Quantum

StorNext 7.2.6 Release Notes

Product/Software Release Date	August, 2026

Contents

What's New in StorNext 7.2.6

Purpose of this Release

This document describes the new features, enhancements, and resolved issues included in StorNext 7.2.6. For more information, see Fixed Issues and Enhancements Addressed in StorNext 7.2.6 on page 9.

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

New Features and Enhancements

Expanded Operating System Support

This release extends client support to the following operating systems:

Operating System	Supported Client Types
RHEL / Rocky Linux / AlmaLinux 9.7 and 9.8	SAN, DLC, MDC, SNSM, DDM, DLS
RHEL / Rocky Linux / AlmaLinux 10.1	SAN, DLC
Ubuntu 26.04 LTS	SAN, DLC
Ubuntu 25.10	SAN, DLC
Debian 13	SAN, DLC (now included in shipping package)

Security Updates

This release includes several security improvements:

- **Appliance Kernel Update (Multiple CVEs)**
  - **Issue:** The appliance kernel contained multiple known security vulnerabilities.
  - **Cause:** The previously shipped kernel version predated patches for several CVEs identified in ongoing security scanning.
  - **Resolution:** The appliance kernel has been updated to version 4.18.0-553.141.1, addressing multiple CVEs.
- **CVE-2026-31431: SSH Authentication Vulnerability on Xcellis Appliances**

- **Issue:** Xcellis appliances were vulnerable to an SSH authentication bypass (CVE-2026-31431, "Copy Fail").
- **Cause:** A vulnerability in the SSH authentication stack allowed potential unauthorized access under specific conditions.
- **Resolution:** The SSH authentication vulnerability has been remediated on Xcellis appliances in this release. Customers are strongly encouraged to upgrade.

- **CVE-2025-14847: MongoDB zlib Memory-Read Vulnerability in UUI**
  - **Issue:** The MongoDB instance bundled with the Unified User Interface (UUI) was vulnerable to a zlib memory-read issue (CVE-2025-14847).
  - **Cause:** The previously shipped version of MongoDB contained a zlib vulnerability that could allow unintended memory reads.
  - **Resolution:** MongoDB in UUI has been updated to a version that addresses CVE-2025-14847.

Upgrade and Installation

- **Mellanox Drivers Not Updated During Upgrade on NVMe-oF Systems**
  - **Issue:** When upgrading from StorNext 7.2.4 to 7.2.5 or 7.2.6 on systems configured for NVMe-oF, Mellanox drivers were not properly updated, which could result in degraded or non-functional NVMe-oF connectivity after upgrade.
  - **Cause:** The upgrade process did not correctly detect and reinstall Mellanox drivers on NVMe-oF-configured systems.
  - **Resolution:** The upgrade process now correctly updates Mellanox drivers on systems configured for NVMe-oF when upgrading from 7.2.4 to 7.2.5 or 7.2.6.
- **Mellanox Driver Recompilation Packages Now Included on Xcellis**
  - **Issue:** If Mellanox drivers needed to be reinstalled after an upgrade, the RPM packages required to rebuild them were not present on the system, requiring a manual download step.
  - **Cause:** The Mellanox driver recompilation RPMs were not included in the Xcellis software bundle.
  - **Resolution:** The RPM packages needed to rebuild Mellanox drivers are now included on Xcellis systems, eliminating the need for a manual download if driver reinstallation is required.
- **NTFS-Formatted USB Drives Now Supported for Platform Upgrades**
  - **Issue:** USB drives formatted with NTFS (Windows format) could not be used to perform platform upgrades via the service menu, requiring customers to reformat the drive before use.
  - **Cause:** The service menu upgrade process did not support reading from NTFS-formatted media.
  - **Resolution:** The service menu upgrade process now supports USB drives formatted with NTFS, in addition to previously supported formats.
- **Firmware Check Error During Upgrade on Systems with Multiple QLogic HBA Cards**
  - **Issue:** On systems with multiple QLogic HBA cards installed, the upgrade process produced a firmware check error and could fail to complete successfully.

