IBM Cloud Object Storage System
Version 3.10.1 for August Maintenance Release

Release Notes

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 3.10.1

IPv6 [600]

The IPv6 feature allows users to utilize IPv4 and/or IPv6 within a system for reading and writing objects, administration, and most external services. IPv6 addresses can be entered using the nut shell. If both IP protocols are configured, users can specify a preference for object operations within the system by using the Manager UI.

If a user intends to use IPv6 within a system, all devices must be upgraded to an IPv6-capable version. Additionally, devices brought in to an IPv6-enabled system, which are not capable of IPv6 may only function if Slicestors within that Vault do not have IPv6 addresses configured. These devices should be set to the same version as the system and might not provide all functionality until this requirement is fulfilled. Additionally, Docker appliances, File Accesser® F5100, and geo-dispersed efficiency upgrades do not yet support IPv6 and will still require IPv4 connections if IPv6 is enabled.

Storage Metrics Center [946]

The Storage Metrics Center (SMC) is a highly performing and scalable solution that is developed to generate accounting metrics and offload administrative queries from the system. The SMC is itself a stateless entity but relies on its distributed database for workflow management around fault tolerance, error handling, scaling, and high availability. Each SMC node has a processing component for generating metrics, database component for storing metrics and REST interface for producing reports. The SMC digests access logs from the system and calculates metrics of interest from the embedded information streams: usage, throughout, operations. The metrics are aggregated and stored in time series in the SMC database. Accounting reports are produced by the SMC Pool that uses the REST interface. Customers can overlay their own billing models in order to produce customer facing invoices.

The Manager REST API Developer Guide and Manager Administration Guide were updated, and the SMC REST API was added as a new document for this feature.

Note: Storage Metrics center is an optional part of the system and runs on a cluster of three or more 1U servers that can be ordered from IBM®. The server model number is IBM COS Accesser® T5100. Also, it should be noted that the SMC will only aggregate metrics from a COS system or portion of the COS system that is operating in container mode.

Enable Container Mode [942]

Enable Container Mode was developed primarily to enable operators of IBM Cloud Object Storage systems to support large scale Storage-as-a Service use cases. Container mode supports millions of containers (buckets) and millions of end users but requires that the operator build their own account management system to interface to the COS system. In addition, a number of features that are supported in vault mode are not supported in container mode and there will be some performance impacts to the system to run in container mode. Important to note that when in container mode the manager no longer has the capability to display any information on containers, storage accounts...etc. Customers who are considering enabling container mode should contact IBM support for guidance.

This features adds the ability to convert standard vaults to container vaults and support mixed-mode operation, with some access pools operating in container mode while others operate in vault mode. Earlier, container mode could only be enabled on a system that had no standard vaults. This meant that container mode could only be enabled on a new system. This feature also adds the capability to convert standard vaults to container vaults. Once container mode has been enabled on the system, the operator is able to select 1-N access pools for conversion. The Manager converts all the standard vaults that are

deployed to the access pools into container vaults (if the standard vaults are deployed to multiple access pools, all access pools must be selected for conversion at the same time). As part of the conversion, each standard vault is converted into a container vault. A container is created for the objects that are currently stored in the standard vault, and the vault name is used for the containers name. Therefore, the IO paths remain the same following the migration. During conversion, storage accounts are created for each manager user that has access to converted vaults.

DLM Architecture [150]

State Changes

The disk states have been updated as part of an initiative to consolidate disk state information into one standard. Operators will notice that there is now only one state representing the health of any given drive. The changes are outlined in the table below.

