



# StorNext 4.0 Upgrade Guide

---

---

<b>Introduction</b>	<b>3</b>
<hr/>	
<b>Getting Ready to Upgrade</b>	<b>3</b>
Upgrade Guidelines . . . . .	4
StorNext Licensing . . . . .	4
Upgrading the Client Software . . . . .	7
DDisk Support Discontinued . . . . .	7
StorNext Media on DVD . . . . .	7
<hr/>	
<b>Running the Pre-Installation Script</b>	<b>7</b>
Before You Begin . . . . .	8
Running snPreInstall . . . . .	9
Interpreting snPreInstall Output . . . . .	10
<hr/>	
<b>Upgrading Non-HA Systems to StorNext 4.0</b>	<b>10</b>
The License Check Option . . . . .	12
Clear Browser Cache . . . . .	13
SSL Certificate Exception . . . . .	13
<hr/>	
<b>Upgrading a Windows Non-HA Client or Server to StorNext 4.0</b>	<b>13</b>
<hr/>	
<b>Upgrading HA Systems To StorNext 4.0</b>	<b>14</b>
HA Upgrade Options . . . . .	14
Upgrading a 3.x HA System to 4.0 HA. . . . .	14
Upgrading a RHEL 4 3.x HA System to 4.0 HA. . . . .	17
Upgrading a Windows 3.x HA System to 4.0 HA. . . . .	20
Upgrading a 4.x HA System to 4.x HA. . . . .	21



---

<b>Upgrading HA Systems To StorNext 4.x (.RPM)</b>	<b>24</b>
Upgrading a 3.x .RPM HA System to 4.x .RPM HA . . . . .	24
Upgrading a 4.x .RPM HA System to 4.x .RPM HA . . . . .	25

---

<b>Upgrading From StorNext 3.x File System + GUI to StorNext 4.0 File System Only</b>	<b>27</b>
---	-----------

---

<b>Upgrading a Linux MDC From 3.x File System Only to StorNext 4.0 File System Only</b>	<b>28</b>
---	-----------

---

<b>Contacting Quantum</b>	<b>29</b>
---------------------------	-----------

---

---

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

**COPYRIGHT STATEMENT**

© 2010 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.

**TRADEMARK STATEMENT**

Quantum, the Quantum logo, DLT, DLTape, the DLTape logo, Scalar, and StorNext are registered trademarks of Quantum Corporation, registered in the U.S. and other countries. Backup. Recovery. Archive. It’s What We Do., the DLT logo, DLTSage, DXi, DXi-Series, Dynamic Powerdown, FastSense, FlexLink, GoVault, MediaShield, Optyon, Pocket-sized. Well-armored, SDLT, SiteCare, SmartVerify, StorageCare, Super DLTape, SuperLoader, and Vision are trademarks of Quantum. LTO and Ultrium are trademarks of HP, IBM, and Quantum in the U.S. and other countries. All other trademarks are the property of their respective companies. Specifications are subject to change without notice.

---

## Introduction

This document provides information about upgrading to StorNext 4.0 from a previous version. Instructions are provided for metadata controllers (MDCs) running StorNext File System (SNFS) and StorNext Storage Manager (SNSM).

---

**Note:** Quantum supports upgrades only from the previous two major releases of StorNext. For example, you can upgrade to StorNext 4.0 from version 3.5 or 3.1, but not from version 3.0 or earlier.

