

StorNext 4.1.2 Release Notes

Product	StorNext® 4.1.2
Date	July 2011

Contents

Purpose of this Release	3
Changes and Considerations	3
Operating System Requirements	3
Supported Libraries and Tape Drives	11
Supported StorNext Upgrade Paths	17
Older Client Interoperability for StorNext 4.1.2	9
Compatibility With Other StorNext Components and Features	18
Configuration Requirements	19

© 2011 Quantum Corporation. All rights reserved.

6-00431-36 Rev B, July 2011

Quantum, the Quantum logo, DLT, DLTtape, the DLTtape logo, SuperLoader, Scalar, StorNext, and DXi are registered trademarks of Quantum Corporation, registered in the U.S. and other countries. All other trademarks are the property of their respective companies. Specifications are subject to change without notice.

StorNext utilizes the following components which are copyrighted by their respective entities:

ACSAPI, copyright © Storage Technology Corporation

Java, copyright Oracle Corporation

LibICE, LibSM, LibXau, LibXdmcp, LibXext, LibXi copyright The Open Group

LibX11 copyright The Open Group, MIT, Silicon Graphics, and the Regents of the University of California, and copyright (C) 1994-2002 The XFree86 Project, Inc. All Rights Reserved. And copyright (c) 1996 NVIDIA, Corp. NVIDIA design patents pending in the U.S. and foreign countries.

Libxml2 and LibXdmcp, copyright MIT

Linter, copyright © Relex Software Corporation

Ncurses, copyright © 1997-2009,2010 by Thomas E. Dickey <dickey@invisible-island.net>. All Rights Reserved.

TCL/TK, copyright © Sun Microsystems and the Regents of the University of California

vixie-cron: copyright Internet Systems Consortium (ISC)

Wxp-tdi.h, copyright © Microsoft Corporation

Zlib, copyright © 1995-2010 Jean-loup Gailly and Mark Adler without notice.



Hardware Requirements	22
Supported System Components	25
Configuring VMWare for StorNext	21
Previous Versions of Release Notes	27
Resolved Issues	27
Known Issues	31
Operating Guidelines and Limitations	41
Documentation	53
Contacting Quantum	53

Purpose of this Release

StorNext 4.1.2 is a maintenance release that includes a small number of important bug fixes. It also includes a set of bug fixes which were made in Limited Customer Releases (LCRs).

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

No Cryptographic Changes

There is no change to cryptographic functionality in StorNext release 4.1.2.

Changes and Considerations

This section contains important things pertaining to this release you should know.

Upgrading AIX, HP-UX and Solaris NFS Servers

When upgrading AIX, HP-UX, or Solaris NFS servers to this release of StorNext, NFS clients accessing these servers must first unmount StorNext file systems before proceeding with the upgrade.

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: StorNext supports all vendor-supplied security patches within the updates and service packs listed in [Table 1](#).

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 4.1.2 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 4.1.2 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:

- * 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 4.1.2 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 / SWAPI	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	✓	✓	✓	✓	✓	✓	✓
	2.6.18-238.EL (Update 6)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓

‡ 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 with `SELINUX=enforcing` or `SELINUX=permissive` to `SELINUX=disabled`. Refer to Red Hat 5 documentation for more information.

StorNext 4.1.2 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 / SWAPI	Distributed Data Mover	Replication / Dedup Server
SLES 10 ***	2.6.16-46-0.12 (SP1) **	x86 32-bit		✓		✓			
	2.6.16.60-0.27 (SP2) **	x86 32-bit		✓		✓			
	2.6.16.60-0.54.5 (SP3)	x86 32-bit		✓		✓			
See Note 2	2.6.16-46-0.12 (SP1) **	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.16.60-0.27 (SP2) **	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
	2.6.16.60-0.54.5 (SP3)	x86 64-bit	✓	✓	✓	✓	✓	✓	✓
SLES 11 *****	2.6.27.19-5	x86 64-bit		✓		✓			
See Note 2	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	Equivalent supported RHEL5 releases	x86 64-bit		✓		✓			
Scientific Linux***	Equivalent supported RHEL5 releases	x86 64-bit		✓		✓			
Oracle Linux***	Equivalent supported RHEL5 releases	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.

**# 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 4.1.2, StorNext 3.5.x SGI IRIX clients may be used with a StorNext 4.1 or 4.2 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 SDD Patch
 - 120998-01 (or greater) SD Headers Patch
-

[Table 2](#) indicates StorNext MDC to Apple Xsan client interoperability.

Table 2 StorNext MDC to Xsan Client Interoperability

StorNext 4.1.2 Client Interoperability			
Apple Xsan Version	Platform	Compatible	Notes
1.4	x86-32 bit	No	See notes 1 and 2
1.4.1	x86-32 bit	No	See notes 1 and 2
1.4.2	x86-32 bit	No	See notes 1 and 2
2.0	x86-32 bit	Yes	See notes 1 and 2
2.1	x86-32 bit	Yes	See notes 1 and 2
2.1.1	x86 32-bit	Yes	See notes 1 and 2
2.2	x86 32-bit	Yes	
	x86 64-bit	Yes	
2.2.1	x86 32-bit	Yes	
	x86 64-bit	Yes	

¹ Apple Leopard machines run with 32-bit kernel, 64-bit user

² Releases earlier than MacOS X 10.5.5 may have limited Windows Access Control Lists (ACL) functionality.

StorNext support for Apple Xsan is not coupled with StorNext releases. The information in [Table 2](#) is current as of the date these release notes were written.

For compatibility information on Apple Xsan, see <http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx>. (Click the Documentation tab and navigate to the "Compatibility Guide" heading.)

Older Client Interoperability for StorNext 4.1.2

[Table 3](#) indicates previous versions of StorNext SAN clients on certain platforms which can interoperate with a StorNext 4.1.2 metadata controller without upgrading the SAN Client.

Note: The table shows only client platforms for which support has been dropped in StorNext 4.1.2. For other platforms, it is expected that down-revved clients will be updated to StorNext 4.1.2.