- **Cause:** The firmware validation logic did not correctly handle configurations with more than one QLogic HBA card.
- **Resolution:** The firmware check logic has been corrected to handle multi-QLogic HBA configurations, allowing upgrades to complete without error.

StorNext File System

- **SMB and NFS Shares Become Unresponsive When Accessing Named Streams**
  - **Issue:** SMB and NFS shares could become unresponsive (hang) when clients accessed files with named streams (resource forks), such as those created by macOS clients. This required a service restart to recover.
  - **Cause:** A deadlock condition existed in the code path handling named stream access, which could cause the share service to stop responding under concurrent access patterns.
  - **Resolution:** The deadlock condition has been resolved. SMB and NFS shares remain responsive when named streams are accessed. Customers using SMB or NFS with macOS clients are strongly encouraged to upgrade to this release.
- **Buffer Cache Flusher Daemons Fail to Start with Many File Systems Using privatebufcache**
  - **Issue:** On clients with a large number of mounted StorNext file systems configured with **privatebufcache**, buffer cache flusher daemons could fail to start. This could result in data not being flushed to disk in a timely manner, increasing the risk of data loss in the event of an unexpected system failure.
  - **Cause:** A resource limit was exceeded when the system attempted to start flusher daemons for a large number of file systems simultaneously, causing some daemons to fail silently.
  - **Resolution:** The flusher daemon startup process has been corrected to reliably start all required daemons regardless of the number of mounted file systems using **privatebufcache**.
- **Assertion Failure (Core Dump) During qrtree Index Compaction Under High Write/Delete Load**
  - **Issue:** On file systems with high write and delete activity, an assertion failure (core dump) could occur during qrtree index compaction, causing the metadata controller to crash and requiring a restart.
  - **Cause:** A race condition existed in the **qrtree** index compaction code path that could be triggered under concurrent high-volume write and delete operations.
  - **Resolution:** The race condition in **qrtree** index compaction has been resolved. Systems under high write and delete load no longer encounter this assertion failure.

Storage Manager

- **vsarchiveconfig Cannot Communicate with Tape Library Devices on RHEL 9.5 and Later**
  - **Issue:** On systems running RHEL 9.5 or later, the **vsarchiveconfig** utility could not communicate with tape library devices, preventing Storage Manager from managing tape archives.
  - **Cause:** A kernel compatibility change introduced in RHEL 9.5 broke the communication interface used by **vsarchiveconfig** to interact with tape library hardware.

- **Resolution: vsarchiveconfig** has been updated to use a compatible communication method on RHEL 9.5 and later kernels. Customers using Storage Manager with tape libraries on RHEL 9.5+ should upgrade to this release.

Tools and Utilities

- **snwebsetup -a Option for OFM Store Actions Does Not Work**
  - **Issue:** The **snwebsetup -a** option, used to enable Offline File Manager (OFM) store actions, did not function correctly, preventing customers from enabling this feature via the command line.
  - **Cause:** A defect in the **snwebsetup** argument handling caused the -a flag to be ignored.
  - **Resolution:** The **snwebsetup -a** option now correctly enables OFM store actions as documented.
- **sntier Input Validation and Documentation Improvements**
  - **Issue:** The **sntier** utility accepted invalid input without error, **--match** and **--exclude** regex pattern behavior was undocumented, and guidance on job suspension timing was absent, leading to user confusion and unexpected behavior.
  - **Cause:** Input validation was incomplete, and documentation for regex options and job suspension behavior had not been written.
  - **Resolution:** Input validation for **sntier** has been improved to reject invalid arguments with a clear error message. The **--match** and **--exclude** regex pattern options are now documented, and guidance on job suspension timing has been added to the man page and help output.
- **snauditlogs CSV Output Incorrectly Formats File Names Containing Commas or Quotes**
  - **Issue:** When using **snauditlogs** to export audit data in CSV format, file names containing commas or quotation marks were not correctly escaped, resulting in malformed CSV output that could not be reliably parsed by spreadsheet applications or log analysis tools.
  - **Cause**: The CSV output routine did not apply RFC 4180-compliant quoting and escaping to field values containing special characters.
  - **Resolution:** The **snauditlogs** CSV output now correctly quotes and escapes file names containing commas or quotation marks, producing well-formed CSV output.
- **StorNext SNMP OIDs Report Incorrect Data Types**
  - **Issue:** Several StorNext SNMP OIDs reported values with incorrect data types, causing compatibility issues with SNMP monitoring tools that perform strict type validation.
  - **Cause:** The SNMP MIB definitions for the affected OIDs specified incorrect data types that did not match the actual values being reported.
  - **Resolution:** The affected SNMP OIDs now report values with the correct data types as defined in the StorNext MIB, improving compatibility with SNMP monitoring and management tools.
- **Windows SNFS Clients Not Visible in Legacy GUI Client Download Area**
  - **Issue:** Windows SNFS clients did not appear in the legacy GUI client download area, preventing administrators from easily locating and downloading Windows client packages through the GUI.

