Quantum

Upgrade Guide

StorNext 4.3.x

Contents

Upgrading HA Systems To StorNext 4.3.x (.RPM)   26

Upgrading Clients To StorNext 4.3.x   28

Contacting Quantum   31

Introduction

This document provides information about upgrading to StorNext 4.3.x from a previous version. Instructions are provided for metadata controllers (MDCs) running StorNext File System (SNFS) and StorNext Storage Manager (SNSM).

For upgrade procedures, or to learn how to contact Quantum for assistance, see the following sections:

- Getting Ready to Upgrade
- Running the Pre-Installation Script
- Upgrading non-HA Linux MDCs to StorNext 4.3.x
- Upgrading HA Systems To StorNext 4.3.x
- Upgrading HA Systems To StorNext 4.3.x (.RPM)
- Upgrading HA Systems To StorNext 4.3.x (.RPM)
- Contacting Quantum

Getting Ready to Upgrade

This section includes important information about upgrading to StorNext 4.3.x. Review the following sections before upgrading:

- Upgrade Guidelines
- StorNext Licensing
- Upgrading the Client Software
- DDisk Support Discontinued
- StorNext Media on DVD
- Post-Upgrade Browser Refresh
- Considerations for Downloading the Installation Files from the Web
- Performing a test conversion of a Linter database to MySQL
- Rolling back a StorNext 4.3.x version to a Linter-based version

Note: You can upgrade directly to StorNext 4.3.x only from StorNext 4.1.x. If you are on a StorNext release earlier than 4.1.x, you must first upgrade to 4.1.x and then to StorNext 4.3.x.

Upgrade Guidelines

Before upgrading StorNext, review the following upgrade guidelines:

> **Note:** For customers with large Linter databases, Quantum strongly recommends consulting Quantum Professional Services prior to attempting to upgrade to StorNext 4.3.x.

> **Note:** The **snPreInstall** script must be run prior to **install.stornext**.

- Beginning with 4.3.0, StorNext uses a MySQL database for information needed internally by Storage Manager.
- The MySQL database requires the following:

  1 A minimum of 400 MB of RAM.

  2 Approximately 2.5 times the amount of disk space of a current Linter database for both the converted database and backups.

  3 During migration, up to 4 times the amount of disk space of the Linter database.

  4 The system being migrated MUST be offline during the migration.

> **Note:** The upgrade process migrates the content of the existing Linter database to a new MySQL database used by this version of StorNext and later, and may take an extended period of time to complete. Figure 1 shows an estimate of the conversion time based on the file size of a current Linter database snbackup.
>
> Alternately, to get a sense of the amount of time required, you may wish to do a test conversion on a separate system. See Performing a test conversion of a Linter database to MySQL for this procedure.

Figure 1 Linter to MySQL conversion estimate based on snbackup file size
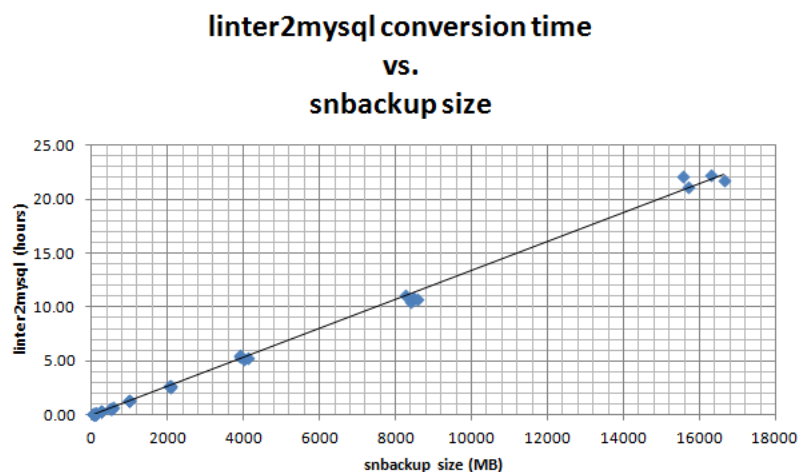
> **Note:** The StorNext MDC cannot be used during the database migration.

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

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

Table 1 Space requirements from **snPreInstall**

Component	Directory Space Required StorNext 4.3.x	Directory Space Required StorNext 4.2.0 (for comparison)
Database	231 MB	64 MB
Journal	11 MB	3 MB
Mapping	157 MB	157 MB
Metadata	333 MB	195 MB
MySQL_tmp	120 MB	N/A
Total	**852 MB**	**419 MB**

- If you are planning on upgrading your RHEL5 to RHEL6, refer to Replacing an MDC in a non-HA environment (non-backup/restore method)" in the "Replacing an MDC" chapter of the StorNext User's Guide.

- Make sure the MDC uses a supported operating system and platform, and if necessary update to a supported release or kernel version before installing StorNext.

Note: A fix is required prior to installing the StorNext file system and Storage Manager. The fix avoids potential clock jumps and temporary system freezes on some systems. This fix only affects RedHat Linux releases 6.0, 6.1, 6.2 and their CentOS counterparts on systems with Intel processor versions identified by the internal codename Nehalem.

Follow the steps below to modify the grub.conf file so that the Intel sleep state is disabled. Making this change could result in increased power consumption, but it helps prevent problems which result in system hangs due to processor power transition.

1 For the above systems, prior to installation:

Add the following text to the "kernel" line in /boot/grub/:

```
grub.conf:Idle=polling intel_idle.max_cstate=0
processor.max_cstate=1
```

2 Reboot the system for the change to take effect.

- Before upgrading, make sure there is no I/O activity. On the StorNext home page, the number of store candidates should be zero. For tape drives, make sure no media is mounted and no read/write operations are being performed. If there is I/O activity, wait for the I/O activity to complete before

upgrading. If you attempt to perform an upgrade during a period of high activity, the upgrade may fail.

- Before upgrading, 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.

- Before upgrading, Quantum recommends reviewing system configuration settings and performing a backup procedure to back up StorNext configuration files.

- Before upgrading, Quantum recommends storing all files. (If you do not store files prior to upgrading, Quantum recommends running a rebuild policy after the upgrade.)

- Before upgrading, Quantum recommends making sure all tape drives are at approved firmware levels. After upgrading, you may not be able to use drives that are not running an approved firmware version.

- A metadata dump is automatically performed during the upgrade process. Depending on the size of the file system, this can extend the time required to complete the upgrade.

