

StorNext 6.3 Release Notes

Contents

What's New in StorNext 6.3	2
Supported StorNext Upgrade Paths and Upgrade Considerations	10
Compatibility Between StorNext and Other Products	11
General Considerations	12
Upgrading Appliances	13
Appliance Release Notes	13
Known Issues	14
Contacting Quantum	30

© 2019 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 6.3

Purpose of this Release

The StorNext 6.3 release provides new features and enhancements listed in the section [New Features and Enhancements in StorNext 6.3 below](#), and also provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 6.3 on page 6](#).

New Features and Enhancements in StorNext 6.3

Support Long Object IDs on Object Import

Prior to StorNext 6.3, objects imported with the command **fsobjimport** were restricted to those with a key length that did not exceed 64 bytes. Beginning with StorNext 6.3, we enhanced the tertiary storage manager (TSM) infrastructure to support object keys with a larger length, up to 1,024 bytes.

Read Stream Reservation Capability

Beginning with StorNext 6.3, you can reserve a subset of the object storage streams for file retrieves.

Use the StorNext GUI to reserve streams for reads when you configure an object storage controller. You can configure the read streams in the **Controller** section of the **Configuration > Storage Destinations > Object Storage** menu. See [Configure Object Storage and Cloud Destinations](#).

You can also use the command line interface (CLI) to run the **fsobjcfg** command with the **-s** option to configure a read stream count. You can exceed the number of streams used for retrieves, but you cannot use the reserved streams for non-retrieve requests.

See the **fsobjcfg** command in the [StorNext 6 Man Pages Reference](#) Guide.

Enhancements to the fsCapacityThreshold Configuration File Parameter

Beginning with StorNext 6.3, StorNext automatically monitors the used capacity of the HA shared file system and generates a RAS event when the file system reaches a capacity of 85% or higher.

i Note: StorNext continues to generate a Severity 2 RAS event once an hour until the used capacity level falls below 85%.

The percent at which a RAS event is generated is controlled by the HA shared file system's configuration file parameter **fsCapacityThreshold**.

i Note: If the **fsCapacityThreshold** parameter is not present in the configuration file or if it is set to zero (0), then the default of 85% is used.

See [How to Configure the fsCapacityThreshold Configuration File Parameter](#) for additional details.

New Extended Attributes

Beginning with StorNext 6.3, a series of new extended attributes have been added, as follows:

Extended Attribute	Description
user.com.quantum.retrieves	This attribute can be read or set . It determines whether a client has the ability to retrieve files or not. To configure, you must write the value enabled (allow retrieves) or disabled (deny retrieves).
user.com.quantum.snsnm	This attribute can be read . It provides the current state of a managed file and also provides different descriptions: <ul style="list-style-type: none"> • disk+archive • store • no_store • archive
com.quantum.rpl	This attribute can be read . It provides a reverse path lookup of a given file.
user.com.quantum.location	This attribute can be read . It provides a json response that includes the media type backing a managed file and also provides the media UUID of the media the file is written to when it is truncated. If the file is an object store file, the object store path is given. If an affinity is assigned to the media of the file, the affinity is included in the json response.
user.com.quantum.offline	This attribute can be read . It returns a true or false response to whether a file is offline.
user.com.quantum.clients [.classname]	This attribute can be read . It returns a json response of client information that have been mounted with a specific mount option.
user.com.quantum.sn_rest_uri	This attribute can be read . It returns the URI for the file system manager's (FSM's) rest interface.
user.com.quantum.sn_rest_challenge	This attribute can be read . It returns the challenge string for the FSM's rest interface.
user.com.quantum.blocklocations	This attribute can be read . It returns the stripe group number of a file's location.
user.com.quantum.affinity	This attribute can be read . It returns the affinity associated with a file.

On systems running Linux, you can run the command **getfattr** to access or write to extended attributes.

Example

```
getfattr -n user.com.quantum.location /path/to/file
```

On systems running macOS, you can may use **xattr** command line tool.

i Note: You must enable **namedstreams** on the file system and you must also remove the user namespace.

Example

```
xattr -p com.quantum.retrieves /path/to/file
```

On systems running Windows, you can use the CvApi external API to develop a command line tool for Windows.