- **Cause:** A display filtering defect in the legacy GUI caused Windows client entries to be omitted from the client download list.
- **Resolution:** Windows SNFS clients now appear correctly in the legacy GUI client download area.

Unified User Interface (UUI)

- **Debian 13 Clients Incorrectly Reported as Unhealthy in UUI Clients Tab**
  - **Issue:** Debian 13 clients were displayed as "unhealthy" in the UUI Clients tab even when operating normally, causing unnecessary alarm for administrators monitoring client health.
  - **Cause:** The UUI health-check logic did not recognize Debian 13 as a supported platform, causing it to default to an unhealthy status for clients running that OS.
  - **Resolution:** The UUI health-check logic has been updated to correctly recognize Debian 13 clients. Debian 13 clients in a healthy state are now reported as healthy in the UUI Clients tab.

Appliance

- **Service Menu SetDateTime Does Not Accept Years Beyond 2025**
  - **Issue:** The **SetDateTime** option in the appliance service menu rejected year values beyond 2025, preventing administrators from setting the correct date and time on systems deployed or serviced after that year.
  - **Cause:** The year validation logic in the service menu contained a hard-coded upper limit of 2025.
  - **Resolution:** The year validation logic has been updated to accept current and future year values. The **SetDateTime** service menu option now correctly accepts years beyond 2025.
- **Incorrect Unified Connector Version Installed During Upgrade on Xcellis Gen 1**
  - **Issue:** When upgrading an Xcellis Gen 1 appliance, the upgrade process installed the wrong version of the Unified Connector, which could result in compatibility issues or degraded functionality.
  - **Cause:** The upgrade logic did not correctly distinguish between Gen 1 and newer Xcellis platforms when selecting the appropriate Unified Connector package, causing the non-Rocky 8 version to be installed on Gen 1 systems.
  - **Resolution:** The upgrade process now correctly identifies Xcellis Gen 1 appliances and installs the Rocky 8 version of the Unified Connector as intended.

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

Client Support

- Added support for **Ubuntu 26.04** clients, including **UUI** download packages and **Unified Connector** compatibility.
- Added **Unified Connector** support for **SLES 16.0**.
- Added **Unified Connector** support for **AlmaLinux**, **Oracle Linux**, **RHEL**, and **Rocky Linux 10.0 and 10.1**.
- Updated **USBE** client packages.

Fixes

- Fixed an issue that prevented updated StorNext client packages from being made available in **/usr/cvfs/CLIENTS**.
- Fixed an issue in the **USBE log server** upgrade script.
- Fixed an issue on the **Clients** tab of the **First Boot Wizard** where the EULA page displayed an outdated date and truncated the word "**Global**".
- Fixed an issue that could prevent the **Software Downloads** page from loading correctly from the **Clients** tab.
- Fixed an issue where the **Software Downloads** page delivered **Unified Connector 7.2.6.1** instead of the current version, **7.2.6.2**.
- Added support for **multi-select LUN assignment** in **Tiering Pools**.

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 7.2.6 only supports FlexSync 3.1.0 (or later). If you want to install StorNext 7.2.6 on your system, or upgrade your system to StorNext 7.2.6 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 7.2.6 Compatibility Guide provides the basic compatibility for StorNext 7.2.6, 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 7.2.6

Ticket Number	Description
SNXT-182	**Issue** The **snhistory** command fails with the error "Request failed - FSM returned result code 4!" when querying a heavily loaded mdarchive with long retention times or when queries span a large time range. **Cause** Long-running **snhistory** queries against large mdarchive databases could exceed internal processing limits, resulting in an FSM error. **Resolution** Fixed the **snhistory** query handling to properly support long-duration queries against large mdarchive databases.

