

## StorNext 7.1.0 Release Notes

### Contents

What's New in StorNext 7.1.0 .....	2
Compatibility and Support .....	6
Fixed Issues and Enhancements Addressed in StorNext 7.1.0 .....	7
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 .....	31

---

© 2022 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. 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 7.1.0

## Purpose of this Release

The StorNext 7.1.0 release provides new features and enhancements listed in the section [New Features and Enhancements in StorNext 7.1.0 on the next page](#), and also provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 7.1.0 on page 7](#).

## 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 (4 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
```

## StorNext Connect No Longer Supported Effective with StorNext 7.1.0

Beginning with StorNext 7.1, the StorNext Connect product is not supported; you can access the StorNext User Interface to manage your system (see [Access the StorNext User Interface](#)). If you purchase a new system, the StorNext Connect product is not installed on your system.

## Linear Tape-Open (LTO) Formats No Longer Supported Effective with StorNext 7.1.0

- LTO-3
- LTO-4

## Operating Systems No Longer Supported Effective with StorNext 7.1.0

- Debian 8
- SuSE 11
- Ubuntu 14
- Ubuntu 16
- Windows 7
- Windows Server 2008

## New Features and Enhancements in StorNext 7.1.0

- Added a new system parameter labeled **ALREADY\_STORED\_FILE\_IS\_ERROR** to display an error when you attempt to store a file that is already stored. To configure the system parameter, edit the respective **ALREADY\_STORED\_FILE\_IS\_ERROR** system parameter located in the file `/usr/adic/TSM/config/fs_sysparm_override`. The default setting for the **ALREADY\_STORED\_FILE\_IS\_ERROR** system parameter is **false**.

- Updated the **IMP\_OBJMETA\_NUM\_THREADS** system parameter from **5** to **15**. To configure the system parameter, edit the respective **IMP\_OBJMETA\_NUM\_THREADS** system parameter located in the file **/usr/adic/TSM/config/fs\_sysparm\_override**.
- Added a new system parameter labeled **DRIVE\_DELAY\_DISMOUNT** to replace the hard-coded internal value, and updated the value from **300 seconds** to **1200 seconds**. To configure the system parameter, edit the respective **DRIVE\_DELAY\_DISMOUNT** system parameter located in the file **/usr/adic/TSM/config/fs\_sysparm\_override**.
- Added a new system parameter labeled **BACKUP\_USE\_PIGZ** to control whether a backup attempts to use **pigz** instead of **gzip** for compression, to improve performance. To configure the system parameter, edit the respective **BACKUP\_USE\_PIGZ** system parameter located in the file **/usr/adic/TSM/config/fs\_sysparm\_override**. The default setting for the **BACKUP\_USE\_PIGZ** system parameter is **true**.

---

**i Note:** A file written by **pigz** can be read by **gzip** and vice versa.

- Added a new system parameter labeled **BACKUP\_EX\_MDA\_UNCOMPRESSED** to control whether a backup process extracts uncompressed copies of metadata archive databases and then separately compresses the databases versus extracting already-compressed copies of the metadata archive databases. To configure the system parameter, edit the respective **BACKUP\_EX\_MDA\_UNCOMPRESSED** system parameter located in the file **/usr/adic/TSM/config/fs\_sysparm\_override**. The default setting for the **BACKUP\_EX\_MDA\_UNCOMPRESSED** system parameter is **false**.
  - The advantage to enabling this parameter is that metadata archive updates are resumed more quickly.
  - The disadvantage to enabling this parameter is that the backup requires more time due to the additional disk I/O required.
- Added a new system parameter labeled **BACKUP\_NUM\_THREADS** to control how many threads are used when **pigz** is used for a backup. It has no effect when the system parameter **BACKUP\_USE\_PIGZ** is disabled (**false**). To configure the system parameter, edit the respective **BACKUP\_NUM\_THREADS** system parameter located in the file **/usr/adic/TSM/config/fs\_sysparm\_override**. The default number of threads for the **BACKUP\_NUM\_THREADS** system parameter is **6**.

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

- StorNext Connect Retirement
  - Unified Connector For Client Monitoring
  - Performance and Capacity/Utilization Monitoring plus Custom Dashboards
  - Client Discovery and Monitoring
  - Client Software Downloads
  - First Boot Wizard In Unified UI
- Node Support Updates

- F/H Series QBSP 2.1 and 2.2 Release Support
- QXS Next Generation Support
- StorNext Feature Management
  - FlexSync 2.3.0
  - StorNext Client Access (snaccess), QoS (QBM), System/Cache, and File System Feature/Option Configuration
- Infrastructure
  - GUI Users Active Directory Integration
  - StorNext 6 Legacy WebServices APIGW Integration (Storage Manager, RAS Alerts)
  - myStorNext SSL Proxy Integration
  - File System Web Services API Completion/Documentation Updates
  - Security and Supportability Updates (Open-Source Updates, Docker Network Customization, etc)
- Selected Customer Found Legacy UI Updates/Fixes

## Enhancement to the Client **qustat** Titles

StorNext 7.1.0 includes an enhancement to the client **qustat** titles to include the client **uuid**, in addition to the existing IP address. You might need to update any script that uses **qustat** to account for the inclusion of the client **uuid**. Prior to StorNext 7.1.0, the format of the **qustat** title for tables associated with **vops**, **tokens**, **dmig** and **url** uses the IP address to identify its clients.

For example, when a client removes a mount to a file system, the **qustat** table contents are printed in the file system **cvlog** with a title that uses an IP address:

```
Table 31: client.10.65.177.246.vops
```

Below is an example of the enhanced **qustat** table title that includes its **uuid**:

```
Table 31: client.10.65.177.246;56bc03be-7583-480d-a7c5-839393e73bc1.vops
```

---

**i Note:** A semicolon separates the IP address from the client **uuid**.

Beginning with StorNext 7.1.0, the following tables are enhanced for clients:

- **vops**
- **tokens**
- **dmig**

- **dmig\_events**
- **brl**

---

## Compatibility and Support

The [StorNext 7.1.0 Compatibility Guide](#) provides the basic compatibility for StorNext 7.1.0, 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 7.1.0 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 7.1.0

Operating System	Change Request Number	Service Request Number	Description
All	52706	3458772, 621720, 667294	Tape drive cleaning error message is misleading
All	61842	656967	sncompare fails on gigafiles system
All	73493	463603, 472099, 496925, 650999	ArcDisp segfault occurs upon looping on attempt to cancel same multimount request, upon a manual TSM start
All	74455	529097, 601617	new version of LTFS library (2.4.0 from 2.2.0)
All	75169	509013, 576274, 606402, 668080	fspolicy crashing w/ unhandled exception in mda4str_proc_cpy_num(), when processing a multi-segment file, all due to segment information missing in filecomp*
All	77392	542765	fsqueue does not show Azure uploads after failures cause a retry on other DDM.
All	77703	n/a	GUI: (branding enhancement) Only display the StorNext name/logo in header
All	77864	570245, 581673, 650671	vsunlock does not accept negative lockid argument
All	77865	552016	Add resolv.conf.save to snapshots
All	77957	575004	daily health check reports Software Error: an unhandled software error has occurred.
All	78049	577711, 562096, 679626, 680572, 680500	Cannot start Storage Manager due to sngateway_install_mysql_tables.pl script mysql error
All	78115	562485, 659262, 679331	fsmedcopy -R command reports success when not all files copied due to tape error
All	78130	577999	enhancement request: add GUI config / file system / edit check box to be able to toggle alloc=true false

Operating System	Change Request Number	Service Request Number	Description
All	78364	582440, 542365, 606008	fs_requestd dumping core on free during delete_list
All	78471	592787	MAX_INPTMPDIR_CLEANTIME readme documentation does not match code.
All	78574	598917	If a configuration file can't be read, we need a RAS.
All	78767	605987	RFE wsapi: allow to define a streams count for fsretrieve fsstore done via the API
All	78786	607276	TSM, while it was stopping, issued misleading "Mover availability error" Admin alert
All	78909	612288	TSM selected a tape from genscratch when many scratch tapes were available in the policy class
All	78914	611455	snbackup will not overwrite partial backup files if Max Inactive Versions set to 0
All	79146	616240	Modify restore_size.sh script to run large tables in separate transactions
All	79178	619726	A typo when running vsarchiveconfig -u made archive unusable
All	79198	623316, 619569	Database error should be more meaningful - Recreate of medium failed. Reason: Database error.
All	79640	639461	fsmedcopy failed with confusing errors "No usable file systems" and "Unable to find usable file system for deviceId"
All	79759	645191	dbcreate interactive mode enters infinite loop if non-numeric character is entered
All	79844	638962	FSM crash during snrecover left file system in a state where it could not start and cvfsck could not run
All	79898	653195, 669166	Qualys Item - EOL/Obsolete Software: Apache Log4j 1.X Detected
All	79950	620930	Consider buffer cache changes to improve compressed video playback



Operating System	Change Request Number	Service Request Number	Description
All	79983	648943	Invalid product key because of cvfsid changing when license.dat gets removed
All	80024	639810, 661605	Modify checkPolicyClassStore health check to verify if snbackup is running or not
All	80129	659778, 671412, 685432	snprobe runs very slowly in some environments
All	80151	657271	BART outputs error every run on Tencents hardware: Illegal ncpu = 144
All	80160	658520	Confusing errors "Unable to get label information" and "Unable to process label information"
All	80207	645927	FSM crash in alloc_btree_delblk() due to corrupted backing_snbt pointer
All	80214	649673	snhistory panics with "Failed to read tablet cluster"
All	80219	664,374,678,308	sgoffload is using 512 block writes to a 4096 block device
All	80222	653533	fsquota errors on file names containing apostrophe
All	80242	664518	TSM incorrectly encodes objectid while handling an AWS url redirection.
All	80244	664301	Rdar 92981986 SEED: Web: Monterey Xsan DLC via 10GbE fails on all hardware
All	80266	667605, 670515	Support for RHEL 8.5 and RHEL 8.6 Clients
All	80273	667927	Enhancement: Modify restore_size.sh script to drop ACTIVEFL table to save time
