IBM Cloud Object Storage System
Version 3.10.0 for May Maintenance Release

Release Notes

IBM

COSRN_3800-20160516L

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

Contents

Support information

For more information on the product or help with troubleshooting, contact IBM Support at IBMCloudStorageSupport@us.ibm.com or visit the Directory of worldwide contacts.

Chapter 1. New Features and Improvements in ClevOS 3.10.0

Storage as a Service (STaaS) [546]

IBM® Cloud Object Storage product is capable of multi-tenancy, meaning a single offering of product, or application, or instance of software, which serves multiple customers, where each customer that uses this service is called a tenant.

With the Storage as a Service (STaaS) feature, a tenant can be a single user, a single enterprise with multiple users, partners, operators, or resellers who offer this service to its own customers in verticals such as medical, archive, and backup as a service. The STaaS providers or customers of the STaaS service provider (tenants) have the ability to customize their application UI and offering to their users by using APIs and logs for account management, authentication, usage information for billing/charge back, issue isolation, and so on, as needed. There are limitations with using existing features that are supported by STaaS feature (also known as container mode) enabled. The following features are not supported in STaaS mode.

- Accesser® Application
- Embedded Accesser
- Versioning
- Delete restrictions
- Mirrors
- Proxies
- Swift, SOH, and DDN WOS API
- Bucket tags

Manager Graph of Ops [42]

The Accesser request graph displays incoming HTTP requests to the Accesser device. Above the graph are controls to switch the graph mode and to filter the graph. The graph has two modes: request and response. In request mode, graph lines are displayed for each request type, and the filter control allows the selection of a particular response code. In response mode, graph lines are displayed for each response code, and the filter control allows the selection of a particular request type.

Improved Support for Unqualified Appliances [846]

In previous releases of ClevOS, system appliances that had not been previously qualified by IBM were required to be reimaged upon formal qualification and might not be upgraded. Unqualified appliances that are imaged with ClevOS 3.10.0 (not including upgrade) might be imaged and upgraded similar to a qualified appliance. If in future releases qualification is performed on said appliance, it benefits from that qualification upon upgrade. This functionality is intended to allow clients the ability to use unqualified appliances without delay and allow for deployment on a broader range of appliances. When using unqualified appliances, there are several things to be aware of when working with them.

Prior to ClevOS 3.10.0, unqualified appliances would not report a model name in the Manager interface. Starting with 3.10.0, unqualified appliances report a model name that is marked as unknown to distinguish it from a qualified appliance. Unqualified appliances that were imaged prior to 3.10.0 need to be reimaged to 3.10.0 before performing any upgrades.

When adding unqualified Slicestor® appliances to a storage pool, the administrator is required to specify the maximum number of data drives that the appliance can possess. This setting cannot be changed later

and require a full reimage to correct. Customer support should be contacted to ensure that the proper value is set to ensure that there are no unexpected issues during upgrade and the possible transition from unqualified to qualified as part of said upgrade.

When upgrading an unqualified Slicestor appliance to a ClevOS version in which said appliance was qualified, all pulled drives should be either replaced or disposed of by using the Manager interface. Failure to do so might result in the absence of the appliance's drive bay diagram in the Manager.

Enhanced Backup/Restore with Vault Config Data [633]

With this feature, configuration changes that are associated with create, update, and delete operations for vaults, mirrors, accounts, and groups are captured in a management vault and the manager database. All other configuration changes will still only be captured in the manager database. The feature is enabled by default. It starts storing statistics and the configuration changes above into a management vault, once the first user vault has been created on a storage pool. If a manager restore is performed, the configuration changes from the management vault since the last backup will automatically be applied. It is important that manager backups continue to be performed daily and when significant configuration changes occur.

Set Replacement/Set Resize [345]

Set Replacement allows the customer to remove a set or sets of Slicestor devices from a given storage pool and replace them with a new set of devices. This functionality allows hardware to be repurposed for other uses or fully retired from the IBM Cloud Object Storage System

Set Resize allows the customer to rebalance an existing storage pool when additional storage capacity is added to the Slicestor devices. In a partially populated set of devices, adding additional capacity and resizing allows the user to take advantage of this new capacity and evenly allocate data between existing device sets.

Add SSH/CAB keys [924]

The Manager now offers the ability to authenticate to a remote SFTP server that uses SSH keys when transferring a Manager backup file. Previously, a user password was required in order to perform the authentication. This functionality is available through the Manager UI and REST API. See the Manager Administration Guide and Manager REST API Guide for details.

