

# StorNext 4.2 Release Notes

<b>Product</b>	StorNext® 4.2
<b>Date</b>	September 2011

## Contents

Purpose of this Release . . . . .	3
Other Changes and Considerations . . . . .	5
Supported Platforms, Libraries and Drives. . . . .	6
Supported StorNext Upgrade Paths. . . . .	8
Compatibility With Other StorNext Components and Features . . . . .	9

© 2011 Quantum Corporation. All rights reserved.

6-67376-01 Rev A, September 2011

Quantum, the Quantum logo, DLT, DLTtape, the DLTtape logo, Scalar, StorNext, the DLT logo, DXi, GoVault, SDLT, StorageCare, Super DLTtape, and SuperLoader are registered trademarks of Quantum Corporation in the U.S. and other countries. Protected by Pending and Issued U.S. and Foreign Patents, including U.S. Patent No. 5,990,810. LTO and Ultrium are trademarks of HP, IBM, and Quantum in the U.S. and other countries. All other trademarks are the property of their respective companies. Specifications are subject to change without notice.

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

ACSAPI, copyright © Storage Technology Corporation

Java, copyright Oracle Corporation

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

LibX11copyright The Open Group, MIT, Silicon Graphics, and the Regents of the University of California, and

copyright (C) 1994-2002 The XFree86 Project, Inc. All Rights Reserved. And

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

Libxml2 and LibXdmcp, copyright MIT

Linter, copyright © Relex Software Corporation

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

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

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

vixie-cron: copyright Internet Systems Consortium (ISC)

Wxp-tdi.h, copyright © Microsoft Corporation

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



Configuration Requirements . . . . .	10
Hardware Requirements . . . . .	13
Supported System Components . . . . .	16
Previous Versions of Release Notes . . . . .	18
Resolved Issues . . . . .	18
Known Issues . . . . .	30
Operating Guidelines and Limitations . . . . .	45
Documentation . . . . .	58
Contacting Quantum . . . . .	58

---

## Purpose of this Release

StorNext 4.2 contains the new features and enhancements listed below, as well as bug fixes described in the [Resolved Issues](#) section.

Visit [www.quantum.com/ServiceandSupport](http://www.quantum.com/ServiceandSupport) for additional information and updates for StorNext.

---

### Allocation Session Reservation

---

StorNext 4.2 adds a new method of allocating space to regular files. *Allocation Session Reservation* (ASR) optimizes on-disk allocation behavior in workflows (such as some rich media streaming applications) which write and read sequences of files of certain sizes in specific directories.

With ASR enabled, file sequences in a directory are usually placed on disk based on the order in which they are written. ASR keeps these files together even if other applications are writing at the same time in different directories or from different StorNext clients. The feature also has the ability to reduce file system free space fragmentation since collections of files that are written together are then typically removed together.

For more information about ASR, refer to the *StorNext File System Tuning Guide*.

---

### WS-API Enhancement

---

Web Services API has been enhanced to allow XML and JSON output for the `fsfileinfo` command.

---

### StorNext GUI Enhancements

---

The following improvements have been added to the StorNext graphical user interface for release 4.2:

- **User Accounts:** When adding new StorNext users you can now select from one of three user profiles with typical permissions predefined: StorNext Admin, StorNext Operator, or StorNext General User. Each user profile contains tasks typically performed by a person in that role, and you can easily add or remove permission for tasks to meet your specific needs.
- **New Reports:** Three new reports have been added to StorNext 4.2:
  - **Files Report:** This report provides general information about selected files, as well as specific details if you require more granular information.
  - **Drives Report:** This report provides a list of drives in your system and enables you to view details about selected drives.
  - **Relations Report:** This report shows the pathname of the managed file system's directory and the corresponding policy class name.
- **Logging:** The Service menu's new Logging option is a robust debugging tool which enables you to turn on tracing for various system components. The feature is useful if you have been asked by Quantum Service personnel to enable debugging for one or more components in order to help them identify and diagnose a particular error.

When logging (debugging) is enabled, information is copied to the same location as regular log files.

- **Capture State for Secondary Nodes:** The StorNext Capture State feature has been expanded to allow capturing information for secondary nodes in a high availability (HA) system.
- **Advanced Reporting Integration:** If you have purchased StorNext Advanced Reporting, you can now access this feature from the StorNext Tools menu.
- **Expanded Truncation Time Parameter:** When you create or edit a storage policy you can now specify the file age before truncation value in seconds, hours, minutes or days.
- **Copy Number and Percentage Used:** These headings have been added to the Media Actions screen, showing you the policy class number on the media and the percentage of the media which is currently used.
- **Improved Default Overriding:** When creating a replication policy, you can now easily override the default values by deselecting the desired box under the **Use Defaults** heading and then manually entering the values.
- **Moving Media Manually:** The Media Action screen now gives you the ability to manually move media from one library to another.

---

## New Supported Platforms and Libraries

---

StorNext 4.2 adds or expands support for the following platforms:

- Red Hat 5 Update 6
- Red Hat 6 and Red Hat 6.1
- SUSE Linux 11 SP1
- SUSE Linux 10 SP4
- Windows 2008 Server R2 SP1
- IBM AIX 7

---

**Note:** StorNext 4.2 has been qualified on AIX 7 using the AIX 6 build. Customers wanting to install StorNext 4.2 on AIX 7 should install the AIX 6 binaries.

---

StorNext 4.2 also adds support for the Quantum DXi 8500 library running in i2000 mode.

For detailed information including which StorNext components are supported for the above operating systems, refer to the "StorNext Supported Platforms" document for StorNext 4.2 available at this location:

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

---

## Xsan Compatibility

---

StorNext 4.2 adds compatibility with Apple Xsan 2.2.2 and 2.3. For more information, refer to the document "StorNext and Apple Xsan Compatibility" available at this location:

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

---

## SNAPI 2.0.2 compatible with 4.2

---

StorNext 4.2 adds support for StorNext API (SNAPI) release 2.0.2, which contains new APIs and other updates. Support SNAPI 2.0.1 is maintained in StorNext 4.2.

---

# Other Changes and Considerations

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

---

## StorNext Upgrades to 4.2

---

StorNext 4.2 supports direct upgrades *only* from StorNext 4.1.x. For more information about upgrade paths, see [Supported StorNext Upgrade Paths](#) on page 8.

---

## Upgrading AIX, HP-UX and Solaris NFS Servers

---

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

---

## Network File System (NFS) Support in StorNext

---

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

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

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

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

---

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

---

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

---

## Red Hat 6 and Virtual Memory

---

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

Red Hat 6 has changed the memory allocator for multi-threaded programs so that each thread can (and many times does) receive its own heap arena. This memory is mapped into virtual memory in large chunks. However, the actual memory allocated by the thread still uses roughly the same amount of memory, as it would have if the heap arenas were shared. The increase in virtual memory size between Linux releases can be ignored.

---

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

---

---

## Replication Report Location on HA Systems

---

Starting with StorNext 4.2, 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 created whenever replication is performed.

If replication was performed prior to upgrading to 4.2.0, 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 4.2.

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

---

## Supported Platforms, Libraries and Drives

The operating systems, kernel versions, and hardware platforms supported by this StorNext 4.2 release are listed in the "StorNext Supported Platforms" document posted at this location:

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

---

**Note:** The compatibility guides on this Web page show the supported platforms, libraries and drives for previous and current StorNext releases, so be sure you locate and view or download the version for StorNext 4.2.

---

This document also shows StorNext client interoperability. Apple Xsan compatibility information is provided in a separate document.

---

## Caveats

---

This section contains some notes and cautions related to platforms which apply to StorNext 4.2 supported platforms.

---

**Note:** When adding StorNext Storage Manager to a StorNext File System environment, the metadata controller (MDC) must be moved to a supported platform. If you attempt to install and run a StorNext 4.2 server that is not supported, you do so at your own risk. Quantum strongly recommends against installing non-supported servers.

---

---

**Note:** For systems running Red Hat Enterprise Linux version 5 or 6, before installing StorNext you must first install the following kernel files:

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

For systems running SUSE Linux Enterprise Server, you must first install the kernel source code (typically shipped as the kernel-source RPM).

---

---

---

**Caution:** Red Hat 5 and 6 ship with Security-Enhanced Linux (selinux) enabled by default. To ensure proper StorNext operation, you must not install Red Hat 5 or 6 with selinux enabled. That is, selinux must be off, or the file system could fail to start.

If Red Hat 5 or 6 have already been installed with SELINUX enabled, edit the file `/etc/selinux/config` and change the line with `SELINUX=enforcing` or `SELINUX=permissive` to `SELINUX=disabled`. Refer to Red Hat 5 or 6 documentation for more information.

---

---

---

**Note:** Linux places both an IPV4 and an IPV6 address in the `/etc/hosts` file, but for StorNext the `/etc/hosts` file must contain **only** the IPV4 loopback.

