

Clouidian HyperStore 8.2.8 Release Notes (2026-07-13)

This is an Early Access (EA) release. The EA classification is given to each feature release of Clouidian software when it is initially posted to the FTP site. This classification indicates that the product has completed a comprehensive development and release procedure including extensive testing for functionality, reliability, and full regression. The release has fully met the Clouidian criteria for production usage by customers if desired. This release may be promoted to GA at a later date. In that case, it will be declared at the Clouidian Support Portal as such.

Note: HyperStore File Services and HyperStore Search Services are not supported on HyperStore 8.2.8. Do not upgrade HyperStore to 8.2.8 if you have File or Search services in your cluster.

Supported OS

- Rocky Linux 8.9
- Rocky Linux 8.10 (recommended)
- Rocky Linux 9.8
- RHEL 8.7, 8.8, 8.9
- RHEL 8.10 (recommended)

HyperStore 8.2 or later cannot be run on CentOS systems.

Installation Notes

NOTE: Starting in 8.2.5, CMC will by default connect to S3 via https (s3.tls.required=true). This is for new installations and upgrades where the 's3.tls.required' config setting has not been changed from the prior default. If your environment requires CMC to connect to S3 via http (not https), run the following commands prior to upgrade:

```
$ hsctl config get s3.tls.required
false
$ hsctl config set s3.tls.required=false
$
```

- For a new install of, or upgrade to, HyperStore 8.2.8 your system must be running Rocky Linux or Red Hat Enterprise Linux (RHEL) 8.9 or 8.10 (recommended).
- If you want to install HyperStore on a later version of Rocky Linux or RHEL please contact Clouidian Support. HyperStore does not support other types of Linux distribution, or non-Linux operating systems.
- If you want to install HyperStore on a QLC system, please see the Clouidian HyperStore Installation Guide.

Upgrade Notes

Supported versions to upgrade from:

- 8.1
- 8.1.x or an 8.1.x patch
- 8.2 or an 8.2.0.x patch
- 8.2.1 or an 8.2.1.x patch
- 8.2.x or an 8.2.x.y patch

If you are running an older version, please contact Support for upgrade information.

Before upgrading to this release we recommend your cluster has 256GB RAM or greater. Contact your Clouidian representative for information about upgrading existing nodes to meet this requirement.

When upgrading from a HyperStore system that had previously received a `shutdownQoS` via JMX for Cassandra, S3 or both, then to preserve this configuration perform the following before upgrade:

- Within the 8.2.5 staging directory, copy file `sample-migrate-template-customizations.yaml` to `migrate-template-customizations.yaml`
- Edit the YAML file and update `cassandra.rateLimit.enabled` and/or `s3.rateLimit.enabled` to value `false`

QoS will now remain disabled during the upgrade, and will be persistent during any future upgrades.

New Features and Improvements

- 45256 - Support for publishing S3 event notifications to Apache Kafka.
- 60127 - Automate unknown use case bucket rules.
- 60128 - Run Major Compactions on a schedule.

- 60749 - Support for changing frontend network interface properties.
- 61866 - Enable Bucket Level Usage Tracking by default.
- 61991 - Improve tiering performance.
- 61995 - Expand nodes by adding disks.
- 62001 - CMC now shows admins the stored object and byte counts for each bucket.
- 62096 - HyperStore to act as a “cache” for data in AWS when used for GDS.
- 62334 - OpenID Connect (OIDC) support.
- 62394 - New integrated hsctl-based installation and add-node validation test suite.
- 62502 - Automate drive replacement process.
- 63801 - Added synchronous deletion mode for DeleteObject operations.
- 63875 - Complete and Abort multipart upload requests are now idempotent.
- 63876 - Reverse tiering now works by migrating metadata from the remote bucket.
- 63903 - CMC now hides the Buckets and Objects menu from system admin users when not applicable.
- 63921 - Added configurable default size for thread pool waiting queues.
- 63925 - Admin and user passwords are now required to disable MFA in all cases.
- 63931 - The CMC group list page now shows the number of users in each group.
- 63941 - Fixed Cassandra busy-writing schema updates issue.
- 64078 - Improved performance of URL-form encoding and decoding in S3.
- 64104 - HyperBalance Fusion relay.
- 64213 - Deprecated the Error Counting Mechanism feature.
- 64297 - Implemented IAM OpenID Connect Provider management APIs.
- 64318 - Cassandra autorepair is now enabled earlier in the startup sequence.
- 64449 - Fixed a cross-site scripting (XSS) vulnerability in the Add a Rating Plan page in the CMC.
- 64474 - A single cluster can now contain a mix of Rocky Linux 8 and Rocky Linux 9 nodes.
- 64836 - Bidirectional and reverse tiering are now available as options for Custom Endpoints.

Bugs Fixed

- 56711 - A disk with very low usage being incorrectly placed in stop-write state.
- 60954 - Clearing alerts in the CMC redirecting to the home page instead of clearing them.
- 61598 - Rate limiting incorrectly throttling S3 operations below the configured threshold.
- 61870 - Persistent false disk error alerts on HyperStore appliances.
- 62311 - Occasionally, the disk audit disabled drives that had no errors.
- 63086 - Uneven data distribution leaving a node at high disk utilization after decommission.
- 63580 - Excessive 503 errors after reducing the maximum requests per connection setting.
- 63611 - Cross-region replication queue failures caused by a metadata parsing error.
- 63752 - Disabling the rate limiting feature not clearing restrictions on existing buckets.
- 63839 - A memory leak causing S3 heap usage to gradually increase over time.
- 64141 - S3 ListParts API returning identical timestamps for all parts of a multipart upload.
- 64155 - SNMP MIB file returning syntax errors when imported into third-party monitoring tools.
- 64196 - False positive OS disk errors on HSA-1700 appliances.
- 64330 - hsstool rdb -compact returning an error when compacting RocksDB.
- 64403 - hsstool cleanup being unable to remove duplicate blobs.
- 64444 - False RAID disk errors showing in the CMC after upgrading.
- 64521 - Repair completion alerts not firing after upgrading.
- 64700 - IAM users being unable to create access keys after upgrading.
- 64752 - Repair status page displaying incorrectly in the CMC.
- 64762 - Missing rollup data after upgrading across multiple versions.
- 64814 - Node agent searching for disk map script in the wrong path on HSA-4300 appliances.
- 64838 - Proactive repair completing in 0 seconds with no data actually repaired.
- 64938 - Redis monitor failing silently after an exception until restarted.
- 64946 - Usage data generation failing when a batch statement exceeded 65,535 entries.
- 64962 - Intermittent performance and memory problems in S3.
- 64995 - Incorrect Japanese translation in the Node Repair Status page.
- 65092 - Auto-expiry jobs were getting stuck without errors.