---

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 Systems to StorNext 4.0](#)
- [Upgrading a Windows Non-HA Client or Server to StorNext 4.0](#)
- [Upgrading HA Systems To StorNext 4.0](#)
- [Upgrading HA Systems To StorNext 4.x \(.RPM\)](#)
- [Upgrading From StorNext 3.x File System + GUI to StorNext 4.0 File System Only](#)
- [Upgrading a Linux MDC From 3.x File System Only to StorNext 4.0 File System Only](#)
- [Contacting Quantum](#)

---

## Getting Ready to Upgrade

This section includes important information about upgrading to StorNext 4.0. Review the following sections before upgrading:

- [Upgrade Guidelines](#)
- [StorNext Licensing](#)
- [Upgrading the Client Software](#)
- [DDisk Support Discontinued](#)
- [StorNext Media on DVD](#)

---

## Upgrade Guidelines

---

Before upgrading StorNext, review the following upgrade guidelines:

- Before upgrading, make sure there is no I/O activity. On the StorNext home page, the number of store candidates should be 0. 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.
- Because the metadata dump format changed starting in StorNext 3.0, 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.
- If you are upgrading from 3.x File System + GUI to StorNext 4.0 + GUI, before you begin the upgrade you must first manually remove affinities from the .cfg file. After you finish the upgrade you can use the StorNext GUI to add the affinities back.
- In the 4.0 release of StorNext, individual features are now being licensed separately. Because of this new behavior license checks of these features are made up front in the upgrade processing. If there are features that are found to be in use but not licensed, the upgrade will abort. Before attempting an upgrade, contact Quantum to get an updated license file for the features currently in use.
- Starting in StorNext 4.0, to be able to perform an upgrade, a site must have a maintenance license with an expiration date that is after the build date of the upgrade release. Contact Quantum for more information or to obtain this license.

---

## 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.)

Here is a list of all the StorNext licenses:

- **File System:** A File System license enables you to create and modify StorNext-supported file systems.

- **LAN Client:** You must have a Distributed LAN Client license for each LAN client you use with StorNext (in addition to any SAN clients).
- **Storage Manager:** A Storage Manager license provides full access to StorNext's Storage Manager features that are not licensed separately.
- **Replication:** A replication license is required if you want to use StorNext's Data Replication feature.
- **Deduplication:** A deduplication license is required if you want to use StorNext's Data Deduplication feature.
- **Vaulting:** A Vaulting license provides the ability to move seldom-used media to a manual archive vault, freeing room for media in the managed archives.
- **Storage Disk:** You must have a Storage Disk license to be able to configure and use StorNext storage disks.
- **Checksum:** A Checksum license enables you to verify data integrity by ensuring that the checksum created when data was stored matches the checksum upon data retrieval.
- **Distributed Data Mover (DDM):** The number displayed is the maximum number of clients that can be used to run mover processes.
- **Failover (HA):** A Failover (High Availability) license is required if you plan to use StorNext's HA failover features.
- **Maintenance:** A Maintenance license verifies that your site has purchased StorNext upgrade licenses, and is required for StorNext upgrades.

## Upgrade Implications

The new licensing implementation affects StorNext upgrades, both for release 4.0 and future releases. Be aware of the following upgrade-related implications and plan accordingly:

- A 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 4.0.**
- 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. This step is done by manually editing the `license.dat` file because the StorNext GUI has not been installed and therefore you cannot enter licenses through the GUI. This applies to upgrades to StorNext 4.0 and also to future upgrades.

**You will not be allowed to proceed with the upgrade if a valid Maintenance license is not in the `license.dat` file.**

---

**Note:** In addition to the Maintenance license required to upgrade, you must enter into the `license.dat` file license strings for ALL the StorNext component and features you have purchased.

---

- For 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.

- The Maintenance license must remain in place even after expiration to allow the StorNext software to run.

## 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  
000099             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, you must contact the Quantum Technical Assistance Center at [licenses@Quantum.com](mailto:licenses@Quantum.com) and give them the following information:

- The serial number from your product DVD or box.
- A list of the components/features you wish to license.
- The number of StorNext SAN clients and distributed LAN clients you want to support, as well as capacity.
- The StorNext server identification number.

After the Quantum Technical Assistance Center receives the above information, a representative will send you license strings for the components/features you specified.

### 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 4.0:

- 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 (`install.stornext -upgrade`) as explained in [Upgrading Non-HA Systems to StorNext 4.0](#) on page 10.
- 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 4.0 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. For more information about installing the client software on Linux and Unix, see the *StorNext 4.0 Installation Guide*.