---

---

**Note:** GNU tar is required on Solaris systems. In addition, for systems running Solaris 10, install the Recommended Patch Cluster (dated March 10, 2006 or later) before installing StorNext.

To enable support for LUNs greater than 2TB on Solaris 10, the following patches are required:

- 118822-23 (or greater) Kernel Patch
  - 118996-03 (or greater) Format Patch
  - 119374-07 (or greater) SD and SDD Patch
  - 120998-01 (or greater) SD Headers Patch
- 

---

## Quantum OS Upgrade Support Policy

---

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

---

## Supported StorNext Upgrade Paths

In general, sites running the following StorNext versions may upgrade directly to StorNext 4.2, assuming that the platform, service pack, architecture (32-bit or 64-bit), and StorNext component are supported in the installed StorNext version and in StorNext 4.2:

- StorNext 4.1
- StorNext 4.1.1
- StorNext 4.1.2

All other versions of StorNext require additional steps to upgrade to StorNext 4.2.

## Compatibility With Other StorNext Components and Features

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

### Data Replication Compatibility

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

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

Table 1 Replication Compatibility Between Releases

Source Release	Target Release							
	4.0 without Dedup.	4.0 with Dedup.	4.0.1 without Dedup.	4.0.1 with Dedup.	4.1/4.2 without Dedup.	4.1/4.2 with Dedup.	4.1.1/4.1.2 without Dedup.	4.1.1/4.1.2 with Dedup.
4.0 without Dedup.	✓	--	See note below	--	See note below	--	See note below	--
4.0 with Dedup.	--	✓	--	Replication rejected	--	Replication rejected	--	Replication rejected
4.0.1 without Dedup.	Replication rejected	--	✓	--	✓	--	✓	--
4.0.1 with Dedup.	--	Replication rejected	--	✓	--	✓	--	✓
4.1/4.2 without Dedup.	Replication rejected	--	✓	--	✓	--	✓	--
4.1/4.2 with Dedup.	--	Replication rejected	--	✓	--	✓	--	✓
4.1.1/4.1.2 without Dedup.	Replication rejected	--	✓	--	✓	--	✓	--

Source Release	Target Release							
	4.0 without Dedup.	4.0 with Dedup.	4.0.1 without Dedup.	4.0.1 with Dedup.	4.1/4.2 without Dedup.	4.1/4.2 with Dedup.	4.1.1/4.1.2 without Dedup.	4.1.1/4.1.2 with Dedup.
4.1.1/4.1.2 with Dedup.	--	Replication rejected	--	✓	--	✓	--	✓

**Note:** StorNext 4.0 to 4.0.1 and 4.1/4.1.1/4.1.2/4.2 without deduplication successfully transfers content, but there may be synchronization issues in situations where content was changed very recently before a scheduled replication. If this occurs, the file may initially appear to be transferred without content, but the content will eventually be transferred in its entirety.

### Partial File Retrieval

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

StorNext Partial File Retrieval releases 1.0 and 1.0.1 are compatible with StorNext 4.2. Upon release, PFR release 1.0.2 will also be compatible with StorNext 4.2.

### StorNext API (SNAPI)

StorNext API (SNAPI) versions 2.0.1 and 2.0.2 are compatible with StorNext 4.2.

## Configuration Requirements

Before installing StorNext 4.2, note the following configuration requirements:

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

- For Windows systems (server and client), the operating system must always be run in U.S. English.

---

**Caution:** If a Library used by StorNext Storage Manager is connected via a fibre switch, zone the switch to allow only the system(s) running SNSM to have access to the library. This is necessary to ensure that a “rogue” system does not communicate with the library and cause data loss or corruption. For more information, see StorNext Product Alert 16.

---

---

## Time Synchronization for Replication and Deduplication

---

If you plan to use the Replication or Deduplication features, ensure that the time on your file system clients is synchronized to your metadata controllers.

The age values for Deduplication and Truncation are based on the clients’ time, so if your clients’ time is different from the MDC’s time you may see files ingested earlier or later than you’ve configured.

---

## Disk Naming Requirements

---

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

---

## SAN Disks on Windows Server 2008

---

SAN policy has been introduced in Windows Server 2008 to protect shared disks accessed by multiple servers. The first time the server sees the disk it will be offline, so StorNext is prevented from using or labeling the disk.

To bring the disks online, use the `POLICY=OnlineAll` setting. If this doesn’t set the disks online after a reboot, you may need to go to Windows Disk Management and set each disk online.

Follow these steps to set all disks online:

- 1 From the command prompt, type **DISKPART**
- 2 Type **SAN** to view the current SAN policy of the disks.
- 3 To set all the disks online, type **SAN POLICY=onlineall**.
- 4 After being brought online once, the disks should stay online after rebooting.
- 5 If the disks appear as “Not Initialized” in Windows Disk Management after a reboot, this indicates the disks are ready for use.

If the disks still appear as offline in Disk Management after rebooting, you must set each disk online by right-clicking the disk and selecting **Online**. This should always leave the SAN disks online after reboot.

---

**Note:** If the disks are shared among servers, above steps may lead to data corruption. Users are encouraged to use the proper SAN policy to protect data

---

6 Perform the following for each disk manifesting the problem:

diskpart.exe:

select disk X

attribute disk clear readonly

EXAMPLE:

```
C:\ >Diskpart
```

```
Microsoft DiskPart version 6.0.6001
```

```
Copyright (C) 1999-2007 Microsoft Corporation.
```

```
On computer: CALIFORNIA
```

```
DISKPART> SAN
```

```
SAN Policy : Offline All
```

```
DISKPART> san policy=onlineall
```

```
DiskPart successfully changed the SAN policy for the current  
operating system.
```

---

## Configuring VMWare for StorNext

---

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

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

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

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

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

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

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

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

## Hardware Requirements

To successfully install StorNext 4.2, the following hardware requirements must be met:

- [StorNext File System and Storage Manager Requirements](#) on page 13
- [StorNext Client Software Requirements](#) on page 14

---

**Note:** The following requirements are for running StorNext only. Running additional software (including the StorNext client software) requires additional RAM and disk space.

---

### StorNext File System and Storage Manager Requirements

The hardware requirements for StorNext File System and Storage Manager are presented in [Table 2](#).

Table 2 File System and Storage Manager Hardware Requirements

No. of File Systems	RAM	File System Disk Space	Storage Manager Disk Space
1–4*	4 GB	2 GB	<ul style="list-style-type: none"> <li>• For application binaries, log files, and documentation: up to 30GB (depending on system activity)</li> <li>• For support directories: 3 GB per million files stored</li> </ul>
5–8**	8 GB	4 GB	

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

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

### Additional Memory and Disk Requirements for Deduplication and Replication

In order to use the data deduplication and replication features in StorNext 4.2, your system must have the following memory and disk capacity **in addition to** the base memory and disk capacity required to run StorNext File System and Storage Manager.

---

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

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

---

### Minimum Additional Disk and Memory Required for a Replication License

- 50 MB available hard disk space

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

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

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

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

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

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

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

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

### Minimum Additional Memory on Replication Target Machines

On replication target machines, Quantum recommends adding a minimum of 1GB additional memory per 22 million files replicated, which is approximately 4 bytes per entry.

---

## StorNext Client Software Requirements

---

To install and run the StorNext client software, the client system must meet the following minimum hardware requirements.

For SAN (FC-attached) clients or for Distributed LAN Clients:

- 1 GB RAM
- 500 MB available hard disk space

For SAN clients acting as a Distributed LAN Server:

- 2 GB RAM
- 500 MB available hard disk space

---

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

---

---

## Library Requirements

---

The following libraries require special configurations to run StorNext.

### DAS and Scalar DLC Network-Attached Libraries

Prior to launching the StorNext Configuration Wizard, DAS, and Scalar DLC network-attached libraries must have the DAS client already installed on the appropriate host control computer.

#### DAS Attached Libraries

For DAS attached libraries, refer to "Installation and Configuration" and "DAS Configuration File Description" in the *DAS Installation and Administration Guide*. The client name is either the default StorNext server host name or the name selected by the administrator.

StorNext can support LTO-3 WORM media in DAS connected libraries, but WORM media cannot be mixed with other LTO media types in one logical library.

To use LTO-3 WORM media in a logical library, before configuring the library in StorNext, set the environmental variable `XDI_DAS_MAP_LTO_TO_LTOW` in the `/usr/adic/MSM/config/envvar.config` file to the name of the library. The library name must match the name given to the library when configuring it with StorNext. If defining multiple libraries with this environmental variable, separate them with a space. After setting the environmental variable, restart StorNext Storage Manager (SNSM).

---

**Note:** SDLC software may not correctly recognize LTO-3 WORM media in the library and instead set it to "unknown media type." In this case you must manually change the media type to "LTO3" using the SDLC GUI.