Table 3 StorNext Client Interoperability

StorNext 4.1.2 Client Interoperability	
StorNext SAN Client Version	Platform
StorNext 3.1.1 and older	Back-revision clients running these StorNext versions are not supported, even during the upgrade process. Clients must be upgraded with MDCs to achieve a compatible back-rev client version.
StorNext 3.1.2 StorNext 3.1.3 StorNext 3.1.4 StorNext 3.1.4.1 StorNext 3.1.5	Solaris 9 (sparc only) AIX 5.3 HPUX 11iv2 SGI IRIX 6.5.30 RHEL4 Itanium SLES10 Itanium SLES10 32-bit Quantum recommends that other clients be upgraded along with the MDC.
StorNext 3.5 StorNext 3.5.1 StorNext 3.5.1.1 StorNext 3.5.2 StorNext 3.5.2.1 StorNext 3.5.3	AIX 5.3 HPUX 11iv2 SGI IRIX 6.5.30 SLES10 Itanium SLES10 32-bit Quantum recommends that other clients be upgraded along with the MDC.
StorNext 4.0 StorNext 4.0.1 StorNext 4.0.1.1 StorNext 4.1 StorNext 4.1.1	Quantum recommends that other clients be upgraded along with the MDC.

Virtual Machine Support

StorNext supports SAN client and DLC clients running within VMware virtual machines on the Windows and Linux operating systems listed in [Table 1](#) on page 4.

Refer to StorNext Product Bulletin 69 for additional details about VMware support.

Supported Libraries and Tape Drives

Libraries and tape drives supported for use with StorNext 4.1.2 are presented in [Table 4](#). Where applicable, minimum firmware levels for libraries are provided.

Table 4 StorNext Supported Libraries and Tape Drives

StorNext 4.1.2 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: Minimum 586G.GS001	IBM LTO-2	i5.1: 8571 i6: A4N0	
			IBM LTO-3	i5.1: 93G0 i6: 93GE	
			IBM LTO-3 WORM		
			IBM LTO-4	i5.1: 94D4 i6: A239	
			IBM LTO-4 WORM		
			IBM LTO-5	i6: A5M0	
			IBM LTO-5 FC	i6: A5M0	
			HP LTO-4 FC	i5.1: H46Z i6: H58Z	
			HP LTO-4 SAS	i5.1: A45Z i6: A55Z	
			HP LTO-5 FC	i6: I39Z Support starts in Quantum i6.x	
	HP LTO-5 SAS	i6: X38Z Support starts in Quantum i6.x			
	Scalar i2000 / i6000 (i6000 branding started at i2k i8)†	Minimum: 120A IBM LTO-3, IBM LTO-3 WORM Minimum: 300A.xxx IBM LTO-4, IBM LTO-4 WORM Minimum 540A.xxx i6.5: 590A i6.6: 595A.01601 i6.7: 596A.GS00301 i8.0(.1): 600A.GS23201 i8.1: 605A.GS07401	IBM LTO-1 FC and SCSI IBM LTO-2 FC and SCSI IBM LTO-3 (2G and 4G) IBM LTO-3 WORM IBM LTO-4 4G IBM LTO-4 WORM IBM LTO-5 HP LTO-3 2G HP LTO-3 4G HP LTO-3 WORM HP LTO-4 4G HP LTO-4 WORM HP LTO-5 FC Quantum DLT-S4 Quantum SDLT 320 SCSI Quantum SDLT 600 FC	5AU1	i6000 branding started at i2000 i8.
				i6.x: 93T0 i8.x: A4N0	
				i6.x: 93G0 i8.x: 93GM	
				i6.x: 94D4 i8: A239	
				i8.1: A5M0 requires i8.1 or later	
				L67Z	
M69Z					
H58Z					
i6.x: I24Z i8.0: I39Z i8.1: I3AZ Requires i6.7 or later					
V42					
V94					
V53					

StorNext 4.1.2 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 i40 / i80	Base (SP5): 105G.GS001 I1 (SP6): 111G.GS003 I2: 120G.GS003	HP LTO-4 FH SAS	A55Z	
			HP LTO-4 FH 4GB FC	H58Z	
			HP LTO-4 HH SAS	U52Z	
			HP LTO-4 HH FC	V52Z	
			HP LTO-5 HH SAS	Z38Z, requires i1	
			HP LTO-5 HH FC	Y23Z, requires i2	
	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		
	Scalar 1000	Minimum: 3.00.0017	IBM LTO-2		Must use SDLC/DAS, SDLC/SCSI Target Mode or Native SCSI
			IBM 3590B1A		
			AIT-1		
	Scalar 10000	Minimum: 110A.00001	IBM LTO-1		Must use SDLC/DAS, SDLC/SCSI Target Mode or Native SCSI
			IBM LTO-2		
			IBM LTO-3	See library firmware requirement	
			IBM LTO-4	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			

‡ 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 4.1.2 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
Dell	PV136T	Minimum: 3.11	IBM LTO-2		
			IBM LTO-3		
			IBM LTO-4		
	PowerVault ML6000 (6010, 6020, 6030)	Minimum: 585G.GS003	IBM LTO3FH SCSI IBM LTO3FH FC	Minimum: 93G6	
			IBM LTO4FH SAS IBM LTO4FH FC	Minimum: A232	
			IBM LTO5FH SAS IBM LTO5FH FC	Minimum: A420	
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	
	ESL G3	Recently tested: 620H	HP LTO-4 4G	Recently tested: H63W	
			HP LTO-4 WORM		
			HP LTO-5 FC	Recently tested: I3FW	
	MSL 6000	Minimum: 5.07	HP LTO-2		
			HP LTO-3	Recently tested: L67W	
			HP LTO-3 WORM		
			HP LTO-4		
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			

StorNext 4.1.2 Supported Libraries and Tape Drives (Continued)					
Vendor Library Family	Libraries	Enforced Minimum / Recently Tested Library Firmware Level	Drive Types	Enforced Minimum / Recently Tested Drive Firmware Level	Notes
IBM	TS3500	Minimum: 7422 Recently Tested: A420	IBM LTO-2		
			IBM LTO-3	Minimum 93GE	
			IBM LTO-4	Minimum A239	
			IBM LTO-5	Minimum A6S0	
			IBM 3592 (J1A and E05)		
	IBM TS1120 (E05)		Same as IBM3592 E05		
	TS3310	Minimum: 587G.GS003	IBM LTO-3	Minimum: 93GE	
			IBM LTO-4	Minimum: A239	
			IBM LTO-5	Minimum: A6S0	
Qualstar	XLS	Minimum: 0880	IBM LTO-3		
			IBM LTO-4		
Sony	Petasite CSM-200	Minimum: 6.30	IBM LTO-4 drive (T1600)		
Spectra Logic	T-Series (T50e, T120, T200, T380, T680, T950, and T-Finity)	Minimum: unknown Recently Tested: 11.4.3	LTO-3	Vendor supported: 93G0	See Bulletin 46 Library firmware is known as BlueScale 11. Both L700 emulation and Native mode are supported In L700 emulation mode, LTO-5 drives report as LTO-4, limiting the capacity of the media.
			LTO-4		
			LTO-5	Recently tested: B170	

StorNext 4.1.2 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	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		
			HP LTO-4		
			IBM LTO-3		
			IBM LTO-4		
	SL3000	Minimum: 0235 LTO-5 requires minimum 2.35	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	Recently tested: I2DS	
			IBM LTO-3		
	SL500	Minimum: 1373 LTO-5 requires minimum 1395	HP LTO-3		
			HP LTO-4		
			HP LTO-5	Recently tested: I2DS	
			IBM LTO-3		
			IBM LTO-4		
	9740	Minimum: 2000	Sun/STK 9840		Obsolete

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".

