

StorNext 6.4.0 Release Notes

Contents

What's New in StorNext 6.4.0	2
Supported StorNext Upgrade Paths and Upgrade Considerations	9
Compatibility Between StorNext and Other Products	10
General Considerations	11
Upgrading Appliances	12
Appliance Release Notes	12
Known Issues	12
Contacting Quantum	27

© 2020 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 6.4.0

Purpose of this Release

The StorNext 6.4.0 release provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 6.4.0 on page 6](#).

Important Information for Storage Manager Users

If you upgrade from StorNext 6.2.0 (or later) to StorNext 6.4.0, there is a possibility that the upgrade fails when you attempt to run Storage Manager.

i Note: This information **only** applies to systems running Storage Manager.

Do the following before you perform an upgrade from StorNext 6.2.0 (or later) to StorNext 6.4.0.

i Note: Do not perform this procedure if you upgrade from StorNext 6.1.1 (or earlier).

1. Log in to the primary MDC.
2. Load the StorNext environment. For example:

```
source /usr/adic/.profile
```

3. Run a check on your MDC to determine if there might be a problem:

```
mysqlcheck sys
```

4. Depending on the output of the check, do one the following:
 - If the check does not report an error, then your system is functioning properly and there is nothing further for you to do.

For example, a non-error output is displayed:

```
sys.sys_config OK
```

- If the check reports an error about the connection, then verify that you are logged in to the primary node where StorNext is running or verify that StorNext is running on this node and repeat the procedure.

For example, a connection error output is displayed:

```
mysqlcheck: Got error: 2002: Can't connect to local MySQL server through
socket '/usr/adic/mysql/config/mysql.sock' (2) when trying to connect
```

- If the check reports an error, then execute the commands below.

For example, an error output is displayed:

```
sys.sys_config
Error: Table 'sys.sys_config' doesn't exist
status: Operation failed
```

Execute the following commands to correct the database entry that is missing:

```
mysql -e "drop database sys"
mysql_upgrade --force
```

Run a check on your MDC to determine if there might be a problem:

```
mysqlcheck sys
```

The check should not report an error:

```
sys.sys_config OK
```

New Features and Enhancements in StorNext 6.4.0

New Default Value in the StorNext File System REST Configuration File

Beginning with StorNext 6.4.0, the default value in the StorNext File System (SNFS) **snfs_rest_config.json** file for the attribute/parameter in the table below has changed.

-
- Note:** The existing **snfs_rest_config.json** file is saved to **snfs_rest_config.json.rpmsave**. If needed, you can modify any settings you might have changed in the new **snfs_rest_config.json** file.

Attribute/Parameter	Old Default Value	New Default Value
<code>data_timeout</code> (for the "fsm" process)	10 seconds	900 seconds

If you upgrade to StorNext 6.4 from a previous version of StorNext, the installation process relocates the `snfs_rest_config.json` file to `snfs_rest_config.json.rpmsave`, and installs the new template with the new default value.

See `snfs_rest_config.json` in the [StorNext 6 Man Pages Reference Guide](#) for more information.

Adding Metadata to Pre-Existing Objects

Beginning with StorNext 6.4.0, you can use the `fsobjmeta` CLI command to set metadata on an object in Object Storage.

The metadata includes the following:

- Copy number
- Last modification time
- File offset
- File key
- Path
- Segment number
- File version

By default, this command will set metadata for all objects in Object Storage. Alternatively, you can select a limited subset of the objects, either by using available options for filtering by media, namespace, class or copy number, or by providing a list of file names, either on the command line or in a file.

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

Multi-Streamed Object Store and Retrieval

Beginning with StorNext 6.4.0, you can enable the use of multi-streams operating in parallel when you store and retrieve large files to/from object storage in a request.

Use the StorNext GUI to enable multi-streams when you create or modify a policy class. See the **Multi-Stream Min Size, PUT Streams** and **GET Streams** section within the **Steering Tab** heading in [Add a Storage Manager Policy](#) (or in [Edit a Storage Policy](#)).

Use the command line interface (CLI) to run the `fsaddclass` or `fsmodclass` command to enable multi-streamed uploads/downloads when you create or modify a class. You can also enable multi-streaming with the `fsstore` and `fsretrieve` commands. See the [StorNext 6 Man Pages Reference Guide](#) for additional details.

Object Metadata

Beginning with StorNext 6.4.0, you can enable the storing of user-defined metadata along with each stored object.

Use the StorNext GUI to enable object metadata when you create a policy class. See the **Metadata** section within the **Steering Tab** heading in [Add a Storage Manager Policy](#).

