StorNext 4.0.1/4.0.1.1 Release Notes

Product	StorNext® 4.0.1/4.0.1.1
Date	August 2010

Contents

Document 6-00431-31 Rev C

Purpose of this Release

StorNext 4.0.1 is a maintenance release that includes bug fixes and adds support for additional libraries and drives.

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.

StorNext 4.0.1.1

StorNext 4.0.1.1 contains everything included in 4.0.1, plus two important fixes (described below) for Windows platforms which address potentially serious error conditions.

Operating System	CR Number	SR Number	Description
Windows	32039	n/a	Corrected a problem in CvUniClone which produced a blue screen error. This problem occurred when CvUniClone() accessed memory beyond the allocated memory amount.
	32041	n/a	Resolved an issue which prevented opening files with names that were 164 characters or more on 32-bit systems, or 194 characters or more on 64-bit systems.

Note: ***Unless otherwise noted, all information in this document pertains to both StorNext 4.0.1 and 4.0.1.1.***

StorNext Documentation

StorNext 4.0.1.1. uses the same documentation set as StorNext 4.0.1. This documentation is available here: http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx

Upgrading from StorNext 4.0.1 to 4.0.1.1

To ensure a successful upgrade from StorNext 4.0.1 to 4.0.1.1, before beginning the upgrade you must remove the contents of the /tmp/stornext directory.

If you do not remove the contents of the /tmp/stornext directory, installation files from the previous release may not be completely overwritten and the upgrade could fail. (For more information about this issue which will be addressed in a future StorNext release, refer to CR 31828.)

Enhancements and Improvements

This section describes the additions and enhancements in StorNext 4.0.1

Added Library Support

StorNext 4.0.1 adds support for the following libraries:

- Quantum Scalar i500 i6, and i6.1
- Quantum Scalar i2000 i6.7
- Quantum Scalar i6000 i8

Added Platform Support (Clients only)

StorNext 4.0.1 adds support for SLES 11 SP1 64-bit.

Added Tape Drive Support

- Quantum Scalar i500 (HP LTO-5 FC supported in i6x only)
- Quantum Scalar i20000/i6000 (HP LTO-5 FC supported in i6.7 and i8 only)
- Support for HP LTO-5 tape drives in the following libraries has been added:
  - HP ESL E Series
  - HP EML E Series
  - HP MSL G3 Series
  - Oracle (Sun/StorageTek) SL3000 and SL500 SCSI/FC direct libraries
- Support for HP LTO-5 WORM tape drives in HP ESL E Series libraries has been added.

StorNext GUI Improvements

A significant number of improvements which make StorNext easier to use have been made to the StorNext GUI for release 4.0.1. Improvements include the following:

- The process for creating a file system has been simplified and streamlined, although the same fields are still available.
- Fields on other screens have been consolidated, in some cases reducing the number of tabs on a screen.
- Both the hostname and IP address are now displayed in the upper right corner of each screen.
- Selecting items on a screen is now easier, in many cases accomplished by checkboxes.
- The **Back** and **Done** buttons are implemented more consistently throughout.

- The **Setup** menu has been renamed **Configuration** menu because its tasks could be ongoing. (“Setup” implies a one-time operation.)

Faster Replication Performance

The algorithm used for StorNext’s replication feature has been modified to improve performance, resulting in faster replication times. Depending on your configuration, number of replicated files and file size, the time savings could be significant.

Linux Device Mapper Multipath Support

StorNext 4.0.1 adds support for the Linux Device Mapper (DM) Multipath driver. This driver provides redundancy and improved I/O performance by taking advantage of multiple paths to storage.

If you plan to use the Linux DM Multipath support with StorNext, be aware of the following:

1. Not all RAIDs work with the DM Multipath Driver. Check with your storage vendor for compatibility.
2. For detailed instructions on installing and configuring the DM Multipath Driver, refer to the RedHat or SuSE documentation provided with your version of Linux.
3. For StorNext to use Linux Device Mapper Multipath devices, you must make two changes to the `/etc/multipath.conf` file .
   a. Set "`user_friendly_names`" to "`yes`"
   b. Quantum recommends that the cvfsctl devices not be included as multipath devices. This can be achieved by including the following in the blacklist entry: `devnode "cvfsctl*"`
4. Using the cvpaths file and udev rules configuration files is typically unnecessary with Linux Device-Mapper in StorNext 4.0.1.

On SuSE Linux Systems

In order to use Linux Device Mapper Multipath with Stornext, `/etc/multipath.conf` must be used because SuSE Linux does not install a multipath.conf, and Novell recommends against using it.