---

#### 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 4.0 when you upgrade.* For assistance removing DDisks, contact Quantum Technical Support.

---

#### StorNext Media on DVD

---

Beginning with StorNext 4.0, 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.

---

## 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 1](#).

Table 1 StorNext Support Directories

Support Directory*	Description
<b>Database</b> /usr/adic/database	Records information about where and how data files are stored.
<b>Journal</b> /usr/adic/database/ journal	Records changes made to the database.
<b>Mapping</b> /usr/adic/TSM/ internal/mapping_dir	Contains index information that enables quick searches on the file system.
<b>Metadata</b> /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 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.

- 5 Run the script. At the command prompt, type:

```
./snPreInstall
```

The pre-installation script runs ([Figure 1](#)).

Figure 1 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 occurred, 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>]:
```

- 6 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 2](#).

---

Table 2 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, the **Metadata** directory should not be located on the same file system as (in order of priority) the **Journal**, **Database**, or **Mapping** directory.

---

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

---

---

## Upgrading Non-HA Systems to StorNext 4.0

Use the StorNext installation script to upgrade to StorNext 4.0 on a metadata controller running StorNext Storage Manager and StorNext File System. The installation script will upgrade both 3.X (3.1.2 and later) full StorNext and 3.X file system plus GUI installations. To launch the script, use the StorNext installation DVD.

The following instructions apply to non-HA servers. For instruction on upgrading HA servers, see [Upgrading HA Systems To StorNext 4.0](#) on page 14.

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

- 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 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.

- 5 As described in [Upgrade Guidelines](#) on page 4, 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**

- 6 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 12.

- 7 Run the script with the **-upgrade** option. At the command prompt, type:

**./install.stornext -upgrade**

The installation script runs ([Figure 2](#)).

Figure 2 StorNext Installation Script

```
StorNext Upgrade 4.0 (6013)  MAIN MENU
1) Installation Configuration
2) Upgrade stornext
3) Show Status
4) Quit
Enter your choice <2>:
```

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

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

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

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

---

## 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.
```

+++++

---

## Clear Browser Cache

---

If you are upgrading to StorNext 4.0 from a previous 3.x release, before logging in to the StorNext GUI for the first time you must first clear your Internet browser's cache. Doing this ensures that any settings from an earlier StorNext release are cleared and will not affect StorNext 4.0 operation.

You don't need to clear the cache if you are logging in using `https:<StorNext MDC IP address>`, but you must clear the cache if the login address is the same as in your 3.x version.

---

## 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 4.0

This section describes how to install StorNext on Windows non-HA clients or servers.

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

To install, run `SnfsSetup.exe` from the StorNext installation disk, and then follow the on-screen instructions.

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, etc 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.

---

# Upgrading HA Systems To StorNext 4.0

This section describes how to upgrade an HA system from StorNext releases 3.x through 4.0.

Refer to the section that applies to your upgrade:

- [Upgrading a 3.x HA System to 4.0 HA](#)
- [Upgrading a RHEL 4 3.x HA System to 4.0 HA](#)
- [Upgrading a Windows 3.x HA System to 4.0 HA](#)
- [Upgrading a 4.x HA System to 4.x HA](#)

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.

Remember, you will need new licenses when upgrading from 3.x to 4.0. If you're unsure whether you have the appropriate licenses, contact the Quantum Technical Assistance center.

---

## HA Upgrade Options

Prior to StorNext 4.0, HA upgrades were performed exclusively by Quantum Professional Services personnel. Beginning with StorNext 4.0, customers now have the option of performing the upgrade themselves, provided they possess the technical expertise required to successfully complete the upgrade steps.

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 4.0 and a general familiarity with the StorNext GUI.

*You should not attempt performing 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 Professional 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 3.x HA System to 4.0 HA

Follow these steps to upgrade a 3.1.X or 3.5.X HA system to a 4.0 HA system. You must be at release 3.1.2 or later or 3.5.0 or later to perform this procedure. If you are using a StorNext version earlier than 3.1.2 or 3.5.0, upgrade to version 3.1.2 or 3.5.0 or later before proceeding.

