

IBM Cloud Object Storage System
Version 3.14.13 September Maintenance
Release

Release Notes



This edition applies to IBM Cloud Object Storage System™ and is valid until replaced by new editions.

© **Copyright International Business Machines Corporation 2020.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

- Support information..... V**
- Chapter 1. New Features and Improvements in ClevOS 3.14.13..... 1**
- Chapter 2. New Features and Improvements in ClevOS 3.14.12..... 2**
- Chapter 3. New Features and Improvements in ClevOS 3.14.11..... 3**
- Chapter 4. New Features and Improvements in ClevOS 3.14.10..... 4**
- Chapter 5. New Features and Improvements in ClevOS 3.14.9..... 5**
- Chapter 6. New Features and Improvements in ClevOS 3.14.8..... 6**
- Chapter 7. New Features and Improvements in ClevOS 3.14.7..... 8**
- Chapter 8. New Features and Improvements in ClevOS 3.14.6..... 10**
- Chapter 9. New Features and Improvements in ClevOS 3.14.5..... 13**
- Chapter 10. New Features and Improvements in ClevOS 3.14.4..... 14**
- Chapter 11. New Features and Improvements in ClevOS 3.14.3..... 17**
- Chapter 12. New Features and Improvements in ClevOS 3.14.2..... 19**
- Chapter 13. New Features and Improvements in ClevOS 3.14.1..... 21**
- Chapter 14. New Features and Improvements in ClevOS 3.14.0..... 24**
- Chapter 15. Interface Modifications..... 25**
- Chapter 16. Resolved Issues..... 33**
 - Resolved issues in 3.14.13 September Maintenance..... 33
 - Resolved issues in 3.14.13 July Maintenance..... 33
 - Resolved issues in 3.14.13..... 33
 - Resolved issues in 3.14.12 33
 - Resolved issues in 3.14.11 34
 - Resolved issues in 3.14.10 34
 - Resolved issues in 3.14.9 March Maintenance..... 35
 - Resolved issues in 3.14.9 February Maintenance..... 36
 - Resolved issues in 3.14.9..... 36
 - Resolved issues in 3.14.8..... 36
 - Resolved issues in 3.14.7 January Maintenance..... 37
 - Resolved issues in 3.14.7 November Maintenance..... 37
 - Resolved issues in 3.14.7 October Maintenance..... 37
 - Resolved issues in 3.14.7..... 37

| | |
|--|-----------|
| Resolved issues in 3.14.6 September Maintenance Release..... | 38 |
| Resolved issues in 3.14.6..... | 38 |
| Resolved issues in 3.14.5 July Maintenance Release..... | 39 |
| Resolved issues in 3.14.5..... | 39 |
| Resolved issues in 3.14.4 July Maintenance Release..... | 40 |
| Resolved issues in 3.14.4..... | 40 |
| Resolved issues in 3.14.3 June Maintenance Release..... | 41 |
| Resolved issues in 3.14.3 May Maintenance Release..... | 41 |
| Resolved issues in 3.14.3 April Maintenance Release..... | 41 |
| Resolved issues in 3.14.3 March Maintenance Release..... | 42 |
| Resolved issues in 3.14.3 | 42 |
| Resolved issues in 3.14.2 February Maintenance Release..... | 42 |
| Resolved issues in 3.14.2 | 42 |
| Resolved issues in 3.14.1 February Maintenance Release | 43 |
| Resolved issues in 3.14.1 January Maintenance Release | 43 |
| Resolved issues in 3.14.1 December Maintenance Release | 44 |
| Resolved issues in 3.14.1..... | 44 |
| Resolved issues in 3.14.0..... | 44 |
| Chapter 17. Known issues..... | 45 |
| Upgrading and Installation..... | 48 |
| Container..... | 48 |
| Alerting and Reporting..... | 49 |
| System Behavior..... | 49 |
| Storage Pools..... | 49 |
| Data Evacuation..... | 49 |
| System Configuration..... | 50 |
| Deleting objects..... | 50 |
| Manager Web Interface..... | 50 |
| Vaults..... | 51 |
| Vault Mirrors..... | 51 |
| Vault migration..... | 51 |
| Chapter 18. Supported Hardware Platforms..... | 52 |
| IBM Cloud Object Storage Appliances..... | 52 |
| Hewlett Packard Enterprise..... | 53 |
| Seagate..... | 53 |
| Cisco..... | 53 |
| Dell..... | 54 |
| Lenovo..... | 54 |
| Quanta Cloud Technology (QCT)..... | 54 |
| Chapter 19. Incompatible Hardware and Firmware with ClevOS..... | 55 |
| Broadcom..... | 55 |
| Hewlett Packard..... | 55 |
| IBM Cloud Object Storage Appliances..... | 55 |
| Seagate..... | 55 |
| Supermicro..... | 56 |
| Notices..... | 57 |
| Trademarks..... | 58 |

Support information

Technical support contacts.

For more information on the product or help with troubleshooting, contact IBM Support at ibm.com/mysupport or visit the [Directory of worldwide contacts](#).

Chapter 1. New Features and Improvements in ClevOS 3.14.13

Enhanced 500/503 error alerting (1507)

This feature modifies the existing IBM COS Manager incident messages for 500 and 503 errors that display in the event console. The HTTP method of a 500 or 503 error incident message is now replaced with the root cause of the error. For example, an error incident that would display as “A 503 error occurred on a PUT request” would now read as “A 503 error occurred due to Write Permit Exhaustion”. Metrics that correspond to the number of errors for each category are now available in the Device API.

Device Memory Graphs (1638)

This feature provides three new graphs for the monitor storage pool and monitor device manager UI pages that are useful for troubleshooting systems. The graphs focus on core process resident memory set size and on-heap/off-heap resource permits. Resource permits referenced for this feature description are a construct used by the system to approximate the on/off heap memory used by the core process. Permits are acquired when memory is used and released when memory is freed.

- Resident Memory Set Size, which is the portion of memory occupied by a process that is held in the main memory (RAM).
- On Heap Resource Permits, which is the ON heap memory usage.
- Off Heap Resource Permits, which is the OFF heap memory usage.

Note: See the section on Performance Graphs in the Manager Administration Guide for more details.

Chapter 2. New Features and Improvements in ClevOS 3.14.12

Optimistic Status Reporting (1519)

This feature extends the period from one minute to five minutes that a device reports status to the IBM COS Manager, and is designed to improve the Manager's performance . However, if the device detects an adverse condition, it reports the status immediately to the IBM COS Manager . If a device is running a software version that is not supported by this feature it does not report the status immediately to the Manager.

Conditional PUTs for multi-part uploads (1610)

Conditional requests (<https://tools.ietf.org/html/rfc7232>) are now supported for multipart uploads. Preconditions are only specified with the initiate request, but are enforced during both initiate and complete processing.

Note: The COS SDK will be updated in a future release to support these new request headers.

Chapter 3. New Features and Improvements in ClevOS 3.14.11

Optimistic Status Reporting (1519)

This feature extends the period from one minute to five minutes that a device reports status to the IBM COS Manager, and is designed to improve the Manager's performance. However, if the device detects an adverse condition, it reports the status immediately to the IBM COS Manager. If a device is running a software version that is not supported by this feature it does not report the status immediately to the Manager.

Mirror Write Threshold Functionality (1590)

When creating and editing mirrors via the Manager REST API and Manager UI, system administrators can no longer configure the mirror synchronous writes property. Instead, the system administrator can choose one of the following options:

- Asynchronous: A response is returned once either vault has confirmed the write successful.
- Synchronous: A response is returned once both vaults have responded. Only one vault needs to confirm the object was written.
- Hard Synchronous: A response is returned once both vaults have responded, and both vaults need to confirm the object was written.

Mirror behavior has not changed. Only a new option is introduced (Hard Synchronous). The case when the Mirror Synchronous Writes property not selected in previous releases of the Manager UI, will now map to the Asynchronous option on the new Manager UI. Conversely, the case when the Mirror Synchronous Writes property was selected in previous releases of the Manager UI, will now map to the Synchronous option on the new Manager UI.

Chapter 4. New Features and Improvements in ClevOS 3.14.10

Device Scalability (1302)

This feature delivers a set of manager performance improvements that are designed to accommodate increasing system sizes.

Organizations (1485)

This feature introduces the concept of an Organization, which is a configurable grouping of vaults and accounts. Administrators can create new vaults and accounts, and associate them with specific Organizations. In addition to the existing capability of allowing storage quotas to be applied to each individual vault, a storage quota can be applied to the collection of vaults owned by the Organization. The Administrator can set the maximum number of vaults to be created under an Organization.

A user belonging to any Organization can have a Vault Provisioner role, allowing them to create new vaults using the S3 API. A user with the Security Officer or Super User role can create one or more Organizations from the top level Security page and then specify storage pool size allotments on the Create or Edit Organization pages under the Security tab. On this page, they can also specify the number of vaults that can be created by a given Organization.

Note: This feature is not supported in container mode, which means standard vaults converted to container vaults lose the configurable grouping of vaults by Organization. This feature is not supported for container vault, service vault or management vaults. LDAP will not be supported for Organizations. A user cannot assign an Organization to groups.

Chapter 5. New Features and Improvements in ClevOS 3.14.9

Log Redaction (1516)

This feature introduces the ability for the system to redact uploaded access logs in management vaults after they are a certain number of days old: for example 365 days. It is not changing what or how redaction occurs. It is only changing when redaction occurs.

This feature applies to only the redaction of access logs in management vaults, which was released earlier as Personally Identifiable Information (PII) for Log Redaction 1, and highlights of its use are noted by three primary benefits.

1. PII is removed from logs. Features, past and future, define what keys are PII and update the redaction framework so that the values of those key are redacted as desired.
2. Only access log objects stored in management vaults, or access logs and HTTP logs collected during Log Collection are affected (if redaction is enabled). Logs on the OS disks of devices are never redacted.
3. Redaction for a given access log object in a management vault happens X days after the access log's rotation, where X is defined by the system administrator using the Manager UI.

Chapter 6. New Features and Improvements in ClevOS 3.14.8

New vaults that are created using this release with index enabled have a new index format that includes object metadata. This release greatly improves LIST performance (up to 1500% better in a controlled test environment) and is required for the Object Expiration On-Prem (1450) feature. Customers with "Index On" that are extremely sensitive to small file write performance, and have no interest in expiry, might want to override this new default setting. If so, contact IBM® support.

Object Expiration On-Prem (1450)

This feature provides S3 Lifecycle Configuration Support for Expiration actions. It allows automatic deletion of certain objects based on object age or a specific date of expiration. Users can configure expiration policies on a bucket level.

This feature requires that buckets have an index with metadata, which is already the case for container mode buckets, and with 3.14.8 is now the case for each newly created vault mode bucket. Existing vault mode buckets cannot be modified and are not support this feature.

Note:

This feature requires that buckets have an index with metadata, which means Object Expiration can be enabled on vaults with only Name Index Format (version 4). See the [Feature Description Document](#) in the IBM Knowledge Center for more information.

Support Notifications for On-Prem Container Mode (1447)

This feature extends the Notifications Feature (1074) by allowing Notifications to be configured for Containers in Container mode systems. Notifications can be configured to be generated for all object write/overwrite/delete operations to specific Containers or all Containers in a Container Vault. A unique Topic may be specified when enabling notifications on a particular Container. See the [Feature Description Document](#) in the IBM COS IBM Knowledge Center for more information.

Support Enabling/Disabling Access Keys (1477)

This feature adds the ability to enable/disable access keys. "Cloud Object Storage" API request that uses disabled access keys are rejected. Security Officers can enable/disable access keys using the Manager's **editAccountAccessKey** UI/REST API. Vault users can enable/disable their access keys using the Manager's **editMyAccountAccessKy** UI/REST API.

SDK support for On-Prem IBM Cloud Object Storage (COS) (1362)

This feature covers the COS SDK languages for Java™, Python, and NodeJS that now support COS systems and services. For application developers targeting the IBM Cloud Object Storage System, IBM provides SDKs that have been purpose that is built for use with IBM COS to best utilize the capabilities available, and provide an alternative to other more generic SDKs.

For detail on the specific IBM COS features supported by these SDKs please refer to the section on [SDKs and Tools by language](#) in the IBM COS IBM Knowledge Center .

Vault Index Version Selection (1596)

Each Standard Vault has an associated Vault Index in an IBM Cloud Object Storage System. Vault Index is a distributed dispersed-data structure that resides in the Vault, and is enabled with an S3 Listing of objects for the vault. Prior to the release of ClevOS 3.14.8, all configured Standard Vaults supported only Name Index Format (version 2) as such just the names of all S3 Objects in the Vault are maintained. All S3 Listing requests for Standard Vault are serviced using the Vault Index (version 2) to get the object names,

followed by the lookup of the object's metadata to retrieve info about the object's last modified time, e-tag, size, storage class and owner information for inclusion in the listing response.

Starting with ClevOS 3.14.8, IBM is introducing feature functionality to support Name Index Format (version 4) and the ability for the user to configure the Index Version for a Standard Vault. This feature allows a default of version 2 or version 4 to be specified at the system level, and optionally at an individual storage pool. When a new Standard Vault is created, the user has the option to override the default and pick the Name Index Format to be either in legacy version 2 or new version 4. Vault Index (version 4) supports storing additional object's metadata information in the Index along with the object names. As part of the S3 Listing requests all the information needed for the responses are available in the Index, as such the listing performance is significantly better, as there is no need for additional lookup of metadata of all objects included in the response. However, and depending upon the workflow for the Standard Vault, the user might encounter a performance degradation with latencies, while servicing S3 PUT requests for small objects (< 1 MB).

