



StorNext 6.0.1 Release Notes

Release	6.0.1
Supported Product	StorNext 6.0.1
Date	October 2017

Contents

What's New in StorNext 6.0.1	2
StorNext Compatibility	3
Supported StorNext Upgrade Paths and Upgrade Considerations	4
Compatibility Between StorNext and Other Products	5
General Considerations	7
Upgrading Appliances	7
Known Issues	8
Contacting Quantum	21

© 2017 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. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexSpace, FlexSync, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, 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. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.

What's New in StorNext 6.0.1

Purpose of this Release

StorNext 6.0.1 is a full software release that includes improvements for any StorNext 6 based configuration.

New Features and Enhancements in StorNext 6.0.1

Notable enhancements delivered with StorNext 6.0.1 include support for [Xcellis systems](#) with Veritone® aiWARE™ — a hybrid on-premise and cloud version of Veritone's best-in-class, cloud-based [AI platform](#). This combination allows you to leverage the power of Veritone's cognitive analytics — along with top cognitive engines in areas such as face detection, object recognition and transcription — to extract new value from their on-premise video and audio content without having to move it to the cloud. As a result, if you have made significant investments in on-premise storage and/or have cost or security concerns about storing your content in the cloud, you will be able to take advantage of Veritone's cognitive capabilities and analytics in a StorNext-based solution.

This release also corrects an issue that affects case insensitive StorNext file systems exported by Linux Network File System (NFS). For complete details, see [Fixed Issues and Enhancements Addressed in StorNext 6.0.1 on the next page](#).

Compatibility and Support

The [StorNext 6 Compatibility Guide](#) provides the basic compatibility for StorNext 6.0.1. The *StorNext Compatibility Guide* includes 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 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

- **Upgrade Paths:** Provides information on what upgrades to this release are supported.
- **Appliance Support:** Provides information on what StorNext and Lattus 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 6.0.1

Operating System	Change Request Number	Service Request Number	Description
All	69374	358555	An issue was fixed that affects case insensitive StorNext file systems exported by Linux Network File System (NFS).

StorNext Compatibility

For information on StorNext compatibility with operating systems, kernel versions, hardware platforms, drives, libraries, StorNext Appliances, StorNext client interoperability, and other compatibility items, see the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

Note: SNAPI and Partial File Retrieval information is provided in separate documents.

Quantum Operating System Upgrade Support Policy

StorNext supports any security or functional bug update that applies to the current StorNext-supported Red Hat update level or SuSE Linux Service Patch. StorNext does not support upgrading to an update level or service patch beyond the currently supported levels shown in the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

StorNext and Linux Interoperability

Newer versions of the Linux `tail` command leverage the `inotify` mechanisms within Linux. The `inotify` mechanisms in Linux are not triggered by file updates coming from other StorNext nodes.

When using the `tail` command on files located in StorNext, Quantum recommends using the following option:

```
---disable-inotify
```

Recommended usage:

```
tail ---disable-inotify -f filename
```

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 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

Considerations for the StorNext File System Directories


On upgrades to StorNext 6.0.1, note that the attributes of many directories in the StorNext file system show much smaller sizes, even zero sizes, where these same directories showed non-zero sizes in previous releases of StorNext. This is expected behavior.

Journal Size Guidelines

The absolute minimum Journal Size in StorNext 6.0.1 is 4 MB. If a file system is configured with a Journal Size smaller than 4 MB, the Journal Size must be increased prior to upgrading. The recommended Journal Size is 64 MB. New file systems must have a Journal Size of 64 MB or larger.

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, etc.) are intended to be done to all systems present in a given deployment. For example, if Xcellis, M660, M440, Pro Foundation, Artico, and G300 are present, they all must be upgraded. One appliance cannot be "left behind".

Considerations When Upgrading NFS Server Nodes to StorNext 6.0.1

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.0.1, 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.

Database Schema Update During Upgrades

Database schema updates are applied to Storage Manager when upgrading from StorNext 4.3.x and StorNext 4.7.x to StorNext 6.0.1. The M660 appliance can achieve approximately one hour for every hundred million entries in the filecomp tables. Smaller appliances and software only configurations might take considerably longer depending on CPU speed and memory availability.

i Note: The database schema update conversion time from StorNext 4.7.x to StorNext 6.0.1 is significantly faster than that from StorNext 4.3.x to StorNext 6.0.1.

StorNext file systems are accessible while the database schema is being updated, but Storage Manager functionality (including stores and retrieves) will be offline.

