

StorNext 4.3.2 Release Notes

Product	StorNext® 4.3.2
Date	January 2013

Made in the USA. Quantum Corporation provides this publication “as is” without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.

COPYRIGHT STATEMENT

© 2013 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.

TRADEMARK STATEMENT

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. Preserving the World's Most Important Data. Yours., Backup. Recovery. Archive. It's What We Do., the DLT logo, DLTSage, Dynamic Powerdown, FastSense, FlexLink, GoVault, MediaShield, Optyon, Pocket-sized. Well-armed, SDLT, SiteCare, SmartVerify, StorageCare, Super DLTtape, and Vision are trademarks of Quantum. 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, 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 © 1994-2002 The XFree86 Project, Inc. All Rights Reserved. And

copyright © 1996 NVIDIA, Corp. NVIDIA design patents pending in the U.S. and foreign countries.

Libxml2 and LibXdmcp, copyright MIT

MySQL, copyright © Oracle Corporation

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

strlcat.c, strlcpy.c copyright © 1998 Todd C. Miller

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



Contents

What's New in StorNext 4.3.2	3
What's New in StorNext 4.3.1	6
What's New in StorNext 4.3	9
Other Changes and Considerations	24
StorNext Compatibility	27
Supported StorNext Upgrade Paths	28
Compatibility With Other StorNext Components and Features	29
Configuration Requirements	31
Hardware Requirements	33
Supported System Components	37
Previous Versions of Release Notes	38
Known Issues	39
Operating Guidelines and Limitations	51
Contacting Quantum	65

What's New in StorNext 4.3.2

Purpose of this Release

StorNext 4.3.2 is a maintenance release of StorNext that adds LTO-6 support, and fixes over 85 issues, including 51 customer-reported issues. See the *StorNext Compatibility Guide* for libraries supported with LTO-6. For details on fixes for customer reported issues, see [Customer Reported Fixed Issues and Enhancements](#) on page 3 in this document.

New Features in StorNext 4.3.2

Support for LTO-6

StorNext 4.3.2 adds support for LTO-6. The tape capacity of LTO-6 is 2.5 TB/tape uncompressed or 6.25 TB/tape compressed. LTO-6 tape drives also have an idle power consumption of just 4 W, resulting in even higher energy efficiency vs spinning disk.

Support for Windows 8 as a client

StorNext 4.3.2 adds support for Windows 8 as a client. See the *StorNext 4.3.2 Compatibility Guide* for details. The *StorNext 4.3.2 Compatibility Guide* is online at: <http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Fourth>

Customer Reported Fixed Issues and Enhancements

The following table lists the customer reported fixed issues and enhancements for StorNext 4.3.2.

Operating System	Change Request Number	Service Request Number	Description
All	32395	1345028	GUI doesn't allow 0 for buffer cache read ahead value
All	32750	1213826	fsrecover behaves differently in report mode and can be confusing
All	33861	1275078, 1340806, 1364246, 1375912, 1380562, 1394616, 1409014, 1410462, 1478136	MSM config_file_<library> files get corrupted

Operating System	Change Request Number	Service Request Number	Description
All	35006	1304018, 1312422, 1312592, 1317006, 1359920, 1403366, 1448800	User's Guide: document procedure for moving HA shared f/s to new RAID
All	35815	1333946, 1432722	Unnecessary, repetitive fs_fmover messages
All	36361	1337734, 1386160	Activating fsm not responsive to voting clients
All	37380	1398402	The GUI job table blows up on the amount of commands sent to the table
All	37397	1391640, 1406444	When DDMs are disabled retrieves fail
All	37451	1397708, 1418042, 1498350	Segmentation fault from network order incorrect for remote client replay error message
All	37543	1401350	fs_fmoverc fails creating parm.res but errno is not printed
All	37946	1389684, 1419568	.auth_secret file on non-Apple clients should be flagged more explicitly
All	38088	1473108	GUI: Browsing directories with more than 500k files will cause GUI error, 1M will cause GUI hang
All	38112	1416246	The "Total Files" listed in the GUI does not match the actual number of files on the filesystem.
All	38464	1483968	Unnecessary RAS Emails
All	38701	1434230	cvfsck failure related to bad_name_list
All	38846	1442074	pse_snapshot now calls cvgather for all file systems and duplicates core files
All	39020	1444998	QoS: rtio client hang can cause the FSM to hit OpHangLimit
All	39217	1456410	DDM does not recognize fsddmconfig'ed MDC hostname reliably
All	39425	1314038	Confusing RAS event caused customer contact support
All	39427	1420742	Release Notes: Cannot mv a file from one directory to another at same level when they have same policy

Operating System	Change Request Number	Service Request Number	Description
All	39446	1460842	Use the MAX_FILES_PER_CLUSTER parameter on retrieves
All	39451	1464156	Putting a drive offline does not remove the SCSI-3 reservation
All	39518	1455128	DDM mover clients are not disabled
All	39522	1466130	Chown command can fail because of limitation on the number of gids associated with a file in AIX
All	39549	1474866	kernel panic: CVFS ASSERTION FAILED: f_rwlck->rw_state & RW_WRITE line 2379 file
All	39551	1471604, 1473518	snmetadump resizing not updating the SG depth field
All	39572	1384234, 1450522	GUI hangs in the StorNext Media Action page when performing media operations
All	39591	n/a	GUI: Add support for policy drive limits
All	39847	1459366, 1468688	SNFS client buffer cache should handle ENOMEM errors better
All	39863	1477098, 1488088	fsrecover can't undelete a deleted file on replication target under SNMS
All	39992	1477208	CLI moves of tapes from vault to library can cause 'Assigned Locations' to increment by 2 not 1
All	40029	1479162	GUI displays incorrect time for "Last Run" on Schedule view page
All	40033	1479612	GUI get Java heap space errors in the file recovery webpage
All	40145	1435656 1484602 1499696	Large DMA I/O spanning many luns hangs
All	40449	n/a	GUI adjust fsretrieve GUI job status and job details
All	40500	1487006	User's Guide: Failure to start a filesystem standby due to license issues went undetected
All	40670	1489608	FILEINFO table primary key missing in MySQL DB after upgrade from Linter DB
All	40936	1504406	Filenames with 249 to 255 characters cause store problems
Linux	38354	1420026, 1428770, 1451220	No I/O via drive, error: "Launch timed out, rtn: 7" for fs_fmover

Operating System	Change Request Number	Service Request Number	Description
Linux	39069	1371426	Linux automount 5.0+ not supported because mount_cvfs does not support -f option
Linux	39317	1454754	Customer's M662 running SNFS 4.3.0 Build 26880 host name was changed by the GUI software to "Click-Scan-to-Refresh"
Linux	39395	1454742	SLES 11 SP2 Direct I/O fails with "large" requests
Linux	39556	1422802	Gateway Proxy server leaks kernel memory for each TCP proxy socket at time of disconnect.
Linux	39573	1469216	Gateway agent not reporting information about Linux device-mapper controlled disk devices
Linux	39695	1464094	TSM-based tape reads may fail with E_SRC_BOF_MS or similar
Linux	40169	1482638, 1494196	Checksum information for file is lost during fsmedcopy . For a detailed description of the problem, and the workaround, see Change Request Number 40169 in Table 7 on page 41.
Linux	40384	1482638	Release Note for 40169: Checksum information for file is lost during fsmedcopy . For a detailed description of the problem, and the workaround, see Change Request Number 40169 in Table 7 on page 41.
Other	34394	1284900, 1493068	CVFS ASSERTION FAILED: f_io->io_components > 0 line 1418 file iodesc.c
Windows	38717	1436466	Windows 7 Ultimate client panic running mount_options_test to M330 MDC
Windows	40058	1472016, 1466368	BreakWriteOplock()/OplockChangeDone() leaks locked memory pages

What's New in StorNext 4.3.1

Purpose of this Release

StorNext 4.3.1 is a maintenance release for StorNext 4.3, with over 40 bug fixes and enhancements, including those for 13 customer-reported issues, listed below.

