



StorNext 4.2.1 Release Notes

Product	StorNext® 4.2.1
Date	December 2011

Contents

Purpose of this Release	3
Other Changes and Considerations	4
Supported Platforms, Libraries and Drives.	5
Supported StorNext Upgrade Paths.	6
Compatibility With Other StorNext Components and Features	6

© 2011 Quantum Corporation. All rights reserved.

6-67376-02 Rev A, December 2011

Quantum, the Quantum logo, DLT, DLTtape, the DLTtape logo, Scalar, StorNext, the DLT logo, DXi, GoVault, SDLT, StorageCare, Super DLTtape, and SuperLoader are registered trademarks of Quantum Corporation in the U.S. and other countries. Protected by Pending and Issued U.S. and Foreign Patents, including U.S. Patent No. 5,990,810. LTO and Ultrium are trademarks of HP, IBM, and Quantum 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, LibXdmp, LibXext, LibXi copyright The Open Group

LibX11copyright 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 LibXdmp, 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

TinyXML, copyright (c) 2000 -- 2002 by Lee Thomason and Yves Berquin

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.



- Configuration Requirements 8
- Hardware Requirements 10
- Supported System Components 14
- Previous Versions of Release Notes 15
- Resolved Issues 16
- Known Issues 19
- Operating Guidelines and Limitations 31
- Documentation 44
- Contacting Quantum 45

Purpose of this Release

StorNext 4.2.1 is a maintenance release which includes important bug fixes described in the [Resolved Issues](#) section.

In addition, some enhancements which improve performance without changing functionality were added to this release.

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

New Supported Platforms, Libraries and Drives

StorNext 4.2.1 adds or expands support for the following platforms, libraries and drives:

- **Red Hat 5 Update 7**
- **Quantum StorNext Q-Series Storage**
- **Oracle/StorageTek T10000 Rev C Tape Drives.** Quantum has tested the T10000C single-drive performance with StorNext Storage Manager, and confirmed that the T10000C Drive can transfer files larger than 100MB at the nominal rate.

For detailed information including which StorNext components are supported for the above operating systems, refer to the "StorNext Supported Platforms" document for StorNext 4.2.1 available at this location:

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fifth#compatibility>

Read-Only File System Checking

You can now check active file systems in read-only mode through the StorNext GUI. Checking in read-only mode allows the check to complete faster, but is typically not as thorough as a check run in non-read only mode.

An advantage of running the check in read-only mode is that the file system can remain active during the check. (Active file systems must be checked in read-only mode, but offline file systems can be checked in either read-only mode or traditional mode.)

Quantum recommends also running the Recover Journal process whenever you check the file system in read-only mode.

Note: Whenever you check a file system in read-only mode on an active file system you might receive false errors because the check is running while the file system is still processing.

Support for Media Capacities Larger than 2TB

Beginning in release 4.2.1, StorNext now supports tape media whose capacities exceed 2TB.

Other Changes and Considerations

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.

Network File System (NFS) Support in StorNext

StorNext supports NFS version 3 (NFSv3). By default, newer versions of Linux such as RedHat Enterprise Linux 6.1 perform NFS mounts using NFS version 4, which is not supported.

Linux NFS clients mounting StorNext file systems with version 4 will experience incorrect results when performing the `readdir(2)` system call. This can lead to commands such as `ls` hanging and consuming large amounts of memory and potentially causing an out-of-memory condition for the entire system.

To avoid this problem, NFSv3 should be configured when using NFS with StorNext. On the Linux NFS client side, this is accomplished by specifying the `"vers=3"` mount option. Alternatively, NFSv4 can typically be disabled on Linux NFS servers by applying the following setting to the file `/etc/sysconfig/nfs`:

```
# Turn off v4 protocol support  
RPCNFSDARGS="-N 4"
```

Note: When disabling NFSv4 on the server side, all exported file systems are affected, not just StorNext volumes.

Consult the documentation for your version of Linux for details. Also refer to the StorNext File System Tuning Guide for further information on configuring StorNext for use with NFS.

Red Hat 6 and Virtual Memory

StorNext users may notice StorNext and other processes on Red Hat Linux 6 systems consuming much more virtual memory than on older releases. The real memory consumption should be approximately the same between the various releases. Extensive testing has NOT identified an increase in memory pressure with this condition.