Security

- 62889 - Updated Spring Framework from 6.1.17 to 6.1.21 to address security vulnerabilities.
- 63150 - Updated Go standard library to address multiple security vulnerabilities.
- 63157 - Updated golang.org/x/crypto to address 10 vulnerabilities including a critical issue.

- 63309 - Updated PostgreSQL driver (pgx) to fix critical and high severity vulnerabilities.
- 63400 - Fixed a server-side request forgery (SSRF) vulnerability in HTTP request handling.
- 63487 - Fixed a GitHub Actions workflow template injection vulnerability.
- 63488 - Fixed unsafe YAML deserialization that could allow remote code execution.
- 63489 - Fixed unsafe subprocess usage in the support tunnel component.
- 63502 - Fixed unsafe subprocess usage that could allow remote code execution.
- 63524 - Fixed an OS command injection vulnerability in the node agent component.
- 63595 - Updated Spring Security Web to address an authorization bypass affecting static resources.
- 63848 - Updated Apache Commons FileUpload from 1.5 to 1.6 to address high severity vulnerabilities.
- 63901 - Fixed a cross-site scripting (XSS) vulnerability when uploading crafted S3 object names.
- 63924 - Fixed an XML External Entity (XXE) injection vulnerability in the XML parser configuration.
- 63926 - [CVE-2025-8885](#) - Addressed Bouncy Castle resource allocation vulnerability.
- 63946 - Updated cloudian-coredns to fix a resource exhaustion denial-of-service vulnerability.
- 64175 - Fixed a cross-site scripting (XSS) vulnerability in the CMC bucket name parameter.
- 64225 - [CVE-2025-48976](#) - Resolved Apache Commons FileUpload denial-of-service vulnerability.
- 64324 - RHSA-2026:0932 - Resolved Java security vulnerability.
- 64386 - Updated gRPC-Go to fix an authorization bypass via improper HTTP/2 path validation.
- 64434 - CWE-425 - Updated Spring Security to address an HTTP response header bypass vulnerability.
- 64764 - [CVE-2026-25243](#) and [CVE-2026-23631](#) - Resolved Redis vulnerabilities.
- 64776 - [CVE-2025-41234](#), [CVE-2025-41242](#), [CVE-2025-41249](#), [CVE-2025-41254](#) - Resolved multiple Spring Framework vulnerabilities.
- 64784 - CWE-287, CWE-863, CWE-303, CWE-770, CWE-337, CWE-400, CWE-367 - Updated cloudian-coredns to fix a resource exhaustion denial-of-service vulnerability.
- 64787 - CWE-1285, CWE-362, CWE-416 - Updated Apache Kafka to address critical and high severity vulnerabilities.

Cloudian HyperStore 8.2.5 Release Notes (2026-02-25)

This is an Early Access (EA) release. The EA classification is given to each feature release of Cloudian software when it is initially posted to the FTP site. This classification indicates that the product has completed a comprehensive development and release procedure including extensive testing for functionality, reliability, and full regression. The release has fully met the Cloudian criteria for production usage by customers if desired. This release may be promoted to GA at a later date. In that case, it will be declared at the Cloudian Support Portal as such.

Installation Notes

NOTE: Starting in 8.2.5, CMC will by default connect to S3 via https (s3.tls.required=true). This is for new installations and upgrades where the 's3.tls.required' config setting has not been changed from the prior default. If your environment requires CMC to connect to S3 via http (not https), run the following commands prior to upgrade:

```
$ hscctl config get s3.tls.required false $ hscctl config set s3.tls.required=false $
```

Supported OS

- Rocky Linux 8.9
- Rocky Linux 8.10 (recommended)
- RHEL 8.7, 8.8, 8.9
- RHEL 8.10 (recommended)

HyperStore 8.2 or later cannot be run on CentOS systems

Installation Notes

For a new install of, or upgrade to, HyperStore 8.2.5 your system must be running Rocky Linux or Red Hat Enterprise Linux (RHEL) 8.9 or 8.10 (recommended).

If you want to install HyperStore on a later version of Rocky Linux or RHEL please contact Cloudian Support. HyperStore does not support other types of Linux distribution, or non-Linux operating systems.

If you want to install HyperStore on a QLC system, please see the Cloudian HyperStore Installation Guide.

Upgrade Notes

Supported versions to upgrade from:

- 8.1
- 8.1.x or an 8.1.x patch
- 8.2 or an 8.2.0.x patch
- 8.2.1 or an 8.2.1.x patch
- 8.2.x or an 8.2.x.y patch

If you are running an older version, please contact Support for upgrade information.

Before upgrading to this release your nodes must have a minimum of 128G of RAM. Contact your Cloudbian representative for information about upgrading existing nodes to meet this requirement.

