StorNext 4.0 Release Notes

Product	StorNext® 4.0
Date	March 2010

Contents

Document 6-00431-27 Rev A

Purpose of this Release

StorNext 4.0 is a major release that provides several important new features, as well as an improved graphical user interface.

This document also lists issues that were resolved for this release, currently known issues, and known operating limitations.

Visit www.quantum.com/ServiceandSupport for additional information and updates for StorNext.

Enhancements and Improvements

This section describes the new features in StorNext 4.0.

New Graphical User Interface

StorNext 4.0 incorporates a new, improved graphical user interface which greatly simplifies and streamlines data entry and configuration. The interface utilizes tabs on various screens to make configuration and data entry more intuitive.

The new interface also includes "mouseover help" which displays helpful hints when you move the mouse over a field. If you need more detailed information, launch the online help system by clicking the question mark icon on any screen. Detailed, context-sensitive help pertaining to that screen appears. In addition, you can also access the regular StorNext documentation from the Help menu.

The Configuration Wizard, which guides you through the process of configuring StorNext and many of its main features, has also been enhanced for release 4.0. As before, each of the Configuration Wizard's steps has an equivalent command which you can access at any time without using the wizard.

Data Replication

Data replication is the process of moving duplicated data from one or more source servers to one or more destination (target) servers. Replication is an important feature for backup and archive purposes.

The Replication Multilink Feature

The StorNext 4.0 Replication feature includes *Multilink* capability, which enables you to configure an aggregation of multiple network interface streams.

This feature provides bandwidth throttling and link aggregation capabilities, and is configured on the StorNext GUI's the Tools > Replication > Bandwidth screen.

The following table provides the details specific to Multilink behavior in the StorNext 4.0 release:

Table 1 Multilink Behavior

Change Request Number	Description	Workaround
30746, 30921	Once enabled, the Multilink feature cannot be disabled from the StorNext GUI.	Contact Quantum Technical Support for assistance with disabling the Multilink feature
30874	Bandwidth control is inaccurate when low throughput rates are specified.	When configuring bandwidth throttling, specify values of at least 40 Mbit/s.
30920	In HA configurations, the StorNext GUI may display errors about virtual IPs (vIPs) when configuring Multilink if the peer system is down.	Ensure the standby MDC is up when configuring Multilink.
30750, 30832	In rare cases, when Multilink is enabled the blockpool process may crash during connection establishment between the replication source and target, leading to replication failure.	If replication was initiated manually, rerun. Otherwise, replication will be restarted automatically according to the configured schedule. Also, inspect `/usr/cvfs/debug` for blockpool core files. These can be identified using `"file /usr/cvfs/debug/core*"` If any core files are found, provide them to Quantum Technical Support for analysis.
29984	The StorNext GUI does not indicate whether Multilink is enabled.	Examine `/usr/cvfs/config/blockpool_config.txt` and look for lines containing `MultiLinkConfig` and `MultiLinkInterface`. Values other than "None" indicate that Multilink is configured.

In addition to the above issues, please note that the link aggregation capability is not fully fault tolerant. If a network error occurs, this will cause an in-progress replication to fail, requiring it to be rerun if it was manually initiated. Or, if a scheduled replication attempt fails, file transfer may be delayed until replication is rescheduled per the specified policy. Quantum recommends using the Multilink feature only in environments with stable networks.

Data Deduplication

Data deduplication is a technology that conserves storage space by eliminating unnecessary and redundant data duplication. When deduplication is enabled, only the blocks of changed data are stored rather than the entire file.

Deduplication can also be used with replication, so that only changed blocks of information are replicated from the source to the target. Depending on your system configuration, the bandwidth and storage savings could be dramatic.

Note: Beginning with StorNext 4.0, with the introduction of the deduplication feature DDisk is no longer supported. *Any existing DDisks will not be automatically transferred to 4.0 when you upgrade, and the installation will fail if there are DDisks.* For assistance transitioning from DDisk to deduplication, contact Quantum Technical Support.

HA Support

The StorNext High Availability (HA) feature has been enhanced for StorNext 4.0. This feature allows you to configure and operate a redundant server that can quickly assume control of the StorNext file systems and management data in the event of certain software, hardware and network failures on the primary server.

Power Bricks No Longer Required for HA

Power bricks are no longer required or used for StorNext HA configurations. If desired, power bricks may be removed after you upgrade to StorNext 4.0. (If you use power bricks for non-HA purposes, there is no need to remove power bricks after upgrading to StorNext 4.0.)

Distributed Data Mover

The Distributed Data Mover (DDM) feature enables you to distribute data movement operations to client machines from the metadata controller, which can improve the overall throughput of data movement to archive tiers of storage. Previously, the data mover process, fs_fmover, ran only on the metadata controller, allowing up to one fs_fmover process per tape drive or storage disk (SDisk) stream to run at one time.

Note: The DDM feature supports only storage disks on StorNext file systems, not on NFS.

Web Services API

Beginning with StorNext 4.0, a Web Services version of StorNext API (WS-API) is installed along with the StorNext 4.0 Linux GUI binary, which is installed by default.

WS-API provides basic control over StorNext Storage Manager systems to track media and drives and to store/truncate/retrieve files from any computer capable of creating a Web Services connection, which includes Windows, Macintosh, and Linux based systems, among others.

WS-API is different from the StorNext API (SNAPI) product, which allows you to run APIs for StorNext File System and Storage Manager. For more information about SNAPI, contact your Quantum Sales representative.

Added Library Support

StorNext 4.0 adds support for the Quantum Scalar i40 and Scalar i80 libraries.

Added Platform Support

StorNext 4.0 adds support for the following platforms since the 3.5.1 release:

- Windows Vista SP 2 x86 32-bit and x86 64-bit
- Windows Server 2008 R2 and SP2 x86 32-bit and x86 64-bit
- Windows 7 x86 32-bit and x86 64-bit
- Red Hat Enterprise Linux 5 Update 4 x86 64-bit
- SLES 10 SP3 x86 32-bit and x86 64-bit
- SLES 11 x86 64-bit
- IBM AIX 6.1 64-bit Power Architecture
- HP-UX 11i v3 Itanium 64-bit
- MacOS X 10.6.1 with Xsan 2.2 and StorNext MCD x86 32-bit and x86 64-bit
- MacOS X 10.6.1 with Xsan 2.2.1 and StorNext MCD x86 32-bit and x86 64-bit

To determine which StorNext components and features are supported for these and all other platforms, refer to Operating System Requirements on page 14.

Equivalent Platforms

StorNext 4.0 also adds support for the following platforms which are equivalent to Red Hat Enterprise Linux releases. These new platforms are supported to the degree that the issue can be demonstrated on the equivalent Red Hat Enterprise Linux release by the customer.

- CentOS (based on RHEL5 Update 4) x86 64-bit
- Scientific Linux (based on RHEL5 Update 3) x86 64-bit
- Oracle Linux (based on RHEL5 Update 2) x86 64-bit

Added Tape Drive Support

Support for the Sun/StorageTek 9840D and T10000 Rev B tape drives has been added. For more information, see table 4 on page 19.

Other Changes in This Release

This section contains a summary of changes in StorNext 4.0 which are not necessarily enhancements or new features.

StorNext Licensing

StorNext licensing has changed significantly from previous releases. Beginning with release 4.0, StorNext licenses are now required and enforced for several features. If you do not have the appropriate license, you will not be able to use that StorNext feature.

Here is a list of StorNext licenses:

- **Failover (HA)**: A Failover (High Availability) license is required if you plan to use StorNext's HA failover features.
- **Maintenance**: A Maintenance license verifies that your site has purchased StorNext upgrade licenses, and is required for StorNext upgrades. It is used at run time to verify the build date in the software matches what was purchased.
- **File System**: A File System license enables you to create and modify StorNext-supported file systems.
- **LAN Client**: You must have a Distributed LAN Client license for LAN clients you use with StorNext (in addition to any SAN clients). A single license specifies how many clients you are authorized to use with StorNext.
- **Storage Manager**: A Storage Manager license provides full access to StorNext's Storage Manager features that are not licensed separately.
- **Replication**: A replication license is required if you want to use StorNext's Data Replication feature.
- **Deduplication**: A deduplication license is required if you want to use StorNext’s Data Deduplication (blockpool) feature.
- **Vaulting**: A Vaulting license provides the ability to move seldom-used media to a manual archive vault, freeing room for media in the managed archives.
- **Storage Disk**: You must have a Storage Disk license to be able to configure and use StorNext storage disks.
- **Checksum**: A Checksum license enables you to verify data integrity by ensuring that the checksum created when data was stored matches the checksum upon data retrieval.
- **Distributed Data Mover (DDM)**: A license is required for using DDM if you plan to use additional machines besides the primary server to distribute data streams.

Some StorNext feature license have a license expiration date and a limit. The expiration date and limit are shown on the StorNext Setup > License screen. Many feature licenses include unlimited usage. For features with limits, the number shown in the **Limit** column on the Setup > License screen refers to the following:

- **File System**: The number displayed is the maximum number of SAN clients allowed
- **LAN Clients**: The number displayed is the maximum number of LAN clients allowed
- **Storage Manager**: The number displayed is the licensed capacity for files being managed by the SNSM software.

  The current used capacity for the Storage Manager feature is determined as follows: It is the sum of all managed files that have been copied to Storage Manager media. This includes tape media and also storage disk. (The capacity calculation does not include files on tape media that have been moved to a vault archive.)

  Note the following about the used space that is tracked:

  - The used space includes the sizes of all files stored to the media. This includes multiple copies of files, if present, as well as old versions available for recovery.
  - The used space includes the size of the header information written to the media with the file. This header contains information used to verify file information at retrieve time, etcetera.
  - If compression is enabled, the used space is determined by the native size of the files and not the compressed size.
  - For tape media, the used space includes "dead space" for files that have been removed or modified, and for which fsclean has been run to clean up knowledge of the file in the Storage Manager database.

  Note: The capacity for storage disks does not include "dead space" but does include *all* data on the file system where the storage disk has been configured, not just files copied to the file system by the Storage Manager. To maximize the licensed Storage Manager capacity, the storage disk file systems should be restricted to Storage Manager data only. If your storage disk contains 'user' data you should consider moving that data to an alternate location prior to performing a StorNext upgrade.

  If you are unsure about the location of the Storage Manager data on a file system run the command `fsdiskcfg` with no arguments. This command produces a report on the configured storage disks and the location of the managed data on each file system. The 'Path' column in the command output indicates the directory containing the managed data.

- **Deduplication**: The number displayed is the size of the blockpool. The capacity is tracked to the nearest terabyte.
- **Distributed Data Mover (DDM)**: The number displayed is the maximum number of clients that can be used to run mover processes.

- **Maintenance**: The number displayed is the expiration date of the service contract which entitles software upgrades. StorNext can be upgraded only to versions that were released by Quantum prior to this expiration date. The software version that is running on the displayed expiration date will continue to run without limitations, but attempts to upgrade to a version with a release date after the expiration date will fail.

If the `license.dat` file does not contain permanent licenses, StorNext produces an auto-generated license with an expiration date for all StorNext products and features except Deduplication. In some cases Quantum may provide evaluation licenses for features. Evaluation licenses also have a fixed expiration date.

Beyond the evaluation period, you must have a permanent license to configure or use StorNext features.

Note: You cannot mix auto-generated/evaluation and permanent licenses. Once you enter one or more permanent licenses in StorNext, all auto-generated/evaluation licenses are deleted. If you are using an auto-generated or evaluation license period to evaluate new features, be aware that you will lose any remaining time on your temporary licenses as soon as you enter one or more permanent license.

After you have entered permanent licenses, you should not install an evaluation license.

For more information about licenses and obtaining a permanent license, see the StorNext User's Guide or the online help.

Licensing and Upgrades

The new licensing implementation affects StorNext upgrades, both for release 4.0 and future releases. Be aware of the following upgrade-related implications and plan accordingly:

- A current (non-expired) Maintenance license is required to perform a StorNext upgrade. This means you must contact Quantum Technical Support for a Maintenance license before you can upgrade to StorNext 4.0.
- The Maintenance license provided by Quantum Technical Support must be put into place prior to the upgrade, or you will not be allowed to proceed with the upgrade. This step is done by manually editing the `license.dat` file because the StorNext GUI has not been installed and therefore you cannot enter licenses through the GUI. This applies to upgrades to StorNext 4.0 and also to future upgrades.
- For future upgrades, for any StorNext feature or component currently in use, you must have a license in place prior to the upgrade. Otherwise, you will not be allowed to proceed with the upgrade.
- For future upgrades, the software build date must precede the Maintenance license's expiration date.
- The Maintenance license must remain in place even after expiration to allow the StorNext software to run, because it indicates the build date of the software purchased.