Although SuSE Linux does not install a multipath.conf file by default, an example file located at `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic` can be copied to `/etc/multipath.conf`.

On RedHat Linux Systems

Red Hat *does* install a multipath.conf file. By default, Red Hat multipath.conf file blacklists all multipath-capable targets. This means "blacklist { devnode "*"} must be commented out.

Instance Handling Enhancement

Beginning with release 4.0.1, the method used by StorNext Storage Manager to handle accessing instances has been enhanced. Previously, some third-party applications which use temporary files and renaming schemes to update files made it difficult or impossible to locate previous file instances. The new method makes it much easier to locate and recover old instances of files.

The fsrecover command has been enhanced to help manage old instances left by deleted or renamed files. For more information, see the man page for fsrecover.

This enhancement is not turned on by default. In order to enable this feature, you must set the MICRO_RENAME value in the TSM/config/fs_sysparm file. For more detailed information, see the description in the TSM/config/fs_sysparm.README file.

Note: StorNext continues to support *versioning*, which is different from instance handling. A *version* should not be confused with an *instance*.

A new *version* of a file is created when a file is modified and then stored again.

A new *instance* is created when a file is removed and a new file with the same name is created. Some applications "modify" a file by creating a new instance and replacing the existing instance with the newly created one. One instance of a file may or may not be related to a previous instance.

End User License Agreement Updated

StorNext 4.0.1 incorporates an updated version of the StorNext End User License Agreement (EULA).

No Cryptographic Changes

There is no change to cryptographic functionality in StorNext release 4.0.1.

Operating System Requirements

Table 1 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 4.0.1 server that is not supported, you do so at your own risk. Quantum strongly recommends against installing non-supported servers.

Table 1 StorNext Supported OSes and Platforms

StorNext 4.0.1 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 / SNAPI	Distributed Data Mover	Replication / Dedup Server
Windows Server 2003	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		✓		✓			

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

* 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 4.0.1 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 / SNAPI	Distributed Data Mover	Replication / Dedup Server
RHEL 4 See Note 1 See Note 2	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 See Note 1 See Note 2	2.6.18-53.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	✓	✓	✓	✓	✓	✓	✓
	2.6.18-194.EL (Update 5)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓

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.

‡ All releases of RHEL4 and RHEL5 except RHEL4U8 and RHEL5U4 / RHEL5U5 have a possible silent data corruption issue as documented in Product Alert #20. Quantum recommends that users migrate to RHEL4U8 or RHEL5U4 / RHEL5U5 or later as soon as possible.

1 The "Xen" virtualization software is not supported for RHEL 4 and RHEL5.

2 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: For systems running Red Hat Enterprise Linux version 4 or 5, before installing StorNext you must first install the following kernel files:

- Base kernel
- Kernel-header
- kernel-devel
- gcc-c development tools

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 4.0.1 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 / SNAPI	Distributed Data Mover	Replication / Dedup Server
SLES 10 *** See Note 2	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 ## *** See Note 2	2.6.27.19-5	x86 64-bit		✓		✓			
	2.6.32.12-0 (SP1)	x86 64-bit		✓		✓			
Sun Solaris 10	Generic 141444-09	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 (See Note 3)	Itanium 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 5	x86 64-bit		✓		✓			
Scientific Linux‡‡‡	Based on RHEL5 Update 5	x86 64-bit		✓		✓			
Oracle Linux‡‡‡	Base on RHEL5 Update 5	x86 64-bit		✓		✓			

SLES10 SP1, and certain SLES10 SP2 releases are sensitive to the silent data corruption issue documented in Product Alert #20. The problem has been fixed in SLES 10 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 work-around at this time.

*** 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.

2 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.

3 HPUX 11iv3 requires the "0909 Patch set"

Note: Although SGI IRIX clients cannot be upgraded to StorNext 4.0.1, StorNext 3.5.1 SGI IRIX clients may be used with a StorNext 4.0.1 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 4.0.1 are presented in Table 2. Where applicable, minimum firmware levels for libraries are provided.