---

#### Scalar DLC Attached Libraries

For Scalar 10K and Scalar 1000 DLC attached libraries, refer to "Installation and Configuration" and "Client Component Installation" in the *Scalar Distributed Library Controller Reference Manual* (6-00658-02).

The DAS client should be installed during the installation of the Scalar DLC attached libraries. Use this procedure to install the DAS client.

- 1 Select **Clients > Create DAS Client**.

The client name is either the default StorNext server host name or the name selected by the administrator.

- 2 When the DAS client is configured in Scalar DLC, select **Aliasing**.

- 3 Select **sony\_ait** as the **Media** aliasing.  
The default value is 8mm.
- 4 Verify that **Element Type** has **AIT** drive selected.
- 5 Click **Change** to execute the changes.

---

## Disk Requirements

---

Disk devices must support, at minimum, the mandatory SCSI commands for block devices as defined by the SCSI Primary Commands-3 standard (SPC-3) and the SCSI Block Commands-2 (SBC-2) standard.

To ensure disk reliability, Quantum recommends that disk devices meet the requirements specified by Windows Hardware Quality Labs (WHQL) testing. However, there is no need to replace non-WHQL certified devices that have been used successfully with StorNext.

Disk devices must be configured with 512-byte or 4096-byte sectors, and the underlying operating system must support the device at the given sector size. StorNext customers that have arrays configured with 4096-byte sectors can use only Windows, Linux and IRIX clients. Customers with 512-byte arrays can use clients for any valid StorNext operating system.

In some cases, non-conforming disk devices can be identified by examining the output of `cvlabel -vvv1`. For example:

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

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

---

## Supported System Components

System components that are supported for use with StorNext 4.2 are presented in [Table 3](#).

---

Table 3 StorNext Supported System Components

Component	Description
Tested Browsers	Internet Explorer 7, 8 or 9 Mozilla Firefox 3 or 4 Other browsers and versions besides the ones tested may work but are not officially supported.

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

---

## Previous Versions of Release Notes

Previous versions of the StorNext release notes contain additional information specific to earlier StorNext releases. You can find previous release notes at the locations below.

Release notes for earlier StorNext releases are available here:

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

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

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

---

## Resolved Issues

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

- [StorNext File System Resolved Issues](#) on page 18
- [StorNext Storage Manager Resolved Issues](#) on page 24
- [StorNext GUI Resolved Issues](#) on page 26
- [StorNext Replication and Deduplication Resolved Issues](#) on page 27
- [StorNext HA, Installation and Other Resolved Issues](#) on page 29

---

### StorNext File System Resolved Issues

---

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

---

Table 4 StorNext File System Resolved Issues

Operating System	CR Number	SR Number	Description
AIX	35025	n/a	Resolved a condition which produced an EFAULT error during I/O on AIX 6.1 client.
Mac OS	32121	1136314, 1190910, 1187818, 1188046	Addressed a condition which caused the Finder to hang on Macintosh clients.

Operating System	CR Number	SR Number	Description
Solaris	34093	n/a	Resolved a condition in which a Distributed LAN Client did not honor a read-only mount option.
	35021	n/a	Corrected a condition which using the mount option <code>io_retry_time</code> rebooted Solaris 10 x86 clients.
Linux	30002	n/a	The <code>listxattr</code> buffer sizing special case now returns results correctly.
	31795	1184942	Corrected an issue related to <code>stat()</code> in which the wrong link count was returned immediately after creating a link to a file.
	33261	n/a	Support was added for loopback devices on Linux clients.
	34020	1153036, 1170810, 1232928, 1258034	Resolved a condition which caused a panic after running <code>fspolicy -b</code> .
	34044	1193788, 1206518, 1271482	Resolved a condition in which a storage disk FSM panic caused a failover.
Windows	31612	n/a	Corrected a blue screen error which occurred due to reconnection issues.
	32022	1169638	Resolved a condition which caused a stop error on a Windows client during shutdown.
	32054	n/a	Addressed a condition which restricted the maximum unicode buffer to 512 characters, which is half of the allowable maximum.
	32203	1287544	Corrected a condition in which <code>CvNtCreate()</code> returned an error if the <code>RelatedFileObject</code> is not a directory.
	32595	1207906, 1241824	Resolved an issue which caused Windows Client 4.0+ to fail when installing on a drive other than "C".
	32865	n/a	Addressed a condition in which changing delete-on-close disposition failed across a CIFS Share.
	33160	1248464, 1057376, 1188830, 1194598, 1222804, 1230408, 1231830, 1241096, 1238878, 1253776	Resolved a condition which caused File Attributes for StorNext directories to be incorrectly set to 0x90.

Operating System	CR Number	SR Number	Description
Windows	33290	1240540, 1259688	Corrected a condition in which an administrative shutdown of only the fsmppm caused the SNFS kernel driver to never reconnect to a new fsmppm.
	33329	n/a	Corrected an issue which prevented unmounting a file system after running IFS tests.
	33866	n/a	Resolved a condition in which the Offline Notification feature blocked access to offline files on a non-StorNext file system.
	34023	1057376, 1188830, 1194598, 1222804, 1230408, 1231830, 1241096, 1238878, 1253776, 1293852, 1288714	Addressed a condition which caused issues accessing files on the StorNext file system related to Server Message Block v2.
	34024	1089486, 1125474, 1198738	Resolved a condition which resulted in incorrect non-realtime credits.
	34040	1133098, 1259688	Corrected an issue which caused intermittent errors during shutdown.
	34041	1185732, 1207432, 1243710	Issues related to the StorNext File System Administrator tool on Windows 2008 have been corrected.
	34062	1196290	User arguments for the device name are now checked and validated before using.
	34091	n/a	Distributed LAN Client proxypath mount options are now supported on Windows.
	34339	1179102	Addressed an issue which prevented a Windows StorNext client (version 4.1.1) from running cvcp commands on Windows 2008 Server.
	35028	n/a	Resolved a condition in which screen information was not properly displayed after refreshing when using the StorNext File System Administrator tool.
	35029	n/a	Corrected a condition which produced a runtime error when checking the clients on a stopped file system from the StorNext File System Administrator tool.
	35138	n/a	Addressed a condition which produced long delays accessing offline files when using the Offline Notification feature.

Operating System	CR Number	SR Number	Description
All	35078	n/a	Corrected a condition which produced the following error: PANIC: fsm "Segmentation Fault Signal." file snfs/fsm/fsm.c, line 622
	32871	n/a	Corrected an issue which caused an FSM panic at startup if a root inode was assigned a security descriptor larger than 4096 bytes.
	30641	n/a	The man page for cvadmin now lists all multipath methods.
	30829	n/a	Resolved a condition which produced incorrect logs when cvadmin started a corrupt *.cfgx file.
	30846	1249250, 1295952	The fsports AltPmap setting now works correctly.
	30955	n/a	StorNext upgrades now automatically detect whether metadumps are required, and then initiates them only if needed.
	31118	1122492, 1250776	The "sticky server" bit for Distributed LAN Client servers threshold values is now configurable instead of hard-coded at 2.
	31982	1178034	Resolved a condition which caused the buffer cache size shown as signed to display incorrectly.
	32218	n/a	Corrected an issue when running on NFS clients which produced sporadic Permission Denied errors incorrectly returned for namespace operations.
	32423	n/a	FSM disk path handling has been improved.
	32440	1198340	Corrected a condition in which cvdb PERF trace has a misleading offset value on proxy servers (aka distributed LAN gateways).
	32446	n/a	Resolved issues which caused a system panic and reboot when running new rename regression tests.
	32857	n/a	Addressed a condition in which StorNext File System wouldn't start when both .cfg and .cfgx files were present, and no diagnostic message were logged.
	33032	1235828	Resolved an issue which caused the counters for TokenSubTypeDataTokenRequestV4 to be erroneously reported as 'Reserved9' when the Allocation Session Reservations feature was enabled.
	33151	n/a	StorNext MD5 code has been optimized to improve performance.
33243	n/a	Addressed an issue which prevented cvdbset from not always setting the requested module bits.	