- If you are ready to upgrade and then notice that the Storage Manager capacity has been exceeded, you can follow the procedure below to free up capacity to bring it under the licensed value. These steps will clean up "dead space" on tape media, and do not apply to storage disks.

  1 Run the `fsmedcopy` command with no arguments. Running command generates a report of media and wasted space. The report looks similar to this:

  ```
  % fsmedcopy
  ------------------------------------------------------------
   Media Fragmentation Report          Thu Feb  4 15:18:14 2010

   Media ID           Fill Level     Wasted Space
  ------------------------------------------------------------
   000099              10.00             0.00
   000098              30.00             5.00
   000096              70.99            44.98
   000099               0.12            99.65
  ```

  The "Fill Level" column shows how full the media is. The "Wasted Space" column indicates how much of the used space is wasted, NOT the percentage of the entire tape.

  Media with high percentages for both values (such as media 000096 in the example) are good candidates for eliminating dead space.

  2 For each media you identify as having significant dead space, run this command:

  ```
  % fsmedcopy -r <mediaid>
  ```

  Note: On a large system, the `fsmedcopy` report can take a long time to complete. If your site is dynamic (that is, it has a lot of file removes and updates,) you can skip step 1 and just pick media you know are full, and then run the command `fsmedcopy -r` against those media.

StorNext Media on DVD

Beginning with StorNext 4.0, installation and upgrade media is shipped on DVDs. (Media was previously shipped on CDs.) Both software and documentation are now stored on a single DVD, eliminating multiple discs for different operating systems.

If you plan to install from media, you must have a DVD ROM drive to perform the installation or upgrade.

Revised Journal Size Requirements

Beginning with StorNext release 4.0, all new file systems must have a journal size that is at least 1024 times the file system blocksize, up to 128MB. For example, a file system with 64k FS Blocksize must have a journal that is at least 64M.

Quantum also recommends adjusting any upgraded file systems to meet this new size requirement as well.

Tape Devices and Persistent SCSI Reserve

StorNext 4.0 incorporates support for SCSI-3 persistent reservations for tape devices. SCSI-3 persistent reservation commands attempt to prevent unintended access to tape drives that are connected by using a shared-access technology such as Fibre Channel. Access to a tape drive is granted based on the host system that reserved the device. SCSI-3 persistent reservation enables access for multiple nodes to a tape device and simultaneously blocks access for other nodes.

Quantum recommends enabling SCSI-3 persistent reservations. Refer to parameter `FS_SCSI_RESERVE in /usr/adic/TSM/config/fs_sysparm.README` to direct the StorNext Manager to use SCSI-3 persistent reservations.

One implication of using SCSI-3 reservations is that all tape devices used must support the PERSISTENT RESERVE IN/OUT functionality as described in SCSI Primary Commands-3 standard (SPC-3).

A supported upgrade to StorNext 4.0 will maintain any existing SCSI tape reservation settings (SCSI-2 reservations or no SCSI tape reservations). A *new* StorNext 4.0 installation will by default use SCSI-3 persistent reservations.

The StorNext Distributed Data Mover feature requires that SCSI-3 persistent reservations be used.

For more information about SCSI reserve and instructions on determining your drive version, refer Chapter 5, “Storage Manager Tasks” in the StorNext User’s Guide.

New “localhost” Requirement

Beginning with StorNext release 4.0, the StorNext GUI requires that “localhost” be a resolvable IP address that resolves to an IP in the range 127.0.0.0/8.

IPV6 Support in StorNext 4.0

StorNext 3.5 provided partial support for IPv6 in StorNext File System. StorNext 4.0 provides the same IPv6 support as in StorNext 3.5, but does not provide IPV6 support for StorNext Storage Manager, replication, deduplication, and other features. Future StorNext releases may included expanded IPv6 support.

XML Configuration File Format for Linux

There is a new file system configuration format written in XML for linux MDCs. This is the default format, and is required to use the StorNext GUI interface. The new format is identified by the suffix .cfgx. The old .cfg format is deprecated on Linux.

Windows continues to use the .cfg format, but will switch to using the XML format in a future StorNext release.

Refer to the StorNext Installation Guide for examples of both .cfgx and .cfg configuration files.

Optimistic Allocation

Starting with StorNext 4.0, the InodeExpand parameters (InodeExpandMin, InodeExpandInc, and InodeExpandMax) in the file system configuration file have been deprecated and replaced by a simple formula that should work better in most cases, especially with very large files.

The InodeExpand values are still honored if they are in the .cfgx file, but the StorNext GUI no longer lets you set these values. Furthermore, when converting to StorNext 4.0, during the .cfg to .cfgx conversion process, if the InodeExpand values in the .cfg file are found to be the default example values, these values are not set in the new .cfgx. Instead, the new formula is used.

The original InodeExpand configuration was difficult to explain, which could lead to misconfigurations that caused either over or under allocations (resulting in wasted space or fragmentation,) which is why the new formula seeks to use allocations that are a percentage of the existing file's size to minimize wasted space and fragmentation.

For more information about Optimistic Allocation, refer to the *StorNext File System Tuning Guide*.

Cvfail and cvpower Scripts Removed

Earlier versions of High Availability StorNext (HA) used the `cvfail` and `cvpower` programs to perform a power-cycle reset of the peer server computer and to start Storage Manager components if they were configured. These scripts are obsolete in StorNext 4.0. The FSM process no longer calls out to the `cvfail` program, and upgrade instructions include removing `cvfail` and `cvpower`.

StorNext 4.0 Server for Linux uses the `/usr/cvfs/lib/snactivated.pl` script to mount file systems and to start the Storage Manager components as needed. This script differs in use from `cvfail` because it is called at the end of every activation of every file system rather than at the beginning of activation of the first file system after a CVFS restart. Proper functioning of HA depends in large part on the correct execution of the snactivated script. Local modification of the script is not advised and not supported.

Unmounting or Stopping a File System is Not Persistent

The file system stop action in the StorNext GUI does not change the configuration. That is, the file system will start again when StorNext is restarted.

This is a change from the 3.X StorNext GUI, which edited the `/usr/cvfs/config/fsmlist` file so that a file system would remain stopped until the configuration was changed to start it again. (As in 3.x, for 4.0 the file system must still be manually unmounted before the file system is stopped).

Entries in the `fsmlist` can be commented-out by inserting a leading '#' character to perpetuate the stoppage of a file system across restarts of StorNext. When this is done, the file-system configuration will not appear in the GUI file-system display.

In HA configurations the file system may continue to run on the peer server computer, or it may start up briefly and then stop as the `fsmlist` is mirrored from the Primary to the Secondary server. One way to avoid this is to follow these steps for changing the `fsmlist` file:

1. Stop CVFS on the Secondary.
2. Unmount and stop the file system on the Primary.

3 Comment-out the file system entry in the `fsmlist` file on both server computers.

4 Start CVFS on the Secondary.

Subtree Check Option in NFS No Longer Supported

Although supported in previous StorNext releases, the `subtree_check` option (which controls NFS checks on a file handle being within an exported subdirectory of a file system) **is no longer supported** as of StorNext 4.0.

No Cryptographic Changes

There is no change to cryptographic functionality in StorNext release 4.0.

Discontinued Support on Some Platforms

Some StorNext services that were supported on various platforms in previous StorNext releases are no longer supported in StorNext 4.0. These services will continue to be supported for previous StorNext releases, but going forward beginning with release 4.0 will not be supported.

Table 2 shows the StorNext services for which support is discontinued as of StorNext 4.0. These services cannot be upgraded to StorNext 4.0.

Table 2 Discontinued Platforms

StorNext 4.0 Discontinued Support By OS Platform		
Operating System	**Platform(s)**	**StorNext Services No Longer Supported**
Windows 2003 Server	X86 32-bit	MDC Server and Distributed LAN Server
Red Hat Enterprise Linux 4 Update 6 (2.6.9-67.EL) and Update 7 (2.6.9-78.EL)	x86 32-bit	MDC Server and Distributed LAN Server
Red Hat Enterprise Linux 4 Update 8 (2.6.9-89 EL)	x86 32-bit	MDC Server, Distributed LAN Server and Storage Manager
Red Hat Enterprise Linux 4 Update 8 (2.6.9-89 EL)	x86 64-bit	Storage Manager
SLES 10 SP1 (2.6.16-46-0.12) and SP2 (2.6.16.60-0.27)	Itanium 64-bit	MDC Server, File System SAN Client and File System LAN Client
SLES 11 2.6.27.19-5	Itanium 64-bit	MDC Server and File System SAN Client
Sun Solaris 10 Generic 120011-14	Sparc 64-bit	MDC Server and Storage Manager
IBM AIX 5.3	64-bit Power Architecture	File System SAN Client
HP-UX 11i v2	Itanium 64-bit	File System SAN Client
SGI IRIX 6.5.30	64-bit MIPS	File System SAN Client

Note: MDC server, Distributed LAN and Storage Manager support has been dropped for *all* 32-bit Red Hat 4 systems (not just those running update 6 and 7).

LIkewise, Storage Manager support has been dropped for all Red Hat 4 systems (not just those running update 8).

Operating System Requirements

Table 3 shows the operating systems, kernel versions, and hardware platforms that support the following:

- MDC Servers
- File System SAN Clients
- Distributed LAN Servers
- File System LAN Clients
- Storage Manager
- Distributed Data Mover
- Replication/Deduplication Server

Note: When adding StorNext Storage Manager to a StorNext File System environment, the metadata controller (MDC) must be moved to a supported platform. If you attempt to install and run a StorNext 4.0 server that is not supported, you do so at your own risk. Quantum strongly recommends against installing non-supported servers.

Table 3 StorNext Supported OSes and Platforms

StorNext 4.0 Supported Operating Systems and Platforms									
Operating System	**Kernel or Release**	**Platform**	**MDC Server**	**File System SAN Client**	**Distributed LAN Server**	**File System LAN Client**	**Storage Manager**	**Distributed Data Mover**	**Replication / Dedup Server**
Windows 2003 Server	R2 SP2*	x86 32-bit		✓		✓			
		x86 64-bit	✓	✓	✓**	✓			
Windows XP	SP2	x86 32-bit		✓		✓			
		x86 64-bit		✓		✓			
	SP3	x86 32-bit		✓		✓			
		x86 64-bit		✓		✓			
Windows Vista	SP1	x86 32-bit		✓		✓			
		x86 64-bit		✓		✓			
	SP2	x86 32-bit		✓		✓			
		x86 64-bit		✓		✓			
Windows Server 2008	SP1	x86 32-bit		✓		✓			
		x86 64-bit	✓	✓	✓**	✓			
	R2	x86 32-bit		✓		✓			
		x86 64-bit	✓	✓	✓**	✓			
	SP2	x86 32-bit		✓		✓			
		x86 64-bit	✓	✓	✓**	✓			
Windows 7		x86 64-bit		✓		✓			
		x86 32-bit		✓		✓			

* StorNext supports and has been tested using R2 SP2 since StorNext release 3.1.2.

** Windows Distributed LAN Server supports up to 128 distributed LAN clients.

StorNext 4.0 Supported Operating Systems and Platforms (Continued)

Operating System	Kernel or Release	Platform	MDC Server	File System SAN Client	Distributed LAN Server	File System LAN Client	Storage Manager	Distributed Data Mover	Replication / Dedup Server
RHEL 4 ‡	2.6.9-67.EL (Update 6) †	x86 32-bit		✓		✓			
	2.6.9-78.EL (Update 7) †	x86 32-bit		✓		✓			
	2.6.9-89 EL (Update 8)	x86 32-bit		✓		✓			
	2.6.9-67.EL (Update 6) †	x86 64-bit	✓	✓	✓	✓			
	2.6.9-78.EL (Update 7) †	x86 64-bit	✓	✓	✓	✓			
	2.6.9-89 EL (Update 8)	x86 64-bit	✓	✓	✓	✓			
RHEL 5 ‡	2.6.18-92.EL (Update 1) †	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.18-92.EL (Update 2) †	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.18-128.EL (Update 3) †	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.18-164.EL (Update 4)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓

† **All releases of RHEL4 and RHEL5 except RHEL4 Update 8 and RHEL5 Update 4 have a possible silent data corruption issue as documented in Product Alert #20. Quantum recommends that users migrate to RHEL4 Update 8 or RHEL5 Update 4 as soon as possible.** Also, note that the "Xen" virtualization software is not supported for RHEL 4 and RHEL 5.

‡ HBA multipath customers: please verify with your HBA vendor that your current multipath driver is supported for any planned Linux OS version/update/service pack level. If your driver is not supported for your planned Linux OS version/update/service pack, the StorNext client or server may not be functional after your Linux upgrade.