Table 1. State Changes

Old State	New State	Description
GOOD	**ONLINE**	Healthy and usable drive.
PULLED	**OFFLINE**	A drive that is physically pulled, or manually detached by an operator.
QUARANTINED	**DIAGNOSTIC**	A drive that has been programmatically determined not able to be used for IO or manually suspended by an operator.
FAILING	**MIGRATING**	A drive undergoing failure migration to move slices away from an unhealthy drive.
FAILED	**FAILED**	A drive that has attempted to migrate all of its slices and is ready for replacement.
INITIALIZING	**INIT**	A drive that is preparing to be used in the system.
UNUSABLE	**UNUSABLE**	A drive that cannot be used by the system due to perceived hardware issues.
FOREIGN	**FOREIGN**	A drive that belongs to a different appliance and cannot be used.
UNKNOWN		Deprecated.

Reason Code Changes

Suspend reason codes have been consolidated and re-mapped as outlined in the following table:

Note: Starting with release 3.10.1 and continuing into later releases, DLM Numeric Codes are not presented on the Manager UI. Only event descriptions are provided.

The table with Reason Code Changes shows a mapping of pre-3.10.1 DLM numeric codes to 3.10.1 event descriptions that are available on the Manager UI.

Table 2. Reason Code Changes

Old Code	Pre 3.10.1 Event Descriptions	New Event Description	Extended Comments
1, 2	A SMART failure A SMART command failure	drive attribute exceeded threshold	A preconfigured or user defined attribute h exceeded the defined threshold. E.g. if SM attribute #5 reaches 1000 reallocated sector

Table 2. Reason Code Changes (continued)

Old Code	Pre 3.10.1 Event Descriptions	New Event Description	Extended Comments
3	User initiated operation	User initiated action.	Self explanatory.
4	Excessive I/O errors on disk	IO errors exceeded threshold.	IO errors reported by dsnet-core have e the allowed threshold (20% by default).
5	Excessive timeouts on disk	IO timeouts exceeded threshold.	IO timeouts reported by dsnet-core hav exceeded the allowed threshold (20% by
6-9,11-13	An invalid internal structure on the data drive	possible storage metadata issues detected	This consolidates all of the invalid inter structure quarantine codes.
		possible file system issues detected	The drive signature did not match the e structure and the drive could not be ver belong to the dsNet.

Event Console

Many events have been deprecated and are now only triggered on disk state changes.

Virtual Appliances - Disk Operations

Most disk operations for virtual devices should now be supported enabling failure migration and disposal operations.

Changes to Disk Operations

- Disposing Drives

  Disposing data drives no longer removes them from the storage list. Instead they are simply marked as failed and automatically removed once they are no longer in the system.
- Blacklisting and Whitelisting

  The blacklist and whitelist operations are no longer available. Operators should follow the state diagram to permanently remove drives from the system.
- Reset Retries

  Only resetting drives attempt to retry their operations should they not complete successfully. This means that should an appliance be powered off during drive initialization, the initialization will continue on the next boot as long as the maximum number of retries is not exceeded. All other operations are not retried.

General Appliance Changes

CLI Utilities

- storagectl and nut-storage
  - Both storagectl and nut storage have equal functionality now. They can be used interchangeably although nut storage is more restrictive.
  - Both storagectl and nut storage are now blocking until the requested operation is completed.
- storagectl interface
  - The storagectl command line interface has been expanded to include additional commands:
    - info:
      - Prints drive statistics (SMART attributes, sysfs counters, etc) as last polled by the service.
    - log:
      - Prints the history of drive statistics over time.
    - list:

    - Print the current drive states as depicted by the system.
- history:
    - Prints the history of drive states over time.

Settled Read Only File Writes [1034]

This new feature is designed to provide a write-once read many times like functionality (Settled Read Only File Writes) for the Network Attached Storage (NAS) interface to the IBM Cloud Object Storage System™.

The Settled Read Only File Writes option allows the system administrators to define a settle duration, and an override POSIX-based user ID or group ID. Once the settle time expires on a file system (file or directory's last modified time), the file/directory becomes Read Only for all users except the override user or a user part of the override group.

This option can only be enabled at file-system creation time, and the settled time cannot be modified once it is set. The override user ID or group IDs can be changed at any point in a filesystem's lifetime.