CvApi Functions	Description
CvApi_GetXattrByName(int fd, XattrRequest_t *xreq)	<p>This API takes a file descriptor and a XattrRequest_t pointer. The XattrRequest_t structure consists of two buffers:</p> <ul style="list-style-type: none">• xattr_name• xattr_value <p>The xattr_name is confined to 117 bytes and the xattr_value is confined to 4,096 bytes.</p> <p>The xattr_value should not have any data written to its buffer.</p> <p>By calling this function, the xattr_value buffer is filled and returned to the user.</p>

CvApi Functions	Description
CvApi_SetXattrByName(int fd, XattrRequest_t *xreq)	<p>This API takes a file descriptor and a XattrRequest_t pointer. The XattrRequest_t structure consists of two buffers:</p> <ul style="list-style-type: none"> • xattr_name • xattr_value <p>The xattr_name is confined to 117 bytes and xattr_value is confined to 4,096 bytes.</p> <p>This function only works with xattrs that are able to be written to.</p> <p>Both buffers must have data in order to properly use this function.</p> <p>The structure is returned back to the user with the same data it was supplied with.</p>

To build a program with Quantum's external API, you must include **extapi.h** and **cvapi.h** in the source. You must also include **cvextapi.lib** when you build the program for Windows.

By using these functions, a Windows user can emulate the use of extended attributes without them being an interposed component of the file system and Windows.

i Note: Not all of the extended attributes are available for use through Windows and not all of the attributes are able to be read or written depending on the user requesting and the settings on the file.

Compatibility and Support

The [StorNext 6.3 Compatibility Guide](#) provides the basic compatibility for StorNext 6.3, including the StorNext components supported, operating systems and service packs, libraries and drives, browsers, virtual machines, and appliance support. Listed below are just a few of the types of information available to you in the [StorNext 6.3 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 6.3

Operating System	Change Request Number	Service Request Number	Description
All	32784	474686, 1223902, 1197170, 1286816, 1559740, 3545946, 3642968	Excessive tac log messages - no media found to meet user criteria
All	62966	3681170, 344347, 478334, 480193	MSM can wrongly identify the running process, leading to system crash
All	65786	503315	Update TSM infrastructure to support long/alternate object names
All	68878	n/a	tac log shows mysql connection error when querying for Qcloud configuration
All	70955	399606 489960	FSM OpHangLimitSecs exceeded VOP-Class-0 Type-4 Subtype-11RevokeWait -> set_wait_racing
All	71103	402207, 504661, 506855	Need to monitor used capacity of HA shared file system more closely
All	73252	n/a	provide streams reservation capability for object get operation
All	73295	424025	Security Scanner cause FSM on secondary MDC to die Weekly
All	73322	463088	HA shared FS expansion via Stripe Group Actions leads to non-recommended configuration, when HA shared FS has a mixed metadata and data SG and no exclusive MD SG
All	73674	478052	Name service attempts to send udp messages to port zero on startup
All	73676	501660	ANTF & LTFS tape fsexport failure after database error

Operating System	Change Request Number	Service Request Number	Description
All	73896	n/a	GUI: Support for read stream reservation count in controller configuration
All	74108	481323	Found several minor issues in wsar_client with the -s option, -l option, and documentation
All	74111	482099	attempting to start second instance of snstatd from command line leads to qustat failures
All	74139	474784	fsm memory bloating: filesystem hit OOM upon starting servicing change_tree REST call from flexsync
All	74164	482099, 502217	fs_resourced locks up and requires restart of TSM
All	74169	477772	fsmedcopy/fsfilecopy from tape to S3-compatible medium fails and marks it as write-protect
All	74174	484016, 494543, 504800	StorNext should warn when metadata space is low
All	74175	484016	cvupdatefs failed to add a new metadata stripe group because there were 0 free inodes
All	74228	495024	'/usr/cvfs/lib/snnas_control start' does not start snnas_service
All	74231	479813	syncha can delete all local config files if the shared filesystem is dismounted while it is running
All	74290	488616, 493109	NAS beyond 2.2.0 and SNFS work -- Windows Directory notifications (samba <-> SNFS) not delivered.
All	74453	490072	fs_feature core dump due to segmentation fault when running archive_cmp
All	74476	480485	Non-existent or truncated display of existing file system quotas in the snquota CLI and StorNext GUI
All	74502	n/a	Add FS_CLUSTER_LIMIT_XXX to fs_syparm.README file
All	74509	491986	syncha too noisy: reports "unable to determine hostdatadir for peer" in single mode
All	74559	490052	Incorrect name server info messages when using fsforeignservers
All	74609	477352	A relocate policy will not clean all candidates that it should