Note: The RHEL and SLES kernel levels listed indicate which kernel levels were used for the majority of testing. In general, other kernel levels within the same service pack are supported unless otherwise noted.

Note: For systems running Red Hat Enterprise Linux version 4 or 5, before installing StorNext you must first install the kernel header files (shipped as the kernel-devel-smp or kernel-devel RPM).

For systems running SUSE Linux Enterprise Server, you must first install the kernel source code (typically shipped as the kernel-source RPM).

Caution: Red Hat 5 ships with Security-Enhanced Linux (selinux) enabled by default. To ensure proper StorNext operation, you must not install Red Hat 5 with selinux enabled. That is, selinux must be off, or the file system could fail to start.

If Red Hat 5 has already been installed with SELINUX enabled, edit the file `/etc/selinux/config` and change the line `SELINUX=enforcing` to either `SELINUX=permissive` or `SELINUX=disabled`. Refer to Red Hat 5 documentation for more information.

StorNext 4.0 Supported Operating Systems and Platforms (Continued)									
Operating System	**Kernel or Release**	**Platform**	**MDC Server**	**File System SAN Client**	**Distributed LAN Server**	**File System LAN Client**	**Storage Manager**	**Distributed Data Mover**	**Replication / Dedup Server**
SLES 10 †† ‡ ***	2.6.16.46-0.12 (SP1)	x86 32-bit		✓		✓			
	2.6.16.60-0.27 (SP2)	x86 32-bit		✓		✓			
	2.6.16.60-0.54.5 (SP3)	x86 32-bit		✓		✓			
	2.6.16.46-0.12 (SP1)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.16.60-0.27 (SP2)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.16.60-0.54.5 (SP3)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
SLES 11 †† ‡ ***	2.6.27.19-5	x86 64-bit		✓		✓			
Sun Solaris 10	Generic 120011-14	Sparc 64-bit		✓					
	Generic 127128-11	Opteron x86 64-bit		✓		✓			
		Intel x86 64-bit		✓		✓			
IBM AIX	6.1	64-bit Power Architecture		✓					
HP-UX	11i v3 (Requires the "0909 Patch set")	Itanium 64-bit		✓					
Apple (using Xsan)	MacOS X 10.5.5 w/Xsan 2.1.1, StorNext MDC	x86 32-bit		✓					
	MacOS X 10.6.1 w/Xsan 2.2, StorNext MDC	x86 32-bit		✓					
		x86 64-bit		✓					
	MacOS X 10.6.1 w/Xsan 2.2.1, StorNext MDC	x86 32-bit		✓					
		x86 64-bit		✓					
The following platforms have equivalent RedHat releases, and are supported only if the issue can be reproduced on the equivalent RedHat release.									
CentOS	Based on RHEL5 Update 4	x86 64-bit		✓		✓			
Scientific Linux‡‡‡	Based on RHEL5 Update 3	x86 64-bit		✓		✓			
Oracle Linux‡‡‡	Base on RHEL5 Update 2	x86 64-bit		✓		✓			

†† SLES 10 SP1 (and earlier) and SP2 kernels earlier than 37 are sensitive to the same silent data corruption issue documented in Product Alert #20. The problem has been fixed in SP2 that includes level 2.6.16.60-0.37_f594963d, in SLES 10 SP3, and in the SLES 11 releases. There is no recommended workaround at this time.

‡ HBA multipath customers: please verify with your HBA vendor that your current multipath driver is supported for any planned Linux OS version/update/service pack level. If your driver is not supported for your planned Linux OS version/update/service pack, the StorNext client or server may not be functional after your Linux upgrade.

*** A "roll" of a particular digit is not indicative that a new SLES service pack has been declared by Novell. The kernel revisions listed in this document are typically (but not always) the first kernel revision of the service pack. Later revisions within the service pack are typically, but not always, supported.

‡‡‡ These platforms are not specifically tested for StorNext releases. Support for these releases will be at the equivalent RHEL or SLES kernel service pack release, and issues reported against these platforms must be reproducible on the equivalent base RHEL or SLES release for additional support to apply.

Note: Although SGI IRIX clients cannot be upgraded to StorNext 4.0, StorNext 3.5.1 SGI IRIX clients may be used with a StorNext 4.0 MDC.

Note: GNU tar is required on Solaris systems. In addition, for systems running Solaris 10, install the Recommended Patch Cluster (dated March 10, 2006 or later) before installing StorNext.

To enable support for LUNs greater than 2TB on Solaris 10, the following patches are required:

- 118822-23 (or greater) Kernel Patch
- 118996-03 (or greater) Format Patch
- 119374-07 (or greater) SD and SSD Patch
- 120998-01 (or greater) SD Headers Patch

Supported Libraries and Tape Drives

Libraries and tape drives supported for use with StorNext 4.0 are presented in Table 4. Where applicable, minimum firmware levels for libraries are provided.

Table 4 StorNext Supported Libraries and Tape Drives

StorNext 4.0 Supported Libraries and Tape Drives

Vendor Library Family	Libraries	Enforced Minimum / Recently Tested Library Firmware Level	Drive Types	Enforced Minimum / Recently Tested Drive Firmware Level	Notes
Quantum / ADIC	Scalar i500	Minimum: 140G	IBM LTO-1		Library firmware upgrade may be required for LTO-3 WORM support 420G.GS00400
			IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
			IBM LTO-3 WORM		
			IBM LTO-4 WORM		
			HP LTO-4		
	Scalar i2000 ‡	Minimum: 120A Minimum (IBM LTO-3, IBM LTO-3 WORM): 300A Minimum (IBM LTO-4, IBM LTO-4 WORM) 540A Minimum: 7404 i/o blades require 590A (library firmware version i6.5)	IBM LTO-1		
			IBM LTO-2		
			IBM LTO-3	See library firmware requirement	
			IBM LTO-4		
			IBM LTO-3 WORM		
			IBM LTO-4 WORM		
			HP LTO-4		
			HP LTO-4 WORM		
			DLT-S4	Minimum: 1F1F	
	Scalar 24	Minimum: 107A.GY0002	IBM LTO-1		Not including WORM
			IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
	Scalar 50	Minimum: 002A	HP LTO-4		
	Scalar 100	Minimum: 2.05.0003	IBM LTO-1		Not including WORM 2.10.0013 is bad firmware
			IBM LTO-2		
			IBM LTO-3		
			AIT-2		
	Scalar 1000	Minimum: 3.00.0017	IBM LTO-2		Must use SDLC/DAS, SDLC/SCSI Target Mode or Native SCSI
			IBM 3590B1A		
			AIT-1		

‡ Before using DLT cleaning with DLT-S4 or SDLT 600 drives, configure the library (Scalar i2000 or PX720) to disable reporting of the media ID. If media ID reporting is not disabled, StorNext will not recognize the cleaning media (SDLT type 1).

StorNext 4.0 Supported Libraries and Tape Drives (Continued)					
Vendor Library Family	**Libraries**	**Enforced Minimum / Recently Tested Library Firmware Level**	**Drive Types**	**Enforced Minimum / Recently Tested Drive Firmware Level**	**Notes**
Quantum / ADIC	Scalar 10000	Minimum: 110A.00001	IBM LTO-1		Must use SDLC/DAS, SDLC/SCSI Target Mode or Native SCSI
			IBM LTO-2		
			IBM LTO-3	See library firmware requirement	
			IBM LTO-4		
			IBM LTO-3 WORM		
			AIT-2		
			AIT-2 WORM		
			IBM 3592		
	Scalar i40	Minimum: 101G.GS005 Recently Tested: 101G.GS005	HP LTO-4		
	Scalar i80	Minimum: 101G.GS005 Recently Tested: 101G.GS005	HP LTO-4		
	PX500	Minimum: 001A	HP LTO-3		Not including WORM 3.0
	PX720 ‡	Minimum 4.00	HP LTO-2		Not including WORM
			HP LTO-3		
			DLT-S4		
	DXI 7500	Minimum: N / A Recently Tested: 05.02.084	Supported emulations include: DLT7000, SDLT320, SDLT600, DLT-S4, Quantum/Certance LTO-2 and, 3, HP LTO-1, 2, 3 and 4, IBM LTO-1, 2, 3 and 4		
Dell	PV136T	Minimum: 3.11	IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
HP	ESL E Series	Minimum: 4.10	HP LTO-3		Including Regular and WORM
			HP LTO-4		
	MSL 6000	Minimum: 0507	HP LTO-2		Including Regular and WORM
			HP LTO-3		
			HP LTO-4		
	MSL G3 Series (2024/4048/8096)	Minimum 2024: 0370 Minimum 4048: 0600 Minimum 8096: 0850	HP LTO-2		
			HP LTO-3		
			HP LTO-4		
	EML E-Series	Minimum: 1070	HP LTO-3		LTO-4 WORM added to matrix at SN 3.5.1; it was validated earlier but missing from earlier release matrices
			HP LTO-4		
			HP LTO-4 WORM		

‡ Before using DLT cleaning with DLT-S4 or SDLT 600 drives, configure the library (Scalar i2000 or PX720) to disable reporting of the media ID. If media ID reporting is not disabled, StorNext will not recognize the cleaning media (SDLT type 1).

StorNext 4.0 Supported Libraries and Tape Drives (Continued)					
Vendor Library Family	**Libraries**	**Enforced Minimum / Recently Tested Library Firmware Level**	**Drive Types**	**Enforced Minimum / Recently Tested Drive Firmware Level**	**Notes**
IBM	TS3500	Minimum: 4680	IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
			IBM 3592 (J1A and E05)		
			IBM TS1120 (E05)		
Qualstar	XLS	Minimum: 0880	IBM LTO-3		
			IBM LTO-4		
Sony	Petasite CSM-200	Minimum: 6.30	IBM LTO-4 drive (T1600)		
Spectralogic	T-Series	Recently Tested: 2000	LTO-3	Vendor supported: 93G0	See Bulletin 46. Library firmware is known as BlueScale 11.
			LTO-4	Recently tested: 97F9	
Sun / StorageTek SCSI/FC Libraries	L180/ L700/ L1400	Minimum: 3.18.02	T9840C		
			T9840D		
			T10000A		See Note 2 below
			T10000B	See Note 4	See Notes 2 and 3
			HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL3000	Minimum: 2.30	T9840C		
			T9840D		
			T10000A		See Note 2
			T10000B	See Note 4	See Notes 2 and 3
			HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL500	Minimum: 1373	HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	9740	Minimum: 2000	Sun/STK 9840		Obsolete
			Sun/STK 9940		

Note 2: Tape cleaning problem introduced with 1.38.209, special fix at RA1.38.209. 1.40.x firmware level where fix was rolled forward has not been determined yet.

Note 3: When using T10000 drives, the STK library parameter "Fastload" must be set to "OFF".

Note 4: Minimum: RA1.38.209, but note that 1.38.209 manifests tape cleaning issue, recently tested: 1.40.208 (with tape cleaning failure)

StorNext 4.0 Supported Libraries and Tape Drives (Continued)					
Vendor Library Family	**Libraries**	**Enforced Minimum / Recently Tested Library Firmware Level**	**Drive Types**	**Enforced Minimum / Recently Tested Drive Firmware Level**	**Notes**
Sun / StorageTek ACSLS 7.3 Libraries See Notes 1 and 5	L180/ L700/ L1400	Minimum: 3.18.02	T9840C		
			T9840D		
			T10000A		See Note 2
			T10000B	See Note 4	See Notes 2 and 3
			HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL3000	Minimum: 0230 Recently Tested: Unavailable	T9840C		
			T9840D		
			T10000A		See Note 2
			T10000B	1.44 See Note 4	See Notes 2 and 3
			HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL500	Minimum: 1373 Recently Tested: Unavailable	HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL8500	Minimum: 4.14 Recently Tested: 4.70	T9840C		
			T9840D		
			T10000A		See Note 2
			T10000B	1.44 See Note 4	See Notes 2 and 3
			HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		

Note 1: The Sun / StorageTek FC and ACSLS sections have been modified to include drive and library permutations that are "paper certified" based on \testing that has been performed and validated by Sun/STK.

Note 2: Tape cleaning problem introduced with 1.38.209, special fix at RA1.38.209. 1.40.x firmware level where fix was rolled forward has not been determined yet.

Note 3: When using T10000 drives, the STK library parameter "Fastload" must be set to "OFF".

Note 4: Minimum: RA1.38.209, but note that 1.38.209 manifests tape cleaning issue, recently tested: 1.40.208 (with tape cleaning failure)

Note 5: ACSLS versions prior to ACSLS 7.3 have not been tested with this release.