Deploy with Missing Slicestor [555]

Prior to this release, if some Slicestor devices were not available during installation, the initial storage pool creation would be blocked. This feature allows an administrator to create or expand a storage pool with a small number of missing Slicestor devices in order to begin storing data. Once the devices become available, they can be placed into the storage pool where rebuilder will then rebuild any missing slice data.

Active Directory SRV Lookup & Logging Support [927]

This feature allows for automatic service discovery of Active Directory / LDAP controllers from standard DNS SRV records. This capability can be enabled with a new check box on the Active Directory configuration page in the Administration section.

Add Read-only mode [929]

This feature provides Read Only Access Control Lists for the existing System Administrator and Security Officer roles. These capabilities can be enabled through the “Read Only” column of the Create/Edit Account and Group pages on the Security tab or from the Manager REST API. Once the Read Only ACLs are assigned to these roles, users will not be able to access the create/update/delete pages of the UI or

APIs but are allowed to view all UI content and use all list/view APIs. For additional information, please see the Manager Administration Guide and the Manager REST API Guide.

Device Set Removal [897]

Set Removal allows the customer to remove a set or sets of Slicestor devices from a given storage pool. This functionality allows hardware to be repurposed for other uses or fully retired from the IBM Cloud Object Storage System.

Manager Session Timeout [910]

When automatic logout is enabled, the session timeout set in the Administration/System Preference section is respected in all sections of the web application except for the Maintenance/Upgrade page and the Administration/Manager Backup and Restore section. A pop-up window appears when the session is approximately one minute from expiring, which gives the user the choice to reset the session timeout. Navigating will also reset the session timeout. The pages in the monitor section will no longer reload every 5 minutes.

Enforcing Strong Encryption [894]

This feature provides Advanced Configuration support to customize Apache web server cipher encryption on the manager. It allows a user to specify cipher encryption, which is passed down using Advanced Configuration. Once passed down, Apache is restarted for the changes to take effect. Contact Customer Support to use this feature, including limitations and supported ciphers.

Drive Power Control by Using Manager [977]

Some Slicestor models support the ability to disable individual drive bays. This allows a drive to be powered down and taken out of service in the same way as physically removing it from the appliance. Previous releases of ClevOS supported the ability to utilize this functionality via the localadmin shell of an individual Slicestor appliance. Additional support has been provided to extend this functionality to the Manager interface. When utilizing appliances with the requisite hardware facilities, drive bays may be disabled from the Manager drive bay view as well as by means of the Manager REST API. At present, the following appliances support this facility:

- IBM Slicestor 2584 Seagate
- OneStor® AP-2584

40 Gbit Networking [218]

Optional 40 Gb networking is now supported on IBM Branded Cloud Object Storage hardware. Features such as PXE boot, network monitoring, and all other previously supported network features are enabled.

Cross Site Request Forgery [1032]

The manager now offers the ability to protect against CSRF attacks. This protection is off by default and can be enabled in the Manager UI in the Administration section.

Slicestor 2212A [571]

IBM Cloud Object Storage Slicestor S2212A is a 2U server that contains 12 HDDs. The HDDs can be 4 TB, 6 TB, 8 TB, or 10 TB in capacity. The S2212A is a medium-powered Slicestor that is ideal for small cloud Object Storage solutions that could be classified as cold or warm storage. In a standard 12-wide solution with an information dispersal algorithm (IDA) of 12/6/8, a system that uses S2212A Slicestors can provide usable capacity ranging from 288 TB to 720 TB. S2212A is a replacement for S2212.

Location Constraint Aliases [1080]

This feature allows an operator the ability to configure aliases for a location constraint provided during bucket creation. This can be used resolve bucket creation issues seen when using certain AWS SDKs that attempt to parse location constraint from the hostname of the accessor. Location constraint alias tables are applicable to Accesser devices and take on the form:

s3.location-constraint-aliases={"alias": {"provisioning-code": "us-standard", "vault-provisioning-group": "us", "default-storage-class": "standard"}}

Where "alias" is the location constraint alias that maps to the specified "provisioning-code".

Support REST API method for removing a device

This feature provides REST API support for removing a device. Refer to the Manager Administration Guide for additional information.