The steps in this section refer to installing from the StorNext 4.0 DVD, but the same steps apply if you are using StorNext 4.0 installation objects downloaded from the StorNext Website.

---

**Note:** The DDisk feature is no longer supported in StorNext 4.0. Before upgrading to StorNext 4.0, first remove any DDisks you have configured. For assistance removing DDisks, contact Quantum Technical Support and reference internal document number 0-12770-01.

---

## Pre-Upgrade Steps

- 1 Run a manual backup to tape from the StorNext GUI.
- 2 Make sure all store/retrieve requests have finished.
- 3 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.

- 4 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.)
- 5 Uninstall the standby server (also called the secondary server) by mounting your StorNext 4.0 installation DVD and running this command:

```
install.stornext -remove
```

---

**Note:** You MUST use the StorNext 4.0 version of `install.stornext -remove` for this step.

---

## Upgrade Installation Steps

- 1 Install StorNext 4.0 on the secondary server by running this command:

```
install.stornext
```

This task requires that you mount your StorNext 4.x DVD.

- 2 Update the license file (`/usr/cvfs/config/license.dat`) on the secondary server as needed.
- 3 Upgrade the primary server by performing these steps:
  - a Make sure the current server is the primary.
  - b Make sure the dedicated shared file system is mounted on `/usr/adic/HAM/shared`.
  - c Replace the old license file with the new one in the `/usr/cvfs/config` directory.

- d Push the license changes to the synchronization directory:

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

- e Run the following command to upgrade to StorNext 4.0:

```
install.stornext -upgrade
```

## Pre-Conversion Steps

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

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

(You should do this every time an MDC restarts.)

- 1 Create the file `/usr/cvfs/config/ha_peer` on the primary node. Put the secondary MDC's IP address in the file.

- 2 Start the cluster HA manager service by running this command:

```
service snhamgr start
```

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

```
snhamgr peerdown
```

```
snhamgr mode=config
```

- 4 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 – `snbackup`. Depending on the number of managed files in the system, this backup may take substantial time to complete.

---

- 5 Run this command to verify that `post_upgrade` has completed on 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.

---

- 6 Verify that all processes are running correctly on the primary server by testing access to all file systems and moving files to/from tapes.

- 7 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 in a Web browser: `https://<primary_server>`
- 2 Choose **Convert** from the **Tools > HA** menu.



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

4 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-Upgrade Steps

- 1 If you are using the DDM feature, make sure that all of the mover clients have been upgrade and then edit `fs_sysparm` or `fs_sysparm_override` to set your preferred DDM mode: All or Threshold. Use the command `adic_control restart TSM` to put this change into effect.
- 2 If you are planning to use Deduplication or Replication, you must create at least one virtual IP address (VIP). The VIPs must be created on both the Primary and Secondary using the command `/usr/cvfs/bin/vip_control -u <vip_str>`

Refer to the `vip_control` man page and the "Converting to HA" chapter in the StorNext User Guide for more information on VIPs.

If you plan to use Deduplication and/or Replication, you should configure these features as described in the "Replication and Deduplication" chapter in the StorNext User's Guide.

---

### Upgrading a RHEL 4 3.x HA System to 4.0 HA

---

HA on RHEL 4 is not supported in StorNext 4.0. In order to upgrade a StorNext 3.x (3.1.2 and later) HA system running on RHEL 4 to StorNext 4.0, both HA MDC servers must first be upgraded to RHEL 5, and then upgraded to StorNext 4.0.

To upgrade from RHEL 4 to RHEL 5, you must have the RHEL 5 StorNext 3.x CD. To upgrade to StorNext 4.0, you must have the StorNext 4.0 DVD *and* new licenses. Refer to [The License Check Option](#) on page 12 if you are not sure you have the necessary licenses for StorNext 4.0.