Since the settle time is based on a file's last modified time, any changes to a file after the file is settled resets the file's settled status. The file has regular POSIX permissions until the settle time expires on the file.

Creating a new file in a directory results in an update to the parent directory's modified time. Any new file that is created in a settled directory resets the settled status on the directory. Regular POSIX permissions apply until the settle time expires.

Share Level Authorized IP Addresses [1099]

This new feature allows system administrators to restrict access to shares for Network Attached Storage (NAS) interface to the IBM Cloud Object. Access can be restricted based on fully qualified domain names, IPv4 Addresses, and wildcards ('*', '?', Character classes []) in domain names.

The restriction can be any of the following:

- An IPv4 address.
- A Classless Inter-Domain Routing string. For example: 192.168.56.101/24.
- An IPV4 address with a subnet address. For example: 192.168.56.101/255.255.255.0.
- A Fully Qualified Domain Name. For example: usil.ibm.com
- A string equal to "*". A '*' implies any system can access the share.
- A "*" in domain names. A name with a "*" is matched as a Posix basic regular expression. For example, *.ibm.com. Note: *.ibm.com matches usil.ibm.com or host1.usil.com.
- One or more '?' in the domain name. A domain name with one or more '?' is also matched as a Posix basic regular expression. For example, usil?.ibm.com matches 'usi.ibm.com' or 'usil.ibm.com'
- Character Classes ([..]) in domain names. For example, [abc]-host.ibm.com. This will match a-host.ibm.com or b-host.ibm.com or c-host.ibm.com
- '*' or '?' or '[]' can be present in the same domain name multiple times.

Note: If an entry in the list contains '?', '*', or '[', then the Manager application will only make sure that the allowed system string is a valid regular expression. This is so that entries such as *.ibm.com, host?.ibm.com, and host[1-9].ibm.com can be entered.

The access restriction can be applied either at create share time, or an existing share can be modified to add/remove restrictions.

Chinese Character Support [1040]

This feature allows Unicode characters in user supplied values throughout the system. The Object Storage interface already supports Unicode in object names and object content. Previously, the management UI and API supported Unicode, but only the subset of characters that lie within the BMP (Basic Multi-Lingual Plane). This feature expands the supported characters to cover the full Unicode character set. Full Unicode input is allowed on all "free form" text fields, such as:

- Account name
- Site name, abbreviation, description, company, and address
- Storage pool name and description
- Access pool name and description
- Device alias and description
- Cabinet name and description
- Vault description
- Mirror description
- Vault template name and description
- Mirror template name and description
- Tag name and description
- System name
- Login banner
- Account passwords

Note: The exceptions are vault and mirror names, which are syntactically restricted to ASCII-only for interoperability purposes. In addition, the device administrative shell (Nut interface) has been modified to support Unicode values when configuring the device's city, state, organization, or organizational unit.

Support OCSP [1042]

The manager has long supported CRLs for users (such as administrators) who are logging in to the manager UI or REST API using client PKI certificates. As of 3.10.1, the manager now also supports revocation checking by way of OCSP.

New Time Zone Support in the Manager UI

Additional time zones have been added to the manager that are accessible through the Preferences Configuration section under the Administration tab.

Remove 'Server' field in un authorized HTTP responses [833]

There are many ways a client application could query the server version anonymous, even if it was not properly authorized. Instead of preventing the Server header from being returned only with 401 Unauthorized or 403 Forbidden responses, IBM COS instead allows whether the Server header is included at all to be configurable.

