



StorNext 3.5.1 Release Notes

Product	StorNext® 3.5.1
Date	July, 2010

Contents

Purpose of this Release	2
Enhancements and Improvements	2
Changes From Previous Releases	4
Configuration Requirements	7
Operating System Requirements	10
Supported Libraries and Tape Drives	12
Minimum Firmware Levels for StorNext Drives	16
Supported System Components	16
Hardware Requirements	18
Resolved Issues	19
Known Issues	25
Operating Guidelines and Limitations	37
Documentation	40
Contacting Quantum	40



Purpose of this Release

StorNext 3.5.1 is a maintenance release that improves behind-the-scenes performance and corrects issues identified in previous releases.

This document 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

Although 3.5.1 is a maintenance release, a few enhancements and improvements were incorporated into the release.

Platform Support

In release 3.5, support for some platforms was discontinued. In StorNext 3.5.1, support has been restored for the following platforms:

- Windows Server 2003 (SP2) for x86 32-bit: MDC Server and Distributed LAN Server
- Red Hat® Enterprise Linux 4 (Update 6 and 7) for x86 32-bit: MDC Server, File System SAN Client, Distributed LAN Server, and File System LAN Client
- Red Hat® Enterprise Linux 4 (Update 6 and 7) for x86 64-bit: MDC Server and Distributed LAN Server
- SUSE Linux Enterprise Server 10 (SP1 and SP2) for x86 32-bit: File System SAN Client and File System LAN Client
- SUSE Linux Enterprise Server 10 (SP1 and SP2) for Itanium 64-bit: MDC and File System SAN Client
- Sun Solaris 10 (Generic 120011-14) for SPARC 64-bit: MDC and Storage Manager

Added Tape Drive Support

Support for the Sun/StorageTek 9840D and T10000 Rev B tape drives has been added. For more information, see [table 2](#) on page 13.

Note: Issue 27942 in the [StorNext Storage Manager Known Issues](#) section describes a condition related to tape drive cleaning that could affect 9840D tape drives. See the issue description for more information.

New Remote Notification Feature

The Windows operating system provides functionality known as Directory Change Notification that allows a process to be informed when updates occur on files within a particular directory hierarchy. The StorNext 3.5.1 release adds partial support for this feature.

The list of events supported includes only `FILE_NOTIFY_CHANGE_FILE_NAME` and `FILE_NOTIFY_CHANGE_DIR_NAME`. Another limitation is that only a specific directory can be monitored, and not the full hierarchy beneath it.

Remote Notification is enabled through the `RemoteNotification` variable in the configuration (`.cfg`) file's global section. The default value is `No`, which disables the feature.

Note: This option is not intended for general use. Use this feature only when recommended by Quantum Support.

New Media Limiting Feature

StorNext users can now minimize the number of new tape media that can be used for stores by a single policy class. In most cases one new media will be used per policy class at a time. A new piece of media must be filled before another new media can be allocated for the policy class.

It is necessary to understand how Storage Manager allocates media to understand when only one piece of media will be used and in what circumstances several media might be used.

Consider the following example. In standard StorNext Storage Manager without the new feature, suppose we have a group of 1200 files that need to be stored for policy class `pc1`. Assume there are four tape drives available, and they are currently idle. Assume also that there is a supply of blank unassigned media, and that there are no partially used media currently assigned to policy class `pc1`. Storage Manager splits its files to be stored into groups of 300 (the default value), so in this example the list of 1200 files to be stored will be split into four groups of 300 files each. Four `fs_fmover` processes can be copying the files to four different newly allocated media. This is good because it achieves parallelization for the copies for class `pc1`, but a side effect is that four different media are now assigned to class `pc1`, and each of the four media might be only partially used.

The new media limiting feature allows the StorNext administrator to select a different media allocation strategy. The alternate allocation strategy will allocate one new media for the policy class, and files will be written only to that media until it fills up. Then a new media will be allocated and stores are written only to that media. And so on. This feature maximizes storage media usage per policy class by eliminating the possible parallelization of copy operations (stores) for the policy class.

When preparing to store a set of files, Storage Manager looks to see if there are any media already assigned to the policy class that can hold at least one of the files in the list. If so, the store operation proceeds, using that piece of media. If there are no such available media, a blank tape is requested. That blank tape will be used for the store operation and will then be owned by the policy class.

To limit the media used per policy class, insert this line:

```
LIMIT_MEDIA_PER_CLASS=y;
```

in the file `/usr/adic/TSM/config/fs_sysparm_override`

You must stop and restart Storage manager for this change to take effect. If you wish to disable the feature, remove that line from `/usr/adic/TSM/config/fs_sysparm_override` and then stop and restart Storage Manager.

Enabling the feature prevents Storage Manager from requesting a new blank tape if there is a piece of media owned by the policy class containing enough space to hold at least one of the files in the current store list. However, if there are several media already assigned to the policy class they can still be used to store files at the same time.

You must be careful if no media are currently assigned to the policy class, because a new piece of media is not really assigned to a policy class until a copy operation to that media has successfully completed.