Ticket Number	Description
SNXT-1537	**Issue** When configuring Mellanox ConnectX-5 or ConnectX-6 network cards for iSCSI using the **qtmcontroller iscsiadm create** command, the network interfaces go offline and the command reports "interface link is down." Users must manually bring the interfaces back up (e.g., **ifconfig <portname> up** or **qtmcontroller -c 'iscsiadm network start'**) and then re-run the configuration command for it to succeed. This affects both Xcellis Gen 2 and Xcellis Gen 3 systems. **Cause** The iSCSI initiator creation process resets the Mellanox network interfaces as part of the configuration work-flow. The CX5 and CX6 cards do not recover the link state in time before the configuration script checks connectivity, causing a temporary link-down condition that is reported as a failure. **Resolution** The iSCSI configuration work-flow now properly handles the interface reset sequence for ConnectX-5 and ConnectX-6 cards, ensuring ports remain online throughout the configuration process without requiring manual intervention.
SNXT-1541	**Issue** When importing media volumes from one StorNext system to another using **fsrecover**, some files report "PATH UNKNOWN" and cannot be recovered to their destination directories. **Cause** An internal message size limitation of 8192 bytes prevented **fsrecover** from processing files with very long path names during media import operations. Files with paths exceeding this limit failed to recover. **Resolution** Increased the internal message size limit to accommodate files with long path names during media import and recovery operations.
SNXT-2287	**Issue** After upgrading to StorNext 7.2.4 or later on Rocky Linux 8, the StorNext SNMP agent does not return the full OID tree. An **snmpwalk** on the StorNext root OID (1.3.6.1.4.1.3697.1.50) returns only top-level OIDs; many expected sub-OIDs are missing. **Cause** A platform compatibility issue specific to Rocky Linux 8 caused the StorNext SNMP agent to return incomplete OID data, whereas the same agent functioned correctly on CentOS 7. **Resolution** Fixed the StorNext SNMP agent to correctly return the full OID tree on Rocky Linux 8 systems.

Ticket Number	Description
SNXT-2329	**Issue** The MongoDB instance used by the StorNext Unified User Interface (UUI) was affected by CVE-2025-14847, a vulnerability in which zlib compressed protocol header length confusion may allow unauthorized memory reads. **Cause** The bundled MongoDB version (7.0.18) included a vulnerable zlib compression handler in its network protocol stack. **Resolution** Updated the MongoDB component in the UUI to a patched version that addresses CVE-2025-14847.
SNXT-2450	**Issue** SMB and NFS shares served by a StorNext gateway can become unresponsive, requiring a restart to restore access. **Cause** A lock ordering violation in the SNFS client occurred when accessing named streams (resource forks) through the extended attribute interface. The code acquired a node-level lock before the node-change lock, which could deadlock with other operations that acquire these locks in the correct order. **Resolution** Corrected the lock acquisition order to prevent the deadlock condition when accessing named streams.
SNXT-2500	**Issue** The **fw_report.sh** and **ClearToShip.sh** scripts fail on systems with two or more QLogic Fibre Channel HBA cards installed. **Cause** The firmware utility scripts assumed device numbering always starts at index 0 for each card. On systems with multiple QLogic FC cards, the second card's devices are numbered starting at a higher index (e.g., 4, 5, 6, 7), causing a look-up failure. **Resolution** Corrected the device enumeration logic in the firmware utility scripts to properly handle multiple QLogic FC HBA cards.