System External APIs:

Support the following options for HTTP Server header:

- Show software and version
- Show software only
- Omit header

This will impact these external APIs in the following way:

Accesser (and Slicestor® when embedded Accesser enabled):

80/443/8080/8443 - HTTP for S3/SOH/Swift API

NAS Filer

80/443 - HTTP for S3 (always omits Server header, not configurable)

Manager

80/443 - HTTP for UI and Manager REST API

Chapter 2. Interface Modifications

API updates for the 3.10.1 release have been referenced in the following documentation:

- Manager REST API
  - Updated Account Management Chapter
  - Added NEW SMC Pool Management Chapter
  - Updated Remove Storage Pool Sets Chapter
  - Updated Mirror Management Chapter
  - Updated Vault Management Chapter
  - Updated Site Management Chapter
  - Updated Organization Management Chapter
  - Updated Reports Chapter
  - Updated Administration Chapter
- OpenStack Object Storage API Developer Guide
  - Updated Differences between OpenStack and IBM Cloud Object Storage System™ OSOS APIs
  - Updated Using the OSOS API
  - Updated Supported subset - Common request headers
  - Updated Storage container services
  - Updated Create or update object
- Device API Developer Guide
  - Updated Statistic Response Parameters
- SMC REST API Development Guide - Added new document containing a section on "Introduction to SMC REST API Details."
  - Command Summary
  - Time Range for a Servlet Query
  - Calculation of Usage Metrics
  - Calculation of byte_hours
  - Calculation of average_byte_used
  - Calculation of operation metrics
  - List metrics

API Changes 3.10.1

Improved the acceserRequest JSON format for Accesser requests' statistics from the device API:

- Changed from format "accesserRequest":{"put”:{“200”: 2486}}
- Changed to format "accesserRequest":{"PUT":{"OBJECT":{"200":"2486"}}}

Handling of the synchronous option is updated in the following APIs:

- CreateMirror API
- EditMirror API
- CreateMirrorTemplate API
- EditMirrorTemplate API

If a customer uses the above APIs, and are passing multipart as an asynchronous option, then they need to update the APIs. As such, this option can no longer be set or configured from Manager API. NOTE: The ClevOS software always creates a mirror with "multipart set to synchronous mode," and "is non configurable."

The containerMode key was removed from the View System and View System Configuration API responses. Any scripts leveraging this key will have to be updated. In the View System and View System Configuration API responses, the vaultPurpose associated with each vault within the system can have one of the following values: standard, management, service, container. The presence of a vaultPurpose set to service or container indicates the system can support container vaults.

Chapter 3. Resolved Issues

Resolved issues in 3.10.1 August Maintenance Release

Table 3. Resolved issues

Issue	Description
COS-23585	Get Bucket versions latency is increased due to secondary look-up not being run in parallel. The fix is to implement the parallel execution of a secondary look-up.
COS-23159	A caching issue that resulted in re-using the same stale revision on retries resulted in a occasional failure to add parts to a multipart transaction. The stale entry is now being evicted from cache before starting the retry.
COS-23832	Selective debug logging is now enabled by default at a selection rate of 0.001 (i.e. 1/1000 requests). Starting with release 3.10.1, a change was made to enable selective debug logging at a low rate by default. Selective logging rate takes precedence over selection frequency unless it is explicitly overridden to a selection rate of 0.0. If a specific selection frequency is desired, the selection rate must be explicitly set to 0.0 to override the default selection rate in order for the selected frequency to take effect. This will be addressed in a future release

Resolved issues in 3.10.1 July Maintenance Release

Table 4. Resolved issues

Issue	Description
COS-23490	Accesser® Devices sometimes crashed due to a runtime exception caused by blocking calls in certain thread pools. This was fixed by moving the blocking call to a thread pool that allows blocking so that the runtime exception does not happen.
COS-23515	S3 listing latency drastically increased in ClevOS 3.10.1.
COS-23078	Index split operations are performed in the background to maintain proper balance in the index structure. Update operations as part of a split could be improperly sequenced, such that failures might leave the index in an internally inconsistent state. This might result in 500 errors for insert, removal, or listing operations for this portion of the index. The fix ensures proper sequencing of the internal updates such that failed updates will always result in a consistent internal structure.
COS-22330	Incorrectly formatted disks causing upgrade failures.
COS-22823	Disks using legacy partition tables with non-4K aligned boundary are quarantined.
COS-22391	Concurrent SES LED operations cause SES device resets.