Consider the example above in which we wanted to store 1200 files. Even with the limiting feature enabled, four media would be used because files have been written to each of the four media. Four store operations (300 files each) would be launched almost simultaneously. Each of these store operations will query the media database and conclude that there are no writable media assigned to the policy class. Each of the four store operations will request a new blank tape. When the operations have completed, four media are now assigned to the policy class.

This behavior can be avoided by seeding the process with one file. If we stored one file initially (with the `fsstore` command, for example) then there would be one piece of media assigned to the policy class. Then, if we stored 1200 files, each of the stores of 300 files would wait for that piece of media. If that piece of media fills, a new piece will be assigned and stores will be able to continue.

Changes From Previous Releases

The following changes were instituted in a previous StorNext release and are listed here as a reminder that important settings have been changed.

Configuration File Changes

The following variables have been deprecated (removed) from the configuration file's "Globals" section:

- AttrTokenSize
- BufferPoolSize
- DirCacheSize
- DirFDCacheSize
- ForceStripeAlignment
- IoHangLimitSecs
- JournalIcBufNum
- JournalIcBufSize
- MaxMBPerClientReserve

- Mbufs
- MbufSize
- ReaddirForcedVersion
- StaticInodes

In addition, the following settings have been removed from the configuration file's "StripeGroup" section:

- StripeClusters
- Type

For detailed explanations regarding why these items were deprecated, see the `cvfs_config(4)` man page.

ONC Portmapper Services Deprecated

ONC Portmapper Service is used by remote network applications to locate the correct port and application on the local computer. This service was used by StorNext version 2.6 and earlier to locate the correct port for FSMPM/Name Services. Since StorNext version 3.5 is not supported in the same network environment as StorNext 2.6, Quantum's ONC Portmapper is no longer installed or enabled.

Although unlikely, there might be other non-Quantum applications which rely on ONC Portmapper Services. In those cases the vendor should provide their own ONC portmapper service. You may need to re-enable their service.

Single LUN Stripe Group I/O Characteristics

The I/O characteristics on a single LUN stripe group have changed in release 3.5. In prior releases all I/Os were broken down into, at most, StripeBreadth-sized I/Os aligned on the stripe group's StripeBreadth value. This was the case even when there was a single LUN in a stripe group. In StorNext 3.5, I/Os are no longer forced into StripeBreadth-sized chunks when there is only one LUN in the stripe group. All I/Os in this case should match the DMA I/O requests, or the I/O requests coming out of the file system's buffer cache. Be aware of the change in behavior when evaluating performance characteristics between differing StorNext releases.

Revised FsBlockSize, Metadata Disk Size, and JournalSize Settings

The FsBlockSize (FSB), metadata disk size, and JournalSize settings all work together. For example, the FsBlockSize must be set correctly in order for the metadata sizing to be correct. JournalSize is also dependent on the FsBlockSize.

For FsBlockSize the optimal settings for both performance and space utilization are in the range of 16K or 64K.

Settings greater than 64K are not recommended because performance will be adversely impacted due to inefficient metadata I/O operations. Values less than 16K are not recommended in most scenarios because startup and failover time may be adversely impacted. Setting FsBlockSize (FSB) to higher values is important for multi-terabyte file systems for optimal startup and failover time.

Note: This is particularly true for slow CPU clock speed metadata servers such as Sparc. However, values greater than 16K can severely consume metadata space in cases where the file-to-directory ratio is low (e.g., less than 100 to 1).

For metadata LUN size, you should have a *minimum* of 25 GB, with more space allocated depending on the number of files per directory and the size of your file system.

The following table shows suggested FsBlockSize (FSB) settings and metadata disk space based on the average number of files per directory and file system size. The amount of disk space listed for metadata is *in addition* to the 25 GB minimum amount. Use this table to determine the setting for your configuration.

Average No. of Files Per Directory	File System Size: Less Than 10TB	File System Size: 10TB or Larger
Less than 10	FSB: 16KB Metadata: 32 GB per 1M files	FSB: 64KB Metadata: 128 GB per 1M files
10-100	FSB: 16KB Metadata: 8 GB per 1M files	FSB: 64KB Metadata: 32 GB per 1M files
100-1000	FSB: 64KB Metadata: 8 GB per 1M files	FSB: 64KB Metadata: 8 GB per 1M files
1000 +	FSB: 64KB Metadata: 4 GB per 1M files	FSB: 64KB Metadata: 4 GB per 1M files

The best rule of thumb is to use a 16K FsBlockSize unless other requirements such as directory ratio dictate otherwise.

This setting is not adjustable after initial file system creation, so it is very important to give it careful consideration during initial configuration.

Example: FsBlockSize 16K

JournalSize Setting

The optimal settings for JournalSize are in the range between 16M and 64M, depending on the FsBlockSize. Avoid values greater than 64M due to potentially severe impacts on startup and failover times. Values at the higher end of the 16M-64M range may improve performance of metadata operations in some cases, although at the cost of slower startup and failover time.

The following table shows recommended settings. Choose the setting that corresponds to your configuration.