Chapter 2. Interface Modifications

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

- Manager REST API - Modified Account management section.
  - Updated List my access keys
  - Updated Create Group
  - Updated Edit Group
  - Updated Delete Group
- Manager REST API continued- Modified Device Management.
  - Updated List Device Registrations
  - Updated List Device
  - Updated Change Device Password
  - Updated Device Disk Nut Storage Action
  - Added Device Drive Bay Nut Enclosure Action
- Manager REST API continued- Access pool management.
  - Updated Create access pool
  - Updated Edit access pool
- Manager REST API continued- Storage pool management.
  - Updated Create storage pool
  - Updated Edit storage pool
  - Updated List storage pools
  - Updated Expand storage pool
  - Updated Delete storage pool
  - Updated Replace storage pool device
  - Added Resize storage pool sets
- Manager REST API continued- Added Remove storage pool sets.
- Manager REST API continued- Mirror management.
  - Updated Create mirror
  - Updated Edit mirror
  - Updated Create mirror from template
- Manager REST API continued- Vault management.
  - Updated Create vault
  - Updated Edit vault
- Manager REST API continued- Administration.
  - Updated Configure system backup settings
  - Updated Edit preference configuration
  - Updated Edit system keystone authentication settings
  - Updated Edit system multi-tenancy configuration
  - Updated System properties configuration
  - Updated View system
  - Updated View System Configuration
- Access Log - Modified Structured logs section.
  - Updated Multi-delete access log entry fields

- Added Storage account usage access log entry
- Added Container usage access log entry
- Updated Error message access log entry
- Updated Retention and purge policy

v CSO API - The Cloud Storage Object (CSO) Application Programming Interface (API) that enables application developers to use existing Amazon Simple Storage Service (S3) applications to access object vaults on a system. No change in the documentation for this release.

v Service API - The Storage as a Service (STaaS) feature delivers a base set of Service APIs that are intended for deployment, system management, and service operator usage. This API covers the interfaces relating to the management of storage accounts. No change in the documentation for this release.

v AWS Credentials Management - The Storage as a Service (STaaS) feature delivers a base set of Service APIs that are intended for deployment, system management, and service operator usage. This API covers the interfaces relating to the management of access keys. No change in the documentation for this release.

API Changes

New REST APIs for Set Replacement and Set Resize. Storage Pool Group response from View System API now includes:

v storagePoolRemoving, following by a list of storage pool IDs for the pool that are being replaced/removed

v eligibleForSetResize, either true or false if the storage pool can be resized

Chapter 3. Resolved Issues

Resolved issues in 3.10.0 May Maintenance Release

Table 1. Resolved issues

Issue	Description
COS-14255	Support REST API method for removing a device
COS-1259/15400	When performing a listing of a storage account by using the service API, the storage account enable/disable state is not displayed.

Resolved issues in 3.10.0 April Maintenance Release

Table 2. Resolved issues

Issue	Description
COS-8630	SNMP service does not start on ClevOS
COS-13361	3.10.0 Manager Restore does not Complete in Some Scenarios In prior 3.10.0 releases, a manager restore issue was observed when (A.) the previous manager had a management vault configured, (B.) a different IP/chassis than the new manager, (C.) on systems where a previous manager replacement had already been performed.
COS-13378	After upgrading, or restarting the Docker Manager, the login page sometimes does not appear.
COS-12285	For storage pools of width greater than 10 Mirror template vault width is not displayed properly at the Manager.
COS-8028	IBM COS Slicestor® 2584 - Failed with timeout after Upgrade from 3.7 to 3.8.
COS-11595	IBM COS Slicestor® 2584 , mpt2sas not all Drives are attached after reboot.
COS-12288-9829	Internal handling of the Prioritization Filter
COS-13786	High latency on HEAD.CONTAINER after upgrade to 3.10. This issue is addressed in this release.
COS-13655	When performing a deletion of large objects with many regions, it is possible to encounter an error condition that triggers a restart of the core process. The core process with automatically restart when this condition is triggered. Implemented stack unrolling within the software for this recursive operation to prevent this condition from occurring.
COS-14306	Docker Manager log rotation has been addressed in this release.

Resolved issues in 3.10.0

Table 3. Resolved issues

Issue	Description
14274	For client workloads that involve multipart upload on vaults deployed on file storage based storage pools, the multipart transaction index is used to determine whether or not a particular upload id exists in the system. Upon upgrading to this release, multipart index operations will use the index delegation feature by default.
14777	Versioning state watcher returns before new state is found in the registry
14591	Kernel panic detected on devices utilizing Avago MegaRAID SAS controllers..

Table 3. Resolved issues (continued)

14639	An issue has been seen where client listing operations, or ongoing Data Migration activities, in the presence of a failed drive can cause listing operations to queue up in memory. Over time, this can cause significant memory to be consumed, leading to out of memory conditions and a core process crash.