Red Hat 6 has changed the memory allocator for multi-threaded programs so that each thread can (and many times does) receive its own heap arena. This memory is mapped into virtual memory in large chunks. However, the actual memory allocated by the thread still uses roughly the same amount of memory, as it would have if the heap arenas were shared. The increase in virtual memory size between Linux releases can be ignored. In other words, even though more virtual memory may get allocated with Red Hat 6, the actual memory pressure on the working memory won't increase noticeably.

Note: This issue currently applies only to Red Hat Linux 6 or higher, using a StorNext MDC on a 64-bit platform.

Replication Report Location on HA Systems

Starting with StorNext 4.2.1, the directory of replication completion reports history is converted to a shared directory on HA shared file systems.

This directory is located under `/usr/cvfs/data/fsname/rep_reports`, where "fsname" is the name of an snpolicy managed file system. The directory is updated whenever replication is performed.

If replication was performed prior to upgrading to 4.2.1, the primary and secondary node may both have part of the replication completion reports. The upgrade preserves the replication reports on the primary node, but discards the reports on the secondary node.

To avoid this issue, Quantum recommends backing up the replication reports for each snpolicy managed file system on the secondary node before upgrading to StorNext 4.2.1.

To work around this issue, manually copy the report content from the machine which was primary before failover to the new primary server.

Supported Platforms, Libraries and Drives

The operating systems, kernel versions, and hardware platforms supported by this StorNext 4.2.1 release are listed in the "StorNext Supported Platforms" document posted at this location:

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fifth#compatibility>

Note: The compatibility guides on this Web page show the supported platforms, libraries and drives for previous and current StorNext releases, so be sure you locate and view or download the version for StorNext 4.2.1.

The "StorNext Supported Platforms" document also shows StorNext client interoperability. Apple Xsan compatibility information is provided in a separate document.

Quantum OS Upgrade Support Policy

StorNext supports any security or functional bug update that applies to the current StorNext-supported Red Hat update level or SuSE Linux Service Patch. StorNext does **not** support updating the update level or service patch beyond the currently supported levels shown in the Supported Platforms document available at this location: <http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fifth#compatibility>

Supported StorNext Upgrade Paths

In general, sites running the following StorNext versions may upgrade directly to StorNext 4.2.1, 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 4.2.1:

- StorNext 4.1
- StorNext 4.1.1
- StorNext 4.1.2
- StorNext 4.1.3
- StorNext 4.2

All other versions of StorNext require additional steps to upgrade to StorNext 4.2.1.

Compatibility With Other StorNext Components and Features

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

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.

Refer to the PFR compatibility matrix at this location for information about compatibility between PFR and StorNext 4.2.1: <http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fifth#compatibility>

StorNext API (SNAPI)

StorNext API (SNAPI) enables you to run third-party APIs with StorNext.

Refer to the SNAPI compatibility matrix at this location for information about compatibility between SNAPI and StorNext 4.2.1: <http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fifth#compatibility>

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.

To ensure maximum replication performance, Quantum strongly recommends that all systems utilizing replication upgrade to StorNext 4.2.x

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

Table 1 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.x/4.2.x without Dedup.	4.1.x/4.2.x with Dedup.
4.0 without Dedup.	✓	--	See note below	--	See note below	--
4.0 with Dedup.	--	✓	--	Replication rejected	--	Replication rejected
4.0.1 without Dedup.	Replication rejected	--	✓	--	✓	--
4.0.1 with Dedup.	--	Replication rejected	--	✓	--	✓
4.1.x/4.2.x without Dedup.	Replication rejected	--	✓	--	✓	--
4.1.x/4.2.x with Dedup.	--	Replication rejected	--	✓	--	✓

Note: StorNext 4.0 to 4.0.1 and 4.1.x/4.2.x 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.

Configuration Requirements

Before installing StorNext 4.2.1, 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 or the tape drives therein. 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 4.2.1, the following hardware requirements must be met:

- [StorNext File System and Storage Manager Requirements](#) on page 10
- [StorNext Client Software Requirements](#) on page 12

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 2](#).

Table 2 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 4.2.1, 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.

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 22 million files replicated, which is approximately 4 bytes per entry.

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 4.2.1 are presented in [Table 3](#).

Table 3 StorNext Supported System Components

Component	Description
Tested Browsers	Internet Explorer 7, 8 or 9 Mozilla Firefox 3 or 4 Other browsers and versions besides the ones tested may work but are not officially supported.
NFS	Version 3 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. 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> . 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.
LDAP	LDAP (Lightweight Directory Access Protocol) support requires Windows Active Directory.