Operating System	CR Number	SR Number	Description
All	33266	n/a	The Windows Delay Start function is now disabled by default.
	33291	1250220	TCP keepalive is now enabled to prevent client socket leaks.
	33308	n/a	Addressed issues which could cause read_data() to return incorrect content when certain I/O patterns are used.
	33666	1270786	Corrected a condition which caused OpHangLimit panics due to deadlock between VopLink and free pending inode thread.
	33699	n/a	Resolved an issue in which stale mount points were seen on NFS clients when a Solaris or HP-UX NFS server was rebooted.
	33765	n/a	Diagnostic log message were added to address fsm pm handle_traffic thread stalls.
	34022	1049222, 1173138, 1257754	Addressed and issue which produced the following error: PANIC: /usr/cvfs/bin/fsm ASSERT failed "pthread_mutex_lock(&op->open_lock) == 0" file inode.c, line 9536
	34037	1168230	Corrected a condition in which attempting to run cvfsck on the file system resulted in a segmentation violation.
	34048	1211262	Added a message that an alsplaytree_insert PANIC is likely caused by metadata corruption, so cvfsck should be run.
	34051	1206682	Corrected a condition in which clients with mounted file systems are unaware of stripe group READ/WRITE access changes until remounting.
	34055	1184566	Resolved a condition in which snmetadump failed to restore journals.
	34058	1226008	Corrected a condition which prevented cvfsck from correctly clearing the root directory.
	34063	1199968, 1264772	Resolved a condition which caused snmetadump allocation to fail when trying to mark space for inode data.
34066	1235828, 1270982, 1240698	Alloc_space() performance now scales appropriately with many stripe groups or large numbers of fragments.	
34068	1221394	Corrected a condition which caused snmetadump failure after pushing all of the working inode data back to the metadump.	

Operating System	CR Number	SR Number	Description
All	34468	n/a	Corrected a condition in which the show fl command in cvfsdb returned no or incomplete results.
	34628	1276898	Resolved a condition which caused DDM hangs on the MDC if Distributed Data Mover on a client is rebooted.
	34728	n/a	Addressed an issue in which LSN reuse lead to the following assertion failure: free_ip->i_idinode.idi_flags & InodeFlagFree
	34803	n/a	Resolved a condition causing a client file system to hang when FSM ungracefully shut down while performing creates.
	34903	1268288	StorNext formatted drives now report the volume serial number correctly.
	34929	1233602 1277592	Corrected an issue in which cvfsck -f used 100% of the CPU and did not complete.
	34979	1295102	CvApiGetDiskInfo() now returns the error number.
	35017	1285178	Resolved a condition which caused backups to fail with the following error message: Sequence number error: cookie 0x800012f34503 dinode type CHILD, dump seqno 0xe5 journal seqno 0x2e5
	35033	1233414, 1292222	Addressed an issue in which the license.dat file was incorrectly generated by snlicense on a client-only system.
	35121	n/a	Corrected a condition in which editing a file system immediately after creation caused the file system to hang.

## StorNext Storage Manager Resolved Issues

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

Table 5 StorNext Storage Manager Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	34045	1197476	The maximum allowed policy class hard limit has been increased from 25,000. The new limit is 65,550.
	34046	1196858, 1172414, 1246590	The RAS event "No cleaning media available" can now be suppressed.
	34060	1229524, 1208816	Addressed an issue related to the MAX_STORE_SIZE parameter which resulted in unclear admin alert messages or warnings.
	34978	n/a	Corrected a condition in which fsdevice -f -c direct failed to process the command.
All	28050	990408, 1055592, 1132632, 1210220, 1243948, 1249176	Resolved an issue which resulted in an fsqueue -f panic.
	31263	1136906	Addressed an issue which caused fs_fmover to retrieve the incorrect media copy during fsretrieve operations.
	31409	n/a	Corrected a condition which prevented TSM from starting if a drive was missing the serial number in the cfgdir configuration directory.
	32751	1217762 1220512	RAS messages for Linter 903 errors have been added to inform StorNext users of potential problems.
	32888	n/a	The Web Services APIs fsfileinfo and fsretrieve now output to XML and JSON formats.
	34021	987894, 1108690, 1110514, 1200698	Resolved a condition which prevented fsrecover from creating empty relation points, and also caused fsclean failure.
	34067	1244034	Resolved a condition which cleared tape drive reservation if the library was powered Off or On.

Operating System	CR Number	SR Number	Description
All	34203	n/a	Corrected a condition which prevented additional storage disk file system capacity from being recognized.
	34544	1271618	Resolved an issue in which a relation point directory could be deleted while media was still associated with store policy.
	34635	1293458	Addressed an issue which caused the Samfs export command to fail dumping large SAM-QFS namespaces.
	34721	1277992	Corrected an issue which resulted in inaccuracies for fill level and fragmentation levels on the Media Fragmentation Report.
	34972	n/a	Corrected a condition which caused truncation policies to incorrectly determine when they have truncated enough data.
	35090	n/a	You can now run two truncation policies for different classes simultaneously.
	29268	1005118	Class-based truncation has been added to fsschedule.
	34035	1156698	Notification has been improved for situations which cause the MSM XdiAMTask process to terminate.
	34338	n/a	Resolved a condition in which running vsaudit -p on ACSLS library generated Error VOL205: hardware communication error.
	34672	1276858	Addressed an issue which caused memory corruption in the ArcDisp process for fsmedcopy mount requests.
	34724	1129698;1204516	Changes were added to MSM to clean up after an "add library" error.
	35069	n/a	Resolved an issue in which running the vsarchiveconfig command failed to update the configuration files to a new (replacement) server.

## StorNext GUI Resolved Issues

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

Table 6 StorNext GUI Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	29648	n/a	Resolved a condition which produced a GUI hung after changing the block size from 64k to 4k.
	30959	n/a	Addressed an issue in which editing a file system erroneously changed the sector count for a LUN in the cfgx configuration file.
	32923	n/a	Corrected a condition which caused the StorNext Configuration Wizard to stop at Storage Policies for an SNFS-only configuration.
	35024	1293270	Resolved a condition which prevented the StorNext GUI from displaying the connected LAN clients.
	29724	n/a	Corrected an issue which produced an internal error when deleting a managed file system
All	31583	n/a	Resolved a condition which prevented Distributed Data Mover configuration from logging entries on the StorNext GUI job page.
	32149	n/a	Corrected a condition which caused HA conversion to fail with a Web service error.
	32542	n/a	GUI Usability issue when adding/removing drives from an archive
	32934	1217762	Resolved a condition in which deleting a Replication/ Deduplication policy did not create a job.
	33124	1247562; 1269314	The Recover File now shows more than just the latest deleted version.
	33564	1269690, 1276702	The Copy Number heading has been added to the Media Actions Page.
	34028	1129698; 1204516	The StorNext GUI now cleans up after an "add library" error occurs.
	34029	1129698	Resolved an issue which required a StorNext GUI restart after receiving an "add library" timeout.

Operating System	CR Number	SR Number	Description
All	34036	1174094, 1300576	Corrected a condition in which the StorNext Home page did not display information if any of the components fail to return data.
	34054	1214540	Corrected an issue which prevented moving media from an offline archive.
	34204	1277000	Addressed issues which cause the Library Operator Interface to fail in some instances.
	35079	1304018	Resolved an issue where the StorNext GUI obtained and cached the HA shared file system name from an incorrect source.

## StorNext Replication and Deduplication Resolved Issues

[Table 7](#) lists resolved issues that are specific to StorNext replication and deduplication.

Table 7 StorNext Replication, Deduplication and Other Resolved Issues

Operating System	CR Number	SR Number	Description
Linux	34861	1298916	Addressed an issue in which namespace collection on the source was missing files that were in directory.
	35082	1236248	Additional threads were added for delete processing on the blockpool.
	35137	1298916	Resolved an issue in which namespace replication generated a bad stream format error from the target.
All	32607	1212124	Addressed an issue in which running <code>snpolicy -removepolicy</code> left a policy key in the memory inode.
	34064	1217762	Addressed an issue which caused replication cleanup to clear two INHERITED events which could cause StorNext to generate Linter 903 errors.
	34074	1239918	Resolved a condition in which ingesting into full blockpool file system left a wrong reference count on blockpool tags.
	34082	n/a	Addressed an issue in which replication failed with errors reporting missing binary objects.

Operating System	CR Number	SR Number	Description
All	34638	n/a	Corrected an issue which caused replication reports to be stored locally even in HA environments.
	34675	n/a	Resolved an issue which caused replication to fail when replicating terabytes of data.
	34864	n/a	Addressed a condition which cause the following error: PANIC: snpolicyd ASSERT failed "stream->rep_ref_cnt > 0" rep_stream.c, line 447
	34867	n/a	Added error messages which should appear when replication failed due to "bad stream format".
	34868	n/a	Added warning messages which should appear when files are skipped in a directory content dump.
	35070	n/a	Resolved an issue which prevented the target's rep_private's 'dir' from being cleaned up.
	35076	n/a	Corrected an issue which produced erroneous replication report data.
	35083	n/a	Addressed an issue which caused data replication to be mixed with namespace replication for the same policy key.
	35119	n/a	Measures were added to prevent failure when upgrading from a truncated generation number for the replication target.
	35202	n/a	Corrected a condition in which calling bp_update on an HA system with a blockpool prevented startup.
	35231	1272478, 1312226	Addressed a condition which prevented replication cleanup_tree from removing the top directory.
32645	n/a	The replication source now retains the namespace target size.	

## StorNext HA, Installation and Other Resolved Issues

[Table 8](#) lists resolved issues that are specific to StorNext HA, installation and other StorNext features or components.

