

StorNext 6.4.1 Release Notes

Contents

What's New in StorNext 6.4.1	2
Supported StorNext Upgrade Paths and Upgrade Considerations	8
Compatibility Between StorNext and Other Products	9
General Considerations	10
Upgrading Appliances	11
Appliance Release Notes	11
Known Issues	11
Contacting Quantum	24

© 2021 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. ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Quantum, The Quantum Logo, QXS, Scalar, StorNext, SuperLoader, Vision, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Quantum specifications are subject to change.

What's New in StorNext 6.4.1

Purpose of this Release

The StorNext 6.4.1 release provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 6.4.1 on page 5](#).

New Features and Enhancements in StorNext 6.4.1

Support for Operating Systems

Beginning with StorNext 6.4.1, the following operating systems are supported:

- Red Hat Enterprise Linux (RHEL) 7 Update 8
- Red Hat Enterprise Linux (RHEL) 7 Update 9
- Red Hat Enterprise Linux (RHEL) 8 Base
- Red Hat Enterprise Linux (RHEL) 8 Update 1
- CentOS 7 Update 8
- CentOS 7 Update 9
- CentOS 8 Base
- CentOS 8 Update 1

For complete details, see the [StorNext 6.4.1 Compatibility Guide](#).

Information About StorNext Upgrades

Background

Due to the potential of the MySQL database corruption when you upgrade the release of the MySQL database, Quantum added logic to the upgrade utilities to ensure that the MySQL database is cleanly shut down on upgrade (*reference StorNext Bugzilla Bug 77325*).

The database is shut down using a "slow shutdown" process in a MySQL setting; the setting used for the database shut down is documented in the MySQL Reference Manual (see [InnoDB shutdown mode](#)).

There are instances where excessive slowdowns on upgrade occur due to the "slow shutdown" of the MySQL database (*reference StorNext Bugzilla Bug 78960*).

Analysis

Quantum assessed the "slow shutdown" process is **ONLY** required when you upgrade from a major release of MySQL to another major release.

- i Note:** A major release of MySQL is signified by a change to the second digit in the release number. For example, MySQL 5.6 to MySQL 5.7.

The table below provides the MySQL release included in each StorNext release.

StorNext Release	MySQL Release
5.1.1 through 6.1.1	5.6
6.2.0 through 7.0.1	5.7

Quantum recommends the following, according to the releases in the table above:

- If you are upgrading from StorNext 6.2.0 through StorNext 7.0.1, then you are **NOT** required to perform the "slow shutdown" process.
- If you are upgrading from StorNext 6.1.1 (or earlier) to StorNext 6.2.0 (and later), then you **MUST** perform the "slow shutdown" process to avoid potential database corruption.

Workaround

The following section describes the remedy for customers doing an upgrade.

Upgrade from StorNext 6.2.0 (and later) to StorNext 6.4.0 through StorNext 7.0.1 Releases

Before you initiate a StorNext upgrade, add a line of text to the `/usr/adic/mysql/bin/mysql_control` script on **both** the **primary** and **secondary** MDCs to enable a fast shutdown immediately before the command to shutdown MySQL (and avoid the slow shutdown).

For example, the highlighted text in the code snippet below illustrates the line to add in the `mysql_control` script:

```

stop)
  stopit=0
  if mysqld_running
  then
    stopit=1
    if [ -f $ENABLE_CLEAN_SHUTDOWN ] ; then
<- Not present in earlier releases
      log_out "Requesting clean database shutdown"
<- Not present in earlier releases
      $MYSQL_DIR/bin/mysql -e "set global innodb_fast_shutdown=0"
<- Not present in earlier releases

```

```

        fi
    <- Not present in earlier releases
        log_out "Stopping mysqld"
        $MYSQL_DIR/bin/mysql -e "set global innodb_fast_shutdown=1"
        $MYSQL_DIR/bin/mysqladmin --defaults-file=$MY_CNF shutdown
        sleep 1
    fi

```

After you add the line of text to the **mysql_control** script, then you can initiate the upgrade process.

-
- i Note:** Do not be concerned about backing out the script change after the upgrade is finished; the StorNext upgrade process overwrites the **mysql_control** script.

Upgrade from StorNext 6.1.1 (and earlier) to StorNext 6.4.0 through StorNext 7.0.1 Releases

Quantum recommends you perform a "slow shutdown", as the upgrade is currently doing. The upgrade process might require additional time due to the clean shutdown of the database during the upgrade.

-
- i Note:** It is not possible to estimate the amount of time required for the upgrade process.

