

Quantum®

StorNext®



Upgrade Guide

6-68044-02, Rev. F

StorNext 6 Upgrade Guide, 6-68044-02, August 2019, Product of 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

© 2019 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

Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, FlexTier, GoVault, Lattus, NDX, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners.

Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners.

Quantum specifications are subject to change.



Contents

Preface	v
Chapter 1: Introduction	1
Introduction to StorNext Upgrades	1
Chapter 2: Getting Ready to Upgrade	3
Get Ready To Upgrade	3
Chapter 3: Run the Pre-Installation Script	12
Run the Pre-Installation Script	12
Switch to a Named Cluster Using the Command Line Interface	16
Chapter 4: Upgrade non-HA Linux MDCs	18
Upgrade non-HA Linux MDCs	18
Chapter 5: Upgrade a Windows Client or Non-HA Server	23
Upgrade a Windows Client or Non-HA Server	23

Chapter 6: Upgrade an HA System	26
Upgrade an HA System	26
Chapter 7: Upgrade an HA System Using RPM Files	33
Upgrade an HA System Using RPM Files	33
Chapter 8: Upgrade Clients	37
Upgrade a Client	37
Chapter 9: Upgrade a StorNext Storage Manager Object Storage Configuration	50
Upgrade a StorNext Storage Manager Object Storage Configuration	50
Chapter 10: Upgrade Firmware	54
Upgrade Firmware	54
Appliance Release Notes	54
Chapter 11: Compression and Encryption Usage Report	56
Compression and Encryption Usage Report	56



Preface

This manual contains the following chapters:

- [Chapter 1: Introduction](#)
- [Chapter 2: Getting Ready to Upgrade](#)
- [Chapter 3: Run the Pre-Installation Script](#)
- [Chapter 4: Upgrade non-HA Linux MDCs](#)
- [Chapter 5: Upgrade a Windows Client or Non-HA Server](#)
- [Chapter 6: Upgrade an HA System](#)
- [Chapter 7: Upgrade an HA System Using RPM Files](#)
- [Chapter 8: Upgrade Clients](#)
- [Chapter 9: Upgrade a StorNext Storage Manager Object Storage Configuration](#)
- [Chapter 10: Upgrade Firmware](#)
- [Chapter 11: Compression and Encryption Usage Report](#)

Audience

This manual is written for StorNext 6 operators, system administrators, and field service engineers.

Notational Conventions


This manual uses the following conventions:

Convention	Example
User input is shown in bold monospace font.	./DARTinstall
Computer output and command line examples are shown in monospace font.	./DARTinstall
User input variables are enclosed in angle brackets.	http://<ip_address>/cgi-bin/stats
For UNIX and Linux commands, the command prompt is implied.	./DARTinstall is the same as # ./DARTinstall
File and directory names, menu commands, button names, and window names are shown in bold font.	/data/upload
Menu names separated by arrows indicate a sequence of menus to be navigated.	Utilities > Firmware

The following formats indicate important information:

 **Note:** Note emphasizes important information related to the main topic.

 **Caution:** Caution indicates potential hazards to equipment or data.

 **WARNING:** Warning indicates potential hazards to personal safety.

- Right side of the system - Refers to the right side as you face the component being described.
- Left side of the system - Refers to the left side as you face the component being described.
- Data sizes are reported in base 10 (decimal) rather than base 2¹⁰ (binary). For example:
10,995, 116,277,769 Bytes are reported as 11.0 TB (decimal/1000). In binary, this value is 10 TiB (binary/1024).

Product Safety Statements

Quantum will not be held liable for damage arising from unauthorized use of the product. The user assumes all risk in this aspect.

This unit is engineered and manufactured to meet all safety and regulatory requirements. Be aware that improper use may result in bodily injury, damage to the equipment, or interference with other equipment.

