

StorNext 6.1.0 Release Notes

Contents

What's New in StorNext 6.1.0	2
StorNext Compatibility	6
Supported StorNext Upgrade Paths and Upgrade Considerations	7
Compatibility Between StorNext and Other Products	8
General Considerations	9
Upgrading Appliances	10
Appliance Release Notes	10
Known Issues	11
Contacting Quantum	24

© 2018 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, FlexSync, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.

What's New in StorNext 6.1.0

Purpose of this Release

The StorNext 6.1.0 release delivers support for a pair of new [non-expiring FlexTier licenses](#), plus software fixes that correct the Spectre and Meltdown security vulnerabilities (variants) listed in the tables below.

Software-only Configuration			
StorNext Release	Operating System	Kernel	Variant Addressed
6.0.6.1	Red Hat EL / CentOS 6 Update 9	2.6.32-696.30.1	1,2,3,4
	Red Hat EL / CentOS 7 Update 4	3.10.0-693.21.1	1,2,3
6.1.0	Red Hat EL / CentOS 6 Update 9	2.6.32-696.30.1	1,2,3,4
	Red Hat EL / CentOS 7 Update 5	3.10.0-862.11.6	

Note: Software-only configurations must get BIOS fixes from their hardware vendor.

Caution: StorNext 6.0.6.1 does not work when used with Red Hat EL / CentOS 7 update 5. If you are an administrator of a non-StorNext Appliance system, then you must be careful when upgrading a Red Hat EL / CentOS operating system. You must only upgrade to levels supported in the [StorNext 6 Compatibility Guide](#).

Quantum Appliance Configuration			
StorNext Release	Operating System	Kernel	Variant Addressed
6.0.6.1	Spectre and Meltdown vulnerabilities are not addressed in StorNext 6.0.6.1.		
6.1.0	Red Hat/CentOS 6 Update 9	2.6.32-696.20.1	1,3
	Red Hat/CentOS 7 Update 4	3.10.0-693.21.1	1,2,3

See [Fixed Issues and Enhancements Addressed in StorNext 6.1.0 on page 5](#).

Important Information Regarding Upgrades to StorNext 6.1.0

WARNING: To implement some functionality and performance improvements to the Metadata Archive (**mdarchive**) in the StorNext 6.2 release, the schema for the database used by **mdarchive** has changed. As a result, upgrading to StorNext 6.1.0 causes FSMs to build new metadata archives, removing all metadata history generated prior to the upgrade. This means that tools **snhistory** and **snaudit** will not report any events that occurred prior to the upgrade and **snrecover** (not to be confused with **fsrecover** for managed file systems) will not recover any files deleted prior to the



New Features and Enhancements in StorNext 6.1.0

Enhancements to the fsfilecopy Command

The duplicate a file copy feature provides the capability to store another copy of a file. This feature is useful when the default copies for a policy class are increased and the desire is for existing files in the policy to take on the new setting. You can use the **fsfilecopy** command to invoke this capability.

The new **-n** option indicates the new copy to create. The specified new copy cannot exceed the policy class **max copies** setting. It is only allowed to be one larger than the current default copies setting for each file being processed. For example, if a file has default copies of 2, then the new copy value can be no larger than 3, so copies are not skipped.

For the new copy operation, logic was added when processing error conditions for a file to determine if they are fatal or if the command should continue processing. Most conditions were considered to be non-fatal. A summary of these conditions with the number of files falling into each category is displayed just before the processing begins to make the new copy of truncated files.

Below is a sample of the summary:

```
FS0355 04 0009486129 fsfilecopy interim: File scan summary
FS0355 04 0009486129 fsfilecopy interim: 4 files already store candidates.
FS0355 04 0009486129 fsfilecopy interim: 3 files the copy already exists.
FS0355 04 0009486129 fsfilecopy interim: 1 file belonged to another class.
FS0355 04 0009486129 fsfilecopy interim: 1 file had no content.
FS0355 04 0009486129 fsfilecopy interim: 1 file is a directory.
FS0355 04 0009486129 fsfilecopy interim: 2 files didn't exist.
FS0355 04 0009486129 fsfilecopy interim: 1 file not in archival directory.
FS0355 04 0009486129 fsfilecopy interim: 1 file exceeded max path length.
FS0355 04 0009486129 fsfilecopy interim: 100 files will be copied.
```

The new limit (**-L**) option allows you to throttle the request by limiting the number of I/O requests that are generated.

i Note: The default limit of I/O requests you can generate is 10. For requests using tape devices, this is the number of queued I/O requests. Tape I/O requests are processed sequentially, so there is no benefit in adjusting this value. Quantum recommends you use the default setting.