Resolved issues in 3.10.1

Table 5. Resolved issues

Issue	Description
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.

Table 5. Resolved issues (continued)

Issue	Description
COS-7370	On occasion, a management vault GET failure event appears in the Event Console on the Manager UI after vault creation.
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::[]
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.
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.
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.
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.
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.
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.
COS-6872	The File Accesser device REST Endpoint does not support SSL Connections (HTTPS), only HTTP is supported.
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.
COS-1259	When performing a listing of a storage account by using the service API, the storage account enable/disable state is not displayed.
COS-2741	Native File Interface and STaaS must not both be enabled on the same system as Native File is incompatible with STaaS vaults.
COS-1478	It has been observed that read operations for large objects (20MB or greater) are degraded (15-20%) relative to prior releases.
COS-6282	In the Accesser request graph on the monitor device page, request data now include the resource for the request.
COS-2263	Removed Read Threshold from all the reports
COS-6798	In the Accesser request graph on the monitor device page, the legend labels now include the total number of requests for the time range shown.
COS-7971	Attempts to change device passwords for docker containers from the maintenance page produces an error message. This operation is not allowed for this type of device, and the fix in addition to providing an error message is to disable the check boxes for the docker devices.
COS-7796	For Accesser devices that have a one drive configuration, the Manager UI erroneously reports the device health as yellow when it should report green.

Table 5. Resolved issues (continued)

Issue	Description
COS-9083	S3 Remote Proxy fails to list delete markers and their retained versions. Fixed this defect as part of listing refactoring in core. In case of a versioned listing a versioned listing request is issued to S3 to respond to the end user. This API call gives us results including the delete markers and their retained versions.
COS-9418	Gateway device upgrade is possible if device is not in a file server pool.
COS-9608	This release introduces changes to the format of the JSON response for Accesser requests' statistics from the device API.
COS-5794	Write Performance Degradation at Low Container Fill
COS-5615	Elevated client latencies in the presence of impaired store due to not reassigning delegation from a slow store.
COS-7255	Out of memory on File Accesser.
COS-5454	Under heavy load conditions where the File Accesser or Accesser devices become overloaded, the client might receive "Remote IO Error" messages.
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.
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.
COS-7256	One Accesser went "unresponsive" after upgrade to 3.8.2.
COS-9534	S3 container creation fails if storage account is updated during creation request.
COS-9922	Nut enclosure bay list - Platform not supported.
COS-9089	dsnet-core flapping on slicestor98.
COS-8630	SNMP service does not start.
COS-8007	Accesser not servicing requests.
COS-11319	In an earlier release, a feature was introduced to provide advanced configuration support to customize Apache web server cipher encryption on the manager. Some weak ciphers were removed in this release. As a result, IE7 and IE8 running on Windows XP may not work properly.
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.
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.
COS-6895	The File Accesser device REST API endpoint does not support authentication requests that use query parameters (AWS V2 style authentication).
COS-12693	In earlier releases for some scenarios, manager backups failed to complete due to backup files consuming too much memory.
COS-11951	In prior releases, a manually failed drive is reported as "Failed" in the Statistic Device API, but the State Device API reported the drive as "Failing." Also, upon device reboot, this entry is missing from the State Device API.
COS-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.
COS-7412	When performing an S3 POST request with and invalid signature and an empty policy, an exception is encountered when parsing the request, resulting in a runtime exception being thrown and the core process on the Accesser to restart. This release introduces a fix for this issue, causing this condition to result in an appropriate HTTP error messages being sent to the client.