FsBlockSize	JournalSize
16KB	16MB
64KB	64MB

This setting is adjustable using the **cvupdatefs** utility. For more information, see the **cvupdatefs** man page.

Example: JournalSize 16M

Configuration Requirements

Before installing StorNext 3.5.1, 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 impact the library and cause data loss or corruption. For more information, see StorNext Product Alert 16.

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/rdisk/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 than one SAN, a conflict will arise if the client sees two disks with the same name.

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 following entries 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
```

Operating System Requirements

[Table 1](#) shows the operating systems, kernel versions, and hardware platforms that support StorNext File System, StorNext Storage Manager, and the StorNext client software.

This table also indicates the platforms that support the following:

- MDC Servers
- Distributed LAN Servers
- File System LAN Clients

Table 1 StorNext Supported OSES and Platforms

StorNext 3.5.1 File System and Storage Manager 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 / SNAP1
Windows Server 2003	R2 SP2**	x86 32-bit	✓	✓	✓*	✓	
		x86 64-bit	✓	✓	✓*	✓	
Windows XP	SP2 SP3	x86 32-bit		✓		✓	
		x86 64-bit		✓		✓	
Windows Vista	SP1	x86 32-bit		✓		✓	
		x86 64-bit		✓		✓	
Windows Server2008	SP1	x86 32-bit		✓		✓	
		x86 64-bit	✓	✓	✓*	✓	
Red Hat Enterprise Linux 4	2.6.9-67.EL (Update 6) 2.6.9-78.EL (Update 7)	X86 32-bit x86 64-bit	✓	✓	✓	✓	
	2.6.9-89 EL (Update 8)	X86 32-bit x86 64-bit	✓	✓	✓	✓	✓ ⁱ
Red Hat Enterprise Linux 5 (all versions, excluding virtualization)	2.6.18-53.EL (Update 1) 2.6.18-92.EL (Update 2) 2.6.9-8x EL (Update 3)	x86 64-bit	✓	✓	✓	✓	✓
SUSE Linux Enterprise Server 10	2.6.16-46-0.12 (SP1) 2.6.16-60-0.27 (SP2)	X86 32-bit		✓		✓	
		x86 64-bit	✓	✓	✓	✓	✓
		Itanium 64-bit	✓	✓			
Sun Solaris 10	Generic 120011-14 Generic 127128-11	SPARC 64-bit	✓	✓			✓
		Opteron x86 64-bit		✓		✓	
		Intel x86 64-bit		✓		✓	
IBM AIX	5.3 ⁱⁱ	64-bit Power Architecture		✓			
HP-UX	11i v2 ⁱⁱ	IA-64		✓			
SGI-IRIX	6.5.30	64-bit MIPS		✓			
Apple Mac OS X	10.5.5 with Xsan 2.1.1, StorNext MDC	x86 32-bit		✓			

Notes:

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 3.5.1 server that is not supported, you do so at your own risk. Quantum strongly recommends against installing non-supported servers.

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

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

‡In order to install StorNext Storage Manager on a Red Hat Enterprise Linux 4 system, you MUST first install 2.6.9-98 EL (Update 8).

‡StorNext support will transition from HP-UX 11i v2 to 11i v3, and from IBM AIX 5.3 to 6.1 on a future date.

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 secure Linux kernel enabled by default. To ensure proper StorNext operation, you must not install Red Hat 5 with secure Linux enabled. The secure Linux kernel must be off, or the file system could fail to start.

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 3.5.1 are presented in [Table 2](#). Where applicable, minimum firmware levels for libraries are provided.

Table 2 StorNext Supported Libraries and Tape Drives

StorNext 3.5.1 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		See library firmware requirement
			IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
			IBM LTO-3 WORM		
			IBM LTO-4 WORM		
			HP LTO-4		
			HP LTO-4 WORM		
	Scalar 24	Minimum: 107A.GY0002	DLT-S4	Minimum: 1F1F	
			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		
			IBM 3590B1A		
AIT-1					
Scalar 10K	Minimum: 110A.00001	IBM LTO-1		See library firmware requirement	
		IBM LTO-2			
		IBM LTO-3			
		IBM LTO-4			
		IBM LTO-3 WORM			
		AIT-2			
		AIT-2 WORM			
PX500	Minimum: 001A	HP LTO-3		Not including WORM 30.0	

Note: 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 3.5.1 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	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			
		LTO-4 WORM			
IBM	TS3500	Minimum: 4680	IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
			IBM 3592		
Qualstar	XLS	Minimum: 0880	IBM LTO-3		
			IBM LTO-4		
Sony	Petasite CSM-200	Minimum: 6.30	IBM LTO-4 drive (T1600)		

Note: 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 3.5.1 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 SCSI/FC Libraries	9740	Minimum: 2000	Sun/STK 9840			
			Sun/STK 9940			
	L700	Minimum: 2.36	T10000 Rev A*			
			T10000 Rev B			
			9840D			
	L180	Minimum: 2.00	T10000 Rev A*			
Sun / StorageTek ACSLS (pre-7.3) Libraries	9310	Minimum: None				
	9710	Minimum: None				
	9740	Minimum: 2000				
	L5500	Minimum: None				
	L700	Minimum: 2.36	Sun/STK 9840, Sun/STK 9940, Sun/STK 9940B, T10000 Rev A*, HP LTO-2, HP LTO-4, IBM LTO-2, IBM LTO-3, IBM LTO-4			
	L180	Minimum: 2.00				
	SL8500	Minimum: N / A				
	SL500	Minimum: 10.67			Not including WORM	
	SL3000	Minimum: 0205	HP LTO-4			
			IBM LTO-4			
T10000 Rev A*						
T10000 Rev B						
9840D						
SL8500	Minimum: None	HP LTO-4				
		IBM LTO-4				
		T10000 Rev A*				
		9840D	Minimum: 1.43 Recently Tested: May 2009			
Sun / StorageTek ACSLS 7.3 Libraries	SL3000	Minimum: 0205	HP LTO-4			
			IBM LTO-4			
			T10000 Rev A*			
			T10000 Rev B			
			9840D			
SL8500	Minimum: None	HP LTO-4				
		IBM LTO-4				
		T10000 Rev A*				
		9840D				

Notes: Sun/Storage Tek no longer supports ACSLS version 6.x. Version 7.0 or higher is required.

* When using T10000 Rev A drives, the STK library parameter Fastload must be set to OFF.

Minimum Firmware Levels for StorNext Drives

