



Release Notes

Release	5.3.1
Supported Product	StorNext 5
Date	April 2016

Contents

What's New in StorNext 5 release 5.3.1	3
Purpose of this Release	3
New Features and Enhancements in StorNext 5 release 5.3.1	4
Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.3.1	4
StorNext Compatibility	6
Quantum OS Upgrade Support Policy	6
StorNext and Linux Interoperability	7
Supported StorNext Upgrade Paths and Upgrade Considerations	7
StorNext Software Upgrade Matrix	7
Considerations for the StorNext File System Directories	7
Journal Size Guidelines	7

© 2016 Quantum Corporation. All rights reserved. Artico, Be Certain, DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, GoVault, Lattus, NDX, the Q logo, the Q Quantum logo, Q-Cloud, Quantum, the Quantum logo, Quantum Be Certain, Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo, 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. Quantum specifications are subject to change.



Distributed Data Mover (DDM) Guidelines	8
Considerations When Upgrading NFS Server Nodes to StorNext 5	8
Database Schema Update During Upgrades	8
Compatibility Between StorNext and Other Products	9
NAS	9
Infiniband	9
Lattus	9
Partial File Retrieval	9
StorNext Web Services	10
StorNext API (SNAPI)	10
Apple Xsan	10
Supported System Components	10
Supported Browsers	10
General Considerations	11
Checksum Performance Considerations	11
Upgrading Appliances	11
Known Issues	12
StorNext File System Known Issues	12
StorNext Storage Manager Known Issues	15
StorNext GUI Known Issues	21
StorNext Installation, Replication, HA and Other Known Issues	22
StorNext 5 Release 5.3.1-Specific Pre-Upgrade Requirement for StorNext Connect	28
Contacting Quantum	30
Getting More Information or Help	30

What's New in StorNext 5 release 5.3.1

Purpose of this Release

StorNext 5 release 5.3.1 delivers an array of innovative new features, important enhancements, support for new tape libraries and drives, updated client support and improved support for video surveillance workloads. For complete details, see [New Features and Enhancements in StorNext 5 release 5.3.1 on the next page](#).

StorNext 5 release 5.3.1 also resolves several customer reported issues listed within the section [Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.3.1 on the next page](#).

Corrections to the StorNext 5 Compatibility Guide

Quantum recommends you download the latest revision of the *StorNext 5 Compatibility Guide* online at <http://www.quantum.com/sn5docs>.

i Note: Your software shipped with an incorrect revision of the *StorNext 5 Compatibility Guide*. The following corrections have been made to the *StorNext 5 Compatibility Guide* corresponding to the file name "6-68043-01_RevW_StorNext_5_Compatibility_Guide.pdf":

Section	Section Title	Page	Description of Correction
6	Supported Operating Systems and Platforms	14	Replaced SLES 11 kernel/release version 3.0.13-0.27.1 (SP2) with 3.0.13-0.27 (SP2) .
6	Supported Operating Systems and Platforms	14	For the SLES 11 operating system, added the following acronyms to the StorNext 5 release 5.3.1 column: "DC, SNSM, DLS, DDM, SC, DLC, FX"
6	Supported Operating Systems and Platforms	14	Removed the entire row for the SLES 11 kernel/release version 3.0.13-0.27 (SP4) .
6	Supported Operating Systems and Platforms	15	Replaced SLES 11 kernel/release version 3.0.101-63 (-SP4) with 3.0.101-63 (SP4) .
13	Supported Quantum Library and Drive List	22	Added the following statement to the Note : "StorNext LTFS is not supported with LTO-7 tape drives."

New Features and Enhancements in StorNext 5 release 5.3.1

- Provides faster and more efficient object store space management with the ability to recursively delete directories.
- Introduces Lattus space reporting.
- Introduces per NAS share data collection generation.
- Introduces a new **Alternate Store Location (ASL)** license. Beginning with StorNext 5 release 5.3.1, the **Alternate Store Location** license is required to enable ASL when configuring Storage Manager policies.
 - For information about the **Alternate Store Location** license, see the **Licenses for Features** section in the *StorNext 5 Licensing Guide* available online at <http://www.quantum.com/sn5docs>.
 - For information about functionality and using the **Alternate Store Location** feature, see the **Alternate Store Location** section in the *StorNext 5 User's Guide* available online at <http://www.quantum.com/sn5docs>, and also the **Configuration > Storage Manager Policies > General** online help page in the StorNext 5 release 5.3.1 software GUI.
- Adds support for Quantum i500, i40 and i80 libraries with IBM LTO-7. For additional information, see the **Supported Quantum Library and Drive List** section in the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

i Note: StorNext 5 release 5.3.0 is “Quattro Ready” today; completing Quantum's final qualification tests with Quattro and Storage Manager to support these new libraries.