Component	Description
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 16
- [StorNext Storage Manager Resolved Issues](#) on page 17
- [StorNext GUI Resolved Issues](#) on page 18
- [StorNext Installation and Other Resolved Issues](#) on page 18

StorNext File System Resolved Issues

[Table 4](#) lists resolved issues that are specific to StorNext File System.

Table 4 StorNext File System Resolved Issues

Operating System	CR Number	SR Number	Description
Windows	31120	1120006, 1146724, 1236692	Corrected an issue which prevented Windows from following Linux symbolic links to a directory.
	33947	1272886	Resolved a condition in which all versions of StorNext CIFS Server may leak IRPs, resulting in memory leaks.
	33871	1237338, 1279306	Corrected a blue screen error in Symbolic Link Processing on Windows clients.
All	33514	1267158, 1293446	Addressed an issue which caused an fsm panic in Named_Inodeupdate.
	35626	n/a	Cvadmin's Error() function now works correctly.
	35637	1327392, 1333350	Resolved an issue which caused StorNext to compute a non-md5-compliant checksum in some cases, resulting in incorrect calculations.
	35792	1333174, 1335996	Corrected an issue which prevented Quotas from working correctly on a Windows MDC.
	35831	n/a	fmt_odd_size function prints wrong wording for TB, PB and EB.
	35877	1331410	CvApi updates were added to ensure that truncation free-space calculations ignore write-disabled stripe groups.

Operating System	CR Number	SR Number	Description
All	35717	n/a	Resolved an issue which resulted in the incorrect tape cartridges being exported/vaulted from Scalar Tape libraries.
	36102	n/a	A new capability bit was added to ensure that Xsan correctly distinguishes between an Xsan MDC which supports case insensitivity versus a StorNext 4.2.0 MDC which does not.

StorNext Storage Manager Resolved Issues

[Table 5](#) lists resolved issues that are specific to StorNext Storage Manager.

Table 5 StorNext Storage Manager Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	35444	1305778, 1334434	Resolved an issue which caused a ArcDisp process to core.
All	34744	1262264	Corrected a condition which caused a core dump after running the vsarchiveconfig command when adding media slots.
	35190	1309138	Resolved a condition in which running the fsretrieve -B (batch) option failed on file names which have spaces in them.
	35704	1331410	Debug logging for truncation policy is no longer incomplete.
	35801	1331410	Truncation free space calculations now ignore write-disabled stripe groups.

StorNext GUI Resolved Issues

[Table 6](#) lists resolved issues that are specific to the StorNext graphical user interface.

Table 6 StorNext GUI Resolved Issues

Operating System	CR Number	SR Number	Description
All	35605	n/a	Corrected a condition which caused the StorNext GUI to hang when displaying RAS tickets containing thousands of event details.
	35795	1336044	File age truncation entered through the StorNext GUI can now exceed six days.
	35832	n/a	Corrected a condition in which the StorNext GUI always passed a read-only flag to cvfsck to check the file system.

StorNext Installation and Other Resolved Issues

[Table 7](#) lists resolved issues that are specific to StorNext installation and other StorNext features or components.

Table 7 StorNext Installation and Other Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	35528	n/a	Corrected a condition in which running <code>/usr/adic</code> and <code>cvfs</code> symlinked; <code>install.stornext -removeall</code> removed only <code>cvfs</code> .

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 19
- [StorNext Storage Manager Known Issues](#) on page 21
- [StorNext GUI Known Issues](#) on page 24
- [StorNext Installation, Replication, HA and Other Known Issues](#) on page 27

Note: If you encounter one or more of the issues listed in this section, please contact Quantum Customer Support and report the issue(s) you encountered. Also inform the support representative whether you were able to successfully work around the issue(s) by using the provided workaround. Doing these things will help Quantum prioritize the order in which known issues are addressed in future StorNext releases.

StorNext File System Known Issues

[Table 8](#) lists known issues specific to StorNext File System.

Table 8 StorNext File System
Known Issues

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.	To work around this issue, you must find the appropriate network interfaces and reconfigure the dpserver.fsname file and then remount the file system.
	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.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
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.
	30696	n/a	Attempting to "loopback" mount an ISO image residing in a StorNext file system resulted in 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.
	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.

StorNext Storage Manager Known Issues

[Table 9](#) lists known issues specific to StorNext Storage Manager.