Ticket Number	Description
SNXT-2503	**Issue** An Admin Alert reporting "Failure detected checking FirmwareBundle updates" is generated during or after upgrade to StorNext 7.2.5 on systems with multiple QLogic cards. **Cause** The firmware version check performed during the upgrade process encounters the same device-indexing error described in **SNXT-2500**, causing the FirmwareBundle verification step to fail and triggering an Admin Alert. **Resolution** Corrected the device enumeration logic in the firmware utility to properly identify all installed QLogic cards during the upgrade firmware check, preventing the erroneous alert.
SNXT-2556	**Issue** When new storage pools are created on an existing file system that already contains files, running **sntier** with the **--fill** option does not relocate any files, even though the pool utilization exceeds the fill threshold. **Cause** The **--fill** policy searches for candidate files by matching the pool affinity. Files that existed before the pools were created have no pool affinity assigned, so they are not detected as candidates for relocation. **Resolution** Fixed the candidate selection logic so that **sntier --fill** correctly identifies and relocates pre-existing files that have no pool affinity assigned. Updated documentation with guidance for this scenario.
SNXT-2557	**Issue** After issuing a suspend command for a running **sntier** job, the job does not suspend promptly and may continue processing for an extended period before stopping. **Cause** The internal suspend flag was only evaluated after completing the current processing loop iteration, which could involve scanning and moving a large number of files. **Resolution** Improved the suspend check frequency so that **sntier** jobs respond to suspend requests more promptly.

Ticket Number	Description
SNXT-2564	**Issue** The **sntier --match** and **--exclude** options accept pattern-matching expressions, but the supported syntax and limitations are not documented, leading to unexpected behavior. **Cause** The man page did not specify the supported regex-style patterns (glob/extended patterns) or clarify that matching applies only to the file's base name (not the full path). **Resolution** Updated the man page to describe the supported pattern syntax, clarify that patterns match against the file name only (not directory paths), and provide usage examples.
SNXT-2580	**Issue** On SNFS clients with many file systems using the **privatebufcache** mount option, flusher daemons may fail to start, leading to potential I/O issues. The message "Can't enqueue CvBCFlusher" appears in the logs. **Cause** Each distinct buffer cache requires its own flusher daemon, but the internal daemon queue was limited in size (25 or fewer slots). Configurations with many private buffer caches could exhaust this queue, preventing flusher daemons from being created. **Resolution** Increased the daemon queue capacity to support configurations with a large number of distinct buffer caches.
SNXT-2594	**Issue** The StorNext metadata archive (mdarchive) process may unexpectedly crash with an assertion failure in the spatial index compaction routine. **Cause** A race condition in the spatial index code allowed concurrent delete operations to under-count stale entries due to a non-atomic counter increment. Over time, this caused the counter to drift from the actual number of stale entries, eventually triggering an assertion failure during index compaction. **Resolution** Made the stale entry counter updates atomic to eliminate the race condition and prevent the assertion failure.

Ticket Number	Description
SNXT-2616	**Issue** Running **cvfsck** on a file system with auditing enabled reports thousands of false errors stating that audit records (SNBT_PAYLOAD_XATTR_SYS) are illegal for named stream inodes, and proposes to delete them. **Cause** The **cvfsck** validation logic did not include audit records in the allowed payload list for named stream (resource fork) inodes, even though these records are valid and functional. **Resolution** Updated **cvfsck** to correctly recognize audit records as valid on named stream inodes, eliminating the false error reports.
SNXT-2637	**Issue** When a Windows SMB client accesses a StorNext file system reshared by a native Windows DLC (Distributed LAN Client), the volume properties display the full file system size instead of the user's quota allocation. **Cause** The quota lookup used the default user ID (60001) instead of the actual Active Directory user's identity when querying quota information over the SMB reshare path. This caused the system to return the default user's quota rather than the connected user's quota. **Resolution** Corrected the quota query logic to properly resolve the connected user's identity, ensuring that the correct quota limits are reported to Windows SMB clients on DLC reshare configurations.
SNXT-2726	**Issue** When attempting to upgrade a Quantum Appliance using a Windows-formatted (NTFS) USB drive connected to an Xcellis node, the **Platform Upgrade Menu** fails to mount the USB device and the upgrade image cannot be loaded. **Cause** The upgrade utility did not support NTFS-formatted USB drives. While the system kernel includes NTFS support via FUSE, the upgrade script's mount logic did not account for this file system type, causing all mount attempts to fail even when the device and its partitions were otherwise accessible. **Resolution** The Platform Upgrade utility has been updated to correctly detect and mount NTFS-formatted USB drives, allowing upgrades from Windows-formatted media to proceed as expected.