A Standard Vault using the Name Index Format (version 4) provides numerous benefits over a Standard Vault using Name Index Format (version 2).

1. Observes improved latencies while servicing all the S3 Listing Requests
2. Supports more S3 listing operations.
3. Observes less CPU and Disk Utilizations on the Slicestor[®] devices. Depending upon the workflow for the Standard Vault, the user might encounter a performance degradation with latencies, while servicing S3 PUT requests for small objects (< 1 MB).

Note: The Object Expiration feature that is also available in this release can be enabled on only Standard Vaults that have Vault Index (version 4) configured. All planned future Data Lifecycle Management features can be supported on only Standard Vaults that have Vault Index (version 4) configured. The exception are vaults for workloads that are sensitive to increase in latency for PUT operations for small objects, it is recommended to enable and use Vault Index (version 4) for all new vaults. See the [Feature Description Document](#) in the IBM Knowledge Center for more information.

Chapter 7. New Features and Improvements in ClevOS 3.14.7

New Manager REST APIs - Removing Mirrored Vaults (1433)

This feature adds two new APIs, one for breaking a mirror and another for destroying a mirror.

1. Breaking a mirror removes the relationship between vaults. This relationship cannot be reestablished. The vaults and object data are not deleted. This behaves the same way as the existing Delete Mirror API.
2. Destroying a mirror deletes the mirror and its associated vaults, but only if it is empty. This is a permanent action. There is no process to recover or restore the deleted vaults or mirror.
3. There should be no active write during a mirror destroy action on standard or protected mirrors, it can result in destruction of a object data. This also applies to the break mirror operation.
4. This feature, particularly the Destroy Mirror API, may not work properly with SOH objects.
5. Destroy action is not supported if both name index and recover listing are disabled.

Interfaces/Manager REST API

New APIs have been added for breaking a mirror (Break Mirror API) and destroying a mirror (Destroy Mirror API). See the Manager REST API Guide for details.

Examples

- Break Mirror API-

```
curl -X POST -u <USERNAME>:<PASSWORD> --header "Content-Type: application/x-www-form-urlencoded" 'https://<manager-ip>/manager/api/json/1.0/breakMirror.adm' -d 'id=<mirror-id>&vaultId=<vault-id>'
```

- Destroy Mirror API-

```
curl -X POST -u <USERNAME>:<PASSWORD> --header "Content-Type: application/x-www-form-urlencoded" 'https://<manager-ip>/manager/api/json/1.0/destroyMirror.adm' -d 'id=<mirror-id>&password=<password>'
```

Note: The functionality for the Break Mirror API is equivalent to the existing Delete Mirror API. Support for the Delete Mirror API continues to exist for now. However, it will be deprecated in a follow-on release. As a result we strongly encourage users to replace the Delete Mirror API with the Break Mirror API.

Compliance logging changes (1453)

This feature adds a new log that records successful and unsuccessful authentication attempts on the Manager. All attempts made to the Manager UI or REST API are consolidated into a single log file located on the Manager device.

See Log File Reference Guide for new content on Additional Structured Logs.

Net Promoter Score Manager GUI Integration (1425)

After successfully upgrading to ClevOS 3.14.7, all Manager users can provide feedback about IBM's Cloud Object Storage System by responding to the modal dialog that appears on the landing page after logging into the system, or by selecting the link from the help menu.

Support in-release vault format increment (1514)

The view system API output has changed slightly for this feature. The format for the "**vaultFormat**" and "**newVaultFormat**" fields have changed from a String to an Integer.

Concurrent Accesser Upgrade (1424)

This feature introduces the ability to concurrently upgrade multiple Accesser® devices in an Access Pool, resulting in significantly faster system upgrades. Previously, Accesser device upgrades were limited to one Accesser device at a time per Vault deployment. By default, the system now allows upgrading of one Accesser device at a time per Access Pool, and per Site when multiple Access Pools are simultaneously submitted for an upgrade.

System administrators can now set a maximum percentage of Accesser devices at an Access Pool level, while taking into account Sites, which can be unhealthy while executing maintenance operations like an upgrade. The system attempts to upgrade as many Accesser devices in an Access Pool as possible, but the percentage of Accesser devices that are unhealthy or upgrading do not exceed the specified percentage of devices per Access Pool and per Site. This percentage can be set as a system default, with an option to override on individual Access Pools.

Interface Modifications:

Manager REST API

- Added new method for **editDefaultMaxAccesserDevicesOfflinePercentage**
- Added new method **editMaxAccesserDevicesOfflinePercentage**

Documentation to reference:

- Manager REST API Guide
- Manager Administration Guide
- Manager UI Help

Programmatic Access Key - Secret Key enable/disable (1477)

This feature adds the ability to enable/disable access keys. "Cloud object storage" API request that use disabled access keys are rejected. Security Officers can enable/disable access keys using the Manager's **editAccountAccessKey** UI/REST API. Vault users can enable/disable their access keys using the Manager's **editMyAccountAccessKy** UI/REST API.

Infiniband Partition Key Support (F1448)

This feature added support to InfiniBand devices for enabling InfiniBand partitions when using PKEYs. Before an InfiniBand device with a PKEY can be assigned to a channel, you must first add a partition to the port using the command "**port <INFINIBAND_DEVICE> add_partition <PKEY>**", and then activate it.

Chapter 8. New Features and Improvements in ClevOS 3.14.6

SNMPv3 (1378)

SNMPv3 support is added for both polling and alert forwarding. In addition, inform support has been added for both SNMPv2c and SNMPv3. If you are using AES256 encryption, you must review the Manager Administration Guide for more information. Please see the Manager REST API Guide and the Interfaces section found in these release notes for specific SNMPv3 changes.

- SNMP configuration
- Alert forward configuration

AONT Performance Improvement (1202)

This feature introduces several changes to implementation and configuration of privacy algorithms for SecureSlice:

- AES-GCM-256 is being added to list of available privacy algorithms. It provides better security and performance than existing algorithms.
- AES-GCM-256 can provide up to 15% improvement for maximum OPS and latency versus current default SecureSlice algorithm AES-128.
- Following completion of upgrade process for all devices in the system, AES-GCM-256 is enabled and set as default / recommended privacy algorithm.
- Functionality to configure privacy algorithm to be used as default when creating new vaults added to Manager.
- Functionality to select privacy algorithm for new vault is added to vault creating page, vault template page, mirror page and mirror template page in Manager, as well as to corresponding Manager API methods.
- Functionality to change privacy algorithm for existing vault added to vault configuration page in Manager as well as to corresponding Manager API method. It's possible to enable privacy for vault, which had it disabled before and vice versa.
- If security algorithm for existing vault has been changed, only new objects written to this vault will be encrypted using such algorithm. Existing objects won't be re-encrypted and won't benefit from improved performance and security.
- Vault Summary report in Manager will provide details about current privacy algorithm for each vault.
- New audit message is generated when privacy algorithm are changed for a vault, which allows Insight® and third-party monitoring tools to track such events.

Interface Modifications

Manager REST API

- New method `configureSystemLevelPrivacyAlgorithm` added.

Notification Service for IBM Cloud Object Storage (1074)

This feature supports the COS Notification Service which integrates the system with the Apache Kafka distribute streaming-platform as a producer. COS publishes a record each time an object is written, overwritten, or deleted. Notable benefits include:

- Supports Apache Kafka clusters for versions 2.1.1 and up
- A highly reliable implementation that survives system and network outages

- Supports multiple and different Kafka clusters
- Notification service is configurable by assigning it to vaults/storage pools/vault templates
- Distributed retry mechanism works around localized network issues
- Manager incident support to track Kafka cluster issues
- Notification content is easily parsed JSON structured text

Notifications cannot be used on the following:

- Container vaults
- Mirrors
- Vaults with proxy
- Vaults with data migration

Manager UI

Configure Management Vault

The Configure Management Vault page's "Backup HTTP access logs" checkbox is now overloaded to include the backup of rotated notification logs.

Device Log Configuration

The Device Log Configuration page's "Rotation Period" dropdown is now overloaded and additionally applies to notification logs.

The Device Log Configuration page's "Retention Period" dropdown is now overloaded and additionally applies to notification logs.

The Device Log Configuration page's "Rotation Period" dropdown's default is now 1 hour. This default only applies to NEW systems.

Manager REST API

Configure Management Vault

The Configure Management Vault (`configureManagementVault.adm`) REST API Endpoint's "accessLogs" property is now overloaded to include the uploading of rotated notification logs.

Device Log Configuration

The Device Log Configuration (`deviceLogConfiguration.adm`) REST API Endpoint's "accessLogRotationTime" property is now overloaded and additionally applies to notification logs.

The Device Log Configuration (`deviceLogConfiguration.adm`) REST API Endpoint's "maximumAccessLogRetentionTime" property is now overloaded and additionally applies to notification logs.

The Device Log Configuration (`deviceLogConfiguration.adm`) REST API Endpoint's "accessLogRotationTime" property's default is now 3600000 (1 hour). This default only applies to NEW systems.

Log Collection (a.k.a. dump-log)

During Log Collection, rotated notification logs are pulled from the device's Management Vault and placed in the `100-managerdaemon-notificationLogs` directory of a log collection bundle.

Call Home (1495)

For IBM branded hardware, hardware (fan and power supply), device-specific (for example, disk management), and system level incidents (for example, storage pool) are supported.

For non-IBM branded hardware, only system level incidents are supported.

Note: The following incidents do not generate support cases:

- **process**
- **managementDaemon**
- **diskUnusable**
- **maintenanceMode**
- **slicestorConnectionsHigh**
- **deviceCommunication**
- **managerDeviceCommunication**
- **dataReallocationPausedEvents**
- **accesserRequests**

Chapter 9. New Features and Improvements in ClevOS 3.14.5

IBM Log Server Integration (1423)

This feature adds new capability that facilitates upload of dump logs to the IBM Log Server, which replaces the existing IBM Log Server. Through this added functionality, the UI interface for log configuration allows users to configure two log-server endpoints as noted below:

- IBM Log Server that requires IBM log server host name, transfer ID and password. Transfer ID and password can be set up using the customer support portal [IBM Support](#)
- Log Destination, a custom or other internal configuration, which can be an SFTP or HTTP server

This benefits users who can choose between the configured log-server endpoints to upload logs using the UI interface for Collect Logs Now. An option to the above is to enter the IBM log-server end-point data on the UI interface for Collect Logs Now, and upload logs to the IBM log server while doing so without configuring server endpoints.

Drive Dashboard and Bulk Resume (1184)

This feature introduces a new Manager UI focused on drive management, which supports monitoring and issuing bulk-resume actions. Benefits include the ability to:

- View all drives in the system, including their properties, such as drive state, firmware, model, and association to a storage pool or set
- Sort tables and filters to easily locate drives of interest
- Aggregate drive information for a given storage pool or set through new graphs
- Issue resume actions to several drives at once by users who have the following roles: super user, system administrator, or operator

Note: This feature is only compatible with devices 3.13.4 or higher.

Chapter 10. New Features and Improvements in ClevOS 3.14.4

IBM Branded Cloud Object Storage (COS) Hardware Refresh (1275)

Cloud Object Storage System Gen2 hardware is both an extension and an enhancement of the present Cloud hardware. It is an extension as it provides the same functionality using the same Cloud Object Storage software and can co-exist with the existing generation. Clients can intermix past and present generations of hardware in the same storage pool and cluster, even at the set level.

The Cloud Object Storage architecture does not change in the Gen2 hardware as the storage system is provided by a Manager, Accesser, and a Slicestor function. The Manager and Accesser appliance are very much consistent with the present Manager and Accesser appliance.

A key difference in Gen2 will be in how the storage layer is architected. In the present generation, the storage layer is provided by storage servers, meaning that the performance components and the storage components in the storage layer are in the same physical box and therefore must be installed, grown, and expanded together. In Gen2, the storage layer will be divided into two separate components, a controller node that contains the performance components and a disk enclosure that contains the storage components.

In both the present generation and in Gen2, the storage layer is called the Slicestor function, and the functional components, component names, and system management remains exactly the same. This ensures that the same version of Cloud Object Storage software is used across both the present generation and Gen2.

The Manager appliance, the Accesser appliance, and the Controller Node appliance are all based on the exact same 1U server. The appliance configuration is determined based on the function performed by the server:

The Manager (Model M10) appliance supports the following configuration:

- 1 x Intel Xeon Silver 4110 CPU
- 96 GB RAM
- 2 x 960 GB OS/Boot drives – SSD
- 1 x Broadcom MegaRAID 9361-4i controller set in RAID 1 managing the OS/Boot Drives
- 2 x 10 GbE ports
- 2 x Redundant power supplies

The Accesser (Model A10) appliance supports the following configuration:

- 1 x Intel Xeon Gold 6126 CPU
- 192 GB RAM
- 2 x 480 GB OS/Boot drives – SSD – Managed by integrated motherboard control with SW RAID
- 2 x 10 GB Ethernet ports
- 2 x Redundant power supplies

The Controller Node (Model C10) appliance supports the following configuration:

- 1 x Intel Xeon Silver 4110 CPU
- 96 GB RAM • 2 x 480 GB OS/Boot drives – SSD – Managed by integrated motherboard control with SW RAID
- 1 x Broadcom 9305-16e in HBA Mode used to manage the HDDs in the disk enclosures
- 2 x 10 GB Ethernet ports

- 2 x Redundant power supplies

| <i>Table 1. Disc enclosures (models J10, J11 and J12)</i> | | | |
|---|--------------|---------------|--------------|
| Size | Small | Medium | Large |
| JBOD total rack space | 2U | 4U | 4U |
| Slicestor total rack space | 3U | 5U | 5U |
| Capacity enclosures (rack space/disks) | 2U/12 | 4U/53 | 4U/106 |
| Drive sizes (TB) | 4, 8, 12 | 4, 8, 12 | 4, 8, 12 |
| Node raw capacity (min/max TB) | 48/144 | 212/636 | 424/1270 |

When a Controller Node and a disk enclosure are combined, they function as and are referred to as a Slicestor appliance, just like the storage appliance is used today. Gen2 hardware brings about a set of improvements, capabilities, and consolidations to Cloud Object Storage, while doing so at a reduced hardware cost and with increased hardware performance.