Table 9 StorNext Storage Manager Known Issues

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>
	29445	n/a	<p>When adding Distributed Data Mover (DDM) mover hosts (either by the StorNext GUI or by the <code>fsddmconfig</code> command,) the same host can be added multiple times. This can impact performance tuning.</p>	<p>The workaround is to avoid re-defining the same host under multiple equivalent identities.</p>

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	32596	1209112	<p>When Distributed Data Mover (DDM) is enabled, Storage Manager may fail to store any file whose size is sufficient to make it a multi-segment file of 3 or more media segments. The actual file size that causes storage failure may vary depending on media capacity and data compression values. Storage failure is not necessarily caused by storing a file of this size, but rather storing a large file that is continually changing.</p> <p>For example, suppose a large file is appended to at regular intervals of 30 minutes. After the file remains unmodified for 5 minutes (the default value,) the Storage Manager assumes it should begin storing the file.</p> <p>While it is storing the file (all segments) one of the 30 minute updates occurs; meanwhile the store processing is still continuing on the first set of segments. Five minutes after the update, the software re-issues another store attempt on the recently modified file, and you end up with wasted resources because the first stores are still occurring even though the data will be discarded at the end due to the file change. So in this example resources are wasted on a file that is not getting stored as re-store attempts begin queuing up.</p>	<p>If you have a large file that is regularly and continuously updated, and that file is large enough that it cannot be stored in the time between regular updates, mark the file so it will never be stored. Since it is never going to get completely stored anyway, this prevents wasting resources. Mark these files via the fschfiat command using the '-se' argument (exclude from store). If you have a large file that is regularly updated but will eventually be complete, you have two options:</p> <ol style="list-style-type: none"> 1 Manually mark the file as excluded from storing via fschfiat when it is created. When the file is done changing, use fschfiat again to clear the store exclusion. It will then be stored by a policy run as usual. 2 If you don't want the manual step (or to make it easier if you have a set of these files), set the minimum store time for the policy class via fsmodclass so that the software will not attempt to store the file until it is done changing. As previously mentioned, the default minimum store time is 5 minutes after the last update. <p>If you know a particular class of files will be large and change periodically over a known time interval, set the store time accordingly.</p> <p>For example, if you have a file that grows every half hour for a week until it is done, you could set the class time to one hour. If it has not been modified in an hour it will be stored. If, however, the file will be changing for a week but at random time intervals, you will want to set the minimum store time to be longer than the longest expected time interval between updates.</p>

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.
	32711	1213924 121777 1288408	When using the StorNext GUI to configure the Distributed Data Mover (DDM) feature, errors can result if you use the IP address when configuring the MDCs or clients.	To avoid this situation, use only the host name when configuring DDM.
	34145	1287532	When using the StorNext GUI to defragment or copy tapes on systems with multiple libraries, the GUI does not allow you to specify the library in which the target tape resides. This issue applies only to systems with multiple libraries, archives or partitions attached.	To work around this issue, use fsmecopy to defragment tapes or to copy media to tapes in the desired library.
	35007	n/a	StorNext does not currently support running two truncation policies for different policy classes at the same time.	To avoid errors, do not run two truncation policies for different classes at the same time.
	36211	n/a	Tape media information provided to the Stornext GUI during library discovery prevents displaying the correct T10000C tape media capacity. The minimum capacity for the T10000 family of tape cartridges is displayed rather than the actual media capacity.	The displayed capacity value is used only during library discovery and does not affect tape media usage.

StorNext GUI Known Issues

[Table 10](#) lists known issues specific to the StorNext GUI process.

Table 10 StorNext GUI Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	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.
	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.
	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: <code>DSM_control start</code> 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.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	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 by running the command <code>service stornext_web start</code> .
	29728	n/a	Due to an error in a third-party component, background pages respond to keyboard input when modal dialogue windows are open.	Avoid entering keyboard input when dialogue windows are displayed.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	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/mimportinfo <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>
	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>
	34989	n/a	<p>If a Web Services URL request contains the string "policy", it fails with a "Requested resource not found!" error.</p>	<p>To avoid this situation, do not use the string "policy" in a pathname that is intended for Web Services use.</p>

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	35496	n/a	When checking for data savings on a deduplication-enabled file system, the data on the Replication / Deduplication Policy Activity Report may be inaccurate.	Until this issue is resolved, do not rely exclusively on the data savings information shown in the Replication / Deduplication Policy Activity Report.
	36208	n/a	The StorNext GUI reports an incorrect media capacity for the T10000C tape cartridges during library discovery. The minimum capacity for the T10000 family of cartridges is displayed rather than the actual media capacity.	The displayed capacity value is used only during library discovery and does not affect the use of the media.