Where applicable, the minimum firmware levels for StorNext-supported drives are shown in [Table 3](#).

Table 3 Minimum Firmware Levels for Drives

StorNext 3.5.1 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 System Components

System components that are supported for use with StorNext 3.5.1 are presented in [Table 4](#).

Table 4 StorNext Supported System Components

Component	Description
Tested Browsers	Internet Explorer 6.0 or later (up to 8.x) Mozilla Firefox 2.0 or later (up to 3.x) (Minimum browser resolution: 800x600) NOTE: Disable pop-up blockers.

Component	Description
LTO-1 Media and LTO-3 or LTO-4 Tape Drive Compatibility	<p>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:</p> <ol style="list-style-type: none"> 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. <p>NOTES:</p> <ul style="list-style-type: none"> • 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.
NFS	<p>Version 3</p> <p>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.</p> <p>This issue may be addressed in a future StorNext release. As a workaround, add the <code>no_wdelay</code> option to each line in the <code>/etc/exports</code> file that references a StorNext file system. For example, typical export options would be <code>(rw, sync, no_wdelay)</code>.</p>
Addressable Power Switch	<p>WTI RPS-10m WTI IPS-800</p> <p>The RPS-10m (master) is supported. The RPS-10s (slave) is not supported.</p>
LDAP	<p>LDAP (Lightweight Directory Access Protocol) support requires Windows Active Directory.</p>

Hardware Requirements

To successfully install StorNext 3.5.1, the following hardware requirements must be met:

- [StorNext File System and Storage Manager Requirements](#) on page 18
- [StorNext Client Software Requirements](#) on page 19

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 5](#).

Table 5 File System and Storage Manager Hardware Requirements

No. of File Systems	RAM	File System Disk Space	Storage Manager Disk Space
1–4*	2 GB	2 GB	<ul style="list-style-type: none">• 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**	4 GB	4 GB	

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

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

Note: If a file system uses deduplicated storage disks (DDisks), note the following additional requirements:

- Requires 2 GB RAM per DDisk in addition to the base RAM noted in [Table 5](#).
 - Requires an additional 5GB of disk space for application binaries and log files.
 - Deduplication is supported only for file systems running on a Linux operating system (x86 32-bit or x86 64-bit).
 - An Intel Pentium 4 or later processor (or an equivalent AMD processor) is required. For best performance, Quantum recommends an extra CPU per DDisk.
-

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.

Resolved Issues

The following sections list resolved issues in this release of StorNext:

- [StorNext File System Resolved Issues](#) on page 19
- [StorNext Storage Manager Resolved Issues](#) on page 22
- [StorNext GUI and Installation Resolved Issues](#) on page 24
- [StorNext Miscellaneous Resolved Issues](#) on page 24

Note: There is no change to cryptographic functionality in StorNext release 3.5.1.

StorNext File System Resolved Issues

[Table 6](#) lists resolved issues that are specific to StorNext File System.

Table 6 StorNext File System
Resolved Issues

Operating System	CR Number	SR Number	Description
SGI IRIX	26976	724406	Support added for RAID/disks with 4KB sectors on IRIX machines.
Solaris	26985	n/a	Corrected a condition that prevented mount failures from being included in system logs on Solaris 10 platforms.

Operating System	CR Number	SR Number	Description
Solaris	27495	n/a	Corrected a condition in which running the cvlabel command wrote VTOC labels that are incompatible by default with recent releases of Solaris 10.
Linux	26274	901844	Corrected a condition in which asyn msync failed on mmapped files when the application is terminated by a signal.
	26656	931830	Resolved a condition in which running llistxattr occasionally failed.
	26854	597852, 888502, 987654	Corrected a condition in which running the ls -l command returned incomplete information on a file.
	26993	933846	Corrected the man page for io_penalty_time.
Windows	26149	914828 928754 941848	Corrected a condition related to StorNext counters in the Windows Performance Monitor in which instperfreg.exe and instperfreg.exe do not put StorNext on the Windows 2008 Perfmon.
	26414	722239	Corrects a condition in which the CvReadDir command entered an infinite loop.
	26852	n/a	Resolved a condition in which CvRtlGetVersion was called at the wrong IRQ level, resulting in a blue screen.
	26951	n/a	Corrected issues with the Windows Performance Monitor on 64-bit platforms.
	27268	971688	Addressed a condition in which Microsoft IIS version 7 Web Hosting software did not operate.
	27472	970490 973766	Resolved a blue screen condition related to OplockOpenChange on newer versions of Windows.
	27726	971688	Corrected a condition in which Windows IIS version 7 Web Hosting software did not work correctly with StorNext File System.
All	26281	932342 959134	Corrected a condition in which cvbtree code failed to insert high keys after certain removal patterns.
	26302	909788	The stripe group number is now displayed in cvfsck when doubly allocated extents are detected.
	26868	955980	Corrected a condition in which cvfsck could overwrite icb when clearing xattr chains.
	26938	867760	Corrected a condition in which searchfs hung when renaming long-named files.
	26946	892030	Addressed a condition in which cvcp -xyz sometimes failed to update mtime correctly.

