

StorNext 6.3.1 Release Notes

Contents

What's New in StorNext 6.3.1	2
Supported StorNext Upgrade Paths and Upgrade Considerations	5
Compatibility Between StorNext and Other Products	6
General Considerations	7
Upgrading Appliances	8
Appliance Release Notes	8
Known Issues	8
Contacting Quantum	23

© 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 6.3.1

Purpose of this Release

The StorNext 6.3.1 release provides new features and enhancements listed in the section [New Features and Enhancements in StorNext 6.3.1 below](#), and also provides software fixes listed in the section [Fixed Issues and Enhancements Addressed in StorNext 6.3.1 on the next page](#).

New Features and Enhancements in StorNext 6.3.1

Support for Operating Systems

Beginning with StorNext 6.3.1, the following operating systems are supported:

- Red Hat Enterprise Linux (RHEL) 7 Update 7
- CentOS 7 Update 7
- Debian 9.4
- Xsan 5.0.1 / macOS 10.15

For complete details, see the [StorNext 6.3.1 Compatibility Guide](#).

Support for FlexTier Vendors and Platforms

Beginning with StorNext 6.3.1, the following FlexTier private cloud vendors and platforms are supported:

- Object Matrix (MatrixStore)
- Wasabi (Hot Cloud Storage)
- Caringo (Swarm)

For complete details, see the [StorNext 6.3.1 Compatibility Guide](#).

Compatibility and Support

The [StorNext 6.3.1 Compatibility Guide](#) provides the basic compatibility for StorNext 6.3.1, 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 6.3.1 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 6.3.1

Operating System	Change Request Number	Service Request Number	Description
All	42076	142079, 590581	cvadmin who command shows same snfs id for all clients
All	71410	417848, 469654, 523117, 529382	TSM retrieves fail after LTFS cannot lock the device when opening the tape drive device
All	72675	450834	Provide a log message that fsm binary isn't executable
All	73199	509542	Running a file system out of space produces cryptic and alarming errors
All	74170	483303	StorNext buffered client I/O can hang for minutes at a time
All	74469	n/a	cvupdatefs should perform UNMAP for any added stripe groups containing thin-provisioned LUNs

Operating System	Change Request Number	Service Request Number	Description
All	74650	488961, 509455	fs_moverd segfault in SNDB::Connection::connect() when attempting to connect over NAS VIP or docker0 interface
All	74764	499502, 1381652, 1481710, 493804	Gui Report "SAN and LAN clients" missing details if FS is running on secondary Node
All	74955	497830	GUI does not allow to set media format for _adic_ backup copy1 after removal of Object Storage
All	75030	523224, 521892, 507994, 507556, 508808, 509216, 512793, 512838, 518689	Downgrade Severity 1 RAS Alert SL_EVT_TAPE_ALERT (Tape alert)
All	75097	508720	mdarchive build error when FSM assigns an invalid inode_id because Inodelookup failed
All	75127	n/a	The fsexport command in File Copy-Export mode should list files that can't be copied due to unavailable source media
All	75167	506855, 511403, 512492, 510010, 514813, 515127, 519025, 523837, 527502	Runaway mdarchive created close to 1 million files and filled up HA shared file system
All	75168	506855	cvadmin "mdarchive rebuild" didn't rebuild runaway mdarchive
All	75207	n/a	snaci is unable to add ACE's with spaces in their names
All	75236	513366	Provide more info in the "Invalid msg length" error message
All	75237	511468, 513161, 524677	snquota -L report is looping endlessly
All	75314	501812, 513189	Direct Cross Launch from SN UI To Connect Does Not Work In 6.2 On M-Series Due to Lack Of Containers On M-Series
All	75461	513366	client disconnect with Invalid msg length
All	75602	514175	Capacity pie chart wrong for library on home page
All	75603	494038	The fsmpm on Windows picks a random IP address to use as pm_ipaddr, also known as "id"

Operating System	Change Request Number	Service Request Number	Description
All	75696	520733, 522482	Open Failures from the MDC occurring when upgrading to 6.3.0 and nss_ctl contains MDC IP Address
All	75920	527400, 503185, 512277, 529382	fs_moverd leaks SCSI file descriptors when constant file read errors occur during stores
Windows	27418	1142606	Windows cannot set the noatime mount option - no gui support
Windows	74339	488689	Windows SAN client crashes with bugcheck error DPC_WATCHDOG_VIOLATION (133) in DropDiskReference()

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 6.3.1 Compatibility Guide](#).

Considerations for the StorNext File System Directories

On upgrades to StorNext 6.3.1, 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 6.3.1 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 6.3.1

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.3.1, 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 *Appliance Controller Compatibility* available online at <https://gsupport.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 (AXR, S3) or P100/X100

See the [StorNext 6.3.1 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus (AXR, S3) or P100/X100, and StorNext 6.3.1.

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 6.3.1, refer to the [StorNext 6.3.1 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 6.3.1 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

For all other components and features, see the [StorNext 6.3.1 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

General Considerations

This section provides information about items to consider for StorNext 6.3.1.