StorNext Installation, Replication, HA and Other Known Issues

[Table 11](#) lists known issues specific to StorNext installations, data replication, HA systems and other areas.

Table 11 StorNext Installation, Replication, HA and Other Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	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.)
	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	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.
	35387	n/a	When the system is configured and licensed for either replication or deduplication without also being licensed for both features, running <code>install.stornext -upgrade</code> fails with a licensing error. The license warning can be ignored, provided that the system has the appropriate licenses installed for the features in use.	When either replication or deduplication is licensed without the other, use the <code>-licforce</code> option when upgrading Stornext . For example: <code>install.stornext -upgrade -licforce</code>
All	29023	n/a	Replication quiesce scripts do not synchronize data on any clients that have open files.	To avoid this issue, close open files prior to running quiesce scripts.
	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 <code>snhamgr status</code> 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.	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. 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.
	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.
	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.
	35579	n/a	The <code>snpolicyd</code> process can crash on a replication-enabled source system when a replication source policy has configured multiple targets and a replication target file system has filled up.	To work around this issue, ensure that there is sufficient space in any replication target file systems to complete all outstanding replication tasks.

Operating System	Change Request Number	Service Request Number	Description	Workaround
All	36207	n/a	During archive discovery and when displaying media information, T10000C cleaning tapes are incorrectly identified as LTO-WORM media.	The incorrect media type does not prevent the cleaning tape from being used to successfully clean a T10000C drive.

Replication Multilink Known Issues

The Replication Multilink Feature

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

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

The following table shows known issues related to Multilink:

Table 12 Multilink Known Issues

Change Request Number	Description	Workaround
30746, 30921	Once enabled, the Multilink feature cannot be disabled from the StorNext GUI.	Contact Quantum Technical Support for assistance with disabling the Multilink feature.
30920	In HA configurations, the StorNext GUI may display errors about virtual IPs (vIPs) when configuring Multilink if the peer system is down.	Ensure the standby MDC is up when configuring Multilink.

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

Operating Guidelines and Limitations

[Table 13](#) lists operating guidelines and limitations for running StorNext. Items are grouped according to operating system.

Table 13 StorNext Operating Guidelines and Limitations

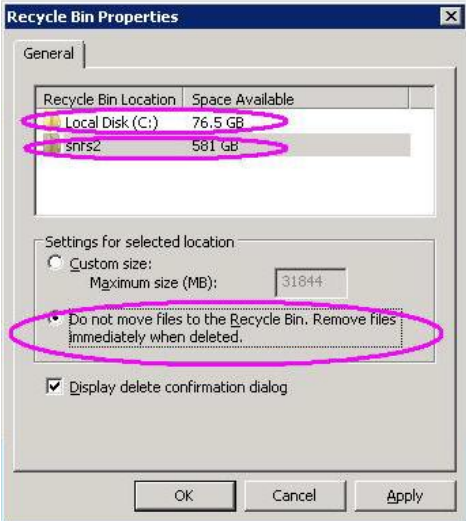
Operating System	Feature or Category	Description
Solaris	StorNext labels	<p>Solaris hosts may need to rescan disk devices after StorNext labels have been applied.</p> <p>In particular, when a StorNext label is put on a LUN less than 1TB in size, Solaris hosts will not be able to use that LUN until they 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>
	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>

Operating System	Feature or Category	Description
Linux	HA	<p>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.</p> <p>When the HaShared or HaManaged 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 (HaUnmanaged) configurations.</p> <p>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.</p> <p>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:</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, Red Hat 6, SuSe 10 and SuSe 11.</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>
	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.