Note: The performance of the storage holding the Storage Manager database and metadumps has an impact on overall StorNext file system performance. Quantum recommends that the Storage Manager database and metadata dumps are stored on high performance storage devices, such as a RAID array. Storing the Storage Manager database on slower storage devices, including local disks, can result in poor performance or time-outs.

Note: The **install.stornext** script will scan the rpms that are to be installed to ensure that all dependencies will be successfully resolved. Unresolved dependencies will result in a warning from the **install.stornext** script, and prompt the user to either quit or continue. Continuing without resolving the dependencies may result in system instability or StorNext not being able to operate. It is strongly recommended that the missing dependencies are resolved prior to installing StorNext.

Consult your operating system's package repository for the missing dependencies. A common missing dependency with 4.3.0 is the **perl-DBI** package which is not installed by default on some Linux distributions. Other required dependencies include **gcc**, **make** and **kernel-headers** on Redhat or **kernel-source** on SuSE.

Note: The UIDs for the **quantumdb** and **tdlm** users and the **adic** group must be identical on both nodes of an HA pair. If the UIDs are not identical, the MySQL database will not start (due to file permission errors), which in turn prevents storage manager from fully starting up. Quantum recommends creating the UIDs for the **quantumdb** and **tdlm** users along with the **adic** group on both nodes prior to running **install.stornext**.

StorNext Licensing

Beginning with StorNext 4.0, licensing has changed significantly. Separate licenses are now required for different StorNext components (such as File System and Storage Manager) and features (such as Replication and Distributed Data Mover.)

For a list of all the StorNext licenses, refer to the current *StorNext Licensing Guide* posted here (click the "Select a StorNext Version" menu to view the desired documents):

http://www.quantum.com/ServiceandSupport/
SoftwareandDocumentationDownloads/SNMS/Index.aspx?whattab=Third

Upgrade Implications

The new licensing implementation affects StorNext upgrades, both for release 4.3.x and future releases. Be aware of the following upgrade-related implications and plan accordingly:

- A non-expired Maintenance license is required to perform a StorNext upgrade. This means you must contact Quantum Technical Support for a Maintenance license before you can upgrade to StorNext 4.3.x.

- The Maintenance license provided by Quantum Technical Support must be put into place prior to the upgrade, or you will not be allowed to proceed with the upgrade. You can either enter licenses through the StorNext GUI or manually edit the `license.dat` file. This applies to upgrades to StorNext 4.3.x and also to future upgrades.

- For any StorNext feature or component currently in use, you must have a license in place prior to the upgrade. Otherwise, you will not be allowed to proceed with the upgrade.

- After an upgrade you will still be allowed to run StorNext if the Maintenance license expires. However, no future upgrades will be allowed.

- The Maintenance license must remain in place even after expiration to allow the StorNext software to run, because it indicates which version of the software was purchased.

- When you extract client/server .bin files, you must read and accept the new StorNext End User License Agreement (EULA).

Freeing Media Capacity

If you are ready to upgrade and then notice that the Storage Manager capacity has been exceeded, you can follow the procedure below to free up capacity to bring it under the licensed value. These steps will clean up "dead space" on tape media, and do not apply to storage disks.

1 Run the **fsmedcopy** command with no arguments. Running command generates a report of media and wasted space. The report looks similar to this:

```
% fsmedcopy
---------------------------------------------------------
 Media Fragmentation Report          Thu Feb  4 15:18:14 2010

 Media ID          Fill Level     Wasted Space
---------------------------------------------------------
```

```
000099              10.00              0.00
000098              30.00              5.00
000096              70.99             44.98
000095               0.12             99.65
```

The "Fill Level" column shows how full the media is. The "Wasted Space" column indicates how much of the used space is wasted, NOT the percentage of the entire tape.

Media with high percentages for both values (such as media 000096 in the example) are good candidates for eliminating dead space.

2 For each media you identify as having significant dead space, run this command:

```
% fsmedcopy -r <mediaid>
```

> **Note:** On a large system, the **fsmedcopy** report can take a long time to complete. If your site is dynamic (that is, it has a lot of file removes and updates,) you can skip step 1 and just pick media you know are full, and then run the command **fsmedcopy -r** against those media.

Steps for Obtaining License Strings

To obtain StorNext component and feature licenses, see the "Obtaining Licenses" section in the *StorNext Licensing Guide*.

Upgrade Summary Steps For Non-HA Systems

Taking into consideration the new licensing procedures, here is a summary of the steps necessary to upgrade to 4.3.x:

1 Obtain the necessary licenses from Quantum Technical Support.

2 Update the `license.dat` file with a current Maintenance license and any other StorNext component/feature licenses you have purchased.

3 Run the upgrade program as explained in .

4 Log in to the StorNext GUI.

5 Use the StorNext Configuration Wizard to finish configuring your system.

Upgrading the Client Software

To upgrade to the StorNext 4.3.x client software on Linux and Unix systems, install the new client software package. There is no need to uninstall previous versions of the client software. (On Solaris systems, you will need to uninstall previous versions before upgrading.) For more information about installing the client software on Linux and Unix, see the current *StorNext Installation Guide*.

The steps required to upgrade clients to StorNext 4.3.x includes the following:

1 Unmount all StorNext file systems.

2 Stopping the cvfs services by running /etc/init.d/cvfs stop, /sbin/init.d/cvfs stop, or some other method.

3 Unpack the new build.

4 Unload and load or upgrade `rpm/package/builds`.

5 Restart services.

6 Check cvadmin/cvlabel.

7 Mount StorNext file systems

DDisk Support Discontinued

Beginning with StorNext 4.0, DDisk is no longer supported due to conflicting technology. *Any existing DDisks must be removed and will not be transferred to 4.3.x when you upgrade.* For assistance removing DDisks, contact Quantum Technical Support. See Contacting Quantum on page 31.

StorNext Media on DVD

Installation and upgrade media is shipped on DVDs. (Media was previously shipped on CDs.) If you plan to install from media, you must have a DVD ROM drive to perform the installation or upgrade.

Post-Upgrade Browser Refresh

Quantum recommends to clear the browser and cookie cache after upgrading a StorNext release before logging into the GUI on the upgraded system.