## Upgrade the Operating System on the First of Two Servers

The two machines in the HA Cluster will have their OS versions upgraded one at a time while the StorNext HA configuration is preserved on the other machine. During the full set of steps to accomplish this, the roles of Primary and Secondary server move back and forth between the machines.

The following instructions refer to the machines as **First** and **Second** to avoid confusion about the machine to which the steps apply.

The First machine to be upgraded is the Standby at the start of the instructions, and likewise, the Second machine is the Primary at the start of the instructions.

## Pre-Upgrade Steps

- 1 Run a manual backup to tape from the StorNext GUI.
- 2 Make sure all store/retrieve requests have finished.
- 3 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.)
- 4 Unmount all file systems from all client machines and stop the SNFS processes on each client machine.
- 5 Copy the `/usr/cvfs/config/license.dat`, `fsnameservers`, and `cvpower` files from the standby server to a safe location such as a USB flash drive.

## Upgrade The First Server

- 1 Perform the following steps to upgrade the First server to RHEL 5"
  - a Make sure the First server is in "standby" mode. That is, FSMs are not activated on this server.
  - b Shut down the First server.
  - c Perform a fresh RHEL 5 installation on this server, making sure the host name and IP addresses are not changed.
- 2 Run this command to install StorNext 3.x for RHEL 5 on the First server:  

```
install.stornext
```
- 3 Perform these steps to convert the First server to HA:
  - a Copy the `license.dat`, `fsnameservers` and `cvpower` files from the safe location (e.g., USB flash drive) back into the `/usr/cvfs/config` directory.

- b** Run the following command:

```
/usr/adic/util/cnvt2ha.sh standby <second server>  
<shared_fs>
```

The `<second server>` parameter should be the same as the host name shown in the Second server's license file. The `shared_fs` parameter is the file system name used as the dedicated shared file system.

- c** Verify that the mirrored configuration files are properly synchronized.

---

**Note:** After you complete this step you'll have the Second server running RHEL 4 StorNext 3.x and the First server running RHEL 5 StorNext 3.x. Technically, this is not a supported configuration since the primary and the standby servers are running different operating systems. However, this condition is temporary during the upgrade process. When the whole upgrade is complete, both servers will conform to the supported configuration.

---

- 4** Fail over to the First server so that it is Primary, and then review the configuration to make sure the HA conversion was done properly. Failover is induced by running the following command on the Second server:

```
cvadmin -e 'fail fsname'
```

where `fsname` is one of your file systems. You'll notice a failover occurs, as the Second server goes offline and the First server activates its FSMs.

- 5** When the Second server has finished rebooting, shut down primary cvfs by running `'service cvfs stop'`
- 6** Copy the `/usr/cvfs/config/license.dat`, `fsnameservers`, and `cvpower` files from the standby server to a safe location such as a USB flash drive.

### Upgrade The Operating System on the Second of Two Servers

- 1** Perform a fresh RHEL 5 installation on the Second server, making sure the host name and IP addresses are not changed.
- 2** Install StorNext 3.x for RHEL 5 on the Second server.
- 3** Perform these steps to convert the Second server to HA:
  - a** Copy the `license.dat`, `fsnameservers` and `cvpower` files from the safe location (e.g., USB flash drive) back into the `/usr/cvfs/config` directory.
  - b** Copy the `cvpower` script from `/usr/adic/util` to `/usr/cvfs/config`, and then edit the script so that it has the proper plug name, password, etc.
  - c** Run the following command:

```
/usr/adic/cnvt2ha.sh standby <first server> shared_fs
```

---

**Note:** The "first server" is the name for the original standby server. It should be the same as the host name shown in the First server's license file.

---

- d Verify that the mirrored configuration files are properly synchronized.
- 4 Fail over to the Second server by running the following command on the First server, and then verify that its HA conversion was done properly:  

```
/usr/cvfs/bin/cvadmin -e "fail fsname"
```

Where `fsname` is one of your file systems. You'll notice a failover occurs, as the First server goes offline and the second server activates its FSMs.
- 5 Log in to StorNext on the Second server, and then review the GUI configuration information to make sure the conversion was done properly.