Operating System	Feature or Category	Description
Linux	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>
	Software Firewalls	<p>Software firewalls such as “iptables” on Linux and Windows Firewall can interfere with the proper functioning of StorNext and result in unexpected errors unless specifically configured for use with StorNext.</p> <p>Quantum strongly recommends that all software firewalls be disabled on systems used as StorNext clients and servers. If required, StorNext can be configured for use with hardware firewalls.</p> <p>For more information, refer to the fsports man-page or help file and the “Ports Used By StorNext” section in the StorNext Tuning Guide.</p>
cpuspeed Service	<p>cpuspeed, an external Linux service on recent Intel processors, is not correctly tuned to allow StorNext to take advantage of processor speed. Suse systems may also be impacted, as may AMD processors with similar capabilities.</p> <p>On processors with a variable clockspeed (turboboost), the cpuspeed service on Redhat controls the actual running speed of the processors based on system load.</p> <p>A workload such as a heavily used FSM and probably Storage Manager does not register as something which needs a faster cpu. Turning off the cpuspeed service has been shown to double metadata performance on affected hardware.</p> <p>Looking at the reported CPU clock speed by doing <code>cat /proc/ cpuinfo</code> while the system is under load shows if a system is impacted by this issue.</p>	

Operating System	Feature or Category	Description
Linux	Power-on Diagnostics	<p>During testing a Quantum PX502 library running Red Hat 6.1 did not finish power-on diagnostics. When the same test was run on a PX502 library running either Red Hat 5.X or SuSE 10 / 11, power-on diagnostics completed and the system initialized without any issues.</p> <p>The workaround for this issue is to disconnect the SAN from the library running Red Hat 6.1. If the library powers on while the SAN is disconnected from the library controller, the library finishes its power-on diagnostics and performs an audit of the library. Subsequently reconnecting the Red Hat 6.1 server to the SAN (library ready) causes the library to perform a new physical audit of the library.</p> <p>Note: Testing was performed on the Red Hat 6.1 system which did not have StorNext loaded or running.</p>
Windows and Linux	Symbolic links to StorNext directories	<p>If you create a symbolic (soft) link in Linux to a directory on a StorNext file system, the link cannot be used by Windows. Windows also cannot process a symbolic link which contains a path to a file in another file system.</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>
	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. 
	Offline Notification Feature	<p>The StorNext Offline Notification feature is intended for single user systems only. Do not install this feature on systems where multiple users might be logged on at the same time.</p>
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	Internet Explorer 8 and 9	<p>When using StorNext with Internet Explorer 8 or 9, warnings about insecure and secure items may be generated. These warnings can be ignored and require no response. As a workaround, follow these steps to prevent the warning messages from appearing:</p> <ol style="list-style-type: none"> 1 Launch Internet Explorer 8 or 9. 2 Choose Internet Options from the Tools menu. 3 Click the Advanced tab. 4 Under the Security heading, make sure the option Display mixed content is checked. 5 Click OK.
	StorNext Licensing	<p>After you purchase or update a feature license and then enter license information through the StorNext GUI, you should restart StorNext services to ensure that your new license is recognized. Certain StorNext features such as replication may not recognize your license until services are restarted.</p>
	Upgrade	<p>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.</p>
	StorNext home page	<p>The capacity indicators on the StorNext home page provide approximations and may not accurately summarize the actual current capacity.</p> <p>If you require accurate, up-to-the-minute capacity information, click the Capacity areas of the home page to view current capacity.</p>
	Backups	<p>Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.</p>
	Tape drives	<p>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.</p>
	Cluster-Wide Central Control	<p>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.</p> <p>Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.</p>

Operating System	Feature or Category	Description
All	Labels	<p>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.)</p> <p>Caution: Modifying the label on an active non-StorNext disk can make the disk unusable. Proceed with caution.</p>
	Xsan	<p>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.</p>
	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 4.2.1 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>

Operating System	Feature or Category	Description
All	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.
	Replication	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 *.o core would exclude all .o files and also files named "core." You could also skip all core files by entering rep_skip=core*.
	Deduplication	If a deduplication candidate is removed before blockpool processing is completed, errors such as the following may be sent to the syslog: Oct 2 15:22:00 orleans Blockpool[16403]: E: [5] (Store Local) Error storing file "/stornext/source/____CVFS_Handle.000474F892EBB65E000E000000000000000000000292BF4". Error opening file "/stornext/source/____CVFS_Handle.000474F892EBB65E000E0000000000000000000000292BF4". No such file or directory. Errors such as these may appear serious, but there is no reason for concern. If you receive these errors, no action is required.
	Deduplication	For small deduplication-enabled configurations (1TB deduplication capacity), the default settings for snpolicyd may potentially cause memory starvation problems when ingesting to or retrieving from the blockpool. To address this, Quantum recommends changing the values of the parameters ingest_threads and event_threads to 4 (from their default values of 8) in the StorNext Replication/Deduplication configuration file (/usr/cvfs/config/snpolicyd.conf).
	HA	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	Feature or Category	Description
All	HA	<p>On HA systems only:</p> <p>The <code>/usr/cvfs/config/ha_peer</code> 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 <code>ha_peer</code> file to the IP address of the peer MDC, and restart StorNext by running the following command:</p> <pre>service cvfs restart</pre> <p>Note: The peer will be Primary after running this command.</p> <p>If the <code>ha_peer</code> file is removed for any length of time while StorNext is running, the <code>snhamgr(1)</code> 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 <code>snhamgr</code> itself. If this occurs, restore the <code>ha_peer</code> file to the IP address of the peer MDC, and then restart the HA Manager service by running the following command: <code>service snhamgr restart</code></p>
	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>.