The steps used to clear the cache depend on the browser:

For Google Chrome:

1 Click the configuration icon and the **Settings option.**

2 Click the "Show advanced settings" option.

3 In the **Privacy** section, click **Clear browsing data**.

4 Select "the beginning of time" option in the "Obliterate the following items from" menu.

5 Select the "Clear browsing history", "Empty the cache", and "Delete cookies and other..." options.

6 Click **Clear browsing data**.

7 Close the **Settings** tab.

For Internet Explorer:

1 Go to **Tools -> Internet Options**.

2 On the **General** tab, look in the **Browsing History** section and click **Delete**.

3 Select the **Temporary Internet Files** and **Cookies** options.

4 Click **Delete**.

5 Click **OK** to close the **Internet Options** dialog.

For Mozilla Firefox:

1 Click **Tools -> Clear Recent History**. (If this option is not visible, right-click the top of the browser window and select the "Menu Bar" option).

2 In the **Time Range to Clear** menu, select "Everything".

3 Click **Details**, and select the **Browsing & Download History**, **Cookies**, and **Cache** options.

4 Click **Clear Now**.

Considerations for Downloading the Installation Files from the Web

The upgrade instructions in this section are for DVD upgrades. If you download your StorNext installation software from the Web, your steps will be slightly different. You will download an executable program that will require you to accept the StorNext end user license agreement and will then extract the StorNext installer.

As part of the installation from a Web download, the software files are extracted into the **/tmp/stornext** directory by default. The software, when extracted, requires approximately 400 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the -d option.

To specify an alternate location before beginning the installation, enter the following command:

```
# ./<installation file> -d <dir>
```

where *<installation file>* is the name of the installation file, and *<dir>* is the directory into which you want to extract the installation files.

Performing a test conversion of a Linter database to MySQL

Note: Prior to performing any rollback, Quantum strongly recommends contacting Quantum Technical Support. See Contacting Quantum on page 31.

StorNext 4.3.x uses a MySQL database to store internal information needed by the Storage Manager layers of StorNext rather than the Linter database used in previous versions of StorNext. As part of the upgrade from previous versions of StorNext to StorNext 4.3.x, existing databases will need to be converted to the new database format.

This procedure explains how to take an existing Linter database on a **production** StorNext system and convert it to a MySQL database on another system, which we will call the **migration** system. The conversion procedure documented here helps determine approximately how much time the database conversion will take when upgrading to StorNext 4.3.x, without impacting the **production** system.

Assumptions

- The Upgrade Guidelines on page 2 have been followed.

- The **production** system is already installed and configured.

- The **production** system should be at StorNext release level that is on the supported upgrade path to 4.3.x (upgrades to StorNext 4.3.x can be done from StorNext 4.1.0 up to 4.2.2).

- The **migration** system's hardware is very similar to the **production** system.

Note: This procedure is intended only as a planning aid. After the test conversion is run, StorNext must be removed from the **migration** system and the migrated database must be discarded.

On the Production System

1 Start by taking a full backup of the **production** system. The time that the backup is taken and its backup ID should be noted:

   snbackup

2 The snbkpreport can be used to determine the latest backup ID if necessary:

   snbkpreport

   Backup files are stored to the file system identified by the BACKUPFS **sysparm**:

   showsysparm BACKUPFS

   The backups are kept in the **$BACKUPFS/.ADIC_INTERNAL_BACKUP** directory.

   Copy the backup from the **production** system to the **migration** system. The **migration** system will need the **db.$ID.0.tgz** and **db.$ID.1.tgz** files where **$ID** is the backup ID from Step 1:

   scp $BACKUPFS/.ADIC_INTERNAL_BACKUP/db.$ID.*.gz \

   $MIGRATION_HOST:/path/to/linter_backup_dir

On the Migration System

3 Install StorNext 4.2.x onto a clean, pristine system that will be used to migrate Linter into MySQL. The version of StorNext must match the version of StorNext installed on the **production** system.

   cd /path/to/4.2.x/phdist/SYSTYPE

   ./install.stornext

   . /usr/adic/.profile

4 Restore the Linter database from the **production** system's backup taken in Step 1 and copied over in Step 2. Enter the backup ID when prompted:

   snrestore -d -r /path/to/linter_backup_dir

   Ignore the following backup manifest errors:

   ERR: No backup manifest file found

   ERR: No device manifest file found

5 **snrestore** will stop all of Storage Manager. Following the **snrestore**, make sure that Linter is started. Linter needs to start and allow for the its journals to be replayed. At the command prompt, type:

   database_control start

   Verify that Linter has completely started up by verifying that the following line is in the **/usr/adic/database/db/linter.out** file corresponding to the process ID of the running linter process. At the command prompt, type:

   # pidof linter

29827

tail /usr/adic/database/db/linter.out

17.07.2012 17:56:39 Kernel system parameters: MBX - "20566", Pid - 29827

17.07.2012 17:56:39Copyright (C) 2000-2004 RelexUS, Inc. All rights reserved.

Copyright (C) 1995-2003 Relex, Inc. All rights reserved.17.07.2012 17:56:39

17.07.2012 17:56:39 *** RDBMS Linter is running

17.07.2012 17:56:39 *** Press <ENTER> for shell prompt

17.07.2012 17:56:39

6 Install the **mysql** component:

cd /path/to/4.3.x/phdist/SYSTYPE

./install.stornext -mysql

7 Shutdown StorNext. Enter the following:

. /usr/adic/.profile

service cvfs stop

8 Verify that the Linter database contains no errors. testdb must be run while Linter is stopped. Enter the following:

testdb -p 10000

If **testdb** finds any errors in the database, it will report some suggested commands to run.

Note: **DO NOT RUN THESE COMMANDS** Instead, contact Quantum Support if **testdb** reports any errors.

If **testdb** runs successfully, it will print a message that states:

Database Testing Report

- Database contains no errors and can be used.

9 Start Linter and MySQL. Enter the following:

database_contol start

mysql_control start

10 Run the migration utility:

Note: Be sure to note the time the conversion starts and the time the conversion completes. This will give you a good approximation of the time it will take to convert the database on your production system.

/usr/bin/time -v /usr/adic/mysql/util/migrate/linter2mysql

Note: If the migration does not complete, you will see a message similar to the following:

```
Database migration error detected. Upgrade aborted see
/usr/adic/install/stornext.install.log for details
```

If this occurs, attempt the migration again. If this fails a second time, contact Quantum Technical Support. See Contacting Quantum on page 31.

