

StorNext 6.2.1 Release Notes

Contents

What's New in StorNext 6.2.1	2
Supported StorNext Upgrade Paths and Upgrade Considerations	3
Compatibility Between StorNext and Other Products	4
General Considerations	5
Upgrading Appliances	6
Appliance Release Notes	6
Known Issues	7
Contacting Quantum	26

© 2019 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexSync, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners. Quantum specifications are subject to change.

What's New in StorNext 6.2.1

Purpose of this Release

The StorNext 6.2.1 release provides important software fixes. See [Fixed Issues and Enhancements Addressed in StorNext 6.2.1 on the next page](#).

 **Caution:** If your system is running StorNext 6.1.x or earlier, see [Product Bulletin 103](#).

Compatibility and Support

The [StorNext 6 Compatibility Guide](#) provides the basic compatibility for StorNext 6.2.1, includes the StorNext components supported, operating systems and service packs, libraries and drives, browsers, virtual machines, and appliance support. Listed below are just a few of the types of information available to you in the [StorNext 6 Compatibility Guide](#).

- **Upgrade Paths:** Provides information on what upgrades to this release are supported.
- **Appliance Support:** Provides information on what StorNext and Lattus appliances are supported with this release or are compatible with it.
- **Operating Systems and Platforms:** Provides information on what StorNext components run on various operating systems and service packs. Also includes which operating systems have been newly added or removed.
- **Client Interoperability:** Provides information on what StorNext clients running other versions of StorNext are compatible with metadata-controllers (MDCs) running this release.
- **Virtual Machine Support:** Provides information on what StorNext components running on selected operating systems and service packs are supported in virtual machines.
- **Compatibility with Other Products:** Provides information on references to additional StorNext sold-separately products that are supported with this release.
- **Browser Support:** Provides information on what versions of browsers are supported with the GUI in this release.
- **Drives and Libraries:** Provides information on what Quantum and 3rd party drives and libraries are supported with this release.

Fixed Issues and Enhancements Addressed in StorNext 6.2.1

Operating System	Change Request Number	Service Request Number	Description
All	74164	482099	fs_resourced locks up and requires restart of TSM
All	74228	495024	'/usr/cvfs/lib/snnas_control start' does not start snnas_service
All	74290	488616, 493109	NAS beyond 2.2.0 and SNFS work -- Windows Directory notifications (samba <-> SNFS) not delivered.
All	74476	480485	Non-existent or truncated display of existing file system quotas in the snquota CLI and StorNext GUI
All	74644	496096	recover_xattr_payloads() does not honor snbt_lookup() API causing snbt_payload_expect_dupes():ASSERT(FALSE)
Windows	74181	479521; 487047	max length russian characters
Windows	74583	480485	Photoshop deletes the original file if Quota is exceeded - FileSetAttr() problem

Supported StorNext Upgrade Paths and Upgrade Considerations

StorNext Software Upgrade Matrix

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

Considerations for the StorNext File System Directories


On upgrades to StorNext 6.2.1, 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 6.2.1 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 6.2.1

Due to the fact that the full 64-bit inode numbers are exposed to Linux after Linux clients are upgraded to StorNext 6.2.1, 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 <http://qsupport.quantum.com/kb/Flare/Content/appliances/ACC/DocSite/Compat.htm>.

Infiniband

StorNext 6.2.1 works with Infiniband SRP (SCSI RDMA Protocol) attached storage for Linux and Windows 2008R2.

Lattus

See the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#) for information about compatibility between Lattus and StorNext 6.2.1.

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

StorNext Web Services

StorNext Web Services enables you to run third-party application program interfaces (APIs) with StorNext. To view the latest commands supported by the StorNext Web Services, refer to the [StorNext 6 Web Services Guide](#) in the [StorNext 6 Documentation Center](#).

Apple Xsan

Xsan is software that enables multiple Mac computers to concurrently access hundreds of terabytes of content on Xserve RAID or Promise RAID storage over high-speed Fibre Channel which allows you to share data faster and consolidate projects. Quantum supplements this solution with StorNext data management software, enabling Apple Xsan customers to use applications running on Windows, Linux, and UNIX with their Xsan and share content across more systems.