Use the command line interface (CLI) to run the **fsaddclass** command or the **fsmodclass** command to enable object metadata when you create or modify a class. See the **fsaddclass** command or the **fsmodclass** command in the [StorNext 6 Man Pages Reference Guide](#).

Support for Amazon Web Services (AWS) Glacier Deep Archive Storage Class

Beginning with StorNext 6.4.0, the AWS Glacier Deep Archive storage class is supported and is a much lower cost option for data archive and protection.

- Containers configured with the **glacier_deep** storage class will restore data from Glacier Deep Archive storage within 12 hours.
- Containers configured with the **glacier_deep_bulk** storage class will restore data from Glacier Deep Archive storage within 48 hours.

Use the StorNext GUI to configure the AWS Glacier Deep Archive storage class when you configure an object storage container. See the **Containers** section in [Configure an AWS Object Storage Destination](#).

Use the command line interface (CLI) to run the **fsobjcfg** command with the **-O** option to configure a Glacier Deep Archive Storage Class. See the **fsobjcfg** command in the [StorNext 6 Man Pages Reference Guide](#).

Compatibility and Support

The [StorNext 6.4.0 Compatibility Guide](#) provides the basic compatibility for StorNext 6.4.0, 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 6.4.0 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 6.4.0

Operating System	Change Request Number	Service Request Number	Description
All	67986	337136, 404325, 467638	Having too many qustat files causes pse_snapshot to miss collecting cvgather bundle
All	69792	451886	Add in logic for LTO-8 for STK SL8500 Library
All	72660	451734, 457348, 482143, 486984, 489143, 501124	mtime in the storecand table does not get updated by rebuild policy
All	72884	538729, 526991	Multi Threading Object PUT and Get.
All	72885	454483	Windows clients have full access to StorNext specific directories.
All	73368	n/a	Support for Google multi-part uploads
All	73576	n/a	NAS beyond 2.2.0 and SNFS work – Windows Mandatory Locks need to work between SMB clients and native SNFS Windows clients
All	73634	n/a	sgadd should perform UNMAP for any added stripe groups containing thin-provisioned LUNs
All	74091	527663, 481993, 456474, 491196, 498960, 500798, 533125, 540767	fsm panic w/ ASSERT failed "!set (rptr->iref_flags & IREF_RETURN_ATTRS) == 0" on unmanaged filesystem
All	74309	538729	Delete object error after copy from tape to object storage fails
All	75270	n/a	selfattr for special attribute names not always working as expected

Operating System	Change Request Number	Service Request Number	Description
All	75411	535900	P100 media remained 100% used after fs-clean removed objects
All	75455	517831; 522496	Need notification or more retries before marking Object Storage media as write-protected
All	75469	518684	sgoffload fails due to autoAffinity with Invalid argument - Error comparing config files
All	75613	543550	Large segmented files retrieve very slowly
All	75684	396771	Allow more than 1 recursive retrieve to work at same time
All	75773	n/a	Merge code from 6.3.0LFR for TS1155 & TS1160 drives to trunk
All	75826	496940	Windows rename not working if the Parent Dir has not the ACL FILE_WRITE_ATTRIBUTES permission
All	75845	528295, 545889	Set system serial number fails for older gateways if installing newer version from ISO
All	75904	530604	snrecover caused FSM segmentation fault because lost+found directory not found
All	76209	529757, 530493	CVFS client hung after directory rename on managed file system under Linux 4.15
All	76259	534915	fs_restd fails to check for allocation errors due to oversized COPY_REQ_IPC_T objects
All	76338	536060	GUI should warn of consequences if a user tries to set the number of read streams to the value of max streams
All	76341	536060	TSM needs to log a descriptive message instead of just "No controllers found" when no streams are available
All	76347	533735	FlexSync can cause the fsm of the file system its replicating from to use up all its memory
All	76348	n/a	Changing number of backup copies does not affect all files
All	76352	537146	DDM: file system internally disabled and Admin Alert when fsretrieve from P100