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 7.2.6 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 7.2.6.

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 7.2.6

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 7.2.6 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 7.2.6 Compatibility Guide in the StorNext Documentation Center for information about compatibility between ActiveScale (AXR, S3) or P100/X100, and StorNext 7.2.6. **Note:** See the Quantum Documentation Portal for more information regarding ActiveScale products and solutions.
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 7.2.6, refer to the StorNext 7.2.6 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 7.2.6 Compatibility Guide in the StorNext Documentation Center.

For all other components and features, see the StorNext 7.2.6 Compatibility Guide in the StorNext Documentation Center.

General Considerations

This section provides information about items to consider for StorNext 7.2.6.

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 on the next page in the StorNext Installation, Upgrade, HA, and Other Known Issues on page 39 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-713	734265	**Issue**
			
			When offloading a stripe group with **sgoffload** (or **sgmanage --offload**), the command may report that the retry limit has been exceeded even though the stripe group has been successfully emptied. The command may terminate with messages similar to the following:
			
			`The stripe group is not empty, N extents remain - retrying...`
			
			`Warning: Retry count exceeded and the stripe group is not empty.`
			
			Cause
			
			This issue occurs because **sgoffload** verifies completion by querying metadata information that may not immediately reflect the results of a large offload operation. After a large number of extents have been moved, there can be a delay before the metadata status is fully updated. As a result, the offload operation may have already completed successfully, or be in the process of completing, but **sgoffload** is unable to verify completion within its fixed number of retry attempts.
			
			Workaround
			
			Re-run the same **sgoffload** command. The operation is safe to run multiple times and will continue from its current state.
			
			• If the reported extent count decreases with each run, the offload is progressing normally. Continue re-running the command until the stripe group is reported as empty.
			• If the reported extent count remains unchanged across multiple runs, contact Quantum Technical Support and reference Jira ticket **SNXT-713**.

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.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2637	00873779	**Issue** In StorNext 7.2.6 and later, on a StorNext file system with quotas enabled, viewing file system properties from a Windows client now reports the user's own quota (allotted capacity and remaining availability) instead of the total capacity and free space of the entire file system. However, if StorNext cannot resolve a caller's uid/gid, that caller is mapped to the user "nobody" and will see the quota amounts configured for "nobody." This can occur when a user logs on to a Windows system as a local user, or as any user whose credentials are valid on a StorNext-connected Windows system but are not stored in Active Directory or another credential repository. Local Windows processes that scan the StorNext file system are also resolved to "nobody," and this information may be surfaced to users. If no quota is configured for "nobody," such a caller will see the full capacity of the StorNext file system. **Cause** When StorNext receives a file system properties request from a Windows client, it attempts to resolve the caller's identity via Active Directory (or another configured credential source). If the lookup fails, for example, because the user is a local account or their credentials are not present in AD, StorNext maps the caller to the default "nobody" uid/gid and returns the quota associated with that identity. Without a quota explicitly set for "nobody," no limit is applied and the full file system capacity is reported. **Workaround** To prevent clients that map to "nobody" from viewing the entire file system capacity, configure a quota for both the uid and gid "nobody" with the desired limit using the **snquota** command. This ensures that any unresolved callers see a controlled capacity value rather than the full file system size.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2758	n/a	**Issue** When a storage pool configured with **affinityPreference = true** (lenient mode) runs out of space, file data is silently redirected to a stripe group in a different pool. However, file system pooling commands (**sntier location**, **sntier pool --list**, **sntier jobs --show**) continue to report the affected files under their originally-designated pool, creating a mismatch between the reported pool affinity and the file's actual physical location. **Cause** The ENOSPC fallback allocator redirects writes to any available stripe group to prevent write failures, but does not update the file's stored affinity tag. Since file system pooling commands read the affinity tag rather than inspecting physical data placement, the reported pool no longer reflects where the file's data actually resides. **Workaround** Use **snfsdefrag -e <file>** to determine a file's true physical location. To reconcile the affinity tag with the data's actual location, or to relocate the data back to the originally-designated pool once space is available; run **snfsdefrag -k <pool> -m 0 <file>** (add **-r** to recurse over a directory).