### Post-Upgrade Steps

At this stage, both servers have been upgraded to RHEL 5 StorNext 3.x HA. You are now ready to follow the instructions described in [Upgrading a 3.x HA System to 4.0 HA](#) on page 14

---

### Upgrading a Windows 3.x HA System to 4.0 HA

---

Follow these steps to upgrade a Windows 3.1.X or 3.5.X HA system to a 4.0 HA system:

- 1 Obtain StorNext licenses, which must include at least Maintenance, Server and Failover licenses for both MDCs.
- 2 Stop StorNext services on all StorNext clients.
- 3 Rename or remove `cvfail.bat` from the StorNext configuration directory on both MDCs.
- 4 Verify that configuration files are correct and identical on both MDCs.
  - a The `*.cfg` files should be identical on both computers.
  - b The `fsnameservers` files should be identical on both computers.
  - c The `fsmlist` files should be identical on both computers.
- 5 Run `SnfsSetup.exe` from the StorNext installation disk on both MDCs. Select "Upgrade/Reinstall". Follow the on-screen instructions to import your license file.
- 6 Enable High Availability (HA) on both MDCs by doing the following:
  - a After installing StorNext, on each MDC run: **Start->All Programs->StorNext->File System Cfg (Advanced)**.
  - b Select **HA Configuration** from the **Tools** menu.
  - c Click "**Enable High Availability (HA)**". Note that you cannot enable HA if there are no file systems configured.
- 7 MDC upgrades are now complete. You may now test failover by running `cvadmin -e "hamon smith"` on the primary MDC. (The Primary MDC is indicated by an asterisk in `cvadmin`.)

---

## Upgrading a 4.x HA System to 4.x HA

---

This procedure describes how to upgrade from a StorNext 4.x HA system to a 4.x HA system. Although this procedure is not required for the initial 4.0 release, the procedure must be followed in the near future for maintenance releases and any HA-related LCRs.

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 12 if you are not sure you have the necessary licenses for StorNext 4.0.

### 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 Uninstall the standby server (also called the secondary server) by running this command:

```
install.stornext -remove
```

- 7 Install the StorNext 4.x on the secondary server by running this command:

```
install.stornext
```

- 8 Add the `license.dat` file into the `/usr/cvfs/config` directory

### 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 Put the new licenses for both servers into `/usr/cvfs/config/license.dat`.
- 7 Run the following command to upgrade to StorNext 4.x:  
`install.stornext -upgrade`

### Pre-Conversion Steps

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

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

- 2 Change the mode of the primary and secondary servers to "default" in preparation for adding the secondary server back into the HA cluster. Run the following commands on the primary server:

```
service cvfs start      (#pushes configuration changes to the mirror)
snhamgr mode=default   (#stops cvfs)
snhamgr peerup
service cvfs start      (#restarts cvfs)
```

---

**Note:** A `post_upgrade` script runs which performs a system backup – `snbackup`. 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:

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

- 4 Verify that all processes are running correctly by testing access to all file systems and moving files to/from tapes.

---

**Note:** Make sure the `snbackup` process has completed prior to continuing.

---

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

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

- 6 Run this command to start StorNext 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.
- 2 Choose **Convert** from the **Tools >HA** 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 HA Systems To StorNext 4.x (.RPM)

This section describes how to upgrade a 3.x HA system to a StorNext 4.x HA system using an RPM-only installation. Instructions are also provided for 4.x to 4.x RPM systems.

Steps are run on both the primary and secondary nodes, and are indicated accordingly.