Operating System	Change Request Number	Service Request Number	Description
All	74618	487263	D2D Relocation - Nightly class based relocation does not work if tiers are named out of order
All	74623	490052	StorNext should alert Customer if current mounted file system volumes are out of date and need to be rescaned
All	74629	493309	fsmpm: PANIC: fsmpm ASSERT failed "ntohs(pmsg->nsm_data16) & NSS_HB_EXTENDED_OK"
All	74632	487263	Fix relocation-by-class policy so that it doesn't overwrite its policy start time
All	74633	487263	Fix fs_tierman so that it honors the configured policy class Reloc Min Time when it is a value less than 1 day
All	74644	496096	recover_xattr_payloads() does not honor snbt_lookup() API causing snbt_payload_expect_dupes():ASSERT(FALSE)
All	74645	494879	kernel panic w/ kernel BUG at fs/dcache.c:667, exception in shrink_dcache_for_umount_subtree() during cvfs umount
All	74646	492529	wsar_agent segfault occurred due to unhandled null pointer exception in wsar_db_CleanJobs()
All	74652	n/a	FS_BLOCKING_READS setting not working
All	74665	498164	immutable bit on root prevented starting file system because qbm create/remove of directory
All	74796	501350, 474104	snaudit "Failed to read tablet cluster" file /scm/nightly/VM-RedHat7-x86-64-SP3-0/sn/snfs/qrtree/qrtree_backend_qrdb.c, line 1058
All	74808	493941	Reading data from Media with large clusters can lead to issues finding the BOF and report Invalid label data
All	74822	498657	After upgrade from 6.1 to 6.2 ala report main: Database error during insert: code: 1406, message on insert into sl_admin_log
All	74842	502251	GUI: deprecated media types in MSM database can cause vault create failure in UI after upgrading to 6.2.0
All	74869	493530	Remote DDM processes fail and generate cores after upgrade to StorNext 6.2.0
All	74871	502252, 497249, 503172, 498392	CPU or system lockup when kernel oplocks enabled in NAS

Operating System	Change Request Number	Service Request Number	Description
All	74873	503173	fsm segmentation fault upon sgmanage resize, due to unhandled exception in move_remainder_space()/inode_btree_start_slice()
All	74889	499858	StorNext close locking causes nfs server to consume millions of file entries
All	74890	503808	fsm panic "OpHangLimitSecs exceeded" whilst vacating stripe group in progress, due to deadlock between free_pending_inode and vacate_stripegroup threads
All	74891	505693	sgmanage displays the wrong size for several counters when the fs block size is not 4k
All	74892	503808	fsm panic upon ASSERT failed "sl < pgroup->pg_slices" snfs/fsm/alloc.c, line 3058, after a stripe group offload/vacate was interrupted
All	74902	n/a	TSM checkMediaAvailabilityTsm health check script must detect non-tape store availability
All	74915	488583	dm_info attributes not set correctly after a file move. endtime of old filecomp not reset
All	74948	503993	snrestore fails due to duplicate entries (but with different media) in the snbackup_manifest file
All	74952	n/a	The new scan functionality being used in fsaddrelation and rebuild policy can hang.
All	74986	505693	sgadd miscalculates stripe breadth when adding a stripe group to a file system with an original block size other than 4k
All	75032	501660	GUI : Export Files job report incorrect : job overall status is reported as success, although errors were encountered
Linux	73452	460749, 465920	Stopping managed FSM hangs Stornext GUI
Linux	74545	501127, 493912, 501127	Ubuntu Kernel 4.4.0.145 new get_user_pages() signature, cvfsbuild fails error: too many arguments to function 'get_user_pages'
Linux	74687	497732	Support for CentOS/RHEL 7.6