New Features in StorNext 4.3.1

StorNext 4.3.1 adds client support for SuSE Enterprise Linux 11 SP 2.

Customer Reported Fixed Issues and Enhancements

The following table lists the customer reported fixed issues and enhancements for this release of StorNext.

Operating System	Change Request Number	Service Request Number	Description
All	29504	820088	Need tool to verify database and filesystem are in sync
All	36324	n/a	Improve logging for MySQL error 2013
All	37619	1383828	Mac clients soft quotas exceed warning messages are not written to the user's console
All	37950	1413766	Document license installation steps for RPM-only install in the Installation Guide
All	38080	1416212	Create procedure in User guide on to return secondary HA MDC from Locked:Stopped
All	38444	1424818, 1419878	GUI appears hung when it can't resolve DLC's hostnames
All	39144	1354190, 1420958	Error using "Delay File Storage" Option in GUI
Linux	32438	1202314, 1205446, 1288386, 978164, 1321718, 1376596, 1348778, 1365298, 1381616, 1407656, 1427076, 1435124, 1449244	Improve StorNext log rolling to reduce likelihood of filling up partition where StorNext is installed. For additional information, see Improved StorNext Log Rolling on page 8.
Linux	38550	1407654	M330 upgrade succeeded on secondary but failed on primary with expired maintenance license
Linux	38987	1443302	G300 Gateway is not forwarding performance metrics to MDC
Linux	38991	1443304	G300 Gateway long reboot time with large number of disk paths

Operating System	Change Request Number	Service Request Number	Description
Windows	37392	1382092	StorNext Windows security issues with traverse checking policy
Windows	38631	1384372	Windows Install: upgrade from 4.1.0 to 4.2.2 fails because registry SID string is longer than buffer.

Improved StorNext Log Rolling

The following information supports **Change Request Number 32438** which is listed in the **Customer Reported Fixed Issues and Enhancements** table above.

StorNext automatically rolls logs via a scheduled cron job. StorNext 4.3.1 enhances log rolling to better manage the file system space used by the rolled log files. An additional log rolling cron job will now run frequently to check for and roll "runaway" log files. Additional enhancements include:

- Compression of rolled log files
- New configuration parameters specifically for space management
- Efficiency improvements for better log rolling performance

The new configuration parameters include:

- **CRITICAL_FILL_LEVEL**: If the file system exceeds this fill level (default: 98%), StorNext will remove rolled files to recover space.
- **CLEANUP_MIN_SIZE**: StorNext will remove rolled files during file system space recovery only if they meet this minimum size (default: 10 MB).

All log rolling configuration parameters are contained in the `sn_log_update.cfg` configuration file located in the following directory:

`/usr/adic/util/sn_log_update.cfg`

Note: The improved StorNext log rolling will not guarantee that the file system will not fill up, given that other non-StorNext files using the same file system may accumulate and fill up the file system. Additionally, the probability exists that a StorNext log may grow at an extraordinarily rapid rate that exceeds the ability for the automatic log rolling to keep up.

StorNext uses a timestamp for the filename extension, instead of a sequential numeric count extension. The rolled files are compressed; a ".gz" extension is appended to the filename. The complete filename format for rolled log files is illustrated in the following example:

`tac_00.08:08:2012:13:00:01.gz`

You may also configure the log rolling cron job to back up rolled log files to a managed file system. A new storage policy class will be needed, which will require additional media or sdisk space. The size of the data stored will depend on system activity levels.

The instructions below will back up all rolled logs. Additionally, the MSM and TSM tac log backups will be compressed to approximately 1/20th of their original size before being stored.

Identify the managed file system containing the `.ADIC_INTERNAL_BACKUP` directory (typically `/stornext/snfs1`). See the `BACKUPFS` environment variable in `/usr/adic/TSM/config/fs_sysparm` to determine the file system name. In the instructions below, change `/stornext/snfs1` as necessary to the name of the managed file system containing the backups.

- 1 If this is a High Availability (HA) configuration, issue the following command on both MDCs.

```
# mkdir -p /stornext/snfs1/.SNSM_LOG_BACKUP/`/bin/hostname`
```

- 2 Issue the following two commands on the active MDC only.

```
# fsaddclass -d 1 -f i -m 5 _snsm_log_backup
```

```
# fsaddrelation /stornext/snfs1/.SNSM_LOG_BACKUP -c  
_snsm_log_backup
```

- 3 Save the existing `tdlm` crontab so it may be restored if an error occurs while updating the crontab.

```
# /usr/bin/crontab -l -u tdlm > /tmp/crontab.tdlm.save
```

- 4 Issue the following command to edit the `tdlm` crontab on the active MDC.

```
# /usr/bin/crontab -e -u tdlm
```

- 5 Append the following text to the end of the existing `sn_log_update` entry. It is all one continuous line. Note that it begins with a space, and there is a space preceding every hyphen, every occurrence of `/usr/adic`, and every occurrence of `/stornext/snfs1`.

```
-s /stornext/snfs1/.SNSM_LOG_BACKUP/`/bin/hostname`; /usr/  
adic/gui/bin/cmdwrap -NO_END_OF_FILE /bin/gzip /stornext/  
snfs1/.SNSM_LOG_BACKUP/`/bin/hostname`/?SM/logs/tac/  
tac_00.??:?:20??:??:??:??
```

A complete **crontab** command is illustrated in the following example:

```
0 1,7,13,19 * * * /usr/adic/gui/bin/cmdwrap -NO_END_OF_FILE  
/usr/adic/util/sn_log_update /usr/adic -s /stornext/snfs1/  
.SNSM_LOG_BACKUP/`/bin/hostname`; /usr/adic/gui/bin/cmdwrap  
-NO_END_OF_FILE /bin/gzip /stornext/snfs1/.SNSM_LOG_BACKUP/  
`/bin/hostname`/?SM/logs/tac/tac_00.??:?:20??:??:??:??
```

What's New in StorNext 4.3

Purpose of this Release

StorNext 4.3 is an important new release of StorNext, with over 630 bug fixes and enhancements including more than 120 customer-reported issues.

New Features and Enhancements

Improved Scalability and Performance of StorNext Storage Manager

This release replaces the database used for StorNext Storage Manager to achieve improved scalability and performance: MySQL replaces Linter.

The larger the file system, the bigger the performance increase.

The StorNext Storage Manager database in StorNext 4.3 has been measured at up to 1000x faster than StorNext 4.2 at truncation in very large file systems. In smaller file systems of 1 million files, truncation has been measured at 1.5x faster than StorNext 4.2.

At one million files, file systems with StorNext Storage Manager create files up to 20% faster in StorNext 4.3 compared to StorNext 4.2.

Conversion from Linter to MySQL happens during conversion, and may take significant downtime. See the StorNext Upgrade Guide for details, including a way to run a trial conversion on a separate system to measure time in customer environments.

Support for 1 Billion Files per MDC

StorNext 4.3 supports up to 1 billion files per MDC, in 1 to 8 file systems.

Faster Directory Lookups, Creates and Deletes

StorNext 4.3 supports up to 8 GB of RAM for the FSM buffer cache, which can significantly boost the speed of directory operations by exploiting the increased memory sizes of modern systems. Prior releases of StorNext limited the FSM buffer cache to 1 GB.

Creating files has been measured at up to 50% faster in StorNext 4.3 than StorNext 4.2.

Deleting files has been measured at up to 30% faster in StorNext 4.3 than StorNext 4.2.

Performance Optimizations for Windows Clients

StorNext Windows clients perform faster with StorNext 4.3 because of new optimizations, especially for "small" I/O sizes like 64 kB reads and writes. Latency in the StorNext code has been reduced by up to 7x, and performance increases of up to 200% have been measured for single-stream 64 kB writes on Windows SAN clients; 70% increases for four-stream 64kB writes on Windows SAN clients. Windows DLC clients also benefit.