For information about compatibility between Apple Xsan and StorNext 6.2.1, refer to the [StorNext 6 Compatibility Guide](#) in the [StorNext 6 Documentation Center](#).

Supported Browsers

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

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

General Considerations


This section provides information about items to consider for StorNext 6.2.1.

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 [73945 on page 23](#) 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>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	67363	n/a	<p>StorNext 5.4.0.x incorrectly allowed the Unix ID Mapping type to be set to none when the Security Model is set to acl. As a result, file systems fail to start when the Unix ID Mapping type is set to none when the Security Model is set to acl.</p> <p>Beginning with StorNext 6, the FSM does not start when this invalid combination of settings is used.</p> <p>Workaround:</p> <p>To prevent this issue, set the Unix ID Mapping to either winbind or algorithmic for any file system where the Security Model is set to acl. You can make the adjustment before or after upgrading.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	74202	n/a	<p>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"> Check to see if StorNext file systems are mounted: <pre>grep cvfs /proc/mounts</pre> Mount the HaSahred file system: <pre>mount /usr/adic/HAM/shared</pre> 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> <p>Maintenance Mode. An extension of Scenario #1, where NAS must vacate the file systems for an extended period of time.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<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>su sysadmin -c 'reg show nas.heartbeat.check_state_secs'</pre> <p>or</p> <pre>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>su sysadmin -c 'reg set nas.heartbeat.check_state_secs X'</pre> <p>or</p> <pre>su sysadmin -c 'reg set stornext_service.stop_period X'</pre> <p>Workaround:</p> <p>For Scenario #1, Quantum recommends that you extend the time period</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<p>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="683 709 1458 779">/usr/cvfs/lib/snnas_control stop</pre> 2. Stop the Controller by executing the following command: <pre data-bbox="683 905 1458 1010">/usr/local/quantum/bin/sml_service_tool stop snnas_controller</pre>
All	74299	n/a	<p>Beginning with StorNext 6.2, when the snfsdefrag command skips any file for any reason and the verbose option is not supplied, the follow message is displayed:</p> <pre data-bbox="626 1178 1458 1283">Processing one or more files failed, run Verbose option for details</pre> <p>When this occurs, the command also returns a non-zero exit status, suggesting that the command encountered an error. However, there are normal conditions where snfsdefrag can skip a file. For example, a file can be skipped if it does not need to be defragmented because its extent count is already low enough. Also, snfsdefrag skips device files and symbolic links. So the fact that snfsdefrag is reporting "failures" in this case is unexpected.</p> <p>Workaround</p> <p>You can ignore the error and non-zero exit status.</p>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	74871	502252, 497249, 503172	<p>If your system meets ALL of the following criteria, you could encounter a locking issue:</p> <ul style="list-style-type: none"> Your system is running StorNext 6.2.0 or StorNext 6.2.1. Your system is running StorNext NAS. Your system is configured with the default NAS configuration of oplocks to yes. <p>There was change in StorNext 6.2.0 that resulted in a locking order issue. The locking order issue results in the Linux kernel watchdog detecting any CPU that loops in the kernel for more than twenty seconds without giving other tasks a change to run. The CPU prints the current stack trace and then the soft lockup is detected.</p> <p>The kernel stack must have a stack trace with FsmOpenChange calling linuxif_has_lease and another stack with smbd (or another process but most likely smbd) doing a call to cvfs_setlease_internal that then calls VerifyCvp and blocks waiting for the cvnode. Once these two threads are in deadlock, other threads can block. The following are examples of Kernel soft lockup call traces:</p> <pre> soft lockup - CPU stuck - call trace #0 [ffffffff8970dc5f] queued_spin_lock_slowpath+0xb/0xf #1 [ffffffff8971b750] _raw_spin_lock+0x20/0x30 #2 [ffffffffffc0a5776e] linuxif_has_lease+0x1e/0x50 [cvfs] #3 [ffffffffffc0ad9ea6] FsmOpenChange+0x746/0x1090 [cvfs] #4 [ffffffff890fbef] ? __getnstimeofday64+0x3f/0xd0 #5 [ffffffffffc0a9d6ed] FsmInputFl+0x5fd/0x880 [cvfs] #6 [ffff90bd395bbe28] DmonWaitReq at fffffffffffc060adc5 [cvfs] #7 [ffff90bd395bbe80] StartDmon at </pre>