Do NOT interrupt StorNext services while the database is being updated. Interrupting the database schema update could result in an inconsistent database, and might require assistance from Quantum Support to repair or restore the database.

Use the following commands to determine the number of filecomp entries on the StorNext primary node:

1. List the managed filesystems configured.

```
mysql -e "select Device_key, Path from tmdb.devdb_v;"
```

2. For each <Device_key> number listed, display a count of the number of entries in the corresponding filecomp table:

```
mysql -e "select count(*) from tmdb.filecomp<Device_key>;"
```

i Note: The query in **Step 2** might require a significant amount of time. Quantum recommends you execute the query before the day of an upgrade.

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 *Appliance Controller Compatibility Guide* available online at http://www.quantum.com/acc_cg.

Infiniband

StorNext 6.0.1 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

Lattus

Refer to the version of the [Lattus Release Notes](#) applicable to your system for information about compatibility between Lattus and StorNext 6.0.1.

Object Storage documentation is available online at <http://www.quantum.com/lattusdocs>.

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 6.0.1, see the *StorNext Partial File Retrieval Compatibility Guide* in the [StorNext 6 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 6 Web Services Guide](#) in the [StorNext 6 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 6.0.1, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

General Considerations

This section provides information about items to consider for StorNext 6.0.1.

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

For instructions on upgrading your firmware, refer to the current Release Notes for your particular appliance:

- For Xcellis Workflow Director, see <http://www.quantum.com/xcelliswfddocs>.
- For Artico, see <http://www.quantum.com/articodocs>.
- For M660, M440, M330 Metadata Appliance and Pro Foundation, see <http://www.quantum.com/snmdcdocs>.
- For G300 Gateway Appliance, see <http://www.quantum.com/sngatewaydocs>.

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	47501	3591086, 3687554, 3669192	<p>If you are upgrading to StorNext 6, the configuration parameter <code>bufferCacheSize</code> might need to be increased.</p> <p>If your file system has ever contained a large directory (millions of entries) and those entries were subsequently removed without also removing the directory, you should confirm that the StorNext 6 upgrade is making progress.</p> <p>Quantum recommends you perform the Workaround below if the following conditions occur.</p> <p>In the log file <code>/usr/cvfs/data/<file system>/log/cvlog</code>, if you see a line like the following:</p> <pre>(Info) Inode conversion will stop at IEL chunk <xxxxx></pre> <p>And you do not see messages like this (shortly thereafter):</p> <pre>(Info) Conversion completed on 100/1079 IEL chunks</pre> <pre>(Info) Conversion completed on 200/1079 IEL chunks [...]</pre> <p>Or, if you begin to see Conversion completed messages, but the messages stop appearing for an extended period of time and you do not also see a message like the following:</p> <pre>(Info) IEL scan for conversion complete</pre> <p>Workaround:</p> <p>Increase the configuration parameter <code>bufferCacheSize</code> by following the procedure outlined online in the StorNext Documentation Center (see the The Metadata Controller System topic, sub-section BufferCacheSize).</p>