Operating System	CR Number	SR Number	Description
All	26953	919864	Addressed an FSM panic in alloc_space_debug.
	26964	946732	Corrected a condition in which running cvfsck did not fix corruption as adequately as the 3.1.2 version.
	26982	751246	Dump_rj was modified to collect information about a specific inode or a range of lsn entries.
	26986	958296	Added the ability to place orphaned inodes into a specified directory inode.
	27002	956004	Added Quality of Service (QOS) support for Distributed LAN clients and servers.
	27009	955194	Addressed an FSM panic that occurred because trans_write: space reservation was too small for type 6.
	27367	962028	Addressed a condition in which data corruption sometimes occurred when bringing UP a DOWN stripe group.
	27484	985224	Corrected an assert that occurred when running snmetadump.
	27772	968024	Corrected a condition which caused metadata dumps to fail on systems with multiple metadata stripe groups.
	27862	871834, 877874	Corrected a condition in which running an external API (ExtApiGetSglInfo) returned a static stripe group status instead of a dynamic stripe group status.
	27932	996254	Corrected a condition in which performing a metadata dump caused problems for directories in multi metadata stripe groups.
	27964	n/a	Corrected a condition in which metadata on a file system with multiple stripe groups could not be correctly restored.
	27967, 27943	n/a	Corrected a system panic caused by multiple clients accessing a file at the same time.
27982	n/a	Resolved a metadata dump failure that occurred during the initial optimization phase.	

StorNext Storage Manager Resolved Issues

[Table 7](#) lists resolved issues that are specific to StorNext Storage Manager.

Table 7 StorNext Storage Manager Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	26753	896404	Corrected a condition in which conflicting CONF manifest numbers prevented StorNext system backup from running successfully.
Windows	27167	832794	Corrected a condition which caused a disconnection and error message due to an erroneous call to FindNicServiceName.
All	25314	871834, 877874	Corrected a condition in which a file system containing several offline stripe groups prevented truncation from occurring. Note: If you applied the workaround for this issue before installing release 3.5.1, you must “undo” the workaround by using the StorNext GUI to reset the truncation values to either their original or default values.
	26054	878862, 842072, 942922	Corrected a condition in which running fsmedcopy -r failed when the destination media was full.
	26292	909624	Corrected a condition in which some large minsetsize settings did not trigger stores.
	26385	887578	Corrected a condition in which Storage Manager filled tapes with wasted space if the inode class was incorrect.
	26388	909878	Max channels has been adjusted in the StorNext database.
	26552	n/a	Corrected a condition which caused a fs_dmapi fault due to long file names.
	26576	927346	Addressed a condition in which LibManager processes could become defunct if SL_TAC_MASK is not set to the default value in /usr/adic/MSM/logs/log_params.
	26579	933846	Corrected a condition in which fs_fmover hung due to a StorNext database issue.
	26624	n/a	Added support for Sun/STK T10000 rev B tape drives.
	26809	947962	Addressed a condition in which fsretrieves from the same media had to wait for the unmount delay.

Operating System	CR Number	SR Number	Description
All	26855	n/a	Addressed snmetadump failures on systems with multiple metadata stripe groups.
	26856	n/a	Corrected a condition in which fsmedlist output showed media as protected when they were not.
	26949	n/a	Addressed a condition in which HA conversion did not move the database after OS conversion.
	26950	948026, 946116, 958332	Corrected a condition that caused StorNext tables to become corrupted. For more information, see StorNext Product Alert 29
	26954	n/a	Support for the Sun/STK SL3000 library has been added to StorNext.
	26957	n/a	Cvgather/snapshot now gathers capacity and usage information.
	26979	795978	Improved logic used when determining the Storage Manager API host name.
	26991	787774	Corrected a condition that prevented the t_parent process from shutting down.
	27178	n/a	Corrected a condition in which the service log closed stdin after a executing a process.
	27485	985224	Addressed an snmetadump ASSERT failure: Processing Error: snmetadump ASSERT failed "max_length < length" file process_metadata.c, line 860 snmetadump error at line 602 of snmetadump.c: Dump file error flag set (16): New metadump needed.
	27746	n/a	Corrected a condition in which tape media was ejected on failover due to a reservation conflicted.
	27762	985494, 983646	Resolved a condition in which fsretrieve never completed.
	27865	961840	Corrected a condition in which files with a percentage symbol (%) in their names could not be recovered.

StorNext GUI and Installation Resolved Issues

[Table 8](#) lists resolved issues that are specific to the StorNext GUI or the installation process.

Table 8 StorNext GUI and Installation Resolved Issues

Operating System	CR Number	SR Number	Description
Windows	27519	978244	Corrected a condition in which popup windows did not work correctly in Internet Explorer 8.
All	26582	779840	Corrected a condition in which the Apache component modified the httpd.conf file to set the user directory to disabled.
	26959	n/a	Addressed a condition in which performing an upgrade from StorNext 3.1.1 build 12 to build 26 failed on Storage Manager.