Table 8 StorNext HA,  
Installation and Other Resolved  
Issues

Operating System	CR Number	SR Number	Description
Linux	34077	1246826 1216514	Corrected an issue which prevented "cp -a ..." on a client from running during MDC failover for a managed file system.
	34975	n/a	Resolved a condition in which HaManaged FSM is started on both MDCs when an unmanaged file stem is changed to managed.
All	29708	n/a	Resolved a condition in which snhamgr stopped working due to MAX_CONNECTIONS.
	34052	1212388	Web Services API must run on a metadata network, port 81.
	34061	1227570	Addressed an issue which caused syncha.pl to use the wrong "filelist".
	34362	n/a	Corrected FSM shutdown errors which occurred after updating the license.dat file.
	34069	1242518, 1259020	Corrected an issue which caused XdiAMTask to hang during MSM startup for an ACSLS configuration.
	34095	1266808	Resolved problems configuring a QualStar XLS tape library and drives, and bringing it online.
	34457	1281618	Cannot configure IBM TS3500 library with IBM 3592E05 drive
	35023	n/a	Resolved an issue which caused importing media on a SpectraLogic library to fail.
	32914	1279850, 1279862	Install.stornext now warns users about missing rpm dependencies before starting the install/upgrade process.
34050	1211912	Addressed an issue with the installer's SDBM_File (Install_stats) which caused the upgrade to fail.	

Operating System	CR Number	SR Number	Description
All	35139	n/a	After a StorNext installation or upgrade successfully completes, install.stornext now automatically exits rather than waiting for the user to explicitly choose Quit from the StorNext installation/upgrade menu.

## 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 30
- [StorNext Storage Manager Known Issues](#) on page 32
- [StorNext GUI Known Issues](#) on page 36
- [StorNext Installation, Replication, HA and Other Known Issues](#) on page 40

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

Table 9 StorNext File System Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	28561	n/a	When a machine is rebooted, in rare situation, the network interfaces may have been changed. If you have configured Distributed LAN Server (DLS), the file system may fail to mount because the DLS may perceive that the associated network interface may have changed or disappeared.	To work around this issue, you must find the appropriate network interfaces and reconfigure the dpserver.fsname file and then remount the file system.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	31964	n/a	If an MDC running Storage Manager mounts a file system being served by another MDC pair, Healthcheck will generate a RAS message containing the following text: Internal Software Error: an unhandled software error has occurred. ERROR: mounted CVFS file system name(/stornext/win_snfs1) not found in list of all CVFS file system names.	In this case the error message is falsely generated and can be safely ignored. Alternatively, the issue can be avoided by not mounting a StorNext file system as a client on the MDC.
Windows	29483	n/a	After changing fsnameservers in the StorNext GUI, the file system failed to mount and returned a "device not connected" error.	The workaround is to stop and start the StorNext services manually from the command line.
	30945	n/a	On Windows Systems, creating a Virtual Hard Disk (VHD) on a StorNext file system will cause a system crash.	Until there is a resolution to this problem, all VHDs must be created on NTFS file systems. This includes both direct creation of VHDs through the Windows Disk Manager on Windows 7 and Windows 2008 R2 systems, as well as any Windows-based product that uses VHDs implicitly. This includes, but is not limited to, the Windows Complete PC Backup and Microsoft Virtual PC products. All of these products should continue to function normally with StorNext installed as long as they are not configured to create their VHDs on StorNext.
	30696	n/a	Attempting to "loopback" mount an ISO image residing in a StorNext file system resulted in an error message and mount failure.	One workaround is to copy the ISO image to a local file system and then perform the loopback mount on the copy.

## StorNext Storage Manager Known Issues

[Table 10](#) lists known issues specific to StorNext Storage Manager.

Table 10 StorNext Storage Manager Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	29445	n/a	When adding Distributed Data Mover (DDM) mover hosts (either by the StorNext GUI or by the <code>fsddmconfig</code> 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.
	28447	n/a	<p>In the event of an environmental issue (such as power interruption) in which both the MDC and Distributed Data Mover mover machines need to restart, there can be a timing situation in which the MDC comes up and tries to initiate TSM before the client has completed its own reboot.</p> <p>TSM will use the <code>fs_fmoverc</code> process for up to five minutes to clean up the status of the DDM mover machines. If the DDM mover machines remain down for an extended period, there can be retries of the five-minute <code>fs_fmoverc</code>.</p>	<p>If TSM has not completed its startup and <code>fs_fmoverc</code> processes persist (as shown by running the <code>tsmup</code> command), use the <code>fsddmconfig</code> command to set the state of the DDM mover machine to disabled.</p> <p>For example, if the hostname of the DDM mover machine is <code>minnesota</code>, run this command on the MDC:</p> <pre>fsddmconfig -u -s d minnesota</pre> <p>Once the DDM mover machine is back up, re-enable it using the StorNext GUI or with this command:</p> <pre>fsddmconfig -u -s e minnesota</pre> <p>Under most circumstances, <code>fs_fmoverc</code> will successfully determine the downed status of a DDM mover machine, and automatically set its status to DISABLED. A RAS ticket is emailed when <code>fs_fmoverc</code> automatically disables a mover, so after receiving such a ticket and rebooting the DDM mover, use the StorNext GUI (or the command <code>fsddmconfig</code> in the second example above) to enable the mover.</p>

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

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	23377	725697 1236258	Errors can result when there are multiple retrieves in progress at the same time for the same multi-segmented file, and then one of the retrieves is canceled when the last segment is being retrieved.	You can avoid this issue by either not requesting multiple retrieves for the same multi-segment file, or by not canceling the subsequent requests.
	32711	1213924 121777 1288408	When using the StorNext GUI to configure the Distributed Data Mover (DDM) feature, errors can result if you use the IP address when configuring the MDCs or clients.	To avoid this situation, use only the host name when configuring DDM.
	34050	n/a	StorNext does not currently support running two truncation policies for different policy classes at the same time.	To avoid errors, do not run two truncation policies for different classes at the same time.
	34145	1287532	When using the StorNext GUI to defragment or copy tapes on systems with multiple libraries, the GUI does not allow you to specify the library in which the target tape resides.  This issue applies only to systems with multiple libraries, archives or partitions attached.	To work around this issue, use fsmecopy to defragment tapes or to copy media to tapes in the desired library.
	35537	n/a	Shell expansion in TSM could cause a segment fault (core dump) due to the way command line wildcards are expanded.	StorNext users will not see this error.

## StorNext GUI Known Issues

[Table 11](#) lists known issues specific to the StorNext GUI process.

Table 11 StorNext GUI Known Issues

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	30842	n/a	<p>In some situations, on HA systems the StorNext GUI exits the config mode (with the redundant server in peerdown mode) without starting Storage Manager.</p> <p>When an HA Cluster is operating without a redundant server for an extended time period, it is best practice to place the cluster into the single/peerdown state to prevent an HA Reset. The StorNext GUI allows transitions from this state into the config/peerdown state for making configuration changes.</p> <p>However, when you use the "Exit Config Mode" button on the StorNext GUI's Tools &gt; HA &gt; Manage screen to transition back from Config to Single mode, the GUI uses a diagnostic method for restarting the system that results in starting the file system without starting the Storage Manager.</p>	<p>Two actions are necessary to correct this issue:</p> <ol style="list-style-type: none"> <li>1 Run the following CLI command: "DSM_control start". This will produce the following warning message, which can be ignored: "fsmgm is already running. Only one fsmgm can be running at a time. Aborting start."</li> <li>2 In the StorNext GUI, go to the <b>Tools &gt; System Control</b> screen and start the Storage Manager.</li> </ol>
	29929	n/a	Due to an error in a third-party component, double-clicking the New button on the File System > New screen returns exceptions.	This issue should be addressed in a future release of the third-party component. Other solutions may be considered for a future StorNext release.
	29038	n/a	Clicking through StorNext GUI logs pages more than once every few seconds causes error.	This error message can be ignored and requires no response. To remove the error message, click <b>Refresh</b> or select another log or page.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
Linux	31722	n/a	You may receive the following error message when attempting to scan the secondary system in an HA configuration: “Host cannot be used as a secondary system. Error getting system descriptor. Unable to get system id for texas < org.apache.axis2.AxisFault: Transport error: 4.0.1 Error Authorization Required.”	This indicates a failure for the StorNext GUI to communicate with the secondary MDC. Check that the StorNext software on the secondary system has been installed correctly, and then reboot the secondary system to make sure the software is running correctly.
	32553	n/a	The Library Operator Interface (LOI) page and the Media Action page may not display immediately if you have thousands of media requiring attention in the LOI page (for example, media moves, media ejects, etc) at the same time. A general estimate is that the delay could be approximately 10 seconds for every 500 media requiring attention on the LOI page.	This situation is extremely unlikely to occur, and most StorNext users will never experience this delay unless they need to export thousands of media at one time. In that situation, you can avoid this issue by breaking up media action requests into smaller chunks.
All	29557	n/a	The StorNext GUI can hang if it is unable to create a file system. Additional validation is necessary.	If the StorNext GUI makes an API call which causes an FSM panic, the StorNext GUI is taken down as well. To work around this condition, restart the StorNext GUI (service stornext_web start).
	29728	n/a	Due to an error in a third-party component, background pages respond to keyboard input when modal dialogue windows are open.	Avoid entering keyboard input when dialogue windows are displayed.

