Quantum

StorNext 7.2.5 Release Notes

Product/Software Release Date	January, 2026

Contents

What's New in StorNext 7.2.5

Purpose of this Release

StorNext 7.2.5 adds support for SMB Multichannel, LTO-10, and new client operating systems. This release also includes critical software fixes listed in the section Fixed Issues and Enhancements Addressed in StorNext 7.2.5 on page 5.

Note: Once you upgrade StorNext software, downgrading or rolling back to a previous version is not supported.

New Features and Enhancements

- NAS Cluster Multichannel Support
  - Added support for **SMB multichannel** in NAS cluster mode, enabling enhanced throughput and redundancy (for more information, see Multichannel Support in the Appliance Controller Documentation Center).
- Support for higher density tape media
  - Support for LTO-10 media technology, enables StorNext HSM to manage 66% more data per media, allowing faster restores
- Improved auditing
  - Continuous log output for monitoring all file system activities
  - Enhanced the uptime for the audit service
- Client OS support
  - Red Hat / Rocky / Alma / Oracle 10.0 (Kernel 6.12.x)
  - Red Hat / Rocky / Alma / Oracle 9.6 (Kernel 5.14.0-600)
  - SUSE 15.7 (Kernel 15.14.x)
  - Debian 13 (Kernel 6.8, 12.11 – Kernel 6.1)
  - Ubuntu 25.04 (Kernel 6.14, 24.04.3 – Kernel 6.8)
  - macOS 26
- Add option to disable the **Recycle Bin** (for Windows) and the **Trash** (for macOS)
- Enable users to recover files from the **Recycle Bin** (for Windows) and **Trash** (for macOS) within a StorNext File System

New Features and Enhancements to the Quantum Unified User Interface (UUI)

- If the primary UUI node is unresponsive for more than five (5) minutes, the standby node automatically becomes the primary. When the original node returns, it is reassigned as the standby node.
- Improved log in experience with a simplified flow and consistent theme support.
- Added a new public API with full documentation (see **API Service** under the **About** menu).
  - StorNext File System Pooling APIs (manage file tiers, policies, pools, and configuration)
  - FlexSync APIs (manage host groups, hosts, and file tasks)
  - UUI Node and Client information APIs
- Introduced a path selection tool to make setting source and destination paths easier when creating FlexSync tasks.
- Added the ability to enable or disable SMB Multichannel.
- The StorNext GUI now features an updated design for a modern appearance.

Information about Pre-Upgrade Checks

Beginning with StorNext version 7.2.0, the pre-upgrade check validation process also checks for RPM files that cannot be upgraded. Upon validation failure, an **Admin Alert** containing a list of RPMs found that are not expected on the node(s), is generated. You must remove these RPMs from the node(s) before you activate the upgrade. If the system you are upgrading includes the **lin_tape** RPMs, you must remove the RPMs before you start the upgrade process, and then install the RedHat 8 version of the RPMs after the upgrade process is complete.

Information About StorNext Unified Connector Upgrades

If your system is running StorNext version 7.1 (or later), you **must** upgrade the StorNext Unified Connector on non-MDC based clients.

Note: The StorNext Unified Connector is automatically upgraded on StorNext appliance based clients as a part of the StorNext software upgrade process. This step does not occur for StorNext software (which includes customer-supplied hardware installations).

Caution: Only the latest version of the StorNext Unified Connector properly displays performance metrics and statistics in the UUI.

To upgrade your StorNext Unified Connector, see Unified Connector Tasks.

Information about FlexSync Support

StorNext 7.2.5 only supports FlexSync 3.1.0 (or later). If you want to install StorNext 7.2.5 on your system, or upgrade your system to StorNext 7.2.5 and want to use FlexSync, then you must install FlexSync 3.1.0 (or later), or upgrade to FlexSync 3.1.0 (or later) on each system using the feature, including the source node, the target node, and the managing appliance.

Caution: Mixed versions of FlexSync daemons are not supported and results in a communication error. You must install the same version of FlexSync, or upgrade to the same version of FlexSync on all the hosts or systems using FlexSync. A newer version of the FlexSync daemon cannot communicate with an older version within a configuration.

Compatibility and Support

The StorNext 7.2.5 Compatibility Guide provides the basic compatibility for StorNext 7.2.5, including the StorNext components supported, operating systems and service packs, libraries and drives, browsers, virtual machines, and appliance support. Listed below are just a few of the types of information available to you in the compatibility guide.

- **Upgrade Paths**: Provides information on what upgrades to this release are supported.
- **Appliance Support**: Details the StorNext appliances that are officially supported or verified as compatible with this release.
- **Operating Systems and Platforms**: Provides information on what StorNext components run on various operating systems and service packs. Also includes which operating systems have been newly added or removed.
- **Client Interoperability**: Provides information on what StorNext clients running other versions of StorNext are compatible with metadata-controllers (MDCs) running this release.
- **Virtual Machine Support**: Provides information on what StorNext components running on selected operating systems and service packs are supported in virtual machines.
- **Compatibility with Other Products**: Provides information on references to additional StorNext sold-separately products that are supported with this release.
- **Browser Support**: Provides information on what versions of browsers are supported with the GUI in this release.
- **Drives and Libraries**: Provides information on what Quantum and 3rd party drives and libraries are supported with this release.

Fixed Issues and Enhancements Addressed in StorNext 7.2.5

Ticket Number	Description
SNXT-200	Fixes a problem where CVFS clients could see confusing I/O error messages after data offload, so file access remains reliable and understandable for admins and users.
SNXT-383	Corrects a reporting issue where files deleted during a store operation still appeared active, so file status and lifecycle information are accurate.
SNXT-465	Ensures vault operations verify that required libraries exist and are online, reducing failed jobs and making archive operations more reliable.
SNXT-812	Prevents the **cvcp copy** command from crashing when given invalid thread settings, so scripted and bulk copy jobs fail safely with clear behavior instead of core dumps.
SNXT-1326	Fixed a misleading error message in **cvadmin** related to the shared file system and **mdarchives** directory to be more specific, distinguishing between the shared file system not being mounted and the **mdarchives** directory missing.
SNXT-1478	Fixed an issue where unprivileged users are not allowed to log in during the upgrade process until the platform upgrade to StorNext 7.2.0 completes.
SNXT-1491	Fixed an issue where environments that include an Xcellis Workflow Extender that originally had StorNext version 6.1.1 or earlier installed, the upgrade will fail.
SNXT-1495	Fixed an issue where environments with a large number of LUNs visible on the Xcellis node, the node might run out of memory after the upgrade to StorNext 7.2.0.
SNXT-1530	Improves diagnostics when the REST configuration file is misconfigured, so **fsmpm** startup failures clearly point to the root cause and are quicker to fix.
SNXT-1536	Fixed an issue where the upgrade to StorNext 7.2.0 might fail in some cases due to a lack of free space available in the /var file system to complete the operating system conversion portion of the upgrade.
SNXT-1738	Fixed an issue where upgrades on systems originating from StorNext version 6.0.6 or earlier might fail in the pre-upgrade check when upgrading to StorNext version 7.2.0 due to the presence the obsolete rpms openssl098e and tg3 being installed.
SNXT-1773	Fixes an upgrade scenario where a service being enabled too early could break upgrades to StorNext version 7.2.x, making upgrades more predictable and resilient.
SNXT-1814	Fixed an issue by improving how the file system handles and cleans up large sets of extended attributes during file deletion, preventing an OpHang timeout and associated panic.