Issue	Description
14770	On the Edit Access Pool page within the Manager user interface, when access device or vault deployment changes are made and the "Update" button is selected on the lower action bar, a popup confirmation dialog appears outside the view of the web browser window.
14581	Device with quarantined drives are not included in the Communication error widget.
14828	When using the AWS .NET SDK to create a vault, the SDK client uses an improper xml schema when performing the put bucket request, causing the request to fail with a HTTP 400® error.
COS-2104	When issuing a PUT Bucket request for a vault or container that already exists, and a query parameter is provided with the request, the Accesser appliance will incorrectly return a HTTP 400 response code instead of a HTTP 409 response code.
COS-4095	The passwordAuthenticationEnabled parameter of the Edit Authentication Mechanism Manager REST API will not take effect when it is the only parameter used in the API. However, if the accessKeyAuthenticationEnabled parameter is used in conjunction with passwordAuthenticationEnabled, the passwordAuthenticationEnabled parameter will work properly.
COS-2824	Locked vaults reporting zero usage can be deleted from the UI/API. For name index enabled locked vaults, even after deleting all objects, vault still report some data due to left over root index node on the vault. In this scenario, user is still allowed to delete locked vault based on the following checks: if reported usage on vault is less than 1MB then a recovery listing is done to get the accurate object count on the vault, only when the count is zero the user is allowed to delete the locked vault.
COS-2983	Nut activation would fail and roll back if no global or local ssh keys were set .
COS-5562	The troubleshooting console may incorrectly filter devices when selecting a storage pool, access pool, or site with an ID >= 10, resulting in a larger number of devices than expected.
COS-5473	The Storage Pool Capacity and Disk Report API output shows the wrong value (always zero) for percentageOfFreeSpace in the All Sets section.
14438	When viewing the Message Acknowledge Time graph with Firefox, a portion of the legend appears outside of the graph area.
14396	When performing multipart upload requests, especially in cases where multiple parts are uploaded in parallel, it is possible to observe degraded performance and metadata contention when the number of parts per multipart transaction grows large.
14425	Drive failure LEDs are not functional on HP Gen9 appliances.
14436	When using AWSv2 authentication via the CSO API, if the user provides an access key id that is either null or an empty string, the system will respond with an HTTP 500 instead of the desired HTTP 403 response code.
14465	If data evacuation is paused and then resumed, the progress bar in the Manager UI does not increase until evacuation is completed.
13533	The IBM Cloud Object Storage System™ CSO API does not enforce that new vaults created through the PUT Bucket API method be created with a naming convention compliant with the Amazon S3 bucket naming restrictions.
COS-4094	PUT.COPY IO requests from the Hadoop / S3a native connector show up as PUT.OBJECTS in the Accesser access logs.
COS-5473	Storage Pool Capacity and Disk Report API output shows the wrong value for percentageOfFreeSpace in the All Sets section.

Table 3. Resolved issues (continued)

Issue	Description
14530	When expanding a storage pool that uses Slicestor devices in locations other than the locations of the current Slicestor appliance, it is likely that the pairing of Slicestor devices for data reallocation would cross data center boundaries that can cause unintended WAN traffic. In ClevOS 3.9.0, when expanding a storage pool with a new set of devices, the Manager user interface displays an expansion confirmation page. The page displays all the existing device sets and the new device set in such a way that the representation visually aligns the sites and devices based on the manner in which data reallocation are performed. This page also shows the user how the data is reallocated from existing devices to new devices, which helps identify reallocation that has high latency, have low bandwidth, span large geographic distances, or cross national boundaries.
COS-1115	Fixed an issue where slow or zombie stores can cause a delegated index operation to be delayed in its execution, and race with a subsequent user operation. Performing concurrent write and delete operations for the same object can result in inconsistencies between the index and the object metadata.
COS-892	When in container mode, the system accepts non-compliant container names. In ClevOS 3.9.0. system can be configured to enforce DNS-compliant container names when creating new containers.
COS-1920	We currently do not support the "encoding-type" header when performing xml-based listing requests.
COS-653	When performing sustained delete operations in the presence of a slow or unresponsive Slicestor device, memory resources on the Accesser appliance might be consumed, resulting in an out of memory exception and restart of dsnet-core process.
COS-867	When a Cloud Storage Object request is signed with AWS Auth SigV4 and includes an empty date header, but a proper x-amz-date header, the request fails with a 403 response.
COS-4857	Presigned V4 GET requests with additional query parameters is rejected with 403 response code. The sorting of Canonical Query Parameters is now changed to a "Case sensitive" sorting which fixes the issue.
COS-7056	When removing a registered appliance from the system, improved messaging has been provided that indicates the device must be re-imaged before using it again.
COS-4163	Added the ability to lock down the listVaults API by default with a checkbox on the Administration -> System Properties page that allows the API to be open again for use by the Accesser Application. Checking the box in the Advanced System Configuration will allow anonymous access to both XML list vault and JSON list Storage Pool Internal Usage that the Accesser Application uses.
COS-6089	Provide single value for current drive Thresholds. The Drive Report API has been updated to include the drive error threshold and drive warning threshold currently in effect for every device in the response. Note that these new values account for hardware-based settings and individual device level settings.
14856	This defect could potentially occur on very old storage pools which were created prior to write threshold becoming a required configuration parameter. If a write threshold was not set, the system would automatically configure one as width - threshold + 1 for the vault's default write threshold. Under certain circumstances involving a combination of very high concurrency (100s to 1000s of threads) and very large object uploads (GB and larger), it is possible that multiple Slicestor appliances may experience disks being quarantined due to IO timeouts simultaneously. This is a direct consequence of the workload being too high for the system and is likely to occur under certain test conditions but is much less likely to occur in a production environment. If this occurs, resume the disks and resume IO but reduce the workload on the system.
COS-2655	Manager UI shows incorrect capacity of drives of File Accesser devices. Instead of reporting the individual drive capacity, aggregate capacity is displayed.
COS-3026	Fixed an issue where the object length field of the access log entry was not being reported correctly for complete multipart upload requests.