Table 2 StorNext Supported Libraries and Tape Drives

StorNext 4.0.1 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 i5.1 (Quantum, Dell, IBM) i6 (Quantum) i6.1 (Quantum, Dell) i6.2 (IBM)	i5.1: 572G.GS002 i6, i6.1: 580G.GS003 i6.2: 587G.GSxxx	7404 I/O blade	i5.1: V05090501 i6, i6.1: V05100201	
			IBM LTO-2	8571, QTM i6 and i6.1 requires 93T0, Dell and IBM i5.1 / i6.1 do not support LTO-2	Library firmware upgrade may be required for LTO-3 WORM support
			IBM LTO-3	93G0, i6,and i6.1 require 93G6	
			IBM LTO-3 WORM		
			IBM LTO-4	94D4, FC requires 97FB, SAS requires 97F0, i6 and i6.1 require A232	
			IBM LTO-4 WORM		
			HP LTO-4 FC	H46Z, i6, i6.1 requires H57Z Only supported in Quantum-branded libraries	
			HP LTO-4 SAS	A45Z, i6, i6.1 requires H57Z Only supported in Quantum-branded libraries	
			HP LTO-5 FC	I24Z Only supported in Quantum i6.x	

StorNext 4.0.1 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	Scalar i2000 / i6000 ‡	Minimum: 120A IBM LTO-3, IBM LTO-3 WORM: 300A IBM LTO-4, IBM LTO-4 WORM 540A i6.5; 590A i6.6: 595A.01601 i6.7: 596A.GS00301	7404 I/O Blade		i6000 branding started at i2000 i8.
			IBM LTO-1 FC and SCSI	5AU1	
			IBM LTO-2 FC and SCSI	93T0	
			IBM LTO-3 (2G and 4G	93G0, i8 requies 93G6	
			IBM LTO-3 WORM		
			IBM LTO-4 4G	94D4, i8 requires A232	
			IBM LTO-4 WORM		
			HP LTO-3 2G	L67Z	
			HP LTO-3 4G	M69Z	
			HP LTO-3 WORM		
			HP LTO-4 4G	H57Z	
			HP LTO-4 WORM		
			HP LTO-5 FC	I24Z, i8 requires TBD Only supported in i6.7, i8	
			Quantum DLT-S4	V42	
			Quantum SDLT 320 SCSI	V94	
			Quantum SDLT 600 FC	V53	
	Scalar i40 / i80 i1 includes HP HH LTO-5 SAS	101G.GS005	HP LTO-4 FH SAS	A54Z	
			HP LTO-4 FH 4GB FC	H57Z	
			HP LTO-4 HH SAS	U517	
			HP LTO-4 HH FC	V51Z	
	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 **NOTE**: 2.10.0013 firmware not to be used.
			IBM LTO-2		
			IBM LTO-3		
			AIT-2		

‡ 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 4.0.1 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	Scalar 10000	Minimum: 3.00.0017	IBM LTO-2		Must use SDLC/DAS, SDLC/SCSI Target Mode or Native SCSI
			IBM 3590B1A		
			AIT-1		
		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	See library firmware requirement	
			IBM LTO-3 WORM	See library firmware requirement	
			AIT-2		
			AIT-2 WORM		
			IBM 3592		
	PX500	Minimum: 001A	HP LTO-3		Not including WORM
	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, 3, HP LTO-1, 2, 3, 4, IBM LTO-1, 2, 3, 4		
Dell	PV136T	Minimum: 3.11	IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
HP	ESL E Series	Minimum: 4.10 Recently tested: 7.50	HP LTO-3	Recently tested: L68W	
			HP LTO-3 WORM		
			HP LTO-4		
			HP LTO-4 WORM		
			HP LTO-5	Recently tested: I25W	
			HP LTO-5 WORM		
	EML E-Series	Minimum: 1070 Recently tested: 1395	HP LTO-3		
			HP LTO-4		
			LTO-4 WORM		
			HP LTO-5	Recently tested: I25S	
	MSL 6000	Minimum: 5.07	HP LTO-2		
			HP LTO-3	Recently tested: L67W	
			HP LTO-3 WORM		
			HP LTO-4		

‡ 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 4.0.1 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
HP	MSL G3 Series (2024/4048/8096)	Minimum 2024: 0370 (3.70) Minimum 4048: 0600 (6.00) Recently tested: 7.20 Minimum 8096: 0850 (8.50)	HP LTO-2		
			HP LTO-3		
			HP LTO-3 WORM		
			HP LTO-4		
			HP LTO-4 WORM		
			HP LTO-5		
IBM	TS3500	Minimum: 4680	IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
			IBM 3592 (J1A and E05)		
			IBM TS1120 (E05)		Same as IBM3592 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	Minimum: unknown Recently Tested: 2000	LTO-3	Vendor supported: 93G0	See Bulletin 46 Library firmware is known as BlueScale 11.
			LTO-4	Recently tested: 97F9	
Oracle (Sun / StorageTek) SCSI/FC Libraries	L180/L700/L1400	Minimum: 3.18.02	T9840C		
			T9840D		
			T10000A		See Note 2
			T10000B		See Note 2
			HP LTO-3		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL3000	Minimum: 0235	T9840C		
			T9840D		
			T10000A	Minimum: 1.40	See Note 2
			T10000B	Minimum: 1.40	See Note 2
			HP LTO-3		
			HP LTO-4		
			HP LTO-5		
			IBM LTO-3		
			IBM LTO-4		

Note 2: When using T10000 drives, the STK library parameter "Fastload" must be set to "OFF".