When upgrading from a HyperStore system that had previously received a `shutdownQoS` via JMX for Cassandra, S3 or both, then to preserve this configuration perform the following before upgrade:

- Within the 8.2.5 staging directory, copy file `sample-migrate-template-customizations.yaml` to `migrate-template-customizations.yaml`
- Edit the YAML file and update `cassandra.rateLimit.enabled` and/or `s3.rateLimit.enabled` to value `false`

QoS will now remain disabled during the upgrade, and will be persistent during any future upgrades.

New Features and improvements

- 46710 - Enable AES-256 by default.
- 47737 - Support for the S3 MFA header for deletes.
- 54136 - Improved security for the indexing and search solution.
- 57491 - Param: `credentials.iamuser.max` (`mts.properties`).
- 58012 - Optimization of Repair operations - `repair/repaiREC`.
- 59712 - Upgrade Runtime to Java 21.
- 60020 - CMC support for blinking drives in HPE appliances.
- 60326 - Optimizations for where Hybrid policies store data.
- 60416 - CMC support for recovery from cybersecurity attack
- 60491 - Notify Bucket Owner when events are detected.
- 60513 - Move RocksDB to NVMe (new clusters).
- 61001 - Add support for repairing a single bucket for replication storage policies
- 61074 - IAM users can now log into CMC.
- 61081 - IAM user can now login using temporary credentials.
- 61082 - CMC page to manage LoginProfile for IAM Users.
- 61162 - Move RocksDB to NVMe (existing clusters).
- 61432 - Move CMC Root User login to use temporary credentials.
- 61492 - Better S3 Credential Protection for Replication.
- 61494 - CMC now uses HTTPS to connect to S3.
- 61613 - Improve Repair efficiency.
- 61679 - Adding a connection pool for Redis to improve performance of QoS operations
- 61756 - [HyperBalance Fusion] Support for multiple VLANs.
- 62155 - Improve `opstatus` for batch delete and tiering
- 62224 - Add maximum number of managed policies limit
- 62320 - Remove deprecated/legacy Bucket Stats feature.
- 62397 - User/Group QoS request rate enforcement for smaller time window
- 62442 - [HyperBalance Fusion] - Ability to exclude newly added nodes from DNS.
- 62469 - [HyperBalance Fusion] Performance improvements large clusters.
- 62532 - S3 API QOS Whitelist/Blacklist S3 API Calls (CMC support).
- 62775 - CMC supports the ability to read and display S3 credentials expiry status.
- 62820 - More logging for the rebalance process.
- 63169 - Terraform: Unable to destroy IAM resource with externally managed credentials [8.2.5]
- 63251 - [HyperBalance Fusion] Allow healthcheck consult probe configuration in DC scope.
- 63437 - Support `-transition-default-minimum-object-size` for `put-bucket-lifecycle-configuration` API.
- 63441 - [Add configuration parameter to control S3 maximum HTTP requests per connection
- 63467 - Improved redundancy for token allocation on disks.
- 63503 - Support multi-threads for BatchDelete to improve performance.
- 63515 - Rename `hsDNS` to HyperBalance Fusion.
- 63520 - Better handling of TCP open connections.
- 63530 - Add new `.tls.client.preferred` options.
- 63615 - Default KMIP profile per user/group.

Bugs Fixed

- 54391 - Salt has been upgraded to 3006.9 to address CVEs.
- 59595 - IAM should not throw an Exception if an IAM Policy is missing records/entries in Cassandra.
- 59876 - “Connection refused” error messages were being displayed while moving the Redis-Mon Primary role.
- 61152 - Multiple disk error alerts were logged after a cluster was upgraded to v8.1
- 61347 - There were issues with the ReplaceDisk process working in CMC.
- 61572 - The netbios name for a File Service pool should be available in the CMC before joining the AD
- 62120 - File Service displays a warning if a RabbitMQ message could not be published.
- 62237 - Threadpool size in Gosbench performance testing was limiting write performance.
- 62264 - Unable to create a group with Spanish/Mexican characters
- 62308 - Changed the IndexSummaryManager alerts from ERROR to INFO to suppress it from alerts.
- 62327 - CMS now identifies disks in a RAID volume that has an error.
- 62546 - Unable to send SMTP notifications despite credentials being correct
- 62630 - Unmounted drives was stopping HyperStore during startup.
- 62722 - On large, busy clusters HyperStore could experience memory leaks.
- 62736 - Japanese translation in Invalid Login message has been improved.
- 62739 - Could not delete non-default system admin user if group listing is disabled for interior CMC
- 62754 - S3 logs sent via rsyslog were incomplete.
- 62814 - CMC had the incorrect HSA-1700 appliance manual.
- 62937 - Creating a User with Spanish/Mexican characters in CMC shows an “illegal characters” error.
- 62972 - Upgrade to 8.2.2 failed due to post upgrade task failure when HTTP proxies were defined.
- 62955 - Upgrade to 8.2.2 failed if a Firewall and remote logging are both enabled.
- 63031 - On a File Services the command “fsck recover-queue” crashed so the archive queue could not be recovered and files were archived.
- 63036 - Files Services archive does not work due to Cassandra invalid query
- 63083 - ReplaceDisk failed on a disabled disk
- 63093 - Usage reports show different results when different granularity option is used
- 63121 - There were occasional failures to upgrade a software-only environment with Salt minions errors.
- 63417 - MPUs were missing fragments for some parts.
- 63494 - There were false down alerts and Cassandra timeout after upgrading to 8.2.2.
- 63587 - Historical usage data per group was lost after upgrading to HyperStore 8.2.2.2.
- 63713 - In some cases the replace disk process would write data to the OS drive.

Security

The following security issues have been addressed in this release:

- CWE-352
- CWE-770
- [CVE-2025-49844](#)
- 63439 - Upgrade spring security library from 6.4.4 to 6.4.10 to address CVEs.
- 62888 - Upgrade OpenJDK 8 version to 1.8.0.462.b08-2 to address CVEs.
- 62888 - Upgrade OpenJDK 21 version to 21.0.8.0.9-1 to address CVEs.

