



StorNext 6 Release Notes

Release	6
Supported Product	StorNext 6
Date	October 2017

Contents

What's New in StorNext 6	3
Purpose of this Release	3
New Features and Enhancements in StorNext 6	3
Fixed Issues and Enhancements Addressed in StorNext 6	8
StorNext Compatibility	17
Quantum Operating System Upgrade Support Policy	17
StorNext and Linux Interoperability	17
Supported StorNext Upgrade Paths and Upgrade Considerations	18
StorNext Software Upgrade Matrix	18

© 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.

Considerations for the StorNext File System Directories	18
Journal Size Guidelines	18
Distributed Data Mover (DDM) Guidelines	18
Considerations When Upgrading NFS Server Nodes to StorNext 6	18
Database Schema Update During Upgrades	19
Compatibility Between StorNext and Other Products	19
Appliance Controller	20
Infiniband	20
Lattus	20
Partial File Retrieval	20
StorNext Web Services	21
Apple Xsan	21
Supported Browsers	21
General Considerations	21
Checksum Performance Considerations	21
Upgrading Appliances	22
Known Issues	22
StorNext File System Known Issues	22
StorNext Storage Manager Known Issues	28
StorNext GUI Known Issues	33
StorNext Installation, Replication, HA, and Other Known Issues	34
Contacting Quantum	35

What's New in StorNext 6

Purpose of this Release

Unrelenting growth in the amount and importance of data has become commonplace. New, often unexpected technical and business challenges are impacting legacy storage solutions. Scaling upward in capacity alone is not enough, as the exponential growth of unstructured data, coupled with new ways of exploiting the value of data, have made efficient data management more important than ever before. Purpose-built for supporting today's most demanding storage requirements, StorNext 6 delivers sustained application-level performance and advanced data protection capabilities – while lowering management complexity.

StorNext 6 is an unparalleled combination of an extensive set of important enhancements with a pair of optionally licensed data protection and sharing features.

New Features and Enhancements in StorNext 6

Notable enhancements that are delivered as a core part of StorNext 6 range from storage Quality of Service (QoS) and file system auditing to offline file management. For example, storage QoS allows a storage administrator to throttle application I/O to ensure it is used by the most important applications. File system auditing, which supports StorNext version 6 and 5 clients' tracks and reports on changes made to metadata or the contents of files. Stripe group management can add, defragment or offload new stripe groups while file systems are online and available.

Offline file management prevents the inadvertent recall of truncated files. Integrated with Apple Finder and supported with other StorNext clients and Appliance Controller, this vital enhancement allows users to easily control file retrieval, store and truncation requests.

Maximizing the value and availability of secondary storage tiers has been improved. StorNext 6 supports that latest version of Open LTFS, and adds support for exporting and importing data from LTFS media. Improvements have been made that boost the performance and efficiency of StorNext Storage Manager, along with ongoing scalability improvements. The new copy expiration feature automates the expiration of file copies by media type – equally important is that copy expiration can be used with existing or newly created Storage Manager policies. Not to be overlooked, administrators can customize the ordering of media used for file retrieval requests to optimize the use of expanded archive tiers: whether cloud and object or legacy tape and disk.

In addition to these core improvements, a new optionally licensed feature is delivered with StorNext 6. StorNext FlexSync™ is a very fast and efficient way to create local or remote replicas of files and metadata. FlexSync is easy to configure and allows users to browse and restore their own files using standard tools.

StorNext 6 makes storage simpler, faster and more efficient.

See [Fixed Issues and Enhancements Addressed in StorNext 6 on page 8](#) for a complete list of the issues addressed by StorNext 6.

i Note: Advances in disk arrays have made the 32-bit Volume Table of Contents (VTOC) disk labels obsolete. As a result, StorNext 6 does not support VTOC labels. Modern disk arrays use Extensible Firmware Interface (EFI) disk labels capable of supporting physical disks that are larger than 2 TB in size. Prior to upgrading to StorNext 6, convert arrays that utilize VTOC labels to the EFI format.