Changed Interfaces:

- Manager Application User Interface
- Manager Daemon Device Application Programming Interface
- Manager Application REST Application Programming Interface
- Manager Application Report CSV/Email Output
- Syslog
- SNMP Traps
- Manager Application Call Home Email Bodies.

Email Bodies of "Call Home" Emails

The Call Home Email functionality has been updated. If an incident has a chassis ID parameter, the chassis serial associated with the chassis ID will be used as the "Vender Serial Number" in the body of the email. If a serial cannot be found, then the serial associated with the hardware profile will be used (matching existing behavior).

Open Incident Upgrade Behavior

Any open Fan/Voltage/PSU incidents will automatically close temporarily upon upgrade. The closing event will have text similar to: "Incident closed automatically due to change in software version <version>".

The Device API ("breaking" / non-backwards compatible change made)

The Device API has been updated. Both the state and statistic endpoints of the Device API have been altered.

- The "state" endpoint

Fan, voltage, and PSU JSON maps are moved within an enclosing "chassis" map where the key is the associated chassis ID. The Device API Developer Guide has been updated to reflect these changes. This is a "breaking" / non-backwards compatible change.

- The "statistic" endpoint

The output of the "extract chassis components" System Advanced Configuration setting (also known as "compatibility mode") for "fan", "voltage", and "psu" maps have been updated. Each key is now prepended with "<chassis-id>." As an example: assume "Fan 1" is in chassis with ID "main", the key would now be "main.Fan 1". The Device API Developer Guide was not updated to reflect these changes given compatibility mode is not referenced in that document.

Fan Graphs on the Monitor Device Page

Fan stat data series names will now be of the form "<chassis id> <fan-name>". The Manager Administration Guide has been updated to reflect this change. Chassis serial is not included on these graphs.

Old Fan Data on the Monitor Device Page

Old fan statistical data will remain on graphs.

Device Health Summary widget on the Monitor Device Page

The monitor device health summary lists failed fans and PSUs by name. The output was updated to be prepended with the failed fan/PSU's chassis id. For example, "Fan 1" may be listed as "Main Fan 1" if the chassis id was "main". The first letter of the chassis id is purposely capitalized to be consistent with the fan graph series names. Chassis serial is not included in the updates to this monitor device health summary widget.

The Manager Event Console and SNMP Traps

Users of the event console (via the UI/API) and SNMP traps will see chassis ID added to "drive reset", "psu failed", "psu ok", "fan speed error", "fan speed ok", "voltage error", and "voltage ok" events.

The Failed FRU Report Columns

The chassis ID/chassis serial of fans/PSUs was added to the FRU report. This includes JSON/XML/CSV/Email/UI views. The JSON/XML views will display the value provided by platform (even null and empty strings). The other views will render null and empty strings as "N/A". Non-empty blank strings will be rendered as non-empty blank strings in all views. Note that chassis ID/chassis serial is not provided by devices in mainline (pre merge of this feature) so null/"N/A" will be rendered until those devices are updated to the build with this feature.

The Failed FRU Report Performance

The FRU report now looks up the chassis serial for each FRU using chassis ID provided by the management daemon. This lookup has impacted the FRU report generation performance negatively. Users may notice longer load times for the FRU report when systems have large numbers of devices.

Chapter 11. New Features and Improvements in ClevOS 3.14.3

Container IP allow/disallow (948)

Container mode bucket quota and bucket configuration service API (1342)

These features extend the Service API with the capability of performing bucket-level operations and setting quota and IP whitelisting and blacklisting in Container Mode such that a user with Manager credential, and Service User role, would be able to configure a bucket on behalf of their clients for both on premise, and public cloud. See Container Mode Guide for more information.

The bucket IP whitelisting in container mode extends the bucket IP enforcement capability to proxy client connections for on-prem deployments in both standard vault mode and container mode.

Access Control Consistency (1395)

The system provides a consistent Bucket Read Access Control that specifically addresses standard vault and container level read permission behavior. Previously, containers and standard vaults within the same system had inconsistent Bucket Read Access Control behavior. In some cases, a Bucket Read Access Control grants "read and listing" and in other cases just "listing."

New systems default Bucket Read Access Control to "read and listing." This protects existing operators from unexpected changes. If an operator wants to use listing only, it is recommended that listing is enabled prior to any vaults being created on the system. Otherwise, some objects may become inaccessible when the change is made. The object owner needs to explicitly add object access controls to fix this issue.

During upgrades, the initial value for Bucket Read Access Control are decided based on the following rules:

1. If the system has only container vaults (not migrated from standard vaults) and management vaults, listing is used.
2. Otherwise, "read and listing" are used.

A new checkbox has been added to the Configure Container Mode UI page (Configure Container Mode API). If checked, bucket level requests (for example, PUT bucket, DELETE bucket, ...) received using the Cloud Object Storage API are rejected. Instead, the Service API must be used to send bucket level requests. This should be enabled when operators want to force clients to use operator portals to execute bucket level requests.

Immutable Object Storage in 2 site/mirrored COS systems (1393)

This release enhances the existing protected mirror functionality to improve system availability during an outage or network partition. Vaults in a protected mirror are no longer designated as primary/secondary, instead writes can complete to either vault in a mirror and are synced. The protected mirror functionality is available only after all the devices in the system are upgraded to 3.14.3. Please refer to our third-party assessment letter from Cohasset Associates that speaks to the two site deployment ability to meet the requirements of SEC 17a-4(f).

Documentation to reference:

- Feature Description Document for Retention Vaults and Protected Mirrors
- IBM COS API Reference Guide
- Manager REST API Guide
- Manager Administration Guide
- Vault Mirror Guide

Interface Changes:

- In the manager REST API, the create mirror/mirror template APIs no longer support the protection state parameter; Instead the mirror type parameter is required

Vault Deletion Authorization (1406)

This feature provides an added layer of security against inadvertent or malicious deletion of data. If enabled, the Vault Deletion Authorization feature prevents a single storage system administrator from being able to delete vaults through the Manager UI and Manager REST API. In order for a vault to be deleted, two users with the System Administrator role must approve the deletion of the vault. Vault Deletion Authorization is disabled by default, and can be enabled by users with the Security Officer role.

Custom HTTPS Certificate Chain for Storage Pools with Embedded Accesser Service enabled. (1427)

This feature allows users of the Embedded Accesser Service to configure a single HTTPS certificate chain for the HTTPS interface running across all of the Slicestor devices in a storage pool. The certificate can be shared across all the devices (rather than requiring a unique certificate for each device), greatly lowering cost and maintenance effort.

Drive Dashboard and Bulk Resume (1184)

This feature introduces a new Manager UI focused on drive management, which supports monitoring and issuing bulk-resume actions. Benefits include the ability to:

- View all drives in the system, including their properties, such as drive state, firmware, model, and association to a storage pool or set
- Sort tables and filters to easily locate drives of interest
- Aggregate drive information for a given storage pool or set through new graphs
- Issue resume actions to several drives at once by users who have the following roles: super user, system administrator, or operator

Note: This feature is only compatible with devices 3.13.4 or higher.

Chapter 12. New Features and Improvements in ClevOS 3.14.2

Device Role-based Access (1286)

This feature adds local and LDAP account authentication to the device console login within a system. Permissions can be assigned to accounts by the Security Officer on the Security tab in the Manager Web Interface for Manager devices, all other devices, specified sites, and regions. User commands through the device console such as device configuration changes and local files accessed are recorded in audit logs and are linked back to their local account or LDAP username. These audit logs are stored in the Management Vault for archival purposes.

Login permissions include

- root access
- read only
- read/write
- no access

Manager UI Global Header and Navigation Changes (1172)

This feature introduces several changes to improve UI scalability and ease of use.

- The navigation tree has been replaced with a more scalable navigation menu
- New summary pages have been provided for access pools and sites
- Search has moved to the header
- Added icon-based menus in the header for help, user profile, and log out

NVMe support for ClevOS devices (1297)

This feature adds support for NVME block devices. Examples include NVME SSD drives and can be used for the operating system or data drive(s). This makes certification as a Slicestor device possible for servers using NVME rather than HDD to store Slice data. It is possible to create a storage pool or change a given storage pool's state such that it contains a mix of NVMe and non-NVMe devices. This state is strongly discouraged but not blocked by the Manager application. With this release, the Dell 740xd NVME server is certified as a Slicestor device.

The troubleshooting console still allows for users to call **cat /proc/scsi/scsi** on NVMe devices; the response does not include information related to NVMe drives. There are no plans to enhance the troubleshooting console to include the ability to grab NVMe drive information.

The Drive Report CSV export had one of its column headings changed from "SCSI Name" to "Device Name".

The Storage Pool Capacity and Disk Report export had one of its column headings changed from "SCSI Name" to "Device Name".

Deprecation:

- The **SCSIName** event parameter is now a deprecated event parameter in favor of the new **deviceName** event parameter.
- The **csTrapSCSIName** is now a deprecated MIB OBJECT-TYPE in favor of the new **csTrapDeviceName** MIB OBJECT-TYPE.

Customer facing API endpoints that have changed:

CLEVERSAFE-TRAP-MIB.txt

- The **csTrapDeviceName** OBJECT-TYPE was added, deprecating the **csTrapSCSIName** OBJECT-TYPE.
- The **csTrapDeviceName** OBJECT-TYPE was appended to notifications that contain **csTrapSCSIName**.
- The descriptions of **csTrapDriveName** and **csTrapDriveID** were modified.

/manager/api/{format}/1.0/eventConsole.adm

- The event parameter **DeviceName** was added, deprecating **SCSIName**.
- Events that send the **SCSIName** parameter will now also send the **deviceName** parameter.

Drive Report:

- CSV Export – The column heading "SCSI Name" was changed to "Device Name."
- JSON/XML view – The key "deviceName" was added. Its value is duplicate to that of "scsiName."

Note: Do not confuse the word "Device" with respect to the "Device Name" column and the "Device Serial #" columns. The former refers to the drive name and the later refers to the drive's device's serial number.

Storage Pool Capacity and Disk Report:

- CSV Export – The column heading "SCSI Name" was changed to "Device Name."
- JSON/XML view – The key "deviceName" was added. Its value is duplicate to that of "scsi."

Support Infiniband Hardware in 3rd Party Servers (1313)

This feature supports the Mellanox ConnectX-5 family of InfiniBand single and dual port cards that utilize the IPoIB protocol. Note the ConnectX-5 family of cards has not been sufficiently tested in Ethernet mode. InfiniBand ports can also have their transition modes set while setting the mtu, which is demonstrated in the following examples:

- port <PORTNAME> mtu connected/65520
- port <PORTNAME> mtu datagram/2044

Chapter 13. New Features and Improvements in ClevOS 3.14.1

Immutable Object Storage (1269)

Object Retention is supported for both Vault and Container Modes. In Vault Mode, you can create retention vaults or protected mirrors with immutable object storage policies and objects that are stored in these retention vaults or protected mirrors have an immutable object storage policy.

For Container Mode, you can create container vaults to allow object retention. When Retention is enabled for a container vault, you can create containers with an immutable object storage policy. Objects that are stored in these protected containers also have an immutable object storage policy.

Objects that are contained within retention vaults or protected containers cannot be deleted or modified until the immutable object storage policy allows for the deletion or overwrite. There are various ways to protect vaults or containers using the IBM Cloud Object Storage System to meet the needs of customers that have strict retention requirements from regulatory entities (such as the Security and Exchange Commission), or customers that might have organizational retention requirements, including finite retention, indefinite retention, permanent retention, and legal holds.

Before you upgrade to the 3.14.1 release, and to find more information on this feature, refer to the documentation listed in the reference table.

New Functionality

- Support for Immutable Object Storage in Container Mode
- Support for Permanent Retention
- System Level Configuration of Retention settings:
 - System Minimum Duration
 - System Maximum Duration
 - System Default Retention duration
 - Allow Permanent Retention
- Allows either Content MD-5 **or** V4 content signing for Write Operations
 - Previous releases required Content MD-5 even if V4 content signature was included
- S3 API updates to support permanent retention
 - Added flag for protection operations at bucket level to denote state of permanent retention
 - Updates to error codes and error messages
 - New error codes and error messages
 - Support of -2 for Object Retention-Period
 - -2 denotes permanent retention of object
- Access Log Updates
 - Additional failure messages included
 - New parameter added for bucket protection information to denote state of permanent retention
- Container Vault
 - Flag added to Container Vault Configuration to enable protection support
 - A protection policy can be added to a container only if the associated container vault is enabled for protection.