-
- WARNING:** Before operating this product, read all instructions and warnings in this document and in the *Quantum Products System, Safety, and Regulatory Information Guide*.
-
- ADVARSEL:** Læs alle instruktioner og advarsler i dette dokument og i *Informationsvejledning vedrørende system-, sikkerheds- og lovbestemmelser for Quantum produkter, før produktet betjenes*.
-
- AVERTISSEMENT :** Avant d'utiliser ce produit, lisez toutes les instructions et les avertissements de ce document et du *Guide d'informations sur le système, la sécurité et la réglementation de Quantum*.
-
- WARNUNG:** Lesen Sie vor der Inbetriebnahme dieses Produkts alle Anleitungen und Warnungen in diesem Dokument und im *System-, Sicherheits- und Betriebsbestimmungen-Handbuch für Quantum-Produkte*.
-
- ADVERTENCIA:** Antes de hacer funcionar este producto, lea todas las instrucciones y advertencias de este documento y de la *Guía de información normativa, del sistema y de seguridad de los productos de Quantum*.
-
- VARNING:** Läs igenom alla instruktioner och varningar i detta dokument och i *Quantums produktsystem, säkerhet och reglerande informationsguide* innan denna produkt används.
-
- ВНИМАНИЕ!** Перед началом эксплуатации данного изделия прочтите все инструкции и предупреждения, приведенные в настоящем документе и в *Руководстве по системе, технике безопасности и действующим нормативам компании Quantum*.
-
- 警告：**本製品を使用される前に、本書と『*Quantum製品システム、安全、規制情報ガイド*』に記載されているすべての説明と警告をお読みください。
-
- 경고:** 본 제품을 작동하기 전에 본 문서와 *Quantum 제품 시스템, 안전 및 규제 정보 설명서*에 있는 모든 지침과 경고를 참조합니다.
-
- 警告：**在操作本产品之前，请阅读本文档和 *Quantum 产品系统、安全和法规信息指南*中的所有说明和警告。
-
- 警告：**操作此產品前，請閱讀本檔案及 *Quantum 產品系統、安全與法規資訊指南*中的指示與警告說明。

אזהרה: לפני ההפעלה של מוצר זה, קרא את כל ההוראות והאזהרות הכלולות במסמך זה וכן במדריך המידע בנושא מערכת, בטיחות ותקינה עבור מוצרי Quantum.

For the most up to date information on StorNext 6, see:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<http://www.quantum.com/aboutus/contactus/index.aspx>

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com

Getting More Information or Help

StorageCare™, Quantum's comprehensive service approach, leverages advanced data access and diagnostics technologies with cross-environment, multi-vendor expertise to resolve backup issues faster and at lower cost.

Accelerate service issue resolution with these exclusive Quantum StorageCare services:

- **Service and Support Website** - Register products, license software, browse Quantum Learning courses, check backup software and operating system support, and locate manuals, FAQs, firmware downloads, product updates and more in one convenient location. Get started at:
<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>
- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can also be used to access Quantum's Knowledge Base, a comprehensive repository of product support information. Get started at:

<http://www.quantum.com/customercenter/>

For further assistance, or for training opportunities, contact the Quantum Customer Support Center:

Region	Support Contact
North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49 6131 324 185
Asia Pacific	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

Worldwide End-User Product Warranty

For more information on the Quantum Worldwide End-User Standard Limited Product Warranty:

<http://www.quantum.com/serviceandsupport/warrantyinformation/index.aspx>



Chapter 1: Introduction

This chapter contains the following topics:

Introduction to StorNext Upgrades	1
---	---

Introduction to StorNext Upgrades

This section provides instructions for upgrading metadata controllers (MDCs) running StorNext File System (SNFS) and StorNext Storage Manager (SNSM) to StorNext from prior StorNext releases.

System Requirements

The operating systems, kernel versions, hardware platforms, hardware, and network requirements, for both StorNext MDCs and StorNext clients are listed in the [StorNext 6 Compatibility Guide](#).

Make sure the system meets all the requirements outlined in the guide, and, if necessary, update your system before upgrading StorNext.