StorNext 4.1.2 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) ACSLS 7.3 ACSLS 7.3.1 ACSLS 8.0.x Libraries See Note 1	L180/L700/ L1400	Minimum: 3.18.02 Recently tested (L700): 3.18	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		
	SL3000	Minimum: 2.35 LTO-5 requires minimum 2.35 Recently tested: 2.35	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 minimum ACSLS 7.3.1
	SL500	Minimum: 1373 LTO-5 requires minimum 1395	IBM LTO-3		
			IBM LTO-4		
			IBM LTO-5		Requires minimum ACSLS 7.3.1
			HP LTO-3		
			HP LTO-4		
	SL8500	Minimum: 4.14 LTO-5 requires minimum 4.73 Recently Tested: 6.02	HP LTO-5	Recently tested: I2DS	Requires minimum ACSLS 7.3.1
			IBM LTO-3		
			IBM LTO-4		
			IBM LTO-5		Requires minimum ACSLS 7.3.1
			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	Recently tested: I2DS	Requires minimum ACSLS 7.3.1			
IBM LTO-3					
IBM LTO-4					
IBM LTO-5		Requires minimum ACSLS 7.3.1			

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”.

Supported StorNext Upgrade Paths

In general, sites running the following StorNext versions may upgrade directly to StorNext 4.1.2, 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 4.1.2:

- StorNext 3.1.2
- StorNext 3.1.3
- StorNext 3.1.4
- StorNext 3.1.4.1
- StorNext 3.1.5
- StorNext 3.5
- StorNext 3.5.1
- StorNext 3.5.1.1
- StorNext 3.5.2
- StorNext 3.5.2.1
- StorNext 3.5.3
- StorNext 4.0
- StorNext 4.0.1
- StorNext 4.0.1.1
- StorNext 4.1
- StorNext 4.1.1

All other versions of StorNext require additional steps to upgrade to StorNext 4.1.2.

Compatibility With Other StorNext Components and Features

This section describes various interactions between this release and StorNext components and features.

Data Replication Compatibility

The following table illustrates compatibility between StorNext releases when using the replication feature. The table differentiates between systems using deduplication and systems not using deduplication.

Note: If a source replication policy uses deduplication, the target policy must also use deduplication.

Table 5 Replication Compatibility Between Releases

Source Release	Target Release							
	4.0 without Dedup.	4.0 with Dedup.	4.0.1 without Dedup.	4.0.1 with Dedup.	4.1 without Dedup.	4.1 with Dedup.	4.1.1/4.1.2 without Dedup.	4.1.1/4.1.2 with Dedup.
4.0 without Dedup.	✓	--	See note below	--	See note below	--	See note below	--
4.0 with Dedup.	--	✓	--	Replication rejected	--	Replication rejected	--	Replication rejected
4.0.1 without Dedup.	Replication rejected	--	✓	--	✓	--	✓	--
4.0.1 with Dedup.	--	Replication rejected	--	✓	--	✓	--	✓
4.1 without Dedup.	Replication rejected	--	✓	--	✓	--	✓	--
4.1 with Dedup.	--	Replication rejected	--	✓	--	✓	--	✓
4.1.1/4.1.2 without Dedup.	Replication rejected	--	✓	--	✓	--	✓	--
4.1.1/4.1.2 with Dedup.	--	Replication rejected	--	✓	--	✓	--	✓

Note: StorNext 4.0 to 4.0.1 and 4.1/4.1.1/4.1.2 without deduplication successfully transfers content, but there may be synchronization issues in situations where content was changed very recently before a scheduled replication. If this occurs, the file may initially appear to be transferred without content, but the content will eventually be transferred in its entirety.

Partial File Retrieval

StorNext Partial File Retrieval (PFR) is a separately available product which enables you to quickly retrieve and utilize segments of large media files— rather than the entire file—based on timecode parameters.

StorNext support for Partial File Retrieval (PFR) is not coupled with StorNext releases, so this compatibility information is not included in [Table 1](#) on page 4.

For compatibility information on PFR see <http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx>. (Click the Documentation tab and navigate to the “Compatibility Guide” heading.)

StorNext Offline Notification

The StorNext Offline Notification feature remains a separate download and uses StorNext 4.1.

Configuration Requirements

Before installing StorNext 4.1.2, 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.

Disk Naming Requirements

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

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: 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

6 Perform the following for each disk manifesting the problem:

diskpart.exe:

```
select disk X
```

```
attribute disk clear readonly
```

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.
```

Configuring VMWare for StorNext

To configure StorNext Distributed LAN Clients in VMWare guests, follow the same procedures you would for a physical system. There are no VMWare-specific requirements or issues.

To configure StorNext SAN clients in VMWare guests, be aware of the following considerations:

- StorNext Data LUNs must be assigned to each StorNext SAN client VM using Raw Device Maps (RDMs) in /Physical Mode/ on a Shared virtual SCSI adapter. Never use /Virtual Mode/ RDMs for StorNext LUNs.

Consult your storage vendor for details on properly configuring the storage for use as VMWare vSphere to use raw LUNs as RDMs.

- On each SAN client, generate a raid-strings file by running the command:

```
cvlabel -R > /usr/cvfs/config/raid-strings
```

Then open /usr/cvfs/config/raid-strings in a text editor and change the third column to JBOD for all storage types. This disables StorNext multi-path handling, which is not needed in a guest. The host will handle multi-pathing.

- Guests running StorNext SAN clients have limited cluster functionality due to the use of RDMs to access storage. In particular, snapshots, vMotion, DRS, and fault tolerance are disabled. If these features are required, then DLC clients should be used.

Hardware Requirements

To successfully install StorNext 4.1.2, the following hardware requirements must be met:

- [StorNext File System and Storage Manager Requirements](#) on page 22
- [StorNext Client Software Requirements](#) on page 23

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 6](#).

Table 6 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	<ul style="list-style-type: none">• For application binaries, log files, and documentation: up to 30GB (depending on system activity)• For support directories: 3 GB per million files stored
5-8**	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 4.1.2, 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.

Note: Additional disk and memory requirements are inclusive, so there is no need to fulfill the requirements for a lesser capacity system on higher capacity systems.

For example, if you have a license for 1 - 10 TB of deduplication data (requiring 6 GB additional RAM and 10 TB available hard disk space), there is no need to *also* fulfill the requirement for 0 - 1 TB of deduplication data (1 GB additional RAM and 1 TB available hard disk space) on top of the 1 - 10 TB data requirements.

Minimum Additional Disk and Memory Required for a Replication License

- 50 MB available hard disk space

Minimum Additional Disk and Memory Required for Systems with a 0 - 1 TB Deduplication Data License

- 1 GB additional RAM
- 1 TB available hard disk space

Minimum Additional Disk and Memory Required for Systems with a 1 - 10 TB Deduplication Data License

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

Minimum Additional Disk and Memory Required for Systems with a 10 - 50 TB Deduplication Data License

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

Minimum Additional Disk and Memory Required for Systems with a 50 - 150 TB Deduplication Data License

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

Minimum Additional Memory on Replication Target Machines

On replication target machines, Quantum recommends adding a minimum of 1GB additional memory per 44 million files replicated.

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.

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.

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 -vvv1`. For example:

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

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

Supported System Components

System components that are supported for use with StorNext 4.1.2 are presented in [Table 7](#).

Table 7 StorNext Supported System Components

Component	Description
Tested Browsers	Internet Explorer 7 or 8 Mozilla Firefox 3 or 4 Other browsers and versions besides the ones tested may work but are not officially supported.

Component	Description
NFS	<p>Version 3</p> <p>An NFS server that exports a StorNext file system with the default export options may not flush data to disk immediately when an NFS client requests it. This could result in loss of data if the NFS server crashes after the client has written data, but before the data has reached the disk.</p> <p>As a workaround, add the <code>no_wdelay</code> option to each line in the <code>/etc/exports</code> file that references a StorNext file system. For example, typical export options would be <code>(rw, sync, no_wdelay, no_subtree_check)</code>.</p> <p>NOTE: Although supported in previous StorNext releases, the <code>subtree_check</code> option (which controls NFS checks on a file handle being within an exported subdirectory of a file system) is no longer supported as of StorNext 4.0. Be sure to specify the <code>no_subtree_check</code> option in the exports file.</p>
LDAP	<p>LDAP (Lightweight Directory Access Protocol) support requires Windows Active Directory.</p>
Mixed-Level Tape Drive Compatibility Within the Same Device Family	<p>LTO-1 media in a library containing LTO-3 or LTO-4 drives are considered for store requests unless they are logically marked as write protected. When LTO-1 media is mounted in an LTO-3 or LTO-4 drive, StorNext marks the media as write protected. Quantum recommends circumventing LTO-1 media for store requests by following this procedure:</p> <ol style="list-style-type: none"> 1 From the SNSM home page, choose Attributes from the Media menu. 2 On the Change Media Attributes window, select the LTO-1 media from the list. 3 Click the Write Protect option. 4 Click Apply to make the change. 5 Repeat the process for each piece of LTO-1 media. <p>NOTES:</p> <ul style="list-style-type: none"> • A similar issue exists for LTO-2 media in a library containing LTO-4 tape drives. • LTO-3 drives can read but not write LTO-1 tapes. • A similar issue exists for LTO-3 media in a library containing LTO-5 drives. • LTO-4 drives can read but not write LTO-2 tapes, and also cannot read LTO-1 tapes at all.

Previous Versions of Release Notes

Previous versions of the StorNext release notes contain additional information specific to earlier StorNext releases. You can find previous release notes at 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 27
- [StorNext Storage Manager Resolved Issues](#) on page 29
- [StorNext GUI Resolved Issues](#) on page 29
- [StorNext Replication, Deduplication and Other Resolved Issues](#) on page 30

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	33781	n/a	Added native support for SpectraLogic T-Series libraries.
Windows	33778	1240540	Corrected a condition which prevented fsmppm from restarting after initialization failure.

Operating System	CR Number	SR Number	Description
All	31157	n/a	Resolved a condition which caused StorNext-truncated files to be incorrectly retrieved due to the stub file size.
	33768	1226008	Resolved a condition in which cvfsck cleared the root directory incorrectly.
	33770	n/a	Corrected a condition which caused FSM panics at startup due to corrupt Windows security descriptors following file system recovery from StorNext backup.
	33772	1235828, 1270982	Resolved a condition which resulted in unnecessary rbtree scans in Alloc_space().
	33780	n/a	Corrected a condition in which a stale value was returned after an ExtRep update or removal operation.
	33784	1267158	Resolved a condition in which cvfsck continued without enough space, causing metadata corruption.
	33785	1268280	Corrected a condition which causes system panics when a process doing I/O to a StorNext file system was killed.
	33787	1269876	Resolved a condition in which NFS clients lost their mount points and access did not return after these servers were rebooted.
	33788	n/a	Corrected a condition in which running cvdbset or cvdb -l caused a system panic.

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
All	33773	1241492, 1258510, 1246554	Corrected a condition which caused a drive to show "IN USE" for an extended period even though no writes were occurring.
	33775	1242518	Fix vsarchiveconfig -u (modify configuration) for multiple ACS numbers
	33776 33777	1249000	Resolved a condition which could cause StorNext commands to fail on or after January 1, 2012.
	33783	n/a	Native support was added for HP ESL G3 libraries.