- Enabling Protection for a container vault does not mean that all containers within that container vault must have protection that is enabled.
- Listing of buckets in a container that uses the service API shows that the protection is enabled/disabled

References to documentation that supports this feature:

| Name | Location |
|--|---|
| Feature Description Document (discusses all features that are related to immutable object storage) | https://www.ibm.com/support/knowledgecenter/STXNRM |
| Manager Administration Guide | https://www.ibm.com/support/knowledgecenter/STXNRM |
| REST API Guide | https://www.ibm.com/support/knowledgecenter/STXNRM |
| COS API | https://www.ibm.com/support/knowledgecenter/STXNRM |

Query number of parts with an MPU object (1176)

This feature provides support for the “part-number” query string for HEAD and GET requests for objects, which were uploaded using Multipart Upload (MPU). It supports querying the number of parts that are associated with an object that have been uploaded using MPU. This enables clients to parallelize large object reads by fetching the component parts in parallel. Additionally, this allows objects written using MPU to be copied while preserving the part boundaries of the original object thus preserving the duplicating etag for this object.

1. The part number query parameter can be provided for GET or HEAD requests.
 - a. For a non-MPU object, a request to read part number 1 should be interpreted as a ranged read request for the entire object.
 - b. For an MPU object, a request to read a part number should be interpreted as a ranged read request for the byte range that is associated with the requested part.
2. Part numbers must be between 1 and 10,000 (inclusive). Any request outside of this range will result in an HTTP 400 error. If a request is made for a part number that is beyond the range of the object, the response will be an HTTP 416 - Requested Range not satisfiable.
3. All ranged read responses must include the Content-Range header consistent with a ranged read response.

Vault Scalability: Support for Slicestor Devices with Large Drive Counts (1219 B)

This feature now supports a maximum of 1500 vaults for systems containing Slicestor devices with large drive counts, for example, those with 60 drives and 96 GB of physical memory. In previous releases, systems with such devices were limited to 1000 vaults.

The actual number of vaults that a system can support will vary based on the following:

- Number of drives within deployed Slicestor devices
- Physical memory present in Slicestor devices
- Manager hardware configuration
- Total number of devices within the system

Additional changes can be required to the Manager and Slicestor device configuration to leverage this feature. Please contact Customer Support if greater than 1000 vault support is required.

Support 9-wide Concentrated Dispersal Mode Device set (1372)

This feature now supports concentrated dispersal mode for a 9-wide storage pool. Unlike concentrated dispersal mode for 3-6 wide device sets, a 9-wide device set will not operate in concentrated dispersal mode by default, since a 9-wide device set is supported as one side of a mirror in standard dispersal mode. Please contact IBM customer support to enable a 9-wide device set in concentrated dispersal mode.

Chapter 14. New Features and Improvements in ClevOS 3.14.0

Indefinite Retention and Event-based Retention capability support (1247)

This feature update is now supporting the following items:

1. The ability to extend retention of an object from the current time using a new header (extend-retention-from-current-time). Refer to COS API documentation.
2. Interpretation of bucket max:

Previous Releases: The total retention period (initial retention period + all subsequent retention extensions) applied to an object cannot exceed the bucket maximum.

Current Release: The retention period being applied to an object in any single request cannot result in the expiration date of that object exceeding the bucket maximum + current time (i.e. cannot extend object beyond bucket maximum from current time)
3. This feature also provides users with the ability to write an object into a bucket with a retention period of -1. This value is used as a placeholder for a user to provide a finite retention period at a later time, through a POST ?extendRetention request. While the retention period of the object is set to -1, the object cannot be deleted or modified. Retention Period of -1 can only be set on the object metadata and can only be configured via an object write operation.
4. The ability for an application to store an object in the IBM Cloud Object Storage System with an indefinite retention period and then allow the object retention to be changed to a finite value. Third party applications can implement Event-based Retention through the use of the indefinite retention API.

Note: See supporting documentation in the Retention Vaults and Protected Mirrors FDD and COS API Guide.

Notification Service for IBM Cloud Object Storage (1074)

This feature supports the COS Notification Service which integrates the system with the Apache Kafka distribute streaming-platform as a producer. COS publishes a record each time an object is written, overwritten, or deleted. Notable benefits include:

- Supports Apache Kafka clusters for versions 0.10.2.1 and up
- A highly reliable implementation that survives system and network outages
- Supports multiple and different Kafka clusters
- Notification service is configurable on a vault-by-vault basis
- Distributed retry mechanism works around localized network issues
- Manager incident support to track Kafka cluster issues
- Notification content is easily parsed JSON structured text

Notifications can not be used on the following:

- Container vaults
- Mirrors
- Vaults with proxy
- Vaults with data migration

Chapter 15. Interface Modifications

API updates for the 3.14.5 release have been referenced in the following documentation:

- REST API Developer Guide
 - NEW section added under Device Management>Bulk drive resume action
New request parameters: **deviceUuid** and **driveUuid**
New response parameters: **deviceStatus**, **driveResults**, **driveUuid** and **driveStatus**

API updates for the 3.14.4 release have been referenced in the following documentation:

- REST API Developer Guide
 - Updated section on Reports>Failed field replaceable unit report>Response>JSON response example
Added new parameters **chassisID** and **chassisSerial** for 1275 hardware feature
 - Updated section on Reports>Field replaceable unit report>Response>Parameters
Added new parameters **chassisID** and **chassisSerial** for 1275 hardware feature

API updates for the 3.14.3 release have been referenced in the following documentation:

- REST API Developer Guide
 - NEW section added under Vault Management>Vault deletion authorization
 - Updated section on Vault Management>Delete a vault
New request parameter: **action**
 - NEW section added under Access Pool Management>Edit storage pool HTTPS certificate chain
New request parameters: **privateKeyPem** and **certificatePem**
 - Revised section heading under Administration>Edit the system TLS Configuration to say Edit the system Network Transport Layer Configuration
Updated request parameter description for parm **clientToAccesserConnectionMode**
 - NEW section added under Access Pool Management>Edit storage pool HTTPS certificate chain
New request parameters: **privateKeyPem** and **certificatePem**

API updates for the 3.14.2 release have been referenced in the following documentation:

- REST API Developer Guide

API updates for the 3.14.1 release have been referenced in the following documentation:

- CSO API Developer Guide
 - Error Codes
 - Add protection to a bucket
MinimumRetention, DefaultRetention, MaximumRetention, EnablePermanentRetention
 - List the protection configuration for a bucket
EnablePermanentRetention
 - Upload a protected object
Retention-Period, Retention-Expiration-Date
 - Upload a protected object using HTML webforms
Retention-Period, Retention-Expiration-Date

- Get the headers of a protected object
Retention-Expiration-Date
- Download a protected object
Retention-Period, Retention-Expiration-Date
- Copy a protected object or copy an object to a protected bucket
Retention-Period
- Extend the retention period of a protected object
Additional-Retention-Period, New-Retention-Period, New-Retention-Expiration-Date, Extend-Retention-From-Current-Time
- List legal holds on a protected object
RetentionExpirationDate
- Upload a part for a protected object
- Complete a multipart upload for protected objects
Retention-Period, Retention-Expiration-Date (edited)
- REST API Developer Guide
 - Updated section on Mirror Management>Create a Mirror
New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled
Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod
 - Updated section on Mirror Management>Create a Mirror Template
New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled
Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod
 - Updated section on Mirror Management>Edit a Mirror
New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled
Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod
 - Updated section on Mirror Management>Edit a Mirror Template
New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled
Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod
 - Updated section on Vault Management>Create a Vault
New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled
Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod
 - Updated section on Vault Management>Create a Vault Template
New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled

Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod

- Updated section on Vault Management>Edit a Vault

New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled

Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod

- Updated section on Vault Management>Edit a Vault Template

New Request parameters: permanentRetentionEnabled and defaultPermanentRetentionDurationEnabled

Revised Request parameters: minimumRetentionPeriod, maximumRetentionPeriod and defaultRetentionPeriod

- NEW section added for Configure Vault Protection

New Request parameters: vaultProtectionEnabled, systemMinRetentionPeriod, systemMaxRetentionPeriod, systemDefaultRetentionPeriod and systemPermRetentionEnabled

API updates for the 3.14.0 release have been referenced in the following documentation:

- CSO API Developer Guide

- Updated section on API reference>Operations on objects

New valid value of -1 for the Retention-Period header, which indicates indefinite retention:

- Requests
 - Upload a protected object
 - Upload a protected object using webforms
 - Get an object's protection configuration
 - Copy a protected object
 - Complete a multipart upload for protected objects
- Responses
 - Download a protect object

New header Extend-Retention-From-Current-Time:

- Requests
 - Extend retention period of a protected object

- REST API Developer Guide

- Added new section on Administration>Add notification service configuration
- Added new section on Administration>Edit notification service configuration
- Added new section on Administration>Delete notification service configuration
- Added new section on Administration>Edit notification service configuration assignment
- Updated section on Vault Management>Create a vault

New Request parameters: notificationServiceTopicOverride and notificationServiceConfigurationId

- Updated section on Vault Management>Create a vault template

New Request parameters: notificationServiceTopicOverride and notificationServiceConfigurationId

- Updated section on Vault Management>Edit a vault

New Request parameters: notificationServiceTopicOverride and notificationServiceConfigurationId

- Updated section on Vault Management>Edit a vault template

New Request parameters: notificationServiceTopicOverride and notificationServiceConfigurationId

API updates for the 3.13 release have been referenced in the following documentation:

Feature Limitations:

COS-31712: If a user uses **createVault** and specifies retention periods, but does not specify the **protectionState** or the **protectionState** is specified as disabled' the user should expect a reject where as in previous releases of the software, the retention periods would have simply been ignored.

COS-34240: Changed **retention-legal-hold-count** header to lower-case for consistency with other retention header responses.

- CSO API Developer Guide

- Mirror-Destination header for GET /bucket, GET /bucket?acl, GET /bucket?cors, GET /bucket?uploads, GET /object, HEAD /object, GET /object?legalhold

- Maximum number of days for retention periods settings is 36159 days

- Value for the "Status" parameter is now "Retention" (it was "Compliance" before)

- New methods:

POST /object (Specify retention periods and add a single legal hold to a protect object with webforms)

POST /object?extendRetention (Extend the retention period of a protected object)

- Device API Guide

- Updated section on Device API Reference>State

New raid section added

State -> raid

Updated JSON and Response Parameters Table to include:

- New Response parameter: raidStatus

- New Response parameter: arrayHealth

- Updated section on Device API Reference>Statistic

Updated JSON and Response Parameters Table to include:

- New Response parameter: applianceLayout

- New Response parameter: applianceType

New Response section: capabilities -> {monitoring, visualization and other capabilities available on the device - see Device API guide for details}

New Response section chassis -> [discrete enclosure units that describes hardware entity information - see Device API guide for details]

New Response section driveThresholds -> { total, warning and error thresholds by drive usage type - see Device API guide for details}

New Response section raid -> arrayHealth parameter

- REST API Developer Guide

- Updated section on Mirror Management>Create a Mirror

New Request parameters: protectionState, minimumRetentionPeriod, maximumRetentionPeriod, defaultRetentionPeriod, and restrictiveAccessControlEnabled

- Updated section on Mirror Management>Create a Mirror Template

- New Request parameters: protectionState, minimumRetentionPeriod, maximumRetentionPeriod, defaultRetentionPeriod, and restrictiveAccessControlEnabled
- Updated section on Mirror Management>Edit a Mirror
 - New Request parameters: protectionState, minimumRetentionPeriod, maximumRetentionPeriod, and defaultRetentionPeriod
- Updated section on Mirror Management>Edit a Mirror Template
 - New Request parameters: protectionState, minimumRetentionPeriod, maximumRetentionPeriod, defaultRetentionPeriod, and restrictiveAccessControlEnabled
- Updated section on Vault Management>Create a Vault
 - New Request parameter: restrictiveAccessControlEnabled
- Updated section on Vault Management>Create a Vault Template
 - New Request parameter: restrictiveAccessControlEnabled
- Updated section on Vault Management>Edit a Vault Template
 - New Request parameter: restrictiveAccessControlEnabled
- NEW section added for Upgrade System Software
- Updated section on Reports>Disk drive and device report>Response
 - Updated JSON
 - New Response parameter: chassisId
 - New Response parameter: enclosureId
 - New Response parameter: slotId
- Updated section on Reports>Failed field replaceable unit report>Response
 - Updated JSON
 - New Response parameter: chassisId
 - New Response parameter: enclosureId
 - New Response parameter: slotId
- Updated section on Reports>Firmware report>Response
 - Updated JSON
 - New Response parameter: chassisId
 - New Response parameter: enclosureId
 - New Response parameter: slotId
- Updated section on Reports>Storage pool capacity and disk report>Response
 - Updated JSON
 - New Response parameter: chassisId
 - New Response parameter: enclosureId
 - New Response parameter: slotId
- Updated section on Administration>View system configuration>Response
 - Updated JSON
 - New Response parameter: driveTotalCount
- Updated section on Device management>Device drive bay nut enclosure action
 - Updated description
 - Updated HTTP

Updated Curl

Response>New Response parameter: chassisId

Response>New Response parameter: enclosureId

Response>New Response parameter: slotId

API Changes 3.14.5

COS-53036: API URLs / parameters related to notification service functionality has changed such that the word "configuration" is no longer present. Old audit action codes were not migrated to remove the word "configuration."

COS-56456: Changed the status code returned for conditions where a bandwidth or operation rate limit is exceeded from HTTP status code 503 to HTTP status code 429.

COS-56281: Device API for Notification Service: Field has a typo for notification blob "**notificationService**" and Fields lack description..