To convert arrays that utilize VTOC labels to the EFI format, use the CLI command **cvlabel1**. For additional information about **cvlabel1**, refer to the **cvlabel1** man page in the **StorNext 6 Man Pages Reference Guide** available on the [PDFs Downloads page](#) in the [StorNext 6 Documentation Center](#).

i Note: StorNext 6 includes changes that affect the configurable file system security parameters. Unless the **unixPermBits** security model is used, support for StorNext file systems created with the configuration parameter **windowsSecurity** set to **false** has been deprecated in favor of other security options. Specifically, Quantum will no longer support permission related problems on StorNext file systems created with these parameters:

- File systems created using StorNext 5.3.x or earlier that have the FSM configuration parameter **windowsSecurity** set to **false**, or
- StorNext file systems created with any version of StorNext 5.4.x that use the configuration parameter **securityModel** set to **legacy** and **windowsSecurity** set to **false**.

Although not changed with StorNext 6, note that setting **windowsSecurity** to **false** when the **securityModel** is set to **legacy** might be removed in a future release.

Quantum recommends customers with file systems created with these parameters convert to one of the supported security models. See [StorNext Security](#) for additional details in the [StorNext 6 Documentation Center](#).

MySQL File-per-table Conversion

The **sn_fpt_convert** script in StorNext 6 allows you to split a global datafile into separate files for each table as well as enable compression for those tables. Quantum recommends converting and compressing database data files for systems using a global datafile. The process allows you to defragment and compress Storage Manager's MySQL database tables, which reduces disk usage and disk i/o while running.

If your system is configured with Storage Manager running MySQL using a global data file (applicable to StorNext releases prior to StorNext 5 release 5.2.x), then Quantum recommends you run the conversion script. For additional information, see [Overview of MySQL File-per-table Conversion](#) in the [StorNext 6 Documentation Center](#).

i Note: Running the conversion script requires Storage Manager be restarted twice, once at the start of the process, then again at the end of the process.

Configurable Retrieve Order

StorNext Storage Manager has been enhanced to provide a means of overriding the default order for

choosing copies on a retrieve. The retrieve order is defined on a per policy basis and can be configured either through the **Steering Tab** of the Storage Manager Policies page or by using the policy class commands, `fsaddclass` and `fsmodclass`. See [About Retrieve Order](#) for additional information in the [StorNext 6 Documentation Center](#).

Enhancements to the Tertiary Storage Manager (TSM)

Beginning with StorNext 6, TSM introduces exclusion behavior by default for all new installations for the following atomic file name patterns:

Exclusion	Exclusion File	Exclusion Type	Exclusion Pattern
Storage	<code>/usr/adic/TSM/config/excludes.store</code>	EXACT	<p><code>._DS_Store</code> and <code>.DS_Store</code></p> <p>i Note: Atomic file names that exactly match <code>._DS_Store</code> and <code>.DS_Store</code> are never stored.</p>
Truncation	<code>/usr/adic/TSM/config/excludes.truncate</code>	CONTAINS	<p><code>/._</code></p> <p>i Note: Atomic file names that start with <code>._</code> are stored, but are never truncated.</p>

If you do not want Tertiary Manager to perform this exclusion behavior, edit the files to remove these lines, or precede each of these lines with the `#` character to disable the exclusion check(s). See [Storage Exclusions](#) or [Truncation Exclusions](#) for additional information in the [StorNext 6 Documentation Center](#).

i Note: Beginning with StorNext 6, you can use the TSM commands without having to set up various environment items. For releases prior to StorNext 6, you must source in the environment or set up various variables to run the TSM commands.

Support for Server-Side Encryption with Amazon Web Services (AWS) Key Management Service (KMS)

With this release, the Storage Manager's server-side encryption capabilities are enhanced to include support for server-side encryption with Amazon Web Services (AWS) Key Management Service (KMS). See [Adding a Storage Manager Policy](#) for additional information in the [StorNext 6 Documentation Center](#).

StorNext Amazon Web Services (AWS) Regions Configuration File

An AWS regions configuration file (`/usr/cvfs/config/awsregions.json`) has been added. By modifying this file, you can allow access to newly defined endpoints and region names supported by the Amazon

Simple Storage Service (S3) and the Amazon Web Services Security Token Service. See [AWS Region Endpoints](#) for additional information in the [StorNext 6 Documentation Center](#).

Enhancements to Directory Quotas

This release includes enhancements to the `quota_report.csv` files found in the `/usr/cvfs/data/` directory. If you have **Directory Quotas** enabled, the **Directory Quota Name Space (DQNS)** path names are now quoted with double-quotes. This allows the CSV files to parse correctly when special characters like commas, quotes, and newlines are present in the directory names.