	27847	993414	The StorNext GUI prevents successfully adding a new disk to an existing stripe group. The configuration file is modified by adding the new disk but not adding it to the stripe group, which causes errors in future work on the file system.

StorNext Miscellaneous Resolved Issues

[Table 9](#) lists miscellaneous resolved issues that are not specific to any of the previous StorNext categories.

Table 9 StorNext Miscellaneous Resolved Issues

Operating System	CR Number	SR Number	Description
All	26415	n/a	On HA systems, cvpower for ips800 now supports multiple ips800's for redundancy.
	26973	876460	Corrected a condition on HA systems in which cvpower hung after upgrading to StorNext 3.1.2 build 10.

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](#) on page 25
- [StorNext Storage Manager Known Issues](#) on page 30
- [StorNext GUI Known Issues](#) on page 33
- [StorNext Installation Known Issues](#) on page 34

StorNext File System Known Issues

[Table 10](#) lists known issues that are specific to StorNext File System.

Table 10 StorNext File System Known Issues

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.
	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 <code>reboot -- -r</code> 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)
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 fsmgm 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
	24890	836284	DMA I/Os may hang when the file offset is not aligned and disks configured with 4096-byte sectors are used. For example: dd if=/stornext/sr83628a/file2g of=/dev/null bs=1880000 count=1 skip=40	This issue will be addressed in a future StorNext release. As a workaround, use the "memalign=4k" mount option. Note: this mount option should NOT be used with file systems containing LUNs configured with 512-byte sectors.
	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 may 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)
Windows	25551	892030	When running the command cvcp -xyz, some of the files copied do not have the mtime updated on the resulting new files.	This issue will be addressed in a future StorNext release. A workaround is to set the environment variable CV_BULK_SIZE=0. This stops cvcp from using bulk create, and the files are created serially.
	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.
	26019	n/a	Changing the Global Log level (Tools -> Global Options) or Name Server settings may result in a software hang or an error message "Error starting services" or the notice "Global options will not be activated until the next reboot..." The system behavior is dependent upon internal timing of stopping SNFS services. StorNext Services are stopped when: <ul style="list-style-type: none"> • Setting Global Options in the Client Configuration tool. • Changing the Name Servers configuration. • Using Start->All Programs->StorNext File System menus "Services Stop" or "Services Stop and Remove". Behavior may appear similar to that of a hung system, but has been seen to eventually recover to the desired state. In some cases a hard reboot has been necessary.	This issue will be addressed in a future StorNext release. Workaround options include disabling DLC Server (if enabled) and unmounting file systems before changing Global Log levels, and then rebooting after this operation. (Rebooting may not be necessary, but is recommended.) Normally the hang is temporary (but can last up to 3-4 minutes) at which point services can be started manually with "Services Start". If the problem persists, a reboot may be necessary. Configuration changes will take affect after the computer reboots. Note: Rebooting after changing settings avoids this issue. Unmounting all file systems prior to changing global setting also reduces the risk of this issue occurring.
	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.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	15032	n/a	<p>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.</p>	<p>Promise is working on a solution to this problem, but in the meantime they have provided the following workaround:</p> <ol style="list-style-type: none"> 1. Install the PerfectPath software on your Windows Server 2008 64-bit system. 2. Restart your system. The login prompt will <i>not</i> 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\DRIVE RS\ 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)
All	25124	833282	<p>A “no space” error appears after you attempt to add a file to a directory known to contain sufficient space.</p> <p>There has been only one known instance of this problem, which is a corrupted B-tree for one directory. Please report any additional instances to Quantum Support.</p> <p>Identify this problem by looking for these indicators:</p> <ol style="list-style-type: none"> 1. Attempts to add a file to an affected directory produce a “no space” error. 2. Files can be added to other directories in the same file system without error. 3. Running the command <code>cvadmin -F <FileSystemName> -e show</code> reports free space in the metadata stripe groups for the file system of the affected directory. 	<p>After collecting debugging information, the following workaround can be applied:</p> <ol style="list-style-type: none"> 1) Rename the corrupt directory <code><dir></code> to <code><dir>.corrupt</code> 2) Create a new directory named <code><dir></code> 3) Move all files from the corrupt directory <code><dir>.corrupt</code> to the newly created directory <code><dir></code> 4) Delete the empty corrupt directory <code><dir>.corrupt</code>