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>i Note: 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>Workaround:</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>
All	67363	n/a	<p>StorNext 5.4.0.x incorrectly allowed the Unix ID Mapping type to be set to none when the Security Model is set to acl. As a result, file systems fail to start when the Unix ID Mapping type is set to none when the Security Model is set to acl.</p> <p>Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used.</p> <p>Workaround:</p> <p>To prevent this issue, set the Unix ID Mapping to either winbind or algorithmic for any file system where the Security Model is set to acl. You can make the adjustment before or after upgrading.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69221	346914, 333021	<p>The sngateway_cleanup crontab entry is missing from systems that were upgraded to StorNext 5.2 or StorNext 5.3. An issue was introduced in StorNext 5.2 that resulted in the removal of crontab entries for the tdlm user not created by PSE_cfg, such as sngateway_cleanup and custom entries.</p> <p>The problem was resolved in StorNext 5.4, such that existing entries were preserved, but it did not re-add the dropped entries for the sngateway_cleanup.pl entry. As a result, customers running with the sngateway metrics feature enabled, might experience some performance degradation and possible space issues on the HaShared file system due to a potential high volume of performance metrics.</p> <p>i Note: All HA systems that had upgraded to StorNext 5.2 or StorNext 5.3 are exposed.</p> <p>Workaround:</p> <p>To correct this, run the following on the primary node:</p> <pre># /usr/cvfs/install/sngateway_install_cron_entry.pl</pre> <p>To verify that the expected crontab entry is present, run the following:</p> <pre># crontab -l -u tdlm grep sngateway_cleanup</pre> <p>As a result, the following entry is expected:</p> <pre>15 1 * * * /usr/adic/gui/bin/cmdwrap -NO_END_OF_FILE /usr/cvfs/bin/sngateway_cleanup.pl</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69697	n/a	<p>In StorNext 6 (and later), file system names might be qualified with cluster information. For example, if the fsmcluster file contains the following:</p> <pre>default_cluster cluster1</pre> <p>Then, snfs1 and snfs1@cluster1 refer to the same file system. This syntax can be used on the Linux mount command and in the Linux /etc/fstab. The syntax fs@cluster is only required to mount file systems in other than your default cluster. It is optional for file systems in your default cluster.</p> <p>Issue:</p> <p>The StorNext GUI does not recognize the syntax <fs>@<cluster. When this syntax is present in /etc/fstab, several screens in the GUI will either fail or provide incomplete information. The following error or similar is displayed:</p> <pre>Error getting data for file system/s (snfs1): Filesystem Name or Mount Point is not provided.</pre> <p>Workaround:</p> <p>Do not use the syntax fs@cluster in /etc/fstab for file systems in your default cluster when you are using the StorNext GUI.</p> <p>Related Issue:</p> <p>Change Request 69572 describes a similar problem, that is, if the syntax fs@cluster is used in /etc/fstab for the HA shared file system, the start of the StorNext services will fail. The workaround is the same, that is, do not use the syntax fs@cluster to describe a file system in your default cluster.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69908, 69910	n/a	<p>Known issues exist that cause problems with several utilities, if you are mounting StorNext file systems with the same name from different clusters.</p> <p>Prior to StorNext 6, all StorNext file systems mounted by a given host were required to have unique names. For StorNext 6 (and later), file systems names can be qualified with cluster information, making them still unique, but allowing the same name to be used. For example, snfs1 in cluster1 and snfs1 in cluster2 represent different file systems with the same name.</p> <p>Change Request 69908:</p> <p>The snquota utility can operate on the wrong file system. For example, the following command could pick either the file system in cluster1 or cluster2, depending on the order in which they are mounted:</p> <pre data-bbox="623 863 1458 934">snquota -F snfs1 -L</pre> <p>Change Request 69910:</p> <p>Similar to Change Request 69908, the following stripe group management utilities could also pick the wrong file system on which to operate:</p> <ul data-bbox="623 1121 792 1255" style="list-style-type: none"> • sgoffload • sgdefrag • sgmanage <p>Workaround:</p> <p>To resolve both issues, make sure that only a single instance of a given file system name is mounted, if you are using any of the above utilities.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Windows	69366	n/a	<p>After installing the StorNext client package on a Windows system, StorNext does not start. If you try to start StorNext, the following error message appears:</p> <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <pre>Error 2: 'Error starting Windows Service 'cvfsfilter'</pre> </div> <p>You might encounter this problem when installing a StorNext 6.0 client-only package under the following circumstances:</p> <ul style="list-style-type: none"> • When you remove the StorNext file system before installing the StorNext client package. • When you upgrade a StorNext client package. <p>Workaround:</p> <p>To correct this, reboot the Windows system and reinstall StorNext.</p>
Mac OS	67871	n/a	<p>Mac OS releases 10.12 (and later) contain an issue in which the operating system crashes and reboots if a rename is performed on a file inside a managed directory.</p> <p>Note: The issue only occurs if rename tracking is disabled on the file system.</p> <p>Workaround</p> <p>To prevent this issue, enable rename tracking on managed file systems with Xsan clients.</p>
Mac OS	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>Workaround</p> <p>If you encounter the errors on an Xsan client, contact Apple.</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>Workaround:</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"> i Note: This workaround might also help improve the performance of the faster LTO drives like LTO-6, and LTO-7 by updating the FS_LTO_BLOCK_FACTORsysparm. i Note: Changes to FS_XXX_BLOCK_FACTOR only affects tapes formatted after the change. <ol style="list-style-type: none"> 1. Change the FS_T10K_BLOCK_FACTORsysparm from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <div data-bbox="678 1031 1458 1100" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>FS_T10K_BLOCK_FACTOR=32;</pre> </div> <ul style="list-style-type: none"> i Note: The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the FS_T10K_BLOCK_FACTORsysparm must not be set to a value that exceeds 32. 2. Restart Storage Manager to ensure the change in Step 1 goes into effect: <div data-bbox="678 1451 1458 1556" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre># tsmstop # tsmstart</pre> </div> 3. Verify the FS_T10K_BLOCK_FACTORsysparm contains the new value:

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre data-bbox="683 331 1458 436"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <ol style="list-style-type: none"> <li data-bbox="618 478 1458 548">4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients. <li data-bbox="618 569 1458 674">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: <pre data-bbox="683 709 1458 856">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> <li data-bbox="618 898 1149 926">6. Unmount and re-mount your file systems. <li data-bbox="618 947 1365 974">7. Use new T10K media to store a copy of the file from the disk. <ul style="list-style-type: none"> <li data-bbox="667 989 1458 1192">i Note: Step 7 is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which fsformat was run before the change will use the block factor in use at that time. This change will not impact those tapes.
All	46693	n/a	<p data-bbox="602 1220 1458 1360">Executing the command snbackup -s 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 data-bbox="602 1367 1458 1436">This is due to the snbackup -s process reading the backup.1f status file while the backup process is updating it.</p> <p data-bbox="602 1442 781 1470">Workaround:</p> <p data-bbox="602 1476 1458 1591">Ignore the message; to clear-up the process, re-execute the command snbackup -s (provided that the backup is not writing to the backup.1f status file while snbackup -s is trying to read it again).</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47833	n/a	<p>When copying files between media using the CLI command fsmedcopy, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p>Note: This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p>Workaround:</p> <p>Currently, a workaround does not exist for this known issue.</p>
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 /usr/adic/TSM/logs/tac which contains the text:</p> <pre>Process fs_moverd on <host> timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.</pre> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Increase the database connection timeout value by adding the following line to /usr/adic/mysql/my.cnf under the section labeled [mysqld] connect-timeout=240. 2. Cycle the Storage Manager in order to pick up the updated timeout value.
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an fsmedread operation of a partial tape block from a full tape block, the operation can fail with errno=12.</p> <p>Note: This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures.</p> <p>Workaround:</p> <p>To correct this, perform an fsmedread operation without the IBM APFO driver.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69889	371926	<p>With StorNext 6, as items are written to LTFS tapes or read from LTFS tapes, the internal driver will output debug information to the fsmpm.out file located in /usr/cvfs/debug/. These LTFS logs are not needed and should not be going to this file. This file does not roll nor is truncated and may result in the file system that the fsmpm.out file resides in filling up.</p> <p>Workaround:</p> <p>If you are reading and writing data from/to LTFS media, the following actions are required:</p> <ul style="list-style-type: none"> • Add cronjob entries, to filter out the LTFS log messages from the fsmpm.out file, on every machine that is running a data mover. This allows you to reduce size of the file, as it grows, and preventing the file system from filling up. <p>On both the primary and secondary MDCs, edit the crontab for the user tdlm so that each hour, the log file is edited to remove the extraneous LTFS lines. Perform the following procedure:</p> <pre>crontab -e -u tdlm</pre> <p>Add the following 2 lines that read:</p> <pre># This will filter out LTFS logs from the fsmpm.out file hourly 0 * * * * /usr/adic/gui/bin/cmdwrap -NO_END_OF_ FILE sed -i '/ LTFS/d' /usr/cvfs/debug/fsmpm.out</pre> <p>i Note: There is a single space between the / and LTFS in the line above.</p> <ul style="list-style-type: none"> • On any other machine that is running a data mover, edit the crontab for the user root so that each hour, the fsmpm.out file is edited to remove LTFS lines. Perform the following procedure: <pre>crontab -e</pre>

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

Add the following 2 lines that read:

```
# This will filter out LTFS logs from the
fspm.out file hourly
0 * * * * sed -i '/ LTFS/d'
/usr/cvfs/debug/fspm.out
```

Note: There is a single space between the / and **LTFS** in the line above.

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 autofs 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 snfs1 is mounted as /stornext/snfs1, then the MDC should not also have an autofs configuration that mounts it on the MDC in another location such as /space/snfs1. Doing so, causes the fsCheckAffinities and fsCheckTsmFilesystemConfig health checks to fail and generate RAS tickets.</p> <p>Additionally, this might cause the StorNext GUI to fail unexpectedly for certain operations.</p> <p>Workaround:</p> <p>There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.</p>

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>Workaround:</p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> 1. Shut down the Safari browser window(s). 2. Restart the Safari browser, and then retry the operation. 3. Uncheck the Verify SMTP Server Connectivity box, and then retry the operation. 4. Set Authentication to NONE, and then retry the operation. 5. Disable the Safari User names and passwords AutoFill under Safari > Preferences > AutoFill, and then retry operation.

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>Workaround:</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>

Contacting Quantum

Contacts

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

<http://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:

<http://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