Chapter 4. Known issues

Table 4. Known Issues

Issue	Failing Condition	Disposition
COS-12691	Instability has been observed when running two 40 Gbit links in LACP mode.	Do not use LACP aggregated links with 40 Gbit Intel Network cards.

Container

Table 5. 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-7089	If a delete container request fails due to an index write failure, subsequent container listing request via the account API will temporarily fail with a 404 error.	The issue will be resolved by index background cleanup, but no immediate recovery action is possible.
COS-5390	The product does not currently support guaranteed delivery of access log or usage log entries to an end consumer.	Contact IBM Customer Support for more information.

Upgrading

Table 6. Upgrading

Issue	Failing Condition	Disposition
	Nothing to report.	

Alerting and Reporting

Table 7. Alerting and reporting

Issue	Failing Condition	Disposition
7598	In the following scenario, a drive is quarantined, pulled, permanently removed, disposed, Slicestor® Node powered down, drive replaced, and Slicestor® Node powered backup. The following incident appears and remains in the Open Incident view of the Manager Web Interface: .Open Incident for Removed and Replaced Drive ==== Disk in drive bay X with S/N Y is a previously removed disk ==== **`Disk in drive bay X with S/N Y is a previously removed disk`** endif::[]	Contact IBM Support to close the incident.
7714	The Storage Pool Capacity and Used graph on the Monitor storage pool page shows a temporary drop in the capacity at times, particularly during upgrade. When upgrading, this is caused by timing issues between the polling of values and when the node values stabilize.	Once node upgrades complete, the capacity returns to normal. The capacity drops can be correlated with upgrade events in the Event Console for nodes in the storage pool.

Table 7. Alerting and reporting (continued)

Issue	Failing Condition	Disposition
11739	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.	Contact IBM Support to confirm and correct the false incident.
12450	If a previously failed disk is reinserted into a Slicestor® Device and the system-core process is running, it generates an incident on the Manager indicating that a previously failed disk was reinserted. Normally, when said disk finally gets replaced, this incident clears. However, if this disk is replaced when the device is powered off or when, for any reason, the system-core process is not running, this incident will not get cleared but remain open forever.	If this situation occurs, contact IBM Support for assistance in manually clearing the incident.
COS-7370	On occasion, a management vault GET failure event appears in the Event Console on the Manager UI after vault creation.	This event can be ignored.
COS-6490	If a device is imaged with a degraded RAID array, no event is presented to the user in the event console. In some cases this can cause no warnings to be shown about a potential problem.	Repair the RAID array by replacing the failing drive.

System Behavior

Table 8. System behavior

Issue	Failing Condition	Disposition
10659	Some drive-related SNMP traps might not appear immediately. If these events occur during early boot of the appliance or while critical system processes are down, there is a delay in the delivery of these SNMP traps.	Once the appliance is fully up and running, the SNMP traps are delivered as expected.
14296	Under conditions where an unresponsive or zombie Slicestor® Device is present in the system, performing multiple large object uploads in parallel might cause uploads to hang. This is caused by a resource starvation issue, in which the outstanding write requests to the zombie store cannot be canceled, and the associated resources that are freed, until the large object upload completes.	This issue is mitigated on IBM Cloud Object Storage Accesser® appliances with larger amounts of memory available. As a workaround, ensure that any potential zombie or unresponsive Slicestor® Devices are dealt with promptly.
14383	It has been observed that for ZTDG Accesser Appliances there are instances where the system time is not properly reported to the application layer, causing negative values to be reported in the stat entry in the device's access log.	This does not affect the proper operation of the system and will be addressed in a future release.
COS-1478	It has been observed that read operations for large objects (20 MB or greater) are degraded (15-20%) relative to prior releases.	There is no mitigation for this issue currently. This issue will be addressed in a future release.