Operating System	Change Request Number	Service Request Number	Description
All	76376	528985, 545298	AWS retrieve from Glacier not optimized and can take a very long time if multiple jobs are started
All	76444	533787	StorNext Client gets errors when trying to get a directory listing
All	76473	538245	snquota has an incorrect check of the hard and soft limit maximum values
All	76543	541186	Installation of Ubuntu 18.04.1 LTS security patches breaks build of cvfs kernel module
All	76726	545381	partial retrieval hang with multi-stream enabled with big start offset number
All	76735	537346, 541579	Line of null characters in a WRF_ file spins fs_policyd and fills tac_00 with errors
All	76792	538729	fsmedcopy terminates on S3 multipart uploads and fails all files in the request
All	76793	538729, 556379	Files marked as failed even though store to S3 bucket was successful
All	76876	547556	mdarchive: sgoffload leaks qtree query transactions causing mdarchive to grow without bound
All	76916	538729	fsmedcopy terminates on curl errors during S3 multipart uploads and fails remaining files in the request
All	76987	546138	metadb rest call files_by_media rest api seems to hang forever
All	77217	538729	add sysparms to tune for wsabi connectivity outages
Linux	75735	492515	snprobe not detecting secondary node as an mdc
Linux	76783	545249	sgoffload fails with extent_swap error when SG being vacated
Linux	76820	547478	quota limits report 16x greater in size after upgrade
Windows	76501	540085	An exception 0xc0000005 has occurred on line 459 of source file S:/00273715_win1064_pkg/snfs/client/vfs/windows/ntif_fastops.c.

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 6.4.0 Compatibility Guide](#).

Considerations for the StorNext File System Directories

On upgrades to StorNext 6.4.0, 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

StorNext 6.4.0 requires a minimum journal size of 4 MB and a recommended size of 64 MB. If you have file systems with journal sizes less than the 4 MB minimum, you must resize your journal size before you upgrade to StorNext 6.4.0.

Use the `cvupdatefs` utility (see the [StorNext 6 Man Pages Reference Guide](#)) or the GUI (see [Edit a File System](#)) to resize your journal size. When you resize your journal size, the new size must be 16 MB or greater. File systems with journals between 4 MB and 16 MB will run with StorNext 6.4.0, but it is recommended that these file systems have their journal resized to the recommended 64 MB.

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 6.4.0

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.4.0, 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 the [StorNext 6.4.0 Compatibility Guide](#).

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 (AXR, S3) or Quantum ActiveScale P100/X100

See the [StorNext 6.4.0 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus (AXR, S3) or Quantum ActiveScale P100/X100, and StorNext 6.4.0.

i Note: Object Storage documentation is available online at <https://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 6.4.0, refer to the [StorNext 6.4.0 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 6.4.0 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6.4.0 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

General Considerations

This section provides information about items to consider for StorNext 6.4.0.

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

⚠ Caution: If you have a Replication, Deduplication, or Object Storage license, see Change Request [Known Issues below](#) in the [StorNext Installation, Replication, HA, and Other Known Issues on page 22](#) 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>
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	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>
All	75633	n/a	<p>A StorNext NAS client cannot rename a file if the file has the read-only attribute set. This problem only affects StorNext NAS clients.</p> <p>Workaround</p> <p>A StorNext NAS client must remove the read-only attribute before it can rename the file.</p>
All	77771	n/a	<p>If you plan to have truncated files on your source system and you are running FlexSync 2.1.x (or any prior release), then the replication process can run slow and might cause the process to timeout when the file system configuration parameter (Metadata Archive Cache Size) is set to 4 GiB or smaller.</p> <p>Workaround</p> <p>Quantum recommends you set the Metadata Archive Cache Size to a minimum of 8 GiB on your source system (the default is 2 GiB). To configure the Metadata Archive Cache Size, see Edit a File System, expand the Manual Configuration drop-down, and then expand the Configuration Parameters > Features Tab drop-down.</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
------------------	-----------------------	------------------------	------------------------

macOS	75819	n/a	An Xsan client cannot mount a StorNext File System volume when a cluster number is included in the local fsnameservers file. An unexpected EOF reading reply error is displayed.
-------	-------	-----	---

```
# xsanctl mount snfs3
xsanctl: unexpected EOF reading reply
```

If you encounter the EOF error on your Xsan client, do the following workaround to prevent the issue.

Workaround

If your MDC **fsnameservers** file includes a cluster number, remove the cluster number (**@_cluster_xx**) from the **mysan.configprofile** file **before** you copy the file to the Xsan client.

For example, change:

```
10.65.181.158@_cluster0
```

to

```
10.65.181.158
```

See [Mount the StorNext File System on Xsan 5.0 \(or later\)](#) for additional information.