Performance Optimizations for Writing Small Files to Tape

StorNext 4.3 boosts the performance of writing small files to tape. StorNext 4.3 writes 1 MB files to tape up to 4x faster than StorNext 4.2.

Support for Project-oriented Quotas on Directories

StorNext customers have asked for an easy way to manage the storage space and number of files used by projects—as defined by directory—rather than managing by user.

StorNext 4.3 augments existing user and group quotas with directory quotas, letting administrators define directory quotas to limit usage in a directory and all its sub directories. Directory quotas have separate hard limits, which cannot be exceeded, and soft limits which can be exceeded for a 'grace period.' Directory quotas can be defined in the StorNext GUI or the StorNext command line.

Directory quotas may also be configured without enforcement to provide a quick way to regularly check on size of selected directories without needing to run expensive "**du -sk**" commands. Directory quotas are logged daily, with administrator control of retention, as text files suitable for easy reporting.

For best results, the Metadata Controller and all clients should be at StorNext 4.3 when using directory quotas.

New **fsactivevault** Command for Easier Active Vault Use

StorNext 4.3 makes using Quantum Scalar Active Vault easier than ever with a new tool to select tapes for vaulting, **fsactivevault**. Policies can pick tapes by age, size, percent used or a combination, and prioritize by sorting on those criteria. Policies operate by high- and low- watermark to automatically free up space in the host-managed partitions of the library. Alternately, users can limit the number of tapes vaulted at any one time with a maximum count. The **fsactive vault** command can move tapes with **vsmove**, or simply create a report.

The command can be scheduled to run automatically.

The **fsactivevault** command can also be used for picking tapes for vaulting outside of a Quantum Scalar Active Vault.

New Tool, **Vidiomap**, for Changing File Sequence on Disk

Customers with files always used in a specific order, like file-per-frame movies in DPX format, benefit from having files arranged sequentially on disk. Already good at arranging storage sequentially-within-a-file, StorNext was extended to support arranging storage sequentially-across-files in StorNext 4.2 with the Allocation Session Reservation (ASR) feature. Once enabled, ASR arranges files on disk in the order of creation on a client.

StorNext 4.3 compliments ASR with a tool, **vidiomap**, to re-order files by name, or by a user-specified list. This allows customers to re-order files on disk that were created before ASR was enabled, or files copied in arbitrary order from another location. **Vidiomap** can also report on the existing file layout.

Improved StorNext Copy Command, **cvcp**

The StorNext 4.2 copy command, **cvcp**, now better supports incremental copy.

In StorNext 4.3, **cvcp** can now be used with both the **-u** and **-z** flags ("update only" and "preserve modification times"); StorNext 4.2 did not support this.

The **cvcp** command has also been enhanced to support long path names with the long path prefix (\\?\).

Problems on Windows with very large directories have been addressed.

Improved Migration Tools

StorNext 4.3 gives users better ability to upgrade the storage hardware for file system metadata. Prior to StorNext 4.3, metadata could only be copied to the same number of LUNs of the same size, meaning that a file system with metadata on four old 100 GB LUNs could only be moved to four new 100 GB LUNs. In StorNext 4.3 the **snmetadump** command has been enhanced to allow recovery to different geometry. See the **snmetadump** man page for details.

Improvements in LDAP Integration

StorNext 4.3 allows installation of a StorNext configuration file for the Windows SAN clients that will allow the administrator to configure the LDAP server/port and LDAP search filter parameters necessary to find the UID and GID values in Active Directory.

Users can configure:

- LDAP server and optional port number that manages the user account information
- LDAP server and optional port number that manages the Unix UID/GID information
- Search base (Distinguished Name) of the AD entry at which to start the search for the user account information
- Search base (Distinguished Name) of the AD entry at which to start the search for the Unix UID/GID information
- Search base (Distinguished Name) of the AD entry at which to start the search for the objectGUID information, which is used to fabricate a UID/GID for compatibility with Apple MacOS clients

Improved Service Ticket Filtering in the StorNext GUI

StorNext 4.3 adds support for extensive filtering of StorNext service tickets (RAS messages) in the StorNext GUI. Users can limit based on priority, component, event, or last update.

Improved File System Check and Repair Tool, **cvfsck**

StorNext 4.3 improves the performance of **cvfsck** when doing directory repair on large directories.

StorNext 4.3 **cvfsck** also reports more summary information on the file system, and how long **cvfsck** took to run.

Control of StorNext GUI Auto-logout

To give administrators and users better control of the balance between security and ease of use in their environments, the auto-logout in the StorNext GUI can be changed from 10 minutes up to 12 hours. Changing the auto-logout time can be only done by administrators and others with the Manage Users privilege.

Improved Error Reporting and Progress Monitoring in StorNext Replication

StorNext replication has been enhanced to provide better progress reporting and error messages.

Updates to Third Party Packages

StorNext 4.3 updates the embedded Apache Tomcat from 5.5.23 to 7.0.23.
StorNext 4.3 updates the embedded version of Java from 1.6.0.21 to 1.6.0.31.

StorNext Man Pages Available as PDF

The most detailed documentation on StorNext commands and configuration files is the StorNext man pages. Formerly reading the man pages required logging into a terminal on a UNIX system. With StorNext 4.3 the man pages are available as a PDF document for convenience. In particular, searching the content of the man pages is significantly easier with the PDF.

Support For No-downtime HA Failover Upgrade For MDCs Without StorNext GUI or Storage Manager Installed

Customers using StorNext High Availability who install StorNext without the StorNext GUI or StorNext Storage Manager can perform a no-downtime upgrade of their Metadata Controller by upgrading the secondary MDC node, failing over, and then updating the other host. See the *StorNext Upgrade Guide* posted here (click the "Select a StorNext Version" menu to view the desired documents):

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

StorNext Administrators can Prevent Clients from Halting a File System with cvadmin

StorNext 4.3 gives administrators the ability to control which clients can halt a file system. See the *StorNext Installation Guide* posted here (click the "Select a StorNext Version" menu to view the desired documents):

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

Support For Apple Named Streams in Mac-only File Systems

StorNext 4.3 adds support for Apple Named Streams in file systems where files will only be accessed from MacOS.

Note: Changing or copying a file in a named streams-enabled file system from a StorNext client or MDC on Linux, Windows or UNIX will generally “slice off” the named stream data, as the application won’t be able to access the Apple Named Stream. File systems using named streams cannot be used with StorNext Storage Manager, replication, or deduplication.

StorNext GUI Reports Un-purchased Licenses in a Neutral Way

StorNext 4.2 displayed un-purchased licenses with a red indicator, causing confusion about whether there was an error. StorNext displays un-purchased licenses with a neutral gray icon.

Added Support

StorNext 4.3 adds support for:

- RedHat Enterprise Linux 5 update 8

Dropped Support

StorNext 4.3 drops support for the following Operating Systems that were susceptible to the silent data corruption bug documented in Product Alert #20:

- RedHat Enterprise Linux 5 updates 1-3
- SuSE 10 service packs 1 & 2

Customer Reported Fixed Issues and Enhancements

The following table lists the customer reported fixed issues and enhancements for this release of StorNext.