Configuration settings

These configuration settings have been added to the hscctl tool. Refer to the user documentation for details on these.

- `admin.tls.client.preferred`
- `admin.tls.client.verifyCA`
- `cassandra.batch.sizeThreshold.fail`
- `cassandra.batch.sizeThreshold.warn`
- `cassandra.memtable.heapSpace`
- `cassandra.snitch.dynamic.enabled`
- `cassandra.trickleFsync`
- `cmc.tls.client.preferred`
- `common.appliance.redfishAPI.auth.password`
- `common.appliance.redfishAPI.auth.username`
- `common.appliance.redfishAPI.enabled`
- `common.node.defaultMode`
- `file.tls.client.preferred`
- `file.tls.client.verifyCA`

- `firewall.services.search.api.https.allow`
- `firewall.services.search.api.https.deny`
- `hyperstore.cassandra.keyspace.defaultCL.read`
- `hyperstore.cassandra.keyspace.defaultCL.userData`
- `hyperstore.cassandra.keyspace.defaultCL.write`
- `hyperstore.cleanup.objectDelete.queue.maxWaitTime`
- `hyperstore.directory.rdb.digest`
- `hyperstore.directory.rdb.wal`
- `hyperstore.disk.fail.error.threshold.count`
- `hyperstore.disk.fail.error.threshold.time`
- `hyperstore.disk.metadata.checkUsage.enabled`
- `hyperstore.disk.metadata.threshold.startWrite`
- `hyperstore.disk.metadata.threshold.stopWrite`
- `hyperstore.redis.connections.minIdle`
- `hyperstore.repair.tokenRange.listing.timeout`
- `hyperstore.rocksDB.maxOpenFiles`
- `hyperstore.rocksDB.readAheadSize`
- `hyperstore.rocksDB.tableBlockCacheSize`
- `iam.maxPolicies`
- `iam.tls.client.preferred`
- `iam.tls.client.verifyCA`
- `iam.user.maxCredentials`
- `lb.fusion.affinity.datacenter.enabled`
- `lb.fusion.affinity.datacenter.groups`
- `lb.fusion.affinity.datacenter.groups.datacenters`
- `lb.fusion.affinity.datacenter.groups.name`
- `lb.fusion.affinity.datacenter.groups.subnets`
- `lb.fusion.affinity.default`
- `lb.fusion.affinity.node.enabled`
- `lb.fusion.affinity.vlan.enabled`
- `lb.fusion.affinity.vlan.groups`
- `lb.fusion.affinity.vlan.groups.interfaces`
- `lb.fusion.affinity.vlan.groups.name`
- `lb.fusion.affinity.vlan.groups.subnets`
- `lb.fusion.debug`
- `lb.fusion.enabled`
- `lb.fusion.healthcheck.probe.consult.adrs`
- `lb.fusion.healthcheck.probe.consult.policy`
- `lb.fusion.healthcheck.probe.interval`
- `lb.fusion.healthcheck.probe.type`
- `lb.fusion.mode`
- `lb.fusion.publish.address`
- `lb.fusion.publish.interface`
- `lb.fusion.ttl`
- `lb.fusion.weights.enabled`
- `rdma.interface`
- `s3.crr.tls.required`
- `s3.hyperstore.maxThreads.delete`
- `s3.hyperstore.maxThreads.read`
- `s3.hyperstore.maxThreads.repair`
- `s3.hyperstore.maxThreads.write`
- `s3.maxRequestsPerConnection`
- `s3.metrics.enabled`
- `s3.qos.rate.expiryTime`
- `s3.tls.client.preferred`
- `s3.tls.client.verifyCA`
- `search.tls.client.preferred`
- `sqs.tls.client.preferred`

These configuration settings have been removed in this release.

- `admin.tls.verifyCA`

- `file.tls.required`
- `file.tls.verifyCA`
- `hyperstore.directory.rdb`
- `hyperstore.threads.maxRead`
- `hyperstore.threads.maxRepair`
- `hyperstore.threads.maxWrite`
- `lb.hsDNS.enabled`
- `lb.hsDNS.healthcheck.probe.consult.addrs`
- `lb.hsDNS.healthcheck.probe.consult.policy`
- `lb.hsDNS.healthcheck.probe.interval`
- `lb.hsDNS.healthcheck.probe.type`
- `lb.hsDNS.locationAffinity`
- `lb.hsDNS.mode`
- `lb.hsDNS.publishBackend`
- `lb.hsDNS.ttl`
- `lb.hsDNS.weights.enabled`
- `s3.tls.verifyCA`
- `s3.usage.bucketLevel`
- `search.tls.required`

These configuration settings defaults have been changed in this release.

- `cmc.tls.required: true -> false`
- `hyperstore.threadPool.repair.session: 50 -> 10`
- `s3.sse.aes256: false -> true`

Clouidian HyperStore 8.2.2 Release Notes (2025-08-15)

This is an Early Access (EA) release. The EA classification is given to each feature release of Clouidian software when it is initially posted to the FTP site. This classification indicates that the product has completed a comprehensive development and release procedure including extensive testing for functionality, reliability, and full regression. The release has fully met the Clouidian criteria for production usage by customers if desired. This release may be promoted to GA at a later date. In that case, it will be declared at the Clouidian Support Portal as such.

Supported OS

- Rocky Linux 8.9
- Rocky Linux 8.10 (recommended)
- RHEL 8.7, 8.8, 8.9
- RHEL 8.10 (recommended)

HyperStore 8.2 or later cannot be run on CentOS systems

Installation Notes

For a new install of, or upgrade to, HyperStore 8.2.2 your system must be running Rocky Linux or Red Hat Enterprise Linux (RHEL) 8.9 or 8.10 (recommended).