11 Remove StorNext from the **migration** system. Enter the following:

```
cd /path/to/4.3.x/phdist/SYSTYPE
./install.stornext -removeall
```

Rolling back a StorNext 4.3.x version to a Linter-based version

This section provides a general overview of the procedure for rolling back a StorNext 4.3.x MySQL-based installation to StorNext 4.1.X or 4.2.X Linter-based code. The purpose of this overview is to give a general high level understanding of what the steps would be to do the rollback. It is not intended to be the actual detailed procedure for how to actually accomplish a rollback.

Requirements

The following items are what would be required before doing the rollback from StorNext 4.3.x:

- The Upgrade Guidelines section recommends a full backup of the previous StorNext system be done prior to upgrading StorNext to 4.3.x. The rollback procedure requires this backup of the Linter database and configuration information necessary to run the StorNext 4.1X or 4.2X system.

- A StorNext 4.1X or 4.2X software distribution.

Rollback tasks

A rollback from StorNext 4.3.x to StorNext 4.1X or 4.2X releases will be basically a disaster recovery type restoration. The rollback would use the same procedures that would be followed for any other StorNext downgrade from one release back to another.

The basic steps involved in the rollback are:

1 Stop the StorNext 4.3.x system(s).

2 Uninstall StorNext 4.3.x and install StorNext 4.1X or 4.2X according to whether it is a HA or non-HA configuration.

3 Restore the configuration and database from the backup done prior to the StorNext 4.3.x upgrade. Note that the managed file systems will still be intact and therefore no need to restore the metadumps from the backup.

4 Run the **fspostrestore** to resync the managed file systems with the recovered Linter database.

5 Complete the startup of all the other Storage Manager services.

Actions on files during fspostrestore

The following actions will be performed on files which are on a managed file system and under the control of Storage Manager:

- The current version of **fspostrestore** will remove **truncated** files from the file system if those files were created/stored/truncated after the database backup creation date.

- If a newer version of a **truncated** file was created and stored after the database backup creation date, the active version of the file will be reverted back to the last version that the restored database knows about.

- Disk-resident files that are either not known to Linter or are a newer version of a file already contained within the restored Linter database will be restored by Storage Manager automatically. This means that the file will be added back to the appropriate store candidates list so that it will be fully stored under Linter.

Precautions

- If a modification was done to a managed file system configuration or to Storage Manager after upgrading to StorNext 4.3.x, there may be issues with rolling back. A configuration change such as adding Directory Quotas will make it difficult or impossible to roll back and/or involve more data loss.

- Be sure that StorNext 4.3.x is functioning in a manner that is acceptable before enabling any StorNext 4.3.x specific features.

- If gateway reporting was running on both the StorNext 4.2X system and StorNext 4.3.x, that information will be lost when a rollback is performed since the gateway database information is not part of a normal StorNext 4.2X backup.

- Rollback from 4.3.x to 4.1X/4.2X is not supported for systems with replication and/or deduplication configured.

Running the Pre-Installation Script

The StorNext pre-installation script (**snPreInstall**) is included on the StorNext installation DVD. When you run **snPreInstall**, you are prompted for information about your system. The pre-installation script uses this information to estimate the amount of local disk space required for SNFS and SNSM support directories. In addition, the script recommends the optimal locations for support directories.

StorNext uses four directories to store application support information. These directories are stored locally on the server.

The StorNext support directories are described in Table 2.

Table 2 StorNext Support
Directories

Support Directory*	Description
Database /usr/adic/database	Records information about where and how data files are stored.
Journal /usr/adic/database/ journal	Records changes made to the database.
Mapping /usr/adic/TSM/ internal/mapping_dir	Contains index information that enables quick searches on the file system.
Metadata /usr/adic/database/ metadumps	Stores metadata dumps (backups of file metadata).

* Depending on file system configuration, the support directory locations may be different. However, symbolic links in **/usr/adic/** will always point to the support directories.

Note: For optimal performance, the **Metadata** directory should not be located on the same file system as (in order of priority) the **Journal**, **Database**, or **Mapping** directory.

Before You Begin

Before running the pre-installation script, be prepared to answer the following questions:

- Is this an upgrade installation?
- What local file systems can be used to store support information?
- What version of StorNext will be installed?
- What is the maximum number of directories expected (in millions)?
- What is the maximum number of files expected (in millions)?
- How many copies will be stored for each file?
- How many versions will be retained for each file?

Note: Keep in mind that storage needs typically grow rapidly. Consider increasing the maximum number of expected directories and files by a factor of 2.5x to ensure room for future growth.

Note: The pre-installation script ignores unmounted file systems. Before running **snPreInstall**, be sure to mount all local file systems that will hold StorNext support information.

Running snPreInstall

To run the pre-installation script, use the StorNext installation DVD.

1 Log on to the MDC as **root**.

2 Mount the StorNext installation DVD and change to the DVD root directory.

> **Note:** When you mount a DVD in a Red Hat 5 system, DVDs are mounted by default with a noexec (non-executable) option which prevents you from proceeding with the installation.
>
> For Red Hat users only, before proceeding you must remount the DVD by typing **mount -o remount, exec ...**

3 List the installation directories on the DVD. At the command prompt, type:

`ls -l`

4 Change to the **stornext_full** directory.

> **Note:** Make sure you are in the **stornext_full** directory and not the **fs_only** directory. The **fs_only** directory installs the StorNext File System only.

5 List the directories in the **stornext_full** directory. At the command prompt, type:

`ls -l`

6 Identify the correct installation directory for your operating system and hardware platform, and then change to that directory.

For example, for Red Hat Linux 5 running on an x86 64-bit platform, change to the **RedHat50AS_26x86_64** directory.

7 Run the script. At the command prompt, type:

`./snPreInstall`

The pre-installation script runs (Figure 2).

Figure 2 Pre-Installation Script

```
----------------------------------------------------------------
This script will determine the disk space requirements for
performing a StorNext install/upgrade. It will then report
whether or not the local system has the required disk space
for the operation. Output will be one of these four types:

   RECOMMEND: A recommendation of how to set up the
              StorNext directories on your system.
   NOTE:      An advisory message indicating that the system
              can run with the given configuration, at the
              desired capacity, but the configuration
              is not optimal.
   SPACE:     Due to space issues the system cannot be
              configured to run at the desired capacity
              without adding more disk space.
   Aborting:  An error has occured, or invalid input has been
              provided, and no further analysis can be done.

The script will exit with one of these values:
     0 - disk space fine, ok to proceed with install/upgrade
     1 - disk space available, but not optimal file system layout
     2 - disk space not available for install/upgrade
     3 - invalid args etc. kept script from completing normally

----------------------------------------------------------------
Are you upgrading ?? [y/<n>]:
```

8 Follow the onscreen instructions, and answer each question when prompted.

The script outputs results based on the information you entered.