API Changes 3.14.4

COS-53785: Fixed an issue where a HEAD request with a ranged read header would ignore the requested range and return the content length of the entire object.

COS-55935: Fixed an issue where a HEAD request for a portions of an object via a part number query would ignore the part information and return the content length of the entire object.

API Changes 3.14.3

COS-49565: Update Manager Data Model and REST API for mirrorType

API Changes 3.14.2

API Changes 3.14.1

COS-48002 : S3 Extended API now supports query parameter of 'extended' in addition to 'pagination' which was previously supported.

COS-42959: The AWS V4 content-sha256 is not always verified when present, and change an error message.

On-prem Vault mode WORM change in behavior:

- For a PUT protection request, either the content-md5 of the request body xml must be provided, or if using a V4 signature, the provided x-amz-content-sha256 must contain the actual hash instead of "UNSIGNED_PAYLOAD".

(The current on-prem vault mode protection does not require content verification (content-md5 or sha256) on the put protection request)

Change in behavior for regular requests with regard to content-sha256 verification:

- If using a V4 signature with a multipart upload PUT part, and the provided x-amz-content-sha256 contains the actual hash instead of "UNSIGNED_PAYLOAD", then that hash will be validated against the payload.
- If using content-md5 with a multipart upload PUT part or a write extent PATCH request, and the content-md5 is valid, but does not match the calculated payload hash, then the error code will be "BadDigest" now instead of "InvalidDigest".

API Changes 3.14.0

COS-42241: Release Note for CSAFE-9996

The 'settings' object in the viewSystem.adm method has been modified. The attributes accessPoolProtocolType, accessServicePorts, certificateExpirationNotificationDays have been removed.

Note: Removed content for the above attributes from the code in View System Configuration>Response>JSON Response Example .

API Changes 3.13.5

COS-42414: DOC UPDATES related to CSAFE-37117

In 3.13.5, code updates to support URL encoding for List Responses is available.

The below feature flag is used currently to disable the feature.

s3.listing-encoding-enabled = false

Once enabled the results for certain response elements will be URL encoded and users need to make corresponding updates if they are using the encoding-type in the requests.

For all the below operations, we now support a method to encode certain response elements using URL encoding in the response being sent. This is in compliance with AWS S3 API Version 2006-03-01. 1.

1. GET BUCKET (List Objects) Version 1

When the Get Bucket list v1 request includes encoding-type element and when the method is set to URL, the response will URL encode the elements - Delimiter, Marker, Prefix, NextMarker and Key.

2. GET BUCKET (List Objects) Version 2

When the Get Bucket list v2 request includes encoding-type element and when the method is set to URL, the response will URL encode the elements - Delimiter, Prefix, ContinuationToken, Key and StartAfter.

3. GET BUCKET Object Versions

When the GET Bucket Object versions request includes encoding-type element and when the method is set to URL, the response will URL encode the elements - Delimiter, Prefix, Key, KeyMarker and NextKeyMarker.

4. LIST MULTIPART Uploads

When the LIST Multipart Uploads request includes encoding-type element and when the method is set to URL, the response will URL encode the elements - Delimiter, Prefix, Key, KeyMarker and NextKeyMarker.

5. LIST PARTS

When the LIST Parts request includes encoding-type element and when the method is set to URL, the response will URL encode the elements - Key Please refer to AWS S3 API reference for detailed notes for the above requests

Note: Please refer to AWS S3 API reference for detailed notes for the above requests.

API Changes 3.13.4

COS-33549: Device API

State API

When a device is upgraded, any existing disabled drive bay power control states in the openExternalEvents object are removed from the State API.

Statistic API

- Several hardware components such as chassis, enclosure, voltage sensors, fan sensors, power supply sensors, and drive configurations are reported in a new format.
- The voltage, fan, and power supply statistics are reported as properties of a **chassis** object instead of the root of the JSON output. However, statistics in the old format are available for backwards compatibility through the advanced configuration settings of the Manager application. For more information on this advanced configuration setting, contact IBM Customer Support.
- For voltage statistics, **maximum_voltage** and **minimum_voltage** readings are removed. Instead, a **status** property is added. The status can be OK, DISABLED, CRITICAL, UNKNOWN, or NOT_PRESENT.
- For fan statistics, **maximum_speed** and **minimum_speed** readings are removed. Instead, a **status** property is added. The status can be OK, DISABLED, CRITICAL, UNKNOWN, or NOT_PRESENT.

- For CPU temperature statistics, **maximum_temperature** has been removed. Instead, a **status** property is reported. The status can be OK, DISABLED, CRITICAL, UNKNOWN, or NOT_PRESENT.
- Drives now report specific usage types. Valid drive usage types are data, os, osSpare, database, and unknown.
- Drives have a new format for reporting bay identifier. It uses the three new identifiers (**chassis_id**, **enclosure_id** and **slot_id**) and concatenates them together to create the drive bay identifier.
- The enclosure object for listing drive bays with power control capability is no longer available in the root of the JSON by default. The drive bay power control statistics can now be found in **chassis[].enclosure[].slots[].phy**. The legacy enclosure object is available for backwards compatibility through the advanced configuration settings of the Manager application. For more information on this advanced configuration setting, contact IBM Customer Support.
- PCI addresses have been removed from network interface sections in device statistic API.

API Changes 3.13.3

Information on the Get Bucket V2 APIs can be found the COS API guide.

Chapter 16. Resolved Issues

Resolved issues in 3.14.13 September Maintenance

| <i>Table 2. Resolved issues</i> | |
|---------------------------------|---|
| Issue | Description |
| COS-75370 | Resolved an intermittent issue where a Slicestor device would occasionally experience hung I/O operations under heavy load, which would result in the storage service being shutdown. |

Resolved issues in 3.14.13 July Maintenance

| <i>Table 3. Resolved issues</i> | |
|---------------------------------|--|
| Issue | Description |
| COS-72178 | Resolved an issue where COS UI is showing error for empty Fan slots 1&2 on a specific accessor host type. |
| COS-13189 | Resolved an issue where the Disk Nut Storage Action API command will no longer accept driveBay or driveName as the API parameter. |
| COS-70308 | Resolved an issue where an API call fails with error during the addition of accessor container to the accesspool using the editDevice.dam. |
| COS-71822 | In prior releases, Mirrors were allowed to be incorrectly seeded with an object expiration enabled vault. This issue is fixed now. |

Resolved issues in 3.14.13

| <i>Table 4. Resolved issues</i> | |
|---------------------------------|---|
| Issue | Description |
| COS-39337 | Resolved an issue where DLM reporting a non-UUID is not handled properly. |
| COS-65143 | Resolved an issue in which retries of in-memory index operations on the Slicestor were either causing the OOM condition or the detection of a deadlock condition triggering a graceful shutdown of the Slicestor appliance. |

Resolved issues in 3.14.12

| <i>Table 5. Resolved issues</i> | |
|---------------------------------|---|
| Issue | Description |
| COS-55331 | Prior to this release but after 3.14.0, when changing the "System Minimum, Maximum, and Default Duration" via the configureVaultProtection.adm page, the original values may be displayed when the page reloads (following the POST). This issue has now been resolved. |

Table 5. Resolved issues (continued)

| Issue | Description |
|-----------|--|
| COS-38264 | In prior releases, disk drive temperature thresholds were set as a single setting per device, which was insufficient for some configurations where drives may be internally located, or where a device contains a mix of drives with different operating temperature ranges. This issue has been addressed by allowing configuration by disk drive model of the alerting temperature. This approach can be used to prevent false-positive alerts being generated for drive models which have higher operating temperatures. For more information, contact IBM Support. |
| COS-69372 | Resolved an issue that caused the Docker Manager to not come up when deployed on hosts that have IPv6 enabled. |
| COS-67879 | In this release, an improved algorithm has been introduced to address system configurations that have storage pools with more than 100 devices actively archiving content to the management vault. This change will result in a lower access log redaction lag on systems that have no redaction backlog. |

Resolved issues in 3.14.11

Table 6. Resolved issues

| Issue | Description |
|-----------|--|
| COS-67668 | Starting with release 3.14.7 and continuing up through this release, when a device associated with an Access Pool or a Storage Pool is re-imaged and approved, and in some scenarios, a warning message may be incorrectly displayed on the system landing (home) page indicating the device is "associated with vaults using a format more advanced than what they can understand. They must be upgraded in order to function as expected." This issue has been resolved. |
| COS-69372 | Resolved an issue that caused the Docker Manager to not come up when deployed on hosts that have IPv6 enabled. |
| COS-69888 | Resolved an issue that would cause certain appliances to hang when rebooting after an upgrade. |

Resolved issues in 3.14.10

Table 7. Resolved issues

| Issue | Description |
|-----------|---|
| COS-61610 | Resolved an issue where the "Chassis Serial" column in the "Disk Drive and Device Report" was not accounting for Gen2 JBOD hardware information. |
| COS-61865 | In previous releases, the Manager database restore may fail on systems with the management vault feature enabled, when old backup files are restored. Specifically, if a new Accesser device is added to the system after the backup was taken and then associated with an Access Pool that has vaults deployed, the restore of the old backup will fail. This issue has now been resolved. |

Table 7. Resolved issues (continued)

| Issue | Description |
|-----------|---|
| COS-58440 | <p>As of 3.14.6, it is possible to change the Secure Slice algorithm on an existing vault. When this is done, existing objects are still encrypted with the original Secure Slice algorithm. The new algorithm only applies to modified and new objects.</p> <p>Prior to 3.14.10, the Manager UI did not include any indication that some objects may be encrypted with a different algorithm. This concern has now been addressed by providing a triangle shaped icon that appears next to the Secure Slice algorithm, and when hovered over indicates objects may be encrypted using different algorithms. Also, new filters "SecureSlice Algorithm" and "SecureSlice Algorithm Modified" have been added to the Vault Summary page. The latter is only present if the Secure Slice algorithm has been modified. In addition, a "SecureSlice Algorithm Modified" column is now present on the Vault Summary Report.</p> |
| COS-65210 | <p>Modified the threshold for which important rebuild is added to the prioritized rebuild queue. This does a better job of controlling the amount of rebuild performed across the set.</p> |
| COS-62567 | <p>For non-versioned delimited listings on vaults that have had versioning enabled, a common prefix is no longer returned in the response if the most recent version of each underlying object within that prefix is a delete marker.</p> |
| COS-66107 | <p>This release fixes a few errors of enforcing bucket-level Access Control List (ACL) after a Service user configured bucket-level ACL using the service API in the Container Mode. The service API now also supports bucket-level WRITE_ACP and READ_ACP permissions. It offers a consistent way for a service provider to optionally revoke the owner's FULL_CONTROL permission using either bucket-level PATCH or PUT service API command.</p> |

Resolved issues in 3.14.9 March Maintenance

Table 8. Resolved issues

| Issue | Description |
|-----------|---|
| COS-66327 | <p>Software issue of Hard Quota enforcement on the destination Vault during Data Migration (DMS) has been resolved.</p> |
| COS-52454 | <p>Resolved an issue that caused invalid temperature alerts on RAID drives on unqualified appliances. Drive temperature alerts will no longer be generated on unqualified appliances.</p> |
| COS-65775 | <p>Resolved an issue where false alerts were being generated for uninstalled fans on HPE Apollo 4200 Gen9 appliances running in a single processor configuration.</p> |
| COS-59447 | <p>Resolved an issue that would cause the rebuilder to not make progress while expanding a storage pool or removing a storage pool set.</p> |

Resolved issues in 3.14.9 February Maintenance

Table 9. Resolved issues

| Issue | Description |
|-----------|--|
| COS-66786 | Resolved an issue with the vault mode Object Expiration feature that can cause an Accesser appliance to gracefully restart. This problem can occur if the Object Expiration feature is enable on the Storage Pool, while there are also vaults within the Storage Pool that have CORS policy configured. |

Resolved issues in 3.14.9

Table 10. Resolved issues

| Issue | Description |
|-----------|--|
| COS-64334 | In earlier releases, when a user would set a vault quota using IEC units from the Manager UI or REST API, the Manager would display a slightly different value, for example, in the Event Console due to incorrect byte-rounding. This issue has now been resolved. Note this is not a concern when SI units are used. |

Resolved issues in 3.14.8

Table 11. Resolved issues

| Issue | Description |
|-----------|---|
| COS-63969 | In previous releases, the mirror management APIs returned an incorrect status code of 403 when an incorrect mirror id was provided. This issue has now been resolved. The APIs correctly return a status code of 422. The following mirror management APIs were impacted: editMirror , editMirrorAccess , editMirrorVaultSettings , editVault , deleteVault , editVaultAccess , and editVaultProxySettings . |
| COS-63717 | The graph "Queue Length (Estimated Data Sources to Rebuild) (sources)" present on the Monitor Storage Pool page has been renamed to "Estimated High Priority Data Sources to Rebuild (sources)." |
| COS-61879 | Fixed an issue where I/O request from end user with valid credentials was failing with HTTP status code 403. With this fix, the Accesser device returns an HTTP status code 500 in case the Slicestor appliance return Authorization error. |
| COS-60789 | In releases after 3.14.3 but prior to this one, and in some circumstances, the graphs that are associated with "Raw Space Used (bytes)" and "Total Number of Slices (Count) per Device" would show a temporary drop during upgrade due to a reporting issue. This issue has now been resolved. |
| COS-59735 | Resolved an issue that had a chance of causing a temporary outage on a Slicestor device when a failing drive became read-only and was subsequently quarantined. Only Slicestor devices that are deployed in storage pools utilizing the Packed Storage Engine were impacted. |
| COS-54455 | Resolved an issue that would cause a Slicestor device upgrade to fail in the event that a previous upgrade of that device had not gracefully terminated. |
| COS-57675 | Drive Model information is now present in certain drive-related Call Home incidents. In particular, incidents that are associated with the following include this information: quarantined drives, OS drive failures, and RAID drive status. |