Table 5. Resolved issues (continued)

Issue	Description
COS-10783	nfsfiler exception when S3 range read is performed on directory.
COS-8204	High load seen on Slicestors during prefix listing when value of max-keys=0 is set.
COS-9270	Rebuilding not making progress - run out of resource_manager.on_heapresources (which causes rebuilder to stop).
COS-11993	Internal server error due to Storage IO error (Couldn't convert string to bytes).
COS-13155	Rebuilder taking very long after disk erasure tests.
COS-8022	Prevent container names from being reused for a short period after they are deleted.
COS-15443	Object Tagging requests not recognized
COS-15891	Accesser device Level API only reports one response code per REST Operation
COS-14806	On-Heap Saturation on Accessers after upgrading to CLevOS 3.10.0.134
COS-15847	Message ack time charts not displayed
COS-11757	Diagnostic incident remains active after reboot
COS-10770	Remove Char Encoding translation support from UI
COS-7589	When a node is removed from file server pool, Corosync resource agent is not deleted
COS-16916	The following changes were introduced in the Manager User Interface (UI) and REST API for storage pools. • Manager UI, the storage engine options for packed storage and file storage appear on the Create Storage Pool page as they did in previous releases. These options are now also available on the Expand Storage Pool page, but with the packed storage option being pre-selected on both pages. A user can change the selection if desired. • REST API, when a storage engine is specified it is used. However, packed storage is used if no storage engine is specified; in this situation no validation error occurs. **Note:** The Expand Storage Pool API changed in 3.10.1. A storage engine property was added that allows a storage engine to be assigned on a per set basis, which is a new capability. In prior releases, storage engine was assigned on a per-storage pool basis.
COS-15179	S3 listing are performed utilizing listing of the index. The accesser caches index nodes to support faster listing operations.
COS-16023	Following the completion of set removal, if the removed devices remain associated with the same dsNet, but are re-used in a different storage pool than the one they were originally removed from, the manager may erroneously display errors on the device communication page indicating that existing devices in the pool are no longer able to communicate with the devices that have been removed.
COS-16915	A set from a storage pool cannot be removed when Slicestor devices are missing (holes exist).
COS-16709	Slicestor appliances that utilize an LSI 9300-8i disk controller for data drives will not quarantine SATA drives in the event of a SMART health failure.
2753	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.
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.
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.

Table 5. Resolved issues (continued)

Issue	Description
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.
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.
COS-9268	When upgrading from a 3.8.x, 3.9.x, or 3.10.0 release, there is a chance that the hard drive temperature graph for all physical devices visible on the Manager interface will show an incorrect reading for up to one hour.

Chapter 4. Known issues

Table 6. Known issues

Issue	Failing Condition	Disposition
COS-6803	For Slicestor devices with multiple OS drives, degradation of OS drives does not affect the device's health on the Monitor device page.	Repair the OS drive or contact IBM Customer Support for more information.
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.
COS-11201	In the Event Console of the Manager User Interface, the event details section for failing disk migration events contains a parameter 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.
COS-15179	S3 listing are performed utilizing listing of the index. The accesser caches index nodes to support faster listing operations.	The accesser index cache is refreshed periodically. After 60 seconds, the node in question will be updated/evicted from the cache and listing operations will succeed. This will be fixed in a future release.
COS-14232	When performing large volumes of delete requests or deletes of large objects, either through individual delete requests or multi-delete operations, it is possible to consume large amounts of memory resources performing the background deletion operations, leading to resource exhaustion and 503 errors.	Once the content deletion completes, the resources will be returned and the 503 errors will cease. If you encounter 503 errors during periods of heavy delete workloads, please contact IBM Customer Support. This will be fixed in a future release.
COS-15399	Following an Accesser OS drive replacement, a new device certificate must be generated for this device, and a whitelist containing this certificate information must be distributed to the other devices in the system which this device will attempt to communicate with.	A core process restart of the Slicestore reporting the authorization error. This will be addressed in a future release.
COS-16023	Following the completion of set removal, if the removed devices remain associated with the same dsNet, but are re-used in a different storage pool than the one they were originally removed from, the manager may erroneously display errors on the device communication page indicating that existing devices in the pool are no longer able to communicate with the devices that have been removed.	This condition will clear once the remaining devices in the set, and any Accesser devices deploying vaults stored on that set, are re-imaged. This will be addressed in a future release.
COS-13575	The "stop migration" operation for failing disk migration on the Manager User Interface (UI) may 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 provided in the Manager Administration Guide under disk lifecycle management.