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

```

ffffffffffc060b1b6 [cvfs]
#8 [ffff90bd395bbec8] kthread at ffffffff84ebb161
    
```

```

soft lockup - CPU stuck - call trace
#0 [ffff90bcf695ba38] __schedule at
ffffffffff85513904
#1 [ffff90bcf695bac8] schedule at
ffffffffff85513f79
#2 [ffff90bcf695bad8] schedule_timeout at
ffffffffff855118e9
#3 [ffff90bcf695bb88] linuxif_sleepwait at
ffffffffffc05f4597 [cvfs]
#4 [ffff90bcf695bc18] MdOsWait at
ffffffffffc06e2662 [cvfs]
#5 [ffff90bcf695bc38] EnqLock at ffffffff80612116
[cvfs]
#6 [ffff90bcf695bcb0] cvAcqReadLock at
ffffffffffc0612732 [cvfs]
#7 [ffff90bcf695bd70] VerifyCvp at
ffffffffffc06772a7 [cvfs]
#8 [ffff90bcf695bdd0] cvfs_setlease_internal at
ffffffffffc06f02b7 [cvfs]
#9 [ffff90bcf695be50] cvfs_setlease at
ffffffffffc05f38aa [cvfs]
#10 [ffff90bcf695be80] __vfs_setlease at
ffffffffff85077f92
#11 [ffff90bcf695be90] fcntl_setlease at
ffffffffff850781c5
#12 [ffff90bcf695bee8] sys_fcntl at
ffffffffff8502ebfb
    
```

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

Workaround

To correct this issue, disable **smb oplocks** by doing the following:

1. Log in to the controller console. See [Access the Appliance Controller Console and Change Passwords](#).
2. At the prompt, enter the following:

```
share change smb global kernel oplocks = no
```

3. Disconnect and reconnect all of the SMB clients to activate the new configuration. If you cannot disconnect and reconnect all of the SMB clients, then you must restart **smbd** on all of the nodes in the cluster by running the command:

```
system restart services smbd
```

If you disable **oplocks**, then Samba does not use a kernel lease to simulate **oplocks** and revert to the behavior prior to Appliance Controller 2.2. When **oplocks** is disabled, a directory is not shared with SMB and NFS at the same time.

macOS	66948	322824, 336945	<p>If you access StorNext file systems from Apple Xsan clients, then you might encounter I/O error messages in the system log that do not contain details about real I/O errors detected on the Xsan client.</p> <p>Workaround</p> <p>If you encounter the errors on an Xsan client, contact Apple.</p>
-------	-------	-------------------	--

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

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

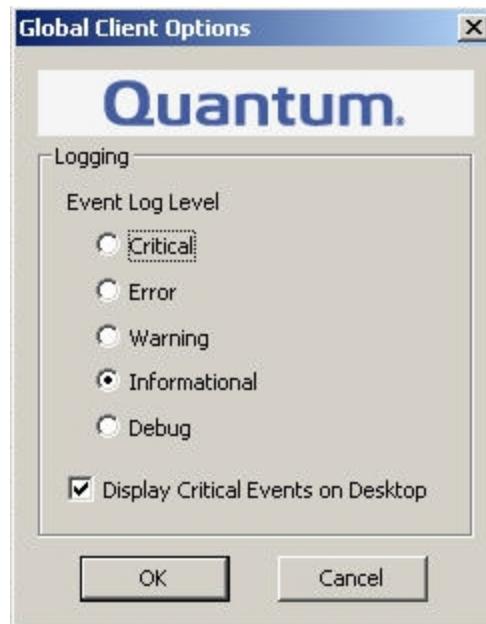
There are two conditions which lead to this issue:

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

Workaround

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

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



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

StorNext Storage Manager Known Issues

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	43320	1581004	<p>File retrieves from media to disk can be suboptimal for fast tape drives like the Oracle STK T10K drives. This scenario can occur when the retrieve event is initiated on a host that is different from the host running the mover process, which requires the use of synchronous direct I/O.</p> <p>Workaround:</p> <p>To work around this issue and achieve optimal performance for both file stores and retrieves with the T10K drives, increase the default I/O size used by the mover process and make the mover process use asynchronous buffered I/O when the use of synchronous direct I/O is not required, using the following steps:</p> <ul style="list-style-type: none"> i Note: This workaround might also help improve the performance of the faster LTO drives by updating the FS_LTO_BLOCK_FACTOR system parameter. i Note: Changes to FS_XXX_BLOCK_FACTOR only affects tapes formatted after the change. <ol style="list-style-type: none"> 1. Change the FS_T10K_BLOCK_FACTOR system parameter from 8 to 32 by adding the following entry to <code>/usr/adic/TSM/config/fs_sysparm_override</code>: <div data-bbox="678 1031 1458 1100" style="background-color: #f0f0f0; padding: 5px; margin: 10px 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 data-bbox="678 1451 1458 1556" style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> <pre># tsmstop # tsmstart</pre> </div> 3. Verify the FS_T10K_BLOCK_FACTORsysparm contains the new value:

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

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

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	72993	452722	<p>When DDM is enabled for non-primary systems, kernel error messages are logged for reservation conflict because the Primary MDC mounts the tape which sets the reservation to itself and the DDM sets the reservation to itself once the tape is ready.</p> <p>This issue causes a benign message in the <code>/var/log/messages</code> file for reservation conflict. Since the unmounting of the tape also resets the reservation back to the primary MDC, this message is generated the next time a DDM (non-primary MDC) accesses a tape.</p> <p>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: <div data-bbox="678 1346 1458 1413" style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre data-bbox="695 1360 1240 1394">/usr/cvfs/data/<fs>/config_history</pre> </div> 2. Run the cvfsck utility to verify and potentially correct the metadata. 3. Start the file system.

Operating System	Change Request Number	Service Request Number	Description/Workaround
Linux	47954	n/a	<p>The Safari browser becomes unresponsive when you attempt to configure an Email server using the StorNext GUI.</p> <p>Workaround:</p> <p>To workaround this issue, perform the following procedure:</p> <ol style="list-style-type: none"> 1. Shut down the Safari browser window(s). 2. Restart the Safari browser, and then retry the operation. 3. Uncheck the Verify SMTP Server Connectivity box, and then retry the operation. 4. Set Authentication to NONE, and then retry the operation. 5. Disable the Safari User names and passwords AutoFill under Safari > Preferences > AutoFill, and then retry operation.

StorNext Installation, Replication, HA, and Other Known Issues

The table below lists known issues specific to StorNext installations, data replication, HA systems, and other areas.

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	68849	n/a	<p>After an appliance firmware upgrade, you might be unable to use previously functioning tape devices because the <code>lin_tape</code> device driver was automatically unloaded during the upgrade.</p> <p>Workaround:</p> <p>To workaround this issue, rebuild the <code>lin_tape</code> device driver as shown in the following example:</p> <pre>rpm -e lin_taped rpm -e lin_tape rpmbuild --rebuild /root/lin_tape-1.76.06-1.src.rpm rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-1.76.06-1.x86_64.rpm rpm -ivh /root/lin_taped-1.76.0-rhel6.x86_64.rpm</pre>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	73557	n/a	<p>Beginning with StorNext 6.2.0, if you change the system clock backwards while MySQL is running and then try to shutdown StorNext, the MySQL database may block when trying to stop until the current time matches the time just prior to changing the time on the system.</p> <p>i Note: If you use a Quantum appliance and use the StorNext GUI to change the time, there is no issue.</p> <p>Workaround:</p> <p>This workaround applies to the following systems:</p> <ul style="list-style-type: none"> • If you use a customer supplied system. • If you use a Quantum appliance and do not use the StorNext GUI to change the system clock. <p>To workaround this issue, perform the following procedure.</p> <ol style="list-style-type: none"> 1. Shut down StorNext. 2. Change the time. 3. Restart StorNext. <p>i Note: If shutting down StorNext is not an option, then wait for the system time to match the time prior to the time change, and then shut down to avoid this issue. For example, if you adjust the system time back one hour, then you must wait one hour before you halt or reboot your system.</p>
All	73945	n/a	<p>If you have a Replication, Deduplication, or Object Storage license installed in <code>/usr/cvfs/config/license.dat</code>, then during HA upgrade time, the first time one of the nodes is activating and becoming primary on the new StorNext version, snpolicyd could be restarted. The restart might take up to 5 minutes to stop.</p> <p>Workaround:</p> <p>Execute the following commands on both nodes in order to shorten the snpolicyd activation time before upgrading to a new version of StorNext:</p> <ol style="list-style-type: none"> 1. <code>mkdir -p /usr/cvfs/update/config</code> 2. <code>touch /usr/cvfs/update/config/objs.conf</code> 3. <code>touch /usr/cvfs/config/objs.conf</code>