Operating System	Change Request Number	Service Request Number	Description
Linux	74809	501686, 507321	snprobe not returning with options -cdlqsC -h localhost  Caution: StorNext 6.2 included new cloud-based monitoring tools to simplify and streamline system management when combined with Web Services. In certain situations, Cloud Based Analytics (CBA) reporting generates repetitive snprobe returns that never complete, resulting in system memory being depleted. If this occurs, Xcellis systems might unexpectedly reboot. See StorNext Product Bulletin 105 for additional details.
Windows	46895	3353988	SNFS latency-test give extremely bad performance on Windows node.
Windows	72817	454915	Stornext Windows GUI need to allow the creation of fsforeignservers
Windows	73687	455957 505429	IsLocalSid() Requires that Administrator account to exist.
Windows	74181	479521, 487047	max length russian characters
Windows	74583	480485	Photoshop deletes the original file if Quota is exceeded - FileSetAttr() problem

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 6.3 Compatibility Guide](#) in the [StorNext 6.3 Compatibility Guide](#).

Considerations for the StorNext File System Directories


On upgrades to StorNext 6.3, 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 6.3 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 6.3

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.3, 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

Infiniband installations require assistance from the Quantum Professional Services team, a Service Partner, or a Quantum Service Provider. For additional information, contact [Quantum Technical Support](#).

Lattus

See the [StorNext 6.3 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus and StorNext 6.3.

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

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 6.3, refer to the [StorNext 6.3 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 6.3 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6.3 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

General Considerations


This section provides information about items to consider for StorNext 6.3.

Checksum Performance Considerations

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

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

Upgrading Appliances

 **Caution:** If you have a Replication, Deduplication, or Object Storage license, see Change Request [Known Issues on the next page](#) in the [StorNext Installation, Replication, HA, and Other Known Issues on page 28](#) section before you upgrade.

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

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	74202	n/a	<p>The command <code>/usr/cvfs/lib/snnas_control stop</code> does not stop the snnas_service.</p> <p>There are two basic scenarios, as follows:</p> <p>Scenario # 1</p> <p>When StorNext stops, it signals NAS to vacate the StorNext file systems that NAS is using. However, the snnas_controller attempts to automatically restart NAS services if these file systems remain available for too long.</p> <ul style="list-style-type: none"> • The restarting of the NAS services can resume use of StorNext resources and interfere with the operation of stopping StorNext. • The interference can result in mounted StorNext file systems that do not have a running fsm process, thereby blocking access to the StorNext file system. • In the event that StorNext fails to stop due to NAS resource usage, you might need to perform the following steps: <ol style="list-style-type: none"> 1. Check to see if StorNext file systems are mounted: <pre>grep cvfs /proc/mounts</pre> 2. Mount the HA shared file system: <pre>mount /usr/adic/HAM/shared</pre> 3. Start the fsm for each StorNext file system mounted on the MDC: <pre>cvadmin -e "start FSNAME on localhost"</pre> <p>Note: You might need to reboot the MDC if StorNext still fails to stop after restarting the required fsm processes.</p> <p>Scenario # 2</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>Maintenance Mode. An extension of Scenario #1, where NAS must vacate the file systems for an extended period of time.</p> <p>By resuming NAS services, those service processes may interfere with maintenance operations.</p> <p>To affect both scenarios, timers exist that you can adjust to extend the amount of time required between the stornext stop operation and before NAS services resume. The timers are controlled by the following Controller Registry values:</p> <ul style="list-style-type: none"> • nas.heartbeat.check_state_secs • stornext_service.stop_period <p>You can modify the timers by using the Controller shell interface reg commands.</p> <p>For example, to observe existing values, execute the following commands:</p> <pre data-bbox="626 953 1458 1024">su sysadmin -c 'reg show nas.heartbeat.check_state_secs'</pre> <p>or</p> <pre data-bbox="626 1146 1458 1218">su sysadmin -c 'reg show stornext_service.stop_period'</pre> <p>Both timers default to 120 (seconds).</p> <p>For example, to set new values, where X is an integer, execute the following commands:</p> <pre data-bbox="626 1423 1458 1495">su sysadmin -c 'reg set nas.heartbeat.check_state_secs X'</pre> <p>or</p> <pre data-bbox="626 1617 1458 1688">su sysadmin -c 'reg set stornext_service.stop_period X'</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>Workaround:</p> <p>For Scenario #1, Quantum recommends that you extend the time period to 10 minutes, a value of 600 for those two variables. The goal is to set a value that exceeds the amount of time necessary for typical StorNext shutdown.</p> <p>For Scenario #2, choosing an X value requires knowing how long maintenance will take and staying under that time. It is difficult to extend the maintenance period once its started. To avoid the complexity of controlling the maintenance period, Quantum recommends you perform the following steps:</p> <ol style="list-style-type: none"> 1. Stop NAS by executing the following command: <div data-bbox="678 793 1458 867" data-label="Code-Block"> <pre>/usr/cvfs/lib/snnas_control stop</pre> </div> 2. Stop the Controller by executing the following command: <div data-bbox="678 989 1458 1094" data-label="Code-Block"> <pre>/usr/local/quantum/bin/sml_service_tool stop snnas_controller</pre> </div>
All	75140	n/a	<p>Exporting an SNFS file system on Ubuntu releases 16.04.2 or later is not supported.</p> <p>Workaround</p> <p>There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.</p>
macOS	66948	322824, 336945	<p>If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client.</p> <p>Workaround</p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>