For example:

```
Type,Name,Hardlimit,Softlimit,Timelimit,CurAllocsize
group,root,0,0,0,722468864
user,root,0,0,0,722468864
dir,"/dqns",1099511627776,966367641600,10080,0
dirfiles,"/dqns",0,0,0,0
dir,"/dqns_with_comma_",2199023255552,2040109465600,20160,0
dirfiles,"/dqns_with_comma_",0,0,0,0
dir,"/dqns_with_quotes_\"",3298534883328,3113851289600,30240,0
dirfiles,"/dqns_with_quotes_\"",0,0,0,0
dir,"/dqns_with_newline_
",0,0,0,0
dirfiles,"/dqns_with_newline_
",0,0,0,0
```

This new format is compatible with Microsoft Excel and other programs that open CSV files. See [Manage Quotas](#) for additional information in the [StorNext 6 Documentation Center](#).

StorNext File System Data Coherency

Beginning with StorNext 6, I/O coherency is handled using “tokens.” The configuration variable, `ioTokens`, can be set to false to re-enable the DMA coherency model. The default is true which allows I/O to use the buffer cache on each node. See [StorNext File System Data Coherency](#) for additional information in the [StorNext 6 Documentation Center](#).

Changes to the Rebuild Policy for StorNext 6

The rebuild policy (`fspolicy -b`) is automatically run on each managed file system on a weekly basis. The purpose of this policy is to find candidates for other policies (for example, store, truncate, and so on) that may have been dropped or lost during normal processing.

i Note: The purpose and schedule for the rebuild policy have not changed but some of the internal processing has been updated.

Prior to StorNext 6, the process of scanning for candidates by the rebuild was built on the Storage Manager (SM) mapping process. The mapping process is time consuming and expensive in terms of resources (for example, CPU and disk space). For StorNext 6, the scanning process, by default, uses the new process, **mdarchive**, for checking for missing candidates. With the change in the scanning mechanism, the rebuild policy runs much more quickly and uses less resources, and also include changes in behavior for the command.

- One of the reasons the command runs quicker is that it does not look for all candidates on a file system, but just a batch.
 - The batch size, by default, is 50,000 but can be updated by using **sysparm: METADB_SM_QRY_LIMIT**.
 - The batches found, are the oldest candidates on the file system, newer candidates may exist, but are ignored.
- In general, the change should not be noticed. Candidates rarely get lost or dropped and when it happens, it is typically a handful and not thousands. The new rebuild picks these up.
 - It is possible that if truncation candidates are lost, it may be some time before a rebuild picks those up. In general, that is not a concern for day-to-day processing, but if there is a reason to believe that significant numbers of truncate candidates have been lost, see the next bullet.
- If candidate lists have become completely destroyed and a complete rebuild of the lists is needed, then use the old mapping process. Set the system parameter, **MAPPING_SOURCE**, to **disk** before running the rebuild by hand. For additional details, refer to the **fspolicy** man page in the **StorNext 6 Man Pages Reference Guide** available on the [PDFs Downloads page](#) in the [StorNext 6 Documentation Center](#).

Sleeping Xsan Clients not Supported

StorNext does not support Xsan clients that sleep. Adjust the **Energy Saver** settings on such systems so that **Computer Sleep** is disabled.

i Note: Manually putting an Xsan client to sleep is not supported.

If you prefer to lock the screen, then use a method that does not cause the system to sleep. For example, on some macOS systems, press Ctrl + Shift while quickly pressing and releasing the power button. See the macOS documentation for additional information.

Beginning with StorNext 6, StorNext MDC appliances send wake-on-LAN packets to Xsan clients to attempt to prevent them from sleeping. This should avoid problems if a system is inadvertently put to sleep. However, if sending wake-on-LAN packets is determined to be problematic, disable the behavior by creating the file **/usr/cvfs/config/DSM_control.conf** on the MDC and add the following line:

```
export FSM_WAKE_UP_MAC_CLIENTS=0
```

i Note: Restart StorNext for the change to take effect.

i Note: In an HA configuration, perform the procedure on both the primary and secondary MDCs.