StorNext GUI Resolved Issues

[Table 11](#) lists resolved issues that are specific to the StorNext graphical user interface.

Table 10 StorNext GUI Resolved Issues

Operating System	CR Number	SR Number	Description
All	33767	1193758, 1255144, 1255348, 1259960, 1265842, 1261440	Performance for the Tools>User Accounts>New feature has been improved.
	33838	1166754	StorNext users are now prompted to confirm twice after changing the login password.

StorNext Replication, Deduplication and Other Resolved Issues

[Table 11](#) lists resolved issues that are specific to StorNext replication, deduplication and other StorNext features or components.

Table 11 StorNext Replication
and Other Resolved Issues

Operating System	CR Number	SR Number	Description
All	33774	1242518, 1259020	Corrected a condition which caused XdiAMTask to hang during MSM startup for an ACSLS configuration.
	33771	1217762	Resolved a condition in which replication cleanup cleared two INHERITED events which could cause StorNext to generate database errors
	33779	1259078	Corrected a condition which prevented data deduplication from occurring without a valid replication license.

Known Issues

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

- [StorNext File System Known Issues](#) on page 31
- [StorNext Storage Manager Known Issues](#) on page 34
- [StorNext GUI Known Issues](#) on page 36
- [StorNext HA and Replication Known Issues](#) on page 39

Note: The issues in [Table 12](#) are the same ones listed in the StorNext 4.1. release notes.

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)
Linux	25978	n/a	Scheduled tasks for "partial backups" and for "rebuild policy" can fail if they overlap.	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	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.)	The workaround is to ensure that your replication workflow does not rely on non-regular file replication.
	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.	To work around this issue, you must find the appropriate network interfaces and reconfigure the dpserver .fsname file and then remount the file system.
	29416	n/a	In an HA configuration, RAS messages are not generated for loss of SAN connectivity by the secondary system.	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.)
	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.	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.
Windows	29483	n/a	After changing fsnameservers in the StorNext GUI, the file system failed to mount and returned a "device not connected" error.	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.	Note: the count of file systems found includes standby file systems.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Windows	30945	n/a	On Windows Systems, creating a Virtual Hard Disk (VHD) on a StorNext file system will cause a system crash.	Until there is a resolution to this problem, all VHDs must be created on NTFS file systems. This includes both direct creation of VHDs through the Windows Disk Manager on Windows 7 and Windows 2008 R2 systems, as well as any Windows-based product that uses VHDs implicitly. This includes, but is not limited to, the Windows Complete PC Backup and Microsoft Virtual PC products. All of these products should continue to function normally with StorNext installed as long as they are not configured to create their VHDs on StorNext.
	32187	1185732	The Windows GUI for StorNext administration (ntcvadmin) may fail to connect to file system servers, including those operating on localhost.	The workaround is to use the command line cvadmin utility instead of the StorNext GUI.
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.	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.
	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.	One workaround is to copy the ISO image to a local file system and then perform the loopback mount on the copy.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	23377	725697 1236258	Errors can result when there are multiple retrieves in progress at the same time for the same multi-segmented file, and then one of the retrieves is canceled when the last segment is being retrieved.	You can avoid this issue by either not requesting multiple retrieves for the same multi-segment file, or by not canceling the subsequent requests.

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	29445	n/a	When adding Distributed Data Mover (DDM) mover hosts (either by the StorNext GUI or by the fsddmconfig command,) the same host can be added multiple times. This can impact performance tuning.	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	<p>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.</p> <p>TSM will use the <code>fs_fmoverc</code> 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 <code>fs_fmoverc</code>.</p>	<p>If TSM has not completed its startup and <code>fs_fmoverc</code> processes persist (as shown by running the <code>tsmup</code> command), use the <code>fsddmconfig</code> command to set the state of the DDM mover machine to disabled.</p> <p>For example, if the hostname of the DDM mover machine is <code>minnesota</code>, run this command on the MDC:</p> <pre>fsddmconfig -u -s d minnesota</pre> <p>Once the DDM mover machine is back up, re-enable it using the StorNext GUI or with this command:</p> <pre>fsddmconfig -u -s e minnesota</pre> <p>Under most circumstances, <code>fs_fmoverc</code> will successfully determine the downed status of a DDM mover machine, and automatically set its status to DISABLED. A RAS ticket is emailed when <code>fs_fmoverc</code> automatically disables a mover, so after receiving such a ticket and rebooting the DDM mover, use the StorNext GUI (or the command <code>fsddmconfig</code> in the second example above) to enable the mover.</p>

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	30842	n/a	<p>In some situations, on HA systems the StorNext GUI exits the config mode (with the redundant server in peerdown mode) without starting Storage Manager.</p> <p>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.</p> <p>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.</p>	<p>Two actions are necessary to correct this issue:</p> <ol style="list-style-type: none"> 1 Run the following CLI command: "DSM_control start". This will produce the following warning message, which can be ignored: "fsmgm is already running. Only one fsmgm 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.
	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.
	29038	n/a	Clicking through StorNext GUI logs pages more than once every few seconds causes error.	This error message can be ignored and requires no response. To remove the error message, click Refresh or select another log or page.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	29413	n/a	Warnings about insecure and secure items are generated when using StorNext with Internet Explorer 8.	These warning can be ignored and require no response. As a workaround, you can configure Internet Explorer to "Display mixed content" under the Security settings. If enabled, the warning message will not be displayed.
	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.AxisFault: Transport error: 4.0.1 Error Authorization Required."	This indicates a failure for the StorNext GUI to communicate with the secondary MDC. Check that the StorNext software on the secondary system has been installed correctly, and then reboot the secondary system to make sure the software is running correctly.