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

Windows	74339	488689	Your Windows SAN client might become unresponsive and crash with bugcheck error DPC_WATCHDOG_VIOLATION (133) .
---------	-------	--------	---

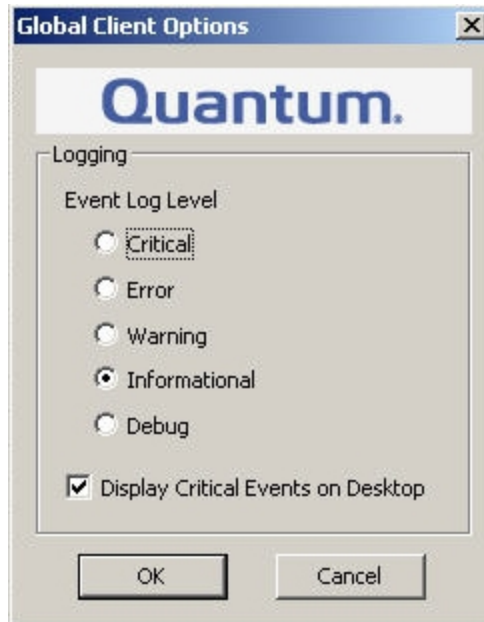
There are two conditions which lead to this issue:

- Your client is performing a reconnect to the FSM.
- Your **Global Client Options** event log level is set to **Debug**.

Workaround

Although you cannot control the reconnect to the FSM, you can change the **Global Client Options** event log level to **Informational**.

1. On the **Client Configuration** dialog box, click **Tools**, and then click **Global Options**. The **Global Client Options** dialog appears.



2. Select **Informational**.
3. Click **OK**.

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 by updating the FS_LTO_BLOCK_FACTOR system parameter. 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 system parameter from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>FS_T10K_BLOCK_FACTOR=32;</pre> </div> <ul style="list-style-type: none"> i Note: The T10K default I/O block size is 512 KB or 8 * 64 KB. With the block factor changed to 32, the new T10K I/O block size will be 2 MB or 32 * 64 KB. Presently, the FS_T10K_BLOCK_FACTOR system parameter must not be set to a value that exceeds 32. 2. Restart Storage Manager to ensure the change in Step 1 goes into effect: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre># tsmstop # tsmstart</pre> </div> 3. Verify the FS_T10K_BLOCK_FACTORsysparm contains the new value:

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre data-bbox="683 331 1458 436"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <ol style="list-style-type: none"> <li data-bbox="618 478 1458 548">4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients. <li data-bbox="618 569 1458 674">5. Modify <code>/etc/fstab</code> on the MDCs and the DDM clients to use the <code>auto_dma_write_length</code> and <code>auto_dma_read_length</code> mount options as follows: <pre data-bbox="683 709 1458 856">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> <li data-bbox="618 898 1149 926">6. Unmount and re-mount your file systems. <li data-bbox="618 947 1458 1192">7. Use new T10K media to store a copy of the file from the disk. <div data-bbox="683 989 1458 1192"> <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> </div>
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>