Ticket Number	Description
SNXT-1848	Add option to disable the **Recycle Bin** (for Windows) and the **Trash** (for macOS).
SNXT-1880	Updated the byte-range lock and reconnect logic so that blocked lock requests are managed safely across client reconnects, preventing lock structure corruption and eliminating these filesystem and gateway panic conditions.
SNXT-1957	Resolves a crash when using the web console on RHEL 8 with StorNext version 7.2.0, so customers can run xdi operations through the UI without unexpected failures.
SNXT-1977	Fixes a crash in the Java-based management path, improving stability for management connections and reducing unexpected service interruptions.
SNXT-2013	Recursive or batch **snretrieve**/**fsretrieve** operations now complete successfully by default when all requested files are already on disk, and an optional system parameter is available for administrators who prefer to keep the previous **no work done = error** behavior.
SNXT-2019	Improved how long-running defragmentation or file system pooling operations are monitored so they no longer get incorrectly treated as a hang, preventing these unnecessary panics and improving overall system stability during heavy defragmentation or file system pooling workloads.
SNXT-2046	Adds an option to accept the StorNext client EULA silently, enabling fully automated and scripted client deployments without manual prompts.
SNXT-2066	Improves how StorNext handles archived data so retrieves for expired Glacier objects fail more predictably and are easier for admins to understand and manage.
SNXT-2068	Enhances **fsqueue -f** output so it includes file names and paths for truncated files, giving better visibility for troubleshooting and audits.
SNXT-2078	This release updates the interaction between **snaccess** and the purge logic so that restricted directories no longer trigger these excessive kernel spins, preventing the soft lockups and improving overall stability for Linux clients using snaccess access controls.
SNXT-2080	Updates lifecycle and end-of-life decisions to rely on current, supported metrics instead of deprecated LP3 counts, so retention and migration behavior better matches expectations.
SNXT-2099	Fixes allocation issues so large files can still be placed efficiently even when free space is highly fragmented, improving capacity usage and reducing allocation failures.
SNXT-2102	Improve the upgrade reliability for systems that were originally installed on older releases and still have legacy or “dangling” Docker images; the upgrade process has been updated so these untagged images are handled correctly, allowing the overlay-to-overlay2 conversion to complete and preventing upgrade failures on systems with old or dangling Docker images.
SNXT-2123	Prevents client crashes when the same metadata controller is configured with multiple IP addresses, improving stability in multi-path and multi-network setups.

Ticket Number	Description
SNXT-2157	Resolves failures in **samfsimport** on Rocky Linux 7.2.x, allowing customers to import data successfully on that platform.
SNXT-2169	Fixes an issue so that StorNext now correctly recognizes this combined power-supply status and raises the appropriate FRU/RAS events and UI alerts whenever a power supply loses AC power or is removed, ensuring that hardware issues are clearly reported to administrators.
SNXT-2211	Fixes a low-level memory error that could cause crashes when opening files, improving overall system stability under load.
SNXT-2240	Corrects a condition where file system pooling could loop on certain zero-length files and consume excessive CPU, improving performance and reducing unnecessary resource usage.
SNXT-2271	Fixed an issue where Xcellis systems running StorNext version 7.2.x systems on Rocky Linux (OS8) did not include the required **postfix** package that provides **/usr/sbin/sendmail**. This prevented **mailx** from sending email notifications (including **sntierd** email delivery). The **postfix** RPM is now installed on both fresh and upgraded Xcellis systems running StorNext version 7.2.5 or later, restoring expected email functionality.

Supported StorNext Upgrade Paths and Upgrade Considerations

StorNext Software Upgrade Matrix

For information on which StorNext versions allow you to upgrade directly to this release, refer to the **StorNext Software Upgrade Matrix** section in the StorNext 7.2.5 Compatibility Guide.

Journal Size Guidelines

As with prior upgrades, if your file system has a journal size less than 4 MiB, you should resize your journal before upgrading to StorNext 7.2.5.

Use the **cvupdatefs** utility (see the StorNext 7 Man Pages Reference Guide) or the GUI (see Edit a File System) to update your journal size. When growing the journal, the new size must be 16 MiB or greater. However, Quantum recommends using a 256 MiB journal as this has been the default size for new file systems since StorNext 7.0.2. In addition, certain operations such as metadata archive restore may encounter issues when using a journal smaller than 256 MiB.

Distributed Data Mover (DDM) Guidelines

Distributed Data Movers (DDMs) must be upgraded to the same version of StorNext that the Metadata Controller (MDC) is running.

WARNING: Upgrades (such as platform, service pack, and so on) are intended to be done to all systems present in a given deployment. For example, if Xcellis, and Artico are present, they all must be upgraded. No appliance can be left behind.

Considerations When Upgrading NFS Server Nodes to StorNext 7.2.5

To prevent issues with mounted NFS clients, NFS clients must be unmounted prior to upgrading StorNext on the NFS server. If unmounting all NFS clients is not an option during the upgrade, Quantum suggests using the "compat32" mount option on NFS servers.

Note: Stop the NFS server before unmounting the **cvfs** mount. After unmounting, update the mount options and then remount it.

Compatibility Between StorNext and Other Products

The following table provides information regarding compatibility between this release and StorNext components and features.

Product	Compatibility
Appliance Controller	To view supported Appliance Controller software configurations, see the StorNext 7.2.5 Compatibility Guide.
Infiniband	Infiniband installations require assistance from the Quantum Professional Services team, a Service Partner, or a Quantum Service Provider. For additional information, contact Quantum Technical Support.
ActiveScale (AXR, S3) or P100/X100	See the StorNext 7.2.5 Compatibility Guide in the StorNext Documentation Center for information about compatibility between ActiveScale (AXR, S3) or P100/X100, and StorNext 7.2.5. **Note:** See the Quantum Documentation Portal for more information regarding ActiveScale products and solutions.