- Adds additional support for Red Hat Enterprise Linux and SUSE Lixux Enterprise Server clients (SAN, DLC and FX). For additional information, see the **Supported Quantum Library and Drive List** section in the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

i Note: Support for Red Hat Enterprise Linux version 7.2 will slightly trail the General Availability date.

- Adds support for the IBM TS4500 library. For additional information, see the **Supported Quantum Library and Drive List** section in the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.
- Adds support for IBM TS1150 “enterprise” tape drives. For additional information, see the **Supported Quantum Library and Drive List** section in the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.
- Enhancements to the file system to support Video Surveillance Solutions.

Customer Reported Issues and Enhancements Addressed in StorNext 5 release 5.3.1

The following table lists the customer reported issues and enhancements addressed in StorNext 5 release 5.3.1.

Operating System	Change Request Number	Service Request Number	Description
All	50792	3448236, 3445930, 3447604	The rasexec syntax does not appear to work
All	57261	3556264	Fixed an incorrect behavior with potential data loss when block size exceeds drive capabilities
All	57943	3573622, 3586078	Disk Proxy Client errors number should be log error messages not error numbers
All	59674	3591506, 3611782, 3611680	Enhancement for implementing API to disable SNFS security check on a per-process basis
All	59806	3597304	cvfsck : Pending inodes cause the hard link count to be wrong
All	59807	3597304	inode expansion with waiting threads leaks free inodes
All	59953	3614876, 3605730	flush_ipaddrs_if_stale() has too small of a timeout
All	60646	3589032, 3568592	Thread processing fallocate() deadlocks with cvfbiod on cv_resextentlock rwlock
All	60664	3608098	Data corruption when DLC client timeout on Proxy server
All	60814	3635792	DSM_control stop logic unmounts HaShared file system too soon
All	60856	3610710	When marking SG up and connecting or labeling LUNs so visible ... clients fail to reconnect.
All	61061	3620852 3630968 3634864 3642114	5.3.0 Inode conversion unable to resume if stop before complete and directory namedStreams issue.
All	61270	n/a	share accounting report scripts
All	61468	3632562	DiskLic: 1st Entry In Disk Catalog Not Recognized
All	61549	n/a	metadump restore can set wrong block count for special file types
All	61567	3616424	MPI or other distributed apps can see stale file size with fseek/tellp

Operating System	Change Request Number	Service Request Number	Description
All	61570	3636214	SNAPI example will not compile due to missing tinyxml static library
All	61774	3639284	inodes with more than 255 hard links cause metadump restore failures
All	62122	3644802	Snbt_rel_descendants() doesn't loop correctly.Â Still allows deadlock in Dirreaddir_old().
Linux	59860	3546218	sncompare "fs" test does not complete when run on a FS containing 120-million directories
Linux	60242	3596060	sncompare "fs" test hangs on msa2dmi_stat() call on a pipe (fifo) file
Linux	61184	3546218	sncompare is killed or hangs if there are too many entries in DIRPATH table
Windows	61000	3623464, 3656700	Win client: critical cvfs errors shown on event viewer on when any action is taken on directory mounted file system
Windows	61884	3646008	Windows 10 Update breaks drive-letter SNFS mounts

StorNext Compatibility

For information on StorNext 5 compatibility with operating systems, kernel versions, hardware platforms, drives, libraries, StorNext Appliances, StorNext client interoperability, and other compatibility items, see the *StorNext 5 Compatibility Guide*. SNAPI, Partial File Retrieval, and Apple Xsan compatibility information is provided in separate documents.

Quantum OS 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 updating the update level or service patch beyond the currently supported levels shown in the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

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 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

Considerations for the StorNext File System Directories


On upgrades to StorNext 5, it may be noted 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 StorNext 4.x. This is expected behavior.

Journal Size Guidelines

The absolute minimum Journal Size in StorNext 5 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 StorNext 5 when the Metadata Controller (MDC) is operating on StorNext 5.

 **WARNING:** Upgrades (such as platform, service pack, etc.) are intended to be done to all of the Lattus and metadata systems present in a given deployment (for example, if M662, A10, C10, S10 are present, they all must be updated, one appliance cannot be "left behind").


Considerations When Upgrading NFS Server Nodes to StorNext 5

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 5, special consideration must be made for Linux NFS servers.