	32553	n/a	The Library Operator Interface (LOI) page and the Media Action page may not display immediately if you have thousands of media requiring attention in the LOI page (for example, media moves, media ejects, etc) at the same time. A general estimate is that the delay could be approximately 10 seconds for every 500 media requiring attention on the LOI page.	This situation is extremely unlikely to occur, and most StorNext users will never experience this delay unless they need to export thousands of media at one time. In that situation, you can avoid this issue by breaking up media action requests into smaller chunks.
All	29557	n/a	The StorNext GUI can hang if it is unable to create a file system. Additional validation is necessary.	If the StorNext GUI makes an API call which causes an FSM panic, the StorNext GUI is taken down as well. To work around this condition, restart the StorNext GUI (service stornext_web start).
	29728	n/a	Due to an error in a third-party component, background pages respond to keyboard input when modal dialog windows are open.	

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	30929	n/a	<p>The StorNext GUI may be inaccessible in a Web browser, with one of the following error messages displayed:</p> <p>Firefox: Unable to connect. Firefox can't establish a connection to the server</p> <p>Internet Explorer: Internet Explorer cannot display the webpage</p>	<p>If you encounter this condition, restart the StorNext GUI on the MDC server by doing the following:</p> <ol style="list-style-type: none"> 1 Open a root UNIX shell window on the MDC. 2 Run the command <code>service stornext_web restart</code> <p>The "service" command will return before the service is ready to be accessed by a browser. Wait a few moments before trying to connect, and then retry if that fails.</p>
	30925	n/a	<p>In systems with archives that have multiple mailboxes available, importing media can fail with the message "No new media found."</p> <p>(This occurs after choosing Storage Destinations > Library > Add Media Mailbox from StorNext's Setup menu.)</p>	<p>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:</p> <ol style="list-style-type: none"> 1 Open up a UNIX root shell on the MDC server. 2 Source the profile by running <code>./usr/adic/.profile</code> 3 Obtain a list of available mailboxes for an archive by running <code>./usr/adic/MSM/bin/mmportinfo <archivename></code> 4 Import media into an archive from a specific mailbox by running <code>./usr/adic/gui/scripts/library.pl add_media --archive=<archivename> --importmethod=mailbox --mailbox=<mailbox></code> <p>Example:</p> <pre>./usr/adic/gui/scripts/library.pl add_media --archive=archive01 --importmethod=mailbox --mailbox=16:LTO:0,0,15,16</pre>
	32626	n/a	<p>Affinities and QoS parameters are not displayed in Configuration > File Systems (Advanced parameters).</p>	

StorNext HA and Replication Known Issues

[Table 15](#) lists known issues that are specific to StorNext HA systems and the replication feature.

Table 15 StorNext HA and Replication Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	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.	The workaround is to manually confirm reachability to the replication source's vIP address on the replication target, and then reconfigure routing, if necessary.
	31959	n/a	After converting to high availability, the blockpool status goes to "Verify pending" state and an error message appears.	The workaround is to stop and then start blockpool services.
All	29023	n/a	Replication quiesce scripts do not synchronize data on any clients that have open files.	
	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: <code>service cvfs restart</code>
	29722	n/a	StorNext API 2.0.1 failed to install on the secondary node of an HA pair.	A workaround is to fail the primary over to the secondary and then install SNAPI.

Operating System	Change Request Number	Service Request Number	Description	Workaround
All	30006	n/a	There is no way to conveniently delete a TSM relation point used for replication.	<p>In the meantime, you can manually delete the relation point by running the command <code>rm -rf /snfs/sn2/tsm/.rep_private</code>, which empties the TSM relation point.</p> <p>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 <code>tsm_dir / .rep_private</code> only after the LAST target policy has been removed from the relation point.</p>
	30817	n/a	When replicating with multiple copies, StorNext can propagate metadata changes into the previous replication copies.	If file attributes on replicated files have been changed since the replication date of the requested copy, verify the retrieved files have the expected file attributes.
	31825	n/a	When running <code>snpolicy -rundelete=mntpnt -wait</code> , <code>snpolicy</code> can run for a long time before returning. While this operation is running other <code>snpolicy</code> commands can block.	The workaround is to allow <code>snpolicy</code> commands to complete before submitting additional <code>snpolicy</code> commands.

Operating Guidelines and Limitations

[Table 16](#) lists operating guidelines and limitations for running StorNext. Items are grouped according to operating system.

Table 16 StorNext Operating Guidelines and Limitations

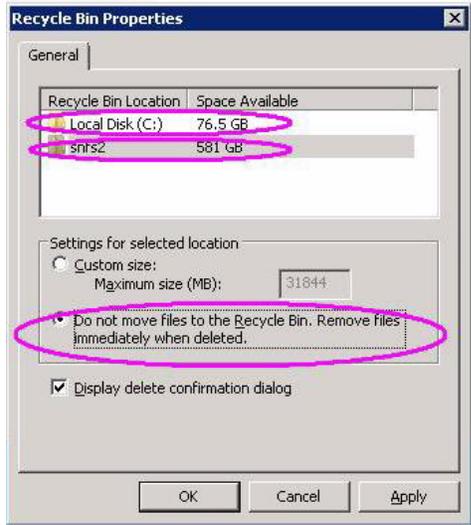
Operating System	Feature or Category	Description
AIX	Clients	<p>Clients on AIX systems may not unmount a file system after running the <code>fsstress</code> command against that file system. The client must then be rebooted to release the mount.</p> <p>This issue occurs on AIX systems when the <code>fsstress</code> command is run using <code>mknod</code> on the same command line. To prevent encountering this behavior, do not include <code>mknod</code> when running <code>fsstress</code>. Otherwise, you will be required to reboot the client without a successful unmount.</p> <p>Note: A StorNext 4.1.2 AIX client supports only AIX version 6.1</p>
Solaris	StorNext labels	<p>Solaris hosts may need to rescan disk devices after StorNext labels have been applied.</p> <p>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:</p> <pre>reboot -- -r</pre> <p>This issue will be addressed in a future StorNext release.</p> <p>In the meantime, work around this issue by rescanning devices using the boot flag <code>reboot -- -r</code></p> <p>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.</p>
Linux	Linux Multipath Support (the <code>rr_min_io</code> setting in the Linux DM Multipath Driver)	<p>Current versions of the Linux DM Multipath driver assign a default value of 1000 for <code>rr_min_io</code> which is too high for most configurations having multiple active paths to storage. Using a smaller value such as 32 will typically result in significantly improved performance. Refer to the RedHat or SuSE documentation provided with your version of Linux for details on how to apply this setting.</p> <p>Note: Experimentation may be required to determine the optimal value.</p>

Operating System	Feature or Category	Description
Linux	StorNext File System	<p>StorNext File System does not support the Linux <code>sendfile()</code> system call.</p> <p>This issue causes Apache web servers to deliver blank pages when content resides on StorNext file systems.</p> <p>This issue also affects Samba servers running on Linux.</p> <p>The workaround is to disable <code>sendfile</code> usage by adding the following entry into the Apache configuration file <code>httpd.conf</code>:</p> <pre>EnableSendfile off</pre> <p>The workaround for Samba servers is to add the following line into the configuration file:</p> <pre>sendfile=no</pre>