Compatibility and Support

The [StorNext 6 Compatibility Guide](#) provides the basic compatibility for StorNext 6 release 6. 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

Operating System	Change Request Number	Service Request Number	Description
All	31624	1153036, 1486162, 3553150, 3553148, 3607120	An issue was fixed to prevent the rebuild policy from taking an excessive amount of time to execute.

Operating System	Change Request Number	Service Request Number	Description
All	39638	1473164	An issue was fixed to prevent the command, cvadmin , from displaying a confusing error message.
All	40906	3509432	An issue was fixed to allow the fsports file MaxPort < MinPort CVFS warning message to alert you of a problem.
All	42071	1545660, 3616844	Enhancements were made to the file, /usr/adic/DSM/examples/nss_cct1.example , in order to provide more client entries.
All	44383	1600844, 1610772, 3351272, 315438	An issue was fixed to prevent the command, fsretrieve , from becoming unresponsive if too many files are specified on the command line.
All	49257	3724506	An issue was fixed to prevent a database duplicate key error from being encountered within the table, filecomp1 , during a file store.
All	52714	3468692, 3702552, 317705	An issue was fixed within the StorNext GUI that prevented you from removing all the tape drives.
All	54341	3502744, 3702062, 3745926, 329649	An issue was fixed to prevent the command, cvfs , from using the byte swapped message length to allocate the qustat response buffer.
All	55758	3412502	An issue was fixed to prevent keywords that are associated with file system names from being used.
All	57699	3562994, 329978	An issue was fixed to prevent the file size in trace_01 from always being zero (0) for LTFS tapes.
All	58070	3578098	Enhancements were made to the command, cvfsck , to include updated naming scheme for orphans.
All	60834	3614180, 328821	Enhancements were made to the StorNext GUI so that the atime of a retrieved file is similar to the output of the command, fsretrieve -a .
All	60916	304538, 292011	An issue was fixed that caused the xdi_Test program to not function properly. The xdi_Test program is used by the Quantum Technical Support team to validate various xdi items.
All	61297	3741336	Enhancements were made to the StorNext GUI so that the latest version of the StorNext PDFs are available from the StorNext Documentation Center .

Operating System	Change Request Number	Service Request Number	Description
All	61390	3625134	Enhancements were made to the StorNext GUI to inform you to click Apply , if changes are made on the page, <u>File Systems > Truncation Parameters</u> .
All	61684	3608098	An issue was fixed that caused data corruption when a DLC client times out on a proxy server.
All	61768	3643696 3732168	An issue was fixed that prevented RAS email authentication.
All	62020	3649742, 3656688	An issue was fixed that prevented replacing a drive in the StorNext GUI, if the serial number changed for the drive.
All	62848	334749	Database insert deadlock in activefl (fs_eventd)
All	63274	3680512	A misspelling was fixed in the VopNotSupported error message, Unsupported Vops message.
All	63847	3652440	A misspelling was fixed in a StorNext Storage Manager log message, which enables support to analyze problems with unmanaged directories.
All	63986	3546218	Enhancements were made to the sncmpare utility to detect files missing several database records and to fix all of the problems in a single pass instead of two passes.
All	63995	3710428	Enhancements were made to all of the Tertiary Storage Manager (TSM) daemons to check if another instance of the daemons are already running when invoked, and safely exit to prevent TSM from crashing or to work incorrectly.
All	64628	3723356	A misspelling was fixed in the LogfailedRevoke Delay message.
All	64754	3731180, 327433	An issue was fixed that prevented the command, fs_stranger_mover , from retrieving from copy2 media if copy1 failed.
All	64782	328821	Enhancements were made to the StorNext GUI to include the new Truncate Directory feature.
All	64820	324979	fsm can panic with OPEN_BUSY set in open_deref_nolock when open_refs hits 0.
All	64959	3733640, 3729298	bug just for tracing code used to debug: fsm can panic with assert "!(pclient->flags & CLIENT_IS_AT)"