---

### Upgrading a 3.x .RPM HA System to 4.x .RPM HA

---

Acquire the StorNext client, server and core 4.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_261621-4.0.0.12838.x86_64.rpm  
snfs-server-SuSE100ES-4.0.0.12838.x86_64.rpm  
snfs-SuSE100ES-4.0.0.12838.x86_64.rpm
```

Make sure you have obtained and installed the proper licenses for both nodes for the targeted build (such as 4.x). Refer to [StorNext Licensing](#) on page 4 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 "service cvfs stop" or /etc/init.d/cvfs fullstop.

#### On The Secondary (Standby) MDC

- 1 Log in to the secondary server (none of its FSMs are activated), and then stop and remove StorNext by running these commands:

```
service cvfs fullstop  
rpm -qa | grep snfs | xargs rpm -ev
```

- 2 Verify that the /usr/cvfs/config/fsnameservers file has the same information as the primary server.
- 3 Remove the /usr/cvfs/config/cvfail and /usr/cvfs/config/cvpower scripts.
- 4 Remove all the \*.cfg files in /usr/cvfs/config

#### On the Primary

- 1 Stop StorNext services by running this command:

```
service cvfs fullstop
```

- 2 Remove the /usr/cvfs/config/cvfail and /usr/cvfs/config/cvpower scripts.



- 3 Upgrade to the StorNext 4.x RPMs you acquired earlier. For example:

```
rpm -Uvh snfs-client-SuSE100ES_261621-  
4.0.0.12838.x86_64.rpm snfs-server-SuSE100ES-  
4.0.0.12838.x86_64.rpm snfs-SuSE100ES-  
4.0.0.12838.x86_64.rpm
```

- 4 Create the `ha_peer` file with the IP address of the **Secondary** in the HA pair.

```
vi /usr/cvfs/config/ha_peer
```

- 5 Enable HA file system monitoring as follows for each file system:

```
/usr/cvfs/bin/sncfgedit <fsname>
```

```
Change <haFsType>HaUnmonitored</haFsType>
```

```
to this: <haFsType>HaUnmanaged</haFsType>
```

- 6 Start StorNext by running this command:

```
service cvfs start
```

## Back On The Secondary MDC

- 1 Create the `ha_peer` file with the IP address of the **Primary** in the HA pair.

```
/usr/cvfs/config/ha_peer
```

- 2 Copy over all `/usr/cvfs/config/*.cfgx` files from the Primary to Secondary server.

- 3 Start up `cvfs` by running this command:

```
service cvfs start
```

- 4 Verify that all file systems are mounted and active.

- 5 Test failover by running the following command on the Primary server:

```
snhamgr force smith
```

## Notes

To turn off HA functionality:

- 1 Stop `cvfs` on the Secondary, and prevent it from starting by removing its configuration files or by uninstalling StorNext.
- 2 On the Primary, change the `haFsType` configuration item to `HaUnmonitored` in every file system's `*.cfgx` file.
- 3 Restart `cvfs` on the Primary.

---

## Upgrading a 4.x .RPM HA System to 4.x .RPM HA

---

Acquire the StorNext client, server and core 4.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_261621-4.0.0.12838.x86_64.rpm
```

```
snfs-server-SuSE100ES-4.0.0.12838.x86_64.rpm
```

```
snfs-SuSE100ES-4.0.0.12838.x86_64.rpm
```

Make sure you have obtained and installed the proper licenses for both nodes for the targeted build (such as 4.x). Refer to [StorNext Licensing](#) on page 4 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.

Log in to the secondary server (none of its FSMs are activated), and then stop and upgrade StorNext by running these (example) commands:

```
service cvfs fullstop  
  
rpm -ev snfs-client-SuSE100ES_261621-  
3.1.x.12838.x86_64.rpm snfs-server-SuSE100ES-  
3.1.x.12838.x86_64.rpm snfs-SuSE100ES-  
4.0.0.12838.x86_64.rpm
```

### On the Primary

- 1 Stop StorNext services by running this command:

```
service cvfs fullstop
```

- 2 Upgrade to the StorNext RPMs acquired earlier by running this (example) command:

```
rpm -Uvh snfs-client-SuSE100ES_261621-  
4.0.0.12838.x86_64.rpm snfs-server-SuSE100ES-  
4.0.0.12838.x86_64.rpm snfs-SuSE100ES-  
4.0.0.12838.x86_64.rpm
```

- 3 Start cvfs by running this command:

```
service cvfs start
```

### Back On The Secondary MDC

Start cvfs by running this command: `service cvfs start`

---

## Upgrading From StorNext 3.x File System + GUI to StorNext 4.0 File System Only

The following procedure is required in order to upgrade from a StorNext 3.x File System plus GUI (i.e., `install.snfs`) configuration to a StorNext 4.0 File System only RPM installation.