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

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	34213	n/a	Some Internet Explorer 9 default security settings could prevent StorNext from operating properly with this browser.	<p>To work around this issue, follow this procedure to enable security options in Internet Explorer 9:</p> <ol style="list-style-type: none"> <li>1 Launch Internet Explorer 9.</li> <li>2 Choose <b>Internet Options</b> from the Tools menu.</li> <li>3 Click the <b>Advanced tab</b>.</li> <li>4 Enable the following: <ul style="list-style-type: none"> <li>• Run ActiveX Controls and Plugins</li> <li>• Script ActiveX Controls Marked Safe</li> <li>• Active Scripting</li> <li>• File Download</li> <li>• Don't Prompt for Client Certificate...</li> </ul> </li> <li>5 Click <b>OK</b>.</li> <li>6 Close Internet Explorer 9 and then open it again to continue.</li> </ol>
	34786	n/a	If you have multiple, active login sessions using the same browser, information on some screens may not be displayed correctly. This behavior was observed on the <b>Reports &gt; Media</b> screen but could affect other pages as well.	To avoid this issue, do not keep open multiple active StorNext sessions which use the same login and the same browser.
	34941	n/a	When you enter a policy class name using upper case characters, the name is displayed in lower case characters after the policy is created.	This issue can be safely ignored. You can avoid this situation by not using upper case characters when entering policy class names.
	34989	n/a	If a Web Services URL request contains the string "policy", it fails with a "Requested resource not found!" error.	To avoid this situation, do not use the string "policy" in a pathname that is intended for Web Services use.
	35496	n/a	When checking for data savings on a deduplication-enabled file system, the data on the Replication / Deduplication Policy Activity Report may be inaccurate.	Until this issue is resolved, do not rely exclusively on the data savings information shown in the Replication / Deduplication Policy Activity Report.
	35506	n/a	Errors could result if you enable quotas on a replication-enabled file system.	To avoid this error, do not enable quotas on files systems which are used for data replication.

Operating System	CR Number	SR Number	Description	Workaround (if applicable)
All	35605	n/a	The StorNext GUI might hang with java out-of-memory errors if you view a RAS ticket which has a large number of event details.	If you encounter this issue, restarting StorNext will allow you to use StorNext, but you will be unable to view RAS tickets. Contact Quantum Support for assistance restoring your capability to view RAS tickets.

## StorNext Installation, Replication, HA and Other Known Issues

[Table 12](#) lists known issues specific to StorNext installations, data replication, HA systems and other areas.

Table 12 StorNext Installation, Replication, HA and Other Known Issues

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	28375	n/a	StorNext doesn't replicate non-regular files. (Currently, StorNext skips files that are UNIX domain sockets, block/char devices, and fifo's. Some of these file types could be included in namespace replication.)	The workaround is to ensure that your replication workflow does not rely on non-regular file replication.
	29416	n/a	In an HA configuration, RAS messages are not generated for loss of SAN connectivity by the secondary system.	A workaround is to activate some of the unmanaged file systems on the secondary metadata controller. This will allow RAS messages from the secondary MDC if there is SAN connectivity loss. (This action would also help with load balancing.)
	29486	n/a	The number of file systems found after clicking the Scan button does not match the number of mountable file systems on an HA system.	Note: the count of file systems found includes standby file systems.

Operating System	Change Request Number	Service Request Number	Description	Workaround
Linux	29678	n/a	Replication may hang if the dedup_bfst parameter ("Address for Replication and Deduplication") on the source has been configured to use an address that is not reachable by the target.	The workaround is to manually confirm reachability to the replication source's vIP address on the replication target, and then reconfigure routing, if necessary.
	31959	n/a	After converting to high availability, the blockpool status goes to "Verify pending" state and an error message appears.	The workaround is to stop and then start blockpool services.
	35387	n/a	Running <code>install.stornext -upgrade</code> fails with a licensing error when the system is configured and licensed for either replication or deduplication without also being licensed for both features. The license warning can be ignored, provided that the system has the appropriate licenses installed for the features in use.	When either replication or deduplication is licensed without the other, use the <code>-licforce</code> option when upgrading Stornext . For example: <code>install.stornext -upgrade -licforce</code>
All	29023	n/a	Replication quiesce scripts do not synchronize data on any clients that have open files.	To avoid this issue, close open files prior to running quiesce scripts.
	29067	n/a	When an MDC in an HA cluster starts up while LUNs are not available, the FSMPM process will try to access the LUNs, but eventually stop trying. The MDC will not provide file system services without intervention to restart StorNext. Although the <code>snhamgr status</code> command reports the MDC as running, it is not performing as a redundant server and will not take control in the event of an HA Reset.	After repairing access to the LUNs, stopping and restarting CVFS by running the following command may correct the problem: <code>service cvfs restart</code>
	29722	n/a	StorNext API 2.0.1 failed to install on the secondary node of an HA pair.	A workaround is to fail the primary over to the secondary and then install SNAPI.

Operating System	Change Request Number	Service Request Number	Description	Workaround
All	30006	n/a	There is no way to conveniently delete a TSM relation point used for replication.	You can manually delete the relation point by running the command <code>rm -rf /snfs/sn2/tsm/.rep_private</code> , which empties the TSM relation point. When running this command, be aware that there may have been several targets being realized with the TSM relation point in question, so you should remove the directory <code>tsm_dir / .rep_private</code> only after the LAST target policy has been removed from the relation point.
	30812	n/a	When a replication source system has a hard reset and reboots, it may require the TCP Keepalive time to expire before replication resumes. On most Linux systems, this time duration defaults to a little over two hours.	The only workaround for this problem is to reset the global TCP Keepalive parameters for your system to lower values which may affect other processes running on the system. You may want to experiment with lower values to see how your system behaves, as higher values of keepalive may be used to protect applications from unstable network links.
	30817	n/a	When replicating with multiple copies, StorNext can propagate metadata changes into the previous replication copies.	If file attributes on replicated files have been changed since the replication date of the requested copy, verify the retrieved files have the expected file attributes.
	31825	n/a	When running <code>snpolicy -rundelete=mntpnt -wait</code> , <code>snpolicy</code> can run for a long time before returning. While this operation is running other <code>snpolicy</code> commands can block.	The workaround is to allow <code>snpolicy</code> commands to complete before submitting additional <code>snpolicy</code> commands.
	35249	n/a	If replication fails due to error "Replication target name collides with existing dir", this indicates that the same directory used to attempt the replication namespace realization already exists on the replication target, but has a different policy key associated with it.	To work around this issue, remove the whole directory and then perform another replication.

Operating System	Change Request Number	Service Request Number	Description	Workaround
All	35334	n/a	<p>Under certain circumstances replication users are susceptible to a functional vulnerability if the source file is accidentally deleted before that file has been replicated on the target.</p> <p>This situation would produce an error message similar to this:  <code>/stornext/source/ __CVFS_Handle.0004AAB6E016CF7D000E00000000000000000000001F0".</code>  <code>No such file or directory.</code>  <code>, errcode: 6806ca1. ]</code></p> <p>This condition could put the target into a loop repeatedly trying to access the now deleted file.</p> <p>This situation applies only if BOTH of the following conditions are met:</p> <ul style="list-style-type: none"> <li>• You have both a managed source file system and a managed target file system</li> <li>• You have upgraded the source to StorNext 4.2, but left the target at an older StorNext release</li> </ul>	If this situation applies to you, Quantum strongly recommends upgrading the target file system to StorNext 4.2.0 as soon as possible.
	35579	n/a	<p>The snpolicyd process can crash on a replication-enabled source system when a replication source policy has configured multiple targets and a replication target file system has filled up.</p>	To work around this issue, ensure that there is sufficient space in any replication target file systems to complete all outstanding replication tasks.

## Replication Multilink Known Issues

### The Replication Multilink Feature

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

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

The following table shows known issues related to Multilink:

Table 13 Multilink Known Issues

Change Request Number	Description	Workaround
30746, 30921	Once enabled, the Multilink feature cannot be disabled from the StorNext GUI.	Contact Quantum Technical Support for assistance with disabling the Multilink feature
30920	In HA configurations, the StorNext GUI may display errors about virtual IPs (VIPs) when configuring Multilink if the peer system is down.	Ensure the standby MDC is up when configuring Multilink.
29984	The StorNext GUI does not indicate whether Multilink is enabled.	Examine <code>/usr/cvfs/config/blockpool_config.txt</code> and look for lines containing <code>MultiLinkConfig</code> and <code>MultiLinkInterface</code> . Values other than "None" indicate that Multilink is configured.