StorNext 4.0.1 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
Oracle (Sun / StorageTek) SCSI/FC Libraries	SL500	Minimum: 1373	HP LTO-3		
			HP LTO-4		
			HP LTO-5		
			IBM LTO-3		
			IBM LTO-4		
	9740	Minimum: 2000	Sun/STK 9840		Obsolete
Oracle (Sun / StorageTek) ACSLS 7.3 ACSLS 7.3.1 Libraries See Note 1	L180/L700/L1400	Minimum: 3.18.02	T9840C		
			T9840D		
			T10000A	Minimum: 1.40	See Note 2
			T10000B	Minimum: 1.40	See Note 2
			HP LTO-3	Recently tested: L6CS	
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL3000	Minimum: 0235	T9840C		
			T9840D		
			T10000A	Minimum: 1.40	See Note 2
			T10000B	Minimum: 1.40 Recently tested: 1.44.210	See Note 2
			HP LTO-3		
			HP LTO-4		
			HP LTO-5	Recently tested: I2DS	Requires ACSLS 7.3.1
			IBM LTO-3		
			IBM LTO-4		
	SL500	Minimum: 1373	HP LTO-3		
			HP LTO-4		
			HP LTO-5		Requires ACSLS 7.3.1
			IBM LTO-3		
			IBM LTO-4		
	SL8500	Minimum: 4.14 Recently Tested: 4.70	T9840C		
			T9840D		
			T10000A	Minimum: 1.40	See Note 2
			T10000B	Minimum: 1.40 Recently tested: 1.44	See Note 2
			HP LTO-3		
			HP LTO-4		
			HP LTO-5		Requires ACSLS 7.3.1
			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: When using T10000 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 4.0.1 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 StorNext Upgrade Paths

In general, sites running the following StorNext versions may upgrade directly to StorNext 4.0.1 or 4.0.1.1, 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 4.0.1 or 4.0.1.1:

- StorNext 3.1.2
- StorNext 3.1.3
- StorNext 3.1.4
- StorNext 3.1.4.1
- StorNext 3.5
- StorNext 3.5.1
- StorNext 3.5.1.1
- StorNext 3.5.2
- StorNext 3.5.2.1
- StorNext 4.0
- StorNext 4.0.1 (upgrading to 4.0.1.1)

All other versions of StorNext require additional steps to upgrade to StorNext 4.0.1 or 4.0.1.1.

Note: HA customers on a StorNext release earlier than 3.5 must first upgrade to StorNext 3.5 before upgrading to StorNext 4.0.1 or 4.0.1.1.

Client Interoperability for StorNext 4.0.1

Table 4 indicates previous versions of StorNext SAN clients on certain platforms which can interoperate with a StorNext 4.0.1 metadata controller without upgrading the SAN Client.

Note: The table shows only client platforms for which support has been dropped in StorNext 4.0.1. For other platforms, it is expected that down-revved clients will be updated to StorNext 4.0.1.

Table 4 StorNext Client Interoperability

StorNext 4.0.1/4.0.1.1 Client Interoperability	
StorNext SAN Client Version	**Platform**
StorNext 3.1.2 StorNext 3.1.3 StorNext 3.1.4 StorNext 3.1.4.1	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 StorNext 3.5.1.1 StorNext 3.5.2 StorNext 3.5.2.1	AIX 5.3 HPUX 11iv2 SGI IRIX 6.5.30 (not supported on SN 3.5) SLES10 Itanium (not supported on SN 3.5) SLES10 32-bit (not supported on SN 3.5)
StorNext 4.0	No platforms were dropped between StorNext 4.0 and this release. Upgrading client is required. Upgrading from 4.0.1 to 4.0.1.1 is supported.

Configuration Requirements

Before installing StorNext 4.0.1, 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 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.