Resolved issues in 3.14.7 January Maintenance

Table 12. Resolved issues

| Issue | Description |
|-----------|--|
| COS-65085 | A memory leak in the Manager Daemon process may cause Accesser appliances to not report UI Chart data and potentially interrupt I/O when the "Backup Access Logs to the Management Vault" functionality on the Manager is enabled. Also, it is likely "not reporting status" incidents will open on the Manager UI. These issues have now been resolved. |

Resolved issues in 3.14.7 November Maintenance

Table 13. Resolved issues

| Issue | Description |
|-----------|---|
| COS-62872 | In previous releases, there was a discrepancy between the Manager UI and the View System API reported device utilization value. Device utilization was being reported only as vault data in the API, and it was the sum of system and vault data on the Manager UI. This issue has now been resolved. The Manager UI and API values are now consistent. |

Resolved issues in 3.14.7 October Maintenance

Table 14. Resolved issues

| Issue | Description |
|-----------|---|
| COS-61622 | In prior releases, the "Authorized IP Addresses" property on the mirror template page was not working. When creating a new Mirror using a mirror template using the UI or API, this field was getting ignored. This issue has now been resolved |

Resolved issues in 3.14.7

Table 15. Resolved issues

| Issue | Description |
|-----------|---|
| COS-56013 | Resolved an issue where S3Proxy endpoints was returning "+" instead of spaces in object names. |
| COS-39321 | Handle hungTasks on Slicestors. |
| COS-46393 | OS drive duplicate entry after OS drive replacement. |
| COS-48121 | RAID Drive replacement with DLM down results in duplicate entry in a bay. |
| COS-57129 | Through the System Properties Configuration on the Manager UI, the user can set how long to wait before generating an event when the rebuilder agent is hung (default is 1 hour). In prior releases, the customization of this value was not working. |
| COS-51588 | Client certificate challenge during TLS authentication is now disabled by default, and automatically enabled on systems with user PKI or user certificates. |
| COS-48768 | Management vault to Accesser device assignment now visible on the access pool page. All management vault to pool or device assignments now available from the Configure Management Vault page. |

Table 15. Resolved issues (continued)

| Issue | Description |
|-----------|--|
| COS-57607 | In prior releases the system returns a HTTP response code 500 when performing a REST.GET.VAULT request with a missing accessKeyId, . Starting with 3.14.7, the system returns a HTTP response code 403, which is the correct behavior. |
| COS-27469 | When performing a PUT-COPY operation, a request header is used to specify the source of the copy operation. If this header is specified, but with an empty value, the request is expected to fail with an HTTP 400 [®] - Bad Request. Instead, the object is being successfully created but with empty content. |

Resolved issues in 3.14.6 September Maintenance Release

Table 16. Resolved issues

| Issue | Description |
|-----------|--|
| COS-60327 | Resolved an issue that was causing devices to run out of memory when TLS was enabled which resulted in process restarts. This issue was more prevalent on devices that had smaller amounts of available memory and were establishing a large number of connections to Slicestor devices. |
| COS-50846 | OS drive OFFLINE manager UI event does not clear after OS disk replacement. |

Resolved issues in 3.14.6

Table 17. Resolved issues

| Issue | Description |
|-----------|---|
| COS-55417 | A "Certificate PEM" box has been added to the SMTP Configuration page on the Manager UI that allows a user to specify their own set of certificates for trusting the connection to the mail server. If no certificates are specified, the default trust store will be used, which contains all the standard global Certificate Authorities. No audits will be generated if certificates are changed, consistent with other features that allow for certificate trust configuration (e.g., LDAP). |
| COS-50636 | In prior releases, when a storage pool is 7-wide or 9-wide and there is no vault on the pool, the storage pool is incorrectly assumed to be a Concentrated Dispersal pool. This issue has now been resolved. |
| COS-51725 | The manager now supports syslog alert forwarding over UDP. Existing syslog configurations will continue to use TCP, with no changes. The syslog alert forwarding settings will default to UDP for any new syslog configurations. If new syslog configurations are required to use TCP, this must be stated explicitly in the IP Addresses/Hostname text box for syslog configuration, for example, tcp://a.b.c.d:port. If this is not specified explicitly for a new syslog configuration, UDP will be used by default. |
| COS-56013 | Resolved an issue where S3Proxy endpoints was returning "+" instead of spaces in object names. |
| COS-46401 | In prior releases, for a protected vault, changes to vault properties, other than retention (for example, description), performed through the edit vault page failed, if the system retention settings were out of range compared to the vault's retention settings. This issue has now been resolved. |

Resolved issues in 3.14.5 July Maintenance Release

Table 18. Resolved issues

| Issue | Description |
|--------------|---|
| COS-57359 | Resolved an issue where a reboot of IBM Cloud® Object Storage Slicestor12, Slicestor53 or Slicestor106 models, the network connectivity is lost and services can not be restarted. On IBM Slicestor 2584 appliances, the serial number is unavailable after a reboot. Recovery of this issue previously required the enclosures to be power cycled. As this issue has been addressed, this is no longer required. |

Resolved issues in 3.14.5

Table 19. Resolved issues

| Issue | Description |
|--------------|---|
| COS-54415 | When using Internet Explorer, the data are cached on first load for the following pages: access pool summary, site summary, and drives. The server now sends cache-control headers for those pages' requests. |
| COS-27974 | Fixed an issue where there is Mismatched Diagnostic reason for file system issues on Manager device summary and open incident. |
| COS-47041 | Fixed an issue where Stale entry of the replaced drive persists after failing disk migration was going on that drive, which was eventually replaced. |
| COS-52236 | Fixed an issue where Notification Service incidents do not open when expected. |
| COS-54475 | Fixed an issue where change in HDD model name results in duplicate drive entries. |
| COS-52443 | Fixed an issue where DLM reports duplicate drives with ONLINE/ONLINE state - None Model Name. |
| COS-43509 | Event message for virtual drives have been updated to display disk status. |
| COS-49906 | Unexpected feature interaction leads to user being forced to enter a notification service topic when they should not be forced to do so. |
| COS-46315 | UI elements related to notification service functionality is now hidden when notification service feature is not in use. |
| COS-52039 | Vault summary report columns related to notification service properties are defective. Changed column headings. Added missing "configured topic" to row. |
| COS-50955 | User must supply name when editing a notification service. If the value is not provided, the old previous value is used. |
| COS-49529 | Altered behavior such that the TLS certificate property of notification services is not required. |

Resolved issues in 3.14.4 July Maintenance Release

Table 20. Resolved issues

| Issue | Description |
|-----------|--|
| COS-57359 | Resolved an issue where a reboot of IBM Cloud Object Storage Slicestor12, Slicestor53 or Slicestor106 models, the network connectivity is lost and services can not be restarted. On IBM Slicestor 2584 appliances, the serial number is unavailable after a reboot. Recovery of this issue previously required the enclosures to be power cycled. As this issue has been addressed, this is no longer required. |

Resolved issues in 3.14.4

Table 21. Resolved issues

| Issue | Description |
|-----------|--|
| COS-49593 | Fixed an issue where a race condition may cause on-heap resource permits to accumulate on an Accesser device and never be released, resulting in HTTP 503 errors. |
| COS-48342 | The Device API (/statistics) reports system load averages with a 100x multiplier in the "loadX" key in the API, which differs from typical output found from "uptime" or "top." This is required for legacy reasons. |
| COS-48308 | When using the DMS or vault proxy features with the system in a mixed version scenario, a background service in the manager may upgrade vault formats prematurely. If this occurs, 500s would occur during writes from the upgraded source vault to the destination vault (DMS), or during reads to the upgraded proxy vault. |
| COS-53463 | In releases after 3.14.1 and prior to this release, for some scenarios, the core software process does not shut down during upgrade within the configured timeout. In these cases, the manager application does not automatically stop the core software process. This issue has now been resolved. For releases without this fix, the workaround is to manually force stop the device by clicking the "Force Quit" button. |
| COS-48947 | Today, we return a status code 507 when a user tries to create more containers in a storage account than is allowed. By comparison, AWS returns a HTTP 400 - TooManyBuckets error: TooManyBuckets You have attempted to create more buckets than allowed. 400 Bad Request Client We should evaluate whether it would be appropriate to change our behavior to match this expectation. A 4xx class "client error" status code seems more appropriate than a 5xx class "server error" status code. |
| COS-53785 | Fixed an issue where a HEAD request with a ranged read header would ignore the requested range and return the content length of the entire object. |
| COS-55935 | Fixed an issue where a HEAD request for a portions of an object via a part number query would ignore the part information and return the content length of the entire object. |

Table 21. Resolved issues (continued)

| Issue | Description |
|-----------|--|
| COS-50845 | In the Manager UI, a new section has been added under the Administration tab in the System Properties Configuration section called Slicestor Storage Engine. This functionality allows a user to set the default storage engine globally. This storage engine is applied to Slicestor devices during approval. Once set, the storage engine cannot be updated without a device reimage. A user can override the default storage engine by selecting the desired storage engine from the drop down available in the device approval section on the Home and top-level Configure pages. |
| COS-40727 | <p>In the Manager UI, the Allocated usage displayed on the Site, System, Storage Pool, and Slicestor Device pages has been updated. In prior ClevOS versions, the Allocated usage represented only the user (vault) data, but starting in ClevOS 3.14.4, the Allocated usage also includes the reserved file system data. In addition, the Manager UI will show an increase in Capacity to account for this change. The Unallocated usage will remain the same.</p> <p>Also, another change introduced in the Manager UI is to display the Reclaimable space on the Slicestor Device page. This usage represents data on disk that has been marked for deletion and is available to be compacted at some point in the future. The COS system has logic in place to determine when this data should be compacted.</p> |

Resolved issues in 3.14.3 June Maintenance Release

Table 22. Resolved issues

| Issue | Description |
|-----------|--|
| COS-56191 | Resolved an issue where incorrect drive model data would result in a drive being shown twice in the drive list. As part of this change, any prior state associated with the invalid drive model data is corrected. |

Resolved issues in 3.14.3 May Maintenance Release

Table 23. Resolved issues

| Issue | Description |
|-----------|--|
| COS-52911 | Fixed an issue where size enforcement was not being performed properly when performing a multipart upload part upload by copying a part from an existing object. |

Resolved issues in 3.14.3 April Maintenance Release

Table 24. Resolved issues

| Issue | Description |
|-----------|--|
| COS-52590 | Resolved an issue with slices being incorrectly processed during reallocation operations. |
| COS-50168 | A STARTTLS option has been added to the "New Configuration" section of the SMTP Configuration page in the Manager. When configured, the Manager follows the STARTTLS protocol, which initiates a secure connection before authenticating to the mail server. |

Resolved issues in 3.14.3 March Maintenance Release

| <i>Table 25. Resolved issues</i> | |
|----------------------------------|--|
| Issue | Description |
| COS-46163 | Fixed an issue where malformed requests with aws-chunked Content Encoding may encounter an exception during processing, causing the core process on the accesser appliance to restart. |

Resolved issues in 3.14.3

| <i>Table 26. Resolved issues</i> | |
|----------------------------------|--|
| Issue | Description |
| COS-39184 | After triggering a storage pool expansion, set replacement or set removal, the audit indicating "The storage was modified. The size was changed from size1 to size2" can show incorrect size values. |

Resolved issues in 3.14.2 February Maintenance Release

| <i>Table 27. Resolved issues</i> | |
|----------------------------------|---|
| Issue | Description |
| COS-47620 | Fixed an issue where Accesser devices may leak a small amount of memory when processing requests via HTTPS, leading to an out of memory condition and core process restart. |

Resolved issues in 3.14.2

| <i>Table 28. Resolved issues</i> | |
|----------------------------------|---|
| Issue | Description |
| COS-46326 | In previous releases the default connection count for all connections between ClevOS devices was 1, and this was only increased when a network latency of greater than 20ms was detected. 3.14.2 has changed the default behavior to always use 8 connections no matter what the detected latency is. This new default is a configuration which was already being manually applied to many customer systems to fix various performance problems. The main impact of this change is that client connections will now allow 8 times more outstanding messages to a server before it begins to queue to that server. In addition, the change allows more parallelism for TLS processing. For these reasons this configuration often fixes situations where requests were responded to with error code 503 unnecessarily. |
| COS-46740 | Deletion of a retention-enabled vault is only allowed if the manager can verify the vault is empty. In prior releases, the check to determine whether the vault is empty would fail if the Access Pool had a customer-supplied certificate chain configured. As a result, the manager would not allow the deletion of a retention-enabled vault, even when empty. This issue has now been resolved. |
| COS-41545 | As part of System NTP Configuration, entering a comma separated list of NTP servers in the External NTP Servers field saves the comma as part of the NTP Server. The NTP server plus comma is rejected as an NTP server, resulting in it not being listed in ntpq -pn output and not taking effect. This issue impacts both the Manager REST API and UI. |