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

## Operating Guidelines and Limitations

[Table 14](#) lists operating guidelines and limitations for running StorNext. Items are grouped according to operating system.

Table 14 StorNext Operating Guidelines and Limitations


Operating System	Feature or Category	Description
AIX	Clients	<p>Clients on AIX systems may not unmount a file system after running the <code>fsstress</code> command against that file system. The client must then be rebooted to release the mount.</p> <p>This issue occurs on AIX systems when the <code>fsstress</code> command is run using <code>mknod</code> on the same command line. To prevent encountering this behavior, do not include <code>mknod</code> when running <code>fsstress</code>. Otherwise, you will be required to reboot the client without a successful unmount.</p> <p><b>Note:</b> A StorNext 4.2 AIX client supports only AIX version 6.1</p>
Solaris	StorNext labels	<p>Solaris hosts may need to rescan disk devices after StorNext labels have been applied.</p> <p>In particular, when a StorNext label is put on a LUN less than 1TB in size, Solaris hosts will not be able to use that LUN until they have done a device rescan. A device rescan is accomplished with a boot flag:</p> <pre>reboot -- -r</pre> <p>This issue will be addressed in a future StorNext release.</p> <p>In the meantime, work around this issue by rescanning devices using the boot flag <code>reboot -- -r</code></p> <p>If the labeling operation was performed on a Solaris host, that host does not need to do the rescan. However, some intermediate versions of the Solaris 10 Kernel Jumbo Patch break the necessary functionality to support this; please be sure you have applied the latest Solaris 10 Kernel Jumbo Patch before labeling any StorNext LUNs.</p>
Linux	Linux Multipath Support (the <code>rr_min_io</code> setting in the Linux DM Multipath Driver)	<p>Current versions of the Linux DM Multipath driver assign a default value of 1000 for <code>rr_min_io</code> which is too high for most configurations having multiple active paths to storage. Using a smaller value such as 32 will typically result in significantly improved performance. Refer to the RedHat or SuSE documentation provided with your version of Linux for details on how to apply this setting.</p> <p><b>Note:</b> Experimentation may be required to determine the optimal value.</p>

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

Operating System	Feature or Category	Description
Linux	Migrating metadata controllers	<p>StorNext users migrating their metadata controllers from Apple Xsan to Linux should be aware of the following upgrade considerations:</p> <ul style="list-style-type: none"> <li>• If the file system is running Xsan 2.1.1 or earlier, it should be a simple upgrade: just replace the MDC.</li> <li>• If the file system is running Xsan 2.2 or later with "NamedStreams No" (which is the default for Xsan 2.2,) it should also be a simple upgrade: just replace the MDC.</li> <li>• If the file system is running Xsan 2.2 or later with "NamedStreams Yes," you must completely remake (reformat) the file system. For obvious reasons, you should do a complete backup before migrating.</li> </ul>
	Subtree Check option	<p>Subtree Check Option in NFS No Longer Supported</p> <p>Although supported in previous StorNext releases, the subtree_check option (which controls NFS checks on a file handle being within an exported subdirectory of a file system) is no longer supported as of StorNext 4.0.</p>
	FQDN	<p>SuSe Linux distributions automatically associate the FQDN of the local machine with the address 127.0.0.2 in the /etc/hosts file. There is no benefit from doing this when the machine is connected to a network that can resolve its name to an IP address.</p> <p><b>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.</b></p>
	Software Firewalls	<p>Software firewalls such as "iptables" on Linux and Windows Firewall can interfere with the proper functioning of StorNext and result in unexpected errors unless specifically configured for use with StorNext.</p> <p>Quantum strongly recommends that all software firewalls be disabled on systems used as StorNext clients and servers. If required, StorNext can be configured for use with hardware firewalls.</p> <p>For more information, refer to the fsports man-page or help file and the "Ports Used By StorNext" section in the <i>StorNext Tuning Guide</i>.</p>
Windows and Linux	Symbolic links to StorNext directories	<p>If you create a symbolic (soft) link in Linux to a directory on a StorNext file system, the link cannot be used by Windows. Windows also cannot process a symbolic link which contains a path to a file in another file system.</p>
Windows	Window backup utility	<p>When a StorNext file system is mounted to a drive letter or a directory, configure the Windows backup utility to NOT include the StorNext file system.</p>

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"> <li>1 Click Start, and then click <b>Run</b>.</li> <li>2 In the Open box, type Regedit and then click <b>OK</b>.</li> <li>3 On the Edit menu, click <b>Find</b>.</li> <li>4 In the Find what box, type Snfs_XXX.dat and then click <b>Find Next</b>.</li> <li>5 If the search result selects a string value called <b>PackageName</b>, continue with these steps. Otherwise, repeat steps 3-4.</li> <li>6 Double-click the <b>PackageName</b> string value.</li> <li>7 In the <b>Value</b> 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.)</li> <li>8 On the Registry menu, click <b>Exit</b>.</li> </ol>
	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>

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"> <li>1 On the Windows client machine, right-click the <b>Recycle Bin</b> icon on the desktop and then click <b>Properties</b>.</li> <li>2 Click <b>Global</b>.</li> <li>3 Click <b>Configure drives independently</b>.</li> <li>4 Click the <b>Local Disk</b> tab that corresponds to the mapped or directory-mounted file system.</li> <li>5 Click the checkbox <b>Do not move files to the Recycle Bin. Remove files immediately when deleted</b>.</li> <li>6 Click <b>Apply</b>, and then click <b>OK</b>.</li> </ol>

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"> <li>1 On the Windows client machine, right-click the <b>Recycle Bin</b> icon on the desktop and then click <b>Properties</b>.</li> <li>2 Click the <b>General</b> tab.</li> <li>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.</li> <li>4 Choose the option <b>Do not move files to the Recycle Bin. Remove files immediately when deleted</b>.</li> <li>5 Click <b>Apply</b>.</li> <li>6 Repeat steps 3-5 for each remaining directory-mounted file system.</li> <li>7 When finished, click <b>OK</b>.</li> </ol> 
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"> <li>• The maximum number of disks per file system is 512</li> <li>• The maximum number of disks per data stripe group is 128</li> <li>• The maximum number of stripe groups per file system is 256</li> <li>• The maximum number of tape drives is 256</li> </ul>
	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	<p>When using StorNext with Internet Explorer 8, 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"> <li>1 Launch Internet Explorer 8.</li> <li>2 Choose <b>Internet Options</b> from the <b>Tools</b> menu.</li> <li>3 Click the <b>Advanced</b> tab.</li> <li>4 Under the <b>Security</b> heading, make sure the option <b>Display mixed content</b> is checked.</li> <li>5 Click <b>OK</b>.</li> </ol>
	StorNext Licensing	<p>After you purchase or update a feature license and then enter license information through the StorNext GUI, you should restart StorNext services to ensure that your new license is recognized. Certain StorNext features such as replication may not recognize your license until services are restarted.</p>
	Upgrade	<p>Before attempting to upgrade from a previous StorNext release, make sure you have free space on the file system. If the file system is nearly full when you begin the upgrade, serious errors may occur or the upgrade could fail. Best practice is to maintain an area on the file system which is not used for data or system files, but is reserved as an empty buffer to ensure that upgrades and other operations complete successfully.</p>
	StorNext home page	<p>The capacity indicators on the StorNext home page provide <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>
	Backups	<p>Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.</p>
	Tape drives	<p>StorNext does not support hot-swapping tape drives. When replacing or adding new tape drives you must first stop StorNext before installing the new drive.</p>
	Cluster-Wide Central Control	<p>The StorNext Cluster-Wide Central Control file (<code>nss_cctl.xml</code>) is used to enforce the cluster-wide security control on StorNext nodes (client nodes, fsm nodes, and nodes running <code>cvadmin</code>). This file is placed on an nss coordinator server.</p> <p>Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.</p>

Operating System	Feature or Category	Description
All	Labels	<p>Disks with existing non-StorNext labels may not show up in the StorNext GUI in order to protect non-StorNext disks from being accidentally overwritten. If you need to label a disk that is not visible in the StorNext GUI, use the <code>cvlabel</code> command to label the disk or use <code>cvlabel -U</code> to remove the existing label from the disks. (Refer to the <code>cvlabel</code> man pages for instructions on labeling and unlabeled drives.)</p> <p><b>Caution:</b> Modifying the label on an active non-StorNext disk can make the disk unusable. Proceed with caution.</p>
	Xsan	<p>It is not possible to delete data within a StorNext policy relation point from an Xsan client via the Finder. Rather, data must be deleted using the shell.</p>
	Replication/ Deduplication	<p>When either the StorNext GUI or the <code>snpolicy</code> command are used to create or modify a replication/deduplication policy, a policy text file is written to the file system.</p> <p><b>Example:</b> 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><b>Note:</b> 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 you run the command <code>snpolicy_gather -b &gt; "some_file"</code></p> <p>This will save a copy of your current configuration. The <code>-b</code> option will also create a copy of policy information in the <code>usr/cvfs/data/fsname/policy_history</code> directory.</p>