Product	Compatibility
Partial File Retrieval	StorNext Partial File Retrieval (PFR) is a product which enables you to quickly retrieve and utilize segments of large media files, rather than the entire file, based on time-code parameters. **Note:** For Quantum Cloud Storage, PFR is not supported for copies with client-side encryption or compression. It is only supported for copies with server-side encryption or without encryption and compression.
StorNext Web Services	StorNext Web Services enables you to run third-party application program interfaces (APIs) with StorNext. To view the latest commands supported by the StorNext Web Services, see StorNext Web Services V2 Commands, Usage, and Descriptions.
Apple Xsan	Xsan is software embedded in macOS that enables Mac clients to connect to the StorNext File System without requiring additional client software. Quantum supplements this solution with StorNext data management software, enabling Apple Xsan customers to use applications running on Windows, Linux, and UNIX with their Xsan and share content across more systems. For information about compatibility between Apple Xsan and StorNext 7.2.5, refer to the StorNext 7.2.5 Compatibility Guide in the StorNext Documentation Center.
Supported Browsers	For information on browsers supported with the StorNext GUI for this release, refer to the StorNext 7.2.5 Compatibility Guide in the StorNext Documentation Center.

For all other components and features, see the StorNext 7.2.5 Compatibility Guide in the StorNext Documentation Center.

General Considerations

This section provides information about items to consider for StorNext 7.2.5.

Checksum Performance Considerations

Note: Generating MD5 checksums is a CPU-intensive operation.

Current StorNext metadata controller and Distributed Data Mover (DDM) hardware is able to calculate MD5 checksums at around 300 MB/s to 500 MB/s. For newer generation tape technology, the maximum throughput might exceed the rate at which the system can generate checksums. In this case, the MD5 checksum calculation will define the throughput of a single data movement operation. With multiple movement streams, MD5 calculations will be done in parallel across the streams and aggregation of performance will be seen.

Upgrading Appliances

Caution: If you have a Deduplication, or Object Storage license, see Change Request Known Issues below in the StorNext Installation, Upgrade, HA, and Other Known Issues on page 26 section before you upgrade.

For instructions on upgrading your firmware, see Upgrade StorNext Software/System Firmware.

Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- Xcellis Workflow Director
- Xcellis Workflow Extender
- Xcellis Foundation
- Artico

Known Issues

The following sections list known issues in this release of StorNext, as well as associated workarounds, where applicable:

Note: If you encounter one or more of the issues listed in this section, please contact Quantum Customer Support and report the issue(s) you encountered. Also inform the support representative whether you were able to successfully work around the issue(s) by using the provided workaround. Doing these things will help Quantum prioritize the order in which known issues are addressed in future StorNext releases.

StorNext File System Known Issues

The table below lists known issues specific to the StorNext File System.

Operating System	Ticket Number	Case Number	Description/Workaround
All	54834	3505208, 3516356	If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file might contain data blocks from old, deleted files. This problem occurs because Windows Explorer sets EOF to the size of the file before it writes the data to the file. This leaves a gap of uninitialized data in the file. **Note:** This problem can also occur with other programs that set EOF beyond the end of data. This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file. **Workaround**: To prevent this problem from occurring on StorNext, you can use the StorNext "client configuration" application's advanced mount option "Restrict Pre-allocation API" on Window systems and the "`protect_alloc=yes`" mount option on Linux systems. This option will set the unwritten parts of the file to zero. When this option is set, non-root users are unable to use the preallocation ioctl. This option also implies `sparse=yes`. For more information on this option, see the man page `mount_cvfs(8)`. The sparse option will introduce some overhead when using Windows Explorer. Before setting the `protect_alloc` option, see the sparse option in `mount_cvfs(8)` for a description of how it changes StorNext behavior.
All	67363	n/a	StorNext 5.4.0.x incorrectly allowed the **Unix ID Mapping** type to be set to **none** when the **Security Model** is set to **acl**. As a result, file systems fail to start when the **Unix ID Mapping** type is set to **none** when the **Security Model** is set to **acl**. Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used. **Workaround**: To prevent this issue, set the **Unix ID Mapping** to either **winbind** or **algorithmic** for any file system where the **Security Model** is set to **acl**. You can make the adjustment before or after upgrading.

Operating System	Ticket Number	Case Number	Description/Workaround
All	75633	n/a	A StorNext NAS client cannot rename a file if the file has the read-only attribute set. This problem only affects StorNext NAS clients. **Workaround** A StorNext NAS client must remove the read-only attribute before it can rename the file.
All	78082	575600	If you set the parameter **audit=true** in the file system configuration file, the result causes all I/O activity, including reads, to be recorded in the **mdarchive**. This can greatly increase the number of updates applied to the **mdarchive** which, in turn, increases the amount of **mdarchive** compaction activity by the FSM. There is one instance where this increased compaction activity caused the **mdarchive** to grow beyond the capacity of the HA shared file system. **Workaround** To work around this issue, Quantum recommends you increase the parameter **metadataArchiveCache** from the default 2 GB to at least 4 GB, but preferably larger when setting audit to **true** in the configuration file.
macOS	66948	322824, 336945	If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client. **Workaround** If you encounter the errors on an Xsan client, contact Apple.

Operating System	Ticket Number	Case Number	Description/Workaround

macOS	75819	n/a	When running older versions of macOS, an Xsan client cannot mount a StorNext File System volume when a cluster name is included in the local **fsnameservers** file. An unexpected EOF reading reply error is displayed.
			
			`# xsanctl mount snfs3`
			`xsanctl: unexpected EOF reading reply`
			
			If you encounter the EOF error on your Xsan client, do the following workaround to prevent the issue.
			
			Workaround # 1
			
			Upgrade the Xsan client to the latest version of macOS supported by StorNext.
			
			Workaround # 2
			
			If upgrading macOS is not possible and your MDC fsnameservers file includes a cluster name, remove the cluster name (**@_cluster_xx** or similar) from the **mysan.configprofile** file **before** you copy the file to the Xsan client.
			
			For example, change:
			
			`10.65.181.158@_cluster()`
			
			to
			
			`10.65.181.158`
			
			See Mount the StorNext File System on Xsan 5.0 (or later) for additional information.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2022	n/a	File System Pooling includes an email notification feature, which is enabled by configuring SMTP settings using the **sntier config** command. These settings are saved in a configuration file within the StorNext configuration directory. As part of recent changes to enhance security, the SMTP password in this file is now masked (replaced with ********) when File System Pooling is restarted. However, the email notification program currently interprets this masked value literally, causing authentication to the user's email account to fail. **Workaround** A script is available to address this issue. Contact Quantum Technical Support and reference Jira ticket **SNXT-2022**.
All	SNXT-2252	n/a	When you create a File System Pooling policy in the UUI, the default interval is set to 0 seconds. This causes the **sntierd** process to trigger the policy in a tight loop, which can impact backend performance. **Workaround:** Set the schedule to a reasonable interval — at least 5 minutes.