Operating System	Change Request Number	Service Request Number	Description
All	65054	3733044, 317792, 319468, 327212, 324979, 333720	An issue was fixed where constant disk rescans on clients with stripe group are marked as down.
All	65289	291964	Enhancements were implemented so that the Altstore daemon discovers file name changes that occur during transfer to allow it to finish the transfer under the new name and avoid a failure from using the old name.
All	65328	329880	An issue was fixed where the file, /usr/cvfs/config/dpserver.<file_system> was created unexpectedly, and overrides the default dpserver settings.
All	65333	264042	An issue was fixed that prevented the command, fsrecover -druat -n , from duplicating the recovered directory structure in the target directory.
All	65356	291603	An issue was fixed that caused LTFS EIO drives to be taken offline.
All	65359	317484	Enhancements were made to allow WindowsSecurity to be turned off after a file system is created.
All	65524	298594, 311826, 311869, 331436	An issue was fixed that caused RAS tickets for the All CVFS file systems are not mounted error message to be automatically delivered to Quantum Technical Support. The RAS tickets are no longer automatically delivered, and you can decide if Quantum Technical Support is notified.
All	65544	297332, 301259	Enhancements were made to include the ability to unmap all unused space using the command, cvmkfs .
All	65595	300308, 333018	An issue was fixed that prevented reserved inodes from being given back when a transaction is aborted.
All	65603	298505, 300800, 302483, 303615, 310863, 333020	Enhancements were made to improve daemon resiliency and to fix the connector errors that resulted from the pace of database activity from the Altstore daemon, which exposed errors in the software connecting the daemon to the database
All	65642	295253, 318754	Enhancements were made to the StorNext GUI to display usage statistics as binary data instead of decimal units.

Operating System	Change Request Number	Service Request Number	Description
All	65702	303104	An issue was fixed that prevented the nssdbg.out from logging anything confirming that the command, fsmpm , found the nss_cctl.xml file.
All	65705	304538	Enhancements were made to the archive_cmp utility to include more descriptive output.
All	65764	309023, 305133, 318227	An issue was fixed where a sleeping Xsan client locked up entire NAS and SNFS clusters with open files.
All	66093	300551	An issue was fixed where SAMFS movers became unresponsive during retrieve operations.
All	66210	314407	Enhancements were made to the Altstore string-printing software so that it does not misinterpret unusual source-file names as string-printing macros.
All	66214	313016	An issue was fixed where the command, fsretrieve , failed reading ANTF format files that had been converted from SAMQFS data
All	66215	314353	fsm panic: /usr/cvfs/bin/fsm ASSERT failed "PIT_IN_TIME_SPLAY(asp)" file /scm/nightly/VM-0-RedHat60AS-26x86-64-SP0/sn/snfs/fsm/alloc_asr.c, line 1699
All	66224	314826, 317863, 319270, 321761, 327347, 330014, 330583, 333304, 338152	Deadlock with Open_clients_lock at 5.4.0.1
All	66225	291864	An issue was fixed where the directory quotacheck namespace check failed.
All	66343	303532, 325222, 338660, 337355	An issue was fixed that caused cvfs_req_data kmem_cache inconsistencies and caused Linux kernel loops.
All	66388	312526	An issue was fixed the caused segments of a multisegment file from not being retrieved.
All	66444	303542	Enhancements were made to the RAS tickets in order to downgrade the SL_EVT_PERF_DEGRADED ticket to severity two (2).

Operating System	Change Request Number	Service Request Number	Description
All	66482	317659	An issue was fixed where recursive fsretrieves fail if any file in recursion is on disk.
All	66489	n/a	An issue was fixed with the Altstore string-printing software so that it does not misinterpret unusual source-file names as string-printing macros.
All	66502	319178	An issue was fixed that prevented the command, snac1 , from properly setting the attribute, READONLY , permission on a Directory.
All	66653	321143	Enhancements were made to improve daemon resiliency and to fix the connector errors that resulted from the pace of database activity from the Altstore daemon, which exposed errors in the software connecting the daemon to the database
All	66678	3692684, 3700462, 317928, 319022, 325160, 312495	An issue was fixed that caused RAS tickets for the No media found to satisfy request error message to be automatically delivered to Quantum Technical Support. The RAS tickets are no longer automatically delivered, and you can decide if Quantum Technical Support is notified.
All	66681	319572, 317526, 329425, 332236	An issue was fixed where disks in use check requested information about clients.
All	66724	317843	Enhancements were made to the fs_foreignMigration utility so that it removes all completed child processes as soon as possible without having to restart the Tertiary Storage Manager. The enhancements also prevent confusion about which child processes are actually running and any potential problems regarding the number of running processes.
All	66783	323799	An issue was fixed to allow the output of the customer-modifiable Altstore transfer script to be recorded at the ERROR logging level to help with debugging.
All	66832	317843	Enhancements were made to the StorNext Storage Manager to increase the maximum number of disk-to-disk movers to ten (10).
All	66848	317843	Enhancements were made to the man page for the fsgetforeign script to warn you that if the script is modified, you need to ensure that special characters (for example, ` and \$) are handled properly. i Note: The fsgetforeign script is used by the Foreign File System Migration feature.