Interpreting snPreInstall Output

After you enter all requested information, the pre-installation script outputs the following results:

- Estimated disk space required for each support directory.
- Recommended file system location for each support directory.

Note: For optimal performance, each support directory should be located on its own local file system, and each local file system should reside on a separate physical hard disk on the server computer.

There are three types of directory recommendations, as indicated by the text preceding the recommendation. The three types of recommendations are described in Table 3.

Table 3 Pre-Installation Script Recommendation Types

Type	Description
RECOMMEND	The suggested directory locations will result in optimal system performance.
NOTE	The suggested directory locations are valid but system performance will not be optimal.
SPACE	Additional hard disk space is required to store the required support directories.

The pre-installation script bases directory location recommendations on the following criteria:

- To aid disaster recovery, the **Database** and **Journal** directories should be located on different file systems.
- For optimal performance, each support directory should be located on its own local file system, and each local file system should reside on a separate physical hard disk in the MDC.

Note: These recommendations apply to non-HA systems only.

Upgrading non-HA Linux MDCs to StorNext 4.3.x

Use the StorNext installation script to upgrade to StorNext 4.3.x on a metadata controller running StorNext Storage Manager and StorNext File System. The installation script will upgrade full StorNext installations. To launch the script, use the StorNext installation DVD.

The following instructions apply to non-HA server upgrades. For instruction on upgrading HA servers only, see Upgrading HA Systems To StorNext 4.3.x on page 21.

As always, you should perform a backup before beginning the upgrade.

> **Note:** You can upgrade directly to StorNext 4.3.x only from StorNext 4.1.x or later. If you are on a StorNext release earlier than 4.1.x, you must first upgrade to 4.1.x and then to StorNext 4.3.x.

1 Log on to the MDC as **root**.

2 Stop Storage Manager. At the command prompt, type:

   `adic_control stop`

3 Verify that the Linter database contains no errors. **testdb** must be run while Linter is stopped. At the command prompt, type:

   `# testdb -p 10000`

   If **testdb** finds any errors in the database, it will report some suggested commands to run.

> **Note:** **DO NOT RUN THESE COMMANDS** Instead, contact Quantum Support if **testdb** reports any errors.

   If **testdb** runs successfully, it will print a message that states:

   `Database Testing Report`

   `------------------------`

   `-  Database contains no errors and can be used.`

4 Mount the StorNext installation DVD and change to the DVD root directory.

> **Note:** When you mount a DVD in a Red Hat 5 or 6 system, DVDs are mounted by default with a noexec (non-executable) option which prevents you from proceeding with the installation.
>
> For Red Hat users only, before proceeding you must remount the DVD by typing **mount -o remount, exec ...**

5 List the installation directories on the DVD. At the command prompt, type:

   `ls -l`

6 Change to the **stornext_full** directory.

> **Note:** Make sure you are in the **stornext_full** directory and not the **fs_only** directory. The **fs_only** directory installs the StorNext File System only.

7 List the directories in the **stornext_full** directory. At the command prompt, type:

   `ls -l`

8 Identify the correct installation directory for your operating system and hardware platform, and then change to that directory.

For example, for Red Hat Linux 5 running on an x86 64-bit platform, change to the **RedHat50AS_26x86_64** directory.

9 As described in StorNext Licensing on page 5, license checks are made as part of the upgrade. Some of these checks require the software to be running, so start it now. At the command prompt, type:

```
service cvfs start
```

10 Run the script with the **-liccheck** option. At the command prompt, type:

```
./install.stornext -liccheck
```

The install script runs and checks for valid licenses. If the license check fails, then contact Quantum before continuing with the upgrade. For sample output, see The License Check Option on page 18.

11 Run the install script. At the command prompt, type:

```
./install.stornext
```

Note: The upgrade process migrates the content of the existing Linter database to a new MySQL database used by this version of StorNext and later, and will take an extended period of time to complete. To get a sense of the amount of time required, you may wish to do a test conversion on an alternate system. See Performing a test conversion of a Linter database to MySQL.

The installation script runs (Figure 3).

Figure 3 StorNext Installation Script


```
    Stornext Upgrade              MAIN MENU
    [U.S. Pat. No. 5,990,810]

    1)   Installation Configuration
    2)   Upgrade stornext
    3)   Show Status
    4)   Quit

Enter your choice <2>: █
```

12 On the Upgrade Menu, type **2** and press **<Enter>**.

The StorNext installation script runs and displays the current progress of the upgrade.

13 When the upgrade is complete, press **<Enter>** to return to the Main Menu.

Note: If the migration does not complete, you will see a message similar to the following:

```
Database migration error detected. Upgrade aborted see
/usr/adic/install/stornext.install.log for details
```

If this occurs, attempt the migration again. If this fails a second time, contact Quantum Technical Support. See <u>Contacting Quantum</u> on page 31.

14 To exit the installation script, type **4** and press **<Enter>**.

15 After running the installer, source the profile on the MDC by running this command:

```
. /usr/adic/.profile
```

The License Check Option

When you run the installation script with the **-liccheck** option, the script checks for valid licenses. Proceed with the upgrade only if the license check is successful.

Note: Even if the license check is successful, messages about unlicensed features may display. However, if these features are not in use, you may proceed with the upgrade.

The license check can fail for the following reasons:

- StorNext or StorNext components are not running, which keeps the licensing status from being determined for a feature. Therefore it is treated as non-licensed even if the license is present.

- The capacity has been reached for a licensed feature so it is treated as non-licensed.

- A feature is found to be in use but not licensed.

- The maintenance license expired prior to the build date of the upgrade release.

See the sections below for sample output for a successful and a failed license check.