For requests that do not involve tape devices, limiting the number of I/O requests also restricts the number of streams that you use. The number of streams that you can use is the lesser of the limit value or the max streams defined for the device. The stream limit is applied to each source and destination device so the

specified limit value is used for each. For example, if you specify **-L 3**, then 6 streams are used; 3 streams are used for the source and 3 streams are used for the destination.

However, if you define the max streams for one of the devices to be smaller, then the request is further restricted. Using the same example, if you configured the source device with max streams of 2, then only up to 2 streams are used for each source and destination devices.

i Note: When you set this value for non-tape devices, keep in mind that **fsfilecopy(1)** requests run at a lower priority than store and retrieve requests, so the requests only use available streams when there are no outstanding store/retrieve requests requiring those resources.

i Note: Prior to StorNext 6.1, you could only allow a list of files to be specified. Beginning with StorNext 6.1, you can include options to operate recursively on a directory or on a batch file containing a list of files.

For additional information, refer to the **fsfilecopy(1)** command in the [StorNext 6 Man Pages Reference Guide](#).

Compatibility and Support

The [StorNext 6 Compatibility Guide](#) provides the basic compatibility for StorNext 6.1.0, 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](#).

- **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 6.1.0

Operating System	Change Request Number	Service Request Number	Description
All	65301	293344, 332924	cvmkfs hangs if DNS is not reachable
All	66283	3748382, 414432	Proxy Client cvfsbufiod spinning in PickPxcon() after adding stripe groups
All	62209	3636874, 414352, 429898	security scan - Autocomplete Not Disabled
All	71237	3584660	provide capability to make a new copy from an existing copy
All	72070	439940	Retrieves from a 100% full non-tape media fail
All	71908	435562 429352, 440980	fs_copymand core caused by fsqueue query due to timing issue dumping the queue as it is cleared
All	71534	423796	xdiamtask core due to unknown check condition
All	71493	423302, 443878, 444771, 445912	QRDB corruption at 2 billion object IDs due to int32 overflow
All	71491	422731	unresolvable hostname in nss_cctl.xml control file permitted retrieves on clients although denyRetrieves value = true
All	71413	419495, 445375	GUI:File and Directory Actions - Retrieve Directory not working
All	71296	417960, 346928, 411188	Multiple OpHangLimitSec panics
All	71061	4020160, 427408	Upgrade issues resulted in removal of TSM libmsc.so library and problems starting TSM fs_resourced daemon
All	71243	416889, 416858, 423186, 426962	fs_resourced segfault occurred upon fsmedin, leading to TSM abnormal termination
All	71212	414396	SNFS: snaudit shows wrong host IP on some rm actions
All	71240	408782	fsimport will not import files with some special characters in them

Operating System	Change Request Number	Service Request Number	Description
All	71102	407649, 403677	snbackup failed with "ERR: Failed to store copy" even though the copy was successful
All	61553	384152	fs_resourced can core and halt TSM when fsmecopy is run
All	71179	380802, 362953	FSM panics for file names with non-ASCII characters
All	70294	378110	snaudit may use all available memory of system on busy system until the kernel kills it
All	70224	376518	fs_moverd: not accessible/working due to bad /etc/hosts value
All	68141	344165	vidio tools for the Mac
All	67827	415475	in the fsm, whenever using a pclient structure, we should make sure it is valid ...
All	71291	n/a	Update fsfileinfo v2 web service call to support object url, encryption and compression information
Linux	71230	414352, 429898	Security scan reports HTTP Header missing on port 443
Linux	71229	414352, 349865	Security scanner showed StorNext GUI allowed connection on TLSv1.2 with DES cipher
Windows	70846	433642	StorNext Gateway Server may cause a system hang on Windows systems

