWS Identity and Access Management (IAM)" [9]. An administrator creates each account's root user with radosgw-admin or the Admin Ops API, and the root user's keys work with the IAM API to create users and roles for the S3 API [9]. Resources that account users create belong to the account, so usage statistics and quotas apply to the account as a whole [9]. Account IDs are "RGW" followed by 17 digits and appear in IAM ARNs [9]. Account users start with no permissions, while users outside accounts can create buckets and upload by default [9].

The IAM API, new in Squid, is a subset of Amazon IAM served from the same HTTP endpoint as S3, and by default only account root users may call it [10]. Its tables list 52 actions: 16 for users, 15 for groups, 16 for roles, 4 for OpenID Connect providers and 1 for accounts, GetAccountSummary [10]. Six managed policies are listed: IAMFullAccess, IAMReadOnlyAccess, AmazonSNSFullAccess, AmazonSNSReadOnlyAccess, AmazonS3FullAccess and AmazonS3ReadOnlyAccess [10]. Squid's notes add that existing users can be adopted into accounts, which is optional but irreversible [20].

Tentacle's v20.2.0 notes deprecate tenant-level IAM in favor of accounts. The features they list for removal "no sooner than the V release" are tenant-level IAM APIs (CreateRole, PutRolePolicy, PutUserPolicy), tenant names in place of accounts in policy, policy without cross-account evaluation, cross-tenant bucket names such as tenant:bucket, and STS Lite with sts:GetSessionToken. Tenants remain for namespace isolation [19].

Ceph implements a subset of STS: AssumeRole, AssumeRoleWithWebIdentity, for OpenID Connect or OAuth 2.0 providers with Keycloak tested, and GetCallerIdentity [11]. DurationSeconds defaults to 3,600 seconds, and a session policy can narrow temporary credentials [11]. STS needs `rgw_s3_auth_use_sts`, false by default, and one `rgw_sts_key` shared by every gateway in a zone and every zone in a realm [11, 21].

6. Policies, ownership and public access

The bucket policy page says RGW "supports a subset of the Amazon S3 policy language applied to buckets" [8]. Its Limitations section lists 50 entries as supported actions, two of which, s3:IPAddress and s3:NotIpAddress, read as condition operators rather than actions (this report's reading), and tagging actions sit in a separate table of object operations [8]. The supported global condition keys are aws:CurrentTime, aws:EpochTime, aws:PrincipalType, aws:Referer, aws:SecureTransport, aws:SourceIp, aws:UserAgent and aws:username [8]. Tentacle 20.2.0 added ArnEquals and ArnLike conditions with their Not and IfExists variants, replication permission checks with s3:ReplicateObject and related actions, `x-amz-confirm-remove-self-bucket-access` on PutBucketPolicy, and RestrictPublicBuckets for PublicAccessBlock [19].

Ceph's pages disagree on who can carry a policy. The bucket policy page says policies cannot be set on users, groups or roles, and that RGW uses the tenant in place of a 12-digit account ID [8]. The User Accounts page, which describes the newer feature, documents iam:PutUserPolicy, group and role policies, and RGW

account IDs [9].

Ownership defaults differ. In Amazon S3, Object Ownership defaults to Bucket owner enforced, so ACLs are disabled for newly created buckets, and AWS recommends keeping them disabled except where access must be controlled per object [32]. Umbrella adds an Object Ownership row to the table, its notes list "S3 Object Ownership (disable ACLs)", and PUT Bucket accepts `x-amz-object-ownership` [2, 6, 18]. S3 Control, new in Umbrella, is a subset of the AWS S3 Control API for PublicAccessBlock on user accounts, applying to every bucket the account owns, and when both levels are set the more restrictive wins. It is enabled by default through `rgw_enable_apis` [16]. awscli and some SDKs put the account ID in the S3 Control host name, as in rgw69573912842483864.s3.example.com, so the page advises wildcard DNS and, with SSL, a matching wildcard certificate. RGW itself does not require the account in the host name [16].