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>: <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_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: <pre style="background-color: #f0f0f0; padding: 5px;"># tsmstop # tsmstart</pre> 3. Verify the FS_T10K_BLOCK_FACTOR system parameter 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 846">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> <li data-bbox="618 898 1149 930">6. Unmount and re-mount your file systems. <li data-bbox="618 951 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="667 993 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 1224 1479 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 1371 1479 1434">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 1444 781 1476">Workaround:</p> <p data-bbox="602 1486 1479 1602">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>

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 1242 1394" data-label="Text"> <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	77468	n/a	<p>This issue only applies to a StorNext software only installation on your own customer-supplied MDC where StorNext is not running at time of upgrade, and StorNext migrations from one system to another where the mysql database is migrated to a system with a different version of StorNext (for example, migrating an M-Series appliance to an Xcellis appliance).</p> <p>Quantum recommends you perform a clean mysql shutdown prior to upgrading StorNext in a Software only environment or when migrating the StorNext mysql database to another system running a different version of StorNext.</p> <p>i Note: The install.stornext script handles this automatically if StorNext is running at the time of upgrade in both software only and Quantum appliance environments and you do not have to do anything further in these cases.</p> <p>Run the following command on the system where mysqld is running at the times to cleanly shutdown the database, prior to SN 6.4.0:</p> <pre># /usr/adic/mysql/bin/mysql -e 'set global innodb_fast_shutdown=0;'</pre> <p>Beginning with StorNext 6.4.0, you can create the file:</p> <pre>/usr/adic/mysql/db/enable_clean_shutdown</pre> <p>After you create the file, you can stop StorNext:</p> <pre># touch /usr/adic/mysql/db/enable_clean_shutdown</pre> <p>Use snhamgr status and mysql_control status to determine which system is currently primary and running mysqld:</p> <pre># /usr/cvfs/bin/snhamgr status # /usr/adic/mysql/bin/mysql_control status</pre>

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

On the primary system, the setting for **innodb_fast_shutdown** in **/usr/adic/mysql/my.cnf** may be set to **0** to ensure that mysqld shuts down cleanly. To do this, add **innodb_fast_shutdown=0** in the **[mysql]** section of the **my.cnf** file.

i Note: This is automatically done if StorNext is running at the time of upgrade.

After starting MySQL following the upgrade or migration, if the **innodb_fast_shutdown** setting is set to **0**, it should be reenabled to its default value of **1** to allow MySQL to shutdown quicker. Leaving the value set to **0** is not harmful, but may result in longer MySQL shutdown times. To change this, set **innodb_fast_shutdown** to **1** in the **my.cnf** file, and apply the setting to the running instance of mysqld by running:

```
# /usr/adic/mysql/bin/mysql -e 'set global innodb_
fast_shutdown=1;'
```

Invoke the following and confirm that the **innodb_fast_shutdown** value is set to **1**:

```
# /usr/adic/mysql/bin/mysql -e "show global
variables like 'innodb_fast_shutdown';"
```


Operating System	Change Request Number	Service Request Number	Description/Workaround
All	77653	563234, 563898	<p>If you upgrade from Stornext 6.2.0 (or later) to StorNext 6.4.0, there is a possibility that the upgrade fails when you attempt to run Storage Manager.</p> <p>i Note: This information only applies to systems running with Storage Manager.</p> <p>Workaround:</p> <p>Do the following before you perform an upgrade from Stornext 6.2.0 (or later) to StorNext 6.4.0.</p> <p>i Note: Do not perform this procedure if you upgrade from StorNext 6.1.1 (or earlier).</p> <ol style="list-style-type: none"> 1. Log in to the primary MDC. 2. Load the StorNext environment. For example: <div data-bbox="678 854 1458 926" data-label="Code-Block"> <pre>source /usr/adic/.profile</pre> </div> 3. Run a check on your MDC to determine if there might be a problem: <div data-bbox="678 1050 1458 1121" data-label="Code-Block"> <pre>mysqlcheck sys</pre> </div> 4. Depending on the output of the check, do one the following: <ul style="list-style-type: none"> • If the check does not report an error, then your system is functioning properly and there is nothing further for you to do. <p><i>For example, a non-error output is displayed:</i></p> <div data-bbox="712 1377 1458 1449" data-label="Code-Block"> <pre>sys.sys_config OK</pre> </div> • If the check reports an error about the connection, then verify that you are logged in to the primary node where StorNext is running or verify that StorNext is running on this node and repeat the procedure. <p><i>For example, a connection error output is displayed:</i></p>

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

```
mysqlcheck: Got error: 2002: Can't connect to local MySQL server through socket '/usr/adic/mysql/config/mysql.sock' (2) when trying to connect
```

- If the check reports an error, then execute the commands below.
For example, an error output is displayed:

```
sys.sys_config
Error: Table 'sys.sys_config' doesn't exist
status: Operation failed
```

Execute the following commands to correct the database entry that is missing:

```
mysql -e "drop database sys"
mysql_upgrade --force
```

Run a check on your MDC to determine if there might be a problem:

```
mysqlcheck sys
```

The check should not report an error:

```
sys.sys_config OK
```

Contacting Quantum

Contacts

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

<https://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:

<https://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