Table 28. Resolved issues (continued)

| Issue | Description |
|-----------|--|
| COS-47362 | In prior releases, incidents did not include severity information in email alerts. Severity and cleared information are now available in the Subject of the email alert. Examples: Subject: [dsnet-alert] Incident updated with critical event for storage pool SP1 Subject: [dsnet-alert] Incident closed with cleared event for storage pool SP1 |
| COS-48911 | In prior releases, email alerts did not include the Access Pool name as part of the Subject. This issue has now been resolved. Example: Subject: [dsnet-alert] Incident opened with critical event for access pool AP1 |
| COS-44824 | In prior releases, the Event Console Manager REST API with the streamTypes parameter set to openIncidentsCurrentState may fail with a 500 HTTP response. This issue has now been resolved. |

Resolved issues in 3.14.1 February Maintenance Release

Table 29. Resolved issues

| Issue | Description |
|-----------|---|
| COS-47620 | Fixed an issue where Accesser devices may leak a small amount of memory when processing requests via HTTPS, leading to an out of memory condition and core process restart. |
| COS-50675 | Resolved an issue where certain mid-stream IO errors that are encountered while using HTTPS were being incorrectly reported as a HTTP 500 error code. |
| COS-48112 | Resolved an issue where slice reallocation between Slicestor devices may be erroneously marked as complete if the storage service was shutdown during the reallocation process. |

Resolved issues in 3.14.1 January Maintenance Release

Table 30. Resolved issues

| Issue | Description |
|-----------|---|
| COS-48885 | In Release 3.13.4 and subsequent releases prior to this release, the Manager UI can report incorrect BGP load balancing router status, inconsistent with the load balancing router status provided by an Accesser device. This issue has now been resolved. |
| COS-48298 | Resolved an issue that would cause various system processes to hang. This could result in components of the system silently becoming inoperative. It is recommended that all affected systems be upgraded. |
| COS-46718 | Resolved an issue where services in the drive monitoring subsystem on Slicestor devices were slow to start. This resulted in errors that prevented the drive monitoring subsystem from initializing properly, which in turn prevented drives from being brought online. |
| COS-46799 | Resolved an issue where Slicestor devices were erroneously detecting USB drives and attempting to process them as data drives. This resulted in existing drives sometimes being reported in invalid states and/or an inability to detect drive removals or insertions. |

Resolved issues in 3.14.1 December Maintenance Release

Table 31. Resolved issues

| Issue | Description |
|-----------|--|
| COS-48002 | S3 Extended API now supports query parameter of 'extended' in addition to 'pagination' which was previously supported. |

Resolved issues in 3.14.1

Table 32. Resolved issues

| Issue | Description |
|-----------|---|
| COS-41430 | If a device doesn't respond to a manager's "Force Kill" request during an upgrade, the manager will no longer initiate upgrades on devices that are waiting in the upgrade queue. The manager will also be unable to remove devices from the upgrade queue. This issue is resolved in this release. |
| COS-45556 | BMC Status Missing from Statistic API. This issue is resolved in this release. |
| COS-43901 | Resolved an issue where Put-Copy request between two different compliance enabled Mirror causes 500 Error. |
| COS-45018 | Resolved an issue where Presigned URL for PUT object and POST(form) returns 403 SignatureDoesNotMatch |

Resolved issues in 3.14.0

Table 33. Resolved issues

| Issue | Description |
|-----------|--|
| COS-41430 | If a device doesn't respond to a manager's "Force Kill" request during an upgrade, the manager no longer initiate upgrades on devices that are waiting in the upgrade queue. The manager is unable to remove devices from the upgrade queue. This issue has now been resolved. |
| COS-12691 | Instability has been observed when running two 40 Gbit links in LACP mode. |
| COS-12983 | Virtual devices running ClevOS within VMware may experience a kernel panic when migrating the virtual machine to a new server using VMware (R) vMotion (tm). |
| COS-16114 | On systems with RAM roughly equal to or greater than the size of the OS drive, a kernel panic may result in the system being in an unusable state. |
| COS-41035 | In 3.13.4 with a mixed release system containing devices on a lower release compared to the Manager, when a drive is failed from the UI, the Monitor Device page displays an incomplete message "diskFailSuccess." |
| COS-1749 | After recovering from an unresponsive IPMI controller, the open incident in the Manager event console sometimes fails to clear. The open incident is misleading, but has no impact on the system operation. |

Chapter 17. Known issues

Table 34. Known issues

| Issue | Failing Condition | Disposition |
|--------------|---|---|
| COS-66590 | Due to an issue with the BMC firmware on IBM Cloud Object Storage Gen2 Hardware (A10/M10/C10) it may result in the node failing to successfully boot up following a reboot. | This issue will be fixed in a future release. |
| COS-67728 | When searching for an Organization via the Manager User Interface, an autocomplete list of items is displayed for all Organization names that match. However, when an item is selected from the list, the user is redirected to the wrong page. | Instead of selecting an item from the autocomplete list directly, hit enter (return) after the item being searched. This will display the Search Results page, where links to all matched items are displayed. From this page, click on the link of interest. |
| COS-56141 | In some circumstances, audit messages that are displayed in the Event Console may not show the correct value for a configuration change, particularly SNMP polling versions and alert forwarding SNMP configuration changes. | A follow-on release will address this issue. |
| COS-50579 | There is a known issue where slice data being reallocated from one Slicestor device to another would not be appropriately removed from the source Slicestor device if the reallocation process was erroneously marked as complete." | This issue still exists in 3.14.3 because the change was reverted in the latest fix. |
| COS-11201 | In the Event Console of the Manager User Interface, the event details section for failing disk migration events contains a parameter that is called Migration Progress. However, it is not clear what this value represents. | This value corresponds to the percentage of failing disk migration that is complete. |
| COS-11355 | Replacing a failed drive with another failed drive results in an inconsistent view on the Manager User Interface. On the Monitor Device page, in the "Summary of device health" section, both the replaced failed drive and the new failed drive are shown. The "Drive Information and Actions" view of the drive layout shows the replaced failed drive. On the Maintenance page, the FRU report contains the replaced failed drive. | Perform another replacement of the failed drive with a good drive. |

Table 34. Known issues (continued)

| Issue | Failing Condition | Disposition |
|-----------|--|---|
| COS-13575 | The "stop migration" operation for failing disk migration on the Manager User Interface (UI) can take ~20 seconds to complete after being initiated by the user. The button continues to be enabled during this time. This issue exists for dispose and reset disk operations as well. | Do not hit the button again until the operation completes. If the drive stays in the same state for more than 20 seconds, perform a refresh of the page. If the drive continues to stay in this state, follow the recommended action that is provided in the Manager Administration Guide under disk lifecycle management. |
| COS-10031 | When resuming a drive in the DIAGNOSTIC state from the Manager User Interface, it can take ~20 seconds to complete. The resume button is not disabled during this time. | Do not hit the resume button until the operation completes. If the drive stays in the DIAGNOSTIC state for more than 20 seconds, perform a refresh of the page. If the drive continues to stay in this state, follow the recommended action that is provided in the Manager Administration Guide under disk lifecycle management. |
| COS-10445 | When using the storage command from the localadmin shell on a Slicestor device, it is possible to resume all drives that are currently in the DIAGNOSTIC state. However, in some cases, this process can take too long, which will cause the command to return an error code -15 due to a timeout. | Despite the error, the resume process is continuing in the background. The storage list command can be used to monitor the progress of resume process. |
| COS-7488 | When performing a storage pool set removal, it is possible that once the reallocation has finished for a source Slicestor device, it can show some small amount of data still present. | No action is required. Once the set removal has completed, all slices have been reallocated to the new storage pool. Any discrepancy in a Slicestor device's used space is generally a result of small inaccuracies that can occur during normal usage of the system. |
| COS-13504 | When failing a quarantined drive, it is possible that after data has been migrated off the failing drive, the Manager event console will report that no data migration was attempted. | No action is required. Despite the event description, data migration will always be attempted unless the user specifically chooses to skip migration via the localadmin shell storage command. |
| COS-22921 | When someone attempts to delete a bucket they first need to determine the assessor that can be used to issue the command. The S3 GET Bucket Location is one means to determine this. However, this command cannot work at every access pool. | Enhancing the S3 GET Bucket Location as a corner case command that can work at any access pool will be addressed in a future release. |

Table 34. Known issues (continued)

| Issue | Failing Condition | Disposition |
|-----------|--|--|
| COS-22990 | <p>The S3 remote proxy implementation of vault proxy has a few limitations that are related to communicating with an Amazon S3 endpoint. The version of the AWS SDK used to communicate to Amazon defaults to using V2 instead of V4 authentication, causing authentication issues when communicating with certain AWS endpoints.</p> | <p>For further assistance in configuring a remote proxy for use with Amazon S3, contact IBM customer support.</p> |
| COS-23025 | <p>SL 4U Slicestor devices, LEDs are incorrectly set.</p> | <p>Recovery Action: The user can use MegaCLI/storcli commands to issue LED actions before performing disk replacements. This will be fixed in a future release.</p> |
| COS-23962 | <p>Vault quotas are static and do not update when storage pool capacities change. If a system expansion, set replacement, or set removal is performed on the storage pool, vault quotas for any vaults on that pool will not update to consider the new capacity.</p> | <p>The user-defined vault quotas work as expected. However, they cannot be consistent with the current storage pool capacity. For example, a vault quota can be higher than total storage pool capacity after a set removal.</p> |
| COS-22924 | <p>When you upgrade the Manager to ClevOS 3.10.1 or newer for the first time, you might not be able to log in immediately. The Manager application might need an extra 20 - 30 minutes to become available due to database schema changes introduced in ClevOS 3.10.1. On systems with large databases, particularly systems with considerable historical event content, the time can be longer.</p> | <p>Contact Customer Support if it takes longer than 30 minutes to successfully log in to the Manager. Do not attempt to restart the Manager while it is upgrading.</p> |
| COS-26214 | <p>Lack of documentation highlighting dependencies of Hadoop-connector package with GA releases.</p> | <p>For legacy customers who are still using Hadoop connector for ClevOS software, please contact IBM customer support to install a new package compatible with latest build.</p> |
| COS-29681 | <p>When using the Microsoft IE9 web browser, certain Manager user interface elements like the left navigation panel and the vault capacity bar charts on the Monitor Vault page cannot appear.</p> | <p>Microsoft has ended support of IE9 and IE10. Users should upgrade to Microsoft IE11 or higher, or use an alternative browser, such as Firefox, Safari, or Chrome.</p> |
| COS-40881 | <p>The Manager REST API Edit Authentication Mechanism does not correctly update the value of the Hiding Secret Access Key flag and returns a status code 200. The flag is visible on the Security tab of the Manager UI.</p> | <p>This issue is resolved in a subsequent release.</p> |
| COS-64358 | <p>If an Accesser device restarts during a cycle , a small number of object deletions may be delayed up to 72 hours.</p> | <p>This issue will be fixed in a future release.</p> |

Table 34. Known issues (continued)

| Issue | Failing Condition | Disposition |
|-----------|---|--|
| COS-65701 | The first time you modify a device, you may see an audit message that indicates the management vault was unassigned. This is benign. Note that modifying a storage pool could also trigger this audit message for all devices in that pool or a subset of them. The Configure Device page on the Manager UI will show the correct management vault. | This issue will resolve itself as configuration changes are made to the devices in the system. |

Upgrading and Installation

Table 35. Upgrading and Installation

| Issue | Failing Condition | Disposition |
|-----------|---|---|
| COS-7126 | When extracting of upgrade file fails when a device is upgrading the failure message "The Selected File cannot be extracted while upgrades are in progress" continue to show if upload is restarted. | Only one upgrade file can be uploaded to the manager at a time. If another file is uploaded during an upgrade, an error message appears until the page is reloaded. |
| COS-15372 | When upgrading from ClevOS 3.8.x, 3.9.x, or 3.10.0 to 3.10.1 or later, all drives not used for Slicestor data (for example, OS drives) will be reported as newly discovered in the Manager event console. | No action is required. |

Container

Table 36. Container

| Issue | Failing Condition | Disposition |
|-----------|--|--|
| COS-1852 | When attempting to write an object to a container that does not exist, the Accesser appliance returns an HTTP 404 response with an error message of NoSuchKey instead of the appropriate NoSuchBucket. This includes cases where the container name includes a "/". | Ensure that your vault or container is successfully created before attempting to write objects to it. If you receive an error message of NoSuchKey for an upload request, verify that the container you are addressing does exist. |
| COS-15401 | If a user attempts to create a management vault by using "manual configuration" (accessed through the Configure Management Vault page) based on an existing vault template, management vault creation fails with the following message: "Cannot create a management vault from this template. It is deployed to access pools with standard vaults" | Use the "automatic configuration" available on the Configure Management Vault page. |
| COS-15218 | Container creation or deletion can sometimes result in 500 error responses when the requests are sent concurrently with other configuration requests to the same storage account. | Retrying the request that received a 500 is a suggested recovery action. It's best to retry the request when not doing other operations on the same storage account. |

Alerting and Reporting

| Table 37. Alerting and reporting | | |
|----------------------------------|--------------------|-------------|
| Issue | Failing Condition | Disposition |
| | Nothing to report. | |