If you want to install HyperStore on a later version of Rocky Linux or RHEL please contact Clouidian Support. HyperStore does not support other types of Linux distribution, or non-Linux operating systems.

To install HyperStore on a QLC system:

Mount Cassandra NVMe drives to `/var/lib/cassandra`

Install HyperStore specifying it is a QLC system:

```
./cloudianInstall.sh configure-hsdns configure-qlc-rdb
```

During install do not change the default cassandra directory from `/var/lib/cassandra/data`.

Upgrade Notes

Supported versions to upgrade from:

- 8.1

- 8.1.x or an 8.1.x patch
- 8.2 or an 8.2.0.x patch
- 8.2.1 or an 8.2.1.x patch

If you are running an older version, please contact Support for upgrade information.

Before upgrading to this release your nodes must have a minimum of 128G of RAM. Contact your Cloudbian representative for information about upgrading existing nodes to meet this requirement.

When upgrading from a HyperStore system that had previously received a `shutdownQoS` via JMX for Cassandra, S3 or both, then to preserve this configuration perform the following before upgrade:

- Within the 8.2.2 staging directory, copy file `sample-migrate-template-customizations.yaml` to `migrate-template-customizations.yaml`
- Edit the YAML file and update `cassandra.rateLimit.enabled` and/or `s3.rateLimit.enabled` to value `false`

QoS will now remain disabled during the upgrade, and will be persistent during any future upgrades.

New Features

- 60348 - Bucket Key for KMIP - caching to improve performance.
- 60764 - Support for Container Object Storage Interface (COSI).
- 61791 - The CoreDNS weighted load balancing feature provides DNS-based load balancing that considers node health and system load. The system adjusts DNS responses based on node load metrics, helping to prevent overloading of individual nodes.
- 62028 - Search service 3 node support.

Improvements

CMC

- 59072 - Users can now choose the storage units (KiB, MiB, GiB and TiB) displayed in CMC for Storage Quota, Data Bytes IN and Data Bytes OUT.
- 61369 - Changes to allow restricting Admins from viewing users S3 Credentials, but still allowing them to administer CMC users passwords. Also, the disk space available alert is no longer active by default. This setting affects new installs, only. The current setting for this alert is preserved on upgrades. "
- 61414 - Preserve CMC branding customizations during HyperStore upgrades.
- 61774 - CMC has a new tab called Active Alerts that displays active alerts in the last 10 minutes.
- 62278 - The `cmc.ui.admin.manageUsers.view.objectDataGrantees` setting can now only be set by root users.
- 62720 - Fix for CMC error messages encoding for duplicated rating plans.

HyperStore

- 52865 - Migrate the current Usage tables from deprecated Super Columns.
- 53100 - In some cases multipart object could not be repaired or downloaded. The S3 service returned 503 errors on getObject.
- 57658 - The raid-check cronjob (`/etc/cron.d/raid-check`) is now scheduled randomly to prevent too much load on the cluster.
- 57942 - Include Disk Slot Location and Top/Bottom Node in Diagnostic and SS logs
- 57991 - Add Node Slot to location to the Appliance section in CMC for the HSA-4200 / 4300.
- 59608 - Stop logging curl ADMIN API sysadmin password to the cron log file (handle rsyslog case)
- 60573 - Support in HSCTL for enabling/disabling S3 QoS.
- 60626 - Report realtime statistics on queue lengths.
- 60886 - Improvements to the performance of common prefix bucket listing.
- 61458 - Improvements to how HyperStore detects disk errors.
- 61536 - For the bucket partitioning service, we improved the Consistency Level to keep metadata consistent in some failure scenarios (such as DC down in multi-DC cluster). Now we apply a strong Consistency Level (QUORUM), instead of using Dynamic Consistency Level in these scenarios. The service will also be able to resume from the previous point of failure even if the fail point partition has been deleted.
- 61625 - Support for running repair for a specified bucket in the CMC Node Advanced page.
- 61642 - Collect number of objects and the number of bytes per bucket in Phone Home.
- 61666 - Upgrade to Java 21 runtime
- 61726 - Improvement to the SQS retry mechanism.
- 61732 - Performance improvements for QoS rate limiting.
- 61844 - Update Tomcat to version 10.1.35.
- 62061 - Updated max value on metadata replicas for distributed EC policy from $2m - 1$ to $2m + 1$. Default selected value will remain $2m - 1$.

- 62353 - Support for AWS managed policies.
- 62420 - Improvements to the Dynamic Consistency Level for Cassandra CQL client.
- 62540 - Performance improvements to the delete process.

Installation/Configuration

- 61790 - Changes to the upgrade process to allow users to select the first DC to be upgraded and the ability to pause the upgrade for up to 48 hours. After that time, the upgrade will continue automatically. For details see the Upgrade instructions in the 8.2.2 HyperStore Admin Guide or in the 8.2.2 HyperStore Help (under /opt/cloudian-staging/8.2.2/doc after unpacking the 8.2.2 product package).

S3

- 61333 - On large, busy clusters the S3 service could become unresponsive.
- 62399 - Support for S3 credentials with expiration date through the admin API.
- 62429 - Support for S3 API QOS Allow/DenyList for S3 API Calls.

Bugs Fixed

CMC

- 50304 - Prevent CMC from sending UninstallNode operation to the current node.
- 52116 - The number of objects on dashboard is negative after deleting all buckets in HyperStore.
- 60836 - Improvements to the information displayed when there was a CMC error.
- 61568 - During an add node CMC reported that chronyd was not running which caused the add node to fail.
- 61688 - When logging.syslog.remote.enabled is set to true, CMC log file ui-action.log was not included in the logs sent to the remote syslog server.
- 62184 - Some clusters reported buckets becoming unlistable and not accessible from the CMC.
- 62560 - The HSX disk display in the CMC was not working.