Minimum Firmware Levels for StorNext Drives

Where applicable, the minimum firmware levels for StorNext-supported drives are shown in Table 5.

Table 5 Minimum Firmware Levels for Drives

StorNext 4.0 Minimum Firmware Levels for Drives		
Drive Type	**Minimum Drive Firmware Level**	**Notes**
IBM LTO-1	25D4	Also known as ULT3580-TD1 and ULTRIUM-TD1
IBM LTO-2	3AY4	Also known as ULT3580-TD2 and ULTRIUM-TD2
IBM LTO-3 IBM LTO-3 WORM	4C17	Also known as ULT3580-TD3 and ULTRIUM-TD3
IBM LTO-4	71G0	Also known as ULT3580-TD4 and ULTRIUM-TD4

Note: When using IBM ULTRIUM-TD3 drives with SUSE Linux Enterprise Server 10, you must upgrade the drive firmware to version 64D0 or later.

Supported StorNext Upgrade Paths

In general, sites running the following StorNext versions may upgrade directly to StorNext 4.0, assuming that the platform, service pack, architecture (32-bit or 64-bit), and StorNext component are supported in the installed StorNext version and in StorNext 4.0:

- StorNext 3.1.2
- StorNext 3.1.3
- StorNext 3.1.4
- StorNext 3.5
- StorNext 3.5.1

All other versions of StorNext require additional steps to upgrade to StorNext 4.0.

Note: HA customers on a StorNext release earlier than 3.5 must first upgrade to StorNext 3.5 before upgrading to StorNext 4.0.

Client Interoperability for StorNext 4.0

Table 5 indicates previous versions of StorNext SAN clients on certain platforms which can interoperate with a StorNext 4.0 metadata controller without upgrading the SAN Client.

Note: The table shows only client platforms for which support has been dropped in StorNext 4.0. For other platforms, it is expected that down-revved clients will be updated to StorNext 4.0.

Table 6 StorNext Client Interoperability

StorNext 4.0 Client Interoperability	
StorNext SAN Client Version	**Platform**
StorNext 3.1.2 StorNext 3.1.3 StorNext 3.1.4	Solaris 9 (sparc only) AIX 5.3 HPUX 11iv2 SGI IRIX 6.5.30 RHEL4 Itanium SLES10 Itanium SLES10 32-bit
StorNext 3.5 StorNext 3.5.1	AIX 5.3 HPUX 11iv2 SGI IRIX 6.5.30 (SN 3.5.1 only) SLES10 Itanium (SN 3.5.1 only) SLES10 32-bit (SN 3.5.1 only)

Configuration Requirements

Before installing StorNext 4.0, note the following configuration requirements:

- In cases where gigabit networking hardware is used and maximum StorNext performance is required, a separate, dedicated switched Ethernet LAN is recommended for the StorNext metadata network. If maximum StorNext performance is not required, shared gigabit networking is acceptable.
- A separate, dedicated switched Ethernet LAN is mandatory for the metadata network if 100 Mbit/s or slower networking hardware is used.
- StorNext does not support file system metadata on the same network as iSCSI, NFS, CIFS, or VLAN data when 100 Mbit/s or slower networking hardware is used.
- The operating system on the metadata controller must always be run in U.S. English.
- For Windows systems (server and client), the operating system must always be run in U.S. English.

Caution: If a Library used by StorNext Storage Manager is connected via a fibre switch, zone the switch to allow only the system(s) running SNSM to have access to the library. This is necessary to ensure that a "rogue" system does not communicate with the library and cause data loss or corruption. For more information, see StorNext Product Alert 16.

Time Synchronization for Replication and Deduplication

If you plan to use the Replication or Deduplication features, ensure that the time on your file system clients is synchronized to your metadata controllers.

The age values for Deduplication and Truncation are based on the clients' time, so if your clients' time is different from the MDC's time you may see files ingested earlier or later than you've configured.

Library Requirements

The following libraries require special configurations to run StorNext.

DAS and Scalar DLC Network-Attached Libraries

Prior to launching the StorNext Configuration Wizard, DAS, and Scalar DLC network-attached libraries must have the DAS client already installed on the appropriate host control computer.

DAS Attached Libraries

For DAS attached libraries, refer to "Installation and Configuration" and "DAS Configuration File Description" in the *DAS Installation and Administration Guide*. The client name is either the default StorNext server host name or the name selected by the administrator.

StorNext can support LTO-3 WORM media in DAS connected libraries, but WORM media cannot be mixed with other LTO media types in one logical library.

To use LTO-3 WORM media in a logical library, before configuring the library in StorNext, set the environmental variable `XDI_DAS_MAP_LTO_TO_LTOW` in the `/usr/adic/MSM/config/envvar.config` file to the name of the library. The library name must match the name given to the library when configuring it with StorNext. If defining multiple libraries with this environmental variable, separate them with a space. After setting the environmental variable, restart StorNext Storage Manager (SNSM).

Note: SDLC software may not correctly recognize LTO-3 WORM media in the library and instead set it to "unknown media type." In this case you must manually change the media type to "LTO3" using the SDLC GUI.

Scalar DLC Attached Libraries

For Scalar 10K and Scalar 1000 DLC attached libraries, refer to "Installation and Configuration" and "Client Component Installation" in the *Scalar Distributed Library Controller Reference Manual* (6-00658-02).

The DAS client should be installed during the installation of the Scalar DLC attached libraries. Use this procedure to install the DAS client.

1. Select **Clients > Create DAS Client**.

   The client name is either the default StorNext server host name or the name selected by the administrator.
2. When the DAS client is configured in Scalar DLC, select **Aliasing**.
3. Select **sony_ait** as the **Media** aliasing.

   The default value is 8mm.
4. Verify that **Element Type** has **AIT** drive selected.
5. Click **Change** to execute the changes.

ACSLS Attached Libraries

Due to limitations in the STK ACSLS interface, StorNext supports only single ACS configurations (ACS 0 only). StorNext support requires that the ACSLS client be installed on the appropriate host machine.

Disk Requirements

Disk devices must support, at minimum, the mandatory SCSI commands for block devices as defined by the SCSI Primary Commands-3 standard (SPC-3) and the SCSI Block Commands-2 (SBC-2) standard.

To ensure disk reliability, Quantum recommends that disk devices meet the requirements specified by Windows Hardware Quality Labs (WHQL) testing. However, there is no need to replace non-WHQL certified devices that have been used successfully with StorNext.

Disk devices must be configured with 512-byte or 4096-byte sectors, and the underlying operating system must support the device at the given sector size. StorNext customers that have arrays configured with 4096-byte sectors can use only Windows, Linux and IRIX clients. Customers with 512-byte arrays can use clients for any valid StorNext operating system (i.e., Windows, Linux, or UNIX).

In some cases, non-conforming disk devices can be identified by examining the output of `cvlabel -vvvl`. For example:

```
/dev/rdsk/c1d0p0: Cannot get the disk physical info.
```

If you receive this message, contact your disk vendors to determine whether the disk has the proper level of SCSI support.

Disk Naming Requirements

When naming disks, names should be unique across all SANs. If a client connects to more that one SAN, a conflict will arise if the client sees two disks with the same name.

SAN Disks on Windows Server 2008

SAN policy has been introduced in Windows Server 2008 to protect shared disks accessed by multiple servers. The first time the server sees the disk it will be offline, so StorNext is prevented from using or labeling the disk.

To bring the disks online, use the `POLICY=OnlineAll` setting. If this doesn't set the disks online after a reboot, you may need to go to Windows Disk Management and set each disk online.

Follow these steps to set all disks online:

1. From the command prompt, type **DISKPART**
2. Type **SAN** to view the current SAN policy of the disks.
3. To set all the disks online, type **SAN POLICY=onlineall**.
4. After being brought online once, the disks should stay online after rebooting.
5. If the disks appear as "Not Initialized" in Windows Disk Management after a reboot, this indicates the disks are ready for use.

   If the disks still appear as offline in Disk Management after rebooting, you must set each disk online by right-clicking the disk and selecting **Online**. This should always leave the SAN disks online after reboot.

Note: NOTE: If the disks are shared among servers, above steps may lead to data corruption. Users are encouraged to use the proper SAN policy to protect data

EXAMPLE:

```
C:\ >Diskpart
Microsoft DiskPart version 6.0.6001
Copyright (C) 1999-2007 Microsoft Corporation.
On computer: CALIFORNIA
DISKPART> SAN
SAN Policy  : Offline All
DISKPART> san policy=onlineall
DiskPart successfully changed the SAN policy for the current
operating system.
```

LDAP Support Requirement

LDAP (Lightweight Directory Access Protocol) support requires Windows Active Directory.

Configuring Quantum Libraries for Solaris 10

To ensure that Quantum libraries are recognized and tape drives function properly, follow the configuration procedure below. This procedure works for the PX502 library and other Quantum tape libraries.

Note: You must be using update 4 or newer for Solaris 10 in order for tape drives to function properly.

1. Edit the `/etc/driver_aliases` file by removing or commenting out the following ST driver entries (if they exist):
   - `"scsiclass,01"`
   - `"scsiclass,08"`
2. Open the `/kernel/drv/sgen.conf` file and verify that the following entries are present in the file, adding them if necessary:
   - `inquiry-config-list="ADIC","*";`
   - `inquiry-config-list="QUANTUM","*";`
   - `inquiry-config-list="HP","*";`
   - `device-type-config-list="changer","sequential";`
3. Reboot the Solaris system to unload any drivers that have been loaded.
4. After rebooting, enter the following commands to configure and load new sgen drivers:
   - `update_drv -a -i '"scsiclass,01"' sgen`
   - `update_drv -a -i '"scsiclass,08"' sgen`
5. Enter the command "`cfgadm -alv`". You should see the entries similar to this in the cfgadm list:

```
c2::500e09e00b40a000           connected    configured
unknown
QUANTUM PX500
unavailable  med-changer  n        /devices/pci@8,700000/
fibre-
channel@3/fp@0,0:fc::500e09e00b40a000
c2::500e09e00b40a010           connected    configured
unknown
HP Ultrium 3-SCSI
unavailable  tape         n        /devices/pci@8,700000/
fibre-
channel@3/fp@0,0:fc::500e09e00b40a010
```

Hardware Requirements

To successfully install StorNext 4.0, the following hardware requirements must be met:

- StorNext File System and Storage Manager Requirements on page 29
- StorNext Client Software Requirements on page 30

Note: The following requirements are for running StorNext only. Running additional software (including the StorNext client software) requires additional RAM and disk space.

StorNext File System and Storage Manager Requirements

The hardware requirements for StorNext File System and Storage Manager are presented in Table 7.

Table 7 File System and Storage Manager Hardware Requirements

No. of File Systems	RAM	File System Disk Space	Storage Manager Disk Space
1–4*	4 GB	2 GB	• For application binaries, log files, and documentation: up to 30GB (depending on system activity) • For support directories: 3 GB per million files stored
5–8**	8 GB	4 GB	

*Two or more CPU cores are recommended for best performance.
**Two or more CPU cores are required for best performance.

Additional Memory and Disk Requirements for Deduplication and Replication

In order to use the data deduplication and replication features in StorNext 4.0, your system must have the following memory and disk capacity **in addition to** the base memory and disk capacity required to run StorNext File System and Storage Manager.

Minimum Additional Disk and Memory Required for the Blockpool

- 50 MB available hard disk space

Minimum Additional Disk and Memory Required for Systems with 0 - 1 TB of Data

- 1 GB additional RAM
- 1 TB available hard disk space (or less, depending on your licensed amount)

Minimum Additional Disk and Memory Required for Systems with 1 - 10 TB of Data

- 6 GB additional RAM
- 10 TB available hard disk space

Minimum Additional Disk and Memory Required for Systems with 10 - 50 TB of Data

- 13 GB additional RAM
- 50 TB available hard disk space

Minimum Additional Disk and Memory Required for Systems with 50 - 150 TB of Data

- 28 GB additional RAM
- 150 TB available hard disk space

StorNext Client Software Requirements

To install and run the StorNext client software, the client system must meet the following minimum hardware requirements.

For SAN (FC-attached) clients or for Distributed LAN Clients:

- 1 GB RAM
- 500 MB available hard disk space

For SAN clients acting as a Distributed LAN Server:

- 2 GB RAM
- 500 MB available hard disk space

Note: Distributed LAN servers may require additional RAM depending on the number of file systems, Distributed LAN Clients, and NICs used. See Distributed LAN Server Memory Tuning in the StorNext User's Guide for Distributed LAN Server memory tuning guidelines.

Supported System Components

System components that are supported for use with StorNext 4.0 are presented in Table 8.