In order 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 5. The M660 appliance can achieve approximately one hour for every hundred million entries in the filecomp tables. Smaller appliances and Software Only configurations may take considerably longer depending on CPU and memory availability.

 **Note:** The database schema update conversion time from StorNext 4.7.x to StorNext 5 is significantly faster than that from StorNext 4.3.x to StorNext 5.

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 may 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** may require a significant amount of time. Quantum recommends you execute the query before the day of an upgrade.

Compatibility Between StorNext and Other Products

This section describes various interactions between this release and StorNext components and features.

NAS

For NAS software compatibility, see the *NAS Compatibility* section in the **NAS Release Notes** document, available online at <http://www.quantum.com/xcelliswfdocs>.

Infiniband

StorNext 5 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

Lattus

Refer to the *Lattus Release Notes* for information about compatibility between Lattus and StorNext 5.3.1.

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

Partial File Retrieval

StorNext Partial File Retrieval (PFR) is a separately available 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: StorNext Partial File Retrieval is not supported for use with Lattus or Q-Cloud.

For information about compatibility between PFR and StorNext 5.3.1, see the *StorNext Partial File Retrieval Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

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 5 Web Services Guide* available online at

<http://www.quantum.com/sn5docs>.

StorNext API (SNAPI)

StorNext API (SNAPI) enables you to run third-party APIs with StorNext.

i Note: SNAPI documentation is applicable for releases prior to StorNext 5 release 5.0. As of StorNext 5 release 5.0, the SNAPI documentation is being maintained in the *StorNext 5 Web Services Guide*.

For information about compatibility between SNAPI and StorNext 5.3.1, refer to the *SNAPI Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

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 5.3.1, refer to the *Xsan Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

Supported System Components

Supported Browsers

For information on browsers supported with the StorNext GUI for this release, refer to the *StorNext 5 Compatibility Guide* available online at <http://www.quantum.com/sn5docs>.

General Considerations

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 may 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

The **Firmware Upgrade** menu option allows you to upgrade Xcellis, Artico, Pro Foundation and G300 systems, and StorNext Metadata Appliances. This includes upgrading the StorNext software and also upgrades firmware for appliance components, if applicable.

i Note: Use the StorNext GUI to perform all firmware upgrades. The **Firmware Upgrade** menu option is **only** available on the systems indicated above.

- For instructions on upgrading your firmware, refer to the current Release Notes for your particular appliance:
 - For the current Xcellis Release Notes, see <http://www.quantum.com/xcelliswfd docs>.
 - For the current Artico Release Notes, see <http://www.quantum.com/articodocs>.
 - For the current M660, M440, M330 Metadata Appliance and Pro Foundation Release Notes, see <http://www.quantum.com/snmdc docs>.
 - For the current G300 Gateway Appliance Release Notes, see <http://www.quantum.com/sngatewaydocs>.

i Note: During the firmware upgrade progress, your login session is terminated when the upgrade process has completed on the other node. Log in to the other node to continue monitoring the upgrade process. The upgrade is complete when full GUI functionality has been restored.

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 15](#)
- [StorNext GUI Known Issues on page 21](#)
- [StorNext Installation, Replication, HA and Other Known Issues on page 22](#)

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

Table 1 on the next page lists known issues specific to the StorNext File System process.

Table 1: StorNext File System Known Issues

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 may 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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	57304	3561252	<p>Due to a limitation in the Linux rpc.mountd process, mounting a StorNext file system over NFS may return an error or hang in certain cases. To encounter the issue, a directory must be exported that is below the root of the file system. For example, if the StorNext file systems is mounted locally on the NFS server as <code>/stornext/snfs1</code>, then the following export is exposed:</p> <pre>/stornext/snfs1/myshare *(rw, sync)</pre> <p>Whereas this export entry is not:</p> <pre>/stornext/snfs1 *(rw, sync)</pre> <p>Also, the NFS server must be running Linux and have StorNext 5.x installed and the version of rpc.mountd on the system must correspond to version of <code>nfs-utils</code> less than 1.9.0. For example, RHEL6 is exposed but RHEL7 is not.</p> <p>Finally, the directory being exported must have an inode number that is greater than or equal to 4294967296 (2^{32}). This can be checked by executing the following command:</p> <pre>ls -id</pre> <p>StorNext file systems having large inode numbers will usually have one or more of the following attributes:</p> <ol style="list-style-type: none"> 1. A file count greater than 16 million. 2. Multiple metadata stripe groups. 3. Stripe groups containing mixed data and metadata. <p>Workaround:</p> <p>If the problem is encountered and the NFS server cannot be upgraded to a version of Linux that is not exposed to the issue, the workaround is to mount the file system on the NFS server using the compat32 SNFS mount option.</p>
Windows	62342	n/a	<p>Windows File Explorer may not display or allow access to a StorNext mounted file system. The specific circumstances causing the issue have not been identified. Testing indicates this issue is more likely to occur on Windows 10 systems that are not joined to a domain. This issue appears to be related to a Microsoft Windows 10 update release in March of 2016, but the relation between the problem and any specific Microsoft update has not been established.</p> <p>Workaround:</p> <p>The issue only occurs on Windows 10 systems with the Microsoft March of 2016 updates and running the StorNext 5 release 5.3.1 release. Do not install StorNext 5 release 5.3.1 on a Windows 10 system under these circumstances.</p>