StorNext Storage Manager Known Issues

The table below lists known issues specific to StorNext Storage Manager.

Operating System	Ticket Number	Case Number	Description/Workaround

All	43320	1581004	File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.
			Workaround:
			To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:
			Note: This workaround might also help improve the performance of the faster LTO drives by updating the **FS_LTO_BLOCK_FACTOR** system parameter.
			Note: Changes to **FS_xxx_BLOCK_FACTOR** only affects tapes formatted after the change.
			1. Change the **FS_T10K_BLOCK_FACTOR** system parameter from 8 to 32 by adding the following entry to `/usr/adic/TSM/config/fs_sysparm_override`:
			`FS_T10K_BLOCK_FACTOR=32;`
			Note: The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the `FS_T10K_BLOCK_FACTOR` system parameter must not be set to a value that exceeds 32.
			2. Restart Storage Manager to ensure the change in **Step 1** goes into effect:
			`# tsmstop`
			`# tsmstart`
			3. Verify the **FS_T10K_BLOCK_FACTOR** system

Operating System	Ticket Number	Case Number	Description/Workaround

			parameter contains the new value:
			
			`# showsysparm FS_T10K_BLOCK_FACTOR`
			`FS_T10K_BLOCK_FACTOR=32`
			
			4. Save the current copies of your `/etc/fstab` on the MDCs and the DDM clients.
			5. Modify `/etc/fstab` on the MDCs and the DDM clients to use the `auto_dma_write_length` and `auto_dma_read_length` mount options as follows:
			
			`snfs1 /stornext/snfs1 cvfs`
			`rw,auto_dma_write_length=16m,auto_`
			`dma_read_length=16m 0 0`
			
			6. Unmount and re-mount your file systems.
			7. Use new T10K media to store a copy of the file from the disk.
			
			Note: Step 7 is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which **`fsformat`** was run before the change will use the block factor in use at that time. This change will not impact those tapes.

Operating System	Ticket Number	Case Number	Description/Workaround
All	46693	n/a	Executing the command **`snbackup -s`** while a full or partial backup is running might result in a message that `/usr/adic/TSM/internal/locks/backup.lf` is in an invalid format. This is due to the **`snbackup -s`** process reading the **`backup.lf`** status file while the backup process is updating it. **Workaround**: Ignore the message; to clear-up the process, re-execute the command **`snbackup -s`** (provided that the backup is not writing to the **`backup.lf`** status file while **`snbackup -s`** is trying to read it again).
All	47833	n/a	When copying files between media using the CLI command **`fsmedcopy`**, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media. **Note:** This behavior might cause file data segment distribution to be sub-optimal on the destination media. **Workaround**: 1. Execute **fsmedinfo -l** to get the key of the file(s) in question. 2. Execute **fslocate -k** to find the path to that(those) file(s) on disk. 3. Execute **fsfilecopy -r** to make replacement copy (copies) on a new tape.

Operating System	Ticket Number	Case Number	Description/Workaround
All	69265	n/a	Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in **/usr/adic/TSM/logs/tac** which contains the text: `Process fs_moverd on <host> timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.` **Workaround**: 1. Increase the database connection timeout value by adding the following line to **/usr/adic/mysql/my.cnf** under the section labeled **[mysqld] connect-timeout=240**. 2. Cycle the Storage Manager in order to pick up the updated timeout value.
All	69341	n/a	If you have the IBM APFO driver installed and configured, then when you perform an **fsmedread** operation of a partial tape block from a full tape block, the operation can fail with **errno=12**. **Note:** This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures. **Workaround**: To correct this, perform an **fsmedread** operation without the IBM APFO driver.

Operating System	Ticket Number	Case Number	Description/Workaround

All	72993	452722	When DDM is enabled for non-primary systems, kernel error messages are logged for **reservation conflict** because the Primary MDC mounts the tape which sets the reservation to itself and the DDM sets the reservation to itself once the tape is ready.
			
			This issue causes a benign message in the **/var/log/messages** file for **reservation conflict**. Since the unmounting of the tape also resets the reservation back to the primary MDC, this message is generated the next time a DDM (non-primary MDC) accesses a tape.
			
			Note: This issue can result in a large amount of log messages on any machine running DDMs.
			
			Workaround:
			
			To workaround this issue, you can filter and drop the messages so they are no longer logged, as follows.
			
			Note: Reservation conflicts that are a problem also produce RAS alerts, so you can also drop these errors.
			
			Create the following rules on all DDM clients, based on the OS, to drop the messages from **rsyslog**:
			
			For RedHat 6
			
			`# echo ':msg, contains, "reservation conflict" ~' > /etc/rsyslog.d/ignore-reservation-conflict.conf`
			
			`# service rsyslog restart`
			
			For RedHat 7
			
			`# echo 'if $programname == "kernel" and $msg contains "reservation conflict" then stop' > /etc/rsyslog.d/ignore-reservation-conflict.conf`
			
			`# systemctl restart rsyslog.service`

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-1765	825681	If metadata restore is interrupted during file system migration or recovery, the FSM process may leave some inodes with the **Restoring** flag. On restart, the FSM will complete the overall restore operation but leave these inodes with the flag set and incomplete metadata. While no data is affected, this flag is not expected during normal file system operation and will result in an FSM panic when an attempt is made to access the affected files. **Workaround:** Quantum recommends you perform a **cvfsck** following metadata restore which will detect this condition. However, if the cvlogs are reviewed and there is no evidence of FSM restart during the restore, then **cvfsck** can be skipped. If **cvfsck** reports the presence of the **Restoring** flag on files, file system restoration should be restarted from scratch by running **cvmkfs -r** using a fresh copy of metadata archive.
All	SNXT-2282	n/a	If a Distributed Data Mover (DDM) goes offline during an active data transfer, Storage Manager does not fail over to another available DDM. Instead, it retrieves data from an alternate vaulted media copy. Automatic DDM failover is not currently implemented. This may result in longer retrieval times if a DDM becomes unavailable during transfer. **Workaround:** There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.

StorNext GUI Known Issues

The table below lists known issues specific to the StorNext GUI.

Operating System	Ticket Number	Case Number	Description/Workaround
All	69360	n/a	Using **autofs** to mount a StorNext file system on an MDC is not supported when the same file system also has a native mount point. For example, if the StorNext file system **snfs1** is mounted as **/stornext/snfs1**, then the MDC should not also have an **autofs** configuration that mounts it on the MDC in another location such as **/space/snfs1**. Doing so, causes the **fsCheckAffinities** and **fsCheckTsmFilesystemConfig** health checks to fail and generate RAS tickets. Additionally, this might cause the StorNext GUI to fail unexpectedly for certain operations. **Workaround**: There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.