Operating System	Change Request Number	Service Request Number	Description
All	66882	322824	MAC client panic "CVFS ASSERTION FAILED: filecookie == ntohq(replydata->td_file_cookie)
All	66888	323136	<p>An issue was fixed where files that do not get stored because they are zero length, match the exclusions pattern, or change before they reach min-time-to-store.</p> <p>The files are now set aside for later processing to allow Altstore processing of files that are ready to transfer to the remote-store location.</p>
All	66928	323136	<p>An issue was fixed where files that do not get stored because they are zero length, match the exclusions pattern, or change before they reach min-time-to-store.</p> <p>The files are now set aside for later processing to allow Altstore processing of files that are ready to transfer to the remote-store location.</p>
All	67047	325833	An issue was fixed where the fs_stranger_mover became unresponsive during the command, fsretrieve (SAMFS/SAMQFS movers).
All	67074	323331	bart success messages logged in OS logs
All	67112	326117	An issue was fixed that prevented the command, snquota , when the file, nss_ctl.xml , was in use.
All	67431	n/a	Enhancements were implemented to include an accounting tool that allows you to report tertiary storage consumption per group ID.
All	67492	324979	fsm can panic in pushing the tail or adding an inode to a transaction or any number of ways (memory corruption)
All	67528	329978	An issue was fixed that prevented the command, cvfsck -i , from writing an error to stderr .
All	67545	331672	An issue was fixed where the SNFS Client directory cache contained stale data.
All	67561	324979	checkin updates to error injection code ... client disconnects in the fsm and other minor non-functional code edits
All	67596	332367, 333713	Enhancements were made to the RAS tickets to downgrade the SR Error Code: IO Error to severity two (2).

Operating System	Change Request Number	Service Request Number	Description
All	67683	324979	fsm can panic with memory corruption usually in a trans_item_t and sometimes in heap
All	67709	324979	A close or inodereturn racing with the client disconnecting with exclusive access or the token should be accepted
All	67716	324979	reconnect handling needs to look for OPEN_BUSY so two resyncs don't trample on each other
All	67747	331396	Enhancements were made to the Web Services to include APIs that allow querying non-Posix attributes without opening a file.
All	67839	324979	A client that just reads can truncate the file back due to a client reconnect.
Linux	38488	1646652, 3489886, 3688742, 3710144, 3712134, 3732728	An issue was fixed that caused the software to not update the path when brought back online, if a device path for a tape library changed and the tape library went offline. Now the tape library path is properly updated when the drive is brought back online.
Linux	51589	3513422	An issue was fixed where the file system space was not freed due to an inconsistent Storage Manager state.
Linux	52333	3468678, 3468678, 3500342, 3543714, 296701	Enhancements were made to the StorNext GUI to include the ability to provide the SMTP Server Port during your configuration of the email server.
Linux	58139	3468678, 3468678, 3500342, 3543714, 296701	Enhancements were made to the StorNext GUI to include the ability to provide the SMTP Server Port during your configuration of the email server.
Linux	62556	3668188	Enhancements were made to the StorNext GUI to prevent you from selecting an HAsHared file system as a mount point for SDISK
Linux	63603	317140	StorNext SAN client support for Ubuntu 16.04 Xenial
Linux	64616	3724654	An issue was fixed to include the full path name to an Admin Alert message when a file fails to store because it is too large. The fix allows you to easily locate the file and take appropriate actions.

