Quantum

StorNext 7.2.4
Release Notes

Original Product/Software Release Date	May, 2024

Contents

What's New in StorNext 7.2.4

Purpose of this Release

StorNext 7.2.4 is a maintenance and compatibility-focused release that enables seamless upgrades across all Xcellis hardware generations (Gen 1, Gen 2, and Gen 3 – see Upgrade Enablement below). This release consolidates critical software fixes listed in the section Fixed Issues and Enhancements Addressed in StorNext 7.2.4 on page 8, and expands hardware and operating system support.

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

Enhancements

Upgrade Enablement

Direct upgrade paths are now available for Xcellis Gen 1 system and Xcellis Gen 2 systems from StorNext version 7.1.0 or later. Xcellis Gen 3 systems can upgrade directly from StorNext version 7.2.2.

Bug Fix Consolidation

Fixes from StorNext version 7.2.2 (previously Xcellis Gen 3-only) are now available on Xcellis Gen 1 and Xcellis Gen 2 platforms.

Device Mappings

Increased the default value of the **max_dev** mount option from 512 KiB to 2 MiB to support larger device mappings.

Enhanced HA Configuration Security

Introduction of a **Secret** field in the GUI to strengthen security by requiring a shared secret value between nodes during HA configuration.

How It Works

The secret value is predefined on the secondary node and must be retrieved by an administrator (via SSH). The value is then entered into the GUI on the primary node.

Validation

The **Scan Node** function confirms the secret value. Incorrect values result in an error, preventing misconfiguration.

Recommended Access Order (RAO) Functionality

The RAO feature is a new capability introduced with LTO-9 tapes and compatible hardware, designed to improve the efficiency of reading files from tape. When supported by the tape library, RAO operates automatically in the background. This functionality is controlled by the TSM system parameter **RAO_ENABLED**, which is set to **true** by default. RAO is available on all operating system versions that support StorNext 7.2.4 and above.

> **Note:** If the system does not support RAO, this setting is ignored, and leaving it enabled has no impact on system behavior.

The only configuration required for this feature is the TSM system parameter **RAO_ENABLED**, which can be set to either **true** or **false**. On systems with RAO-enabled hardware, setting **RAO_ENABLED** to **true** allows StorNext to utilize the tape drive's RAO functionality for more efficient file retrieval. If set to **false**, the RAO feature is disabled and StorNext will retrieve files using the standard method. To configure the system parameter, modify the sysparm override file located at **/usr/adic/TSM/config/fs_sysparm_override** and add an entry that is **RAO_ENABLED=true;** or **RAO_ENABLED=false;**. For more information see Tools > Storage Manager > System Parameters.

> **Note:** If data compression is on, then RAO does not work as well and may give less enhancement than if the data was not compressed. There is no risk of data loss or errors if this happens. The perfomance degredation is only on the RAO feature, so the system performs as well as a system that does not have RAO enabled.

Upgrade Restrictions

Beginning with StorNext version 7.2.4, you cannot upgrade Xcellis Gen 1, Xcellis Foundation, and Artico (R630) systems that have a Mellanox ConnectX-3 expansion card installed. Quantum no longer provides kernel patches, as the Mellanox ConnectX-3 expansion card is no longer supported by Mellanox.

> **Note:** If your system does **not** have a Mellanox ConnectX-3 expansion card installed, you can upgrade to StorNext version 7.2.4 (or later).

Hardware Support

Newly Supported on Xcellis Gen3 Systems

- Emulex 3700 HBA (32 Gb FC)
- Mellanox CX-6 CDAT (100 GbE)

Library and Drive Support

Supported

- i7 tape library with RAO for LTO9

Not Supported

- LTO10 drives (pending future StorNext update)

OS Compatibility

Customer Supplied Hardware Support

- RHEL/Rocky/Alma/Oracle: 8.10, 9.3

Client OS Support

- RHEL/Rocky/Alma/Oracle: 8.10, 9.3–9.5
- Windows 11 LTSC
- Ubuntu 24.04, 24.10
- macOS 15.1

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

Multi-Factor Authentication (MFA)

- Purpose
  - Adds an extra layer of log in security for StorNext UUI users.
- Configuration
  - Enabled per user by an administrator.
  - Supports both email-based and app-based (e.g., Google Authenticator) MFA methods.
- Applicability
  - Can be applied to both local and Active Directory users.

- User Experience
  - Once configured, users are prompted for a verification code during log in.
  - MFA setup is a one-time process per device.

