

## StorNext 7.2.0 Release Notes

**Original Product/Software Release Date**      May, 2024

### Contents

What's New in StorNext 7.2.0 .....	2
Compatibility and Support .....	7
Fixed Issues and Enhancements Addressed in StorNext 7.2.0 .....	8
Supported StorNext Upgrade Paths and Upgrade Considerations .....	11
Compatibility Between StorNext and Other Products .....	12
General Considerations .....	13
Upgrading Appliances .....	14
Appliance Release Notes .....	14
Known Issues .....	14
Contacting Quantum Support .....	37

---

© 2024 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. Myriad, ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Quantum, The Quantum Logo, QXS, Scalar, StorNext, SuperLoader, Vision, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Quantum specifications are subject to change.

---

# What's New in StorNext 7.2.0

## Purpose of this Release

The StorNext 7.2.0 release provides new features and enhancements, and also provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 7.2.0 on page 8](#).

This release also includes a significant upgrade to the underlying operating system of the Quantum appliances; specifically, the operating system is upgraded to Rocky 8.7 from CentOS 7.9.

---

**i Note:** Upgrades are only allowed from StorNext version 7.1.0 (or later) to StorNext version 7.2.0.

## Information about Pre-Upgrade Checks

Beginning with StorNext 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 manually 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 manually 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 7.1 (or later), you **must** upgrade the StorNext Unified Connector on non-MDC based clients.

---

**i 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 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, 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:

**i 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
```

## New Operating Systems Supported Effective with StorNext 7.2.0

- Debian version 12 (clients only)
- Apple macOS version 14 (Sonoma)
- RedHat, Alma, Rocky, Oracle version 8.9 (clients only)
- RedHat, Alma, Rocky, Oracle version 9.3 (clients only)
- Ubuntu 22.10 (clients only)
- Ubuntu 23.04 (clients only)
- Windows 11 (clients only)

## Operating Systems No Longer Supported Effective with StorNext 7.2.0

- RedHat, Rocky, Alma, CentOS 7.7 and earlier
- RedHat, Rocky, Alma, Oracle 8.0 – 8.4
- Windows Server 2008
- Windows Server 2012
- Ubuntu 20.04 and earlier
- Debian 9 and earlier
- SLES 11 and earlier

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

- First User Wizard
  - Removed **Password checked** success message that appears after **Step1** succeeds
  - Removed **Link Existing MyStorNext Account** tab
- Relocated **API GW** and **Users** menu items to 3 dots menu
- Changed **Update Available Notification** for **Updates** button to a badge instead of changing icon color
- Replaced **myStorNext portal cloud** button with **myQuantum Portal** button
- About
  - Removed browsers version from supported browsers list. Just browsers names are mentioned now
  - Changed the name of the section from **Unified User Interface Browser Support Versions** to **Unified User Interface Supported Browsers**
  - Replaced **myStorNext Portal URL** label and link with **MyQuantum Portal URL** label and link
- Login Screen
  - Removed **Select the link at the left and accept the SSL certificate as a possible solution** and **You may need to accept the SSL certificates at the links below first:** messages
- Nodes
  - Node details
    - Removed **CBA ID** field
    - Removed Eye button if password is stored in db for H/F Series nodes (removed an ability to see

password that was previously entered). "\*\*\*\*\*" is showed when clicking on Edit password button now in this case.

- StorNext
  - Changed menu style to appear as breadcrumb
  - StorNext Access
    - Changed **create/edit rule filesystem** field to be dropdown of valid filesystems instead of textfield
  - H/F - series
    - Added **Test** button to test e-mail settings
- Services
  - Added **Stripe Groups - Usage Dashboard**
  - Removed **Inode Usage Dashboard**
  - NAS
    - Added Support localsam auth method
  - Tiering
    - Added warning that tiering service should be restarted for the changes to take effect when saving Configuration if there were changes in Jobs or System sections
- Clients
  - Added Indicator That File System Has Unmounted
  - Added **Client Details** dialog
  - Added sorting by name ASC for file systems in File Systems column
  - Software Downloads
    - Updated supported versions for all OS types
    - Added download packages for Windows Server
- Metrics
  - Performance
    - Changed the scale for User Data Average Read+Write Latency Per I/O graph. The Y scale adjusted to use seconds, microseconds and milliseconds
    - Changed error message if could not talk to client from **Client not found: <client id>** to **Client not found**
    - Changed error message when client is no longer exist from **Client is no longer exists** to **Client no longer exists**
  - User Dashboards
    - Included an information that table, chart or graph do not have data to show in dashboard request response to avoid redundant rendering of objects

- Utilization
  - Added **OBJECTS OR SEGMENTS MOVED** and **OBJECTS OR FILES STORED** charts
  - Removed the **Storage Manager Activity**, **OBJECTS OR SEGMENTS MOVED** and **OBJECTS OR FILES STORED** widgets
- Logs
  - Added **Last 2 Days** and **Last Week** options to **Time Period** filter
  - Added logging of user activity (audit logs)
  - Changed Time Occurred column format. Milliseconds part has been removed
  - Logs Settings
    - Changed items in **Select Period** list. Spelling mistakes have been removed, overall text style have been changed
    - Changed label from **Age logs older than** to **Delete logs older than**
- Settings
  - Added **Register with Remote API-Gateway** section
  - Added **Upload of SSL Certificate** section
  - Removed **MyStorNext Portal** section
- Users
  - Changed behavior when clicking on parent role. Now list of subroles automatically opens
  - Changes in user roles management
- E-mail notifications
  - Removed StorNext Unified UI reference from emails. Quantum Unified UI is mentioned instead

## Features No Longer Supported Effective with StorNext 7.2.0

- Dynamic Application Environment (DAE); DAE-supported systems:
  - aiWARE for Xcellis Embedded
  - Xcellis Foundation
  - Xcellis Workflow Director (Gen2/R630)
  - Artico (R630)

## Information about FlexSync Support

StorNext 7.2.0 only supports FlexSync 3.1.0 (or later). If you want to install StorNext 7.2.0 on your system, or upgrade your system to StorNext 7.2.0 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.

---

## Compatibility and Support

The [StorNext 7.2.0 Compatibility Guide](#) provides the basic compatibility for StorNext 7.2.0, 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 7.2.0 Compatibility Guide](#).