7. Versioning, object lock and Ceph extensions

Versioning takes the states Enabled and Suspended, set by the bucket owner, and a bucket where it was never set has no versioning state [6]. Object lock is requested at creation with `x-amz-bucket-object-lock-enabled`, and its configuration takes Mode GOVERNANCE or COMPLIANCE with Days or Years [6]. Retention and legal hold have their own object operations [7].

The sources disagree on existing buckets. Umbrella's bucket operations page says the object lock configuration call fails with 409 InvalidBucketState if object lock was not enabled at bucket creation [6]. The Squid v19.2.3 notes of 28 July 2025 and the Tentacle v20.2.0 notes say PutObjectLockConfiguration can enable Object Lock on an existing versioning-enabled bucket created without it [19, 20]. AWS requires versioning and Object Lock before objects are locked, and once enabled, Object Lock cannot be disabled and versioning cannot be suspended [34]. AWS's configuration guide also covers enabling Object Lock on an existing bucket, variable retention periods and releasing an event hold [34]. Ceph's object lock request entities name the two modes with Days or Years, and the page does not mention variable retention periods or event holds [6].

Ceph adds Append Object. A PUT with `?append&position=` creates an appendable object and returns `x-rgw-next-append-position`. It cannot be used on a bucket whose versioning is enabled or suspended, and compression and encryption are disabled for appendable objects [7]. Conditional headers are documented for reads (if-match, if-none-match, if-modified-since, if-unmodified-since), for copies (`x-amz-copy-if-*`) and for deletes (`x-amz-delete-if-unmodified-since`) [7]. 27 tests in `test_s3.py` carry the conditional_write marker [25]. Point releases have fixed conditional writes: Squid v19.2.7, of 12 October 2026, "Fixed conditional Delete, MultiDelete, and Put handling", and Tentacle v20.2.5, of 2 November 2026, fixed conditional writes on buckets with versioning suspended [19, 20].

8. Encryption

RGW supports server-side encryption with three key-management options. Requests must use HTTPS, or `rgw trust forwarded https` behind a TLS proxy, and keys are 256-bit, in base64 [12].

Mode	Specification followed	Key source on RGW
SSE-C	Amazon's SSE-C	The client's own key, with no special Ceph configuration
SSE-KMS	Amazon's SSE-KMS	Barbican, Vault or KMIP
SSE-S3	Amazon's SSE-S3	HashiCorp Vault only

Table 3. Server-side encryption on RGW, from Ceph's Umbrella encryption page [12].

PutBucketEncryption, GetBucketEncryption and DeleteBucketEncryption manage a bucket's SSE-S3 or SSE-KMS default, and the `rgw crypt default encryption key` option is for testing only [12].

Umbrella adds a choice of algorithm. AES-256-CBC is the legacy mode, the default for backward compatibility, with no built-in integrity check. AES-256-GCM is authenticated and recommended for new deployments [12, 21]. GCM adds a 16-byte tag to each 4 KB chunk, about 0.4% overhead, and S3 responses report the plaintext size [12]. Ceph advises keeping CBC until every gateway is upgraded [12]. Tentacle v20.2.5 also gained AES-256-GCM, and Umbrella's notes add "CopyObject now supported for encrypted objects" [18, 19].

Amazon S3's default differs. Since 5 January 2023, every new object upload to Amazon S3 is encrypted automatically with SSE-S3 (AES-256), for new buckets and for existing buckets without default encryption [35].

9. Checksums and current AWS SDKs

AWS's SDK guide says the SDKs used to make integrity checks opt-in and now calculate CRC-based checksums by default on uploads, which Amazon S3 validates and stores [31]. For SDK or CLI versions from before December 2024, S3 computes a CRC64NVME checksum itself [31]. The settings are `request_checksum_calculation` and `response_checksum_validation` in the shared config file, or `AWS_REQUEST_CHECKSUM_CALCULATION` and `AWS_RESPONSE_CHECKSUM_VALIDATION` in the environment. Both default to `WHEN_SUPPORTED`, and `WHEN_REQUIRED` computes a request checksum only when the operation requires it [31]. AWS's table gives CRC64NVME as the default for AWS CLI v2 and the C++ SDK, and CRC32 for the Go V2, Java 2.x, JavaScript 3.x, Kotlin, .NET, PHP, Python (Boto3), Ruby, Rust, Swift and PowerShell SDKs [31].