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 4.0.1, the following hardware requirements must be met:

- StorNext File System and Storage Manager Requirements on page 24
- StorNext Client Software Requirements on page 25

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*	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 4.0.1, 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 4.0.1 are presented in Table 6.

Table 6 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`).
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.

Virtual Machine Support

StorNext supports VMware virtual machines on SAN and LAN clients. The following table shows the configurations which have been tested and are currently supported by Quantum:

Table 7 Validated VMware Configurations

SAN Clients in VMware Virtual Machines **(StorNext SAN Client with SAN access to shared storage pool)**	**LAN Clients Accessing the Shared Storage Pool from VMware Virtual Machines** **(StorNext LAN Client with LAN access to a StorNext SAN client configured as a Gateway)**
SLES10 x86	SLES10 x86
SLES10 x86-64	SLES10 x86-64
Windows 2003 x86	

Note: The VMware host runs VMware ESX Server.

Frequently Asked Questions About VMware Support

Is StorNext certified with VMware?

Quantum has self-validated StorNext running on VMware virtual machine products and will provide full customer support for StorNext in the validated configurations shown in the table.

Why do you support only one SAN client in any given VMware ESX Server?

This is a VMware limitation. VMware currently supports LUN sharing (a pre-requisite for multiple SAN clients on one ESX Server) only for Microsoft Cluster Services, not for distributed file systems such as StorNext.

Can I install an MDC in a virtual machine?

Quantum currently does not support installing a metadata controller in a virtual machine.

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 the locations below.

Release notes for earlier 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

Resolved Issues

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

- StorNext File System Resolved Issues on page 28
- StorNext Storage Manager Resolved Issues on page 30
- StorNext GUI and Installation Resolved Issues on page 32
- StorNext HA, Replication and Other Resolved Issues on page 33

StorNext File System Resolved Issues

Table 8 lists resolved issues that are specific to StorNext File System.

Table 8 StorNext File System Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	31345	923816, 1029238, 1127766, 1107372	Corrected an ASSERT error that occurred when defragmenting a StorNext file system.
	31432	1100192	Resolved a condition in which StorNext appeared to misinterpret Linux error codes in the buffered I/O path.

Operating System	CR Number	SR Number	Description
Linux	31537	1146668, 1146198, 1168692	Corrected a condition which resulted in login failure and the following error message after entering the correct StorNext GUI login password on the primary MDC server in an HA configuration: `"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.
Windows	31249	n/a	Corrected a condition which caused an FSM panic when running the ExtApiSetNTSecurity api on a file system which doesn't have WindowsSecurity enabled.
	31536	1075052	Corrected a condition which caused a reboot after performing file system expansion.
	31555	1137250 1149524 1145636	Resolved a condition which prevented renaming files and directories on a directory-mounted file system in Window clients.
	31830	n/a	Corrected a condition on Windows Vista systems which resulted in an access denied error when saving or editing files within Microsoft Office 2007 on a client interacting with a StorNext file system.
	31969	1174018	Resolved a condition which prevented the file and folder size from displaying.
All	31255	n/a	Resolved a condition in which deleting a file or directory prevented files from being deleted due to a NOSPACE condition.
	31325	1102024, 1109122	FSM now generates an error condition instead of a PANIC when I/Os are attempted on invalid stripe groups.
	31339	n/a	Corrected a condition in which a stub file test incorrectly retrieved the whole file when a 1024 byte block was read from disk.
	31341	1119366	Corrected a condition which caused data corruption on systems with large mixed data/metadata stripe groups. For more information about this condition, see StorNext Product Alert 35, available here: http://downloads.quantum.com/support_bulletins/6-00960-89_SN_ProdAlert_35.pdf.
	31343	1119366	Resolved a condition which caused an endless loop when running cvfsck.
	31365	1078890	Corrected a condition which caused a core dump in a library containing several very large directories.

Operating System	CR Number	SR Number	Description
All	31540	1078134	Corrected a condition which occurred after an FSM with a writing application was disconnected from the network. The application was blocked for a short interval after a new FSM started.
	31541	1050240	Resolved a condition which prevented clients from reconnecting after failover.
	31549	1147310	Running FsmGetClntQos now returns the maximum number of clients instead of the current number of clients.
	31570	n/a	Resolved a condition in which incorrect results or errors were seen in directories or the inode free-list.

StorNext Storage Manager Resolved Issues

Table 9 lists resolved issues that are specific to StorNext Storage Manager.