StorNext Storage Manager Known Issues

Table 2 on the next page lists known issues specific to StorNext Storage Manager.

Table 2: StorNext Storage Manager Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p data-bbox="620 378 1474 512">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 data-bbox="620 520 786 552">Workaround:</p> <p data-bbox="620 560 1474 730">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> <p data-bbox="620 739 1474 840">i Note: This workaround may also help improve the performance of the faster LTO drives like LTO-6, and LTO-7 by updating the FS_LTO_BLOCK_FACTORsysparm.</p> <ol data-bbox="620 848 1474 1785" style="list-style-type: none"> <li data-bbox="620 848 1474 999">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>: <pre data-bbox="669 966 1052 999">FS_T10K_BLOCK_FACTOR=32;</pre> <li data-bbox="620 1008 1474 1142">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 should not be set to a value that exceeds 32. <li data-bbox="620 1150 1474 1486">2. Restart Storage Manager to ensure the change in Step 1 goes into effect: <pre data-bbox="669 1243 831 1310"># tsmstop # tsmstart</pre> <li data-bbox="620 1495 1474 1570">3. Verify the FS_T10K_BLOCK_FACTORsysparm contains the new value: <pre data-bbox="669 1411 1214 1478"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <li data-bbox="620 1579 1474 1785">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 data-bbox="669 1688 1442 1780">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 may 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 may 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	25506, 41413	976344, 1504258, 1635952, 3471000, 3526376	<p>In certain conditions when Storage Manager receives an early end of tape indication, SCSI sense messages may be logged by the fs_fmover process. These the messages will have a format similar to the following: Dec 28 20:38:00 MDC-Hostname sntsm fs_fmover[29837]: E1201(8)<1034815852>:fsScsi1311: {2}: Check condition: op=0Ah key=00h asc=00h ascq=02h END OF PARTITION/MEDIUM DETECTED</p> <p>Workaround:</p> <p>Early EOT reporting can cause tape I/O performance problems due to reporting of and handling of the check conditions. This is more prevalent on HP tape drives but may also occur on IBM tape drives. The workaround is to reduce the usable capacity by up to 1% for IBM drives or 5% for HP drives.</p> <p>To do so, adjust the PERCENT_FULL_TO_MIGRATE system parameter in <code>/usr/adic/TSM/config/fs_sysparm_override</code>.</p> <hr/> <p>i Note: The Early EOT reporting can also cause extreme clutter in some TSM log files.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	56261	n/a	<p>The Storage Manager <code>fsCheckTsmFilesystemsConfig</code> health check is a process that runs as part of the regularly scheduled health checks to help make sure that all managed file systems listed in the <code>/usr/cvfs/config/fsmlist</code> are mounted.</p> <p>If a managed file system appears in the <code>/usr/cvfs/config/fsmlist</code> but is not mounted, the health check process will issue a Service Ticket indicating something like the following:</p> <pre>Quantum SERVICE REQUEST INFORMATION SR Ticket Number: 1234 SR Serial Number: SV123436789 SR Problem Summary: QUANTUM software: Storage Manager Component fsCheckTsmFilesystemsConfig: Software fault SR Problem Code: 010C: 8 SR Error Code: Software fault SR Severity: 1 SR Notes: Internal Software Error: an unhandled software error has occurred. The following TSM managed file systems do not have storageManager enabled: /stornext/FSNAME</pre> <p>There is however an issue in the processing of the mounted Storage Manager enabled file systems that contain an "_" (underscore) in the file system name which can cause the <code>fsCheckTsmFilesystemsConfig</code> to skip over the counting of that file system even though it is mounted. This will then generate a Service Ticket. Systems where the file system name of a Storage Manager enabled file system matches the name of another StorNext file system following the underscore character may be exposed.</p> <p>For example, if there exists a managed file system named "snfs1" that is mounted and there is another mounted file system named "another_snfs1", the health check might incorrectly report that the file system "snfs1" is not mounted.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Verify that the file system is mounted. <p>-or-</p> <ol style="list-style-type: none"> 2. Sometimes changing the order of the file system in the <code>/etc/fstab</code> can reduce the likelihood of hitting this condition by changing the order that the file system appears in <code>/proc/mounts</code>. StorNext may need to be restart for these changes to take affect.