Future StorNext Releases

To prevent manual intervention in the upgrade process, Quantum plans to add logic to the StorNext upgrade process, that verifies the MySQL version (*reference StorNext Bugzilla Bug 79053*).

Compatibility and Support

The [StorNext 6.4.1 Compatibility Guide](#) provides the basic compatibility for StorNext 6.4.1, 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 6.4.1 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 6.4.1

Operating System	Change Request Number	Service Request Number	Description
All	73496	465444, 577160	ArcDisp hit OOM condition and got killed, leading to MSM abnormal termination
All	75143	525455	ENOENT returned for lookups on linux 7.6 with storage manager
All	75213	526974, 555991, 566835, 570951, 570482	Misleading message on MDC w.r.t san_client DOWN disk devices
All	75468	512218, 569557	fsmimport fails to import media with the hardware write lock set
All	76615	560280	Auditing generates a nodechange to every client on a close
All	76970	546119, 554609, 576026, 578397, 576097, 571413, 597177	fsm panic upon ASSERT failed "pclient->cl_rsvd_counted == 0"
All	76997	543055, 571150, 576036, 576037, 577537, 569740, 578514, 578812, 580923, 578146, 586085, 585472, 592724, 599629	Long running transaction in UI logic - this MAY result in growing ibdata1 file
All	77221	553837, 580930	wsar_agent is running into heap exhaustion due to MariaDB mem_root implementation
All	77283	555650	GUI: modify policy "Copy Expiration" warning to be more explicit about file deletion

Operating System	Change Request Number	Service Request Number	Description
All	77481	589728	find command return IO error when running prodcon
All	77515	560062 549992	FSM: PANIC: /usr/cvfs/bin/fsm ASSERT failed "IP_IS_SPACE_TREE (ip) this.idiext_frblock + extsize <= next.idiext_frblock"
All	77547	549992	FSM: PANIC: /usr/cvfs/bin/fsm ASSERT failed "IP_IS_SPACE_TREE (ip) this.idiext_frblock + extsize <= next.idiext_frblock"
All	77646	559566, 570899	Is of a directory failing on a client with "Input/output error" while mdc is creating many files in it.
All	77653	563234, 563898	Mysqld fails to start after upgrading to StorNext version 6.4.0
All	77678	563234, 563898	file-per-table conversion needs to account for sys_config.ibd file
All	77730	564850	Snapshot: Include contents of files in the /usr/adic/TSM/internal/locks directory
All	77746	n/a	Red Hat (RHEL) 7 update 8 Client Support
All	77771	567119	Reduce the processing time of request api files_by_media with small number of truncated files
All	77812	n/a	GUI: Red Hat (RHEL) 7 update 8 Client Support
All	77861	569557	The fsmescan util is producing extraneous fields that causes fsmimport to fail
All	77879	561495	Disks and Path of multipath devices set to standby

Operating System	Change Request Number	Service Request Number	Description
All	77944	587504	Google multi-part upload fails with bad request error
All	78022	563167	RPM script change needed to modify max_user_watches to 32768 in sysctl.conf
All	78031	574044, 581478, 578897	fs_moverd segmentation fault occurs on multipart upload to Google Cloud Storage, with full payload disabled
All	78051	577213	install.stornext will not upgrade components if previous install did not fully complete
All	78105	579144	snrecover process deadlocks restoring files that have been deleted
All	78116	578751, 583673	fsretrieve -R fails if the files are stored on SAMFS tapes
All	78146	580579	fs_moverd segmentation fault occurs during Google delete batch request, due to null-string boundary
All	78188	581921, 589950	fs_resourced crashing frequently after reporting many "No controllers found" errors
All	78292	583924	fsimport can fail to import tapes if a directory has a single quote in the name
All	78315	585257	fsrecover is unable to recover files from a long directory path
All	78508	591254	Google objects left orphaned after parallel upload errors
All	78518	n/a	Update quantum_disk_catalog to include F-Series and H-Series entries for Quantum Branded disk.
All	78520	588842	CVFS down_write()'s i_rwsem without holding the semaphore

Operating System	Change Request Number	Service Request Number	Description
All	78523	n/a	When Windows DLC clients run continuous I/O to a file system that has been remade and resized, mdc kernel panics upon mounting the new FS
All	78601	591208	Windows Explorer large file copy takes several seconds to report any progress when buffer cache >= 4GB
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	77804	n/a	6.4.x Client Support for CentOS/RHEL 8
Linux	77932	n/a	GUI: 6.4.x Client Support for CentOS/RHEL 8
Linux	77934	n/a	Include CentOS/RHEL 8 Client package in 6.4.1
Linux	78070	577099	Ubuntu: mounting an Image as loopback result in "mount: /dev/loop0: can't read superblock"

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 6.4.1 Compatibility Guide](#).