Successful License Check Sample Output

```
+++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
- The server license status is: Good
- The proxy license status is: Good
- The dedup license status is: Unused
- The replication license status is: Unused
- The failover license status is: Unused
- The manager license status is: Unused
- The checksum license status is: Unused
- The vaulting license status is: Unused
- The sdisk license status is: Good
- The ddm license status is: Good
```

```
   - The maintenance license status is: Good
License check done.
+++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
```

Failed License Check Sample Output

```
+++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
   - The server license status is: Good
   - The proxy license status is: Good
   - The dedup license status is: Good
   - The replication license status is: License Missing
   - The failover license status is: Unused
   - The manager license status is: Over Capacity
   - The checksum license status is: Unused
   - The vaulting license status is: Unused
   - The sdisk license status is: Good
   - The ddm license status is: UNKNOWN
   - The maintenance license status is: Good
   ----------------------------------------------------------------------
   Improperly licensed features have been detected on the system. Please
   contact Support to get licenses for the features indicated above that
   are not properly licensed and are in use.
   ----------------------------------------------------------------------
      --- Number of improperly licensed StorNext features: 3
      --- Verify the following two items before contacting Support. If
          these things were not done, correct and try again.
          -> The license file is installed and up to date.
          -> The software is up and running as this is needed
             for some license checks.

+++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
```

SSL Certificate Exception

When you log into StorNext for the first time, you might see a message warning you about a security certificate. Refer to the Quantum Knowledge Base and search for "SSL" for a permanent workaround to this issue.

For a temporary solution, create a certificate exception that will allow you to log into the StorNext GUI without seeing the warning message during subsequent logins.

Upgrading a Windows Non-HA Client or Server to StorNext 4.3.x

This section describes how to install StorNext on Windows non-HA clients or servers. During an upgrade, the old software is removed and the new software is installed. Your previous configuration settings remain intact.

For StorNext clients you do not require licenses. For StorNext MDCs you will need Maintenance, Server and (optionally) Distributed LAN client licenses.

> **Note:** You can upgrade directly to StorNext 4.3.x only from StorNext 4.1.x or later. If you are on a StorNext release earlier than 4.1.x, you must first upgrade to 4.1.x and then to StorNext 4.3.x.

1 If any StorNext file systems are currently shared, temporarily disable sharing.

2 Copy one of the following files from the StorNext installation DVD and place it on the system onto which you want to install it. Be sure to select the file from the **fs_only** directory and not the **stornext_full** directory.

- **fs_only/snfs_client_Windows_i386.zip** (for 32-bit systems, only available for client upgrades)

- **fs_only/snfs_client_Windows_x86_64.zip** (for 64-bit systems)

3 Unzip the file.

4 Open the unzipped folder and double-click the client software installer file. This file is named **SnfsSetup.ex**e.

The **StorNext Installation** screen appears (Figure 4).

Figure 4 StorNext Installation Window: Upgrade/Reinstall StorNext
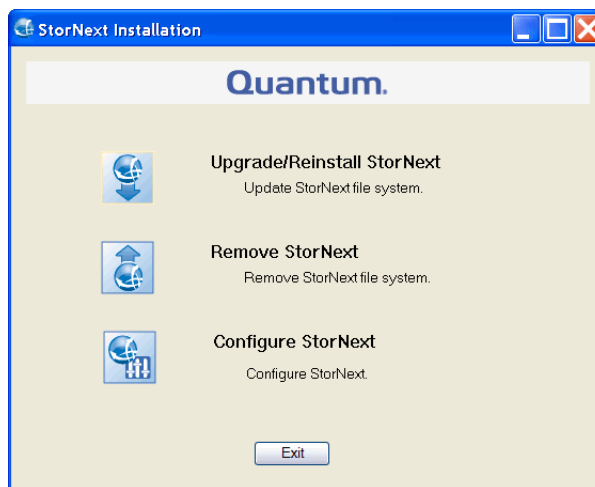
5 Click **Upgrade/Reinstall StorNext**.

The **StorNext Component Selection** window appears (Figure 5).

> **Note:** If you are upgrading the client, you will not have the **Server** or **Install Licenses** options.

Figure 5 Component Selection

6 Select the check boxes of the components you want to include in the upgrade. Clear the check boxes of the components you want to exclude. Click **Next** to proceed.

The StorNext End User License Agreement (EULA) appears.

7 Read the EULA. If you accept, select "I accept the terms in the License Agreement" and then click **Next** to continue the install.

It make take several minutes for the selected components to finish being added or removed.

8 When a message informs you that the installation is finished, click **OK**.

9 In some cases you may be required to reboot your system. In these situations StorNext attempts to complete the installation automatically when the system starts back up. However, some platforms, virus checkers, et.,c prohibit the automatic restart feature. If the installation is not automatically completed (i.e., there is no **Start->All Programs->StorNext** menu), restart the **SnfsSetup.exe** program manually and follow the on-screen directions.

10 After the upgrade is complete, re-enable sharing on your StorNext file systems.

Upgrading HA Systems To StorNext 4.3.x

This section describes how to upgrade an HA system to StorNext 4.3.x.

Before you begin this or any other upgrade procedure, **verify that you have the necessary licenses required to upgrade**. If you do not have the required licenses, the upgrade may fail and you will be unable to continue.

HA Upgrade Options

The upgrade procedure requires a certain level of technical expertise, including running commands from the command line. You should also have a solid understanding of the HA feature in StorNext 4.3.x and a general familiarity with the StorNext GUI.

You should not attempt to perform this procedure unless you have carefully read the upgrade steps and are confident you can complete the steps successfully.

If you are uncertain about your ability to perform the upgrade unassisted, please contact Quantum Professional Services for assistance in upgrading either remotely or on site.

IMPORTANT: If you decide to perform the upgrade yourself and you encounter an error during the upgrade which is not caused by a bug in the upgrade installer, Quantum's Profession Services group must be consulted to come to your location to correct the issue, and you will be billed for their services. Note that Professional Services representatives cannot travel without a purchase order, and normal scheduling applies. Please consider this before you decide to perform an HA upgrade on your own.

Upgrading a Linux 4.1.x or later HA System to 4.3.x HA

This procedure describes how to upgrade from a StorNext 4.1.x or later HA system to a 4.3.x HA system.

As always, before you begin the upgrade make sure you have obtained the proper licenses for both nodes for the targeted build. Refer to The License Check Option on page 18 if you are not sure you have the necessary licenses for StorNext 4.3.x.

Pre-Upgrade Steps

1 Start StorNext on both MDC servers.

2 Run a manual backup to tape from the StorNext GUI.

3 Make sure all store/retrieve requests have finished.

4 If you are using Distributed Data Mover (DDM) feature, make a notation of the value of the `DISTRIBUTED_MOVING` parameter (either `All` or `Threshold`) in `/usr/adic/TSM/config/fs_sysparm` (or `fs_sysparm_override`).

   Use a text editor to set `DISTRIBUTED_MOVING` to None.

   Use the `adic_control restart` TSM command to put this change into effect.