Component	Operating System	Change Request Number	Service Request Number	Summary
GUI	All	29381	n/a	GUI could use a refresh button when displaying disk lists
GUI	All	33093	1245612	GUI is unable to order the luns in a Stripe Group for optimal performance
GUI	All	33449	1264450	GUI restarted file system without a warning

Component	Operating System	Change Request Number	Service Request Number	Summary
GUI	All	33456	1252504	GUI should not allow to set RestoreJournal to Yes on unmanaged FS
GUI	All	35641	1310620	GUI displays incorrect number of unformatted media
GUI	All	35675	1330008, 1348494	Media report does not show cleaning media
GUI	All	35832	1296510	GUI: Gui always passes read-only flag to cvfsck to check the filesystem
GUI	All	35970	1343326	DDM config failure error message is misleading/not descriptive enough
GUI	All	36142	1351256	Tape consolidation values not updating when fs_sysparm_override resides in StorNext
GUI	All	36285	1358052	There is no indication in the StorNext 4.x GUI that SDISK is set to a media status of "unavailable"
GUI	All	36379	1349980	Capacity utilization for storage manager
GUI	All	36462	1369498, 1369854	Edit of filesystem via GUI loses custom changes from fstab
GUI	All	37347	1344818	GUI does not give an option to update expired license when Stornext in HA mode failed to start properly
GUI	All	38045	1415842	Add StorNext version check to convert to HA process when scanning peer for eligibility in GUI
GUI	All	36437	1359522	Tools -> File Systems -> Truncation Parameters; display errors
GUI	Linux	34775	1297088	Configuring Delay Minimum File Size in GUI displays and may set incorrect value
GUI	Linux	35627	1359522	GUI inserts previously selected file system values in truncation parameters field when different file system is selected
GUI	Linux	36025	1346704 1352690	.csv file from GUI media report page always shows 0% used

Component	Operating System	Change Request Number	Service Request Number	Summary
GUI	Linux	36660	1376722	GUI Reports > File Systems should display Stripe Group index instead of ordinal
GUI	Linux	36834	n/a	StorNext: GUI: RAID Type: LSI: default Segment Size is incorrect
GUI	Linux	36846	n/a	StorNext: GUI: Libraries: Name with spaces causes an archive hardware error
GUI	Linux	36939	1375706	StorNext GUI displays inconsistent Stripe Group status in different screens
GUI	Linux	37759	n/a	Web Services enhancement for StorNext: Add a new web service to support moving files using the fsrelocate command; Enhance the getSMQueue web service API to support the "-f" and "-r" options
GUI	Linux	36970	n/a	StorNext: GUI: cannot create copy 2 on tape
Install	All	35618	1326894	Upgrade fails when chkconfig not found because /sbin not in PATH
Install	All	36589	n/a	Upgrade from 4.2 to MySQL Beta failed. The root file system went to 100%
Install	All	36735	1369648	User was able to upgrade only the snfs RPM's without being warned in a full management suite StorNext environment
Install	All	37846	n/a	install.storNext should report progress of linter2mysql conversion
License	All	35088	1389820	Error opening license file: SNFS Cfg: Error(2) opening '/usr/cvfs/config/license.dat'
License	All	36133	n/a	Allow cvfskey to handle more serial number formats
MSM	All	36110	n/a	vsintransitqry man page exists, but there is no such command
MSM	All	36339	n/a	Device Qual - SpectraLogic T-Finity Library with IBM TS-1140 Tape Drives

Component	Operating System	Change Request Number	Service Request Number	Summary
MSM	Linux	31144	1121576 1174998, 1232118	filelist modified when StorNext experiences an un-graceful shutdown
MSM	Linux	35616	n/a	Add script to vault MIGRATE class media
MSM	Linux	37393	n/a	StorNext: Storage Manager: vsclareject: Segmentation fault: core dumped
MSM	Linux	37410	n/a	StorNext: Storage Manager: vsarchiveconfig: a man page doesn't exist
MySQL	All	33972	MySQL-BETA	MYSQL: linter2mysql needs to disable triggers and stored procedures on import into mysql
MySQL	All	36323	n/a	MYSQL: linter2mysql incorrectly attempts to process tables that are unknown to StorNext
MySQL	All	37275	n/a	MYSQL: Queries that do full table scans of large table can cause MySQL to use up available resources
MySQL	All	37706	n/a	Optimize filecomp table queries used by fsmedcopy
MySQL	All	36324	n/a	MYSQL: linter2mysql fails due to MySQL errors
NXDI	All	36125	n/a	Library/drive taken offline while library coming ready
NXDI	All	37336	1395540	GUI reported error on the library partition instead of displaying xdi_sncli errors
PSE	All	31775	1165240	Starting PSE manually can prevent filesystem from unmounting
PSE	All	32990	1214424	pse_snapshot needs to use the filelist in /usr/cvfs/config
RepDedup	Linux	35222	1312226 1371020	cleanup_tree failed due to file names with special characters
RepDedup	Linux	35632	1325082	Print out key file contents as part of snpolicy_gather output
SNFS	All	37260	1391838	Large DMA I/Os break up into fs-blocksize device transfers unnecessarily affecting I/O performance

Component	Operating System	Change Request Number	Service Request Number	Summary
SNFS	All	14982	1292772	space in fsmlist causes SNSM to not start
SNFS	All	26300	909788	cvfsck runs extremely slowly with big directories containing many orphan entries
SNFS	All	26653	946732	cvfsck should print useful information before panicking
SNFS	All	27723	993534, 1027142, 1144340, 1183432, 1351220, 1373028	Failed cvupdatefs can prevent subsequent cvupdatefs run
SNFS	All	28591	1157710 1183706 1341218	Heavy buffered write activity completely stalls buffered reads
SNFS	All	28606	1045638, 1234184, 1311296, 909788	Failure in Inode_fl_scan due to inodes with InodeFlagFscckFree
SNFS	All	30641	1073338	cvadmin help doesn't list all multipath methods
SNFS	All	31532	1134538, 1308996	I/O retries cannot continue since the diskinfo is no longer valid
SNFS	All	32605	1198636	SGI IS600 RAID should be added to the default list of supported RAID types in the code
SNFS	All	32732	1365542	Truncating file to non-zero length results in TSM not retrieving stored file
SNFS	All	32840	1298900	SN 4.1 snfsdefrag with debug option gives IOCTL error
SNFS	All	33038	1261448	missing stripe group indices cause runtime hang
SNFS	All	33197	1241818	man page for mount_cvfs has wrong maximum for cachebufsize
SNFS	All	33447	1214916, 1318556	cvfsck loops trying to delete a directory entry
SNFS	All	33515	1267158	clients reconnect after cvfsck has modified metadata

Component	Operating System	Change Request Number	Service Request Number	Summary
SNFS	All	33647	1266468	snmetadump core dumps because of large ABM
SNFS	All	33713	1255358	need cvcp to be able to use -u and -z at the same time
SNFS	All	33713	1255358	need cvcp to be able to use -u and -z at the same time
SNFS	All	33726	298870, 1242764, 1315306, 1321626	Directory Quotas
SNFS	All	33727	n/a	meta-data ops start great and eventually tank
SNFS	All	33853	1272994	Logic to keep smesg log from logging too many messages not working in LogPathError()
SNFS	All	34352	1197140	Need a tool to detect btree inconsistencies
SNFS	All	34356	1287634	Unexplained RPL_Upgrade failure - use 'Checked Build' to be more robust
SNFS	All	34419	1291478, 1314356	Request for improvement of snmetadump man page / documentation
SNFS	All	34882	1267158	clients should not be able to reconnect after cvfsck has modified metadata
SNFS	All	35311	1295044	Using a debugger on a process waiting in the stornext client causes interrupted system call
SNFS	All	35356	1303678	Warning should be much more emphatic when journal size is too small
SNFS	All	35415	1311296, 1318638, 1412884	cvupdatefs fails when adding stripe group because fl_remove fails to find a free inode
SNFS	All	35446	1267158	document clients need to be unmounted when running cvfsck in repair mode
SNFS	All	35472	1314564, 1054688	Client licensed connection counters are incorrect
SNFS	All	35518	1321740	Code path overhead limits performance when small I/Os are used