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 21](#) 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	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: <pre data-bbox="678 793 1458 867">/usr/cvfs/lib/snnas_control stop</pre> 2. Stop the Controller by executing the following command: <pre data-bbox="678 989 1458 1094">/usr/local/quantum/bin/sml_service_tool stop snnas_controller</pre>
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
------------------	-----------------------	------------------------	------------------------

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>: <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_FACTOR system parameter contains the new value:

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre data-bbox="695 352 1247 422"># showsysparm FS_T10K_BLOCK_FACTOR FS_T10K_BLOCK_FACTOR=32</pre> <p data-bbox="634 485 1446 667">4. Save the current copies of your <code>/etc/fstab</code> on the MDCs and the DDM clients.</p> <p data-bbox="634 569 1438 667">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 data-bbox="695 730 1385 831">snfs1 /stornext/snfs1 cvfs rw,auto_dma_write_length=16m,auto_dma_read_length=16m 0 0</pre> <p data-bbox="634 898 1146 930">6. Unmount and re-mount your file systems.</p> <p data-bbox="634 951 1365 982">7. Use new T10K media to store a copy of the file from the disk.</p> <p data-bbox="683 993 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.</p>
All	46693	n/a	<p data-bbox="602 1230 1458 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 1458 1440">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 1451 781 1482">Workaround:</p> <p data-bbox="602 1493 1458 1591">Ignore the message; to clear-up the process, re-execute the command snbackup -s (provided that the backup is not writing to the backup.1f status file while snbackup -s is trying to read it again).</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	47833	n/a	<p>When copying files between media using the CLI command fsmedcopy, the file is not re-segmented to match the segment size of the destination media. Rather, the original segments are copied to the target media type and the distribution of segments across destination media will, therefore, be the same as the distribution on the source media.</p> <p>i Note: This behavior might cause file data segment distribution to be sub-optimal on the destination media.</p> <p>Workaround:</p> <p>Currently, a workaround does not exist for this known issue.</p>
All	69265	n/a	<p>Your DDMs might experience a timeout if you try to connect to the database. The issue is identified by an error log in /usr/adic/TSM/logs/tac which contains the text:</p> <pre>Process fs_moverd on <host> timed out trying to connect to the database. This usually indicates network connectivity trouble. Try increasing the timeout value by setting the connect_timeout value in /usr/adic/mysql/my.cnf. The default setting is 10 seconds so the new value should be larger.</pre> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Increase the database connection timeout value by adding the following line to /usr/adic/mysql/my.cnf under the section labeled [mysqld] connect-timeout=240. 2. Cycle the Storage Manager in order to pick up the updated timeout value.
All	69341	n/a	<p>If you have the IBM APFO driver installed and configured, then when you perform an fsmedread operation of a partial tape block from a full tape block, the operation can fail with errno=12.</p> <p>i Note: This issue affects all IBM APFO versions 3.0.19 and earlier, and has an impact primarily on disaster recovery procedures.</p> <p>Workaround:</p> <p>To correct this, perform an fsmedread operation without the IBM APFO driver.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	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>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>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
All	76582	n/a	<p>The Debian 9.0 client software download package is not displayed when you navigate to the Tools > Client Download GUI page.</p> <p>Workaround:</p> <p>To fix this issue, you must edit the snfsClients.xml file to include the Debian 9.0 client software download package entry.</p> <ol style="list-style-type: none"> Open an SSH connection to the server using the IP address assigned to the server on the StorNext management/public network. <ol style="list-style-type: none"> Log in to the command line using the following credentials: <ul style="list-style-type: none"> User name: stornext Password: <stornext user account password> Enter sudo rootsh to gain root user access. Enter the password for the stornext user account again. Edit the snfsClients.xml file located in the following directory: <pre style="background-color: #f0f0f0; padding: 5px;">/usr/adic/tomcat/webapps/ROOT/WEB-INF/classes/snfsClients.xml</pre> In the snfsClients.xml file, add a section for Debian 9.0 in lines 58 through 61. For example, see the highlighted section below: <pre style="background-color: #f0f0f0; padding: 5px;"> <resource> <os>Linux Debian 8.0 (64 bit)</os> <file>snfs_client_Debian8_x86_64.bin</file> </resource> <resource> <os>Linux Debian 9.0 (64 bit)</os> <file>snfs_client_Debian9_x86_64.bin</file> </resource> <os>Linux Ubuntu 14 (64 bit)</os></pre>

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

```

        <file>snfs_client_Ubuntu14_x86_64.bin</file>
    </resource>
    
```

- Using the StorNext GUI, reload the **Tools > Client Download** GUI page.

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"> Shut down the Safari browser window(s). Restart the Safari browser, and then retry the operation. Uncheck the Verify SMTP Server Connectivity box, and then retry the operation. Set Authentication to NONE, and then retry the operation. 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>
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>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"> Shut down StorNext. Change the time. Restart StorNext. <p>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>

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

Contacting Quantum

Contacts

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

<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