StorNext Compatibility

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

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

Quantum Operating System Upgrade Support Policy

StorNext supports any security or functional bug update that applies to the current StorNext-supported Red Hat update level or SUSE Linux Service Patch. StorNext does not support upgrading to an update level or

service patch beyond the currently supported levels shown in the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

StorNext and Linux Interoperability

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

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

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

Recommended usage:

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

Supported StorNext Upgrade Paths and Upgrade Considerations

StorNext Software Upgrade Matrix

For information on which StorNext versions allow you to upgrade directly to this release, refer to the **StorNext Software Upgrade Matrix** section in the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

Considerations for the StorNext File System Directories


On upgrades to StorNext 6.1.0, 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 6.1.0 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 6.1.0

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.1.0, special consideration must be made for Linux NFS servers.

To prevent issues with mounted NFS clients, NFS clients must be unmounted prior to upgrading StorNext on the NFS server. If unmounting all NFS clients is not an option during the upgrade, Quantum suggests using the "compat32" mount option on NFS servers.

Compatibility Between StorNext and Other Products

The following sections provide information regarding compatibility between this release and StorNext components and features.

Appliance Controller

To view supported Appliance Controller software configurations, see *Appliance Controller Compatibility* available online at <http://qsupport.quantum.com/kb/Flare/Content/appliances/ACC/DocSite/Compat.htm>.

Infiniband

StorNext 6.1.0 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

Lattus

See the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus and StorNext 6.1.0.

 **Note:** 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 6.1.0, 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 6.1.0, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

Supported Browsers

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

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

General Considerations

This section provides information about items to consider for StorNext 6.1.0.

Checksum Performance Considerations

i Note: Generating MD5 checksums is a CPU-intensive operation.

Current StorNext metadata controller and Mover hardware is able to calculate MD5 checksums at around 300 MB/s to 500 MB/s. For newer generation tape technology, the maximum throughput might exceed the rate at which the system can generate checksums. In this case, the MD5 checksum calculation will define the throughput of a single data movement operation. With multiple movement streams, MD5 calculations will be done in parallel across the streams and aggregation of performance will be seen.

Upgrading Appliances

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* (www.quantum.com/ApplianceInfoHub).

Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- [Xcellis Foundation](#)
- [aiWARE for Xcellis](#)
- [Xcellis Workflow Extender](#)
- [Xcellis Workflow Director](#)
- [Artico](#)
- [Pro Foundation](#)
- [G300](#)
- [M660](#)
- [M440](#)

Known Issues

The following sections list known issues in this release of StorNext, as well as associated workarounds, where applicable:

- i Note:** If you encounter one or more of the issues listed in this section, please contact Quantum Customer Support and report the issue(s) you encountered. Also inform the support representative whether you were able to successfully work around the issue(s) by using the provided workaround. Doing these things will help Quantum prioritize the order in which known issues are addressed in future StorNext releases.

StorNext File System Known Issues

The table below lists known issues specific to the StorNext File System.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	54834	3505208, 3516356	<p>If a file is being copied to the StorNext file system using Windows Explorer and Windows Explorer crashes before it finishes copying all the data, the file might contain data blocks from old, deleted files. This problem occurs because Windows Explorer sets EOF to the size of the file before it writes the data to the file. This leaves a gap of uninitialized data in the file.</p> <p>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	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	70306	n/a	The snrecover utility only recovers files and directories that have been previously deleted. That implies that it cannot be used to find files within an existing directory and then recover some or all deleted contents within the directory. For example, the following removes all contents from directory bar and then attempts to recover the deleted contents:
-----	-------	-----	---

```
# rm -rf /stornext/snfs1/bar/*
# ls -l /stornext/snfs1
total 0
drwxr-xr-x 4 root root 0 Nov 13 13:34 bar
drwxr-xr-x 2 root root 0 Nov 13 13:34 lost+found
# snrecover -p /stornext/snfs1/bar -t 2017-12-03:04:05:06
```

In the example, **snrecover** is attempting to recover a directory that exists (**bar**). This is not a valid **snrecover** request so **snrecover** will return an error. If a directory exists, then it cannot be recovered. The deleted content of directory **bar** can be recovered using successive calls to **snrecover**, one for each deleted file or directory.