Algorithm	In Amazon S3's list [30]	Marked as an AWS type in Ceph v21.2.0 [22]
CRC64NVME	Yes, the default	Yes

Algorithm	In Amazon S3's list [30]	Marked as an AWS type in Ceph v21.2.0 [22]
CRC32	Yes	Yes
CRC32C	Yes	Yes
SHA1	Yes	Yes
SHA256	Yes	Yes
MD5	Yes, supplied with `x-amz-checksum-md5` and not computed automatically by the SDKs	No
XXHASH64, XXHASH3, XXHASH128	Yes	No
SHA512	Yes	No

Table 4. Checksum algorithms. Ceph's v21.2.0 table also holds XXH3, SHA512 and BLAKE3 without the AWS flag [22].

Every SDK default in AWS's table is among the five types Ceph 21.2.0 marks as AWS checksums (compared for this report from [22, 31]). Pull request 71484, "rgw: support new checksum types to match aws s3", merged into Ceph's main branch on 8 September 2026, changing `rgw_cksum.h` and adding tests to `test_s3.py` [23]. It is not part of the v21.2.0 tag, whose table shows five AWS types [22].

Issue 71183, opened on 3 May 2025, reports that the Go SDK v2's config package from v1.29.0 adds new headers and a trailer to upload requests, and that RGW logs "failed to authorize request" and returns 400 InvalidArgument [26]. Its fix, pull request 63326, is recorded as released in v21.2.0, with backports to Tentacle and Squid marked and the issue in status Backporting on 8 November 2026 [26]. Umbrella's changelog lists it as "rgw: recognize checksum from x-amz-checksum-{type} alone" [18].

Issue 79572, opened on 18 August 2026 and updated on 29 October 2026, is marked a regression. After an upgrade from Squid, which had no multipart checksum validation, to Tentacle or main, multipart uploads that send only `x-amz-checksum-algorithm` CRC32C fail CompleteMultipartUpload with 400 BadDigest [27]. AWS defaults CRC32 and CRC32C multipart checksums to COMPOSITE when `x-amz-checksum-type` is omitted, and CRC64NVME is FULL_OBJECT only, while Ceph treats CRC32 and CRC32C as full-object [27]. Backports are listed for Tentacle, Squid and Umbrella, the fix is pull request 71127, and the status read Fix

Under Review on 8 November 2026 [27]. At v21.2.0, 12 tests in `test_s3.py` carry the checksum marker, among them CRC64NVME, SHA256, CRC32, CRC32C and SHA1 multipart helpers and a GetObjectAttributes checksum test [25].

10. Storage classes, notifications and S3 Select

Placement targets choose a bucket's pools when it is created and cannot be changed. Storage classes, added in Nautilus, place object data, every placement target has STANDARD, and lifecycle rules transition objects between classes. A client chooses a class with the `X-Amz-Storage-Class` header [15]. With AWS SDKs such as boto3, class names should match AWS's, or the SDK drops the request and raises an exception [15]. Some clients expect AWS behavior for names starting with GLACIER and fail against RGW, so Ceph advises other names, including INTELLIGENT-TIERING, STANDARD_IA, REDUCED_REDUNDANCY and ONEZONE_IA. Ceph accepts custom names such as CHEAPNDEEP, and some clients may refuse them [15].

AWS sends bucket notifications to SNS, SQS and Lambda. RGW sends them to HTTP/S, Kafka and AMQP endpoints represented with SNS ARNs, and does not support the QueueConfiguration and CloudFunctionConfiguration tags, treating every destination as SNS [13]. Ceph extensions delete one notification or all of them, fetch one by name, and filter by regular expression, metadata and object tags, which need SDK extensions [13]. The event record's `requestParameters.sourceIPAddress` is sent empty, and the topic API takes only CreateTopic, DeleteTopic and ListTopics from SNS, with extensions [13]. ObjectCreated and ObjectRemoved types are supported, lifecycle, sync and restore events are listed with several marked "Ceph extension", and s3:ReducedRedundancyLostObject is marked "Not applicable" [13].