Table 9 StorNext Storage Manager Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	31356	836242, 1076354	Resolved a condition in which fsmedcopy failed when new tape media became full.
	31355	1107366, 1130960, 1132012	Information has been added to the man pages and the StorNext User's Guide describing how to clean up the .gz metadump file after a failed backup.
	31703	n/a	Corrected a condition in which putting a forward slash (/) at the end of a directory recovery path causes failure.
	31704	n/a	StorNext now supports removing an instance of file from a directory and then putting a new instance in the same location.
All	31276	1113942	Resolved an issue in which using large values for POL_STORE_RESP_ABORT_TIME caused integer overflow.
	31336	1100934	Corrected a condition which caused MSM commands issued through SNAPI to fail due to RPC errors.

Operating System	CR Number	SR Number	Description
All	31337	981350, 1045554	Corrected a condition which caused the ArcDisp process to terminate with SIGSEGV.
	31317	1033736, 1061364	Corrected processing in which applications could create a working file and then move that file over the top of the original file.
	31348	1127100	Corrected a condition in which a tape drive with no serial number was improperly mapped.
	31349	1113942	Resolved a condition in which large values for POL_STORE_RESP_ABORT_TIME caused integer overflow.
	31350	1119366	Corrected an issue which caused fsrmcopy -c ... to fail if the SDISK file was missing.
	31351	1113942	Resolved a condition which prevented retrieving files from either copy stored on SDISK (2 copies)
	31352	1081602, 1098924	Resolved an incorrect interpretation of a drive tape alert which caused drives to go offline.
	31353	n/a	A RAS ticket is now generated when an SDisk is full.
	31354	1041806 1051478 954296 1060196	Corrected a condition in which files marked as "TRUNC_IMMED" remained as candidates rather than being truncated immediately.

	31395	n/a	Two new arguments have been added to the fsrecover command which allow you to do the following:
			• Recover a directory as it appeared at a specific point in time.
			• Recover the files to a new name space (location) so that they are 'new' files.
			For example:
			`fsrecover filename... [-p] [-t starttime [endtime]]`
			`fsrecover dirname... -d [-p] [-r] [-a [-t dirtime]]`
			`fsrecover [RM_time::]filepathname... -u [-v]`
			`fsrecover dirpathname... -u -d [-r][-v][-a[-t dirtime[-n dest_dir]]]`

	31453	n/a	Resolved a condition in which a request to store files did not complete successfully due to fs_resource failure.

StorNext GUI and Installation Resolved Issues

Table 10 lists resolved issues that are specific to the StorNext GUI or the installation process.

Table 10 StorNext GUI and Installation Resolved Issues

Operating System	CR Number	SR Number	Description
All	31315	1129698	Corrected a condition which caused the StorNext GUI to time out when adding a large ACSLS library.
	31469	n/a	Corrected a condition in which the StorNext GUI did not recognize an IP addresses for converting to HA page until the next restart.
	31473	n/a	The StorNext GUI now safeguards against users applying policies on the blockpool directory.
	31476	n/a	Resolved a condition which produced a validation error immediately after secondary HA conversion.
	31516	n/a	Corrected an issue which caused page navigation to fails if a table on the destination page contained special characters.
	31517	n/a	Resolved a condition in which the StorNext GUI used an embedded application.properties file instead of the installed WEB-INF version.
	31521	n/a	Corrected a condition which allowed converting to HA to pair with an already configured system in a different HA pairing.
	31527	n/a	Corrected an issue which caused spurious errors on primary during HA conversion.
	31539	n/a	The StorNext GUI now displays the hostname correctly when the hostname contains hyphens.
	31605	n/a	Corrected a condition which prevented RAS tickets from properly displaying the most recent tickets in spite of the filter settings.
	31675	1155024	When viewing files on a tape (Tools > File and Directory Actions > View File Info > Browse), the list of files is no longer limited to 100 rows of files.
	31958	n/a	Database queries for media have been improved for better performance.

StorNext HA, Replication and Other Resolved Issues

Table 10 lists resolved issues that are specific to StorNext HA, replication/deduplication and other features or processes.

Table 11 StorNext HA, Replication and Other Resolved Issues

Operating System	CR Number	SR Number	Description
All	31247	n/a	Resolved a condition which caused RAS messages to be lost after hardware failure.
	31357	n/a	Corrected a condition in which StorNext installation failed on an HA system if the HA shared system is ordinal 10 or higher.
	31640	1150350	Replication performance has been improved for large namespaces.
	31646	n/a	Corrected a condition in which Snpolicy could not reliably identify a Storage Manager relation point during the replication process.
	31767	n/a	Resolved a condition in which directory fragmentation caused large-scale replication policies to slow down.
	31721	n/a	Corrected a condition which resulted in some files lacking data blocks after replication.