- 1 Create a backup of the following files and directories:

```
/etc/fstab  
/usr/cvfs/config  
/usr/cvfs/data  
/usr/cvfs/debug
```

- 2 Change to the directory that contains the 3.X install media. (If you are installing from disc, you must first mount the DVD.)

```
cd /mnt/cdrom/phdist/RedHat50AS_26x86_64
```

- 3 Remove the software:

```
./install.snfs -remove
```

- 4 Install the 3.X snfs rpms from the installation media you mounted in step 2:

```
cd /mnt/cdrom/phdist/RedHat50AS_26x86_64/DSM  
rpm -hiv snfs-client*.rpm snfs-server-*.rpm
```

- 5 Restore the configuration files you backed up as part of step 1.

- 6 Verify that all the file systems mount and are accessible:

```
service cvfs start
```

- 7 Stop the file systems:

```
service cvfs fullstop
```

- 8 Update the file `/usr/cvfs/config/license.dat` to include the licensed StorNext 4.0 features. At a minimum, a Maintenance and Server license are required for StorNext 4.0. A Proxy license may also be necessary if Distributed LAN clients were configured in StorNext 3.X

- 9 Change to the directory containing the StorNext 4.0 RPMs.

- 10 Upgrade the RPMs to StorNext 4.0. You will need to upgrade three RPMs at the same time:

```
rpm -Uvh snfs-RedHat50AS*.rpm snfs-client-*.rpm snfs-server-*.rpm
```

- 11 Verify that all the file systems mount and are accessible:

```
service cvfs start
```

---

## Upgrading a Linux MDC From 3.x File System Only to StorNext 4.0 File System Only

This section describes how to upgrade from a Linux StorNext 3.x metadata controller with File System only to StorNext 4.0 with File System only.

- 1 Acquire the StorNext client, server and core 4.X RPMs for your server OS architecture from the StorNext installation DVD, or by downloading the tarball from the Quantum web site. For example:

```
snfs-client-SuSE100ES_261621-4.0.0.12838.x86_64.rpm
```

```
snfs-server-SuSE100ES-4.0.0.12838.x86_64.rpm
```

```
snfs-SuSE100ES-4.0.0.12838.x86_64.rpm
```

- 2 Make sure you have obtained and installed the proper licenses for the targeted build (4.x). Refer to [StorNext Licensing](#) on page 4 for more information about license requirements.
- 3 Unmount all file systems from all client machines, and then stop the StorNext File System processes on each client machine by running "service cvfs stop" or /etc/init.d/cvfs fullstop.
- 4 Stop StorNext services on the MDC by running this command:

```
service cvfs fullstop
```

- 5 Run this (example) command to upgrade to the StorNext RPMs you acquired earlier:

```
rpm -Uvh snfs-client-SuSE100ES_261621-  
4.0.0.12838.x86_64.rpm snfs-server-SuSE100ES-  
4.0.0.12838.x86_64.rpm snfs-SuSE100ES-  
4.0.0.12838.x86_64.rpm
```

- 6 Start cvfs by running this command:

```
service cvfs start
```

---

## 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 the Quantum Technical Assistance Center:

<b>United States</b>	800-284-5101 (toll free) 949-725-2100
<b>EMEA</b>	00800-4-782-6886 (toll free) +49 6131 3241 1164
<b>APAC</b>	+800 7826 8887 (toll free) +603 7953 3010

For worldwide support:

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

---

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