All	80285	669072, 680687, 683026	FSMPM Heap Corruption Panics since upgrade to 7.0.3
All	80287	SR671325	FSM PANIC after upgrade to 7.x due to bad/old sb_FsStatus flag
All	80304	672015	man pages: sntier pool --delete removes all pools from the file system mount point

Operating System	Change Request Number	Service Request Number	Description
All	80305	672015	sn-tier pool report shows all pools as exclusive if configured using the UUI
All	80308	671024	Media not marked suspect when drive hit E_DEST_POS_MS ending up in a loop
All	80377	675798	Data loss can occur when deferring writing of LTFS index data (Baidu specific issue)
All	80382	674761 681644	mdarchive quits processing apply pending bundles
All	80407	678423	man pages: (TSM recreate tape) clarification needed on -R and -v options
All	80411	680727	PS is migrating to an H4000 cvfsck runs clean on the old system 7+ million duplicate inodes on the new
All	80449	683041	man pages: update fsschedlock (with no option) to explicitly state what you can get
All	80557	649673	snhistory panics with "Failed to read tablet cluster" - follow up fix
Linux	43593	649407	vsping loops forever if rpc service is not running
Linux	54486	3503416, 343638, 427587, 581916	ArcDisp core dumps when cancelling a multimount command awaiting tape entry to library
Linux	76969	551378	Support for SLES 12 SP5
Linux	80335	675152	some files are not returned within api/metadb/v1/change_list, and don't show on the path array. However some seem to be accounted in the sum 'result'.
Linux	80338	675203	Support for RHEL 9.0 Clients
Linux	80415	650990	Remove RAS message issued for Quota update would have caused current size to be negative
macOS	79886	639456	Customer wants to run qbm_mio script with client being a Mac OS system but script doesn't work

# 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 7.1.0 Compatibility Guide](#).

## Journal Size Guidelines

StorNext 7.1.0 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 7.1.0.

---


 **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 7.1.0, 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 7.1.0

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 7.1.0, 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 7.1.0 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](#).

## ActiveScale or Lattus

See the [StorNext 7.1.0 Compatibility Guide](#) in the [StorNext Documentation Center](#) for information about compatibility between ActiveScale or Lattus, and StorNext 7.1.0.

---

**i Note:** ActiveScale and Lattus documentation is available online at <https://www.quantum.com/documentation>.

## 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 7.1.0, 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 7.1.0, refer to the [StorNext 7.1.0 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 7.1.0 Compatibility Guide](#) in the [StorNext Documentation Center](#).

For all other components and features, see the [StorNext 7.1.0 Compatibility Guide](#) in the [StorNext Documentation Center](#).

---

# General Considerations

This section provides information about items to consider for StorNext 7.1.0.

## 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 Director](#)[Xcellis Workflow Director](#)
- [Xcellis Workflow Extender](#)[Xcellis Workflow Extender](#)
- [Xcellis Foundation](#)[Xcellis Foundation](#)
- [Artico](#)[Artico](#)

---

## 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_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.</p> <p>For more information on this option, see the man page <code>mount_cvfs(8)</code>. The sparse option will introduce some overhead when using Windows Explorer. Before setting the <code>protect_alloc</code> option, see the sparse option in <code>mount_cvfs(8)</code> 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="760 472 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:</li> </ol> <pre>yum install java-1.8.0-openjdk</pre> <ol style="list-style-type: none"> <li>2. Restart the Apache Tomcat service:</li> </ol> <pre>systemctl restart stornext_web</pre>
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>

## 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>

---

# Contacting Quantum

## Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<https://www.quantum.com/aboutus/contactus/index.aspx>

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

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

## Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

[doc-comments@quantum.com](mailto:doc-comments@quantum.com)

The Quantum logo is rendered in a bold, blue, sans-serif font. The background of the page features a series of diagonal stripes in various shades of blue and purple, creating a dynamic, modern aesthetic.

# 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