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>-or-</p> <p>3. Disable the health check by commenting out the fsCheckTsmFilesystemsConfig health check from the /usr/adic/TSM/config/filelist file. Lines beginning with "#" are ignored. For example:</p> <pre>\$ grep fsCheckTsmFilesystemsConfig /usr/adic/TSM/config/filelist # health_check : 0 : Config : fsCheckTsmFilesystemsConfig : 0</pre>
Linux	45718	n/a	<p>NFS clients using certain versions of Linux are exposed to a software defect in the Linux kernel, that, in a rare race condition, may lead to an application seeing an incorrect size when accessing a file immediately after it is written.</p> <p>This problem is described by RedHat Bugs 663068 and 672981.</p> <p>https://bugzilla.redhat.com/show_bug.cgi?id=663068</p> <p>https://bugzilla.redhat.com/show_bug.cgi?id=672981</p> <p>However, the Bug impacts releases of other Linux distributions as well. While the problem is unlikely to occur frequently even on systems running affected releases, customers having applications that may be impacted should contact their Linux vendors for guidance on which OS versions they should run to avoid the issue.</p> <p>Workaround:</p> <p>Contact your Linux vendors for guidance on which OS versions you should run to avoid the issue.</p>
Linux	54280	n/a	<p>Device paths to tape libraries may change when the scsi generic kernel module is loaded. This may happen on system upgrades, reboots, or failovers to a peer node. The TSM component maintains a table mapping a device to its scsi generic device path (for example, /dev/sg64). When TSM starts up, the fsconfig utility will update the device paths in the database. Prior to starting TSM, the MSM component will run the the /usr/adic/MSM/util/archive_cmp utility. This utility will check the current device paths to the potentially stale mapping in TSM. An admin alert will be issued if there is a device path mismatch between the host and the TSM database.</p> <p>Workaround:</p> <p>Verify the archive_cmp admin alert is valid as the starting of TSM might have resolved the issue. Run the /usr/adic/MSM/util/archive_cmp utility to verify if there are still problems that need to be updated. If the archive_cmp still reports a mismatch, run /usr/adic/MSM/util/archive_cmp -u to update the device mapping.</p>

StorNext GUI Known Issues

Table 3 below lists known issues specific to the StorNext GUI process.

Table 3: StorNext GUI Known Issues

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 operation. 3. Uncheck the Verify SMTP Server Connectivity box, and then retry the operation. 4. Set Authentication to NONE, and then retry operation. 5. Disable the Safari User names and passwords AutoFill under Safari > Preferences > AutoFill, and then retry operation.
All	57856	n/a	<p>The possibility of having different NIC cards installed in the same slots across boots results in having the same Ethernet alias names being used for the network interfaces of different NICs with different speeds (1G/10G).</p> <p>However, the Ethernet alias names depicted in the StorNext Metrics GUI page do not reflect this possible change of the network device representing the alias.</p> <p>Workaround:</p> <p>There is currently no workaround for Change Request Number 57856.</p> <p>i Note: StorNext Metrics data is only kept for 30 days; the StorNext Metrics Report auto-corrects any port-discrepancies 30 days after the configuration is changed.</p>
All	60659	n/a	<p>There is a known GUI issue with managing ACSLS tape libraries on the Configuration > Storage Destinations > Libraries page. If you attempt to configure a new ACSLS tape library, or edit an existing ACSLS tape library, the GUI will become unresponsive until you navigate to another page.</p> <p>Workaround:</p> <p>To workaround this issue, contact Quantum Support for a software patch that will resolve this issue if you require ACSLS tape library support in the GUI.</p>

StorNext Installation, Replication, HA and Other Known Issues

Table 4 on the next page lists known issues specific to StorNext installations, data replication, HA systems and other areas.

