

## StorNext 7.0.2 Release Notes

### Contents

What's New in StorNext 7.0.2 .....	2
Information About StorNext Firmware Upgrades .....	3
Information about LTO-9 .....	4
Information about the Appliance Controller/NAS Object Store Plugin .....	4
Compatibility and Support .....	4
Fixed Issues and Enhancements Addressed in StorNext 7.0.2 .....	5
Supported StorNext Upgrade Paths and Upgrade Considerations .....	9
Compatibility Between StorNext and Other Products .....	10
General Considerations .....	11
Upgrading Appliances .....	12
Appliance Release Notes .....	12
Known Issues .....	12
Contacting Quantum .....	29

---

© 2021 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 7.0.2

## Purpose of this Release

The StorNext 7.0.2 release provides new features and enhancements listed in the section [New Features and Enhancements in StorNext 7.0.2 below](#), and also provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 7.0.2 on page 5](#).

## New Features and Enhancements in StorNext 7.0.2

### Capacity Based Subscription Software only StorNext

Beginning with StorNext 7.0.2, you can purchase and install StorNext software on your own system.

- Same subscription license options as a StorNext 7 appliance.
- Linux only, MDCs must be running RHEL 7 update 9.
- All client types supported as per the [StorNext 7.0.2 Compatibility Guide](#).
- Includes StorNext Unified User Interface (UUI) GUI that you can deploy as a container on an MDC or on a separate system.
- Includes features such as Storage Manager, DDM, StorNext File System Pooling, and so on.
- Does not include Appliance Controller/NAS or Cloud-Based Analytics (CBA).

### Usability

- Major StorNext UUI GUI updates.
- Includes StorNext 7 Web Services API.

## Security and Client Updates

- Client support up to RHEL 8 update 4 (current), Debian 10.x and Ubuntu 20.04.2 LTS.

---

# Information About StorNext Firmware Upgrades

If you use the StorNext GUI to upgrade your firmware, do the following to upgrade your appliance to StorNext 7.0.2.

- ⚠ 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 **6442450944**.

### Example

```
<!-- Max. file size to upload (4 GB) -->
<context-param>
  <param-name>com.icesoft.faces.uploadMaxFileSize</param-name>
  <param-value>6442450944</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
```

---

# Information about LTO-9

**i Note:** StorNext does not support the LTFS media format with LTO-9.

## Cartridge Initialization

Cartridge initialization is a new feature performed on the first use of your LTO-9 media and requires additional time; the initial load performs a cartridge initialization and a cartridge characterization.

**i Note:** The cartridge initialization is important for Transverse Dimensional Stability (TDS) and Archiving Storage Mode.

The cartridge initialization process moves tape at a controlled speed and tension, end to end four (4) times, measuring lateral servo band to band dimensions for all four data bands and information is stored in the Cartridge CM, which is used as reference for Shingled Magnetic Recording including Archival Storage. The time required for the first load from the **LOAD** command to SCSI Status is anticipated to be a minimum of 12 minutes, although actual numbers might be up to 120 minutes depending on the environmental conditions of your cartridges and your system.

For additional information, see the [Quantum LTO-9 Media Calibration Frequently Asked Questions \(FAQ\)](#) document.

## Unload Delays

The time required for an unload depends on how the cartridge has been used during the mount (for example, the current position and how far from robot the tape has been pulled since the mount). If the tape has not moved, then the time is anticipated to be on the order of 27 seconds. However, if the worst-case scenario is met, then the maximum time is anticipated to be on the order of four (4) minutes, with an average of two (2) minutes.

---

# Information about the Appliance Controller/NAS Object Store Plugin

Beginning with Appliance Controller/NAS 3.1.0, the [Object Store Plugin](#) (which you purchased and installed separately) is no longer available. If you were previously using the plug-in, then contact Quantum Technical Support (see [Contacting Quantum on page 29](#)).

---

# Compatibility and Support

The [StorNext 7.0.2 Compatibility Guide](#) provides the basic compatibility for StorNext 7.0.2, 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 7.0.2 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 7.0.2

Operating System	Change Request Number	Service Request Number	Description
All	71190	394274, 506722, 550448, 580930, 615083, 609812	Raise an explicit RAS alert when a S3 media is set write protected
All	72022	482099	copymand cores during shutdown
All	76830	542765	Retry object storage downloads and part uploads on transient errors
All	78025	576602, 595342, 598265	FSM process requests a socket, it fails with address already in use.

Operating System	Change Request Number	Service Request Number	Description
All	78356	585335, 606467, 603227	Unable to add Azure Data Boxes into an existing Storage Destination
All	78357	n/a	Request for Debian 10 client support
All	78387	587802	fsmpm can crash on libxml free if multiple fsms are started
All	78398	586609	Running cvfsck -q can set the FSCK_NEEDED flag in the super block
All	78470	591234	cvadmin displays component sizes in units of 1024^N but uses TB/GB/MB size label
All	78489	591398	fs_moverd leaking sockets on multi-streamed I/O requests
All	78502	585487	fsimport fails if any filenames from fsmedscan can't be inserted into the fileinfo database table
All	78508	591254	Google objects left orphaned after parallel upload errors
All	78518	642584	Update quantum_disk_catalog to include F-Series and H-Series entries for Quantum Branded disk.
All	78520	588842	CVFS down_write()'s i_rwsem without holding the semaphore
All	78523	n/a	When Windows DLC clients run continuous I/O to a file system that has been remade and resized, mdc kernel panics upon mounting the new FS
All	78550	576870, 582980, 535952	Windows SN client ldap issues can delay or prevent FS mount operation
All	78551	591254	LCR Container: Google objects left orphaned after parallel upload errors
All	78555	597882	fsretrieve from AWS S3 bucket incorrectly reports error on the destination instead of the source
All	78557	593782, 599000	cnvt2ha failed with ERROR: Fsmlist HaShared mismatch. Got " expecting 'snsn_ha1'
All	78575	n/a	Allow LAN clients to override server_conn_count