- **Upgrade Paths:** Provides information on what upgrades to this release are supported.
- **Appliance Support:** Provides information on what StorNext appliances are supported with this release or are compatible with it.
- **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 7.2.0

ID	Description
SNXT-122	fslocate -k file_key fails on some files in mdarchive query
SNXT-139	Improve fsmodclass & fsexpcopy man pages to include info about rebuild candidate list
SNXT-179	Persistent LOG_INFO message in query_ldap() routine produces a large number of entries in /var/log/messages
SNXT-193	sgadd: when --template used for stripe group with exclusive affinity, sgadd issues an error message and exits: *Error*: You must specify data, metadata, or both
SNXT-194	cvadmin -H remote_host -e \"disks refresh\" do not trigger on remote_host but on local_host
SNXT-195	cvadmin -H remote_host -e \"disks\" reports wrongly \"Disks known to FSMPM on host \"local_host\" instead of \"remote_host\"
SNXT-200	CVFS Client hits \"VOP error on cookie <> cvfs error \"IO error\" (0x3) device <> after sgoffload
SNXT-210	TSM recreate tape: misleading error messages , e.g. Some files could not be recreated from alternate media due to segment size mismatches with source media
SNXT-317	Dynamic Library Pooling ( DLP ) doesnt cope well when one library of the DLP Pool runs out of available media
SNXT-335	Storage Manager Restore Procedures: missing crucial tips and warning for pitfall
SNXT-377	fsretrieve can error when object are pulled back from archive tier to cool tier
SNXT-390	A tiering job or policy should default to only affect files
SNXT-446	Fsmprm has sigabort due to accept failure in libevent code
SNXT-457	zombie fsprobe processes pile up on MDCs due to file descriptor leak in parse_json_cli()
SNXT-475	Bug 78307 - modified files that are removed before being re-stored cannot be recovered
SNXT-493	fsmedcopy recreate fails indicating there are no files to copy when there are.
SNXT-536	snprobe -a can stall in environments with large numbers of DLC connections
SNXT-581	sncmpare causes false positive with SDISK and Tape Copy will suggest to remove ALL_COPIES made on All Entries
SNXT-594	fsmedcopy reconciliation to help with - Deleting file during store causes database to show endtime as active.



ID	Description
SNXT-602	cvupdatefs() should not run 'sn_dmap -c' on LUN's that have been upgraded from "Short Labelled"
SNXT-620	XWE Gateway Agent is sending invalid file system metrics for MDC file system when affinity/tiering is enabled
SNXT-643	Fsm Assert Panic - ASSERT failed "!set    (rptr->iref_flags & IREF_RETURN_ATTRS) == 0
SNXT-661	CCTL logs flooding the nssdbg
SNXT-670	snaccess test functions incorrectly states subdirectory access is allowed
SNXT-678	Bug 73290 - File specific xrep error should not fail entire mover request seen on a XWD running SN7.0.3
SNXT-684	(CVE-2023-1945) - firefox-91.13.0-1.el7.centos - Unexpected data returned from the Safe Browsing API could have led to memory corruption and a potentially exploitable crash.
SNXT-685	add mechanism to disable hardcoded mover obj store retry logic
SNXT-688	(CVE-2022-37434) zlib: heap-based buffer over-read and overflow in inflate() in inflate.c via a large gzip header extra field
SNXT-693	(CVE-2023-0767) NSS Tools nss: Arbitrary memory write via PKCS
SNXT-694	escaped chars can lead to >255 name file, silently preventing store, failing policy involving subsequent candidates
SNXT-705	7.1.1 upgrade rendered Mellanox Infiniband devices undetectable.
SNXT-714	Compatibility Guide - Support for IBM TS4300 + IBM LTO 9 HH Fibre channel drives
SNXT-731	snstatd appears to leak memory
SNXT-734	/opt/DXi/scripts/collect-local: line 557: /usr/bin/pg_dump: No such file or directory
SNXT-740	transcribe operation to LTFS tape varies media to suspect if the destination media reach EOT/EOM
SNXT-766	cvfsck -E "FS" fails and coredump if "FS" has vacant SG
SNXT-768	fsexport -b cannot copy contents from LTO to LTOW (WORM) Media
SNXT-774	Without "exec" and "suid" XML tags in nss_ctl.xml, the values are set to false by default
SNXT-775	Man page update cvfsck -A only showing results in read/write mode
SNXT-819	ConvertAllRsvdExtents() triggers: PANIC: "BUG: unable to handle kernel NULL pointer dereference at 0000000000000004"

ID	Description
SNXT-878	Bug 76850 - FSM selecting a Remote NDIS Compatible Device causing connection issues (Clients interlocking because registered with same internal IP)
SNXT-881	Schema update failure when upgrading from SN pre-5.2.0 to post 6.4.0 in 7.1 code
SNXT-891	Documentation: Update the text for "Online" library action to reflect that it updates the device path and the mailbox slot count
SNXT-912	Documentation: Need to document when dbupgrade is run and what changes are being made to avoid unexpected downtimes or delays during Firmware upgrades
SNXT-937	sntier - Make sizes human readable and clarify what is meant by size in sntier command outputs. It is incredibly difficult to keep translating numbers like 47398723478732348 into terabytes or how many "seconds old" something is
SNXT-946	tiering - not enough info provided when a job failed due to estimate bigger than free space in pool