Table 4: StorNext Installation, Replication, HA and Other Known Issues

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47041	n/a	<p>A database index named <code>classndxatimeme</code> will be automatically added to the <code>tmdb.tier000files%</code> and <code>tmdb.tier001files%</code> tables upon starting TSM for the first time after upgrading to specific StorNext 5 releases:</p> <hr/> <p>⚠ Caution: For upgrades from 4.3.2, 4.3.3, 4.7.0, 4.7.0.1, 4.7.1, 4.7.2, 4.7.3, and StorNext 5 Release 5.0, DO NOT use this script for direct upgrades to StorNext 5 Releases 5.0.1, 5.1, 5.1.1, 5.2, 5.2.0.1, 5.2.0.2, 5.2.1, 5.2.2, 5.3, and 5.3.1.</p> <p>This index improves the performance of certain operations such as truncation policies. However, the creation of this index can take multiple hours for very large databases. TSM will be unavailable after upgrading until the indexing has completed.</p> <p>Workaround:</p> <p>To minimize TSM downtime after upgrade, the <code>classndxatimeme</code> index can be created prior to performing the upgrade using the <code>index_tierfiles.pl</code> PERL script. This file is available in the StorNext installation directory by opening a support ticket and requesting the file. (Quantum service and service partners can obtain this file from the StorNext Metadata Appliances page on CSWeb.) The script can be run while TSM is running, although it may impact the performance of other operations while the index is being added to the database.</p> <p>To manually add the index, you must have the <code>index_tierfiles.pl</code> script. Then do the following:</p> <p>Login to the primary server node, and access the command line of the system:</p> <ol style="list-style-type: none"> 1. Login to the primary server node. 2. Enter the following to source the profile: <pre style="background-color: #f0f0f0; padding: 5px;">. /usr/adic/.profile</pre> <ol style="list-style-type: none"> 3. Change to the directory where <code>install.stornext</code> resides on the installation media. For example: <pre style="background-color: #f0f0f0; padding: 5px;">cd /tmp/stornext/stornext_full/RedHat60AS_26x86_64</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>4. Verify that the database is up by running:</p> <pre data-bbox="678 407 1458 478">mysql_control start</pre> <p>5. Execute the PERL script:</p> <pre data-bbox="678 600 1458 722">./TSM /index_tierfiles.pl</pre> <p>The procedure is complete.</p>
All	53933	n/a	<p>The StorNext GUI has always supported SSL for https connections while scanning namespaces. Due to SSLv3 Poodle Vulnerability, the StorNext GUI now requires to support TLSv1.2 protocol.</p> <p>In order to run StorNext 5 release 5.2 HTTPS with Lattus 3.5.1 and Lattus 3.4.4 with LIN031 Poodle WAR applied see the Workaround.</p> <p>Workaround:</p> <p>The application of LIN031 - Poodle WAR secures the system by disabling SSLv2 and SSLv3. This WAR can be applied to Lattus 3.4.4 and 3.5.1.</p> <p>i Note: In Lattus 3.6.X this WAR is not required.</p> <p>StorNext 5 release 5.2 functions as designed with Lattus 3.6.0. If you are using HTTPS, StorNext 5 release 5.2 requires the Poodle WAR be applied to Lattus 3.5.1 and Lattus 3.4.4. However, due to the restrictions on the cipher suites available on Lattus with this WAR, the list(scan), create namespace and buckets feature in the StorNext GUI does not work in HTTPS. It does work in HTTP.</p> <p>If you want to run in HTTPS mode, you can list, create namespace and buckets in HTTP mode. When you done, you can enable HTTPS</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	57789	n/a	<p>When running StorNext replication from a deduplication-enabled StorNext filesystem to a deduplication-enabled HA StorNext filesystem, corruption is sometimes seen in files on the HA target if HA failovers occur during the replication. Attempts to read the contents of such files report EIO (5), input/output error.</p> <p>Examination of such files using snpolicy's report directive show no TAG or BLK_TAG_PRESENT flags in the file inode.</p> <p>It is not known if this problem is an alternate manifestation of CR 58814, which may also occur in these configurations.</p> <p>Workaround:</p> <p>Avoid HA failovers/reconfiguration while StorNext replication of deduplicated files is active. See also the Release Note item for CR 58814.</p>
All	60463	n/a	<p>In some error conditions, the <code>suspect count</code> for Object Storage media is incremented by 1. When the <code>suspect count</code> reaches 3, the media is not allocatable. In this situation the stream count for the controller is not decremented correctly. This has the effect of falsely reducing the total number of available streams. The only way to clear this condition is to reset the media suspect count back to 0 and restart Storage Manager.</p> <p>Workaround:</p> <p>Set the <code>suspect count sysparm</code> parameter to a higher value (for example, <code>MEDIA_SUSPECT_THRESHOLD=0x7fff;</code>) in the <code>/usr/adic/TSM/config/fs_sysparm_override</code> file.</p> <p>Also refer to the <code>MEDIA_SUSPECT_THRESHOLD</code> in the <code>/usr/adic/TSM/config/fs_sysparm.README</code> file.</p>
All	58814	n/a	<p>When running replication from a deduplicated file system to another deduplicated file system, corruption is sometimes observed in large files. Examination of the replicated files using snpolicy's <code>report</code> directive shows that there are holes in the replicated file's extended representation blobmap.</p> <p>These areas read as zeroes rather than the appropriate data.</p> <p>Workaround:</p> <p>Avoid replicating deduplicated files, especially when deduplication is in progress. The problem is more likely to be seen on very large files (tens or hundreds of gigabytes or larger). It is also more likely to be seen when replication is run on a very short schedule (minutes). It is believed to occur mostly or only when deduplication and replication are running concurrently.</p>