DVD Directories

Files on the StorNext installation DVD are located in the following locations, depending on whether a media kit DVD or a user-created DVD from an .iso image for a web-release download is used:

Directory Location	Directory	Contents
Media Kit DVD	/fs_only	This contains the installation binaries for supported platforms for StorNext file system-only and file system-only clients and clients plus Distributed Data Mover (DDM), which are designated with a _plus_mover file name. This also contains the StorNext client installation binaries for supported platforms.
	/stornext_full	This contains the installation binaries for supported platforms for full installation of StorNext, including both the File System and Storage Manager.
	/stornext_full/clients	This contains the StorNext client installation binaries for supported platforms.
	/offlinenotification	This contains the installation binaries for supported platforms for StorNext offline notification installations.
Web Release DVD	/fs_only	This contains the installation binaries for supported platforms for StorNext Filesystem-only and Filesystem-only clients and clients plus Distributed Data Mover (DDM), which are designated with a _plus_mover file name. This also contains the StorNext client installation binaries for supported platforms.
	/stornext_full	This contains the installation binaries for supported platforms for full installation of StorNext, including both the File System and Storage Manager.
	/stornext_full/clients	This contains the StorNext client installation binaries for supported platforms.
	/offlinenotification	This contains the installation binaries for supported platforms for StorNext offline notification installations.

Upgrade Your Appliance Controller Console or StorNext NAS

If you need to upgrade to a newer version of the Appliance Controller Console or StorNext NAS, see the [Upgrades](#) section of the *Appliance Controller Documentation Center*.



Chapter 2: Getting Ready to Upgrade

This chapter contains the following topics:

Get Ready To Upgrade 3

Get Ready To Upgrade

Review the sections below, which include important information about upgrading.

Upgrade Considerations

Before you begin the upgrade you should note the following considerations so you can plan accordingly. Consider the following prior to upgrading:

- Not all StorNext releases may be upgraded to a given StorNext release. As a result, an upgrade to the current version of StorNext may require multiple, incremental upgrades, depending on the version of StorNext currently installed.

For information about supported upgrade paths, consult the [StorNext 6 Compatibility Guide](#). If your system is running a StorNext release prior to the supported upgrade releases for a given StorNext 6 (and later) release, consult an earlier version of the [StorNext 6 Compatibility Guide](#) that applies to your specific upgrade, and the dependencies for StorNext Clients in the environment.

- Firmware upgrade installation files must first be acquired from Quantum.
- If you attempt to perform a full backup operation (for example, by using the command **snbackup**)

immediately after upgrading, the backup operation will fail.

-
- i Note:** The backup operation fails due to the metadata archive build process, and will not succeed until the metadata archive build is complete. The amount of time required for the metadata archive build to complete depends on the size of the file system (for example, the number of files and directories); however, other file system actions are allowed (for example, adding new files, modifying files, and so on).

Run the following command to check the status of the metadata archive build process:

```
/usr/adic/DSM/bin/cvadmin -F <file_system_name> -e "mdarchive status"
```

When the status returns complete, you can perform a full backup operation.

Upgrade Guidelines

Before upgrading StorNext review the following upgrade guidelines:

-
- i Note:** This upgrade guide documents supported upgrades to StorNext 6 (and later) only. For customers with older versions of StorNext, consult the *StorNext Upgrade Guide* that corresponds to the version of StorNext you are upgrading to.
-
- i Note:** StorNext 6 requires a minimum journal size of 4MB and a recommended size of 64MB. If you have file systems with journal sizes less than the 4MB minimum, you must resize their journals before upgrading to StorNext 6. Use the **cvupdatefs** utility or the GUI to resize the journal. When the journal is resized, the new size must be 16MB or greater. File systems with journals between 4MB and 16MB will run with StorNext 6, but it is recommended that these file systems have their journal resized to the recommended 64MB.
-
- i Note:** For non-Windows upgrades, the **snPreinstall** script must be run prior to **install.stornext**.
- Beginning with StorNext 4.3.0, StorNext uses a MySQL database for information needed internally by Storage Manager.
 - The MySQL database requires a minimum of 8 GB of RAM.
 - Before attempting to upgrade from a previous StorNext release, make sure you have free space on the file system. If the file system is nearly full when you begin the upgrade, serious errors may occur or the upgrade could fail. Best practice is to maintain an area on the file system which is not used for data or system files, but is reserved as an empty buffer to ensure that upgrades and other operations complete successfully.
 - After a successful upgrade from StorNext 5, you should move the files under **/usr/adic/database/metadumps** off of the file system or delete the files, as they are no longer used by StorNext 6 and might consume a significant amount of disk space.
-
- i Note:** StorNext 6 uses the directory **/usr/adic/database/mdarchives** to hold external copies of file system metadata.