System Behavior

| Table 38. System behavior | | |
|---------------------------|---|---|
| Issue | Failing Condition | Disposition |
| COS-2498 | The usage of a disk is counted while the disk is offline. However, its capacity is not counted. | No action. Awareness of limitation. If necessary, a restart of core would fix the usage values. Limit DLM events. |
| COS-2128 | In a GDG configuration with high request latency to the remote stores and low latency to local stores, an Accesser Appliance opens multiple connections to the remote stores and a single connection to local stores. Large bursts of IO can overwhelm the single local connection, resulting in elevated response times and operation latencies. | Using the System Advanced Configuration framework, the Accesser Appliance can be configured to open multiple connections to local stores, allowing it to better handle burst of IO activity. The parameter to configure appropriately is network.connection-profile. Please refer to section 3 of the Advanced System Configuration guide for more details. |
| COS-1920 | Support for "encoding-type" header when performing xml-based listing requests is not currently provided. | This feature is not currently supported. |

Storage Pools

| Table 39. Storage pools | | |
|-------------------------|--|--|
| Issue | Failing Condition | Disposition |
| COS-2642 | On the *Monitor Storage Pool Page, the Reallocation Progress graph, which displays historical data, is inaccurate when a device is down or statistics are not collected for a window of time. | The Data Reallocation progress bar, available at the top of the *Monitor Storage Pool Page, is always accurate. This view reflects the status and should be used to monitor progress of the data reallocation activity. |

Data Evacuation

| Table 40. Data evacuation | | |
|---------------------------|--------------------|-------------|
| Issue | Failing Condition | Disposition |
| | Nothing to report. | |

System Configuration

Table 41. System configuration

| Issue | Failing Condition | Disposition |
|-------|--------------------|-------------|
| | Nothing to report. | |

Deleting objects

Table 42. Deleting objects

| Issue | Failing Condition | Disposition |
|-------|---|--|
| 9444 | If a system is 100% full, customers might encounter an HTTP 500 error if they attempt to delete objects larger than the embedded content threshold (<1MB S3, >4MB SOH for default segments size). This issue has existed since release 3.0. It occurs because deleting large objects causes an intermediate write that appears larger to a Slicestor® Node, causing that node to fail the request due to an insufficient space error. | Contact IBM Support. They must use a development-provided procedure to free up disk space. |

Manager Web Interface

Table 43. Manager Web Interface

| Issue | Failing Condition | Disposition |
|-----------|---|---|
| COS-13189 | For drives that do not have a SCSI name, some Disk Lifecycle Management (DLM) actions, such as resume and fail, performed through the Manager User Interface (UI) will fail. | Use drive serial number to perform the action from the command line. Obtain drive serial number information by executing (see SERIAL column): # storage list Perform the operation based on the drive serial number (Z29010L5), for example: # storage fail Z29010L5 |
| COS-10031 | When resuming a drive in the DIAGNOSTIC state from the Manager User Interface, it may take ~20 seconds to complete. The resume button is not disabled during this time. | Do not hit the resume button until the operation completes. If the drive stays in the DIAGNOSTIC state for more than 20 seconds, perform a refresh of the page. If the drive continues to stay in this state, follow the recommended action that is provided in the Manager Administration Guide under disk lifecycle management. |
| COS-23764 | Upon network failure while going through the one time setup process in the manager, a network error page appears. When the network comes back, reload the page, at which point an internal server error page appears in some scenarios. | Log out of the internal server error page and log back into the manager, which will take you through one time setup again. |

Vaults

| <i>Table 44. Vaults</i> | | |
|-------------------------|--------------------------|--------------------|
| Issue | Failing Condition | Disposition |
| | Nothing to report | |

Vault Mirrors

| <i>Table 45. Vault mirrors</i> | | |
|--------------------------------|--|---|
| Issue | Failing Condition | Disposition |
| COS-7019 | When performing IO against a vault mirror with synchronous writes disable, HEAD requests performed against a successfully written object may return an HTTP 404 response. | If an HTTP 404 is returned for a HEAD request for a recently written object, please retry your request. |
| COS-13370 | Through the Manager User Interface (UI), after creating a mirror from a mirror template that has Authorized IP Addresses populated, the mirror does not contain the specified IPs. | Perform the following workaround. After the mirror is created, add the IPs using the Edit Mirror Access Control page. |

Vault migration

| <i>Table 46. Vault migration</i> | | |
|----------------------------------|--|--------------------|
| Issue | Failing Condition | Disposition |
| COS-12442 | When a vault migration finishes the work that is contained in its TODO queue, it kicks off a process to calculate the exact count of the number of objects that are migrated as part of the migration. This process of calculating the exact size is performed by each device in the target pool, and can take a long time to complete for large migrations. | |

Chapter 18. Supported Hardware Platforms

IBM Cloud Object Storage Appliances

Table 47. Minimum Version of ClevOS Compatible with Cleversafe Hardware Platforms

| Appliance | Product | Minimum ClevOS |
|---------------------------|---------|----------------|
| System Manager Appliance | M2100 | ≤2.7.0 |
| System Manager Appliance | M2105 | 3.2.2 |
| System Manager Appliance | M3100 | 2.7.0 |
| IBM COS Accesser® Device | A2100 | ≤2.7.0 |
| IBM COS Accesser® Device | A3100 | ≤2.7.0 |
| IBM COS Slicestor® Device | S1440 | ≤2.7.0 |
| IBM COS Slicestor® Device | S2104 | 3.2.1 |
| IBM COS Slicestor® Device | S2212 | 3.2.1 |
| IBM COS Slicestor® Device | S2440 | 3.0.1 |
| IBM COS Slicestor® Device | S4100 | 3.1.0 |

Table 48. Minimum Version of ClevOS Compatible with IBM Hardware Platforms

| Product Name | Machine Type (1Yr/3Yr Warranty) | Model | Minimum ClevOS |
|----------------------------------|------------------------------------|---------|----------------|
| IBM COS Accesser® 3105 | 3401/3403 | A00 | 3.8.1 |
| IBM COS Accesser® 4105 | 3401/3403 | A01 | 3.8.1 |
| IBM COS Accesser® 3110 | 4958/4957 | A10 | 3.14.4 |
| IBM COS Manager™ 2105 | 3401/3403 | M00 | 3.8.1 |
| IBM COS Manager™ 3105 | 3401/3403 | M01 | 3.8.1 |
| IBM COS Manager™ 3110 | 4958/4957 | M10 | 3.14.4 |
| IBM COS Slicestor® 2212 | 3401/3403 | S00 | 3.8.1 |
| IBM COS Slicestor® 2448 | 3401/3403 | S01 | 3.8.1 |
| IBM COS Slicestor®3448 | 3401/3403 | S02 | 3.8.3 |
| IBM COS Slicestor®2584 (AP-TL-1) | 3401/3403 | S03 | 3.8.1 |
| IBM COS Slicestor®2584 (AP-LS-1) | 3401/3403 | S03 | 3.13.1 |
| IBM COS Slicestor®2212A | 3401/3403 | S10 | 3.10.0 |
| IBM COS Slicestor®12 | 4958/4957 | C10/J10 | 3.14.4 |
| IBM COS Slicestor®53 | 4958/4957 | C10/J11 | 3.14.4 |
| IBM COS Slicestor®106 | 4958/4957 | C10/J12 | 3.14.4 |

Note: □ Requires RPQ

Hewlett Packard Enterprise

Table 49. Minimum Version of ClevOS Compatible with Hewlett Packard Enterprise Hardware

| Appliance | Model | Minimum ClevOS |
|-------------------------------|-------------------|----------------|
| Manager Appliance | DL360P Gen8 | 3.2.1 |
| Manager Appliance | DL360 Gen9 | 3.5.0 |
| Manager Appliance | DL380 Gen9 | 3.5.0 |
| Accesser [®] Device | DL360P Gen8 | 3.2.1 |
| Accesser [®] Device | DL360 Gen9 | 3.5.0 |
| Accesser [®] Device | DL380 Gen9 | 3.5.0 |
| Slicestor [®] Device | SL4540 Gen8 | 2.9.0 |
| Slicestor [®] Device | DL380 Gen9 | 3.5.0 |
| Slicestor [®] Device | Apollo 4200 Gen9 | 3.6.0 |
| Slicestor [®] Device | Apollo 4510 Gen9 | 3.6.0 |
| Slicestor [®] Device | Apollo 4510 Gen10 | 3.14.0 |
| Slicestor [®] Device | Apollo 4530 Gen9 | 3.6.0 |

Seagate

Table 50. Minimum Version of ClevOS Compatible with Seagate Hardware

| Appliance | Model | Minimum ClevOS |
|------------------------------|-------------------|----------------|
| Seagate OneStor [®] | AP-2584 1 AP-TL-1 | 3.4.2 |

Cisco

Table 51. Minimum Version of ClevOS Compatible with Cisco Hardware

| Appliance | Model | Minimum ClevOS |
|-------------------------------------|---------------------------------------|----------------|
| Cisco Slicestor [®] Device | UCS C3260 | 3.7.4 |
| Cisco Slicestor [®] Device | UCS S3260 (Single Node) | 3.12.0 |
| Cisco Slicestor [®] Device | UCS S3260 (Dual Node) | 3.12.0 |
| Cisco Slicestor [®] Device | UCS S3260 M5 (56 drive configuration) | 3.13.1 |
| Cisco Slicestor [®] Device | UCS S3260 M5 (60 drive configuration) | 3.14.3 |
| Cisco Manager Appliance | UCS C220 M4 | 3.12.0 |
| Cisco Accesser [®] Device | UCS C220 M4 | 3.12.0 |
| Cisco Manager Appliance | UCS C220 M5 | 3.13.6 |
| Cisco Accesser [®] Device | UCS C220 M5 | 3.13.6 |
| Cisco Slicestor [®] Device | UCS C240 | 3.13.6 |

Dell

Table 52. Minimum Version of ClevOS Compatible with Dell Hardware

| Appliance | Model | Minimum ClevOS |
|------------------------|------------------------|----------------|
| Dell Slicestor® Device | DSS 7000 | 3.10.1 |
| Dell Slicestor® Device | R740xd w/ HDD Support | 3.14.1 |
| Dell Slicestor® Device | R740xd w/ NVMe Support | 3.14.2 |

Lenovo

Table 53. Minimum Version of ClevOS Compatible with Lenovo Hardware

| Appliance | Model | Minimum ClevOS |
|--------------------------|----------|----------------|
| Lenovo Manager Appliance | X3550 M5 | 3.10.1 |
| Lenovo Accesser® Device | X3550 M5 | 3.10.1 |
| Lenovo Manager Appliance | X3650 M5 | 3.10.1 |
| Lenovo Manager Appliance | SR630 | 3.13.6 |
| Lenovo Accesser® Device | SR630 | 3.13.6 |
| Lenovo Slicestor® Device | SR650 | 3.13.6 |

Quanta Cloud Technology (QCT)

Table 54. Minimum Version of ClevOS Compatible with QCT Hardware

| Appliance | Model | Minimum ClevOS |
|-----------------------|-----------------------|----------------|
| QCT Manager Appliance | QuantaGrid D51PH-1ULH | 3.13.4 |
| QCT Accesser® Device | QuantaGrid D51PH-1ULH | 3.13.4 |
| QCT Slicestor® Device | QuantaGrid D51PH-1ULH | 3.13.4 |

Chapter 19. Incompatible Hardware and Firmware with ClevOS

The hardware components running firmware revisions listed below are incompatible with ClevOS due to the possibility of unexpected behavior.

Note: If you have any hardware on this list running the firmware revisions listed, please contact L3 support immediately to create an upgrade plan. You can determine your firmware revisions using the Firmware Report that is found under the Maintenance menu.

Broadcom

| <i>Table 55. Broadcom Hardware and Firmware Incompatibility with ClevOS</i> | | |
|---|---------------------------|-------------------|
| Type | Model | Firmware affected |
| RAID Controller | Broadcom MegaRAID 9361-8i | 4.650.00-6121 |

Hewlett Packard

| <i>Table 56. HP Hardware and Firmware Incompatibility with ClevOS</i> | | |
|---|-----------------------|-------------------|
| Type | Model | Firmware affected |
| RAID Controller | HP-SL4540 Smart Array | 6.64 |
| iLO | HPE SL4540 Gen 8 | 2.30 |

IBM Cloud Object Storage Appliances

| <i>Table 57. IBM COS Hardware and Firmware Incompatibility with ClevOS</i> | | |
|--|--|------------------------|
| Type | Model | Firmware affected |
| USM | IBM COS Slicestor®2584 (AP-TL-1) 3401/3403 S03 | 4.1.7 |
| BMC | A3105, A4105, M3105, S2212A, S2448 | 1.0.125362, 1.0.135362 |

Seagate

| <i>Table 58. Seagate Hardware and Firmware Incompatibility with ClevOS</i> | | |
|--|-----------------------------|-------------------|
| Type | Model | Firmware affected |
| HDD | Seagate ST1000NM0033-9ZM173 | SN04 |

Supermicro

Table 59. Supermicro Hardware and Firmware Incompatibility with ClevOS

| Type | Model | Firmware affected |
|-------------|------------------------------|--------------------------|
| BMC | Supermicro SSG-6048R-E1CR60N | 3.60 |

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785*

US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Accesser[®], Cleversafe[®], ClevOS[™], Dispersed Storage[®], dsNet[®], IBM Cloud Object Storage Accesser[®], IBM Cloud Object Storage Dedicated[™], IBM Cloud Object Storage Insight[™], IBM Cloud Object Storage Manager[™], IBM Cloud Object Storage Slicestor[®], IBM Cloud Object Storage Standard[™], IBM Cloud Object Storage System[™], IBM Cloud Object Storage Vault[™], SecureSlice[™], and Slicestor[®] are trademarks or registered trademarks of Cleversafe, an IBM Company and/or International Business Machines Corp.

Other product and service names might be trademarks of IBM or other companies.



Printed in USA