Component	Operating System	Change Request Number	Service Request Number	Summary
SNFS	All	35533	1318404	snfsdefrag is unlinking the tmp file incorrectly.
SNFS	All	35573	1371020	Corruption of idi_nchildren with fs_racer
SNFS	All	35695	1314228	cvfsck hangs while clearing an inode with a corrupt extent
SNFS	All	35825	1390636	Windows client config tool displays buffer cache size incorrectly when size >= 4G
SNFS	All	35988	1321740	iozone errors found when running on Windows
SNFS	All	36008	1116778, 1314356	snmetadump -x hits ASSERT(lcb!=NULL) - don't assert, put out a msg.
SNFS	All	36131	n/a	The instperfreq.exe and instpxperfreq.exe tools and not documented
SNFS	All	36158	1416746	Extapi/CvApi library needs to be included on all platforms
SNFS	All	36428	1349924	buffer cache optimization: directly submit READ I/Os
SNFS	All	36556	1374336, 1375442	cvfsck segfaults if nt security descriptor inode is completely lost
SNFS	All	36593	1375442	cvfsck doesn't repair bad NT security index on big inode filesystem
SNFS	All	36684	1364546, 1399632	Platform Bug 27484 FSM runs out of connections due to deadlock
SNFS	All	36684	1364546, 1399632	DXi Bug 27484 FSM runs out of connections due to deadlock
SNFS	All	36729	1375442	cvfsck needs to run in verbose mode by default and always log verbose to trace file
SNFS	All	36731	1375442	cvfsck message about icb mismatch should direct customer to stop
SNFS	All	36792	1381662	cvfsck asserts in add_to_bad_name_list
SNFS	All	36995	1387182, 1391128, 1418302	hostname resolution delay led to FSM connection timeouts

Component	Operating System	Change Request Number	Service Request Number	Summary
SNFS	All	37138	1385988, 1396806	SnAdConConnect overwrites reason for FsmPMConnect failure
SNFS	All	37259	1378248	Remove warn away text in cvupdatefs manpage for rename and fix directory problem
SNFS	All	37881	1412216	DLC write performance is suboptimal when 10Gbit network interfaces are used on client
SNFS	All	36545	1372090, 1367328	No fsm restart with 2 core's in 1 hour
SNFS	All	36581	1372846, 1371886, 1399632	invalid inode messages emitted to cvlog when Apple clients are used
SNFS	Linux	27236	1179834	Thread pool exhausted for queue MsgDecodeQueueEVENTS
SNFS	Linux	35214	n/a	cvupdatefs -R Fails
SNFS	Linux	35344	1295044	Client panics purging inode that earlier encountered VerifyCvp() failure during cvfs_delete_inode()
SNFS	Linux	35622	1324652	SNFS client unnecessarily invalidating directory cache because of atime updates
SNFS	Linux	35700	1328880	PANIC: /usr/cvfs/bin/fsm "Inode not found in transaction"
SNFS	Linux	35863	1340610, 1362342, 1364548, 1372106	Messages file contains Failed to prune qustat archive
SNFS	Linux	37068	1356878	Configuring VLAN interfaces on DL gateways produces mount failures and unexpected results
SNFS	Linux	37279	1394706	HA can get into primary-primary mode
SNFS	Linux	37365	1420364	Rolling journal files are causing backups to fail
SNFS	Linux	37555	1402406, 1417180	HA: When in config mode, rebooting the secondary causes "snhamgr status" to hang
SNFS	Linux	36736	1364546, 1399632	Deadlock between free_pending_inode_thread and a thread calling InodeClose

Component	Operating System	Change Request Number	Service Request Number	Summary
SNFS	Mac OS	36527	n/a	Xsan 3.0 (Radar:10678636) zip containing an alias of a dir does not get restored correctly using Xsan 3.0 on Mac
SNFS	Mac OS	36587	n/a	Xsan 3.0(Radar 10722157)- Many Inode invalid type messages while creating file from Xsan 3.0 MDC
SNFS	Mac OS	37488	n/a	Xsan 3.0(Radar 11146558) -Xsan 3.0 volume not accessible from Xsan MDC after Windows client creates file
SNFS	Mac OS	37495	n/a	Xsan 3.0(Radar 11155117) - cversions for Xsan drop xsanfs-Zin-472.2.tgz displays Revision 4.2.1
SNFS	Mac OS	37561	n/a	Xsan 3.0(Radar 11174753) -Xsan 3.0 Client unable to connect/mount Xsan 3.0 MDC volumes.
SNFS	Other	35405	1271998	AIX kernel panics when doing /etc/rc.cvfs fullstop with I/O running
SNFS	Windows	26176	895984, 925846, 1084432, 1163864, 1318868	Modify snmetadump -r to allow restore of a FS to any disk geometry
SNFS	Windows	31120	1120006, 1146724, 1236692	Windows cannot follow Linux symbolic link to a directory
SNFS	Windows	32682	1217028, 1258150	Client are unable to find a domain controllers from time to time which leads to mount problems
SNFS	Windows	33828	1254042	F_GETRTIO_V3 external API command does not work on Windows
SNFS	Windows	35480	n/a	Enhance the Windows SAN Client LDAP search for dual AD forests
SNFS	Windows	36065	1217028, 1193804	Add logging for problem about Windows Server 2003 Clients will occasionally stall/hang
SNFS	Windows	36359	1318642	Windows fsmppm port being reallocated to other application. Resulting in a conflict and the DLC not connecting.

Component	Operating System	Change Request Number	Service Request Number	Summary
SNFS	Windows	36458	1363178	Windows users with read-only permissions are able to delete files from SNFS
SNFS	Windows	36859	n/a	CvfsPM service fails to start when LogOn is set to a non-local account
SNFS	Windows	36859	n/a	CvfsPM service fails to start when LogOn is set to a non-local account
SNFS	Windows	37197	1389820	Win: client-only: fsmppm generating license.dat files
SNFS	Windows	38505	1420194	Qustat: Windows Install of StorNext on non-default directory fails to create qustats directory
SNFS	Windows	36752	1374644	StorNext Quota enforcement is not working correctly
TSM	All	10992	1117786, 1329476	Address truncation policy inefficiencies especially for large amounts of candidates
TSM	All	15932	1049958, 1253946, 1329702	TSM: fsmedcopy fails due to internal timing issues
TSM	All	29201	1041212, 1222092, 1303038, 1378968	TSM checking of media usability marks the drives as cause of error
TSM	All	31261	1131828, 1138456, 1223964, 1224508	fs_fcopyman delays stores until SDISK cleanup is complete
TSM	All	34145	1287532	tape defragmentation does not allow specification of media/library/drivepool
TSM	All	34994	1344394	A database duplicate key error was encountered within table filecomp1. See TSM tac logs for more information
TSM	All	35406	1152424	TSM: fsmedcopy - No informative error given - only internal processing failure
TSM	All	35745	1319804	SAM-QFS ACU mover code fails to handle tape positioning error
TSM	All	37135	1385988, 1396806, 1398906	The fs_fmoverp is not handling a failing fsmppm connection correctly

Component	Operating System	Change Request Number	Service Request Number	Summary
TSM	All	37263	1394496, 1422456	Running "fsddmconfig -u -s e <active-MDC>" fails if DDM license count already reached
TSM	All	37748	1408724	scsi reservations made on Stornext luns when scsi2 in fs_sysparm
TSM	All	36588	n/a	Running multiple "fsmedinfo -l" might have caused an HA failover
TSM	All	37289	1390454	Drives remain in DISMOUNT state for extended periods of time due to lots of bmap retrieves
TSM	Linux	26377	all, 1049958	fsloglevel -s on -l tac -a hangs
TSM	Linux	35716	1297088	Configuring Delay Minimum File Size in GUI displays and may set incorrect value
TSM	Linux	32596	1209112	TSM fails to store large files (2TB) with DDM enabled.
TSM	Linux	36977	n/a	StorNext: Storage Manager: System backup failed: Store files to media failed
TSM	Other	32711	1213924, 1217774, 1288408, 1378572	Configuring MDCs by IP address to use DDM feature can cause movers to be used on MDCs

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

Other Changes and Considerations

StorNext Gateway Terminology

For the purposes of this document, we will use the following terminology:

StorNext Gateway Term	Description	Historical Customer-configured Gateway Equivalent Terminology
StorNext Gateway License*	A 'per gateway' license for StorNext software that enables LAN-based client connectivity to a StorNext SAN	N/A

StorNext Gateway Term	Description	Historical Customer-configured Gateway Equivalent Terminology
N/A	An individual 'per client' license for StorNext software that enables LAN-based client connectivity to a StorNext SAN	Proxy License; LAN Client License; DLC License; LAN Client License
StorNext Gateway	A StorNext Gateway is a StorNext SAN Client which allows LAN-based client connectivity to a StorNext File System.	Gateway Server; Server; LAN server, LAN-based server; DLC Gateway server; Clustered Gateway; DLC Gateway; DLS
StorNext LAN Client	A LAN-connected computer attached to a StorNext Gateway that has shared access to a StorNext SAN.	StorNext DLC
StorNext Gateway Metrics	A performance reporting and monitoring software module for StorNext Gateways.	N/A, newly created for StorNext Gateways

How the StorNext Gateway license is enabled depends on the current configuration:

- The StorNext G300 Gateway Appliance and the StorNext M660 Metadata Appliance have a "per Gateway" license model. This license allows you to add clients without having to purchase additional individual licenses.
- For new customers with no existing StorNext components, the license comes from the factory pre-installed and enabled for use with a new customer-configured MDC, StorNext M330 or StorNext M660.
- For customers with an existing StorNext M330, StorNext M660 or customer-configured MDC, if you choose to install the StorNext G300 Gateway Appliance in the same StorNext configuration as a customer-configured DLC gateway, you will be limited to the existing client DLC license count.