Operating System	Change Request Number	Service Request Number	Description/Workaround
All	74822	498657	<p>A stored procedure in the MySQL database was using too small of a data type when constructing an admin alert email notification. If the system hostname exceeded 28 characters, an overrun would occur resulting in a failed update into the database preventing an email notification from being generated. Prior to StorNext 6.2.0, this overrun would be silently converted to the correct size and insert without issue. However, as of StorNext 6.2.0, this overrun is no longer allowed by the updated version of the percona-server MySQL database, and will result in an error in /usr/adic/mysql/logs/ODBCError.log that prevents the email notification from being generated.</p> <p>Workaround:</p> <p>To work around the issue, the stored procedure needs to be updated to use the correct size for the email subject line, and then the stored procedure needs to be reloaded. Run the following steps on the MDC. In the case of an HA environment, the following procedure needs to be run on the primary MDC:</p> <ol style="list-style-type: none"> 1. Verify that the MDC is primary and that MySQL is running using: <pre># /usr/cvfs/bin/snhamgr status # /usr/adic/mysql/bin/mysql_control status</pre> 2. Make a backup of the existing procedure file: <pre># cp /usr/adic/SRVCLLOG/dbSchema/Admin_Log_Insert.sql /usr/adic/SRVCLLOG/dbSchema/Admin_Log_Insert.sql.orig</pre> <p>Note: Although it might not appear as such in this column, the command is a continuous string. If you copy and paste from this column, you must verify that the command does not have any breaks. This applies to any command string in subsequent steps.</p> <ol style="list-style-type: none"> 3. Modify the procedure:

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

```
# sed -i 's/VARCHAR(64)/VARCHAR(256)/'
/usr/adic/SRVCLLOG/dbSchema/Admin_Log_
Insert.sql
```

- Verify that the change looks correct. **There should only be one change!**

```
# diff -u /usr/adic/SRVCLLOG/dbSchema/Admin_
Log_Insert.sql.orig
/usr/adic/SRVCLLOG/dbSchema/Admin_Log_
Insert.sql
--- /usr/adic/SRVCLLOG/dbSchema/Admin_Log_
Insert.sql.orig 2019-06-07 11:26:38.376994486
-0500
+++ /usr/adic/SRVCLLOG/dbSchema/Admin_Log_
Insert.sql 2019-06-07 11:28:02.178203087 -0500
@@ -7,7 +7,7 @@
tproc:BEGIN
declare l_body VARCHAR(4000);
- declare l_subject VARCHAR(64);
+ declare l_subject VARCHAR(256);
declare l_attachment VARCHAR(1024);
declare l_subtype VARCHAR(32);
declare l_beg_pos INT;
```

- Drop and reload the stored procedures:

```
# /usr/adic/SRVCLLOG/bin/dbUtil -r
/usr/adic/SRVCLLOG/dbSchema
```

- Restart **SRVCLLOG**:

Operating System	Change Request Number	Service Request Number	Description/Workaround
			<pre># /usr/adic/SRVCLLOG/bin/SRVCLLOG_control restart</pre>
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>Note: This does not otherwise affect the functionality of the HGST ActiveScale™ P100 (Quantum Lattus P100) integrated object storage system.</p> <p>Workaround: There is currently no workaround for this issue. If you experience this issue, contact Quantum Technical Support.</p>

Contacting Quantum

Contacts

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

<http://www.quantum.com/aboutus/contactus/index.aspx>

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

Region	Support Contact
For worldwide support: http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support	

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com