Operating System	Ticket Number	Case Number	Description/Workaround
All	69958	373823	There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation. Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space. The other way to add space to a file system is to create a new stripe group and add this to the file system. When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs **cvupdatefs**. If **cvupdatefs** takes more than 5 minutes to complete, the GUI kills the **cvupdatefs** utility and reports an error. At this point the file system does not start because the configuration file does not match the current state of the metadata. **Workaround**: To address this issue, you can do one of two things: • Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system. • If stripe group expansion is deemed necessary, use the **cvupdatefs** CLI directly instead of through the GUI. This method does not have any time limitations. In the case that the StorNext GUI attempt was made and hit the 5 minute timeout, file system operations can be resumed by performing the following: 1. Restore the previous version of the configuration file. This can be found in the following directory: `/usr/cvfs/data/<fs>/config_history` 2. Run the **cvfsck** utility to verify and potentially correct the metadata. 3. Start the file system.

Operating System	Ticket Number	Case Number	Description/Workaround

All	79549	n/a	You cannot use the StorNext GUI to configure or modify a policy class with the **GOOGLES3** media type. The process fails with an error similar to the following:
			
			`Failed to create storage policy or some of its attributes, e.g. associated directories and schedules.`
			
			`Failed to update steering parameters for Copy [1] for Media Type [GOOGLES3]. Verify that the Media Type and Media Format is correct for each copy defined on the Steering tab.`
			
			Workaround:
			
			To workaround the issue, Quantum recommends you use the **GOOGLE** media type rather than the **GOOGLES3** media type, not only because the **GOOGLES3** media type provides very limited functionality compared to the **GOOGLE** media type, but also because you cannot use the StorNext GUI to configure or modify a policy class with the **GOOGLES3** media type. However, if you have to use the **GOOGLES3** media type, then use the CLI **fsaddclass** command to configure a policy class and the **fsmodclass** command to modify an existing policy class.

Operating System	Ticket Number	Case Number	Description/Workaround

All	80068	n/a	The status and capacity monitors for file systems, libraries, storage disks, and tape drives does not display data on the StorNext GUI home page. This issue is caused as a result of missing font RPM packages on your system.
			Workaround:
			To workaround the issue, Quantum recommends you do the following:
			1. Install the **java-1.8.0-openjdk** package on your system:
			`yum install java-1.8.0-openjdk`
			2. Restart the Apache Tomcat service:
			`systemctl restart stornext_web`

All	80600	n/a	If you have an open StorNext user interface browser session, then the browser page might flash due to an ICEfaces error. This occurs when the browser page displays partial data rows, empty data, or all the data. **Workaround**: To workaround the issue, Quantum recommends you do the following: 1. Log out of your StorNext user interface browser session(s). 2. Close all browser tabs and windows. 3. Open a new browser window and access the StorNext user interface (see Access the StorNext User Interface (new UI introduced in StorNext 7.0.1) or Access the StorNext Software GUI).

Operating System	Ticket Number	Case Number	Description/Workaround
Linux	47954	n/a	The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI. **Workaround**: To workaround this issue, perform the following procedure: 1. Shut down the Safari browser window(s). 2. Restart the Safari browser, and then retry the operation. 3. Uncheck the **Verify SMTP Server Connectivity** box, and then retry the operation. 4. Set **Authentication** to **NONE**, and then retry the operation. 5. Disable the Safari **User names and passwords** AutoFill under **Safari > Preferences > AutoFill**, and then retry operation.
All	SNXT-1329	n/a	When you use the StorNext GUI to configure the **Timeserver** or **Pool** (located within the **Network Time Protocol (NTP)** section on the **Configuration > System > Date & Time** page), and enter multiple comma separated IP addresses in the **Timeserver Pool** field, and then click **Test NTP**, the NTP test hangs and prevents you from accessing the GUI. **Example of Timeserver Pool Entry** `10.65.162.1,10.65.162.2` **Workaround**: Do the following to workaround this issue: 1. Restart the StorNext GUI. 2. Do **not** enter multiple IP addresses in the **Timeserver Pool** field. You can only enter one IP address.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2278	00861871	When you use the StorNext GUI to add a new bucket to an existing namespace, the StorNext GUI incorrectly displays all I/O paths as **AXR**, even when the underlying configuration (verified via **fsobjcfg -l**) is correctly set to **S3COMPAT**. **Workaround:** Use the CLI and execute the following command to add a bucket: `fsobjcfg -add -bucket [bucket_name] -namespace [namespace_name]`

StorNext Installation, Upgrade, HA, and Other Known Issues

The table below lists known issues specific to StorNext installations, upgrade, HA systems, and other areas.

Operating System	Ticket Number	Case Number	Description/Workaround

All	68849	n/a	After an appliance firmware upgrade, you might be unable to use previously functioning tape devices because the `lin_tape` device driver was automatically unloaded during the upgrade.
			Workaround:
			To workaround this issue, rebuild the `lin_tape` device driver as shown in the following example:
			`rpm -e lin_taped`
			`rpm -e lin_tape`
			`rpmbuild --rebuild /root/lin_tape-1.76.06-1.src.rpm`
			`rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-1.76.06-1.x86_64.rpm`
			`rpm -ivh /root/lin_taped-1.76.0-rhel6.x86_64.rpm`

Operating System	Ticket Number	Case Number	Description/Workaround

All	78382	n/a	Quantum installs a self signed certificate (valid for 365 days) in **/usr/cvfs/config/certs/** only when installing a snfs-common RPM and if a preexisting certificate file does not exist.
			
			If you use monitoring software (for example, Zabbix), an expired self signed certificate is flagged after it expires.
			
			Workaround:
			
			To workaround this issue, do the following to update an expired self signed certificate:
			
			1. Stop CVFS on the affected client:
			
			`# service cvfs stop`
			
			2. Update the certificate:
			
			`# cd /usr/cvfs/config/certs`
			
			`# mv server.crt server.crt.orig`
			
			`# mv server.key server.key.orig`
			
			`# openssl req -x509 -newkey rsa:2048 -keyout server.key -out server.crt -days 365 -subj "/C=US/ST=California/L=SanJose/O=Quantum/OU=Demo/CN=`hostname`" -nodes`
			
			3. Start CVFS on the affected client.
			
			`# service cvfs start`

Operating System	Ticket Number	Case Number	Description/Workaround

All	SNXT-389	n/a	The rpm database might corrupt if an rpm operation or a query is unexpectedly interrupted by a signal, reboot, or a system crash on an Xcellis XWD environment in particular, but could also impact a StorNext MDC in a customer-supplied environment and anywhere that the Appliance Controller software or the Unified User Interface (UUI) software runs.
			