Operating System	Change Request Number	Service Request Number	Description
All	78594	598991, 626788, 631664	Multiple threads in fsmppm_about_service() can cause core dump of fsmppm
All	78601	591208	Windows Explorer large file copy takes several seconds to report any progress when buffer cache >= 4GB
All	78608	596346	documentation: tiering: when a pool is called default stntier shows a strange behavior
All	78634	601547, 601380, 603772	Admin alerts triggered by failed retrieves create email spam storm
All	78636	601655	cvfsck -x should use quoting to avoid parsing issues with file names containing commas
All	78637	602110	syncha failed with no error causing hours of debugging on the Customer system
All	78655	601534	fs_moverd dump core on large fsfilecopy job
All	78671	601396, 606929, 608506	JQuery pre v3.5 Vulnerability CVE-2020-11022
All	78676	622718, 627647, 637791	AWS multipart initialization fails to return upload ID
All	78681	598531	Stores to GoogleS3 storage destination hang if multi-stream is configured
All	78684	603227	Tiering failure when trying to move files
All	78693	604182, 625224	FSM hitting ASSERT(amt_diff <= (uint64_t)4096) in RTQOS code
All	78697	605033, 613892, 614914, 621235, 621288, 624383	Directory listing reports I/O error and displays question marks instead of protection, ownership, size and time
All	78762	605987	fsretrieve use only 1 stream when policy configured for 8 streams, fsretrieve -S 8 work as expected
All	78772	576870, 582980, 535952	Workaround for bug 78550 to clean the registry
All	78781	603227	Tiering failure when trying to move files

Operating System	Change Request Number	Service Request Number	Description
All	78827	603227	Provide a message sn-tier job hasn't found any candidates
All	78831	604568	CvApi_ExtendedStatByName segv with f_name argument containing utf-8 characters far exceeding LINKINFO_NAME_MAX
All	78912	610058	Do not retrieve a file from Lattus or P100 if the FILECOMP "Objid" field is blank
All	78974	611674	"fsclean -D -a -R" does not only check file system where specified directory resides
All	78976	612690	Modify cvcp to log file read/write errors and continue
All	79021	615259	fsmpm can panic with ASSERT failed 0
All	79053	612947, 612900	Lengthy database shutdown during upgrade - Correct so that for 7.0.2 we check mysql version
All	79074	616811	Deadlock-like situation for NamedStream_remove() and Quota operations
All	79108	n/a	ActiveVault cannot function without a licensed capacity limit to work against
All	79145	596955, 602014, 611758, 621413, 638010	qevhttp subsystem: memory or fd exhaustion causing FSMPM Panics
All	79162	621280	After HA fail-over DDM hosts are disabled (connected to another StorNext system error)
All	79170	621280	fs_moverd: Add additional data to the fsloglevel -d output
All	79180	618708	Webservices job info call can report syntax errors for failed fsmecopy jobs
All	79258	613084	snquota cli -L report is reporting an error on directory with special char
All	79338	629680	snweb: java core in SNAdmin_DebugInit and tomcat down after voluntary failover
All	79349	631020	Password length limits must be the same across all forms (login/change)



Operating System	Change Request Number	Service Request Number	Description
All	79362	630355	wsar_agent not queuing more than 100 jobs to TSM when worker-job_* parameters set > 100
All	79490	635421	Object still exists in Google namespace after store operation was canceled
All	79502	618931	Create mdarchive check for invalid node id
All	79573	n/a	Multi-streaming configuration is ignored for files in archive storage
Linux	79497	636255	Bandwidth expansion of Stripe Group can take a very long time to complete
Mac OS	78517	637845	Mac OS Big Sur MDC without license can connect to FX clients

## 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 7.0.2 Compatibility Guide](#).

### Journal Size Guidelines

StorNext 7.0.2 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 7.0.2.


**⚠ 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 6 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 7.0.2, 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 7.0.2

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 7.0.2, 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 7.0.2 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 7.0.2 Compatibility Guide](#) in the [StorNext Documentation Center](#) for information about compatibility between ActiveScale or Lattus, and StorNext 7.0.2.

---

**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 7.0.2, 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 7.0.2, refer to the [StorNext 7.0.2 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 7.0.2 Compatibility Guide](#) in the [StorNext Documentation Center](#).

For all other components and features, see the [StorNext 7.0.2 Compatibility Guide](#) in the [StorNext Documentation Center](#).

---

# General Considerations

This section provides information about items to consider for StorNext 7.0.2.

## 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 26](#) 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 Extender](#)
- [Xcellis Foundation](#)
- [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_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>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; border: 1px solid #ccc;"> <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	78777	n/a	<p>If you use the StorNext GUI to upgrade your firmware, you cannot upgrade an appliance to StorNext 7.0.2 due to a known file size limitation.</p> <p>The upload of the second firmware file fails with the following error:</p> <div data-bbox="760 558 1458 695" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>File upload failed. The request was rejected because its size exceeds allowed range.</p> </div> <p><b>Workaround:</b> To workaround this issue, see <a href="#">Information About StorNext Firmware Upgrades on page 3</a>.</p>

---



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
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/serviceandsupport/get-help/index.aspx#contact-support>

## 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)



# 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