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

All	62454	n/a	On upgrades to StorNext 5 Release to 5.3.1 from StorNext releases prior to 5.0 on HA systems (including M660, M440 and M330 Metadata Appliances), the directory /usr/adic/wsar_agent/tmp is not created on the nodes during the upgrade process. Because of this, async web services will not function correctly.
-----	-------	-----	--

A log message similar to the following will be displayed in the **wsar_agent.log** file:

```
> [0328 13:05:35.878] 0x7f36b1346700 ERR
[wsarutils.c:107] [wsar_run_cmd] [22] wsar_run_
cmd: req_id=22: failed to open(/usr/adic/wsar_
agent/tmp/wsar.22.out), errno 2
```

Workaround:

To create the **/usr/adic/wsar_agent/tmp** directory on both nodes:

1. Log into one of the server nodes.
2. Verify the following directory exists on the server node:

```
/usr/adic/wsar_agent/tmp
```

3. If needed, create the following directory:

```
/usr/adic/wsar_agent/tmp
```

4. Repeat the previous steps for the other server node.

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	62301	n/a	<p>Starting the cvfs service outside of systemd on RHEL 7, CentOS 7, and SLES 12 environments can impact shutdown and reboot.</p> <hr/> <p>i Note: The following only applies to systems running systemd (for example RHEL 7, CentOS 7, and SLES 12).</p> <p>The systemd environment is new to RHEL 7, CentOS 7, and SLES 12 distributions, and replaces the traditional SysV init system for startup and shutdown.</p> <p>The systemctl command is used to communicate to the systemd daemon to start and stop units. Implicit units are created by systemd for services in the <code>/etc/init.d</code> directory.</p> <p>Systemd managed the starting and stopping of units on a system used to start and stop services. The cvfs service may still be running even if the unit is in-active, and the service may be stopped even if systemctl reports that the unit is active. If systemd already has a record that a unit is running, it will not start another unit of the same name. Likewise if systemd has a record that the unit is already stopped, it will not by default try to reissue the stop operation for that unit.</p> <p>If the cvfs service is started outside of systemctl (for example, service cvfs start), systemd may not try to stop the cvfs service on shutdown or reboot. If the cvfs service is stopped outside of systemctl, systemd may not try to start the cvfs service again later until the cvfs unit is marked as in-active. Starting a service outside of systemd causes the running processes to be added to a user session slice instead of the system slice. On shutdown or reboot, processes that are running in the user slice are sent the SIGTERM signal, whereas units running in the system slice will try to be gracefully stopped by systemd. If the cvfs service is started in a user slice instead of the system slice, a shutdown or reboot operation might trigger an unexpected SMITH or a hang on shutdown due to systemd not being able to gracefully shutdown the cvfs service.</p> <p>Workaround:</p> <p>The following commands may be used to manage the cvfs service and unit:</p> <ul style="list-style-type: none"> • Use systemd-cgls to determine if the cvfs service is running in the system slice or a user slice. • Use systemctl status cvfs to determine if the cvfs unit is active or in-active in the system slice. • Use /usr/cvfs/bin/DSM_control status to determine if the

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>cvfs service is running or stopped.</p> <ul style="list-style-type: none"> • Use systemctl start cvfs and systemctl stop cvfs to start and stop the cvfs unit running in the system slice. • Use service cvfs stop to stop cvfs if it is running in the user slice prior to shutting down or rebooting a system. • Following the removal of StorNext, run systemctl stop cvfs to mark the cvfs unit as in-active.