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	72993	452722	<p>When DDM is enabled for non-primary systems, kernel error messages are logged for reservation conflict because the Primary MDC mounts the tape which sets the reservation to itself and the DDM sets the reservation to itself once the tape is ready.</p> <p>This issue causes a benign message in the <code>/var/log/messages</code> file for reservation conflict. Since the unmounting of the tape also resets the reservation back to the primary MDC, this message is generated the next time a DDM (non-primary MDC) accesses a tape.</p> <p>i Note: This issue can result in a large amount of log messages on any machine running DDMs.</p> <p>Workaround:</p> <p>To workaround this issue, you can filter and drop the messages so they are no longer logged, as follows.</p> <p>i Note: Reservation conflicts that are a problem also produce RAS alerts, so you can also drop these errors.</p> <p>Create the following rules on all DDM clients, based on the OS, to drop the messages from rsyslog:</p> <p>For RedHat 6</p> <pre># echo ':msg, contains, "reservation conflict" ~' > /etc/rsyslog.d/ignore-reservation-conflict.conf # service rsyslog restart</pre> <p>For RedHat 7</p> <pre># echo 'if \$programname == "kernel" and \$msg contains "reservation conflict" then stop' > /etc/rsyslog.d/ignore-reservation-conflict.conf # systemctl restart rsyslog.service</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	75127	n/a	<p>When fsexport in file copy-export mode fails for an entire source media id, then you cannot determine the list of files to be copied from each failed source medium.</p> <p>Workaround:</p> <p>i Note: The workaround applies ONLY to those source media where the entire source media fails to copy; in other words, ALL files on that source medium fail to copy. The source medium is reported as a failure, but not each file to be copied from that medium.</p> <p>For other source media where some files succeed to copy and other files fail, the individual failures ARE already reported to console. That is why the workaround applies ONLY to those source media that completely fail.</p> <p>Do the following to determine the list of files that failed to copy:</p> <ol style="list-style-type: none"> 1. Add the following line to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <pre>DBG_EXPORT_DISABLE_REMOVE_TRACKING_FILES=true;</pre> 2. Re-run the failed fsexport command in report mode (add the <code>-r</code> option). <ol style="list-style-type: none"> a. If you are unsure of the exact command for fsexport, examine the following file: <pre>/usr/adic/TSM/history/hist_01</pre> b. Select the correct fsexport command. c. Call this command orig-fsexport-cmd. d. Execute the same command, adding the <code>-r</code> option for report mode: <pre>orig-fsexport-cmd -r</pre> 3. For each failed source media id <mediald> from the original

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

fsexport do the following:

- a. To obtain a list of the file name containing all of the failed paths, execute the following:

```
ls -1rt /usr/adic/TSM/internal/fsexport/*_
<mediaId>_1.src | tail -n 1
```

- b. Call this file **<filename>**.
- c. To obtain a list of all of the failed paths, execute the Perl script:

```
decoder.pl <filename>
```

i Note: You must run the Perl script **decoder.pl**. To obtain the script, contact Quantum Technical Support and reference StorNext **Change Request** number **75191**.

4. When all of the investigation and work is complete, remove the line that was added in **Step 1** from **/usr/adic/TSM/config/fs_sysparm_override**:

```
DBG_EXPORT_DISABLE_REMOVE_TRACKING_FILES=true;
```

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="695 1360 1458 1413" data-label="Code-Block"> <pre>/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
All	73557	n/a	<p>Beginning with StorNext 6.2.0, if you change the system clock backwards while MySQL is running and then try to shutdown StorNext, the MySQL database may block when trying to stop until the current time matches the time just prior to changing the time on the system.</p> <p>i Note: If you use a Quantum appliance and use the StorNext GUI to change the time, there is no issue.</p> <p>Workaround:</p> <p>This workaround applies to the following systems:</p> <ul style="list-style-type: none"> • If you use a customer supplied system. • If you use a Quantum appliance and do not use the StorNext GUI to change the system clock. <p>To workaround this issue, perform the following procedure.</p> <ol style="list-style-type: none"> 1. Shut down StorNext. 2. Change the time. 3. Restart StorNext. <p>i Note: If shutting down StorNext is not an option, then wait for the system time to match the time prior to the time change, and then shut down to avoid this issue. For example, if you adjust the system time back one hour, then you must wait one hour before you halt or reboot your system.</p>
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:</p> <p>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