S3 Select runs SQL-like queries inside RGW over an object, pushing the filter to the gateway. The page describes CSV, JSON and Parquet readers sharing one SQL engine, says S3 Select "has a definite set of functionalities compliant with AWS", and says RGW runs the query on each fetched chunk of up to 4 MB [14].

11. Size and listing limits

Limit	RGW default at v21.2.0 [21]	Amazon S3 [29]
Part size, upper bound	5 GiB (`rgw_max_put_size`)	5 GiB
Part size, lower bound	5 MiB (`rgw_multipart_min_part_size`)	5 MiB, no minimum for the last part
Parts per upload	10,000 (`rgw_multipart_part_upload_limit`)	10,000, numbered 1 to 10,000
Maximum object	5 GiB x 10,000, about 48.8 TiB (computed)	48.8 TiB
Parts per ListParts	Up to 5,000 (`rgw_max_listing_results`)	1,000

Limit	RGW default at v21.2.0 [21]	Amazon S3 [29]
Uploads per ListMultipartUploads	Up to 5,000 (`rgw_max_listing_results`)	1,000

Table 5. Upload and listing limits. The maximum object for RGW is computed for this report from [21] and matches the figure in [29].

`rgw_max_put_size` sets the limit for a regular, non-multipart upload and also caps each part, and the maximum composite object is that value times `rgw_multipart_part_upload_limit` [21], 50,000 GiB at the defaults (computed for this report). `rgw_max_listing_results` accepts 1 to 100,000 and caps max-keys for ListObjects and ListObjectVersions as well as max-uploads and max-parts [21]. `rgw_delete_multi_obj_max_num` defaults to 1,000 objects per multi-object delete request [21].

12. Point releases that changed S3 behavior

Squid v19.2.2, of 10 April 2025, was a hotfix for an RGW data-loss bug when CopyObject copies an object onto itself, which S3 clients do to change metadata, and its notes say not to upgrade to v19.2.1 [20]. A SigV4 regression in 19.2.6 and 20.2.4 rejected requests with an unsigned Content-Type and broke presigned PUT uploads. Umbrella's notes record the fix and that `x-amz-content-sha256` is no longer required in SignedHeaders, and Squid v19.2.7 and Tentacle v20.2.5 carry the same fixes [18, 19, 20]. Tentacle v20.2.5 also matches CORS rules the way Amazon S3 matches them, adds a global CORS rule and makes S3 GetUsage read the user usage log [19]. Squid v19.2.6 and Tentacle v20.2.4 fixed the two CVEs in Finding 3 [19, 20].

Two facts frame the choice of release line. Ceph's releases page gives Squid's estimated end of life as 2 December 2026 [17]. Umbrella's container images are based on Rocky Linux 10 and need a CPU at the x86-64-v3 level [18], which matters for a site that moves to Umbrella for the checksum fix in Finding 9.

13. Testing with s3-tests

s3-tests describes itself as "a set of unofficial Amazon AWS S3 compatibility tests" for software that exposes an S3-like API, using Boto2 and Boto3 [24]. Some tests carry attributes based on reliability and on AWS not enforcing its spec strictly, and runs can filter on them, for example on not fails_on_aws [24]. Umbrella vendored the suite into the Ceph repository under `src/test/rgw/s3-tests`, and Ceph's QA suites run it from there [18]. The STS tests need `rgw_sts_key` and `rgw_s3_auth_use_sts=true` on the test cluster [24].

Marker in `test_s3.py` at v21.2.0	Tests
All test functions	770
fails_on_aws	200

Marker in `test_s3.py` at v21.2.0	Tests
conditional_write	27
fails_on_rgw	15
checksum	12

Table 6. Test functions and markers, counted for this report from [25].

The comments beside several fails_on_rgw tests point to tracker issues or to TODOs in the tests themselves [25]. The markers record known differences between the suite, AWS and RGW. They give no pass rate, and this report computes none.