- Prior to StorNext 6, keeping an external copy of metadata for unmanaged file systems was rarely done because the use cases were very limited. However, beginning with StorNext 6, there are several features that require or benefit when you enable **Metadata Archive**. These include, but are not limited to the following features:
 - [FlexSync](#)
 - [Stripe group management \(sgmanage\)](#)
 - [Auditing](#)
 - [snhistory](#)
 - [snrecover](#)

i Note: If you enable **Metadata Archive** for non-managed file systems, additional space is consumed on the HA shared file system. The amount is approximately one gigabyte per million files.

- StorNext 6 (and later) and later has the following additional disk space requirements:

Component	Directory Space Required
Database	231 MB
Journal	11 MB
Mapping	157 MB
Metadata	333 MB
MySQL_tmp	120 MB
Total	852 MB

* Space requirements based on the following:

- 250, 000 files
- 10, 000 directories
- 1 copy per file
- 1 version per file

Additional space will be required as Storage Manager manages more files, directories, copies and versions.

- 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.

Important

There is a potential problem with the system clock on Nehalem Intel processors that can cause the clock to jump and temporarily freeze the system. The problem 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. The fix involves changing the kernel parameters in the `grub.conf` file.

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`:

```
grub.conf:idle=poll intel_idle.max_cstate=0 processor.max_cstate=1
```

In some cases, performance can also be improved by adjusting the idle kernel parameter. However, care should be taken when using certain values. For example, `idle=poll` maximizes performance but is incompatible with hyperthreading (HT) and will lead to very high power consumption. For additional information, refer to the documentation for your version of Linux.

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

i Note: 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.

! Caution: As the File System Manager (FSM) supports over 1000 clients (with more than 1000 file requests per client), the resource limits of your MDC may be exhausted with additional load from other processes. Exceeding the file descriptor limit will cause errors to your system. Quantum recommends you **not** run additional applications on the MDC.

i Note: The performance of the storage holding the Storage Manager database and metadata archives has an impact on overall StorNext file system performance. Quantum recommends that the Storage Manager database and metadata archives 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.

i 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. Required dependencies include **gcc**, **make**, and **kernel-headers** on Redhat or **kernel-source** on SUSE.

i 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

Separate licenses are required for different StorNext components (such as File System and Storage Manager) and features (such as Replication, HA and Distributed Data Mover).

For a list of all the StorNext licenses, see the [StorNext 6 Licensing Guide](#).

Disk licensing has been added to StorNext in order to help quantify customers' configurations. Disk licenses are required on all StorNext Metadata Controllers (MDCs) but may be updated after upgrade. For additional information on disk licensing, see the [StorNext 6 Licensing Guide](#).

Upgrade Implications

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 StorNext.
- 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.
- 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 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 displays 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>
```

i 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.

Upgrade from a snsm_was License to a object_storage License

If you are upgrading to StorNext 6 (and later), and you have an existing snsm_was license:

1. Obtain a new license file for your system that includes the new object_storage license.
2. Manually create a staging area for the license update at /usr/cvfs/update/config. Do this on both the primary and secondary nodes in an HA environment.

3. Manually copy in the new `license.dat` file to `/usr/cvfs/update/config`. Do this on both the primary and secondary nodes in an HA environment.
4. Upgrade to StorNext 6 (and later) per the applicable procedure.

Upon a successful upgrade, use the StorNext GUI to verify the license `object_storage` is installed.

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 6 (and later):

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 [Upgrade non-HA Linux MDCs on page 18](#).
4. Log in to the StorNext GUI.
5. Use the StorNext Configuration Wizard to finish configuring your system.

Upgrade the Client Software

To upgrade to the StorNext 6 (and later) client software on Linux 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, see the current *StorNext Installation Guide*.

The steps required to upgrade clients to StorNext 6 (and later) includes the following:

1. Unmount all StorNext file systems.
2. Stopping the `cvfs` services by running `/etc/init.d/cvfs fullstop`, `/sbin/init.d/cvfs fullstop`, 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 the StorNext file systems.

StorNext Media on DVD

Installation and upgrade media is shipped on DVDs. If you plan to install from media, you must have a DVD ROM drive to perform the installation or upgrade. See [DVD Directories on page 1](#) for information on the directories used on the DVD.

Post-Upgrade Browser Refresh

Quantum recommends to clear the browser and cookie cache after upgrading to 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 the following sections 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.

i Note: If you are performing an upgrade on a single MDC node system using RPM files, see **Installing StorNext File System-Only from RPMs** in the current *StorNext Installation Guide*.

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 800 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.

i Note: When selecting install packages to download from the web, observe that the **stornext_full** directory includes the Storage Manager install packages and the **fs_only** directory installs the StorNext File System only.

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.



Chapter 3: Run the Pre-Installation Script

This chapter contains the following topics:

Run the Pre-Installation Script	12
Switch to a Named Cluster Using the Command Line Interface	16

Run the Pre-Installation Script

StorNext has several dependency checks to make sure required RPM packages are installed. The dependency checks include environmental requirements, such as:

- Disk space
- LTFS appliance

The StorNext pre-installation script (**snPreInstall**) is included on the StorNext installation DVD (for non-Windows installations only). 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 the table below.

Table 1: 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/mdarchives	Stores metadata archives (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.

i Note: For optimal performance in non-HA configurations, 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?

i 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.

i 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.

Run 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.

i 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.

i 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 7 running on an x86 64-bit platform, change to the **RedHat7_x86_64** directory.
7. Run the script. At the command prompt, type:

```
./snPreInstall
```

The pre-installation script runs.

```
-----  
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>]:
```

8. Follow the onscreen instructions, and answer each question when prompted. The script outputs results based on the information you entered.

Interpret the 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.

i 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 the table below.

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, 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.

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

Switch to a Named Cluster Using the Command Line Interface

After an initial upgrade or install of StorNext 6 (or later), you may wish to switch from the default cluster to a named cluster.

File Systems (or more specifically File System Managers - FSMs) run in a specific cluster. Clients mount file systems, so a mount is also associated with a specific cluster.

In preparation for switching to a named cluster, all clients should un-mount the file systems, FSMs should be stopped and services stopped on MDCs and coordinator nodes. It is common to run MDCs as coordinators, but this section uses external coordinator nodes as examples.

The following example illustrates how to switch to cluster "**cluster1**", leaving the administrative domain as the default setting. On the coordinator nodes, you update the **fsnameservers** and create the **fsmcluster** in **/usr/cvfs/config**.

fsnameservers

old

10.0.0.1

10.0.0.2

new

10.0.0.1@cluster1

10.0.0.2@cluster1

fsmcluster

default_cluster cluster1

Start the services on the coordinator nodes. On the MDC, with services stopped, make the same changes to **fsnameservers** and **fsmcluster**. Do not change **fsmlist**. Make the following change to **/etc/fstab**:

old

snfs1 /stornext/snfs1 cvfs rw,noauto 0 0

new

```
snfs1@cluster1 /stornext/snfs1 cvfs rw,noauto 0 0
```

Start the services on the MDC.

On Linux client nodes, make the changes to **fsmcluster**, **fsnameservers** and **/etc/fstab** and start the services. Creating the **fsmcluster** file is optional on client nodes.

For Windows clients, changes to the **fsnameservers** file are made with the client configuration GUI. At this point start the services. The file systems appear with cluster information and can be mounted. The **fsmcluster** file can also be created on Window clients, but, like Linux clients, the creation of this file is optional.



Chapter 4: Upgrade non-HA Linux MDCs

This chapter contains the following topics:

Upgrade non-HA Linux MDCs	18
---------------------------------	----

Upgrade non-HA Linux MDCs

Use the StorNext installation script to upgrade 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.

i Note: If you are performing an upgrade on a single MDC node system using RPM files, see **Installing StorNext File System-Only from RPMs** in the current *StorNext Installation Guide*.

The following instructions apply to non-HA server upgrades. However, they may also be used to upgrade RPM-only configurations to a full StorNext configuration. For instruction on upgrading HA servers only, see [Upgrade an HA System on page 26](#).

This procedure describes how to upgrade a StorNext non-HA system on the Linux operating system. See the [StorNext 6 Compatibility Guide](#) for more information about valid upgrade paths.

Pre-Upgrade Steps

1. Quantum recommends that you run and complete a StorNext backup prior to starting the upgrade. If you have not completed a backup, you can initiate one from the StorNext GUI (see [System Backup](#)).

2. Make sure all store and retrieve requests have finished.
3. If you are using DDM, disable it on all the hosts by using the StorNext GUI or the `fsddmconfig` command. Storage Manager requires that the MDC and DDM clients are running the same version of the software. Refer to the section on **Distributed Data Mover** in the *StorNext User's Guide* for information on DDM.

Upgrade the MDC

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

i Note: When you mount a DVD in a Red Hat system, DVDs are mounted by default with a `noexec` (non-executable) option which prevents you from proceeding with the installation. For complete details on the supported versions of Red Hat, see the [StorNext 6 Compatibility Guide](#).

For Red Hat users only, before proceeding you must remount the DVD by typing `mount -o remount, exec ...`

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