Table 8 StorNext Supported System Components

Component	Description
Tested Browsers	Internet Explorer 7.x or 8.x Mozilla Firefox 3.x Other browsers are not recommended.
NFS	Version 3 **NOTE:** An NFS server that exports a StorNext file system with the default export options may not flush data to disk immediately when an NFS client requests it. This could result in loss of data if the NFS server crashes after the client has written data, but before the data has reached the disk. This issue will be addressed in a future StorNext release. As a workaround, add the `no_wdelay` option to each line in the `/etc/exports` file that references a StorNext file system. For example, typical export options would be (`rw,sync,no_wdelay`).

Component	Description
LDAP	LDAP (Lightweight Directory Access Protocol) support requires Windows Active Directory.
Mixed-Level Tape Drive Compatibility Within the Same Device Family	LTO-1 media in a library containing LTO-3 or LTO-4 drives are considered for store requests unless they are logically marked as write protected. When LTO-1 media is mounted in an LTO-3 or LTO-4 drive, StorNext marks the media as write protected. Quantum recommends circumventing LTO-1 media for store requests by following this procedure: **1** From the SNSM home page, choose **Attributes** from the **Media** menu. **2** On the **Change Media Attributes** window, select the **LTO-1** media from the list. **3** Click the **Write Protect** option. **4** Click **Apply** to make the change. **5** Repeat the process for each piece of LTO-1 media. **NOTES:** • A similar issue exists for LTO-2 media in a library containing LTO-4 tape drives. • LTO-3 drives can read but not write LTO-1 tapes. • LTO-4 drives can read but not write LTO-2 tapes, and also cannot read LTO-1 tapes at all.

Previous Versions of Release Notes

Previous versions of the StorNext release notes contain additional information specific to earlier StorNext releases. You can find previous release notes at this location:

Release notes for previous StorNext releases are available here:

http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx#Documentation

Release notes and other documentation for previous StorNext releases which are no longer supported are available here:

http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/ArchivedManuals/Index.aspx

Known Issues

The following sections list known issues in this release of StorNext, as well as associated workarounds, where applicable:

StorNext File System Known Issues

Table 9 lists known issues that are specific to StorNext File System.

Table 9 StorNext File System Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Solaris	24563	n/a	Solaris hosts may need to rescan disk devices after StorNext labels have been applied. In particular, when a StorNext label is put on a LUN less than 1TB in size, Solaris hosts will not be able to use that LUN until they have done a device rescan. A device rescan is accomplished with a boot flag: reboot -- -r	This issue will be addressed in a future StorNext release. In the meantime, work around this issue by rescanning devices using the boot flag reboot -- -r If the labeling operation was performed on a Solaris host, that host does not need to do the rescan. However, some intermediate versions of the Solaris 10 Kernel Jumbo Patch break the necessary functionality to support this; please be sure you have applied the latest Solaris 10 Kernel Jumbo Patch before labeling any StorNext LUNs.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Solaris	24331	755956	Running an anonymous FTP server inside of a StorNext file system could cause a system to crash with the following error: `CVFS ASSERTION FAILED: f_cvp->cv_opencnt == 0`	This issue will be addressed in a future StorNext release. To work around this issue, install Very Secure FTP Daemon (vsftpd) and use it instead of the FTP daemon (`in.ftpd`) that is shipped with Solaris.
HP-UX	24309	n/a	If the cvpaths file contains an invalid wild-card specification (that is, a wild-card pattern that does not include a leading '/' character), the fsmpm process could panic and the cvlabel command might fail with a core dump.	This issue will be addressed in a future StorNext release.

Linux	23661	958244	StorNext File System does not support the Linux `sendfile()` system call.
			This issue causes Apache web servers to deliver blank pages when content resides on StorNext file systems.
			This issue also affects Samba servers running on Linux.
				The workaround is to disable sendfile usage by adding the following entry into the Apache configuration file `httpd.conf`:
				`EnableSendfile off`
				The workaround for Samba servers is to add the following line into the configuration file:
				`sendfile=no`

	25864	n/a	An NFS server that exports a StorNext file system with the default export options may not flush data to disk immediately when an NFS client requests it. This could result in loss of data if the NFS server crashes after the client has written data, but before the data has reached the disk.	This issue will be addressed in a future StorNext release. As a workaround, add the `no_wdelay` option to each line in the `/etc/exports` file that references a StorNext file system. For example, typical export options would be (`rw,sync,no_wdelay`).
	26321	n/a	Due to the way Linux handles errors, the appearance of SCSI “No Sense” messages in system logs can indicate possible data corruption on disk devices. This affects StorNext users on Red Hat 4, Red Hat 5, SuSe 9, and SuSe 10.	This issue is not caused by StorNext, and is described in detail in StorNext Product Alert 20. For additional information, see Red Hat 5 CR 468088 and SuSE 10 CR 10440734121.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	29416	n/a	In an HA configuration, RAS messages are not generated when the secondary system loses SAN connectivity.	This issue will be addressed in a future StorNext release. A workaround is to activate some of the unmanaged file systems on the secondary metadata controller. This will allow RAS messages from the secondary MDC if there is SAN connectivity loss. (This action would also help with load balancing.)
	28231	n/a	The `cvfsck -s` option should be extended to handle deduplicated files.	This issue will be addressed in a future StorNext release.
	28278	n/a	Replication needs an efficient way to clean up namespaces.	This issue will be addressed in a future StorNext release.
	28345	n/a	Most of the `snpolicy` commands require a 'pathname' argument which is sometimes used as an actual pathname, and is other times used only to determine the file system. snpolicy should be modified such that when the pathname is really a mount point, it should be called that in the help text, and should fail if the given path is not a StorNext file system mount point.	This issue will be addressed in a future StorNext release.
	28375	n/a	StorNext doesn't replicate non-regular files. (Currently, StorNext skips files that are UNIX domain sockets, block/char devices, and fifo's. Some of these file types could be included in namespace replication.)	This issue will be addressed in a future StorNext release.
	28415	n/a	StorNext should support mixed mode replication (dedup and non-dedup).	This issue will be addressed in a future StorNext release.
	28561	n/a	When a machine is rebooted, in rare situation, the network interfaces may have been changed. If you have configured Distributed LAN Server (DLS), the file system may fail to mount because the DLS may perceive that the associated network interface may have changed or disappeared.	This issue will be addressed in a future StorNext release. To work around this issue, you must find the appropriate network interfaces and reconfigure the `dpserver.fsname` file and then remount the file system.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	28587	n/a	Replication/Deduplication keeps its configuration data in a data stripe group in the file system. If this stripe group is down or lost, Replication/ Deduplication truncated files become inaccessible, as the policies that say where to retrieve the data will likewise be inaccessible.	This issue will be addressed in a future StorNext release.
	28944	n/a	Running the `cvmkfs` command doesn't check if the assigned disks for the file system are already being used by another file system.	This issue will be addressed in a future StorNext release.

	27174	n/a	Changing the `haFsType` parameter in a file system configuration file to one of the HA types, and then (re)starting its FSM enables HA-specific features that change the functionality of StorNext.
			When the `HaShared` or `HaManaged` types are configured, other changes must be completed by successfully running the `cnvt2ha.sh` script, which is indicated by the creation of the `/usr/adic/install/.snsm_ha_configured` touch file (`$SNSM_HA_CONFIGURED` environment variable). No conversion is done or necessary for SNFS only (`HaUnmanaged`) configurations.
			If the conversion is not successfully completed, the HaManaged FSMs will not start, and the HaShared FSM will cause an HA Reset when it is stopped.
				This issue will be addressed in a future StorNext release.
				To remedy this situation, edit every FSM configuration file to set its `haFsType` parameter to `HaUnmonitored`, then run the following commands to avoid the HA Reset in this special case only:
				`touch /usr/cvfs/install/.vip_down_hint`
				`service cvfs stop`

	29678	n/a	Replication may hang if the `dedup_bfst` parameter ("Address for Replication and Deduplication") on the source has been configured to use an address that is not reachable by the target.	This issue will be addressed in a future StorNext release. The workaround is to manually confirm reachability to the replication source's vIP address on the replication target, and then reconfigure routing, if necessary.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	30885	n/a	After entering the correct StorNext GUI login password on the primary MDC server in an HA configuration, the login failed with a message, "Login Incorrect, please try again." This error occurred because the StorNext GUI database failed to start due to database synchronization issues in the HA environment.	To recover from this condition, follow these steps: First, attempt to restart StorNext on the primary MDC server by doing the following: **1** Open a unix root shell window on the primary MDC server. **2** Restart the GUI application by running `service stornext_web restart` If you are still unable to log in, the database files may have been improperly synchronized in the HA environment. If so, you have two options for restoring the database: **1** If you are running Storage Manager, you can manually restore the GUI database from a backup. **2** Recreate the database from scratch. (You must do this if the backup database is corrupted.) If you previously changed the login password, you must redo this step. Also, if you configured replication targets, you must re-scan and add back the desired targets.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	30885 (continued)	n/a	(See previous page.)	**Option 1**. To manually restore the GUI database from backup: **1** Open a UNIX root shell on the primary MDC server in the HA environment. **2** Stop the GUI application by running `service stornext_web stop` **3** Remove the GUI database files by running `rm -rf /usr/adic/gui/database/guiPrivateDb` **4** Get mount point for backups by running `grep BACKUPFS /usr/adic/TSM/config/fs_sysparm` **5** Go to the backup directory by running `cd <mount point>/.ADIC_INTERNAL_BACKUP` **6** Find the most recent configuration file backup by running `ls –lrt conf* | tail -1` **7** The file may have been truncated, so retrieve it by running `fsretrieve conf.<n>.<m>.tgz` **8** Extract the files by running `gzip -d -c conf.<n>.<m>.tgz | (cd /;tar xf - /usr/adic/gui/database)` **9** Truncate the file again by running `fsrmcopy conf.<n>.<m>.tgz` **10** Start the StorNext GUI application by running `service stornext_web start`

Operating System	CR Number	SR Number	Description	Workaround (if applicable)