Implications

Choose the release line for the S3 features in use. Account-wide PublicAccessBlock and Object Ownership appear in Umbrella's table, AES-256-GCM in Umbrella and Tentacle v20.2.5, and the Go SDK checksum fix is recorded as released in 21.2.0 with Tentacle and Squid backports marked on 8 November 2026 [2, 12, 19, 26]. Ceph estimates Squid's end of life at 2 December 2026 [17], and Umbrella's images need x86-64-v3 CPUs [18].

Put DNS and certificates into the design. Virtual-hosted-style buckets need `rgw_dns_name` or zonegroup host names and a wildcard certificate, S3 Control clients may put account IDs in host names, and website endpoints need a domain of their own [5, 16].

Build new identity designs on accounts. Accounts carry the IAM API, and Tentacle deprecated tenant-level IAM and STS Lite [9, 10, 19]. Squid's notes call adopting existing users into accounts irreversible [20]. STS needs one shared key across the gateways of a zone [11].

Choose the encryption back end before data lands. SSE-S3 on RGW means HashiCorp Vault, and SSE-KMS means Barbican, Vault or KMIP [12]. Amazon S3 encrypts new uploads automatically [35], and on RGW a bucket default is set through PutBucketEncryption [12]. A move from CBC to GCM waits until every gateway is upgraded [12].

Check the SDK checksum settings. Clients on current SDKs send CRC32 or CRC64NVME by default [31]. `WHEN_REQUIRED` limits request checksums to the operations that require them [31], and multipart uploads with CRC32 or CRC32C need a test of their own on the release in use [27].

Use storage class names that SDKs accept. Ceph advises AWS names and advises against GLACIER prefixes [15].

Read Ceph's conflicts against the running version. Object lock on existing buckets and policies on users, groups and roles are described differently on different Umbrella pages and in the Squid and Tentacle notes [6, 8, 9, 19, 20]. A test on the release in use settles which description applies.

Test before moving data. Run the client's own SDK, versions, settings and operations against the exact Ceph release, with the relevant s3-tests markers, and repeat after each point release, since Squid and Tentacle point releases have changed copy, signature, CORS and conditional-write behavior [18, 19, 20, 24, 25].

Limits of this analysis

The report reads documentation, source files and tracker pages. No gateway was run, so every statement is about documented behavior at the versions named.

The releases index and release-notes pages carry the docs site's banner "This document is for a development version of Ceph", although 21.2.0 was released on 28 October 2026 [17, 18]. Ceph's documentation pages carry no publication dates, and Amazon's pages are current versions, also undated.

Tracker statuses, backport fields and pull request states are as viewed on 8 November 2026 [23, 26, 27]. Which point releases carry the Squid and Tentacle backports of the fix in issue 71183, and when issue 79572 closes, are outside what the sources say.

Counts of rows, actions, policy entries and test markers were made by hand from the pages and files named. A test marker is a different unit from a test result.

Where Ceph's pages disagree, the report gives both and does not settle which one matches the code.

The report measures no performance and compares no throughput or latency.

Sources

1. Ceph Object Gateway. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/

2. Ceph Object Gateway S3 API. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3/

3. Ceph Object Gateway S3 API. Ceph Documentation (Tentacle), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/tentacle/radosgw/s3/

4. Authentication and ACLs. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3/authentication/

5. Common Entities. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3/commons/

6. Bucket Operations. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3/bucketops/

7. Object Operations. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3/objectops/

8. Bucket Policies. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/bucketpolicy/

9. User Accounts. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/account/

10. Ceph Object Gateway IAM API. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/iam/

11. STS in Ceph. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/STS/

12. Encryption. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/encryption/

13. S3 Bucket Notifications Compatibility. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3-notification-compatibility/

14. Ceph s3 select. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3select/

15. Pool Placement and Storage Classes. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/placement/

16. S3 Control. Ceph Documentation (Umbrella), Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/umbrella/radosgw/s3/s3control/

17. Ceph Releases (index). Ceph Documentation, Ceph Foundation, viewed 8 November 2026. https://docs.ceph.com/en/latest/releases/

18. Umbrella (v21.2.0 release notes). Ceph Documentation, Ceph Foundation, 28 October 2026, viewed 8 November 2026. https://docs.ceph.com/en/latest/releases/umbrella/