Table 6. Known issues (continued)

Issue	Failing Condition	Disposition
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 provided in the Manager Administration Guide under disk lifecycle management.
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).	Should this occur when migrating a VMware virtual device using vMotion, a cold migration should be used instead such that the virtual machine is offline during the migration.
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. In some cases however, this process may 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-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.	Contact IBM customer support to help correct the situation.
COS-7488	When performing a storage pool set removal, it is possible that once the reallocation has finished for an source Slicestor device, it may show some small amount of data still present.	No action is required. Once the set removal has completed, all slices will 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 may 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-16915	A set from a storage pool cannot be removed when Slicestor devices are missing (holes exist).	Slicestor devices must be added to the storage pool set so that no devices are missing (no holes exist) prior to performing set removal.

Container

Table 7. 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-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.

Table 7. Container (continued)

Issue	Failing Condition	Disposition
COS-15401	If a user attempts to create a management vault using "manual configuration" (accessed through the Configure Management Vault page) based on an existing vault template, management vault creation will fail 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.

Upgrading and Installation

Table 8. Upgrading and Installation

Issue	Failing Condition	Disposition
COS-7126	When unzipping of upgrade file fails when a device is upgrading the failure message "The Selected File can not be unzipped 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 will appear until the page is reloaded.
627	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.
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 (e.g. OS drives) will be reported as newly discovered in the Manager event console.	No action is required.
COS-15642	When upgrading devices that contain logical RAID drives, the Manager event console will show a drive offline event immediately followed by a drive online event for each physical drive that is part of a logical RAID drive.	No action is necessary. These events are simply representative of a transition phase of the RAID drives during the startup sequence and will be removed in a future release.
COS-22924	When the Manager is upgraded to ClevOS 3.10.1 or newer for the first time, immediate log in might not be possible. The Manager application may 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.	Contact IBM 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.
COS-15370	A number of changes were made to the metadata storage schema for 3.10.1 to improve performance and overall scalability.	Upgrading a File Accesser to 3.10.1 running versions of ClevOS 3.10.0.x or earlier with or without existing Filesystems/data to be preserved requires contacting customer care to assist in properly migrating metadata (Filesystems, Shares etc.) to the new format. **Note:** The migration should be done prior to the upgrade. This only applies to cases where no existing data is to be preserved.

Table 8. Upgrading and Installation (continued)

Issue	Failing Condition	Disposition
COS-22994	In a system with a Manager device on release 3.10.1 or greater, and containing SMC devices, any Slicestor devices or Accesser devices on a release lower than 3.10.1 will not be able to communicate with the Manager.	Upgrade any Slicestor devices or Accesser devices on a release lower than 3.10.1 to the same release as the Manager.
COS-23615	When a set replacement or removal is performed on a storage pool, the set numbering on the upgrade page will differ from the correct set numbering on all other pages.	The ordering of the devices and sets is still correct on the upgrade page, instead the issue is an inconsistency in the set label. To see the correct set numbering, view the monitor storage pool page.

Alerting and Reporting

Table 9. Alerting and reporting

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 Customer Support to confirm and correct the false incident.
COS-6490	If a manager appliance 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.
COS-16709	Slicestor appliances that utilize an LSI 9300-8i disk controller for data drives will not quarantine SATA drives in the event of a SMART health failure.	Prior to upgrading to the 3.10.1 release clients utilizing appliances with an LSI9300-8i disk controller (e.g. Slicestor 2448, Slicestor 2212A, etc.) should contact IBM Customer Support to assess impact.