5 Unmount all file systems from all client machines, and stop the SNFS processes on each client machine. (On Linux platforms, do this by running the `service cvfs stop` command.)

6 Insert the StorNext 4.3.0 or later DVD into the secondary MDC's DVD drive.

7 Use the StorNext 4.3.0 or later installer from the DVD to uninstall Stornext. Run this command:

   `install.stornext -remove`

8 Install StorNext 4.3.x on the standby server by running this command:

```
install.stornext
```

Note: The upgrade process migrates the content of the existing Linter database to a new MySQL database used by this version of StorNext and later, and will take an extended period of time to complete. To get a sense of the amount of time required, you may wish to do a test conversion on an alternate system. See Performing a test conversion of a Linter database to MySQL.

9 Replace the existing /usr/cvfs/config/license.dat file on the standby server with the license.dat file from the primary MDC.

Upgrade the Primary Server

1 Source the profile on both MDCs by running this command:

```
. /usr/adic/.profile
```

2 Make sure the current server is the primary.

3 Make sure the dedicated shared file system is mounted on /usr/adic/ HAM/shared.

4 Set the primary node to "config" mode and the peer node to "peerdown" mode by running the following commands:

```
snhamgr peerdown
```

```
snhamgr mode=config
```

5 Run the command snhamgr status to check the snhamgr status. Status should look similar to this:

```
LocalMode=config
```

```
LocalStatus=primary
```

```
RemoteMode=peerdown
```

```
RemoteStatus=unknown
```

6 Verify that the licenses for both servers in /usr/cvfs/config/ license.dat match.

7 Stop Storage Manager.

Note: The HaShared file system needs to remain mounted for the next step.

At the command prompt, type:

```
adic_control stop
```

8 Verify that the Linter database contains no errors. **testdb** must be run while Linter is stopped. At the command prompt, type:

```
# testdb -p 10000
```

If **testdb** finds any errors in the database, it will report some suggested commands to run.

> **Note:** **DO NOT RUN THESE COMMANDS** Instead, contact Quantum Support if **testdb** reports any errors.

If **testdb** runs successfully, it will print a message that states:

```
Database Testing Report

------------------------

-  Database contains no errors and can be used.
```

9 Run the following command to upgrade to StorNext 4.3.x:

```
install.stornext
```

> **Note:** The upgrade process migrates the content of the existing Linter database to a new MySQL database used by this version of StorNext and later, and will take an extended period of time to complete. To get a sense of the amount of time required, you may wish to do a test conversion on an alternate system. See Performing a test conversion of a Linter database to MySQL.

Pre-Conversion Steps

1 Source the profile on both MDCs by running this command:

```
. /usr/adic/.profile
```

2 Push configuration changes to the mirror and start StorNext on the primary server by running the following command:

```
service cvfs start
```

> **Note:** A post_upgrade script runs which performs a system backup using snbackup if Storage Manager has been configured on this system. Depending on the number of managed files in the system, this backup may take substantial time to complete.

3 Run the following command to confirm that post_upgrade has completed on the primary server:

```
ps -ea | grep post_upgrade
```

The script is complete if there is no output from this command.

> **Note:** Make sure the snbackup process has completed prior to continuing, if Storage Manager has been configured on this system.

4 Verify that all processes are running correctly on the primary server by testing access to all file systems and moving files to/from tape or disk storage.

5 Set the primary node to "default" mode and the peer node to "peerup" mode by running the following commands:

```
snhamgr mode=default
```

```
snhamgr peerup
```

> **Note:** Transitioning from config mode to default mode results in cvfs being stopped.

6 Run this command to start StorNext on the primary server:

```
service cvfs start
```

7 Run the following command to mirror GUI changes to the synchronization directory on the primary server:

```
/usr/adic/util/syncha.sh -primary
```

8 Run this command to start Stornext Web Services on the primary server:

```
service stornext_web start
```

Convert the Secondary Server to HA

1 Log in to the StorNext GUI on the primary server by entering the following into a Web browser:

```
https://<primary_server>
```

2 Choose **Convert** from the **Tools >High Availability** menu.

3 At the **System Name** field, enter the secondary server's address and then click **Scan Host**.

4 If you use Virtual IP (vIP), select the Peer IP interface on the Secondary server.

5 Click **Convert**.

When the confirmation message appears, click **Yes** to convert the secondary server to HA.

Failover Verification

After upgrading you should conduct a failover to verify that the standby server is running properly, and also verify that all processes are running correctly. Test access to all file systems and move files to/from tapes.

The steps to take for testing failover are:

1 On the Primary, run snhamgr -m status, and then check to make sure the output is :default:primary:default:running:

2 Also on the Primary, issue the following command to induce an HA Reset: snhamgr force smith

After you test failover, do the following:

1 Fail back to the primary server if you conducted a failover to the secondary server.

2 Start the SNFS processes on all client machines.

3 Mount the SNFS file systems on each client machine.

4 Verify that all clients have full access.

Post-Conversion Steps

If you are using the DDM feature, follow these steps:

1 If you use the secondary server as a DDM mover, make sure the file systems are mounted.

2 Edit `fs_sysparm` or `fs_sysparm_override` to use your preferred DDM mode: `All` or `Threshold`. Use the command `adic_control restart TSM` to put this change into effect.

Upgrading a Windows 4.1.x or later HA System to 4.3.x HA	To upgrade a Windows HA server to the latest version of StorNext: **1** Navigate to Windows Services, and click **Services** > **Stop and Remove** to stop the File System Managers (FSMs) on the standby machine. **2** Upgrade the standby machine to the new version of StorNext. **3** Click **Services** > **Start** to start the FSMs. **4** From the command line, type the following for all the file systems to fail over the FSMs. `cvadmin fail <fs_name>` **5** Now the upgraded MDC should now be the primary MDC. **6** Repeat above steps 1-3 on the secondary MDC. **7** You can initiate a failover of the MDCs to verify HA is correctly failing over if desired.

Upgrading HA Systems To StorNext 4.3.x (.RPM)

This section describes how to upgrade to a StorNext 4.3.x HA system using an RPM-only installation. Steps are run on both the primary and secondary nodes, and are indicated accordingly.