SNXT-956	Tracking Bug: Radar 119052712 Xsan clients fail to mount at boot
SNXT-1005	CVFS Client is leaking kmalloc-96 objects
SNXT-1019	Segmentation Fault in sn_dirent_key_init dir_btree.c:79
SNXT-1062	CVFS PANIC: _cvRelLock: unlocking a non-locked lock
SNXT-1116	/usr/adic/PSE/bin/query_tape_libraries is missing suspect count in output
CON-10602	MDC node with tiering recently disabled still shows the 'Pools' feature, which is confusing to customers
CON-10604	Unified Connector - Alma 8.7 and Rocky 8.7 are not supported by UC - minimum is 9
CON-10619	StorNext 7.1 + UUI/USB/E/API GW 7.1.1: 500 error on Alerts page when with any action
CON-10689	MacOS Unified Connector Package installer is blocked because it is not from an identified developer
CON-10701	Metrics page is not displaying graphs and statistics due to exceeding qustats data length
CON-10737	Adding second separate MDC node or HA pair pointing to same FS name server can lead to unhealthy MDC nodes status.
CON-10791	USUI – Portswigger scan flags /about/api RESPONSE for sensitive version and OS info
CON-10804	USUI - Nodes - Health status is changed incorrectly from times to times for metadata controller nodes

ID	Description
CON-10902	Linking an existing mystornext portal account to the UUI during user-register fails due to CORS policy
CON-10988	USUI - Support of large integer value(Petabytes) in license report generated and downloaded from the UUI
CON-11015	email test fails if "Use Authentication" is enabled
CON-11051	Unified Connector Installation Challenges on Windows
CON-11082	USUI-UC On newer MacOS can fail to get the hostname if TCPv4/v6 dual stack is active
CON-11168	Enhancement request - UUI - display details about the required password complexity
CON-11172	Health status of H-Series nodes become 'indeterminate' multiple times each day.

---

# 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 7.2.0 Compatibility Guide](#).

## Journal Size Guidelines

StorNext 7.2.0 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 7.2.0.


**⚠ 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 7.2.0, 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. One appliance cannot be "left behind".

## Considerations When Upgrading NFS Server Nodes to StorNext 7.2.0

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 7.2.0, 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 7.2.0 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 7.2.0 Compatibility Guide](#) in the [StorNext Documentation Center](#) for information about compatibility between Lattus/ActiveScale (AXR, S3) or P100/X100, and StorNext 7.2.0.

---

**i 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.

---

**i 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 7.2.0, 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 that enables multiple Mac computers to concurrently access hundreds of terabytes of content on Xserve RAID or Promise RAID storage over high-speed Fibre Channel which allows you to share data faster and consolidate projects. 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 7.2.0, refer to the [StorNext 7.2.0 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 7.2.0 Compatibility Guide](#) in the [StorNext Documentation Center](#).

For all other components and features, see the [StorNext 7.2.0 Compatibility Guide](#) in the [StorNext Documentation Center](#).

---

# General Considerations

This section provides information about items to consider for StorNext 7.2.0.

## Checksum Performance Considerations

---

**i 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 Replication, Deduplication, or Object Storage license, see Change Request [Known Issues below](#) in the [StorNext Installation, Replication, HA, and Other Known Issues on page 28](#) section before you upgrade.

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* ([www.quantum.com/ApplianceInfoHub](http://www.quantum.com/ApplianceInfoHub)).

---

## Appliance Release Notes

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

- [Xcellis Workflow DirectorXcellis Workflow Director](#)
- [Xcellis Workflow ExtenderXcellis Workflow Extender](#)
- [Xcellis FoundationXcellis Foundation](#)
- [ArticoArtico](#)

---

## Known Issues

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

**i 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	<p>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.</p> <p><b>i Note:</b> This problem can also occur with other programs that set EOF beyond the end of data.</p> <p>This problem does not occur if Windows Explorer encounters an error while writing the file; Windows Explorer will delete the partially written file.</p> <p><b>Workaround:</b></p> <p>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_allloc=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.</p> <p>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_allloc option, see the sparse option in mount_cvfs(8) for a description of how it changes StorNext behavior.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	67363	n/a	<p>StorNext 5.4.0.x incorrectly allowed the <b>Unix ID Mapping</b> type to be set to <b>none</b> when the <b>Security Model</b> is set to <b>acl</b>. As a result, file systems fail to start when the <b>Unix ID Mapping</b> type is set to <b>none</b> when the <b>Security Model</b> is set to <b>acl</b>.</p> <p>Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used.</p> <p><b>Workaround:</b></p> <p>To prevent this issue, set the <b>Unix ID Mapping</b> to either <b>winbind</b> or <b>algorithmic</b> for any file system where the <b>Security Model</b> is set to <b>acl</b>. You can make the adjustment before or after upgrading.</p>
All	75140	n/a	<p>Exporting an SNFS file system on Ubuntu releases 16.04.2 or later is not supported.</p> <p><b>Workaround</b></p> <p>There is currently no workaround for this issue. If you experience this issue, contact <a href="#">Quantum Technical Support</a>.</p>