			Querying the database might result in errors and report that an installed package is not installed. For example:
			
			`$ rpm -q snfs-server`
			`error: rpmdb: BDB0113 Thread/process 7492/140149665155136 failed: BDB1507 Thread died in Berkeley DB library`
			`error: db5 error(-30973) from dbenv->failchk: BDB0087 DB_RUNRECOVERY: Fatal error, run database recovery`
			`error: cannot open Packages index using db5 - (-30973)`
			`error: cannot open Packages database in /var/lib/rpm`
			`error: rpmdb: BDB0113 Thread/process 7492/140149665155136 failed: BDB1507 Thread died in Berkeley DB library`
			`error: db5 error(-30973) from dbenv->failchk: BDB0087 DB_RUNRECOVERY: Fatal error, run database recovery`
			`error: cannot open Packages database in /var/lib/rpm`
			`package snfs-server is not installed`
			
			These query errors could result in the inability to start StorNext or other Quantum services.
			
			On rare occasions, you might see corruption of the rpm database after a reboot of the system while an rpm activity occured. StorNext, Appliance

Operating System	Ticket Number	Case Number	Description/Workaround

			Controller, and UUI invoke periodic rpm queries that could be running when a reboot occurs that could lead to corruption. The rpm database corruption is more likely to be seen in the event of an ungraceful reboot resulting from a kernel panic, power outage, or SMITH reset, or upon sending the kill signal to a running rpm command which can happen via systemd as part of a normal shutdown/reboot sequence.
			Workaround:
			Run the following command to repair the rpm database:
			`$ rpmdb --rebuilddb`
			After you rebuild the database, run a query to verify the database is repaired:
			`$ rpm -q snfs-server`
			`snfs-server-7.1.1-91E.RedHat7.x86_64`

All	SNXT-547	n/a	You might see the following error messages when the Rocky 8 kernel first boots:
			`DMAR: [Firmware Bug]: No firmware reserved region can cover this RMRR [0x000000006f760000-0x000000006f762fff], contact BIOS vendor for fixes`
			`DMAR: [Firmware Bug]: Your BIOS is broken; bad RMRR [0x000000006f760000-0x000000006f762fff]`
			`DMAR: [Firmware Bug]: RMRR entry for device 1a:00.0 is broken - applying workaround`
			Workaround:
			You can ignore the error messages.
			Note: Your system might require an updated BIOS from the manufacturer to fix and eliminate the warnings.

Operating System	Ticket Number	Case Number	Description/Workaround

All	SNXT-1537	n/a	If your system is running StorNext version 7.2.0 (or later) and Appliance Controller version 4.1.0, you might receive the following operation failure error when you run the **qtmcontroller iscsiadm** command, and your card link goes down:
			
			`Operation failure: Initiator created, but p3p1.100:isca interface link is down, p3p2.100:iscb interface link is down (E-2003)`
			
			Workaround:
			
			To workaround this issue, run the following Appliance Controller command:
			
			`qtmcontroller -c iscsiadm network start`

Operating System	Ticket Number	Case Number	Description/Workaround

All	SNXT-1738	n/a	**Caution: Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.**
			
			Upgrades on systems originating from StorNext version 6.0.6 or earlier might fail in the pre-upgrade check when upgrading to StorNext version 7.2.0 due to the presence the obsolete rpms **openssl098e** and **tg3** being installed. The following message might appear in the **/var/log/DXi/upgrades/Node1.preupgradechk.out** or **/var/log/DXi/upgrades/Node2.preupgradechk.out** log file if these rpms are present:
			
			`Found unexpected RPMs installed in the system: openssl098e tg3`
			
			Note: This issue only affects upgrades to StorNext version 7.2.0 and is fixed in StorNext version 7.2.4. This issue does not impact StorNext version 7.2.2 as that release is only for fresh installations on Xcellis Workflow Director Gen 3 and Xcellis Workflow Extender Gen 3 systems.
			
			Workaround:
			
			To workaround this issue, run the following commands to manually remove the **openssl098e** and **tg3** rpms, if they are installed prior to upgrading to StorNext version 7.2.0:
			
			`# rpm -e openssl098e`
			`# rpm -e tg3`

Operating System	Ticket Number	Case Number	Description/Workaround

All	SNXT-2254	n/a	When performing an upgrade on multiple Xcellis Workflow Extender nodes within a NAS cluster **at the same time**, you may encounter the following error messages:
			
			`Waiting to join NAS cluster`
			`timeout while waiting to join NAS cluster`
			`Failed to join to NAS cluster`
			`Incomplete processing of SNFS NAS packages; Error occurred during processing`
			`The following problems have to be resolved before an upgrade:`
			`The following nodes are not joined to a NAS cluster: 10.xx.xxx.102 (E-7045)`
			
			Workaround:
			
			To workaround this issue, do the following to performing an upgrade on multiple Xcellis Workflow Extender nodes within a NAS cluster:
			• **Upgrade all non-master nodes first.**
			• Verify that each upgraded node has successfully rejoined the cluster before proceeding.
			• After all non-master nodes are back online and healthy, **upgrade the cluster master node last**.
			
			This sequence ensures cluster stability and prevents service disruption during the upgrade process.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2279	n/a	Xcellis Workflow Director and Xcellis Workflow Extender fail to mount SMB/CIFS shares **Workaround:** Starting with StorNext version 7.2.0, mounting SMB/CIFS shares on Xcellis Workflow Director and Workflow Extender platforms is no longer supported. This change occurs because Mellanox OFED drivers replace the default cifs.ko kernel module with a dummy module, preventing SMB/CIFS mounts. This behavior is by design and will not be modified in future releases.
All	SNXT-2317	n/a	The **snauditlogs** utility does not properly handle file and path names that contain the comma (**,**) symbol. If a file or path contains the comma symbol, the resulting line in the csv log displays incorrect values in some fields. **Workaround:** A patch is available. Contact Quantum Support (see Contacting Quantum Support on page 44) and reference Jira ticket **SNXT-2317**.

Quantum UUI Known Issues

The table below lists known issues specific to the Quantum UUI.

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-11364	n/a	The **CLIENTS** page of the UUI displays duplicate clients. **Workaround**: To workaround this issue, do the following. 1. On the **APIGW** page of the UUI, click **Deregister Endpoint** for the endpoint of the duplicate client. 2. Restart the **Unified Connector** on the client that contains the duplicate entries.