	HA	<p>Changing the <code>haFsType</code> 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.</p> <p>When the <code>HaShared</code> or <code>HaManaged</code> types are configured, other changes must be completed by successfully running the <code>cnvt2ha.sh</code> script, which is indicated by the creation of the <code>/usr/adic/install/.snm_ha_configured</code> touch file (<code>\$SNSM_HA_CONFIGURED</code> environment variable). No conversion is done or necessary for SNFS only (<code>HaUnmanaged</code>) configurations.</p> <p>If the conversion is not successfully completed, the <code>HaManaged</code> FSMs will not start, and the <code>HaShared</code> FSM will cause an HA Reset when it is stopped.</p> <p>To remedy this situation, edit every FSM configuration file to set its <code>haFsType</code> parameter to <code>HaUnmonitored</code>, then run the following commands to avoid the HA Reset in this special case only:</p> <pre>touch /usr/cvfs/install/.vip_down_hint service cvfs stop</pre>
	System logs	<p>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.</p> <p>This affects StorNext users on Red Hat 4, Red Hat 5, SuSe 9, and SuSe 10.</p> <p>This issue is not caused by StorNext, and is described in detail in StorNext Product Alert 20.</p> <p>For additional information, see Red Hat 5 CR 468088 and SuSE 10 CR 10440734121.</p>

Operating System	Feature or Category	Description
Linux	Migrating metadata controllers	<p>StorNext users migrating their metadata controllers from Apple Xsan to Linux should be aware of the following upgrade considerations:</p> <ul style="list-style-type: none"> • 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.
	Subtree Check option	<p>Subtree Check Option in NFS No Longer Supported</p> <p>Although supported in previous StorNext releases, the subtree_check option (which controls NFS checks on a file handle being within an exported subdirectory of a file system) is no longer supported as of StorNext 4.0.</p>
	FQDN	<p>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.</p> <p>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.</p>
Windows	Window backup utility	<p>When a StorNext file system is mounted to a drive letter or a directory, configure the Windows backup utility to NOT include the StorNext file system.</p>

Operating System	Feature or Category	Description
Windows	Upgrades on Windows Vista	<p>StorNext upgrades on Vista machines can fail in the middle of installation. This problem is caused by the way Windows Vista handles software upgrades. A related error is described in Microsoft article 263253.</p> <p>Microsoft has a utility called the Windows Installer Cleanup Utility that removes files left behind by incomplete installations. Access the Microsoft website and search for article ID 290301.</p> <p>To work around this issue, follow these steps:</p> <ol style="list-style-type: none">1 Click Start, and then click Run.2 In the Open box, type Regedit and then click OK.3 On the Edit menu, click Find.4 In the Find what box, type Snfs_XXX.dat and then click Find Next.5 If the search result selects a string value called PackageName, continue with these steps. Otherwise, repeat steps 3-4.6 Double-click the PackageName string value.7 In the Value data box, change the installation directory path to the new pathname. For example if the old installation directory path contained OCT10, change that to the current path (e.g, NOV12.)8 On the Registry menu, click Exit.

Operating System	Feature or Category	Description
Windows	Recycle bin	<p>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.</p> <p>You must disable the Recycle Bin for the drive on which a StorNext file system is mounted. Also, each occurrence of file system remapping (unmounting/mounting) will require disabling the Recycle Bin. For example, if you mount a file system on E: (and disable the Recycle Bin for that drive) and then remap the file system to F:, you must then disable the Recycle Bin on the F: drive.</p> <p>As of release 3.5, StorNext supports mounting file systems to a directory. For Windows Server 2003 and Windows XP you must disable the Recycle Bin for the root drive letter of the directory-mounted file system. (For example: For C:\MOUNT\File_System you would disable the Recycle Bin for the C: drive.)</p> <p>For Windows Server 2003 or Windows XP:</p> <ol style="list-style-type: none"> 1 On the Windows client machine, right-click the Recycle Bin icon on the desktop and then click Properties. 2 Click Global. 3 Click Configure drives independently. 4 Click the Local Disk tab that corresponds to the mapped or directory-mounted file system. 5 Click the checkbox Do not move files to the Recycle Bin. Remove files immediately when deleted. 6 Click Apply, and then click OK.

Operating System	Feature or Category	Description
Windows	Recycle bin (cont.)	<p>(Disabling the Recycle Bin, Continued)</p> <p>For Windows Server 2008, Windows Vista and Windows 7 systems, you must disable the Recycle Bin on C: and the File system name:</p> <ol style="list-style-type: none"> 1 On the Windows client machine, right-click the Recycle Bin icon on the desktop and then click Properties. 2 Click the General tab. 3 Select the mapped drive that corresponds to the StorNext mapped file system. For directory-mounted file systems, select the file system from the list. 4 Choose the option Do not move files to the Recycle Bin. Remove files immediately when deleted. 5 Click Apply. 6 Repeat steps 3-5 for each remaining directory-mounted file system. 7 When finished, click OK. 
All	File systems and stripe groups	<p>Be aware of the following limitations regarding file systems and stripe groups:</p> <ul style="list-style-type: none"> • The maximum number of disks per file system is 512 • The maximum number of disks per data stripe group is 128 • The maximum number of stripe groups per file system is 256 • The maximum number of tape drives is 256
	Managed file systems	<p>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.)</p>

Operating System	Feature or Category	Description
All	Upgrade	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.
	StorNext home page	The capacity indicators on the StorNext home page provide <i>approximations</i> and may not accurately summarize the actual current capacity. If you require accurate, up-to-the-minute capacity information, click the Capacity areas of the home page to view current capacity.
	Backups	Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.
	Tape drives	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.
	Cluster-Wide Central Control	The StorNext Cluster-Wide Central Control file (<code>nss_cctl.xml</code>) is used to enforce the cluster-wide security control on StorNext nodes (client nodes, fsm nodes, and nodes running <code>cvadmin</code>). 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.
	Labels	Disks with existing non-StorNext labels may not show up in the StorNext GUI in order to protect non-StorNext disks from being accidentally overwritten. If you need to label a disk that is not visible in the StorNext GUI, use the <code>cvlabel</code> command to label the disk or use <code>cvlabel -U</code> to remove the existing label from the disks. (Refer to the <code>cvlabel</code> man pages for instructions on labeling and unlabeled drives.) Caution: Modifying the label on an active non-StorNext disk can make the disk unusable. Proceed with caution.
	Xsan	It is not possible to delete data within a StorNext policy relation point from an Xsan client via the Finder. Rather, data must be deleted using the shell.
	Replication/ Deduplication	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	Feature or Category	Description