All	75633	n/a	<p>A StorNext NAS client cannot rename a file if the file has the read-only attribute set. This problem only affects StorNext NAS clients.</p> <p><b>Workaround</b></p> <p>A StorNext NAS client must remove the read-only attribute before it can rename the file.</p>



Operating System	Change Request Number	Service Request Number	Description/Workaround
All	78082	575600	<p>If you set the parameter <b>audit=true</b> in the file system configuration file, the result causes all I/O activity, including reads, to be recorded in the <b>mdarchive</b>. This can greatly increase the number of updates applied to the <b>mdarchive</b> which, in turn, increases the amount of <b>mdarchive</b> compaction activity by the FSM.</p> <p>There is one instance where this increased compaction activity caused the <b>mdarchive</b> to grow beyond the capacity of the HA shared file system.</p> <p><b>Workaround</b></p> <p>To work around this issue, Quantum recommends you increase the parameter <b>metadataArchiveCache</b> from the default 2 GB to at least 4 GB, but preferably larger when setting audit to <b>true</b> in the configuration file.</p>
macOS	66948	322824, 336945	<p>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.</p> <p><b>Workaround</b></p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
macOS	75819	n/a	<p>An Xsan client cannot mount a StorNext File System volume when a cluster number is included in the local <b>fsnameservers</b> file. An unexpected EOF reading reply error is displayed.</p> <pre># xsanctl mount snfs3 xsanctl: unexpected EOF reading reply</pre> <p>If you encounter the EOF error on your Xsan client, do the following workaround to prevent the issue.</p> <p><b>Workaround</b></p> <p>If your MDC <b>fsnameservers</b> file includes a cluster number, remove the cluster number (<b>@_cluster_xx</b>) from the <b>mysan.configprofile</b> file <b>before</b> you copy the file to the Xsan client.</p> <p>For example, change:</p> <pre>10.65.181.158@_cluster0</pre> <p>to</p> <pre>10.65.181.158</pre> <p>See <a href="#">Mount the StorNext File System on Xsan 5.0 (or later)</a> for additional information.</p>

## 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	<p>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.</p> <p><b>Workaround:</b></p> <p>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:</p> <ul style="list-style-type: none"> <li><b>i Note:</b> This workaround might also help improve the performance of the faster LTO drives by updating the <b>FS_LTO_BLOCK_FACTOR</b> system parameter.</li> <li><b>i Note:</b> Changes to <b>FS_XXX_BLOCK_FACTOR</b> only affects tapes formatted after the change.</li> </ul> <ol style="list-style-type: none"> <li>1. Change the <b>FS_T10K_BLOCK_FACTOR</b> system parameter from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre style="background-color: #f0f0f0; padding: 5px;">FS_T10K_BLOCK_FACTOR=32;</pre> <ul style="list-style-type: none"> <li><b>i Note:</b> 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 <b>FS_T10K_BLOCK_FACTOR</b> system parameter must not be set to a value that exceeds 32.</li> </ul> </li> <li>2. Restart Storage Manager to ensure the change in <b>Step 1</b> goes into effect: <pre style="background-color: #f0f0f0; padding: 5px;"># tsmstop # tsmstart</pre> </li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>3. Verify the <b>FS_T10K_BLOCK_FACTOR</b> system parameter contains the new value:</p> <pre data-bbox="818 436 1458 541"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <p>4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</p> <p>5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows:</p> <pre data-bbox="818 814 1458 953">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_ dma_read_length=16m 0 0</pre> <p>6. Unmount and re-mount your file systems.</p> <p>7. Use new T10K media to store a copy of the file from the disk.</p> <p><b>i Note: Step 7</b> 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 <b>fsformat</b> was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	46693	n/a	<p>Executing the command <b>snbackup -s</b> while a full or partial backup is running might result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p>This is due to the <b>snbackup -s</b> process reading the <b>backup.1f</b> status file while the backup process is updating it.</p> <p><b>Workaround:</b></p> <p>Ignore the message; to clear-up the process, re-execute the command <b>snbackup -s</b> (provided that the backup is not writing to the <b>backup.1f</b> status file while <b>snbackup -s</b> is trying to read it again).