19. Tentacle (release notes, v20.2.0 to v20.2.5). Ceph Documentation, Ceph Foundation, 18 November 2025 to 2 November 2026, viewed 8 November 2026. https://docs.ceph.com/en/latest/releases/tentacle/

20. Squid (release notes, v19.2.0 to v19.2.7). Ceph Documentation, Ceph Foundation, 26 September 2024 to 12 October 2026, viewed 8 November 2026. https://docs.ceph.com/en/latest/releases/squid/

21. rgw.yaml.in (RGW configuration options), tag v21.2.0. Ceph project, GitHub, viewed 8 November 2026. https://raw.githubusercontent.com/ceph/ceph/v21.2.0/src/common/options/rgw.yaml.in

22. rgw_cksum.h, tag v21.2.0. Ceph project, GitHub, viewed 8 November 2026. https://raw.githubusercontent.com/ceph/ceph/v21.2.0/src/rgw/rgw_cksum.h

23. rgw: support new checksum types to match aws s3, pull request 71484. Ceph project, GitHub, merged 8 September 2026, viewed 8 November 2026. https://github.com/ceph/ceph/pull/71484

24. S3 compatibility tests, README, vendored at tag v21.2.0. Ceph project, GitHub, viewed 8 November 2026. https://raw.githubusercontent.com/ceph/ceph/v21.2.0/src/test/rgw/s3-tests/README.rst

25. test_s3.py, s3-tests functional suite, tag v21.2.0. Ceph project, GitHub, viewed 8 November 2026. https://raw.githubusercontent.com/ceph/ceph/v21.2.0/src/test/rgw/s3-tests/s3tests/functional/test_s3.py

26. Golang AWS SDK now adds new checksumming trailers, causing most requests to fail, issue 71183. Ceph issue tracker, opened 3 May 2025, viewed 8 November 2026. https://tracker.ceph.com/issues/71183

27. rgw: multipart CRC32/CRC32C defaults to FULL_OBJECT instead of AWS COMPOSITE, issue 79572. Ceph issue tracker, opened 18 August 2026, updated 29 October 2026, viewed 8 November 2026. https://tracker.ceph.com/issues/79572

28. Amazon S3 (actions supported by Amazon S3). Amazon Simple Storage Service API Reference, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/AmazonS3/latest/API/API_Operations_Amazon_Simple_Storage_Service.html

29. Amazon S3 multipart upload limits. Amazon S3 User Guide, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/AmazonS3/latest/userguide/qfacts.html

30. Checking object integrity in Amazon S3. Amazon S3 User Guide, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/AmazonS3/latest/userguide/checking-object-integrity.html

31. Data Integrity Protections for Amazon S3. AWS SDKs and Tools Reference Guide, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/sdkref/latest/guide/feature-dataintegrity.html

32. Controlling ownership of objects and disabling ACLs for your bucket. Amazon S3 User Guide, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/AmazonS3/latest/userguide/about-object-ownership.html

33. Amazon S3 Path Deprecation Plan: The Rest of the Story. Jeff Barr, AWS News Blog, Amazon Web Services, 8 May 2019 (updated 23 September 2020), viewed 8 November 2026. https://aws.amazon.com/blogs/aws/amazon-s3-path-deprecation-plan-the-rest-of-the-story/

34. Configuring S3 Object Lock. Amazon S3 User Guide, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-configure.html

35. Default encryption FAQ. Amazon S3 User Guide, Amazon Web Services, viewed 8 November 2026. https://docs.aws.amazon.com/AmazonS3/latest/userguide/default-encryption-faq.html

About this report

Published by Rillor on 14 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. Ceph S3 compatibility: what the RADOS Gateway supports. 14 November 2026. https://rillor.com/insights/ceph-s3-compatibility

Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

NVIDIA, AMD, Ceph, Kubernetes and related product names are trademarks of their owners and are used only to identify products. No affiliation or endorsement is implied.

Work with Rillor

Send what the storage must hold, how fast it grows, how it is read and written, and where it runs. Contact us at rillor.com/contact?type=storage or contact@rillor.com. More on distributed storage: rillor.com/compute/storage