Upgrading a 4.1.x .or later RPM HA System to 4.3.x .RPM HA	Acquire the StorNext client, server and core 4.3.x RPMs for your server OS architecture from the installation DVD, or by downloading the tarball from the Quantum web site. For example:
	
	`snfs-client-SuSE100ES_261621-4.3.x.12838.x86_64.rpm`
	
	`snfs-server-SuSE100ES-4.3.x.12838.x86_64.rpm`
	
	`snfs-SuSE100ES-4.3.x.12838.x86_64.rpm`

Make sure you have obtained and installed the proper licenses for both nodes for the targeted build (such as 4.3.x). Refer to StorNext Licensing on page 5 for more information about license requirements.

Unmount all file systems from all client machines and stop the StorNext File System processes on each client machine by running /etc/init.d/cvfs fullstop.

On The Secondary MDC

1 Stop StorNext services by running this command:

   service cvfs fullstop

2 Upgrade the StorNext RPMs.

   Example:

   rpm -Uvh snfs-client-<platform/revision>.rpm snfs-server-<platform/revision>.rpm snfs-<platform/revision>.rpm

3 Start cvfs by running this command:

   service cvfs start

4 Verify that all SNFS mounts completed successfully.

On the Primary MDC

1 Stop StorNext services by running this command:

   service cvfs fullstop

2 Verify that the Secondary MDC is now serving the configured file systems.

3 Upgrade to the StorNext RPMs acquired earlier.

   Example:

   rpm -Uvh snfs-client-<platform/revision>.rpm snfs-server-<platform/revision>.rpm snfs-<platform/revision>.rpm

4 Start cvfs by running this command:

   service cvfs start

5 Verify that all SNFS mounts completed successfully

Back On The Secondary MDC

Test failover of the HA system by running the following command:

   /usr/cvfs/bin/cvadmin -e "hamon smith"

If the previous command does not verify the SMITH failover, use the following command to verify that FSM processes failover to the primary MDC.

   /usr/cvfs/bin/cvadmin -e "fail <file system name>"

Upgrading Clients To StorNext 4.3.x

This section describes how to upgrade clients to StorNext 4.3.x. Refer to the instructions pertaining to your particular operating system.

(Procedures for upgrading Windows clients and servers for non-HA systems are described in <u>Upgrading a Windows Non-HA Client or Server to StorNext 4.3.x</u> on page 19.)

Upgrading Linux Clients

Follow these steps to upgrade a Linux client:

1   Stop StorNext services by using this command: `service cvfs stop`. This command unmounts the file systems as well.

2   Extract the appropriate .bin client package (for example, `./snfs_client_RedHat50AS_26x86_64.bin`), and then agree to the StorNext End User License Agreement (EULA).

3   Change to the directory where the new StorNext client package is located.

4   Upgrade the two rpms by using this command: `rpm -U <the stornext rpms>`

5   Restart StorNext services by using this command: `service cvfs start`

Upgrading Solaris Clients

Follow these steps to upgrade a Solaris client:

1 Un-mount the StorNext file system(s) and remove the driver by using this command: `/etc/init.d/cvfs fullstop`.

2 List installed StorNext package with this command:
   `pkginfo | grep ADIC`

   The package is normally called **snfs**.

3 Remove the installed StorNext client package by using this command:
   `pkgrm snfs`

4 Verify StorNext client package has been removed by using this command:
   `pkginfo | grep ADIC`

   No results should be returned.

5 Extract the .bin client package (for example,
   `./snfs_client_SunOS5100i386_64.bin`), and then agree to the StorNext End User License Agreement (EULA).

6 Change to the directory where the new StorNext client package is located. The default location is:
   `/tmp/stornext/`

7 Install the new StorNext client package as described in the StorNext Installation Guide by using this command: `pkgadd -d .`

   Note: The period is part of the command.

8 Reboot the system.

9 Verify the StorNext file system is mounted by running the `mount` command.

Upgrading AIX Clients

Follow these steps to upgrade an AIX client:

1 Un-mount the StorNext file system(s) and unload the driver by using this command: `/etc/rc.cvfs fullstop`

2 Make a copy of the `/etc/filesystems` file to preserve StorNext mount configuration information.

   For example: `cp /etc/filesystems /etc/filesystems.bak`

3 List installed StorNext client package by using this command: `lslpp -L | grep snfs`

   Only one client package should be listed.

4 Remove StorNext client package by using this command: `installp -u -g <file name listed in previous step>`

   For example: `installp -ug snfs_4.3.x.0b21125.rte`

   Note: As part of this step, mount configuration information for StorNext file systems is removed from `/etc/filesystems`.

5 Verify under the Summary section of the removal that the event header displays "de-installed" and the results header displays "success."

6   Use the following command to list installed StorNext client package again to verify all packages has been removed: `lslpp –L | grep snfs`

No client package should be listed.

7   Extract the .bin client package i.e. (./snfs_client_AIX61_powerpc_64.bin) and agree to the StorNext End User License Agreement (EULA).

8   Change to the directory where the new StorNext client package is located.

9   Install the new StorNext client package as described in the StorNext installation Guide by using this command: `installp –ac –d <package name> all`

10 Edit `/etc/filesystems` and re-add mount configuration information for StorNext file systems. Consult the previously saved file `/etc/filesystems.bak`.

11 Start the new StorNext client software by using this command: `/etc/rc.cvfs start`

Alternatively, you can reboot the system.

12 Verify the StorNext file system is mounted by using the `mount` command.

Upgrading HP-UX Clients

Follow these steps to upgrade an HP-UX client:

1   Unmount StorNext file systems by using this command: `umount /stornext/snfs1`

2   Stop all StorNext services by using this command: `/sbin/init.d/cvfs stop`

3   Verify all services are stopped.

4   Uninstall the existing StorNext installed build by using this command: `swremove -x mount_all_filesystems=false snfsclient`

5   Extract the .bin client package (for example, ./snfs_client_HPUX1131_ia64.bin) and then agree to the StorNext End User License Agreement (EULA).

6   Change to the directory where the new StorNext client package is located.

7   Install the new build by using this command: `swinstall -s <location of .depot file> -x mount_all_filesystems=false \*`

For example: `swinstall -s /tmp/stornext/snfs-client_hpux_B.11.31-ia64-4.3.x.0-8872.depot -x mount_all_filesystems=false \*`

8   Verify that StorNext services are running. If necessary, start services by using this command: `/sbin/init.d/cvfs start`

9   If not already mounted, mount the StorNext file systems on the HP-UX client.

Contacting Quantum

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

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

North America	1+800-284-5101 Option 5
EMEA	00800 7826 8888 49 6131 3241 1164
Online Service and Support	www.quantum.com/OSR
Worldwide Web	www.quantum.com/ServiceandSupport

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