Note: For the StorNext G300 and StorNext M660, the license is populated on the gateway server. Also, to determine whether existing StorNext Gateway licenses are enabled, click the **Connected License Gateways** link at the bottom of the associated MDC StorNext license screen.

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). Beginning with StorNext 4.3, NFSv4 is also supported on a limited basis:

- NFSv4 is supported only when using Linux NFS servers and clients and has been tested with SLES 10, SLES 11, RHEL5, and RHEL6.
- NFSv4 delegations are not supported.

- NFSv4 ACLs are not supported.

Customers intending to use NFSv4 should be aware of the following administrative details (that are not specific to StorNext):

- 1 Exporting file systems over NFSv4 requires changes to `/etc/exports` due to the use of a hierarchical export tree.
- 2 To ensure that permissions work correctly, `idmapd` must be enabled and configured properly on NFS clients and servers. If this is not done, files will appear to be owned by the user `nobody`. Also note that `idmapd` may not cache id mappings. This can lead to poor performance unless the name server cache daemon (`nscd`) is also deployed.
- 3 Following a reboot, NFSv4 servers will delay startup for 90 seconds to allow for leases to expire. On some versions of Linux, this delay may be configured to a smaller value but not less than 10 seconds.

Consult the documentation for your version of Linux for additional specifics.

In addition to the above, when NFS servers running SLES10 or RHEL5 are used to export StorNext file systems with NFSv4, leases must be disabled system-wide by adding the following line to `/etc/sysconfig.conf`:

```
fs.leases-enable=0
```

Note: If the Samba daemon (`smbd`) is also running on the server, disabling leases will prevent the use of the Samba “kernel oplocks” feature. In addition, any other application that make use of leases through the `F_SETLEASE` `fcntl` will be affected. Leases do not need to be disabled in this manner for NFS servers running SLES11 or RHEL6.

Due to the additional complexity of NFSv4, you are advised to continue to use NFSv3 unless you have a specific need.

Note: NFS clients running some newer versions of Linux (such as RHEL6.1, RHEL6.2, and corresponding CentOS releases) may attempt to use NFSv4 by default. To force the use of NFSv3 from the client side, use the “`vers=3`” mount option. Alternatively, NFSv4 can typically be disabled on the server side by applying the following setting to the file `/etc/sysconfig/nfs`:

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

This setting will affect all exported file systems, not just StorNext volumes.

Refer to the *StorNext File System Tuning Guide* for additional 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 4.3.0, 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 4.3.0, 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 4.3.0.

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

StorNext Compatibility

For information on StorNext 4.3.x compatibility with operating systems, kernel versions, hardware platforms, drives, libraries, StorNext Appliances, StorNext client interoperability, and other compatibility items, see the "StorNext Compatibility Guide" posted here:

[http://www.quantum.com/ServiceandSupport/
SoftwareandDocumentationDownloads/SNMS/Index.aspx](http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx)

Note: The compatibility guides on this Web page contain information for previous and current StorNext releases, so be sure you locate and view or download the version for StorNext 4.3.x.

SNAPI, Partial File Retrieval, and Apple Xsan compatibility information is provided in separate documents.

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>

Supported StorNext Upgrade Paths

Sites running the following StorNext versions may upgrade directly to this release assuming the platform, service pack, architecture (32-bit or 64-bit), and StorNext component(s) are supported in this:

- StorNext 4.1.0
- StorNext 4.1.1
- StorNext 4.1.2
- StorNext 4.1.3
- StorNext 4.2.0
- StorNext 4.2.1
- StorNext 4.2.1.0.1
- StorNext 4.2.2
- StorNext 4.2.2.0.1 (appliance only release)
- StorNext 4.3.0
- StorNext 4.3.1
- StorNext 4.3.2-LTO-6

All other versions of StorNext require additional steps to upgrade to this release.

Upgrade Guidelines

Before upgrading StorNext, review the following upgrade guidelines:

- Beginning with 4.3.0, StorNext uses a MySQL database for information needed internally by Storage Manager. The MySQL database requires:
 - 1 A minimum of 400MB of RAM.
 - 2 Approximately 2.5 times the amount of disk space of a current Linter database for both the converted database and backups.
 - 3 During migration, up to 4 times the amount of disk space of the Linter database.

Note: Plan ahead for the increased disk space requirements during the migration of your Linter database, and for the new MySQL database and backups.

- Additionally, StorNext 4.3 has the following additional disk space requirements:

Table 1 Space requirements from snPreInstall

Component	Directory Space Required StorNext 4.3.0	Directory Space Required StorNext 4.2.0 (for comparison)
Database	231 MB	64 MB
Journal	11 MB	3 MB
Mapping	157 MB	157 MB
Metadata	333 MB	195 MB
MySQL_tmp	120 MB	N/A
Total	852 MB	419 MB

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 4.3.0.1:

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx>

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 StorNext 4.3.0:

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/SNMS/Index.aspx>

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 or higher.

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

Table 2 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/4.3.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/4.3.x with Dedup.	✓	Replication rejected	✓		✓	

Note: StorNext 4.0 to 4.0.1 and 4.1.x/4.2.x/4.3.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 4.3.0, note the following configuration requirements:

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

Caution: If a Library used by StorNext Storage Manager is connected via a fibre switch, zone the switch to allow only the system(s) running SNSM to have access to the library 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 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 4.3.0, the following hardware requirements must be met:

- [StorNext File System and Storage Manager Requirements](#) on page 33
- [StorNext Client Software Requirements](#) on page 35

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 3](#).

Table 3 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 4.3.0, your system must have the following memory and disk capacity **in addition to** the base memory and disk capacity required to run StorNext File System and Storage Manager.

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 StorNext LAN Clients:

- 1 GB RAM
- 500 MB available hard disk space

For SAN clients acting as a StorNext Gateway Server:

- 2 GB RAM
- 500 MB available hard disk space

Note: Gateway servers may require additional RAM depending on the number of file systems, StorNext LAN Clients, and NICs used. See [Gateway Server Memory Tuning](#) in the StorNext User's Guide for 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.

M660 and M440 Memory Requirements

The M660 and M440 memory allocation settings that are supported for use with StorNext 4.3.2 are presented in [Table 4](#).

Table 4 M660 nd M440
Memory Allocation Settings

	FSM BufferCacheSize	FSM InodeCacheSize	MySQL innodb_buffer_pool_size	Reserved for General Usage	Total
M660	8 GB per file system (e.g. 8 x 8 GB = 64 GB)	512 K inodes per file system (e.g. 8 x ~1 GB = 8 GB)	40 GB	32 GB	144 GB
M440	2 GB per file system (e.g. 4 x 2 GB = 8 GB)	256 K inodes per file system (e.g. 4 x ~0.5 GB = 2 GB)	10 GB	28 GB	48 GB

Supported System Components

System components that are supported for use with StorNext 4.3.0 are presented in [Table 5](#).