Operating System	Feature or Category	Description
All	Replication/ Deduplication	If you are using the Deduplication or Replication feature, part of the installation process is to update the on-disk index. The time required to complete this part of the installation process times may vary depending on the size of your licensed blockpool, drive performance, and other factors. As a general guideline, allow approximately five minutes for a 10TB blockpool.
	Replication	When creating or editing a replication storage policy, there is a field on the Outbound Replication tab called " <b>Filenames Excluded from Replication.</b> " This field allows you to exclude specific files from the replication process.  This field works the same way as a UNIX shell which lets you pattern match names. For example, entering *.Ø core would exclude all .o files and also files named "core." You could also skip all core files by entering rep_skip=core*.
	Deduplication	If a deduplication candidate is removed before blockpool processing is completed, errors such as the following may be sent to the syslog:  Oct 2 15:22:00 orleans Blockpool[16403]: E: [5] (Store Local) Error storing file "/stornext/source/____CVFS_Handle.000474F892EBB65E000E0000000000000000000292BF4".  Error opening file "/stornext/source/____CVFS_Handle.000474F892EBB65E000E0000000000000000000292BF4".  No such file or directory.  Errors such as these may appear serious, but there is no reason for concern.  If you receive these errors, no action is required.
	HA	When exiting HA Config mode, StorNext will be stopped, which will also 'fuser' any processes which have files open on the file system from either node.  Prepare systems and users for this eventuality before entering HA Config mode.

Operating System	Feature or Category	Description
All	HA	<p>On HA systems only:</p> <p>The <code>/usr/cvfs/config/ha_peer</code> file supports some essential HA features by providing an address for HA administrative communications between the MDCs in an HA Cluster. If CVFS is started without this file having correct information, the probability of an HA Reset increases. To correct this condition, restore the <code>ha_peer</code> file to the IP address of the peer MDC, and restart StorNext by running the following command:</p> <pre>service cvfs restart</pre> <p>Note: The peer will be Primary after running this command.</p> <p>If the <code>ha_peer</code> file is removed for any length of time while StorNext is running, the <code>snamgr(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>snamgr</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 snamgr restart</code></p>
	HA	<p>On HA systems only:</p> <p>You may receive the following incorrect error message when scanning for a secondary MDC from the StorNext Convert to HA page:</p> <pre>WARN com.quantum.qutosgui.jsf.ha.HaMBean - doScanHost: Secondary system cannot be same as the primary system.</pre> <p>This message is generated if <code>/usr/adic/util/cnvt2ha.sh</code> fails for any reason (for example, if the file system exists on the secondary, if a shared file system can't mount, etc). Upon secondary conversion failures, StorNext resets the <code>ha_peer</code> file to <code>255.255.255.255</code> on the secondary. Since the conversion fails, the primary <code>ha_peer</code> file is not updated and faulty comparison logic causes the erroneous error message (<code>255.255.255.255 == 255.255.255.255</code>).</p> <p>The workaround consists of two steps:</p> <ol style="list-style-type: none"> <li>1 Remove the <code>/usr/cvfs/config/ha_peer</code> file from the secondary system.</li> <li>2 Reset the StorNext processes on the secondary system by running <code>/etc/init.d/stornext_web restart</code>.</li> </ol>

Operating System	Feature or Category	Description
All	HA	<p>On HA systems only:</p> <p>When a non-managed file system is converted to a managed file system in an HA pair, it is possible for the 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"> <li>1 Complete the configuration changes</li> <li>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</li> <li>3 Restart CVFS on the secondary by issuing "service cvfs restart"</li> <li>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)"</li> </ol>
	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><b>Note:</b> The above applies only when the IP address of the VIP falls into the subnet defined by the base interface's IP address and mask.</p>

Operating System	Feature or Category	Description
All	HA	<p>Understanding the performance of FSM failover in StorNext High Availability installations:</p> <p>When a failover of any file system occurs, the new FSM notices if any clients had a file exclusively opened for writes, and waits up to 35 seconds for those clients to reconnect. In the case of an HA Reset of the Primary MDC, that MDC is not going to reconnect, so the failover to FSMs on the Secondary MDC and the promotion of that MDC to Primary status can be delayed by 35 seconds.</p> <p>The StorNext system exclusively opens files on the HaShared file system, but assumes that only the Primary MDC does this and waives the delay for that one file system. Quantum advises against running user processes other than StorNext processes on HA MDCs for performance, reliability and availability reasons. In the event that processes running on the Primary MDC have files exclusively open for writes on other file systems, the availability of those file systems to all clients will be delayed by 35 seconds following an HA Reset event.</p>
	Quotas	<p>When you enable or disable quotas using the CLI <code>cvadmin</code> command, the change does not persist after rebooting. In order to permanently enable or disable quotas, you must modify the Quotas parameter of the file system config file.</p>
	fsretrieve	<p>If you run multiple <code>fsretrieve</code> commands simultaneously to find files (for example, <code>find -type -f   xargs fsretrieve</code>), you might receive error messages because doing this taxes system resources.</p> <p>Instead, use the recursive retrieve command. When you use this command the files under a directory are retrieved in batches, and more sorting is done to put files in tape order for increased performance. Run recursive retrieve by entering <code>% fsretrieve -R .</code></p>
	Distributed LAN	<p>Distributed LAN Clients in HA Environments:</p> <p>Each HA node must have its own <code>dpserver</code> files detailing the NICs on that node. The <code>dpserver</code> files are not synchronized between HA pairs. If the Distributed LAN Server is configured after converting to HA, the file system(s) running as Distributed LAN servers must be unmounted and mounted again to service DLC requests.</p> <p>When deduplication/replication is enabled, one or more Virtual IP Addresses (VIPs) provides access to the Primary MDC (where the blockpool server is running). In StorNext startup and failover situations, the VIP is dynamically associated with a physical address on the Primary server. Do not use VIP interfaces when setting up the <code>dpserver</code> configuration file, or it will not be available when the node is running as Secondary. The physical interface and IP address should be used in this situation.</p>

Operating System	Feature or Category	Description
All	Stripe group expansion	StorNext does not support expansion on stripe groups containing mixed-sized LUNs. For example, if you create a file system that has two different-sized disks in a userdata only stripe group and then attempt to add a new disk to that stripe group and expand it, the expansion will fail.
	dpserver	<p>In some cases the physical IP address must be included in the dpserver file in addition to the interface name. Note these conditions:</p> <ul style="list-style-type: none"> <li>• When there is one IP address associated with a NIC interface, the interface name alone is a sufficient identifier</li> <li>• If there are multiple IP addresses associated with a NIC interface, one IP address is required in addition to the interface name</li> <li>• On HA systems, the physical IP address is required if virtual IP is configured for the NIC interface. (See also the following entry, "Distributed LAN Clients in HA Environment.")</li> </ul>
	Truncation	By design, replication or deduplication must be completed before data files can be truncated if these files are associated with both a replication/dedup policy and a Storage Manager policy. Even if the Storage Manager policy is configured with the "Truncate Immediately" option, the truncation may not occur at store time unless the file has been replicated or deduplicated.
	DXi Virtual Tape Library Compatibility	<p>Note the following recommendations and limitations for using DXi as a virtual tape library for StorNext:</p> <ul style="list-style-type: none"> <li>• Recommended library emulation: "ADIC Scalar i2000"</li> <li>• Recommended tape drive emulation: "IBM LTO-x" or "HP LTO-x"</li> <li>• DDM (Distributed Data Mover): This feature is currently not supported due to lack of full SCSI3 support in DXi.</li> </ul>
	Affinities	<p>When a file system with two affinities is to be managed by the Storage Manager, the GUI forces those affinities to be named tier1 and tier2. This will cause an issue if a site has an existing unmanaged file system with two affinities with different names and wants to change that file system to be managed. There is a process for converting a file system so it can be managed but it is non-trivial and time consuming. Please contact Quantum Support if this is desired.</p> <p><b>Note:</b> 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>.</p>

## Documentation

The following documents are currently available for StorNext products:

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

## Contacting Quantum

More information about this product is available on the Quantum Service and Support website at [www.quantum.com/ServiceandSupport](http://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](http://www.quantum.com/ServiceandSupport/Upgrade/Index.aspx). For further assistance, or if training is desired, contact Quantum Global Services:

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

For the most updated information on Quantum Global Services, please visit: [www.quantum.com/ServiceandSupport](http://www.quantum.com/ServiceandSupport)