	25663	n/a	<p>The configuration file contains a MediaType section that has not been used for several StorNext releases. Leaving this section in the configuration file(s) for StorNext 3.5 could cause parser breaks.</p>	<p>If any configuration files contain a MediaType section, remove or comment out the section.</p>
	25670	n/a	<p>Using the values “up” or “down” in affinities or any other string-type keyword in the configuration file caused a parser error.</p>	<p>To avoid parser errors, do not use “up” or “down” when naming items in the configuration file.</p>
	26114	n/a	<p>Running <code>cvfsck</code> can crash when a data file is empty.</p>	<p>This issue will be addressed in a future StorNext release.</p>
	27462 27592	976842	<p>The default inode expansion value might be too low, which could lead to fragmentation.</p>	<p>This issue may be addressed in a future StorNext release.</p> <p>The workaround is to increase the inode expansion value <code>InodeExpand()</code>. A general guideline is to make the inode expansion value equal to approximately 10% of the file size.</p>

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

StorNext Storage Manager Known Issues

[Table 11](#) lists known issues that are specific to StorNext Storage Manager.

Table 11 StorNext Storage Manager Known Issues

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

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	24649	n/a	<p>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.</p>	<p>There are two possible workarounds, which also apply to versions of StorNext prior to 3.1.2.</p> <p>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 <code>/usr/adic/TSM/util/fs_scsi</code> on the metadata controller which owns the reservation:</p> <pre># /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)</pre> <p>OR</p> <p>2) Add the following setting to the <code>/usr/adic/TSM/config/fs_sysparm_override</code> file: <code>FS_SCSI_RESERVE=none;</code> ("none" means don't try to reserve tape drives.) StorNext must be restarted for this change to take effect.</p> <p>WARNING: This workaround could leave tape drives exposed to unexpected access by other systems, which could lead to data loss.</p>
Windows	30112	1078374	<p>On a Windows 2003 R2 SP2 system, disabling the network connection to induce failover resulted in the MDC locking up. This behaviour does not occur after inducing failover by unplugging the network cable.</p>	<p>The workaround is to fail over by physically unplugging the network cable rather than simply disabling the network connection.</p>

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
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%.
	25665	897730	Large files could cause a store policy to hang.	This issue will be addressed in a future StorNext release.
	25743	n/a	Deleting ddisk does not kill the blockpool process when disk proxy is enabled.	This issue will be addressed in a future StorNext release.
	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.
	26947	951446	Storage Manager does not detect when its drive configuration becomes out of sync with the Archives internal configuration.	This issue will be addressed in a future StorNext release.
	26115	n/a	Failover did not execute properly after an upgrade.	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	27942	n/a	<p>On T9840 tape drives: After cleaning finishes, the cleaning media ejects but remains in the throat of the drive for twenty minutes, at which time a RAS ticket indicating a failure is generated.</p> <p>However, after the twenty minute interval passes and the RAS ticket is generated, the cleaning media is dismounted and returned to its home slot, and the StorNext GUI indicates the cleaning operation was successful.</p>	<p>This issue will be addressed in a future StorNext release.</p> <p>You can verify whether the drive cleaning was successful by checking the file <code>/usr/adic/TSM/logs/tac/tac_00</code>.</p> <p>When cleaning completes successfully you should see messages similar to this:</p> <pre>fs_drv[497]: E1202(9)<1914158667>: msc9drv84: DRIVE CHECK: Media CLN<mediaID>: Drive / dev/<device> fs_resource[6442]: E1202(9)<1914158667>:msc9cl n506: {ProcessID}: RESOURCE CLEAN DRIVE FINISHED; PROCESSING REQUEST <RequestId></pre>

StorNext GUI Known Issues

[Table 12](#) lists known issues that are specific to the StorNext GUI process.

Table 12 StorNext GUI Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	24326	n/a	On an HA system, deleting a file system through the StorNext GUI could not successfully complete due to contradictory constraints in the GUI.	<p>This issue will be addressed in a future StorNext release.</p> <p>The workaround is to manually delete the file system instead of using the StorNext GUI. Contact Quantum Technical Support for assistance.</p>
	25685	n/a	<p>The StorNext GUI will label all disks in a stripe group according to the single label type (EFI or VTOC) specified per stripe group.</p> <p>The GUI will overwrite any pre-existing labels if (and <i>only</i> if) the label type is changed from VTOC to EFI or vice versa.</p>	<p>This issue will be addressed in a future StorNext release.</p>

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	26364	n/a	After installing the full SN3.5 installation disks, all of the client download files are present in the directory /usr/cvfs/CLIENTS, but the AIX client is not recognized by the StorNext GUI so cannot be downloaded.	This issue will be addressed in a future StorNext release. In the meantime, the workaround is to use the following commands: <pre>cd /usr/adic/www/downloads/ touch sn_dsm_irix65m_client.tar</pre>

StorNext Installation Known Issues

[Table 13](#) lists known issues that are specific to the StorNext installation process.