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 12 lists known issues that are specific to StorNext File System.

Table 12 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)
Macintosh	31669	1160770	Renaming files between policies is not currently supported. Files cannot be moved into and out of policies, nor between different policies. This restriction may cause problems when dragging files inside a replication policy to the trash can, and also when using some GUI-based applications. For example, it has been reported that programs such as TextEdit on a Macintosh cannot be used to create files in a directory that has a replication policy.	This issue may be addressed in a future release. The workaround for application users who are trying to do a "safe save" and move the old file out of the way during the save process is to save their file somewhere else and then move it under the replication point.

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	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`

	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.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	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.
	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.
	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.)
	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.
	31964	n/a	If an MDC running Storage Manager mounts a file system being served by another MDC pair, Healthcheck will generate a RAS message containing the following text: `Internal Software Error: an unhandled software error has occurred. ERROR: mounted CVFS file system name(/stornext/ win_snfs1) not found in list of all CVFS file system names.`	This issue will be addressed in a future StorNext release. In this case the error message is falsely generated and can be safely ignored. Alternatively, the issue can be avoided by not mounting a StorNext file system as a client on the MDC.

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.
	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.
	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.

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.
	31981	n/a	After viewing the 'Properties' information of a file in a StorNext File System, 'Accessed' information is updated. (Other files in NTFS don't change the accessed time when 'Properties' information is viewed.)	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	28573	1111662	If the permissions on an SMB-shared folder do not allow READ access to the SYSTEM user, the folder will become unshared following system reboot.	This issue will be addressed in a future StorNext release. The workaround is to adjust the permissions on the folder to allow READ access to the SYSTEM user. In a future release, additional logging will be added to make the source of the problem more clear.
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.
	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.
	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.
	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

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

StorNext Storage Manager Known Issues

Table 13 lists known issues that are specific to StorNext Storage Manager.

Table 13 StorNext Storage Manager Known Issues

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.

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

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)
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.

	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.000474F892EBB65E000E0000000000000000000292BF4".`
			`Error opening file "/stornext/source/__CVFS_Handle.000474F892EBB65E000E0000000000000000000292BF4".`
			`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	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.
	31615	n/a	Under certain conditions, files flagged for truncation remain on the truncation candidates list indefinitely.	This issue will be addressed in a future StorNext release.

StorNext GUI Known Issues

Table 14 lists known issues that are specific to the StorNext GUI process.

Table 14 StorNext GUI Known Issues

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

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

Linux	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 running this command:
				`. /usr/adic/.profile`
				6 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`

	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.

	31599	n/a	After performing an upgrade, the StorNext GUI may not display correctly due to the web browser reading old cached data.	This issue may be addressed in a future StorNext release. The workaround is to clear your browser's cache and remove cookies, and then restart StorNext.
	31722	n/a	You may receive the following error message when attempting to scan the secondary system in an HA configuration: `"Host cannot be used as a secondary system. Error getting system descriptor. Unable to get system id for texas < org.apache.axis2.AxiaFault: Transport error: 4.0.1 Error Authorization Required."`	This error message will be rewritten to provide greater clarity. If you receive this message, the primary MDC in the HA configuration can obtain the secondary system's system ID, although the message suggests otherwise.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	31724	n/a	If converting to an HA secondary fails, you may have to log in to the secondary system's GUI and view the web server log files to find out why the conversion failed.	This issue will be addressed in a future StorNext release. In the meantime, the workaround is to determine the cause of the failure by logging in to the secondary system and then viewing the log files by using the StorNext GUI's **Reports > Logs** feature.
All	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.
	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.
	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.
	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.
	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.