File Services

- 61441 - File Service installation fails when the region name is “no”
- 62117 - Files were not archived if RabbitMQ reached the default memory limit.
- 62221 - Addressed issue with files not being archived when the number of files to archive is large.
- 62249 - The File Services updatedb command was taking hours to run and consuming a lot of memory.

HyperStore

- 57844 - HyperStore was verifying licensing even after the system Tiering license was deactivated.
- 59586 - Auto repairs for Cassandra were not starting.
- 60118 - Stopwrite not subtracted from diskAvail in SmartSupport.
- 60813 - Improve IAM API performance for validating the requested resources.
- 60901 - Multi Add node failed with a “Failed to stop purgeTombstones” error.
- 60967 - There was an issue with the metadataDB batch operations imposing stricter requirements than quorum, thus breaching quorum guarantees. This would cause a batch sent to the multi-node DC to fail even though quorum was met.
- 61070 - There was an issue with resuming failed rebalances, this has been fixed.
- 61310 - RedisMON service had a problem connecting to JMX service
- 61438 - On systems running repaircassandra sometimes did not start after an addnode.
- 61513 - There were intermittent errors when trying to download some objects.
- 61552 / 61668 - Rebalance stream sessions marked FAILED due to reaching TokenRangeListingTimeout
- 61602 - Adding control over whether certain parts of HyperStore should or should not send logs to remote log server
- 61692 - HyperStore couldn’t list some bucket folders and objects.
- 61796 - Failed RF-Repair items were not logged to cloudian-hyperstore-repair-failure.log
- 62193 - The CN of the CSR is now added to the DNS alternative names along with all the other DNS names.
- 62239 - In 8.2 disruptive service action “restart” was run in parallel by default, instead of sequential and node-by-node. This has been addressed and whether running service actions from within the installer menu or “hscctl service” directly, service actions will now be sequential again (unless using the -parallel flag, or service action “status” which is always parallel).
- 62352 - If a Cassandra exception occurred while creating new storage policy, the storage policy became unuseable.
- 62369 - Failed drive sometimes led to data consumption on HyperStore appliance OS drives.

Installation / Configuration

- 60740 - In some large clusters nodes were marked as temporarily down which caused 503s errors on upgrade.

- 61057 - Could not complete an upgrade due to Salt issues on a large cluster.
- 61443 - Intermittent issues re-running the patch install have been fixed.
- 62580 - Improvements to preserve the non-default System Admin Root Account MFA credentials during an upgrade.

S3

- 59766 - In some cases multipart object could not be repaired or downloaded. The S3 service returned 503 errors on getObject.
- 61701 - Support CRC64NVME checksums to be compliant with latest AWS SDKs.

Security

The following CVEs have been addressed in this release:

- [CVE-2013-6440](#)
- [CVE-2014-0114](#)
- [CVE-2014-3603](#)
- [CVE-2015-1796](#)
- [CVE-2016-1000031](#)
- [CVE-2016-3092](#)
- [CVE-2019-10086](#)
- [CVE-2020-13956](#)
- [CVE-2020-15250](#)
- [CVE-2020-1953](#)
- [CVE-2020-8908](#)
- [CVE-2021-29425](#)
- [CVE-2021-37533](#)
- [CVE-2021-40690](#)
- [CVE-2021-42550](#)
- [CVE-2022-25647](#)
- [CVE-2022-311](#)
- [CVE-2022-31159](#)
- [CVE-2022-40152](#)
- [CVE-2023-24998](#)
- [CVE-2023-2976](#)
- [CVE-2023-34453](#)
- [CVE-2023-34454](#)
- [CVE-2023-34455](#)
- [CVE-2023-39410](#)
- [CVE-2023-43642](#)
- [CVE-2023-44483](#)
- [CVE-2023-46120](#)
- [CVE-2023-6378](#)
- [CVE-2024-11187](#)
- [CVE-2024-12798](#)
- [CVE-2024-12801](#)
- [CVE-2024-29131](#)
- [CVE-2024-29133](#)
- [CVE-2024-38819](#)
- [CVE-2024-38820](#)
- [CVE-2024-38828](#)
- [CVE-2024-47554](#)
- [CVE-2024-47561](#)
- [CVE-2025-21502](#)
- [CVE-2025-21587](#)
- [CVE-2025-22228](#)
- [CVE-2025-30691](#)
- [CVE-2025-30698](#)
- [CVE-2025-31650](#)
- [CVE-2025-31651](#)
- [CVE-2025-48734](#)

Configuration settings

These configuration settings have been added to the hscctl tool. Refer to the user documentation for details on these. *

```
cassandra.compaction.throughput * cassandra.rateLimit.enabled * cmc.ui.branding.enabled * common.upgrade.datacenter
* common.user.allowDeleteWithKMIPProfiles * lb.hsDNS.weights.enabled * s3.listObjects.commonPrefixes.optimize
* s3.rateLimit.enabled * \<servicename>.logging.syslog.remote.enabled
```

These configuration settings have been renamed in this release.

- `cmc.applicationName1` renamed to `cmc.ui.branding.applicationName`
- `cmc.ui.admin.manageUsers.edit.credentialsGrantees` renamed to `cmc.ui.admin.manageUsers.edit.credentials.grantees` and `cmc.ui.admin.manageUsers.edit.credentials.s3Grantees`
- `lb.hsDNS.public` renamed to `lb.hsDNS.mode`

Clouidian HyperStore 8.2.1 Release Notes

This is an Early Access (EA) release. The EA classification is given to each feature release of Clouidian software when it is initially posted to the FTP site. This classification indicates that the product has completed a comprehensive development and release procedure including extensive testing for functionality, reliability, and full regression. The release has fully met the Clouidian criteria for production usage by customers if desired. This release may be promoted to GA at a later date. In that case, it will be declared at the Clouidian Support Portal as such.