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.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2726	n/a	**Issue** The Quantum Appliance **Platform Upgrade Menu > Upgrade from local media** option in the **Service Menu** (**/opt/DXi/scripts/service.sh**) could fail to mount a USB drive in two scenarios: • **NTFS-formatted drives.** NTFS and exFAT share the same MBR partition-type ID, so the mount logic misidentified NTFS partitions as exFAT and always attempted an exFAT-specific mount with no fallback, causing genuine NTFS partitions to fail to mount even though the same drive mounted successfully when done manually. • **Partition tables with gaps or partition numbers beyond the checked range.** Partition detection either counted the number of partition device nodes present or checked only a fixed range of partition numbers (1 through 4), so a partition could be skipped entirely if its number fell outside that range or the count didn't line up with the actual numbering. **Cause** The Platform Upgrade utility does not support NTFS-formatted USB drives. When a Windows-formatted (NTFS) USB drive was inserted, the utility attempted to mount the device using only the filesystem types natively supported by the upgrade environment, which does not include NTFS. As a result, all mount attempts failed and the upgrade image could not be loaded from the USB drive. **Workaround** Reformat the USB drive as **exFAT**, or as a Linux file system (for example, **ext4**) on a partition table with no gaps in partition numbering.

Operating System	Ticket Number	Case Number	Description/Workaround
All	SNXT-2828	881120	**Issue** Commands in **/usr/adic/TSM/bin** that require database access may crash with a segmentation fault when run outside an elevated shell. Running a command from **/usr/adic/TSM/bin** that requires database access (for example, **fsmedinfo**) directly from a standard user account, rather than from within an elevated administrative shell (**sudo rootsh**), may cause the command to terminate with a segmentation fault instead of completing normally or returning an error message. **Cause** These commands locate the local MySQL installation under **/usr/adic/mysql** using the **MYSQL_HOME** environment variable, which is normally established by the StorNext administrative environment. If **MYSQL_HOME** is not set, the command instead attempts to detect the MySQL installation itself, but that detection only succeeds for root or members of the **adic** group, since **/usr/adic/mysql/bin** is restricted to those accounts. For any other user without **MYSQL_HOME** already set, this detection fails, leaving an internal database-connection routine in an inconsistent state and causing the command to crash rather than return a clean error. **Workaround** Any of the following will avoid the issue: • Run the affected command from within an elevated administrative shell (**sudo rootsh**) rather than from a standard user account. • Add the affected user account to the **adic** group. Active sessions for that user must be logged out and back in before the group membership change takes effect. • Before running the affected command, source the StorNext environment profile: `source /usr/adic/.profile`

Operating System	Ticket Number	Case Number	Description/Workaround
			Note: This may produce **Permission denied** errors when accessing **/usr/adic/MSM/.profile** and **/usr/adic/wsar_agent/.profile**. These messages can be safely ignored as long as the **MYSQL_HOME** environment variable is set after setting up the environment.

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-2103	n/a	**Issue** On Xcellis systems configured for high availability (HA), if an upgrade to StorNext 7.2.4 or later is interrupted, the upgrade may fail to restart. The GUI may display the error **Failed to connect to Host on Port 81: Transport error: 401 Error** on the **System > Network** page. **Cause** An interruption during the upgrade process can cause the **/usr/adic/tomcat/secrets/wsuser** file to fall out of sync between HA nodes. When the wsuser secrets are mismatched, the nodes can no longer authenticate with each other's upgrade web services, blocking the upgrade from being resumed. **Workaround** Resynchronize the wsuser secrets using one of the following methods, depending on the StorNext version running on each node: • **If the primary node runs a StorNext version before 7.2.4:** On the standby node, run **/usr/adic/gui/scripts/generate_wsuser_password.sh -p wsuser**. It is expected to report that the **Password update failed on** the primary node's IP address since the primary is running a version that does not support this function. • **If the primary node runs StorNext 7.2.4 or later:** On the standby node, run **/usr/adic/gui/scripts/generate_wsuser_password.sh -f**. If the HA shared file system is unavailable, manually copy **/usr/adic/tomcat/secrets/wsuser** from the primary node to the same location on the standby node. Once the wsuser secrets are resynchronized and the underlying upgrade failure is resolved, the upgrade can be retried.