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-11386	n/a	If your system is running StorNext 7.1.1 or earlier, the label **N/A** is displayed in the **Reserved** column when you display stripe group information in the UUI. **Workaround**: Upgrade to StorNext 7.2.0 or later to display the **Reserved** value in the UUI.
All	CON-11396	n/a	When you upgrade the UUI software, the upgrade process might appear to be complete but the **Quantum Unified UI Software Upgrade** page does not update due to an **update-service** request that continuously fails with a **404** status code. **Workaround**: To workaround this issue, execute the following command as **root** on the node where the UUI is running: `/opt/quantum/uui-common/bin/uui upgrade remote -y`
All	CON-11487	n/a	When you navigate to the **Metrics > Performance and Metrics > Utilization** UUI page, you might experience a long delay for the page to display properly. **Workaround**: To workaround this issue, refresh your browser page.
All	CON-11504, CON-11505, CON-11506	n/a	If you navigate to the **Metrics > Utilization** page, and attempt to download a **Historical Usage** graph, the image file for the graph might not display properly. **Workaround**: Do the following to workaround this issue: 1. Click the **Print chart** icon in the UUI. 2. Save the document as a PDF.

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-11508	n/a	When you navigate to the **Metrics > Performance and Metrics > Utilization** UUI page, you might receive an error when you view a widget in full screen mode and attempt to print or export a data chart. **Example** `Trouble loading Capacity data. Request aborted.` **Workaround**: To workaround this issue, close your browser window and then reopen a new browser window.
All	CON-11541, CON-11542	n/a	If you navigate to the **Metrics > Performance** page, and attempt to save any data graph as a PNG file or a JPEG file, the image for the graph contains a black background with black text, which results in an illegible data graph. **Workaround**: To workaround this issue, download the data graph in a different format, such as a PDF document, or an SVG vector image.
All	CON-11572	n/a	When you upgrade the UUI software, the upgrade process might appear to be complete but the UUI login page displays an error, preventing you from accessing and logging in to the UUI. **Workaround**: To workaround this issue, execute the following command: **`/opt/quantum/usui/docker/usui-init-cron.sh --force`**

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-11587	n/a	During a fresh installation of the StorNext 7.2 software on the secondary node of an Xcellis system, the UUI Unified Connector client registration is missing after StorNext is converted to an HA pair. On the top navigation menu of the UUI, click **CLIENTS** to verify the issue; you might notice a client registration for both the primary node and the secondary node. **Workaround**: If you are missing the client registration for the secondary node, Quantum recommends you perform either of the following options to workaround this issue. **Workaround Option 1** As the **root** user, SSH to the secondary node and execute the following command: `systemctl restart unified-connector` **Workaround Option 2** Reboot the secondary node.

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-11597	n/a	In the UUI (**Metrics > Dashboards**), the error below might appear in your widget when you perform the following steps: 1. On the top navigation menu, click **Metrics**, and then click **Dashboards** on the left navigation menu. 2. Create a new dashboard. 3. Add a widget for a file system to the dashboard you created in **Step 2**. 4. Stop and start the file system. 5. On the top navigation menu, click **Metrics**, and then click **Dashboards** on the left navigation menu. **Example of error** `File System <name of file system> is no longer exist.` **Workaround**: To workaround this issue, remove the existing widget from dashboard and add a new widget.
All	CON-11618	n/a	If you install the UUI software on a customer-supplied system, the following error message appears in the **./uui-install.sh** script output when the installation completes. **Example of error** `Error: failed to remove gateway config: gateway config empty` **Workaround**: You can ignore the error message.

Operating System	Ticket Number	Case Number	Description/Workaround

All	CON-12453	n/a	When performing a single to dual-server upgrade and restoring the UUI on the secondary node, all USBE requests may return a 404 error.
			
			This affects service availability for components relying on USBE endpoints.
			
			Workaround:
			
			To workaround this issue, do the following to make the secondary node become the primary node UUI.
			
			1. On the **secondary** node, enter the following CLI commands to stop any UUI services:
			
			`/opt/quantum/uui-common/bin/uui autostart disable`
			
			`/opt/quantum/uui-common/bin/uui stop`
			
			2. On the **primary** node, enter the following CLI command to force the UUI to synchronize its configuration:
			
			`/opt/quantum/uui-common/bin/uui_ha_mgr.sh --sync --force`
			
			3. On the **secondary** node, enter the following CLI command to make it the primary node UUI:
			
			`/opt/quantum/uui-common/uui_ha_mgr.sh --primary`
			
			Note: If USBE requests return a 404 error, contact Quantum Support (see Contacting Quantum Support on page 44).

Operating System	Ticket Number	Case Number	Description/Workaround

All	CON-12615	n/a	In the UUI under **System Settings > System Settings**, the description of the **Stale Endpoint Expiration Days** field is not correct. The description appears as:
			
			`0 = default expires in 7 days`
			
			Workaround:
			The correct description of the **Stale Endpoint Expiration Days** field is:
			
			`0 = disabled`

All	CON-12633	n/a	In the UUI under **System Settings > Security Settings**, the default value of the **Automatic User Logout Inactivity Time (minutes)** field on a **fresh installation** is not correct. The value appears as:
			
			`60`
			
			Workaround:
			The correct default value of the **Automatic User Logout Inactivity Time (minutes)** field on a **fresh installation** is:
			
			`15`

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-12640	n/a	After upgrading, some API Gateway (APIGW) endpoints may continue to display outdated version information for the Unified Connector (UC) and Storage Manager. **Workaround**: To workaround the issue, do the following: **For the Unified Connector** • Enter the following to restart the Unified Connector service: `service unified-connector restart` **For Storage Manager** 1. Unregister Storage Manager from the API Gateway. 2. Allow approximately 10 minutes for Storage Manager to automatically re-register with the API Gateway.
All	CON-12658	n/a	The **Last Errored Jobs** functionality does not display an accurate representation of your File System Pooling jobs. **Workaround**: Use the **Jobs** or the **Job History** interface to view an accurate representation of your File System Pooling jobs.
All	CON-12661	n/a	There is an issue that prevents a LAN client from appearing on the **Clients** page for up to 20 minutes or more. **Workaround**: There is currently no workaround for this issue. A fix will be implemented in a future release of the Quantum UUI.

Operating System	Ticket Number	Case Number	Description/Workaround
All	CON-12685	n/a	Users that are specifically assigned the **quantum-flexsync-admin** role are currently unable to view or browse the contents of the **File Browser**. When attempting to access the **File Browser**, the following error is displayed: `Unable to get Folder Content: Request failed with status code 403` This issue appears to be related to insufficient permissions associated with the role. **Workaround**: If possible, assign the user a broader composite role (for example, **quantum-admin** or **quantum-superadmin**) that is known to have full access to the **File Browser**.
All	CON-12739	n/a	When downloading and extracting the Unified Connector package for Windows Server OS from the **Clients > Software Downloads** page, you may encounter the following error: `Error 0x80010135: Path too long` This occurs because the **.msi** file within the archive has a file name or path that exceeds the Windows maximum path length limit, preventing successful extraction. **Workaround:** To avoid this issue, download the Unified Connector package for Windows OS instead. This version contains shorter file paths and extracts successfully on Windows systems.