</p>
All	47833	n/a	<p>When copying files between media using the CLI command <b>fsmecopy</b>, 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.</p> <p><b>i Note:</b> This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p><b>Workaround:</b></p> <p>Currently, a workaround does not exist for this known issue.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in <b>/usr/adic/TSM/logs/tac</b> which contains the text:</p> <div data-bbox="750 470 1458 779" style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <pre>Process fs_moverd on &lt;host&gt; 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.</pre> </div> <p><b>Workaround:</b></p> <ol style="list-style-type: none"> <li>1. Increase the database connection timeout value by adding the following line to <b>/usr/adic/mysql/my.cnf</b> under the section labeled <b>[mysqld] connect_timeout=240</b>.</li> <li>2. Cycle the Storage Manager in order to pick up the updated timeout value.</li> </ol>
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an <b>fsmedread</b> operation of a partial tape block from a full tape block, the operation can fail with <b>errno=12</b>.</p> <p><b>i Note:</b> This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures.</p> <p><b>Workaround:</b></p> <p>To correct this, perform an <b>fsmedread</b> operation without the IBM APFO driver.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	72993	452722	<p>When DDM is enabled for non-primary systems, kernel error messages are logged for <b>reservation conflict</b> 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.</p> <p>This issue causes a benign message in the <code>/var/log/messages</code> file for <b>reservation conflict</b>. 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.</p> <p><b>i Note:</b> This issue can result in a large amount of log messages on any machine running DDMs.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, you can filter and drop the messages so they are no longer logged, as follows.</p> <p><b>i Note:</b> Reservation conflicts that are a problem also produce RAS alerts, so you can also drop these errors.</p> <p>Create the following rules on all DDM clients, based on the OS, to drop the messages from <b>rsyslog</b>:</p> <p><b>For RedHat 6</b></p> <pre># echo ':msg, contains, "reservation conflict" ~' &gt; /etc/rsyslog.d/ignore-reservation-conflict.conf  # service rsyslog restart</pre> <p><b>For RedHat 7</b></p> <pre># echo 'if \$programname == "kernel" and \$msg contains "reservation conflict" then stop' &gt; /etc/rsyslog.d/ignore-reservation-conflict.conf  # systemctl restart rsyslog.service</pre>

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



Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69958	373823	<p>There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation.</p> <p>Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space.</p> <p>The other way to add space to a file system is to create a new stripe group and add this to the file system.</p> <p>When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs <b>cvupdatefs</b>. If <b>cvupdatefs</b> takes more than 5 minutes to complete, the GUI kills the <b>cvupdatefs</b> 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.</p> <p><b>Workaround:</b></p> <p>To address this issue, you can do one of two things:</p> <ul style="list-style-type: none"> <li>• Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system.</li> <li>• If stripe group expansion is deemed necessary, use the <b>cvupdatefs</b> CLI directly instead of through the GUI. This method does not have any time limitations.</li> </ul> <p>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:</p> <ol style="list-style-type: none"> <li>1. Restore the previous version of the configuration file. This can be found in the following directory: <pre style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;">/usr/cvfs/data/&lt;fs&gt;/config_history</pre> </li> <li>2. Run the <b>cvfsck</b> utility to verify and potentially correct the metadata.</li> <li>3. Start the file system.</li> </ol>

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 <b>GOOGLES3</b> 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	<p>The status and capacity monitors for file systems, libraries, storage disks, and tape drives does not display data on the StorNext GUI <a href="#">home page</a>. This issue is caused as a result of missing font RPM packages on your system.</p> <p><b>Workaround:</b></p> <p>To workaround the issue, Quantum recommends you do the following:</p> <ol style="list-style-type: none"> <li>1. Install the <b>java-1.8.0-openjdk</b> package on your system: <div data-bbox="818 716 1458 789" data-label="Code-Block"> <pre>yum install java-1.8.0-openjdk</pre> </div> </li> <li>2. Restart the Apache Tomcat service: <div data-bbox="818 911 1458 984" data-label="Code-Block"> <pre>systemctl restart stornext_web</pre> </div> </li> </ol>