Workaround:

There is no workaround for this issue, as the **snrecover** utility is working as designed. The utility cannot be used to search a subtree to find files and directories to recover. As the **snrecover** man page shows in the example section, you can use **snhistory** to find files and directories that were removed in a subtree of the file system and then use **snrecover** to recover them.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	71147	n/a	<p>If a client is not directly connected to the metadata network, then the file systems are not visible to that client. As a result, the client does not have access to mount any file systems.</p> <p>A client is directly connected to the metadata network when one of its interfaces is in the same subnet as the metadata network. The metadata network is defined by the addresses of the name servers and is configured within the fsnameservers file.</p> <p>Workaround:</p> <p>To correct this, use the foreign servers feature of StorNext to locate file system services. This feature allows a client to connect directly to a file system service by placing the address of the MDC (not the name service, although they can be the same) in the fsforeignservers file.</p> <p>See Configuration > Name Servers, or the fsforeignservers command in the StorNext 6 Man Pages Reference Guide for more information.</p> <p>i Note: The foreign servers feature is supported on all StorNext client platforms.</p>
All	72233	n/a	<p>There is an issue when building the metadata archive, while clients are concurrently deleting files and directories. There is a very small possibility that the metadata archive build will fail to complete, if a client delete of a file or directory occurs at the exact same time that the FSM is attempting to add that file or directory to the metadata archive. When this occurs, the FSM will abort the mdarchive build and issue the following error message in the system log file:</p> <pre>Metadata archive scan failed, error 26!</pre> <p>Workaround:</p> <p>To correct this, restart the build of the mdarchive. Using the CLI, execute the following command:</p> <pre>cvadmin -F <fs_name> -e "mdarchive rebuild"</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
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>
Mac OS	67871	n/a	<p>macOS releases 10.12 through 10.12.6 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>
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> <pre>Error 2: 'Error starting Windows Service 'cvfsfilter'</pre> <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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Windows	70190	n/a	<p>If you are using a Windows MDC, you have one or more StorNext file systems configured in HA mode (haFsType set to HaUnmanaged) and do not have a valid ha_peer file, then any file systems configure in HA mode will fail to start when the MDC is upgraded to StorNext 6.0.5. StorNext 6.0.5 enforces having a valid ha_peer file for HA configured file systems.</p> <p>Workaround:</p> <p>To correct this, you can either create a valid ha_peer file or disable HA mode. For more information on configuring HA on a Windows MDC, see Configuring a High Availability System on Windows.</p>

StorNext Storage Manager Known Issues

The table below lists known issues specific to StorNext Storage Manager.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p>Workaround:</p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <ul style="list-style-type: none"> i Note: This workaround might also help improve the performance of the faster LTO drives like LTO-6, and LTO-7 by updating the FS_LTO_BLOCK_FACTORsysparm. i Note: Changes to FS_XXX_BLOCK_FACTOR only affects tapes formatted after the change. <ol style="list-style-type: none"> 1. Change the FS_T10K_BLOCK_FACTORsysparm from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre style="background-color: #f0f0f0; padding: 5px;">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_FACTORsysparm must not be set to a value that exceeds 32. 2. Restart Storage Manager to ensure the change in Step 1 goes into effect: <pre style="background-color: #f0f0f0; padding: 5px;"># tsmstop # tsmstart</pre> 3. Verify the FS_T10K_BLOCK_FACTORsysparm contains the new value:

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre data-bbox="683 331 1458 436"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <ol style="list-style-type: none"> <li data-bbox="618 478 1458 548">4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients. <li data-bbox="618 569 1458 674">5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows: <pre data-bbox="683 709 1458 856">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> <li data-bbox="618 898 1149 926">6. Unmount and re-mount your file systems. <li data-bbox="618 947 1458 1192">7. Use new T10K media to store a copy of the file from the disk. <ul style="list-style-type: none"> <li data-bbox="667 989 1458 1192">i Note: Step 7 is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which fsformat was run before the change will use the block factor in use at that time. This change will not impact those tapes.
All	46693	n/a	<p data-bbox="602 1220 1463 1360">Executing the command snbackup -s while a full or partial backup is running might result in a message that <code>/usr/adic/TSM/internal/locks/backup.1f</code> is in an invalid format.</p> <p data-bbox="602 1367 1463 1436">This is due to the snbackup -s process reading the backup.1f status file while the backup process is updating it.</p> <p data-bbox="602 1442 781 1478">Workaround:</p> <p data-bbox="602 1484 1463 1591">Ignore the message; to clear-up the process, re-execute the command snbackup -s (provided that the backup is not writing to the backup.1f status file while snbackup -s is trying to read it again).</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47833	n/a	<p>When copying files between media using the CLI command fsmedcopy, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p>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>
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in /usr/adic/TSM/logs/tac which contains the text:</p> <pre>Process fs_moverd on <host> timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.</pre> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Increase the database connection timeout value by adding the following line to /usr/adic/mysql/my.cnf under the section labeled [mysqld] connect-timeout=240. 2. Cycle the Storage Manager in order to pick up the updated timeout value.
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an fsmedread operation of a partial tape block from a full tape block, the operation can fail with errno=12.</p> <p>i 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	70991	n/a	<p>After importing an LTFS tape whose data was not originally written by StorNext, a multi-file retrieve command successfully retrieves the first file in the list, but fails to retrieve the remaining files.</p> <ul style="list-style-type: none"> This issue only affects LTFS tapes whose data was NOT originally written by StorNext. Tapes that were written by StorNext and subsequently exported and re-imported are not affected by this issue. This issue only affects multi-file retrieves (including recursive retrieves). Single-file retrieves (including event-driven retrieves) are not affected. This issue only affects files that were imported with the "media ingest" function of fsimport. Any files imported with the "file ingest" function are not affected because they are imported in whole to the disk file system, and the media is not kept in the system. <p>Workaround:</p> <ul style="list-style-type: none"> For files affected by this issue, retrieve files one at a time. When importing non-StorNext LTFS tape data, consider using the "file ingest" function of fsimport. This function ingests the tape data, but not the tape itself. The tape is ejected from the system upon completion of the fsimport operation.

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
All	69958	373823	<p>There is a known issue where stripe group expansion using the StorNext GUI can fail and requires manual intervention to restore file system operation.</p> <p>Stripe group expansion allows an additional disk (LUN) to be added to an existing stripe group, growing the file system user data space.</p> <p>The other way to add space to a file system is to create a new stripe group and add this to the file system.</p> <p>When the GUI executes stripe group expansion, it stops the file system, modifies the configuration file and runs cvupdatefs. If cvupdatefs takes more than 5 minutes to complete, the GUI kills the cvupdatefs utility and reports an error. At this point the file system does not start because the configuration file does not match the current state of the metadata.</p> <p>Workaround:</p> <p>To address this issue, you can do one of two things:</p> <ul style="list-style-type: none"> • Optimally, you would not attempt the stripe group expansion at all. Instead add a new stripe group to the file system. • If stripe group expansion is deemed necessary, use the cvupdatefs CLI directly instead of through the GUI. This method does not have any time limitations. <p>In the case that the StorNext GUI attempt was made and hit the 5 minute timeout, file system operations can be resumed by performing the following:</p> <ol style="list-style-type: none"> 1. Restore the previous version of the configuration file. This can be found in the following directory: <div data-bbox="678 1346 1458 1413" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre data-bbox="695 1360 1240 1394">/usr/cvfs/data/<fs>/config_history</pre> </div> 2. Run the cvfsck utility to verify and potentially correct the metadata. 3. Start the file system.

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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	70282	n/a	<p>The Stornext Connect Utilization App Version 1 does not recognize the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system, and does not incorporate capacity or data movement associated with the HGST ActiveScale™ P100 (Quantum Lattus P100) system as a target.</p> <p>i Note: This does not otherwise affect the functionality of the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system.</p> <p>Workaround: There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.</p>

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