Operating System	Ticket Number	Case Number	Description/Workaround

All	CON-12741	n/a	When uninstalling the Unified Connector for Windows via the user interface, you may encounter the following warning:
			
			`The following applications should be closed before continuing the install:`
			`Search`
			`Start`
			`New notification`
			
			Workaround:
			You can safely click **Ignore** to proceed with the uninstallation. This warning does not impact the removal process.

All	CON-12742	n/a	When installing or uninstalling the Unified Connector for Windows via the command-line interface (CLI), the package downloaded from **Clients > Software Downloads** contains incorrect script files. **Expected files:** • uc-install.ps1 • uc-uninstall.ps1 **Actual files included:** • uc-install.sh • uc-uninstall.sh **Note:** These are intended for Linux systems and are not compatible with Windows. **Workaround**: To avoid this issue, use the UI for installation and uninstallation on Windows (see Install or Upgrade the Unified Connector).

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2031	n/a	After upgrading, the UUI browser console may not display any devices upon login. **Workaround**: Execute the following command to restart the UUI service on the StorNext node where the UUI service is running: `systemctl restart usbe`
All	CON-12687	n/a	Users that are specifically assigned the **quantum-alert-admin** role are currently unable to remove alerts. This issue appears to be related to insufficient permissions associated with the role. **Workaround**: If possible, assign the user a broader composite role (for example, **quantum-admin** or **quantum-superadmin**) that is known to have full access to alerts.
All	CON-13519	n/a	In older QBSP software versions of the H-Series or F-Series systems (prior to QBSP version 2.0), the **Health** status may appear grayed out even when the system is online and healthy; when you navigate to the **Nodes** UUI page, the health indicator for your node may display an incorrect status. This issue can occur if **Controller B** is registered with the API gateway. **Workaround**: To workaround this issue, ensure that **Controller A** is registered with the API gateway (see API Gateway Registration and Configuration).

Operating System	Ticket Number	Case Number	Description/Workaround

All	CON-13555	n/a	Switching UUI HA primary from one node to another could result in an error if the usbe service had not fully initialized after failover. This issue is caused by the **/usr/src/app/var/cache** directory not being reinitialized during failover, which causes missing proxy files.
			Example of error:
			`Can not get nodes: Request failed with status code 500`
			Workaround:
			To workaround this issue, use the following command to restart the usbe service and restore functionality:
			`systemctl restart usbe`

Contacting Quantum Support

Below is information related to contacting Quantum Support as well as steps to improve your Quantum customer journey.

Chatbot

An AI driven Quantum Chatbot is available to ask product support questions, open a service case, or chat with a call center agent. Locate the Q box on the bottom right of a Quantum web page, such as https://www.quantum.com/en/service-support/.

Note: Some ad blockers might interfere.

Open a Service Case

Use any of the following methods to open a service case:

- AI driven Quantum Chatbot. Locate the Q box on the bottom right of a Quantum web page.
- Visit the MyQuantum portal (for more information, see Use MyQuantum Service Delivery Platform below).

  Note: The MyQuantum portal is the most efficient and preferred method to open a service case.

- Visit the Service & Support page.

  Note: You can also access other Support related services.

- Call Quantum Support (see Service & Support).

Use MyQuantum Service Delivery Platform

MyQuantum is a single portal for everything Quantum. You can view assets, open service cases, receive real-time updates, and search the Knowledge Base and documentation, all through a secure, online portal.

1. Create an account and log in to the MyQuantum Service Delivery Platform.
2. Register a product on MyQuantum.

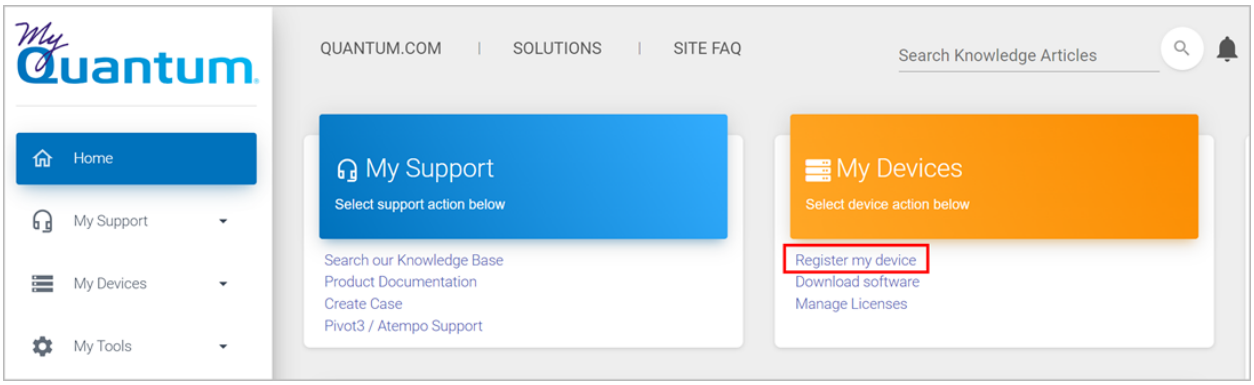

3. Request site access to the Cloud-Based Analytics (CBA) monitoring portal and follow the instructions to set up product(s) to connect to CBA. You can use CBA to monitor Quantum products remotely, from a single dashboard, and Quantum Support can use it to help troubleshoot products more efficiently.

Use Cloud Based Analytics (CBA)

Quantum products are equipped with a Cloud Based Analytics (CBA) agent that can provide log files and snapshots to Quantum CBA servers that are running in the cloud.

CBA enables Quantum systems to collect data regarding system and environment performance. The collected data is bundled and uploaded to the remote CBA server for analysis. You can access Quantum system performance and health results on the CBA dashboard (at https://insight.quantum.com) or through the MyQuantum Service Delivery Platform.

The CBA dashboard displays the analytic results of the uploaded CBA data using flexible charting tools, along with an overall health score of each Quantum system configured for the CBA account.

Refer to product documentation for product-specific information related to CBA.

Refer to the Quantum CBA website for general information about CBA.

Escalate a Service Case

To escalate a service case, follow the process documented here: https://www.quantum.com/en/service-support/resources/escalation/

Contact Quantum Sales

https://www.quantum.com/en/company/contact-us/

Quantum technology, software, and services provide the solutions that today's organizations need to make video and other unstructured data smarter – so their data works for them and not the other way around. With over 40 years of innovation, Quantum's end-to-end platform is uniquely equipped to orchestrate, protect, and enrich data across its lifecycle, providing enhanced intelligence and actionable insights. Leading organizations in cloud services, entertainment, government, research, education, transportation, and enterprise IT trust Quantum to bring their data to life, because data makes life better, safer, and smarter. Quantum is listed on Nasdaq (QMCO) and the Russell 2000® Index. For more information visit www.quantum.com.

www.quantum.com | 800-677-6268