Table 5 StorNext Supported
System Components

Component	Description
Tested Browsers	Internet Explorer versions 7 through 9, Quantum recommends IE9 Mozilla Firefox versions 4 through 11, Quantum recommends Firefox 11 Chrome version 18 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?whattab=Fourth>

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

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

Known Issues

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

- [StorNext File System Known Issues](#) on page 39
- [StorNext Storage Manager Known Issues](#) on page 41
- [StorNext GUI Known Issues](#) on page 43
- [StorNext Installation, Replication, HA and Other Known Issues](#) on page 47

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 6](#) lists known issues specific to StorNext File System.

Table 6 StorNext File System
Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround (if applicable)
All	39298	n/a	An FSM panic may occur when real-time I/O is configured and SNFS clients request more real-time I/Os than are reserved by the RtiosReserved setting.	To work around this issue, increase the RtiosReserved to a value larger than the peak real-time I/Os per second expected for the file system.

Operating System	Change Request Number	Service Request Number	Description	Workaround (if applicable)
Linux	40756	n/a	NFS clients with mounted StorNext file systems may experience "stale NFS file handle" errors if the NFS server is rebooted.	<p>Restarting NFS server processes will resolve the issue until the next reboot. Remounting the file system on NFS clients also serves as a temporary workaround. The problem can be avoided permanently by the following:</p> <ol style="list-style-type: none"> 1 Reordering service startup on the NFS server so that the "cvfs" service is started before NFS server processes. Refer to the Linux chkconfig(8) man-page for details. When making this change, note that if the "cvfsbuild" or "snhamgr" services are configured, they must continue to start before the "cvfs" service. 2 Implementing the workaround for RedHat Bug 645632: https://bugzilla.redhat.com/show_bug.cgi?id=645632 (Edit /etc/init.d/nfs and replace instances of "/var/lock/subsys/nfsd" with "/var/lock/subsys/nfs").

StorNext Storage Manager Known Issues

[Table 7](#) lists known issues specific to StorNext Storage Manager.

Table 7 StorNext Storage
Manager Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround (if applicable)
Linux	40169	1482638	Prior to this release, checksum information for files being copied by fsmedcopy or fsfilecopy would not be propagated to the new media. This only affects the copy of the file which was being operated on and not any other copies of the file. For example, if a fsmedcopy command was invoked on a copy 1 media, then the checksum information for copy 2, 3 and 4 would still be intact.	<p>In order for the changes to be effective, upgrade the clients that are used for Distributed Data Movers.</p> <p>In order to restore checksums for files on the transcribed media, there are a couple of alternatives that could be used:</p> <ul style="list-style-type: none"> • Perform another fsmedcopy operation with the new -G option which would generate checksums for the files which are being transcribed. This option will override any policy class or system parameter settings. It will not generate a new checksum for files which already have them. The -G option was also added to the fsfilecopy command. • The files could be retrieved back to disk, then the copy of the file that is missing a checksum could be removed using fsrmcopy -c. When that copy of the file is stored again a new checksum will be generated.
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	Change Request Number	Service Request Number	Description	Workaround (if applicable)
Linux	36141	1400700, 1412946	Warning message about missing rpm dependencies may occur after the <code>install.stornext</code> script is run.	As of 4.3.0 the perl-DBI rpm is required to be installed prior to running <code>install.stornext</code> . Consult your operating system's package repository for the perl-DBI package if it is not already installed on the system.
All	38632	n/a	A core is observed with <code>fspolicy</code> and <code>fs_cpyresp</code> . It occurs when the processes are exiting under normal conditions.	<p>No workaround is available since it does not effect runtime behavior.</p> <p>The impact on appliances: The cores are detected by the <code>trimcores</code> utility which will generate RAS tickets for the cores. Disk space consumption is not an issue for the appliances because the <code>trimcores</code> utility processes the core files.</p> <p>The impact on other StorNext installations: The cores are undetected unless someone is physically looking for them, and may impact disk space consumption over time. Fortunately, this does not occur frequently, and will require time for disk space to be affected.</p>

Operating System	Change Request Number	Service Request Number	Description	Workaround (if applicable)
All	40803	n/a	The StorNext GUI does not auto-configure media correctly when a library contains mixed media types.	<p>The workaround is to follow the steps below to manually add any media which fail auto-configuration:</p> <ol style="list-style-type: none"> 1 Identify the incorrectly configured media using the command: vsmedqry -a Media listed as belonging to the ADDBLANK media class will need to be manually added to the Tertiary Manager system. 2 Use command fsmedin to add failed media to the Tertiary Manager system: fsmedin -b -w mediaId... <p>Note: Only media of the same type can be specified using a single command.</p>

StorNext GUI Known Issues

[Table 8](#) lists known issues specific to the StorNext GUI process.

Table 8 StorNext GUI Known Issues

Operating System	Change Request Number	Service Request 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.
Linux	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.

Operating System	Change Request Number	Service Request 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.
Windows	36764	n/a	On IE7, the gateway report only renders the first time. If you navigate away from the report and come back, the report does not render and creates JavaScript errors. IE8 also has JavaScript errors, but it does render.	The workaround is to clear the browser and cookie cache after a SN software upgrade.
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> .
All	38414	n/a	In the GUI, the reset button on the Configuration > System -> Network tab doesn't always restore current network configuration values.	When this happens, navigate away from the page by selecting any other menu item and navigate back again to reload the current network settings.
All	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	Change Request Number	Service Request Number	Description	Workaround (if applicable)
Linux	37538	1398524	GUI is unable to down a stripe group when LUNs are unavailable	Mark stripe groups down in the GUI before taking the stripe group's disks offline. If that is not possible, set the stripe group down directly through the FSM configuration file and restart the FSM. See the snfs_config(5) man page or the MAN Pages Reference Guide for details.
Linux	32553	n/a	<p>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.</p> <p>A general estimate is that the delay could be approximately 10 seconds for every 500 media requiring attention on the LOI page.</p>	<p>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.</p> <p>In that situation, you can avoid this issue by breaking up media action requests into smaller chunks.</p>

Operating System	Change Request Number	Service Request 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>
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>

StorNext Installation, Replication, HA and Other Known Issues

[Table 9](#) lists known issues specific to StorNext installations, data replication, HA systems and other areas.

Table 9 StorNext Installation, Replication, HA and Other Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround
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.
Linux	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.
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.
Linux	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.
Linux	35969	n/a	The Linter database is not removed from the HaShared file system after upgrading to 4.3.0 in HA environments.	After successful upgrade to 4.3.0, manually remove /usr/adic/HAM/shared/database/db and /usr/adic/HAM/shared/database/journal. Also, if present, remove the symbolic links /usr/adic/database/db and /usr/adic/database/journal.

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	39034	n/a	<p>The section titled “Performing a test conversion of a Linter database to MySQL”, on page 9 of the StorNext 4.3.x Upgrade Guide, contains the following incorrect information:</p> <p>On the Migration System Install StorNext 4.3.x onto a clean, pristine system that will be used to migrate Linter into MySQL. The version of StorNext must match the version of StorNext installed on the production system.</p> <pre>cd /path/to/4.3.x/phdist/ SYSTYPE ./install.stornext . /usr/adic/.profile</pre>	<p>It is critical that on the external system, the user installs a previous version of StorNext - 4.2.x which matches their current system's version. The correct information is:</p> <p>On the Migration System Install StorNext 4.2.x onto a clean, pristine system that will be used to migrate Linter into MySQL. The version of StorNext must match the version of StorNext installed on the production system.</p> <pre>cd /path/to/4.2.x/phdist/ SYSTYPE ./install.stornext . /usr/adic/.profile</pre>