Supported OS

- Rocky Linux 8.9
- Rocky Linux 8.10 (recommended)
- RHEL 8.7, 8.8, 8.9
- RHEL 8.10 (recommended)

HyperStore 8.2 cannot be run on CentOS systems

Installation Notes

If your cluster is running HyperStore 8.2.0.x and has objects with additional checksums, then you may need to run a tool to correct the returned algorithm type. Note that the checksum values are correct, only the algorithm name is wrong in some cases. This affects bucket listing and related listing operations. Data integrity is not affected. Contact support for more information.

For a new install of, or upgrade to, HyperStore 8.2.1 your system must be running Rocky Linux or Red Hat Enterprise Linux (RHEL) 8.9 or 8.10 (recommended).

If you want to install HyperStore on a later version of Rocky Linux or RHEL please contact Clouidian Support. HyperStore does not support other types of Linux distribution, or non-Linux operating systems.

To install HyperStore on a QLC system:

Mount Cassandra NVMe drives to `/var/lib/cassandra`

Install HyperStore specifying it is a QLC system:

```
./cloudianInstall.sh configure-hsdns configure-qlc-rdb
```

During install do not change the default cassandra directory from `/var/lib/cassandra/data`.

Upgrade Notes

Supported versions to upgrade from:

- 8.1
- 8.1.1 or an 8.1.1.x patch
- 8.2 or an 8.2.0.x patch

If you are running an older version, please contact Support for upgrade information.

Before upgrading to this release your nodes must have a minimum of 128G of RAM. Contact your Clouidian representative for information about upgrading existing nodes to meet this requirement.

S3 Notes

NOTE: The legacy ‘Bucket Stats’ (s3.usage.bucketLevel) feature is being deprecated in favor of the new per bucket usage tracking (via s3.qos.bucketLevel)

These APIs are affected and will be removed in future releases: POST /usage/bucket GET /usage with ‘bucket’ parameter. GET /usage will be maintained in the future, except for the ‘bucket’ parameter, which will no longer be supported POST /usage/repair/bucket

Note that GET /system/bytecount, GET /system/objectcount, GET /system/groupbytecount, and GET /system/groupobjectcount are the newer recommended APIs for getting per bucket usage information.

Improvements

CMC

- 62061 - In CMC the maximum Metadata Replicas has increased from $2m - 1$ to $2m + 1$ for distributed EC policies.

Bugs Fixed

Installation/Configuration

- 62239 - In 8.2 disruptive service action “restart” was run in parallel by default, instead of sequential and node-by-node. This has been addressed and whether running service actions from within the installer menu or “hscctl service” directly, service actions will now be sequential again (unless using the `-parallel` flag, or service action “status” which is always parallel).

S3

- 62182 - Some objects may return the wrong checksum algorithm names for GET/HEAD and bucket listing.
- 62189 - In-progress multi part uploads may fail after upgrade.

Cloudian HyperStore 8.2 Release

8.2 New Features

- 55267 - KMIP Multi-KMS Support
- 57212 - Support for Glacier Deep Archive storage class.
- 59770 - Proxy tiering - support size based tiering for QLC devices.
- 59825 - Support for viewing data between 2 points in time with the option of undoing changes made during a specific time period.
- 59851 - S3 select support.
- 60121 - Support for Cyberstorage putObjectCopy blockOverwriteEncryption feature
- 60504 - Support for built-in load balancer (HSDNS)
- 60598 - QLC support.
- 61185 - Platform support for 4300 family appliances.
- 61713 - Support CRC64NVME checksums to be compliant with latest AWS SDKs

8.2 Improvements

Installation/Configuration

- 59372 - Cleanup Tombstone Service should be disabled during upgrade.
- 60621 - Support for forwarding logs to a remote syslog server.

CMC

- 58549 - CMC bucket listing performance improvements
- 59091 - Better special character handling for static website names.
- 59498 - CMC login via SAML
- 59917 - CMC should not show rack title when there’s only a single rack
- 59965 - CMC System Admin login via SAML

IAM

- 54688 - Prevent user deletion when user tied with IAM users and policies.
- 59967 - Add redis lock to support create, modify, and delete IAM entities concurrently.

Admin

- 59978 - Creating IAM users and access keys concurrently with inconsistent results.
- 60382 - Add a new parameter to enable or disable deleting users that have IAM resources.

HyperStore

- 47888 - HyperStore Firewall improvements, including support for custom ACLs and services.
- 50731 - Repair feature for RF-repair has been added.
- 55860 - Encode metadata stored in RocksDB.
- 57039 - An improvement to add a hsstool opstatus command to diagnose the single job and also to provide an ability to retry an individual job
- 57078 - Clusterwide RF-repair with hybrid storage policy.
- 57752 - Support for concurrent disk rebuild for replica data.
- 58497 - Enable HyperSearch in CMC
- 58545 - Add OpStatus for batch delete operation
- 58546 - Add OpStatus for Bucket Partitioning
- 58575 - Repairec - Add an option to repair a specific bucket
- 59707 - hsstool whereis command now displays unbalanced fragments.
- 59708 - Proactive Repair - Enable proactive repair during cluster expansion.
- 59709 - Proactive Repair - Prioritize repair based on drive failure events.
- 59761 - SQS enhancements
- 59950 - More granular control of the how you run the hsstool repairqueue command.
- 60167 - Cleanup now starts with the fullest drives first.
- 60325 - Allow a disk rebuild to take place while doing a re-balance.
- 60330 - Smart support now tracks the number of Search installs.
- 60333 - RocksDB commit logs were moved to NVMe disks.
- 60997 - Implementing new threshold limit calculator for Metadata DB & S3 QoS

Phone Home/Data Collector

- 60775 - Datacollector logs error “ClouidianDataCollector:Checking serviceStatus has failed” For Cassandra Port 7199

Search Service

- 58967 - Update Search client from Elastic Search 7.7.1 to Open Search 2.10.0.