Table 8. System behavior (continued)

Issue	Failing Condition	Disposition
COS-5539	If a storage account is deleted and re-created with the same name, usage updates that are associated with the previous account might be applied to the new account.	Preventive Action: Always create accounts with unique IDs. Solution: Accounts will have an extra UUID to uniquely identify accounts, and usage updates will only be applied when the UUID matches the expected value. This change will be made in a future release.
14714	When performing heavy write IO to an empty vault, with index enabled and index delegation that is enabled, the index insertion operations on the index take priority over asynchronous split operations, possibly causing the nodes in the index to become large.	The user can avoid this issue by pre-filling the vault with objects. If this condition persists, this can lead to increased latencies and IO failures. Contact IBM Support to confirm scenario.
14783	Problem Under certain circumstances, Slicestor® Device devices that use MegaRAID SAS disk controllers might not immediately detect the removal or replacement of a drive.	A replacement drive should not be added until 90 seconds after the original drive is removed to allow the kernel time to finalize the removal of the device.
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-9955	Under certain circumstances involving a combination of high concurrency (100 s to 1000 s of threads) and large object uploads (GB and larger), it is possible that multiple Slicestor appliances might experience disks being quarantined due to IO timeouts simultaneously.	This is a direct consequence of the workload being too high for the system and is likely to occur under certain test conditions but is much less likely to occur in a production environment. If this occurs, resume the disks and resume IO but reduce the workload on the system.
COS-5794	Performance for object PUTs can sometimes be temporarily degraded when using heavy PUT workflows on containers with minimal amounts of fill.	Recovery Action: The degraded performance is temporary, so continued IO will eventually push through the problem. Temporarily stopping or slowing IO should help make the degraded performance go away faster. This problem has mostly been seen to occur with heavy IO load on many containers (around 40,000). Lighter loads, or loads with less containers result in less occurrences of this problem.
COS-2128	In a GDG configuration with high request latency to the remote stores and low latency to local stores, an Accesser Appliance will open 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-9922	Running the "nut enclosure bay list" command via the Manager troubleshooting console on 3.10.0 devices will result in "No such command" instead of "Platform not supported.."	Upgrading to next release will handle this issue.

Storage Pools

Table 9. Storage pools

Issue	Failing Condition	Disposition
12355	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 10. Data evacuation

Issue	Failing Condition	Disposition
13774	It has been observed that after data evacuation completes, the total evacuation bytes in the event console message indicating X out of Y evacuated isn't always byte-accurate.	Look at the destination Slicestor® Devices to see how much data is stored on it.

System Configuration

Table 11. System configuration

Issue	Failing Condition	Disposition
11405	On certain classes of drives (desktop), it has been observed that the drives can transition to a read-only state when quarantined. If this occurs, a subsequent attempt to fail the drive and migrate its slices to adjacent drives are unsuccessful. Under normal circumstances this is not a major concern since the slices from that drive will be rebuilt. However, on some systems that are experiencing higher than usual rates of drive failure, this might cause reliability concerns.	If a drive quarantines immediately after being resumed, call IBM support to verify whether it is safe to try to fail the drive and migrate its slices. IBM support checks the state of the drive and also assess the health of the system to confirm that the potential loss of slices from that drive will not impact data reliability.
13738	Because data is reallocated between Slicestor® Devices during system expansion, it is preferred that the new Slicestor® Devices are physically located in the same sites as the existing sets. If this is not possible, slices that need to be reallocated might need to cross WAN links between sites. This can result in a slower reallocation rate and a longer reallocation phase. Additionally, the higher latency that typically exists when traversing these links can result in a greater request latency for requests that the source store must proxy.	If this situation arises, contact Customer Support to discuss the proposed system expansion. We work with you to ensure that the new set of devices is provisioned in such a way that the inter-site traffic is minimized.
COS-5615	With index delegation enabled, index operations are delegated to slicestores to perform the update. In cases where a store is slow or impaired, index operations delegated to it may take longer to complete, resulting in elevated client latencies.	Improved handling of delegation requests will be addressed in a subsequent release.

Deleting objects

Table 12. 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 13. Manager Web Interface

Issue	Failing Condition	Disposition
10648	On the edit cabinet page, when unassigned nodes exist, it is not possible to move nodes to the bottom of the cabinet because the page does not scroll automatically.	Change the zoom so that all the cabinet slots are visible and then move the node to the desired slot.