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-2838	n/a	**Issue**
			The Offline File Manager (OFM) application may fail to connect to the API Gateway from a remote host, preventing the file system pooling functionality from being available.
			Cause
			During High Availability (HA) setup, the secondary node registers with the API Gateway using the localhost address (**127.0.0.1**). Because the gateway configuration is stored on the HA shared filesystem, an extra localhost entry can persist when the secondary node becomes the primary, confusing OFM when it attempts to connect remotely.
			Workaround
			As a workaround, remove the existing gateway registration and then register the MDC gateway again.
			1. Remove the current gateway registration:
			`snrest gateway remove 127.0.0.1`
			2. Register the gateway using the MDC IP address:
			`snrest gateway register <MDC_IP_address>`
			Note: Replace **<MDC_IP_address>** with the IP address of your MDC.

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.
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.

Operating System	Ticket Number	Case Number	Description/Workaround
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.
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.

Operating System	Ticket Number	Case Number	Description/Workaround
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`
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 57).

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`

Operating System	Ticket Number	Case Number	Description/Workaround

All	CON-14056	n/a	**Issue**
			
			After upgrading a StorNext Appliance to version 7.2.6, the Unified User Interface (UUI) may display a "**Cannot get nodes: Request failed with status code 500**" error upon login, preventing the Nodes page from loading.
			
			Cause
			
			During the upgrade process, certain UUI backend service (USBE) software files stored in a Docker volume may become missing or corrupted (for example, due to an unclean shutdown), causing the UUI service to fail when handling requests.
			
			Workaround
			
			Run the following commands as **root** on the system where the UUI is running to regenerate the missing USBE files:
			
			`rm -f /var/lib/docker/volumes/usbe_data/_data/usbe.tar.gz.md5`
			
			`/opt/quantum/usbe/docker/post-install.sh`
			
			The UUI should resume normal operation after the **post-install.sh** script completes.

Contacting Quantum Support

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

Open a Service Case

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

- 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.

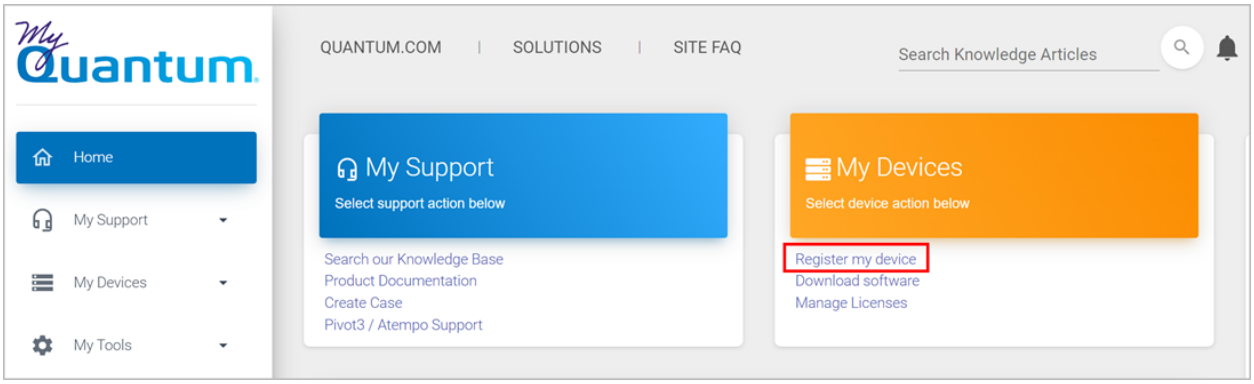

3. Request site access to the Cloud-Based Analytics (CBA) monitoring portal and follow the instructions to set up product(s) to connect to CBA. You can use CBA to monitor Quantum products remotely, from a single dashboard, and Quantum Support can use it to help troubleshoot products more efficiently.

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