Linux	30885 (continued)	n/a	(See previous page.)	**Option 2**. If the backup database backup is also corrupted, follow these steps to reinstall the GUI database:
				1 Open a UNIX root shell on the primary MDC server in the HA environment.
				2 Stop the StorNext GUI application by running `service stornext_web stop`
				3 Remove the GUI database files by running `rm -rf /usr/adic/gui/database/guiPrivateDb`
				4 Source the profile to setup the application environment by running `. /usr/adic/.profile`
				5 Recreate the database by running:
				`/usr/adic/java/bin/java`
				`-cp`
				`/usr/adic/tomcat/common/lib/derby.jar:/usr/adic/tomcat/common/lib/derbytools.jar`
				`-Dij.database="jdbc:derby:/usr/adic/gui/database/guiPrivateDb;create=true"`
				`-Dderby.system.home=/usr/adic/install`
				`org.apache.derby.tools.ij </usr/adic/gui/sql/gui_private_create.sql`
				6 Start the StorNext GUI by running `service stornext_web start`
				7 Advance through the Configuration Wizard steps by clicking Next.
				8 If applicable, change the password on the Tools > User Access screen.
				9 If applicable, rescan for replication targets on the Setup > Storage Destinations > Replication Targets screen.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	25707	n/a	Running the command `df` showed mapped drives but not the mapped folders and directories for cvfs and ntfs file systems.	The workaround is to use the command `mountvol`, which shows the directories and folders in the file system that are mount points.
	26252	n/a	Backup and other applications could fail to function properly when a file system is mounted on a directory using StorNext.	This issue will be addressed in a future StorNext release. In the meantime, do not mount file systems on a directory that uses StorNext.
	26142	n/a	The Windows Client Configuration tool's Global Options dialog does not have a "Notice" log level setting. If you select "Informational" you receive both "Notice" and "Informational" messages.	This issue will be addressed in a future StorNext release.
	29410	n/a	Opening the Windows Client Configuration tool can result in the following error message: `"Error loading mount information from registry-Error (1450/0x5AA - Insufficient system resources exist to complete the requested service".`	This issue will be addressed in a future StorNext release.
	29483	n/a	After changing `fsnameservers` in the StorNext GUI, the file system failed to mount and returned a "`device not connected`" error.	This issue will be addressed in a future StorNext release. The workaround is to stop and start the StorNext services manually from the command line.
	29486	n/a	The number of file systems found after clicking the Scan button does not match the number of mountable file systems on an HA system.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	15032	n/a	If you are using Promise RAID controllers on a Windows Server 2008 64-bit system, you must install Promise's PerfectPath software. If you do not install this software, you will be unable to use your Windows Server 2008 system.	Promise is working on a solution to this problem, but in the meantime they have provided the following workaround: 1. Install the PerfectPath software on your Windows Server 2008 64-bit system. 2. Restart your system. The login prompt will *not* appear after you restart. Instead, the **Windows Boot Manager** screen appears showing an error message: "Windows cannot verify the digital signature for this file" (`\Windows\system32\DRIVERS\perfectpathdsm.sys`) 3. From the **Windows Boot Manager** screen, press **Enter** to continue. A second **Windows Boot Manager** screen appears, asking you to choose an operating system or specify an advanced option. 4. On the second **Windows Boot Manager** screen, press **F8** to specify advanced options. The **Advanced Boot Options** screen appears. 5. On the **Advanced Boot Options** screen, use the arrow keys to choose the option **Disable Driver Signature Enforcement**. Choosing this option will cause the system to display the login prompt normally after you reboot. 6. Restart your system.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	30945	n/a	On Windows Systems, creating a Virtual Hard Disk (VHD) on a StorNext file system will cause a system crash.	Until there is a resolution to this problem, all VHDs must be created on NTFS file systems. This includes both direct creation of VHDs through the Windows Disk Manager on Windows 7 and Windows 2008 R2 systems, as well as any Windows-based product that uses VHDs implicitly. This includes, but is not limited to, the Windows Complete PC Backup and Microsoft Virtual PC products. All of these products should continue to function normally with StorNext installed as long as they are not configured to create their VHDs on StorNext.
All	25836	898484	Failover on stripe groups is not currently supported.	This enhancement request may be considered in a future StorNext release.
	26114	n/a	Running `cvfsck` can crash when a data file is empty.	This issue will be addressed in a future StorNext release.
	27483	983534	The management utility `cvadmin` allows downing the metadata and journal stripe groups while the file system is active. This allows files to be created on a client until the FSM needs to access the metadata, at which point an `ENOSPACE` error is generated.	This issue will be addressed in a future StorNext release to prevent `cvadmin` from allowing stripe groups that have metadata and journaling associated with them to be downed while the file system is active.
	15383	n/a	Running the utility `dev_rescan.sh` (or modifying FC zone) caused I/O disruption on the metadata controller during heavy client load.	This issue will be addressed in a future StorNext release.
	28116	n/a	PunchHole interactions with Replication/Deduplication files must be taken into consideration.	This issue will be addressed in a future StorNext release.
	29023	n/a	Replication quiesce scripts do not synchronize data on any clients that have open files.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30049	n/a	The StorNext GUI allows you to enter a replication target that does not exist, which results in a replication job starting which never completes.	This issue will be addressed in a future StorNext release. (Additional validation may be added when adding replication targets.)
	29557	n/a	The StorNext GUI can hang if it is unable to create a file system. Additional validation is necessary.	This issue will be addressed in a future StorNext release.
	28665	n/a	Spurious errors may appear in the StorNext GUI display on the primary machine during HA conversion. For example, you might receive an error saying that the file system is not started, or that the Storage Manager is offline because required components are not started.	This issue will be addressed in a future StorNext release.
	26300	909788	The `cvfsck` command runs extremely slowly with big directories containing many orphan entries.	Use the following workaround until this issue is resolved: **10** Mount the file system. **11** Create a temporary directory that has the same owner, group and permissions as the big directory. **12** Move as many files and subdirectories out of the big directory into the temporary directory, leaving behind only the orphan directory entries. **13** Unmount the file system. **14** Run the `cvfsck` command. **15** Mount the file system. **16** Remove the original directory (which should be empty after running the `cvfsck` command). **17** Rename the temporary directory to match the original directory's name.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	29753	n/a	When an FSM is selected in `cvadmin` and that FSM subsequently quits, a `cvadmin` command that needs to communicate with the FSM fails with the message, "`Connection reset by peer.`" If the selection of the FSM is still unchanged, the next and subsequent commands report "`broken pipe.`"	This issue will be addressed in a future StorNext release. To work around this issue, re-select an FSM.
	29762	n/a	When running under VMware on SuSE Linux, a system shutdown may stop the network before shutting down StorNext. This is due to a bug in the VMware tools which do not interoperate correctly with the SuSE shutdown system if `RUN_PARALLEL` is enabled.	To address this issue, update to the latest VMware tools.
	28474	n/a	The event file directory is now a configurable value (`eventFileDir`) within the file system configuration file, and could cause problems if the default value is changed. For HA environments the event file directory must be in a location that is visible by both metadata controllers. Additionally, it cannot reside in a file system configured for Storage Manager or deduplication/ replication. The default value (`/usr/adic/TSM/internal/event_dir`) is relocated to the shared HA file system as part of the HA setup. If an event directory is not accessible, TSM will not start. If the value is changed, any unprocessed event files in the old event directory will not get processed unless they are moved to the new event directory. In order to do this, all file systems managed with deduplication/replication and Storage Manager should first be unmounted. The setting can then be changed, and the unprocessed event files can be moved to the new location.	This issue will be addressed in a future StorNext release. In the meantime, avoid changing the default value for the event file directory. (There is no need for a workaround if the value is not modified or set to a value adhering to the constraints.)

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30623	n/a	When an SNFS-exporting NFS server is rebooted, applications running on NFS clients may experience “stale NFS handle” errors if changes are made on the NFS server regarding the number of StorNext file systems mounted, or if the cvfs file system mount order changes on the NFS server.	This issue will be addressed in a future StorNext release. In the meantime, if you are upgrading to StorNext 4.0 you should first unmount file systems from NFS clients that use StorNext servers. Otherwise, the applications running on those NFS clients will experiences “stale NFS handle” errors.
	30696	n/a	Attempting to “loopback” mount an ISO image residing in a StorNext file system resulted in a an error message and mount failure.	This issue will be addressed in a future StorNext release. One workaround is to copy the ISO image to a local file system and then perform the loopback mount on the copy.
	30674	n/a	In rare cases, the StorNext GUI cvadmin can have conflicting views of disk states. This can occur if you have already labeled disks from some other node, make them available to a system after installation or startup, and then try to create a file system using those disks.	This issue will be addressed in a future StorNext release. The following workarounds are available, listed here in order of invasiveness to the system: • call `cvadmin -e 'disks refresh"` from the CLI • Stop/start SNFS from the StorNext GUI • Reboot the system
	30765	n/a	In an HA system, RAS events that appear on the Primary machine before failover are duplicated and appear on the Secondary machine after failover.	This issue will be addressed in a future StorNext release.
	30817	n/a	When replicating with multiple copies, StorNext can propagate metadata changes into the previous replication copies.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30812	n/a	When a replication source system has a hard reset and reboots, it may require the TCP Keepalive time to expire before replication resumes. On most Linux systems, this time duration defaults to a little over two hours.	This issue will be addressed in a future StorNext release. The only workaround for this problem is to reset the global TCP Keepalive parameters for your system to lower values which may affect other processes running on the system. You may want to experiment with lower values to see how your system behaves, as higher values of keepalive may be used to protect applications from unstable network links.
	30874	n/a	The replication bandwidth throttling feature may not perform optimally at low rate limits.	This issue will be addressed in a future StorNext release. To work around this issue, Quantum recommends setting the rate to at least 40Mb/s.
	30982	n/a	StorNext replication will skip files that have been modified shortly before replication begins. These files may not be replicated in subsequent replications unless the source has changed.	This issue will be addressed in a future StorNext release. You have two options for avoiding this issue: 1) Ensure that files are not modified in the few minutes prior to their scheduled replication. 2) Monitor the replication reports and manually touch a file on the source if any files are skipped, ensuring that the next replication runs.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30738	n/a	During an upgrade and other cases where StorNext's file system component is stopped, the blockpool or snpolicyd may take longer than expected to exit. If this happens, the stop will fail, and you will need to determine what caused the blockpool or snpolicyd to stop, and then take appropriate actions to stop them.	This issue will be addressed in a future StorNext release. Follow these steps to correct the problem: **1** Verify that the fsm for the blockpool's file system is running. If it isn't, start it and then try the operation again. **2** If this doesn't fix the problem, verify that the disks for the blockpool's file systems are available and operating normally. If they are not, correct the problem and then try the operation again. **3** If neither of these cases apply, it may be that the blockpool or snpolicyd is simply busy, and waiting for it to slow down will cause the stop to succeed. Retry the operation after waiting a few minutes. If none of these options works, contact Quantum Support for assistance.

StorNext Storage Manager Known Issues

Table 10 lists known issues that are specific to StorNext Storage Manager.

Table 10 StorNext Storage Manager Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	26279	945058	Running an SQL query did not return the desired results.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	25978	n/a	Scheduled tasks for “partial backups” and for “rebuild policy” can fail if they overlap.	This issue will be addressed in a future StorNext release. The default scheduler value for a partial backup is two hours. If you have a large managed file system you might need to adjust schedules to permit longer times if your partial backups require more than two hours to complete. Changing the allotted time will ensure that the partial backup completes before the rebuild policy task starts.
	30858	n/a	The `/etc/init.d/cvfs` script can report that MSM or TSM are not ready, and return that a failure occurred if MSM or TSM have not come online after 100 seconds. However, MSM and TSM will still continue to start up in the background, and should eventually come online.	This issue will be addressed in a future StorNext release. The workaround is to manually run "`MSM_control status`" and "`TSM_control status`" to poll StorNext to see if it is completely online. Verify that there are no errors and that no corrective action is required if MSM and TSM continue to fail to start up.
	25151	n/a	StorNext Storage Manager components `DSM_control`, `TSM_control` and `MSM_control` start operations sometimes return the error code `16 (EBUSY)`.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)

Linux	24649	n/a	StorNext Storage Manager in a High Availability configuration could encounter a problem reserving tape drives following a failover. A “target reset” is used by the newly activated metadata controller to release “scsi reserve” device reservations made by the former metadata controller. The target reset operation might fail due to device driver problems on systems running SUSE Linux Enterprise Server 10 with tape drives attached via an LSI fibre host bus adapter. Any such reserved drives will not be accessible by the new metadata controller.	There are two possible workarounds, which also apply to versions of StorNext prior to 3.1.2.
				1) Following a failover, release any tape drive reservations held by the former metadata controller. This must be done for each tape drive still reserved by the former metadata controller by running `/usr/adic/TSM/util/fs_scsi` on the metadata controller which owns the reservation:
				`# /usr/adic/TSM/util/fs_scsi`
				`Choice==> 10 (list drives)`
				`Choice==> 1 (select drive, e.g. /dev/sg0)`
				`Choice==> 3 (select Test Ready)`
				`Choice==> 24 (select Release)`
				`Choice==> 0 (select Quit to exit fs_scsi)`
				OR
				2) Add the following setting to the `/usr/adic/TSM/config/fs_sysparm_override` file: `FS_SCSI_RESERVE=none;`. (“none” means don't try to reserve tape drives.) StorNext must be restarted for this change to take effect.
				WARNING: This workaround could leave tape drives exposed to unexpected access by other systems, which could lead to data loss.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)

Linux	28447	n/a	In the event of an environmental issue (such as power interruption) in which both the MDC and Distributed Data Mover mover machines need to restart, there can be a timing situation in which the MDC comes up and tries to initiate TSM before the client has completed its own reboot.
			TSM will use the `fs_fmoverc` process for up to five minutes to clean up the status of the DDM mover machines. If the DDM mover machines remain down for an extended period, there can be retries of the five-minute `fs_fmoverc`.
				This issue will be addressed in a future StorNext release.
				If TSM has not completed its startup and `fs_fmoverc` processes persist (as shown by running the `tsmup` command), use the `fsddmconfig` command to set the state of the DDM mover machine to disabled.
				For example, if the hostname of the DDM mover machine is `minnesota`, run this command on the MDC:
				`fsddmconfig -u -s d minnesota`
				Once the DDM mover machine is back up, re-enable it using the StorNext GUI or with this command:
				`fsddmconfig -u -s e minnesota`
				Under most circumstances, `fs_fmoverc` will successfully determine the downed status of a DDM mover machine, and automatically set its status to `DISABLED`. A RAS ticket is emailed when `fs_fmoverc` automatically disables a mover, so after receiving such a ticket and rebooting the DDM mover, use the StorNext GUI (or the command `fsddmconfig` in the second example above) to enable the mover.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	29445	n/a	When adding Distributed Data Mover (DDM) mover hosts (either by the StorNext GUI or by the `fsddmconfig` command,) StorNext does not check whether the same host has already been added under an equivalent identity - IP address versus non-qualified hostname versus fully qualified hostname. If you define the same host more than once and are setting maximum mover process counts, each definition of the host gets its own count against the maximum mover limit and thus the host can run more `fs_fmover` processes than intended. This can impact performance tuning.	This issue will be addressed in a future StorNext release. The workaround is to avoid re-defining the same host under multiple equivalent identities.
All	12321	n/a	Removing affinities does not unconfigure them from managed file systems.	This issue will be addressed in a future StorNext release.
	25506	n/a	Early Warning End-Of-Medium check conditions cause undesirable side effects.	This issue will be addressed in a future StorNext release. A workaround is to not use tapes beyond the point at which drives begin to report the check conditions. To help you stay within the limit, add the following to `/usr/adic/TSM/config/fs_sysparm_override`: `PERCENT_FULL_TO_MIGRATE=95` Making this change prevents undesirable side effects, but the trade off is that the last 5% of a tape is not used rather than the last 0.1%.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	25837	836242	Running `fsmedcopy` fails if the output medium does not have sufficient space to hold the copies from the input. However, the copies successfully transferred to the new media can be identified with the `fsfileinfo` command.	This issue will be addressed in a future StorNext release. As a workaround, a few `fsmedcopy` runs will get all the files from the larger input media to the new (smaller) media.