UUI Failover and Disaster Recovery

- Previous Limitation
  - UUI was only hosted on the secondary node of an MDC pair, causing downtime if the secondary node failed.
- New Capability in StorNext 7.2.4
  - UUI is now installed on both the primary and the secondary node.
  - Supports failover to the operational node in the event of a standby node failure.
- Best Practice
  - Use a virtual IP (vIP) for seamless client access during failover.

For additional information, see Perform a UUI Failover/Disaster Recovery.

Additional Improvements

- Numerous bug fixes and minor UI enhancements included in this release.

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 Firmware Upgrades

If your system is running StorNext 7.0.1.1 (or earlier) and you use the StorNext GUI to upgrade your firmware, do the following to upgrade to StorNext 7.0.2 (or later).

⚠️ **Caution:** Due to a known file size limitation (10 GB), if you attempt to upgrade your firmware using the StorNext GUI, the upload of the second firmware file fails with the following error:

```
File upload failed. The request was rejected because its size exceeds allowed
range.
```

Do the following to workaround this issue:

1. Modify the **com.icesoft.faces.uploadMaxFileSize** parameter in the **web.xml** file (located in **/usr/adic/tomcat/webapps/ROOT/WEB-INF/web.xml**) to a value of **10737418240**.

   Example

   ```
   <!--  Max. file size to upload (10 GB)   -->
   <context-param>
   <param-name>com.icesoft.faces.uploadMaxFileSize</param-name>
   <param-value>10737418240</param-value>
   </context-param>
   ```

2. Save your changes to the **web.xml** file.

3. Open a root UNIX shell window on your appliance, and run the following command to restart the StorNext GUI:

   > **Note:** Wait a few minutes before you try to access the StorNext GUI, and then retry the command if it fails.

   ```
   # service stornext_web restart
   ```

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 MDC based clients as a part of the StorNext software upgrade process.

> **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 7.2.4 only supports FlexSync 3.1.0 (or later). If you want to install StorNext 7.2.4 on your system, or upgrade your system to StorNext 7.2.4 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 destination, the target destination, and the managing appliance.

> ⚠️ **Caution:** Mixed versions of FlexSync daemons is 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 of the hosts or systems using Flexsync. A newer version of the Flexsync daemon cannot communicate with an older version within a configuration, or on another host or system.

Manual Installation Required: Xcellis Workflow Extender (XWE) Unified Connector Package Not Automatically Installed in StorNext 7.2.4

In StorNext 7.2.4, the XWE Unified Connector package is **not installed automatically** and must be installed manually via the UUI. Follow the steps below **after StorNext 7.2.4 is installed and configured on the XWE**:

Installation Steps

1. **Register the XWE with the UUI.** On the XWE, run the following command, replacing the placeholder with the UUI's hostname or IP address:

   ```
   /usr/cvfs/bin/snrest gateway register [hostname-or-ip-address-of-UUI]
   ```

2. **Access the UUI.** Log in to the UUI and navigate to the **Clients** page.

3. **Download the Unified Connector Package.** Locate the XWE in the **Clients** list. Under **Actions**, click the **Download** link to retrieve the Unified Connector package.

4. **Install the Package on the XWE.** Follow the on-screen instructions to complete the installation using the downloaded package.

Compatibility and Support

The StorNext 7.2.4 Compatibility Guide provides the basic compatibility for StorNext 7.2.4, 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 StorNext 7.2.4 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 7.2.4

ID	Description
SNXT-795	frequent lines like crond[216059]: pam_sss(crond:session): Request to sssd failed. Connection refused found in /var/log/secure
SNXT-948	--norequiredmedia description is missleading in fsimport man page
SNXT-1006	sntier relies on mountpoint causing issues when filesystem is stopped /removed but mounted
SNXT-1044	fsaddclass and fsmodclass drive pool name are not checked against the actual drive pools names and naming rules are inconsistent between fs and vs
SNXT-1161	'error 37' during mdarchive restore when 'quotas' is enabled
SNXT-1168	man page in sntier should mention --file to prevent directory affinitty move
SNXT-1234	Health Check (Media) is reporting 'Not enough LTO media' alerts despite there are unused blank tapes assigned to the policy.
SNXT-1408	UUI: Support running UUI on node 1 of MDC HA pair
SNXT-1478	Unprivileged users are not allowed to login after the Leapp upgrade completes and the rest of the platform upgrade is running