StorNext 5 Release 5.3.1-Specific Pre-Upgrade Requirement for StorNext Connect

i Note: Only apply this procedure if your M440, M660, Artico, Pro Foundation, or Xcellis is managed by StorNext Connect.

You should update the Connector on all StorNext Connect-managed systems that you plan to upgrade to StorNext 5 Release 5.3.1. This includes the StorNext Connect system (Node 1). This upgrade **MUST** be done **BEFORE** upgrading the appliance firmware in order to avoid a situation where the StorNext Connect Volume Storage widget on the Dashboard does not display data for StorNext Connect monitored systems.

How to Install the StorNext Connect Connector

i Note: If you did not install the new StorNext Connect Connector prior to applying the StorNext 5 Release 5.3.1 Firmware Upgrade, see [Repair a StorNext Connect System After Firmware Upgrades on page 30](#).

Here are the steps you will need to download and install a new StorNext Connect Connector:

1. Download the StorNext Connect Connector. Do one of the following, depending on your StorNext Connect system's connection to the Internet:
 - [To Download the Connector on a StorNext Connect System With an Internet Connection on the next page](#)
 - OR
 - [To Download the Connector on a StorNext Connect System Without an Internet Connection on the next page](#)
2. Install the new StorNext Connect Connector. See [Install the new StorNext Connect Connector on the next page](#).

To Download the Connector on a StorNext Connect System With an Internet Connection

To download the new Connector from the StorNext Connect App Store when the StorNext Connect system is connected to the Internet:

1. Log into StorNext Connect as an administrator.
2. Go to App Store.
3. Click **Administer** on the **Discover Components** tile.
4. Click the **Download** button next to the “Available Connector.”
5. Continue to [Install the new StorNext Connect Connector below](#).

To Download the Connector on a StorNext Connect System Without an Internet Connection

To download the StorNext Connect Connector when the StorNext Connect system is not connected to the Internet:

1. Log into the StorNext Connect web site at:
<https://stornextconnect.quantum.com/>
2. Click the **Updates** button for the system running StorNext Connect.
3. Under **Offline items available for download**, click **Discover content** and download the tar.gz file.
4. Log into the StorNext Connect system as an administrator.
5. Go to the App Store.
6. In the Connect Apps tile, click **Install File**.
7. In the Connect Apps screen, click **Choose file**.
8. Select the tar.gz file you downloaded earlier and click **Open** to unpackage the file.
9. Continue to [Install the new StorNext Connect Connector below](#).

Install the new StorNext Connect Connector

1. Click **Upload and install file**.
2. Install the new StorNext Connect Connector on the appliance being upgraded:
 - a. Click the StorNext Connect logo in the upper left corner to go to the application tiles.
 - b. Click **Discover Components**.
 - c. Click the “Add or update Connectors” link.
 - d. Click the “update” link under the **Connector action** column for the server node you are updating. Follow the on-screen instructions to complete the update.

Repair a StorNext Connect System After Firmware Upgrades

If you do not install the Connector before you applied the M440, M660, Artico, or Pro Foundation StorNext 5 Release 5.3.1 firmware upgrade, Node 1 of your appliance pair will no longer be managed by StorNext Connect. If this happens, you must take the following steps on Node 1:

1. Open an SSH connection to the MDC node using either IP address **10.17.21.1** (Node 1) or **10.17.21.2** (Node2) on the Metadata network.

Note: Use the IP addresses assigned if different from the defaults used here.

2. Enter **stornext** for the username at the prompt.
3. Enter the password for the **stornext** user account. The default password is “password”, but may have been changed after initial configuration.
4. At the command prompt enter the following to gain root user access:

```
sudo rootsh
```

5. Enter the password for the **stornext** user account again when prompted.
6. Press Enter.
7. Determine the IP of the StorNext Connect system and ensure it can be reached by Node 1. If you are unsure, you can log onto Node 2 and check the **/opt/quantum/connector/etc/connecthostname** file on that system.
8. Edit **/opt/quantum/connector/etc/connecthostname** and change the IP address (10.1.1.2) to the IP of your StorNext Connect system.
9. Adjust the mintd configuration by running the following command, substituting your StorNext Connect IP for W.X.Y.Z.

```
/opt/quantum/mintd/mintd_control.py set --minthost W.X.Y.Z
```

Contacting Quantum

Contacts

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

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

Getting More Information or Help

For further assistance, or if training is desired, 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/index.aspx>

Comments

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

doc-comments@quantum.com