All	Replication/ Deduplication	<p>When either the StorNext GUI or the <code>snpolicy</code> command are used to create or modify a replication/deduplication policy, a policy text file is written to the file system.</p> <p>Example: Suppose that <code>/stornext/photos/</code> is the mount point for file system named photos. If a policy named <code>pol_replicate_1</code> is created in that file system, a text copy of the policy information called <code>/stornext/photos/.rep_private/config/pol_replicate_1</code> is created.</p> <p>If the file system is damaged and has to be recreated, the policy must also be recreated. This is simpler to do beginning with the StorNext 4.1 release because a backup copy of the policy text file is saved whenever a policy is created or updated. (The backup copy is saved as a file named <code>/usr/cvfs/data/fsname/policy_history/policyname.date_time</code>.)</p> <p>In the previous example, the file system name (<code>fsname</code>) is photos and the policy name is <code>pol_replicate_1</code>. So, the backup copy would have a name like this:</p> <pre>/usr/cvfs/data/photos/policy_history/ pol_replicate_1.2010-10-29_14-07-13</pre> <p>Note: The backup copy directory is not in the same file system is photos.</p> <p>If Storage Manager is used on the machine, all the policy backup files will be backed up along with the regular Storage Manager backups.</p> <p>Quantum suggests that after upgrading to StorNext 4.1.2 you run the command <code>snpolicy_gather -b > "some_file"</code></p> <p>This will save a copy of your current configuration. The <code>-b</code> option will also create a copy of policy information in the <code>usr/cvfs/data/fsname/policy_history</code> directory.</p>
	Replication	<p>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.</p> <p>This field works the same way as a UNIX shell which lets you pattern match names. For example, entering <code>*.0 core</code> would exclude all <code>.o</code> files and also files named "core." You could also skip all core files by entering <code>rep_skip=core*</code>.</p>

Operating System	Feature or Category	Description
All	HA	<p>On HA systems only:</p> <p>You may receive the following incorrect error message when scanning for a secondary MDC from the StorNext Convert to HA page:</p> <pre>WARN com.quantum.qutosgui.jsf.ha.HaMBean - doScanHost: Secondary system cannot be same as the primary system.</pre> <p>This message is generated if <code>/usr/adic/util/cnvt2ha.sh</code> fails for any reason (for example, if the file system exists on the secondary, if a shared file system can't mount, etc). Upon secondary conversion failures, StorNext resets the <code>ha_peer</code> file to <code>255.255.255.255</code> on the secondary. Since the conversion fails, the primary <code>ha_peer</code> file is not updated and faulty comparison logic causes the erroneous error message (<code>255.255.255.255 == 255.255.255.255</code>).</p> <p>The workaround consists of two steps:</p> <ol style="list-style-type: none"> 1 Remove the <code>/usr/cvfs/config/ha_peer</code> file from the secondary system. 2 Reset the StorNext processes on the secondary system by running <code>/etc/init.d/stornext_web restart</code>.
	HA	<p>On HA systems only:</p> <p>When a non-managed file system is converted to a managed file system in an HA pair, it is possible for the FSM on the secondary MDC to continue to operate this FSM as non-managed, which incorrectly allows the FSM to start on the secondary MDC.</p> <p>Restarting the CVFS service corrects the problem. Quantum recommends taking the following steps as a temporary workaround after converting any non-managed file systems to managed file systems:</p> <ol style="list-style-type: none"> 1 Complete the configuration changes 2 Make sure that CVFS is running on the secondary MDC, and wait 120 seconds to be sure that the configuration-file changes have been synchronized to the secondary MDC 3 Restart CVFS on the secondary by issuing "service cvfs restart" 4 Issue the command "cvadmin -e fsmist" on the secondary MDC, and make sure that the output shows the FSM as "State: Blocked (waiting for MDC to become HA primary)"

Operating System	Feature or Category	Description
All	HA	<p>Use caution when configuring the netmask for the HA Virtual Interface (VIP).</p> <p>The VIP is an alias IP address that is associated with a real interface. For example, if the VIP is based on eth0, eth0:ha will be created as the VIP.</p> <p>The netmask you associate with the VIP should generally be the same as that of the base interface, but in no case should it be more specific. For example, if the netmask on eth0 is 255.255.224.0 (a /19), then configuring the VIP netmask as anything more than a /19, such as a /24 (255.255.255.0) would be incorrect.</p> <p>Using the same /19 mask on both eth0 and eth0:ha is the correct approach.</p> <p>Note: The above applies only when the IP address of the VIP falls into the subnet defined by the base interface's IP address and mask.</p>
	HA	<p>Understanding the performance of FSM failover in StorNext High Availability installations:</p> <p>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.</p> <p>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.</p>
	Quotas	<p>When you enable or disable quotas using the CLI <code>cvadmin</code> command, the change does not persist after rebooting. In order to permanently enable or disable quotas, you must modify the Quotas parameter of the file system config file.</p>

Operating System	Feature or Category	Description
All	Distributed LAN	<p>Distributed LAN Clients in HA Environments:</p> <p>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.</p> <p>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.</p>
	Stripe group expansion	<p>StorNext does not support expansion on stripe groups containing mixed-sized LUNs. For example, if you create a file system that has two different-sized disks in a userdata only stripe group and then attempt to add a new disk to that stripe group and expand it, the expansion will fail.</p>
	dpserver	<p>In some cases the physical IP address must be included in the dpserver file in addition to the interface name. Note these conditions:</p> <ul style="list-style-type: none"> • 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.")
	Truncation	<p>By design, replication or deduplication must be completed before data files can be truncated if these files are associated with both a replication/dedup policy and a Storage Manager policy. Even if the Storage Manager policy is configured with the "Truncate Immediately" option, the truncation may not occur at store time unless the file has been replicated or deduplicated.</p>

Documentation

The following documents are currently available for StorNext products:

Document Number	Document Title
6-01658-16	<i>StorNext User's Guide</i>
6-00360-25	<i>StorNext Installation Guide</i>
6-01376-20	<i>StorNext File System Tuning Guide</i>
6-01620-20	<i>StorNext Upgrade Guide</i>
6-01688-16	<i>StorNext CLI Reference Guide</i>
6-67041-04	<i>StorNext File System Quick Reference Guide</i>
6-67042-04	<i>StorNext Storage Manager Quick Reference Guide</i>
6-66851-05	<i>StorNext HA Quick Reference Guide</i>
6-66852-05	<i>StorNext Replication Quick Reference Guide</i>
6-01641-10	<i>StorNext Licensing</i>

Contacting Quantum

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

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

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

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