S3

- 54689 - S3 API QoS
- 59954 - Improvements to S3 services performance by increasing concurrency.
- 59770 - Proxy tiering - support size based tiering.
- 60743 - Partial support for Veeam S3 API Calls.
- 61726 - fix SQS message retry mechanism

8.2 Bugs Fixed

Installation/Configuration

- 48571 - IAM CreateAccessKey API did not work asynchronously
- 59967 - Improvements to how IAM entities are created.
- 59978 - Creating IAM users and access keys concurrently gave inconsistent results.
- 60775 - The Datacollector was intermittently unable to connect to Metadata DB port 7199

CMC

- 59685 - CMC: LDAP auth field names are interchanged
- 59822 - Make System User optional for System Admin via LDAP

- 60736 - Once either an addNode or uninstallNode has started, any other cluster modification operations cannot be started until the addNode or uninstallNode has completed.
- 61269 - If s3.tls.required is set to True, buckets could not be created in a non-default region from CMC.
- 61446 - Interoperability problem between MFA and FIPS when setting up MFA on CMC.
- 59685 - CMC: LDAP auth field names are interchanged

HyperStore

- 57063 - Redismon was going down several times per day on some clusters.
- 58560 - The add DC operation did not update /etc/hosts file on existing cluster (non-master) nodes with new nodes IP addresses.
- 60069 - Node loses all content in /etc/fstab and can only boot into emergency mode
- 60300 - CMC allowed invalid characters for the Index Document name of a Static Website. This caused bucket listing to fail.
- 60304 - For replicated EC storage policies, with the CL setting is Local Quorum, the TRMAP wasn't created correctly.
- 60320 - Under some conditions multiple disks were being disabled.
- 60459 - Improvements to HyperStore handling of metadata resurrection
- 60537 - get-bucket-encryption does not report correctly if encryption is enabled on the storage policy
- 60696 - The ledmon process was consuming large amounts of memory.
- 60740 - In some large clusters nodes were marked as temporarily down which caused 503s errors on upgrade.
- 60759 - Add a lock in hsstool decommission command to prevent parallel decommissions.
- 60825 - Improvements to the mountpoint RepairEC process.
- 60845 - There were multiple disk error alerts at the same time after the cluster was upgraded to v8.1.
- 60977 - Rebalances were not progressing.
- 60996 - BLC had no available host.
- 61262 - Some large clusters were experiencing HyperStore java and kernel out of memory conditions.
- 61263 - Some large, busy clusters were experiencing out of memory conditions.
- 61413 - Repairs/Cleanups were causing a large number of connections to be stuck in a CLOSE_WAIT status, which could degrade performance.

File Services

- 58852 - File becomes inaccessible if process terminates during archival
- 61183 - File Services volumes were crashing after a node in the cluster had out of memory errors.
- 61308 - Chrony was not configured on File Service nodes.
- 61325 - Users could not uninstall File Services if a previous cleanup run was not finished successfully.
- 61435 - LockManager may skip flushing metadata when releasing an inode lock.
- 61545 - Segmentation fault before starting copy from write to read cache
- 61558 - Data loss when file write, on the same cluster node, completes after archiving completes
- 61612 - Prevent further use of a file deleted on a different cluster node

S3

- 59480 - PutBucketReplication S3 API failed if HTTP 80 port was blocked by a firewall
- 59658 - Static websites were not accessible if they used s3endpoints instead of static website endpoints.
- 60935 - CORS headers are not being sent back in the response for some request/response types
- 60996 - BLC had no available host.
- 61053 - Using s3 copyobject returns invalid xml (chunked) for objects above a certain size.
- 61265 - MPU object results in status code 500s on get.
- 61306 - In some clusters object restores were failing when trying to restore an encrypted object.
- 61053 - Using s3 copyobject returns invalid xml (chunked) for objects above a certain size
- 61265 - MPU object results in status code 500s on get.
- 61306 - In some clusters object restores were failing when trying to restore an encrypted object.

Security

- 58405 - "HTTP Security Header Not Detected" errors for CMC endpoints has been addressed
- 60273 - Security enhancements to the error messages displayed in CMC.
- 60337 - Alert sysadmin on cybersecurity attack.

The following CVEs have been addressed in this release:

- [CVE-2019-20444](#)
- [CVE-2020-9480](#)

- [CVE-2022-1471](#)
- [CVE-2022-25168](#)
- [CVE-2023-22946](#)
- [CVE-2023-32697](#)
- [CVE-2024-21011](#)
- [CVE-2024-21012](#)
- [CVE-2024-21068](#)
- [CVE-2024-21085](#)
- [CVE-2024-21094](#)
- [CVE-2024-21131](#)
- [CVE-2024-21138](#)
- [CVE-2024-21140](#)
- [CVE-2024-21144](#)
- [CVE-2024-21145](#)
- [CVE-2024-21147](#)
- [CVE-2024-4032](#)
- [CVE-2024-6232](#)
- [CVE-2024-6923](#)

Configuration settings

These configuration settings have been added to the hscctl tool. Refer to the user documentation for details on these.

- `common.user.allowDeleteWithIAMResources`
- `common.xForwardedHeaders.enabled`
- `common.xForwardedHeaders.trusted`
- `iam.xForwardedHeaders.enabled`
- `iam.xForwardedHeaders.trusted`
- `s3.cyberstorage.putObjectCopy.blockOverwriteEncryption`
- `s3.getBucketEncryption.allowPolicyResponse`
- `s3.knip.manageProfiles.grantees`
- `s3.xForwardedHeaders.enabled`
- `s3.xForwardedHeaders.trusted`
- `sqs.xForwardedHeaders.enabled`