ID	Description
SNXT-1489	The grub settings intremap=no_x2apic_optout and nox2apic are not getting set on software RAID based XWD and XWE systems on upgrades to 7.2.0
SNXT-1491	XWEGen1 upgrade failed in ExtraRPMs with missing dependencies
SNXT-1530	Misconfigured snfs_rest_config.json file caused the fsmpm to fail to start without a good error output pointing to the cause.
SNXT-1536	Update the preupgrade check for the free space in /var for OS conversion upgrades from 10GB to 16GB
SNXT-1561	License: Invalid product code 0xB70(=DAE)
SNXT-1583	Consider having configurable use-at-your-own-risk revoke timeout
SNXT-1621	Fresh install of XWD or XWE results in missing /usr/cvfs/config/deviceparams file and system defaults for the I/O scheduler and nr_requests
SNXT-1656	'Balance' allocation strategy does not seem to work with default 'Inode Stripe Width' size
SNXT-1659	Add support for LTO RAO capability
SNXT-1670	XWE upgrade from 6.4.1 through to 7.2.0 fails the MellanoxUpdate checkpoint on 7.2.0 due to missing the kernel-modules-extra RPM
SNXT-1701	Add a new utility script to be used on Gen2 SWRAID based Xcellis systems to map the current boot drives to their physical slot numbers
SNXT-1758	fsmedcopy can write the EOD to the beginning of the tape if the source can't be read.
SNXT-1784	ActiveScale Cold Storage team would like REST API to trigger Audit
SNXT-1799	tracking bug for rdar://140352101 (acfs: heap overflow in dmfs_read_reply)
SNXT-1833	After upgrade to StorNext 7.2.0 the Notification (sl_noti_email_monitor) will not stay started
SNXT-1838	Need instrumented FSM code to resolve NFL issue
SNXT-1846	USBE container log grows to fill/scratch
SNXT-1939	ACLs on NFSv4: the users with read-only access should not be able to modify settings using nfs4_setacl command
SNXT-1957	xdi crashes with SEGFAULT in SN7.2.0 with RH8 when using the webconsole

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 7.2.4 Compatibility Guide.

Journal Size Guidelines

StorNext 7.2.4 requires a minimum journal size of 4 MB (Quantum recommends you configure your journal size to 64 MB). If your file system has a journal size less than the 4 MB minimum, you must resize your journal size before you upgrade to StorNext 7.2.4.

⚠️ **Caution:** If you upgrade and the journal size is less than the 4 MB minimum, then the FSM does not start and an error is logged.

Use the `cvupdatefs` utility (see the StorNext 7 Man Pages Reference Guide) or the GUI (see Edit a File System) to resize your journal size. When you resize your journal size, the new size must be 16 MB or greater. File systems with journals between 4 MB and 16 MB run with StorNext 7.2.4, but Quantum recommends you configure your journal size to 64 MB.

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, M660, M440, Pro Foundation, G300, and Artico are present, they all must be upgraded. No appliance can be left behind.

Considerations When Upgrading NFS Server Nodes to StorNext 7.2.4

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 7.2.4, special consideration must be made for Linux NFS servers.

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.

Compatibility Between StorNext and Other Products

The following sections provide information regarding compatibility between this release and StorNext components and features.

Appliance Controller

To view supported Appliance Controller software configurations, see the StorNext 7.2.4 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.

Lattus/ActiveScale (AXR, S3) or P100/X100

See the StorNext 7.2.4 Compatibility Guide in the StorNext Documentation Center for information about compatibility between Lattus/ActiveScale (AXR, S3) or P100/X100, and StorNext 7.2.4.

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.

For information about compatibility between PFR and StorNext 7.2.4, see the *StorNext Partial File Retrieval Compatibility Guide* in the StorNext Documentation Center.

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, refer to the StorNext Web Services Guide in the StorNext Documentation Center.

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 7.2.4, refer to the StorNext 7.2.4 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 7.2.4 Compatibility Guide in the StorNext Documentation Center.

For all other components and features, see the StorNext 7.2.4 Compatibility Guide in the StorNext Documentation Center.

General Considerations

This section provides information about items to consider for StorNext 7.2.4.

Checksum Performance Considerations

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

Current StorNext metadata controller and Mover 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 27 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:

ℹ️ **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	Change Request Number	Service Request 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.

(Description/Workaround, continued)

ⓘ **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.

Operating System	Change Request Number	Service Request Number	Description/Workaround
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.
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.