All	80600	n/a	<p>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.</p> <p><b>Workaround:</b></p> <p>To workaround the issue, Quantum recommends you do the following:</p> <ol style="list-style-type: none"> <li>1. Log out of your StorNext user interface browser session(s).</li> <li>2. Close all browser tabs and windows.</li> <li>3. Open a new browser window and access the StorNext user interface (see <a href="#">Access the StorNext User Interface (new UI introduced in StorNext 7.0.1)</a> or <a href="#">Access the StorNext Software GUI</a>).</li> </ol>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> <li>1. Shut down the Safari browser window(s).</li> <li>2. Restart the Safari browser, and then retry the operation.</li> <li>3. Uncheck the <b>Verify SMTP Server Connectivity</b> box, and then retry the operation.</li> <li>4. Set <b>Authentication</b> to <b>NONE</b>, and then retry the operation.</li> <li>5. Disable the Safari <b>User names and passwords</b> AutoFill under <b>Safari &gt; Preferences &gt; AutoFill</b>, and then retry operation.</li> </ol>
All	SNXT-1329	n/a	<p>When you use the StorNext GUI to configure the <b>Timeserver</b> or <b>Pool</b> (located within the <b>Network Time Protocol (NTP)</b> section on the <b>Configuration &gt; System &gt; Date &amp; Time</b> page), and enter multiple comma separated IP addresses in the <b>Timeserver Pool</b> field, and then click <b>Test NTP</b>, the NTP test hangs and prevents you from accessing the GUI.</p> <p><b>Example of Timeserver Pool Entry</b></p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p><b>10.65.162.1,10.65.162.2</b></p> </div> <p><b>Workaround:</b></p> <p>Do the following to workaround this issue:</p> <ol style="list-style-type: none"> <li>1. Restart the StorNext GUI.</li> <li>2. Do <b>not</b> enter multiple IP addresses in the <b>Timeserver Pool</b> field. You can only enter one IP address.</li> </ol>

## StorNext Installation, Replication, HA, and Other Known Issues

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	68849	n/a	<p>After an appliance firmware upgrade, you might be unable to use previously functioning tape devices because the <code>lin_tape</code> device driver was automatically unloaded during the upgrade.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, rebuild the <code>lin_tape</code> device driver as shown in the following example:</p> <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	<p>Quantum installs a self signed certificate (valid for 365 days) in <b>/usr/cvfs/config/certs/</b> only when installing a snfs-common RPM and if a preexisting certificate file does not exist.</p> <p>If you use monitoring software (for example, Zabbix), an expired self signed certificate is flagged after it expires.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, do the following to update an expired self signed certificate:</p> <ol style="list-style-type: none"> <li>1. Stop CVFS on the affected client: <pre># service cvfs stop</pre> </li> <li>2. Update the certificate: <pre># 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</pre> </li> <li>3. Start CVFS on the affected client. <pre># service cvfs start</pre> </li> </ol>

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

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

reboot of the system while an rpm activity occurred. StorNext, Appliance 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.

**i Note:** Your system might require an updated BIOS from the manufacturer to fix and eliminate the warnings.



## 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	<p>The <b>CLIENTS</b> page of the UUI displays duplicate clients.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, do the following.</p> <ol style="list-style-type: none"> <li>1. On the <b>APIGW</b> page of the UUI, click <b>Deregister Endpoint</b> for the endpoint of the duplicate client.</li> <li>2. Restart the <b>Unified Connector</b> on the client that contains the duplicate entries.</li> </ol>
All	CON-11386	n/a	<p>If your system is running StorNext 7.1.1 or earlier, the label <b>N/A</b> is displayed in the <b>Reserved</b> column when you display strip group information in the UUI.</p> <p><b>Workaround:</b></p> <p>There is currently no workaround for this issue. If you experience this issue, contact <a href="#">Quantum Technical Support</a>.</p>
All	CON-11396	n/a	<p>When you upgrade the UUI software, the upgrade process might appear to be complete but the <b>Quantum Unified UI Software Upgrade</b> page does not update due to an <b>update-service</b> request that continuously fails with a <b>404</b> status code.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, execute the following command as <b>root</b> on the node where the UUI is running:</p> <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;">/opt/quantum/uui-common/bin/uui upgrade remote -y</pre>
All	CON-11487	n/a	<p>When you navigate to the <b>Metrics &gt; Performance and Metrics &gt; Utilization</b> UUI page, you might experience a long delay for the page to display properly.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, refresh your browser page.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11504, CON-11505, CON-11506	n/a	<p>If you navigate to the <b>Metrics &gt; Utilization</b> page, and attempt to download a <b>Historical Usage</b> graph, the image file for the graph might not display properly.</p> <p><b>Workaround:</b></p> <p>Do the following to workaround this issue:</p> <ol style="list-style-type: none"> <li>1. Click the <b>Print chart</b> icon in the UUI.</li> <li>2. Save the document as a PDF.</li> </ol>