Vaults

Table 14. Vaults

Issue	Failing Condition	Disposition
	Nothing to report	

Vault Mirrors

Table 15. Vault mirrors

Issue	Failing Condition	Disposition
10788	If an extreme network bandwidth imbalance exists between two sites in a mirrored vault configuration, and total load on the system exceeds the capacity of the slower site, traffic to both sites might experience a "sawtooth" pattern with alternating periods of high and low throughput. Additionally, pending writes to the slower site prevent writes to the faster site from proceeding. This occurs even if synchronous write is disabled.	During normal operation, disabling synchronous write allows requests to return to a user as soon as the fastest site returns. Reducing average throughput demand over time to be lower than the throughput capacity of the slower site will remove the "sawtooth" I/O pattern and will allow bursts of I/O to occur at the speed of the fastest site.
12854	When performing writes of small objects to a vault mirror, and synchronous writes are disabled, it is possible to queue a large number of operations in the Accesser® Node memory. If this condition persists, it is possible for the Accesser® Node to run out of memory.	To mitigate this issue from occurring, customers should ensure that they are not uploading objects at a rate greater than the slower site can handle.
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.

Vault migration

Table 16. Vault migration

Issue	Failing Condition	Disposition
14403	When a vault is configured with an internal proxy configuration, there is an inconsistency in the way the client-accesser versus accesser Slicestor® Device throughput is represented.	All inbound and outbound traffic is included in the client-accesser graph, but only the traffic to the backing vault is included in the accesser- Slicestor® Device throughput graph
14450	In cases where the target vault of an active vault migration goes below threshold or becomes unavailable, the migration progress bar displayed in the manager might erroneously jump to 100% completed. In this condition, the migration will still be active, and any unmigrated objects will still be migrated.	The migration completion event in the manager will only trigger once the migration has fully completed, irrespective of the status reported in the progress bar. Therefore, the completion of a migration should be judged by the migration completion event in the manager.
14484	When performing a vault migration, it has been observed that it is possible for the migration activity to halt and not make any progress.	There is no workaround or mitigation identified for this issue at this time. If you are performing a vault migration and progress has halted, contact IBM support.

Installation

Table 17. Installation

Issue	Failing Condition	Disposition
9465	When installing ClevOS using a physical or virtual CD drive, the appliance might reboot or hang while booting.	Use a USB storage device to perform the installation.

Native File

Table 18. Native File

Issue	Failing Condition	Disposition
COS-7255	Certain workload conditions that involve reading content through NFI might result in OOM issues on the filer process.	Nfsfiler process is automatically restarted on F5100 devices. No manual intervention required. Lower concurrency workloads recommended to avoid triggering race condition.
COS-4269	Filesystem directories might become unresponsive if a large number of files are stored in a single directory (50,000+) or files are frequently deleted and added.	Contact IBM support to assist in DB compaction / cleanup. Avoid filling directories with large numbers of files or workloads that repeatedly create and delete files within a directory.
COS-5896	File Accesser devices only support hardware Accesser devices. Docker Accesser installations are not supported.	Deploy F5100 devices for use only with physical Accesser devices.
COS-5454	Under heavy load conditions where the File Accesser or Accesser devices become overloaded, the client might receive "Remote I/O Error" messages.	Application level retry of operation required. Reduce concurrency to File Accesser devices or add more Accessers and retry workload.
COS-6872	The File Accesser device REST Endpoint does not support SSL Connections (HTTPS), only HTTP is supported.	Use HTTP for communication with the File Accesser REST endpoint.

Table 18. Native File (continued)