Operating System	Change Request Number	Service Request Number	Description/Workaround
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.

macOS	75819	n/a	An Xsan client cannot mount a StorNext File System volume when a cluster number 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
			
			If your MDC **fsnameservers** file includes a cluster number, remove the cluster number (**@_cluster_xx**) from the **mysan.configprofile** file **before** you copy the file to the Xsan client.
			
			For example, change:
			
			`10.65.181.158@_cluster0`
			
			to
			
			`10.65.181.158`
			
			See Mount the StorNext File System on Xsan 5.0 (or later) for additional information.

Operating System	Change Request Number	Service Request 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
			After each File System Pooling service restart, reconfigure the SMTP email settings using the **sntier config** command. This ensures that the correct password is applied and email notifications function properly.

StorNext Storage Manager Known Issues

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

Operating System	Change Request Number	Service Request 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
```

Operating System	Change Request Number	Service Request Number	Description/Workaround

			3. Verify the **FS_T10K_BLOCK_FACTOR** system parameter contains the new value:
			
			```# showsysparm FS_T10K_BLOCK_FACTOR\nFS_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\nrw,auto_dma_write_length=16m,auto_\ndma_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	Change Request Number	Service Request 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**: Currently, a workaround does not exist for this known issue.

Operating System	Change Request Number	Service Request 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	Change Request Number	Service Request 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
```

StorNext GUI Known Issues

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

Operating System	Change Request Number	Service Request 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	Change Request Number	Service Request 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	Change Request Number	Service Request 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	Change Request Number	Service Request 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	Change Request Number	Service Request 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.

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	Change Request Number	Service Request 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:
			
			<pre>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</pre>

Operating System	Change Request Number	Service Request 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.

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	Change Request Number	Service Request 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	Change Request Number	Service Request 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```

Operating System	Change Request Number	Service Request Number	Description/Workaround
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	Change Request Number	Service Request Number	Description/Workaround

All	SNXT-1443	n/a	⚠️ **Caution: Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.**
			
			The operating system conversion portion of the upgrade to StorNext 7.2.0 might fail if non-UTF8 files or directories exist in the **/** directory.
			
			Workaround:
			
			Before you initiate the upgrade on your system, create and execute the python script below on each of the nodes in the system.
			
			1. Create a file labeled, **checkNonUtf.py** in the **/tmp** directory with the following contents:
			
			`#!/usr/bin/python3`
			`import os`
			`for path in os.listdir('/'):`
			`        try:`
			`                path.encode('utf-8')`
			`        except UnicodeEncodeError:`
			`                # print path, or attempt to rename file`
			`                print('Non-UTF8 file or directory found on "/"')`
			`                exit(1)`
			`print('No non-UTF8 files or directories found on "/"')`
			
			2. Run the following command on each of the nodes in the system to execute the script:
			
			`python3 checkNonUtf.py`
			
			3. If the script reports **Non-UTF8 file or directory found on /**, then move these files or directories out of the **/** directory or rename them to be UTF8 named files.

Operating System	Change Request Number	Service Request Number	Description/Workaround
			Note: If these files/directories are required on the system, move/rename these back after the upgrade.
All	SNXT-1473	n/a	⚠ **Caution: Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.** After the upgrade to StorNext 7.2.0, attempts to use the email notification option to the **sntier** command fail due to the **/bin/mail** application not being installed on the system. **Workaround:** If your require this feature, you can download and install the Rocky 8 version of the mailx RPM on the Xcellis nodes from the following URL: https://dl.rockylinux.org/vault/rocky/8.8/BaseOS/x86_64/os/Packages/m/mailx-12.5-29.el8.x86_64.rpm

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	SNXT-1478	n/a	⚠️ **Caution:** **Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.**
			Unprivileged users are not allowed to log in during the upgrade process until the platform upgrade to StorNext 7.2.0 completes.
			Workaround:
			Before you initiate the upgrade on your system, do the following:
			1. Run the following commands as a **root** user on the node:

```
cp /usr/lib/tmpfiles.d/systemd-nologin.conf \
/etc/tmpfiles.d/
```

2. After the upgrade is complete, remove the file in **/etc/tmpfiles.d/**:

```
rm /etc/tmpfiles.d/systemd-nologin.conf
```

ℹ️ **Note:** Note: The upgrade to StorNext 7.2.0 might require up to three hours to complete on an Xcellis node. After the initial part of the upgrade which performs the operating system conversion, you might want to connect to the node from an SSN terminal as user **stornext** and monitor the upgrade logs for the remaining portion of the platform upgrade. Performing this workaround before the upgrade is initiated allows this.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	SNXT-1491	n/a	⚠️ **Caution: Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.** If your environment includes an Xcellis Workflow Extender that originally had StorNext version 6.1.1 or earlier installed, the upgrade will fail. **Workaround:** Quantum recommends you contact Quantum Support (see Contacting Quantum Support on page 49) for a software patch before you upgrade to StorNext version 7.2.0.
All	SNXT-1495	n/a	⚠️ **Caution: Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.** On 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. **Workaround:** Before you initiate the upgrade on your system, set the **nr_requests** value in **/usr/cvfs/config/deviceparams** to a value of **256**. The value should persist after the upgrade is complete. **Example** `nr_requests=256`

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	SNXT-1536	n/a	⚠️ **Caution:** **Before you upgrade to StorNext version 7.2.x, Quantum strongly recommends you review this known issue and applicable workaround.** 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. **Workaround:** Before you initiate the upgrade, verify the **/var** file system contains at least 16 GB of free space on each node.

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	Change Request Number	Service Request 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
```

ℹ **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
```

Quantum UUI Known Issues

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

Operating System	Change Request Number	Service Request 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	Change Request Number	Service Request 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	Change Request Number	Service Request 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	Change Request Number	Service Request Number	Description/Workaround
All	CON-11592	n/a	In the UUI (**Clients > Software Downloads**), the **Software Downloads** dialog displays a DDM client binary file that is not supported for Rocky 9, Alma 9, and Oracle 9 clients in this release of StorNext. **Example** ```snfs_client_plus_mover_RedHat9_x86_64.bin``` **Workaround**: To workaround this issue, ignore the DDM client binary file that appears in the UUI and do not attempt to download the binary file. The DDM client is not supported for Rocky 9, Alma 9, and Oracle 9 clients in this release of StorNext.
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.

Operating System	Change Request Number	Service Request Number	Description/Workaround
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	Change Request Number	Service Request 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
```

ℹ **Note:** If USBE requests return a 404 error, contact Quantum Support (see Contacting Quantum Support on page 49).

Operating System	Change Request Number	Service Request 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	Change Request Number	Service Request Number	Description/Workaround
All	CON-12640	n/a	After upgrading to StorNext 7.2.4, 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	Change Request Number	Service Request 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	Change Request Number	Service Request 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: ```\nThe following applications should be\nclosed before continuing the install:\nSearch\nStart\nNew notification\n``` **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	Change Request Number	Service Request Number	Description/Workaround
All	SNXT-2031	n/a	After upgrading to StorNext 7.2.4, 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
```

Contacting Quantum Support

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

- Chatbot below
- Open a Service Case below
- Use MyQuantum Service Delivery Platform on the next page
- Use Cloud Based Analytics (CBA) on the next page
- Escalate a Service Case on page 51
- Contact Quantum Sales on page 51

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/.

ℹ️ **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.
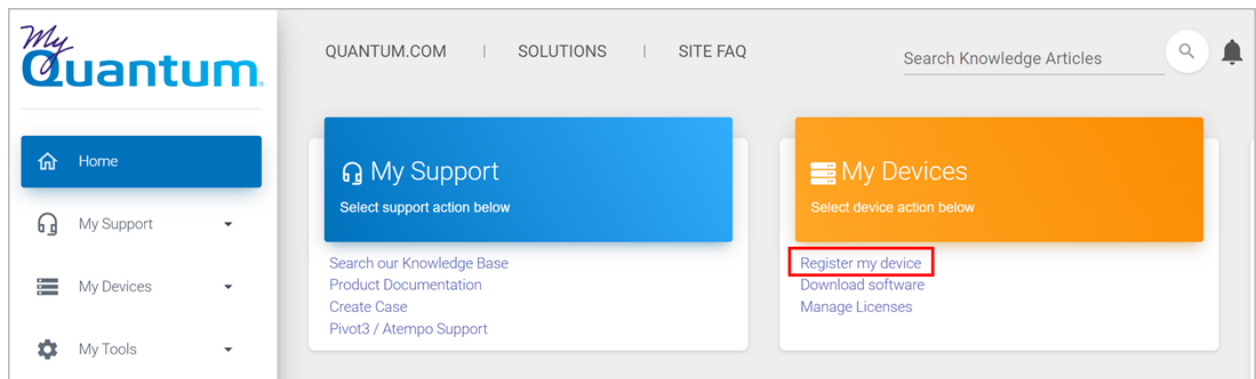
- Visit the Service & Support page.

  > **ⓘ 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

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