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	40812	1496086	<p><code>install.stornext</code> can fail during an upgrade to StorNext 4.3.x, while installing the <code>mysql</code> component, if the directory permissions are too restrictive.</p>	<p>The MySQL daemon runs as the <code>quantumdb</code> user which is a member of the <code>adic</code> group. Make sure that the <code>adic</code> group is capable of read and execute permissions on the following directories:</p> <ul style="list-style-type: none"> <code>/usr</code> <code>/usr/adic</code> <code>/usr/adic/mysql</code> <code>/usr/adic/HAM/shared/mysql</code> (In the case of HA) <p>In addition to this, the <code>quantumdb</code> user needs to have read, write and execute permission to the following directories:</p> <ul style="list-style-type: none"> <code>/usr/adic/mysql/db</code> <code>/usr/adic/mysql/tmp</code> <code>/usr/adic/mysql/journal</code> <p>The <code>install.stornext</code> script may place these directories in different locations on the system and link to them via symbolic links. If any of these directories are a symbolic link to another directory, the full path to that directory needs to have read and execute permission by the <code>adic</code> group.</p> <p>If <code>install.stornext</code> fails to install the <code>mysql</code> component due to incorrect directory permissions, correct the permissions and try again. If <code>install.stornext</code> continues to fail and you are upgrading from StorNext 4.2.1 or earlier, uninstall the <code>mysql</code> component by running <code>./install.stornext -remove -mysql</code> and then re-run <code>install.stornext</code>.</p>

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.
All	37757	n/a	StorNext is unable to automatically propagate geometry changes made to the HaShared file system from the primary system to the secondary system. Updates made to the HaShared file system while the cluster is in config mode need to be manually copied over to the secondary system. Failure to do so may result in a failure to start StorNext services on the secondary system.	Manually copy the <code>.cfgx</code> file from <code>/usr/cvfs/confid/HASHAREDFS.cfgx</code> on the primary to <code>/usr/cvfs/confid/HASHAREDFS.cfgx</code> on the secondary where HASHAREDFS is the name of the HaShared file system.

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.

This release does not include any known issues related to Multilink.

In addition to any listed 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 10](#) lists operating guidelines and limitations for running StorNext. Items are grouped according to operating system.

Table 10 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>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	Migrating metadata controllers	<p>StorNext users migrating their metadata controllers from Apple Xsan to Linux must disable Spotlight indexing through the SNFS config file. See the "enableSpotlight" section in the snfs_config man or help page.</p>
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>
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>

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>
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>
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>.

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>
All	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.
All	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 "FSBlockSize, Metadata Disk Size, and JournalSize Settings" in Appendix A of the <i>StorNext User's Guide</i>.</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>

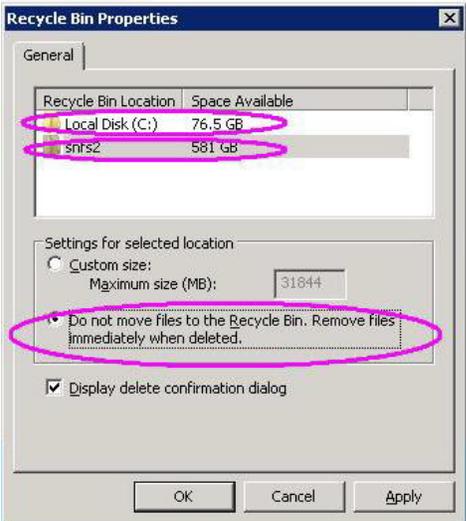
Operating System	Feature or Category	Description
Linux	StorNext File System	<p>StorNext File System does not support the Linux <code>sendfile()</code> system call.</p> <p>This issue causes Apache web servers to deliver blank pages when content resides on StorNext file systems.</p> <p>This issue also 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>
Linux	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</p> <pre>/usr/adic/install/.snsm_ha_configured</pre> <p>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>
Linux	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 Change Request 468088 and SuSE 10 Change Request 10440734121.</p>

Operating System	Feature or Category	Description
Linux	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>
Linux	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 <i>StorNext Tuning Guide</i>.</p>
Linux	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	fsclean command	<p>The behavior of the fsclean command has been modified. The command which is used to clean up the TSM database for removed files and old versions of files basically has two portions:</p> <ol style="list-style-type: none"> 1 The file component and media cleanup, and the purging of the database namespace for these old files. 2 The namespace purge has been made an optional portion of the command. <p>The fsclean command modification was done for the following reasons:</p> <ul style="list-style-type: none"> • The file component and media cleanup that is needed for a media to be made blank, to recoup license space, etc. This functionality is not dependent on the namespace purge. • The purge of the namespace was the most time consuming portion of the command. It was thought there was no need to make a user wait on this portion unless desired. <p>There is a new option to the fsclean command, "-P" which indicates the purge of the namespace should be performed. If this is missing, then the namespace purge is bypassed.</p> <p>Note: The regularly scheduled (weekly) "cInver" feature was updated to do the purge as a part of its normal processing. If there are clean operations run during a week with no purge, the database will be cleaned up at this regular time.</p> <p>If you remove this scheduled item, or lock it out then you should be using the "-P" on fsclean commands that are run by other means. You do not have to run it on every command, but it will be needed on a regular basis.</p> <p>With this new behavior for fsclean, the one thing that will likely be noticed is that fsrecover will report files as being available to recover even though the containing media have been cleaned. Once the next fsclean with the "-P" provided is run, the fsrecover report will be correct. If you do happen to attempt a recovery of one or more of these files that are actually not available, the fsrecover command will report missing component information for these files.</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>
Windows	Anti-virus software	<p>Virus-checking software can severely degrade the performance of any file system, including a StorNext file system. If you have anti-virus software (McAfee, etc.) running on a Windows machine, Quantum recommends you configure the software to not check (ignore) 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. 
Windows	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
All	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.
All	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>
All	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>
All	StorNext home page	<p>The capacity indicators on the StorNext home page provide <i>approximations</i> 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>
All	Backups	<p>Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.</p>
All	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>
All	Cluster-Wide Central Control	<p>The StorNext Cluster-Wide Central Control file (nss_cctl.xml) is used to enforce the cluster-wide security control on StorNext nodes (client nodes, fsm nodes, and nodes running cvadmin). This file is placed on an nss coordinator server.</p> <p>Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.</p>
All	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>

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 4.2.1.0.1 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>
All	Replication/ Deduplication	<p>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.</p>
All	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	Deduplication	<p>If a deduplication candidate is removed before blockpool processing is completed, errors such as the following may be sent to the syslog:</p> <pre>Oct 2 15:22:00 orleans Blockpool[16403]: E: [5] (Store Local) Error storing file "/stornext/source/ __CVFS_Handle.000474F892EBB65E000E000000000000000000000002 92BF4". Error opening file "/stornext/source/ __CVFS_Handle.000474F892EBB65E000E000000000000000000000002 92BF4". No such file or directory.</pre> <p>Errors such as these may appear serious, but there is no reason for concern.</p> <p>If you receive these errors, no action is required.</p>
All	Deduplication	<p>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).</p>
All	HA	<p>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.</p> <p>Prepare systems and users for this eventuality before entering HA Config mode.</p>
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 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 fsmlist" 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>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>
All	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>
All	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>
All	StorNext LAN	<p>StorNext 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 StorNext Gateway Server is configured after converting to HA, the file system(s) running as Gateway servers must be unmounted and mounted again to service StorNext LAN 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.
All	dpserver	In some cases the physical IP address must be included in the dpserver file in addition to the interface name. Note these conditions: <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 "StorNext LAN on page 63.")
All	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.
All	Affinities	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. 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 <i>aff1</i> and <i>aff2</i> , the GUI would restrict any new file systems to have those affinity names as opposed to <i>tier1</i> and <i>tier2</i> .
All	Converting file systems	StorNext does not currently support converting from a managed file system to an unmanaged file system.

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 the Quantum Technical Assistance Center:

North America	+1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+00-800-7826-8888 (toll free) +49-6131-3241-1164
APAC	+1-800-7826-8887 (toll free) +603-7953-3010
Online Service and Support	www.quantum.com/OSR
Worldwide Web	www.quantum.com/ServiceandSupport

(Local numbers for specific countries are listed on the Quantum Service and Support Website.)