Issue	Failing Condition	Disposition
COS-6851	Using Filesystem or Share names with capital letters might prevent some S3 clients from accessing content properly by using the File Accesser device REST API.	Create Filesystems and Shares by using only lower case letters or avoid use of S3 clients that force lowercase referencing of bucket names.
COS-6805	File Accesser devices that are configured within a single IBM COS installation can support a maximum of 100,000 shares and filesystems across all devices.	Avoid creating more than 1000 Filesystems and 100 Shares per Filesystem.
COS-6895	The File Accesser device REST API endpoint does not support authentication requests that use query parameters (AWS V2 style authentication).	Use credentials that are supplied as headers.
COS-6305	If multiple File Accesser devices in a File Server Pool are down only one of them displays the "Not Actively Participating" message in the Manager UI.	Address the problem causing any File Accesser device to go offline as soon as possible.
COS-7349	Filename character set conversion between Windows-1252 and UTF-8 does not handle extended ASCII characters properly and results in filenames with extended ASCII characters represented by "?" for files that are written by Windows and read by Linux (or vice versa). Standard ASCII conversion works properly. This only applies when the "Character Encoding" of a share is modified from the default utf-8 encoding.	Leave "Character Encoding" Share option set to the default of UTF-8 encoding and limit filenames copied from Windows clients to the lower ASCII set of characters.
COS-7611	Upgrading File Accesser Devices using HA from version 3.8.3 to 3.10.x may experience unexpected behavior during and post upgrade including high CPU usage by HA process.	Contact IBM Support for more information if required. Remove any assigned VIPs from File Accesser devices before performing upgrade. Re-assign after successful completion of upgrade to 3.10.0. Do not run File Server Pools with both 3.8.3.x devices and 3.10.0.x devices with HA configured.
COS-7497	When performing large file writes in excess of 1TB through the NFS gateway appliance, the write operation will fail to complete and return an error.	Avoid writing files in excess of 1TB, and break up large files into multiple smaller files.
COS-2741	Native File Interface and STaaS must not both be enabled on the same system as Native File is incompatible with STaaS vaults.	Contact IBM Support for more information if required.
COS-7898	An abrupt shutdown of a File Accesser device can cause issues with the storage database (cassandra) upon restart.	Contact IBM Customer Support and run "nodetool repair" on the effected device. Use a graceful shutdown of a File Accesser device whenever possible.
COS-10951	Host name change causes Activator script errors in nfsfiler.log.	Do not rename host systems after bringing into service.
COS-10783	Ranged Reads are not permitted through the S3 API on File Accesser devices.	Avoid Ranged Read requests when using the File Accesser S3 API.
COS-10774	On a freshly restarted File Accesser device, simultaneous requests for the same resource through a mount may result in a nfsfiler process hanging.	Preventive Action: Mount a Filesystem share upon reboot before allowing production access. Recovery Action: Restart File Accesser
COS-10195	Extended Characters in filename do not convert properly between windows and linux clients.	Do not set character encoding from default (UTF-8). Transformations may not work properly.

Table 18. Native File (continued)

Issue	Failing Condition	Disposition
COS-7783	In process I/O may fail in the event of any File Accesser device going off line if that File Accesser is receiving a metadata update at the time of the outage.	Resend of failed data write.

Chapter 5. Supported Hardware Platforms

IBM Cloud Object Storage Appliances

Table 19. Minimum Version of ClevOS Compatible with Listed Hardware Platforms

Appliance	Model	Minimum ClevOS
IBM COS Manager™ Appliance	M2100	≤2.7.0
IBM COS Manager™ Appliance	M2105 (3401-M00, 3403-M00)	3.2.2
IBM COS Manager™ Appliance	M3100	2.7.0
IBM COS Manager™ Appliance	M3105 (3401-M01, 3403-M01)	3.7.2
IBM COS Accesser® Device	A2100	≤2.7.0
IBM COS Accesser® Device	A3100	≤2.7.0
IBM COS Accesser® Device	A3105 (3401-A00, 3403-A00)	3.7.2
IBM COS Accesser® Device	A4105 (3401-A01, 3403-A01)	3.7.2
IBM COS Accesser® Device	F5100 (3401-A02, 3403-A02)	3.8.3
IBM COS Slicestor® Device	S1440	≤2.7.0
IBM COS Slicestor® Device	S2104	3.2.1
IBM COS Slicestor® Device	S2212 (3401-S00, 3403-S00)	3.2.1
IBM COS Slicestor® Device	S2440	3.0.1
IBM COS Slicestor® Device	S2448 (3401-S01, 3403-S01)	3.7.2
IBM COS Slicestor® Device	S4100	3.1.0
IBM COS Slicestor®Device	S3448 (3401-S02, 3403-S02)	3.8.3
IBM COS Slicestor®Device	S2584 (3401-S03, 3403-S03)	3.8.1
IBM COS Slicestor®Device	S2212A(3401/3403-S10)	3.10.0

Hewlett Packard

Table 20. Minimum Version of ClevOS Compatible with Hewlett Packard 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	3.6.0
Slicestor® Device	Apollo 4510	3.6.0

Table 20. Minimum Version of ClevOS Compatible with Hewlett Packard Hardware (continued)

Appliance	Model	Minimum ClevOS
Slicestor® Device	Apollo 4530	3.6.0

Seagate

Table 21. 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 22. Minimum Version of ClevOS Compatible with Cisco Hardware

Appliance	Model	Minimum ClevOS
Cisco Slicestor® Device	UCS C3260	3.7.4

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.

IBM®

Printed in USA