Operating System	Feature or Category	Description
All	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 FSMPM 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 fsmlist" on the secondary MDC, and make sure that the output shows the FSM as "State: Blocked (waiting for MDC to become HA primary)"
	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>

Operating System	Feature or Category	Description
All	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>
	fsretrieve	<p>If you run multiple <code>fsretrieve</code> commands simultaneously to find files (for example, <code>find -type -f xargs fsretrieve</code>), you might receive error messages because doing this taxes system resources.</p> <p>Instead, use the recursive retrieve command. When you use this command the files under a directory are retrieved in batches, and more sorting is done to put files in tape order for increased performance. Run recursive retrieve by entering <code>% fsretrieve -R .</code></p>
	Distributed LAN	<p>Distributed LAN Clients in HA Environments:</p> <p>Each HA node must have its own <code>dpserver</code> files detailing the NICs on that node. The <code>dpserver</code> 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 <code>dpserver</code> 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>

Operating System	Feature or Category	Description
All	Stripe group expansion	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.
	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	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.
	DXi Virtual Tape Library Compatibility	<p>Note the following recommendations and limitations for using DXi as a virtual tape library for StorNext:</p> <ul style="list-style-type: none"> • Recommended library emulation: "ADIC Scalar i2000" • Recommended tape drive emulation: "IBM LTO-x" or "HP LTO-x" • DDM (Distributed Data Mover): This feature is currently not supported due to lack of full SCSI3 support in DXi.
	Affinities	<p>When a file system with two affinities is to be managed by the Storage Manager, the GUI forces those affinities to be named tier1 and tier2. This will cause an issue if a site has an existing unmanaged file system with two affinities with different names and wants to change that file system to be managed. There is a process for converting a file system so it can be managed but it is non-trivial and time consuming. Please contact Quantum Support if this is desired.</p> <p>Note: The restriction is in the StorNext GUI because of a current system limitation where affinity names must match between one managed file system and another. If a site was upgraded from a pre-4.0 version to post-4.0, the affinity names get passed along during the upgrade. For example, if prior to StorNext 4.0 the affinity names were aff1 and aff2, the GUI would restrict any new file systems to have those affinity names as opposed to tier1 and tier2.</p>

Operating System	Feature or Category	Description
All	Converting file systems	StorNext does not currently support converting from a managed file system to an unmanaged file system.
	JournalSize Setting	<p>The optimal settings for <code>JournalSize</code> are in the range between 16M and 64M, depending on the <code>FsBlockSize</code>. Avoid values greater than 64M due to potentially severe impacts on startup and failover times. Values at the higher end of the 16M-64M range may improve performance of metadata operations in some cases, although at the cost of slower startup and failover time. Existing file systems managed by StorNext Storage Manager MUST have a journal size of at least 64M. The TSM portion of SNSM may not start if the journal size for a managed file system is less than 64M.</p> <p>For more information about <code>JournalSize</code>, refer to “<code>FsBlockSize</code>, <code>Metadata Disk Size</code>, and <code>JournalSize Settings</code>” in Appendix A of the StorNext User’s Guide.</p>

Documentation

The following documents are currently available for StorNext products:

Document Number	Document Title
6-67370-02	StorNext User’s Guide
6-67366-02	StorNext Installation Guide
6-67367-02	StorNext File System Tuning Guide
6-67368-02	StorNext Upgrade Guide
6-67371-02	StorNext CLI Reference Guide
6-67374-02	StorNext File System Quick Reference Guide
6-67380-02	StorNext Storage Manager Quick Reference Guide
6-67399-02	StorNext File System API Guide
6-67369-02	StorNext Licensing

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