Journal Size Guidelines

StorNext 6.4.1 requires a minimum journal size of 4 MB (Quantum recommends you configure your journal size to 64 MB). If your file system has a journal size less than the 4 MB minimum, you must resize your journal size before you upgrade to StorNext 6.4.1.

⚠ Caution: If you upgrade and the journal size is less than the 4 MB minimum, then the FSM does not start and an error is logged.

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 run with StorNext 6.4.1, but Quantum recommends you configure your journal size to 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 6.4.1

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.4.1, 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 6.4.1 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 6.4.1 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus (AXR, S3) or Quantum ActiveScale P100/X100, and StorNext 6.4.1.

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 6.4.1, refer to the [StorNext 6.4.1 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 6.4.1 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6.4.1 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

General Considerations

This section provides information about items to consider for StorNext 6.4.1.

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 23](#) 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 Workflow Director](#)
- [Xcellis Workflow Extender](#)
- [Xcellis Foundation](#)
- [Artico](#)

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_allloc=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 mount_cvfs(8). The sparse option will introduce some overhead when using Windows Explorer. Before setting the protect_allloc option, see the sparse option in mount_cvfs(8) 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>
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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
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>
All	78082	575600	<p>If you set the parameter audit=true in the file system configuration file, the result causes all I/O activity, including reads, to be recorded in the mdarchive. This can greatly increase the number of updates applied to the mdarchive which, in turn, increases the amount of mdarchive compaction activity by the FSM.</p> <p>There is one instance where this increased compaction activity caused the mdarchive to grow beyond the capacity of the HA shared file system.</p> <p>Workaround</p> <p>To work around this issue, Quantum recommends you increase the parameter metadataArchiveCache from the default 2 GB to at least 4 GB, but preferably larger when setting audit to true in the configuration file.</p>
All	78470	591234	<p>The cvadmin output displays the storage size unit as KB/MB/TB (<i>base 10</i>).</p> <p>Workaround</p> <p>To work around this issue, you should treat the storage size unit as KiB/MiB/TiB (<i>base 2</i>).</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
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>
macOS	75819	n/a	<p>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.</p> <pre># xsanctl mount snfs3 xsanctl: unexpected EOF reading reply</pre> <p>If you encounter the EOF error on your Xsan client, do the following workaround to prevent the issue.</p> <p>Workaround</p> <p>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.</p> <p>For example, change:</p> <pre>10.65.181.158@_cluster0</pre> <p>to</p> <pre>10.65.181.158</pre> <p>See Mount the StorNext File System on Xsan 5.0 (or later) for additional information.</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 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>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># tsmstop # tsmstart</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>3. Verify the FS_T10K_BLOCK_FACTOR system parameter contains the new value:</p>
			<pre># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre>
			<p>4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</p>
			<p>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:</p>
			<pre>snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_ dma_read_length=16m 0 0</pre>
			<p>6. Unmount and re-mount your file systems.</p>
			<p>7. Use new T10K media to store a copy of the file from the disk.</p>
			<p>i Note: Step 7 is very important; when the new copy is made to the new tapes, the new tapes are labeled with a 2 MB block size, which is used for subsequent writes or reads to and from the media. Tapes on which fsformat was run before the change will use the block factor in use at that time. This change will not impact those tapes.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	46693	n/a	<p>If you execute the command snbackup -s while a full or partial backup is running, then you might receive a message that indicates the following file is in an invalid format:</p> <pre data-bbox="764 474 1458 543" style="background-color: #f0f0f0; padding: 5px;">/usr/adic/TSM/internal/locks/backup.1f</pre> <p>This is due to the snbackup -s process reading the backup.1f status file while the backup process is updating it.</p> <p>Workaround:</p> <p>Ignore the message; to clear-up the process, re-execute the command snbackup -s (provided that the backup is not writing to the backup.1f status file while snbackup -s is trying to read it again).</p>
All	47833	n/a	<p>When copying files between media using the CLI command fsmedcopy, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p>Note: This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p>Workaround:</p> <p>Currently, a workaround does not exist for this known issue.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in /usr/adic/TSM/logs/tac which contains the text:</p> <div data-bbox="750 472 1458 779" style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> <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> </div> <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: <pre style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;">/usr/cvfs/data/<fs>/config_history</pre> 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>

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

Region	Support Contact
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