	29273	n/a	If a deduplication candidate is removed before blockpool processing is completed, errors such as the following may be sent to the syslog:
			`Oct 2 15:22:00 orleans Blockpool[16403]: E: [5] (Store Local) Error storing file "/stornext/source/__CVFS_Handle.000474F892EBB65E000E000000000000000000000292BF4".`
			`Error opening file "/stornext/source/__CVFS_Handle.000474F892EBB65E000E000000000000000000000292BF4".`
			`No such file or directory.`
			Errors such as these may appear serious, but there is no reason for concern.
				This issue will be addressed in a future StorNext release.
				If you receive these errors, no action is required.

	30006	n/a	There is no way to conveniently delete a TSM relation point used for replication.	This issue will be addressed in a future StorNext release. In the meantime, you can manually delete the relation point by running the command `rm -rf /snfs/sn2/tsm/.rep_private`, which empties the TSM relation point. When running this command, be aware that there may have been several targets being realized with the TSM relation point in question, so you should remove the directory `tsm_dir /.rep_private` only after the LAST target policy has been removed from the relation point.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	28336	n/a	When using the Distributed Data Mover feature, media and drives were excessively marked as failed.	This issue will be addressed in a future StorNext release. Use the following workaround until this issue is resolved: Re-mark media as available, and drives as online.
	30867	n/a	When the `cvfs init` script runs on an HA system, the following error message appears on the console after booting up: `Starting cvfs: F1001(1)<2004944918>:The /usr/adic/TSM/config/ fs_sysparm system parameter file could not be opened: No such file or directory.`	This issue will be addressed in a future StorNext release.

StorNext GUI Known Issues

Table 11 lists known issues that are specific to the StorNext GUI process.

Table 11 StorNext GUI Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	27705	n/a	Lagging GUI behavior may be observed with Internet Explorer 6.	This issue will be addressed in a future StorNext release.
	29114	n/a	The Apply button is always active on the Setup > Storage Pools > Replication Targets screen even when there is nothing to save or apply.	This issue will be addressed in a future StorNext release.
	27803	n/a	The StorNext GUI does not correctly display IP address changes for the HA Convert screen until next logout.	This issue will be addressed in a future StorNext release. In the meantime, you can work around this issue by logging off of StorNext and then logging back in.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	29827	n/a	When updating the "global" replication/deduplication storage policy, you cannot select a directory to which to assign the policy, yet there is a warning about there being no assigned directories.	This issue will be addressed in a future StorNext release.
	29929	n/a	Due to an error in a third-party component, double-clicking the New button on the File System > New screen returns exceptions.	This issue should be addressed in a future release of the third-party component. Other solutions may be considered for a future StorNext release.
	30573	n/a	When attempting to change an unmanaged file system to a managed file system, the procedure fails and times out with the message, "`Failed to modify file system.`"	This issue will be addressed in a future StorNext release. The workaround is to repeat the procedure a second time, which should result in success.
	30783	n/a	In some situations, when using the Configuration Wizard the "Next" button does not appear on the License screen.	This issue will be addressed in a future StorNext release. The workaround is to log off StorNext and then log back on.

	30989	n/a	In some situations, the StorNext Home Page Refresh icon spins for a long time, and the truncation candidates count is not updated during a Refresh operation.
			On systems with a large number of truncation candidates, the calculation process times out after 60 seconds, and the StorNext Home Page is not updated with the correct information.
				Use the following workaround to determine the truncation candidates:
				1 Open a root UNIX shell on the MDC server.
				2 Source the profile to set up the application environment by ruinning this command:
				`. /usr/adic/.profile`
				3 Run the following command to determine the truncation candidates. This can be a resource intensive command to run, and can take a very long time to complete.
				`/usr/adic/TSM/util/showc -t`

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	30842	n/a	In some situations, on HA systems the StorNext GUI exits the config mode (with the redundant server in peerdown mode) without starting Storage Manager. When an HA Cluster is operating without a redundant server for an extended time period, it is best practice to place the cluster into the single/peerdown state to prevent an HA Reset. The StorNext GUI allows transitions from this state into the config/peerdown state for making configuration changes. However, when you use the "Exit Config Mode" button on the StorNext GUI's Tools > HA > Manage screen to transition back from Config to Single mode, the GUI uses a diagnostic method for restarting the system that results in starting the file system without starting the Storage Manager.	This issue will be addressed in a future StorNext release. Two actions are necessary to correct this: **1** Run the following CLI command: "`DSM_control start`". This will produce the following warning message, which can be ignored: "`fsmpm is already running. Only one fsmpm can be running at a time. Aborting start.`" **2** In the StorNext GUI, go to the Tools > System Control screen and start the Storage Manager.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)

Linux	30959	n/a	If a StorNext file system is configured with short-labeled disks in the file system configuration but not the label on-disk, the StorNext GUI adjusts the short-labeled diskType in the file system configuration when editing any value in it.
			This prevents the file system from starting (and also causes an FSM core,) and produces the following RAS message:
			`SR Notes: hostname fsm[PID=21636]: fs "fsname": PANIC: /usr/cvfs/bin/fsm "  The stripe group "StripeGroup1" size is 28441664, expected 28466296. You must maintain the original configuration or re-initialize the file system. File system "fsname" not started. " file alloc.c, line 3183`
			`Ticket creation time: 03/01 10:32:53 CST`
				This issue will be addressed in a future StorNext release.
				Follow these steps to resolve this issue:
				1 Restore the configuration file from the `/usr/cvfs/data/<fsname>/config_history/` directory
				2 Use cvlabel to adjust the disk label(s) sector count to match what is in the file system configuration.
				3 Start the file system from the StorNext GUI.
				4 Verify file system starts up properly.
				5 Edit the file system configuration file again, and redo the changes that were made previously.

All	30146	n/a	None of the StorNext reports currently includes complete media attribute information, including all files on the media.	This issue will be addressed in a future StorNext release.
	28401	n/a	By default, the StorNext GUI chooses the first free port starting at port 81. It would be convenient to have the option of specifying a particular port on which to run StorNext.	This issue will be addressed in a future StorNext release.
	28760	n/a	When using the Storage Pools > Replication / Deduplication feature, StorNext allows changing the Blockpool File System after the block pool has been created.	This issue will be addressed in a future StorNext release.
	28792	n/a	StorNext GUI Menu Navigation sometimes requires multiple mouse clicks.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	29038	n/a	Clicking through StorNext GUI logs pages more than once every few seconds causes error.	This issue will be addressed in a future StorNext release.
	29108	n/a	Setting up blockpool on a shared file system before conversion to HA makes the blockpool fail to start.	This issue will be addressed in a future StorNext release. To avoid this issue, do not install the blockpool on a shared file system.
	30091	n/a	The capacity indicators on the StorNext home page provide *approximations* and may not accurately summarize the actual current capacity.	This issue will be addressed in a future StorNext release. In the meantime, if you require accurate, up-to-the-minute capacity information, click the Capacity areas of the home page to view current capacity.
	28617	n/a	The Storage Policy configuration screen does not display information if a new replication/ deduplication policy is created from the command line.	This issue will be addressed in a future StorNext release. The workaround is to avoid creating replication/ deduplication storage policies from the command line.
	29413	n/a	Warnings about insecure and secure items are generated when using StorNext with Internet Explorer 8.	This issue will be addressed in a future StorNext release.
	29728	n/a	Due to an error in a third-party component, background pages respond to keyboard input when modal dialog windows are open.	This issue should be addressed in a future release of the third-party component. Other solutions may be considered for a future StorNext release.
	29836	n/a	When converting an HA secondary, the StorNext GUI displays an error explains that the secondary name validation failed. Another related error says the Storage Manager is not running.	This issue will be addressed in a future StorNext release.
	30227	n/a	Some errors that occur during bandwidth expansion are not caught by the StorNext GUI, which results in a file system that will not mount.	This issue will be addressed in a future StorNext release. To work around this, run the command `cvupdatefs <fsname>`. After the command completes successfully, start the FSM and mount the file system.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30511	n/a	The StorNext GUI does not support setting an affinity to a directory.	This issue will be addressed in a future StorNext release.
	30544	n/a	After upgrading to StorNext 4.0, StorNext did not launch correctly in the Internet browser. This failure could be caused by settings from an earlier StorNext release remaining in the browser's cache.	This issue will be addressed in a future StorNext release. The workaround is to clear your Internet browser cache and refresh the screen by pressing Control R.
	30884	n/a	If you specify a trailing '/' for the mount point option when creating a file system on the Setup > File System > New page, the StorNext GUI will incorrectly show the file system as not mounted.	This issue will be addressed in a future StorNext release. To correct the StorNext GUI display, manually edit the `/etc/fstab` file and remove the trailing '/' from the mount point entry. To update the GUI display, click the Refresh button on the Setup > File System page . Best practice is to omit the trailing '/' when specifying the file system mount point.

	30929	n/a	The StorNext GUI may be inaccessible in a Web browser, with one of the following error messages displayed:
			`Firefox: Unable to connect. Firefox can't establish a connection to the server`
			`Internet Explorer: Internet Explorer cannot display the webpage`
				If you encounter this condition, restart the StorNext GUI on the MDC server by doing the following:
				1 Open a root UNIX shell window on the MDC.
				2 Run the command `service stornext_web restart`

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30925	n/a	In systems with archives that have multiple mailboxes available, importing media can fail with the message "No new media found." (This occurs after choosing Storage Destinations > Library > Add Media Mailbox from StorNext's Setup menu.)	To fix this problem, try putting the media in one of the other mailboxes and then re-run the import. If the operation still fails, you can run the import manually by performing these steps: **1** Open up a UNIX root shell on the MDC server. **2** Source the profile by running `. /usr/adic/.profile` **3** Obtain a list of available mailboxes for an archive by running `/usr/adic/MSM/bin/mmportinfo <archivename>` **4** Import media into an archive from a specific mailbox by running `/usr/adic/gui/scripts/library.pl add_media --archive=<archivename> --importmethod=mailbox --mailbox=<mailbox>` Example: `/usr/adic/gui/scripts/library.pl add_media --archive=archive01 --importmethod=mailbox --mailbox=16:LTO:0,0,15,16`

StorNext Installation Known Issues

Table 12 lists known issues that are specific to the StorNext installation process.