Table 13 StorNext Installation Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	23486	729897	For sites that have an AML library, when converting to HA the archives aren't accessible when the secondary server takes over.	This issue will be addressed in a future StorNext release. The workaround is to make sure entries are made in the sn_client_map.txt file for both peers. An entry must exist for each library peer combination.
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: <pre># mkdir /mnt/MOUNT_PATH # mount /dev/cdrom /mnt/MOUNT_PATH # cd /mnt/MOUNT_PATH</pre>

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux (RHEL5 and SUSE10 only)	25611	n/a	<p>In rare instances, problems in the SCSI generic driver included with RedHat Enterprise Linux 5 (RHEL5) and SUSE Linux Enterprise Server 10 (SLES10) could prevent data from being written to tape.</p> <p>This includes RHEL5 update 2 (RHEL5u2) and SLES10 service pack 1 (SLES10sp1). The problem is fixed in SLES10sp2. RHEL5u2 and SLES10sp2 are the most recent versions as of the date this document was written.</p>	<p>StorNext 3.5 includes changes that address this issue.</p> <p>When StorNext 3.5 starts up, the following parameters unique to the SCSI generic (sg) device driver are set for systems running RHEL5 and SLES10 (including SLES10 sp2):</p> <ul style="list-style-type: none"> • /sys/module/sg/parameters/allow_dio is set to 1 • /sys/module/sg/parameters/def_reserved is set to 524288 <p>These settings do not affect the standard "scsi tape" or "scsi disk" device drivers. This issue is addressed at length in StorNext Product Alert 15. This alert also provides a workaround for versions of StorNext prior to 3.1.2.</p>
Windows	25777	n/a	<p>Including the .auth_secret file in a non-XSan environment causes the file system to not communicate with the FSMPM.</p>	<p>This issue will be addressed in a future StorNext release.</p> <p>The workaround is to remove or change the name of the .auth_secret file, providing the client is connecting to a non-Xsan metadata controller. This workaround will not work if the client is connecting to an XSan and a non-XSan metadata controller.</p>
	27129	n/a	<p>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.</p> <p>Running Altiris Agent as a service can cause the installation to fail and then automatically reboot the machine.</p>	<p>This issue may be addressed in a future StorNext release.</p> <p>The workaround is to disable the Altiris Agent service before installing StorNext.</p> <p>Altiris Agent is typically located at C:\Program Files\Altiris\Altiris Agent\AeXNSAgent.exe.</p>

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.	<p>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.</p> <p>To work around this issue, follow these steps:</p> <ol style="list-style-type: none"> 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.
All	25192	n/a	VMware snapshots may not be used for virtual machines running StorNext. StorNext does not currently process the loss of state synchronization when a snapshot is restored, so incorrect behavior may result.	This issue may be addressed in a future StorNext release.

Operating Guidelines and Limitations

[Table 14](#) lists updated information and guidelines for running StorNext, as well as known limitations in this release.

Table 14 StorNext Operating Guidelines and Limitations

Operating System / Affected Component	Description
Windows	<p>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.</p> <p>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.</p> <p>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.</p> <hr/> <p>When you are upgrading to StorNext 3.5 from a release prior to version 3.0, you must uninstall StorNext before installing. After uninstalling you must reboot, install StorNext 3.5.</p>

Operating System / Affected Component	Description
Windows	<p>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.</p> <p>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.</p> <p>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.</p> <p>For Windows Server 2003 or Windows XP:</p> <ol style="list-style-type: none"> 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. <p>For Windows Server 2008 and Windows Vista:</p> <ol style="list-style-type: none"> 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
Windows	<p>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.")</p> <p>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.</p>
All	<p>Be aware of the following limitations regarding file systems and stripe groups:</p> <ul style="list-style-type: none"> • 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 <p>For managed file systems only, the maximum recommended directory capacity is 50,000 files per single directory. (This limitation does not apply to unmanaged file systems.)</p> <p>Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.</p> <p>The StorNext Cluster-Wide Central Control file (nss_ctl.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.</p> <p>Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.</p>

Documentation

The following documents are currently available for StorNext products:

Document Number	Document Title
6-01658-07	<i>StorNext User's Guide</i>
6-00360-16	<i>StorNext Installation Guide</i>
6-01376-11	<i>StorNext File System Tuning Guide</i>
6-01620-10	<i>StorNext Upgrade Guide</i>
6-01688-07	<i>StorNext CLI Reference Guide</i>
6-00361-31	<i>StorNext File System Quick Reference Guide</i>
6-00361-32	<i>StorNext Storage Manager Quick Reference Guide</i>

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