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 4.0.1, 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`
	31723	n/a	When using Internet Explorer versions 7 or 8, you may experience slow performance during some tasks, particularly when creating a new Storage Manager storage policy.	This issue will be addressed in a future StorNext release. The workaround is to us a browser other than IE7 or IE8 when creating new storage policies, or if you experience sluggish performance.
	31728	n/a	When viewing reports and other screens which have navigation controls to view the next and previous pages, clicking the double-arrow icon may not work as designed. (Clicking the double-arrow should increment the currently displayed page by ten.)	This issue will be addressed in a future StorNext release.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	31730	n/a	On the screen **Tools > File and Directory Actions > View File Info**, the generated list of files appears to be sortable but is not. (The default sort order is according to file name.)	This issue will be addressed in a future StorNext release.
	31780	n/a	The StorNext GUI auto-detects the drive type and adds the supported media classes to the archive, but it fails to add the WORM media class for new SCSI and ACSLS archives, resulting in any WORM medium brought in as 'Unknown'.	This issue will be addressed in a future StorNext release. The workaround is to follow the steps below to manually add the WORM media class to the archive before importing WORM media. **1** Add import media class for the WORM media type to the archive using this command: `vsarchiveconfig –u scsi –n <archive name> –i LTO~F0_LTO_ADDBLANK,LTOW~F0_LTOW_ADDBLANK` **2** If WORM media have been added as 'Unknown,' contact Quantum Support for their assistance cleaning up the database. This could happen if WORM media existed in the archive during the initial configuration of the library. **3** Re-import the WORM media using the StorNext GUI or the command `vsenter/vsaudit`.
	31983	1166844	In some circumstances, the StorNext Storage Manager Media Actions page may not display all available media in the system.	This issue may be addressed in a future StorNext release. The workaround is to use the command line interface instead of the StorNext GUI to perform the appropriate media action. For example, to move a piece of media to another archive, use the vsmove CLI command: `vsmove –a <archivename> <mediaid>` (For more information about using the appropriate CLI commands, refer to the man pages.) If this workaround does not work, contact Quantum support for further assistance.

StorNext Installation Known Issues

Table 15 lists known issues that are specific to the StorNext installation process.

Table 15 StorNext Installation Known Issues

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 and Replication Known Issues

Table 16 lists known issues that are specific to StorNext HA systems and the replication feature.

Table 16 StorNext HA Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	31662 and 31663	n/a	When converting to HA, mounting a shared file system on the secondary system failed. A related issue may cause the primary MDC to fail after the secondary MDC successfully converts.	This issue will be addressed in a future StorNext release. The workaround is to try converting again. If that fails, try manually mounting the file system.
	31959	n/a	After converting to high availability, the blockpool status goes to "Verify pending" state and an error message appears.	This issue will be addressed in a future StorNext release. The workaround is to stop and then start blockpool services.
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.

Operating System	Change Request Number	Service Request Number	Description	Workaround
All	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.
	32046	n/a	Replication using large policies of several million files can require the ophanglimit config option on the replication target to be increased from the default value. A value of up to 10 minutes may be required for policies replicating 10 million files or more.	This issue will be addressed in a future StorNext release. The workaround is to increase the ophanglimit on the replication target.

Operating Guidelines and Limitations

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

Table 17 lists legacy items from previous releases that still pertain to StorNext 4.0.1.

Table 17 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, Windows Vista or Windows 7, 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 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**.

Operating System / Affected Component	Description
Windows	(Disabling the Recycle Bin, Continued) For Windows Server 2008, Windows Vista and Windows 7 systems, you must disable the Recycle Bin on C: and the File system name: **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**. 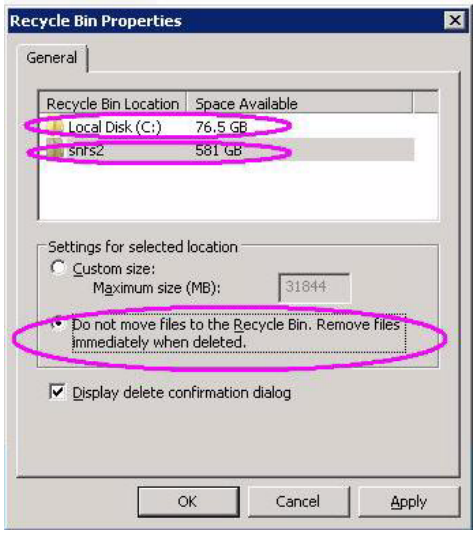

Operating System / Affected Component	Description
All	Be aware of the following limitations regarding file systems and stripe groups: • The maximum number of disks per file system is 512 • The maximum number of disks per data stripe group is 128 • The maximum number of stripe groups per file system is 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 18 lists new guidelines and limitations, many of which are specific to StorNext 4.0.1.

Table 18 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*
6-67041-01	*StorNext File System Quick Reference Guide*
6-67042-01	*StorNext Storage Manager Quick Reference Guide*
6-66851-02	*StorNext HA Quick Reference Guide*
6-66852-02	*StorNext Replication 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:

Quantum Technical Assistance Center in the USA:	+1 800-284-5101
For additional contact information:	www.quantum.com/support
To open a Service Request:	www.quantum.com/osr

For the most updated information on Quantum Global Services, please visit: www.quantum.com/support