Table 12 StorNext Installation Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)

Linux (RHEL4 and RHEL5 only)	24692	n/a	When you mount a CD in a Red Hat 4 or 5 system, CDs are mounted by default with a noexec (non-executable) option which prevents you from proceeding with the StorNext installation.	Remount the CD by typing `mount -o remount, exec ...`
				Alternatively, mount the CD to a different directory by typing the following:
				`# mkdir /mnt/MOUNT_PATH`
				`# mount /dev/cdrom /mnt/MOUNT_PATH # cd /mnt/MOUNT_PATH`

Windows	27129	n/a	There is an incompatibility between the StorNext 3.1.2, 3.1.3, and 3.5 installers and a third-party application called Altiris Agent. Running Altiris Agent as a service can cause the installation to fail and then automatically reboot the machine.	This issue will be addressed in a future StorNext release. The workaround is to disable the Altiris Agent service before installing StorNext. Alteris Agent is typically located at `C:\Program Files\Altiris\Altiris Agent\AeXNSAgent.exe.`

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	25866	n/a	StorNext upgrades on Vista machines can fail in the middle of installation. This problem is caused by the way Windows Vista handles software upgrades. A related error is described in Microsoft article 263253.	Microsoft has a utility called the Windows Installer Cleanup Utility that removes files left behind by incomplete installations. Access the Microsoft website and search for article ID 290301. To work around this issue, follow these steps: 1. Click Start, and then click Run. 2. In the Open box, type Regedit and then click OK. 3. On the Edit menu, click Find. 4. In the Find what box, type Snfs_XXX.dat and then click Find Next. 5. If the search result selects a string value called PackageName, continue with these steps. Otherwise, repeat steps 3-4. 6. Double-click the PackageName string value. 7. In the Value data box, change the installation directory path to the new pathname. For example if the old installation directory path contained OCT10, change that to the current path (e.g, NOV12.) 8. On the Registry menu, click Exit.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)

All	30837	n/a	In some situations, running `install.stornext -upgrade` can lead to an abort during the upgrade, leaving the system without the necessary snfs rpms installed.	This issue will be addressed in a future StorNext release.
				If the snfs, snfs-client and snfs-server rpms are not present following an `install.stornext -upgrade`, use the following workaround to complete the upgrade:
				1 Power cycle the system.
				2 Use the `lsmod` command to verify that the cvfs rpm is not loaded.
				3 Unload the cvfs module with this command: `rmmod cvfs`
				4 Manually install the rpms from the install media using this command:
				`rpm -force -hiv /PATH/TO/MEDIA/phdist/SYSTYPE/DSM/*.rpm`
				5 Rerun `install.stornext` with the `--force` option:
				`install.stornext --force -upgrade`

StorNext HA Known Issues

Table 13 lists known issues that are specific to StorNext HA systems.

Table 13 StorNext HA Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround
All	29067	n/a	When an MDC in an HA cluster starts up while LUNs are not available, the FSMPM process will try to access the LUNs, but eventually stop trying. The MDC will not provide file system services without intervention to restart StorNext. Although the snhamgr status command reports the MDC as running, it is not performing as a redundant server and will not take control in the event of an HA Reset.	After repairing access to the LUNs, stopping and restarting CVFS by running the following command may correct the problem: `service cvfs restart`
	29099	n/a	When exiting HA Config mode, StorNext will be stopped, which will also 'fuser' any processes which have files open on the file system from either node.	Prepare systems and users for this eventuality before entering HA Config mode.
	29184	n/a	Virtual IPs (VIPs) are left behind after uninstalling StorNext.	This issue will be addressed in a future StorNext release.
	29722	n/a	StorNext API 2.0.1 failed to install on the secondary node of an HA pair.	This issue will be addressed in a future StorNext release. A workaround is to fail the primary over to the secondary and then install SNAPI.

Operating Guidelines and Limitations

This section contains operating guidelines and limitations for running StorNext. Items are grouped according to operating system.

Table 14 lists legacy items from previous releases that still pertain to StorNext 4.0.

Table 14 StorNext Legacy Operating Guidelines and Limitations

Operating System / Affected Component	Description
Windows	In StorNext releases prior to 3.5, the StorNext Windows client attempted to keep the UNIX uid, gid and mode bits synchronized with similar fields in the Windows security descriptor. However, these Windows and UNIX fields were often not synchronized correctly due to mapping and other problems. One consequence of this problem was that changing the owner in Windows incorrectly changed the UNIX uid and file permissions and propagated these errors into sub-directories. In release 3.5, the StorNext Windows client sets the UNIX uid, gid and mode bits only when Windows creates a file. The StorNext Windows client will no longer change the Unix uid, gid or mode bits when a Windows user changes the Windows security descriptor or Read-Only file attribute. If you change the UNIX mode bits and the file is accessible from Windows, you must change the Windows security descriptor (if Windows Security is configured On) or Read-Only file attribute to ensure the change is reflected on both Windows and UNIX.
	As of StorNext release 3.5 the Authentication tab has been removed from the Windows Configuration utility. (For several previous StorNext releases a message warned that this tab would be removed in an upcoming release: “WARNING: Active Directory will be the only mapping method supported in a future release. This dialog will be deprecated.”)
	When a StorNext file system is mounted to a drive letter or a directory, configure the Windows backup utility to NOT include the StorNext file system.

Operating System / Affected Component	Description
Windows	If you are using the StorNext client software with Windows Server 2003, Windows Server 2008, Windows XP, or Windows Vista, turn off the Recycle Bin in the StorNext file systems mapped on the Windows machine. You must disable the Recycle Bin for the drive on which a StorNext file system is mounted. Also, each occurrence of file system remapping (unmounting/mounting) will require disabling the Recycle Bin. For example, if you mount a file system on E: (and disable the Recycle Bin for that drive) and then remap the file system to F:, you must then disable the Recycle Bin on the F: drive. As of release 3.5, StorNext supports mounting file systems to a directory. For Windows Server 2003 and Windows XP you must disable the Recycle Bin for the root drive letter of the directory-mounted file system. (For example: For C:\MOUNT\File_System you would disable the Recycle Bin for the C: drive.) For Windows Server 2008 and Windows Vista you must disable each directory-mounted file system. For Windows Server 2003 or Windows XP: **1**On the Windows client machine, right-click the **Recycle Bin** icon on the desktop and then click **Properties**. **2**Click **Global**. **3**Click **Configure drives independently**. **4**Click the **Local Disk** tab that corresponds to the mapped or directory-mounted file system. **5**Click the checkbox **Do not move files to the Recycle Bin. Remove files immediately when deleted**. **6**Click **Apply**, and then click **OK**. For Windows Server 2008 and Windows Vista: **1**On the Windows client machine, right-click the **Recycle Bin** icon on the desktop and then click **Properties**. **2**Click the **General** tab. **3**Select the mapped drive that corresponds to the StorNext mapped file system. For directory-mounted file systems, select the file system from the list. **4**Choose the option **Do not move files to the Recycle Bin. Remove files immediately when deleted**. **5**Click **Apply**. **6**Repeat steps 3-5 for each remaining directory-mounted file system. **7**When finished, click **OK**.

Operating System / Affected Component	Description
All	Be aware of the following limitations regarding file systems and stripe groups: • The maximum number of disks per file system is 512 • The maximum number of disks per data stripe group is 128 • The maximum number of stripe groups per file system is 256 • The maximum number of tape drives is 256
	For managed file systems only, the maximum recommended directory capacity is 50,000 files per single directory. (This recommendation does not apply to unmanaged file systems.)
	Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.
	The StorNext Cluster-Wide Central Control file (`nss_cctl.xml`) is used to enforce the cluster-wide security control on StorNext nodes (client nodes, fsm nodes, and nodes running `cvadmin`). This file is placed on an nss coordinator server. Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.

Table 15 lists new guidelines and limitations, many of which are specific to StorNext 4.0.

Table 15 StorNext New Operating Guidelines and Limitations

Operating System / Affected Component	Description
AIX	Clients on AIX systems may not unmount a file system after running the `fsstress` command against that file system. The client must then be rebooted to release the mount. This issue occurs on AIX systems when the `fsstress` command is run using `mknod` on the same command line. To prevent encountering this behavior, do not include `mknod` when running `fsstress`. Otherwise, you will be required to reboot the client without a successful unmount.
Linux	StorNext users migrating their metadata controllers from Apple Xsan to Linux should be aware of the following upgrade considerations: • If the file system is running Xsan 2.1.1 or earlier, it should be a simple upgrade: just replace the MDC. • If the file system is running Xsan 2.2 or later with "NamedStreams No" (which is the default for Xsan 2.2,) it should also be a simple upgrade: just replace the MDC. • If the file system is running Xsan 2.2 or later with "NamedStreams Yes," you must completely remake (reformat) the file system. For obvious reasons, you should do a complete backup before migrating.
	SuSe Linux distributions automatically associate the FQDN of the local machine with the address 127.0.0.2 in the /etc/hosts file. There is no benefit from doing this when the machine is connected to a network that can resolve its name to an IP address. **However, the existence of this entry can sometimes cause a failure of configuration synchronization within and between the server computers in an HA configuration. For this reason, the 127.0.0.2 entry should be deleted from the /etc/hosts file.**
All	StorNext does not support hot-swapping tape drives. When replacing or adding new tape drives you must first stop StorNext before installing the new drive.
	If you are using the Deduplication or Replication feature, part of the installation process is to update the on-disk index. The time required to complete this part of the installation process times may vary depending on the size of your licensed blockpool, drive performance, and other factors. As a general guideline, allow approximately five minutes for a 10TB blockpool.

Operating System / Affected Component	Description
All	On HA systems only: The `/usr/cvfs/config/ha_peer` file supports some essential HA features by providing an address for HA administrative communications between the MDCs in an HA Cluster. If CVFS is started without this file having correct information, the probability of an HA Reset increases. To correct this condition, restore the `ha_peer` file to the IP address of the peer MDC, and restart StorNext by running the following command: `service cvfs restart` Note: The peer will be Primary after running this command. If the `ha_peer` file is removed for any length of time while StorNext is running, the snhamgr(1) HA Manager subsystem could stop functioning, which impacts the GUI HA Manage status page and the starting and stopping of CVFS, as well as any command line use of snhamgr itself. If this occurs, restore the `ha_peer` file to the IP address of the peer MDC, and then restart the HA Manager service by running the following command: `service snhamgr restart`
	In some cases the physical IP address must be included in the `dpserver` file in addition to the interface name. Note these conditions: • When there is one IP address associated with a NIC interface, the interface name alone is a sufficient identifier • If there are multiple IP addresses associated with a NIC interface, one IP address is required in addition to the interface name • On HA systems, the physical IP address is required if virtual IP is configured for the NIC interface. (See also the following entry, “Distributed LAN Clients in HA Environment.”)
	Distributed LAN Clients in HA Environments: Each HA node must have its own dpserver files detailing the NICs on that node. The dpserver files are not synchronized between HA pairs. If the Distributed LAN Server is configured after converting to HA, the file system(s) running as Distributed LAN servers must be unmounted and mounted again to service DLC requests. When deduplication/replication is enabled, one or more Virtual IP Addresses (VIPs) provides access to the Primary MDC (where the blockpool server is running). In StorNext startup and failover situations, the VIP is dynamically associated with a physical address on the Primary server. Do not use VIP interfaces when setting up the dpserver configuration file, or it will not be available when the node is running as Secondary. The physical interface and IP address should be used in this situation.

Operating System / Affected Component	Description
All	When creating or editing a replication storage policy, there is a field on the Outbound Replication tab called "**Filenames Excluded from Replication**." This field allows you to exclude specific files from the replication process. This field works the same way as a UNIX shell which lets you pattern match names. For example, entering `*.0 core` would exclude all .o files and also files named "core." You could also skip all core files by entering `rep_skip=core*`.
	Understanding the performance of FSM failover in StorNext High Availability installations: When a failover of any file system occurs, the new FSM notices if any clients had a file exclusively opened for writes, and waits up to 35 seconds for those clients to reconnect. In the case of an HA Reset of the Primary MDC, that MDC is not going to reconnect, so the failover to FSMs on the Secondary MDC and the promotion of that MDC to Primary status can be delayed by 35 seconds. The StorNext system exclusively opens files on the HaShared file system, but assumes that only the Primary MDC does this and waives the delay for that one file system. Quantum advises against running user processes other than StorNext processes on HA MDCs for performance, reliability and availability reasons. In the event that processes running on the Primary MDC have files exclusively open for writes on other file systems, the availability of those file systems to all clients will be delayed by 35 seconds following an HA Reset event.
	The `subtree_check` option controls NFS checks on a file handle being within an exported subdirectory of a file system. This option should be turned off in NFS exports.
	Before attempting to upgrade from a previous StorNext release, make sure you have free space on the file system. If the file system is nearly full when you begin the upgrade, serious errors may occur or the upgrade could fail. Best practice is to maintain an area on the file system which is not used for data or system files, but is reserved as an empty buffer to ensure that upgrades and other operations complete successfully.

Documentation

The following documents are currently available for StorNext products:

Document Number	Document Title
6-01658-09	*StorNext User's Guide*
6-00360-18	*StorNext Installation Guide*
6-01376-13	*StorNext File System Tuning Guide*
6-01620-12	*StorNext Upgrade Guide*
6-01688-09	*StorNext CLI Reference Guide*
6-00361-35	*StorNext File System Quick Reference Guide*
6-00361-36	*StorNext Storage Manager Quick Reference Guide*

Contacting Quantum

More information about this product is available on the Quantum Service and Support website at www.quantum.com/ServiceandSupport. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs). You can also access software, firmware, and drivers through this site.

To request a software upgrade, visit www.quantum.com/ServiceandSupport/Upgrade/Index.aspx. For further assistance, or if training is desired, contact Quantum Global Services:

North America	+1 800-284-5101
UK, France, and Germany	00800 4 QUANTUM
EMEA	+44 1256 848 766
World Wide Web	www.quantum.com/ServiceandSupport