All	CON-11508	n/a	<p>When you navigate to the <b>Metrics &gt; Performance and Metrics &gt; Utilization</b> 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.</p> <p><b>Example</b></p> <div data-bbox="760 873 1458 978" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Trouble loading Capacity data. Request aborted.</p> </div> <p><b>Workaround:</b></p> <p>To workaround this issue, close your browser window and then reopen a new browser window.</p>
All	CON-11541, CON-11542	n/a	<p>If you navigate to the <b>Metrics &gt; Performance</b> 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.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, download the data graph in a different format, such as a PDF document, or an SVG vector image.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11572	n/a	<p>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.</p> <p><b>Workaround:</b></p> <p>To workaround this issue, execute the following command:</p> <pre data-bbox="764 590 1458 695">/opt/quantum/usui/docker/usui-init-cron.sh --force</pre>
All	CON-11587	n/a	<p>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 <b>CLIENTS</b> to verify the issue; you might notice a client registration for both the primary node and the secondary node.</p> <p><b>Workaround:</b></p> <p>If you are missing the client registration for the secondary node, Quantum recommends you perform either of the following options to workaround this issue.</p> <p><b>Workaround Option 1</b></p> <p>As the <b>root</b> user, SSH to the secondary node and execute the following command:</p> <pre data-bbox="764 1297 1458 1371">systemctl restart unified-connector</pre> <p><b>Workaround Option 2</b></p> <p>Reboot the secondary node.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11592	n/a	<p>In the UI (<b>Clients &gt; Software Downloads</b>), the <b>Software Downloads</b> 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.</p> <p><b>Example</b></p> <pre data-bbox="764 558 1458 663">snfs_client_plus_mover_RedHat9_x86_64.bin</pre> <p><b>Workaround:</b></p> <p>To workaround this issue, ignore the DDM client binary file that appears in the UI 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.</p>
All	CON-11597	n/a	<p>In the UI (<b>Metrics &gt; Dashboards</b>), the error below might appear in your widget when you perform the following steps:</p> <ol data-bbox="776 1003 1425 1339" style="list-style-type: none"> <li>1. On the top navigation menu, click <b>Metrics</b>, and then click <b>Dashboards</b> on the left navigation menu.</li> <li>2. Create a new dashboard.</li> <li>3. Add a widget for a file system to the dashboard you created in <b>Step 2</b>.</li> <li>4. Stop and start the file system.</li> <li>5. On the top navigation menu, click <b>Metrics</b>, and then click <b>Dashboards</b> on the left navigation menu.</li> </ol> <p><b>Example of error</b></p> <pre data-bbox="764 1436 1458 1541">File System &lt;name of file system&gt; is no longer exist.</pre> <p><b>Workaround:</b></p> <p>To workaround this issue, remove the existing widget from dashboard and add a new widget.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	CON-11618	n/a	<p>If you install the UUI software on a customer-supplied system, the following error message appears in the <code>./uui-install.sh</code> script output when the installation completes.</p> <p><b>Example of error</b></p> <pre>Error: failed to remove gateway config: gateway config empty</pre> <p><b>Workaround:</b> You can ignore the error message.</p>

---


## Contacting Quantum Support

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

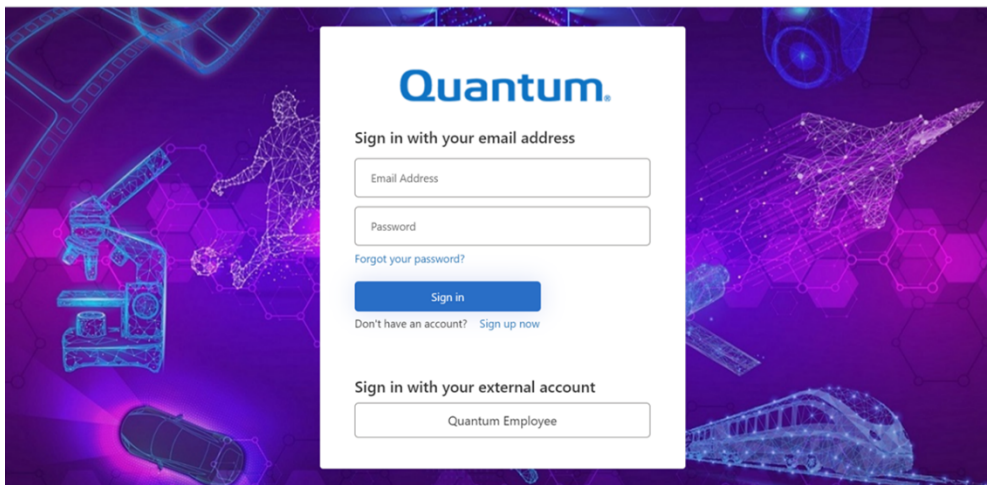
- [Submit a Ticket \(Service Request\) below](#)
- [Use MyQuantum Service Delivery Platform on the next page](#)
- [Use Cloud Based Analytics \(CBA\) on page 39](#)
- [Escalate a Case on page 39](#)
- [Contact Quantum Sales on page 39](#)

### Submit a Ticket (Service Request)

If you need to submit a ticket or speak to Quantum technical support, go to the Support page at <https://www.quantum.com/en/service-support/>



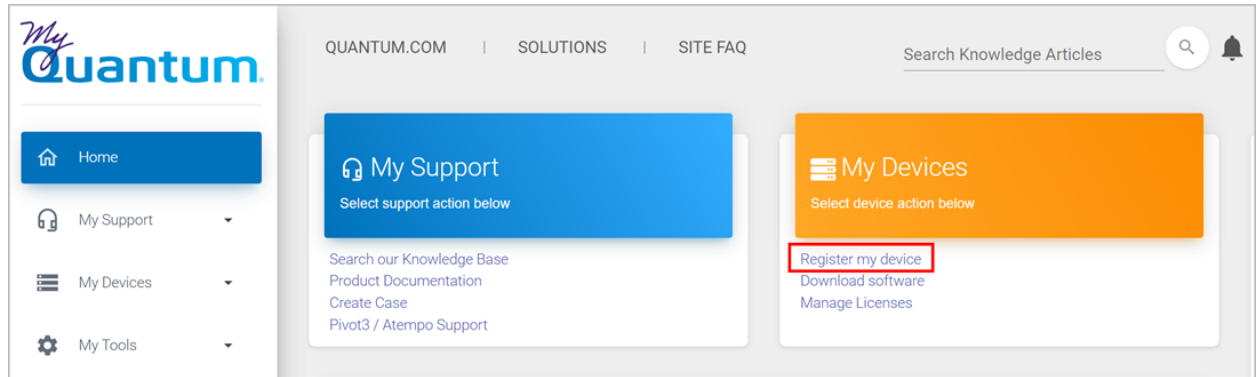
To start the process with Quantum Technical Support, click **Submit a Ticket**. From here, sign in to the MyQuantum Service Delivery Platform or create an account. For more information, refer to the [Use MyQuantum Service Delivery Platform below](#) section below.



## Use MyQuantum Service Delivery Platform

MyQuantum is a single portal for everything Quantum. You can view assets, open support 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.

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

## 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.

## Escalate a Case

To escalate a 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<sup>®</sup>

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<sup>®</sup> Index. For more information visit [www.quantum.com](http://www.quantum.com).

[www.quantum.com](http://www.quantum.com) | 800-677-6268