System Behavior

Table 10. System behavior

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.
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
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.

Table 10. System behavior (continued)

Issue	Failing Condition	Disposition
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-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 11. 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.
COS-16664	For the "Storage Pool Capacity and Disk Report" accessed through the Maintenance tab of the Manager User Interface (UI), sorting for drive category columns do not work with the Safari browser.	Use an alternative browser, such as Chrome or Firefox.

Data Evacuation

Table 12. 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 13. 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.

Table 13. System configuration (continued)

Issue	Failing Condition	Disposition
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.

Deleting objects

Table 14. 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 15. 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.
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 provided in the Manager Administration Guide under disk lifecycle management.
COS-9268	When upgrading from a 3.8.x, 3.9.x, or 3.10.0 release, there is a chance that the hard drive temperature graph for all physical devices visible on the Manager interface will show an incorrect reading for up to one hour.	Wait an hour after the device upgrade has completed and confirm hard drive temperatures are being reported.

Vaults

Table 16. Vaults

Issue	Failing Condition	Disposition
	Nothing to report	

Vault Mirrors

Table 17. 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" IO pattern and will allow bursts of IO to occur at the speed of the fastest site.
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

Table 18. Vault migration

Issue	Failing Condition	Disposition
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.
COS-12442	When a vault migration finishes the work contained in its TODO queue, it kicks off a process to calculate the exact count of the number of objects 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.	

Native File

Table 19. Native File

Issue	Failing Condition	Disposition
COS-15370	Upgrading a File Accesser to 3.10.1 from ClevOS 3.10.0.x or earlier releases will fail.	Contact IBM Customer Support for providing a work around.
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-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-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-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-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.
COS-16461	The System Advanced Configuration page in the Manager User Interface (UI) includes an "Existing Detailed Configuration Rules" section which lists existing advanced configuration settings on individual devices, storage pools, or access pools. However, File Server Pools and SMC Pool configuration settings do not appear in this section.	The configuration settings for File Server Pools and SMC Pools work fine. The parameters are visible on the individual pool pages.
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 20. 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
Accesser Device	A2100	≤2.7.0
Accesser Device	A3100	≤2.7.0
Accesser Device	S1440	≤2.7.0
Accesser Device	S2104	3.2.1
Accesser Device	S2212	3.2.1
Accesser Device	S2440	3.0.1
Accesser Device	S4100	3.1.0

Table 21. 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® F5100	3401/3403	A02	3.8.3
IBM COS Accesser® T5100	3401/3403	A02	3.10.1⌂
IBM COS Manager™ 2105	3401/3403	M00	3.8.1
IBM COS Manager™ 3105	3401/3403	M01	3.8.1
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	3401/3403	S03	3.8.1
IBM COS Slicestor®2212A	3401/3403	S10	3.10.0

Note: ⌂ Requires RPQ

Hewlett Packard

Table 22. 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

Table 22. Minimum Version of ClevOS Compatible with Hewlett Packard Hardware (continued)

Appliance	Model	Minimum ClevOS
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
Slicestor® Device	Apollo 4530	3.6.0

Seagate

Table 23. 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 24. Minimum Version of ClevOS Compatible with Cisco Hardware

Appliance	Model	Minimum ClevOS
Cisco Slicestor® Device	UCS C3260	3.7.4

Dell

Table 25. Minimum Version of ClevOS Compatible with Dell Hardware

Appliance	Model	Minimum ClevOS
Dell Slicestor® Device	DSS 7000	3.10.1

Lenovo

Table 26. Minimum Version of ClevOS Compatible with Lenovo Hardware

Appliance	Model	Minimum ClevOS
Lenovo Manager Appliance	X3550 M5	3.10.1
Lenovo Manager Appliance	X3650 M5	3.10.1

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