Operating System	Change Request Number	Service Request Number	Description
Linux	64652	3729580, 318887	Deadlock between inode_delete_space and VopRename
Linux	64758	3727008, 304461, 304936, 309057, 317498, 317142, 321954, 324079, 339523, 342546	An issue was fixed that was causing a kernel crash due to cvfs_proxy_1stn .
Linux	65190	3745562	An issue was fixed that caused files including the apostrophe (') symbol in the file name to be listed with additional backslash (\) symbols in the file name when the commands ,fsfileinfo and fsmedinfo -l , are used.
Linux	65237	3736606	An issue was fixed that caused the StorNext GUI to become unresponsive when Admin alerts reached 150,000 or more.
Linux	65242	3733306	FSM panic - PANIC: /usr/cvfs/bin/fsm ASSERT failed "free_ip->i_idinode.idi_flags & InodeFlagFree" sn/snfs/fsm/inode.c, line 2302
Linux	65446	296744	An issue was fixed that caused adding a new tape drive in the same device path when a tape drive is taken offline.
Linux	66642	0319572, 317526	Using cvmkfs does not scale well in large system customer environments
Windows	40608	3648662, 3688918	LDAP messages fill-up nssdbg.out log
Windows	65339	284815, 291659, 291968, 304409	An issue was fixed that caused StorNext 5 Windows client to generate exceptions.
Windows	65737	303132	unixpermbits/mdc does not work for Windows client using local account
Windows	65781	305320, 3661514	Enhancements were made to the RAS tickets to lower the severity for the ERR GetLdapMap: cond WAIT error to Informational .
Windows	66811	320059	An issue was fixed to correct the Windows log message's error number in Cannot RegisterEventSource (9) .
Windows	68205	334035	An issue was fixed that prevented the SMB2 from working correctly, if the file system is configured with the legacy setting, No Security .

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](#).

i 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, 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 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

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6, 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. 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 is significantly faster than that from StorNext 4.3.x to StorNext 6.

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

See the following sections for information regarding compatibility between this release and StorNext components and features.

- [Appliance Controller below](#)
- [Infiniband below](#)
- [Lattus below](#)
- [Partial File Retrieval below](#)
- [StorNext Web Services on the next page](#)
- [Apple Xsan on the next page](#)
- [Supported Browsers on the next page](#)

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

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 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.

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, 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, 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](#).

General Considerations

This section provides information about items to consider for StorNext 6.

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:

- [StorNext File System Known Issues below](#)
- [StorNext Storage Manager Known Issues on page 28](#)
- [StorNext GUI Known Issues on page 33](#)
- [StorNext Installation, Replication, HA, and Other Known Issues on page 34](#)

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 data-bbox="623 764 1458 867">(Info) Inode conversion will stop at IEL chunk <xxxxx></pre> <p>And you do not see messages like this (shortly thereafter):</p> <pre data-bbox="623 989 1458 1092">(Info) Conversion completed on 100/1079 IEL chunks</pre> <pre data-bbox="623 1142 1458 1245">(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 data-bbox="623 1436 1458 1507">(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	In StorNext 6 (and later), file system names might be qualified with cluster information. For example, if the fsmcluster file contains the following:
-----	-------	-----	--

```
default_cluster cluster1
```

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.

Issue:

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:

```
Error getting data for file system/s (snfs1):
Filesystem Name or Mount Point is not provided.
```

Workaround:

Do not use the syntax **fs@cluster** in **/etc/fstab** for file systems in your default cluster when you are using the StorNext GUI.

Related Issue:

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.

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_FACTOR sysparm. 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_FACTOR sysparm from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre>FS_T10K_BLOCK_FACTOR=32;</pre> <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_FACTOR sysparm must not be set to a value that exceeds 32. 2. Restart Storage Manager to ensure the change in Step 1 goes into effect: <pre># tsmstop # tsmstart</pre> 3. Verify the FS_T10K_BLOCK_FACTOR sysparm contains the new value: <pre># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> 4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients. 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>snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<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>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.</p>
All	46693	n/a	<p>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>This is due to the snbackup -s process reading the backup.1f status file while the backup process is updating it.</p> <p>Workaround:</p> <p>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>
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>i 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>

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 /usr/adic/TSM/logs/tac which contains the text:</p> <pre data-bbox="618 474 1446 711">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 data-bbox="618 800 1446 989" 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	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.
-----	-------	--------	--

Workaround:

If you are reading and writing data from/to LTFS media, the following actions are required:

- 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.

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:

```
crontab -e -u tdlm
```

Add the following 2 lines that read:

```
# 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
```

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

- 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:

```
crontab -e
```

Add the following 2 lines that read:

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

```
# This will filter out LTFS logs from the
fsmppm.out file hourly

0 * * * * sed -i '/ LTFS/d'
/usr/cvfs/debug/fsmppm.out
```

i 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