```
ls -l
```

3. Change to the `stornext_full` directory.

i 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.

4. List the directories in the `stornext_full` directory. At the command prompt, type:

```
ls -l
```

5. Identify the correct installation directory for your operating system and hardware platform, and then change to that directory. For example, for Red Hat Linux 7 running on an x86 64-bit platform, change to the `RedHat7_x86_64` directory.
6. As described in [Get Ready To Upgrade on page 3](#), 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
```

7. 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 the next page](#).

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

```
./install.stornext
```

The installation script runs, as illustrated below.

```
Stornext Upgrade 6.2.0(79220E) MAIN MENU

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

Enter your choice <2>:
```

9. On the Upgrade Menu, type **2** and press **<Enter>**. The StorNext installation script runs and displays the current progress of the upgrade.
10. When the upgrade is complete, press **<Enter>** to return to the Main Menu.

i Note: A database schema update is applied to Storage Manager when upgrading from StorNext 4.3.x to StorNext 5. This update takes approximately one hour for every hundred million entries in the filecomp tables.

In a non-HA environment, database schema updates will be applied during the first start of StorNext from **install.stornext** following the upgrade to the new version of the StorNext software. StorNext file systems are accessible while the database schema is being updated, but Storage Manager functionality (including stores and retrieves) will be offline.

Do NOT interrupt StorNext services while the database is being updated. Interrupting the database schema update could result in an inconsistent database, and may require assistance from Quantum Support to repair or restore the database.

11. To exit the installation script, type **4** and press **<Enter>**.
12. After running the installer, source the profile on the MDC by running this command:

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

Post-Upgrade Steps

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

1. Upgrade the DDM hosts to same version of StorNext as the MDC.
2. Enable the DDM host by executing the `fsddmconfig` command or the StorNext GUI.

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.

i 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.



Chapter 5: Upgrade a Windows Client or Non-HA Server

This chapter contains the following topics:

Upgrade a Windows Client or Non-HA Server	23
---	----

Upgrade a Windows Client or Non-HA Server

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

i Note: Depending on your security settings, you must be **Administrator** or have **Administrator** privileges/rights to install StorNext software.

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

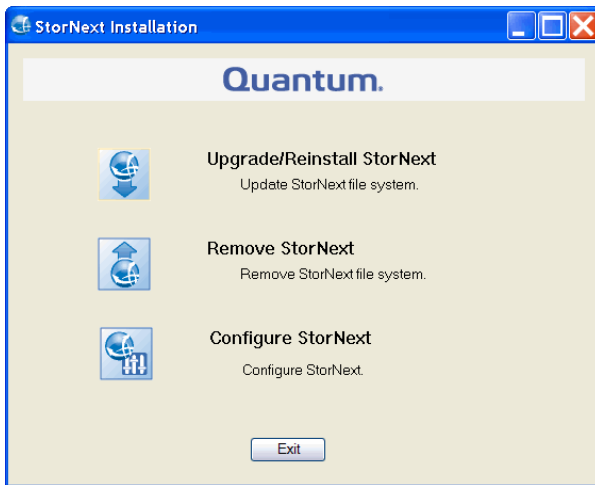
i Note: See the *StorNext Compatibility Guide* for supported upgrade versions for a given StorNext release.

Install StorNext on Windows clients or non-HA servers

Follow the procedure below to install StorNext on Windows clients or non-HA servers.

1. If any StorNext file systems are currently shared, temporarily disable sharing.
2. Copy the **snfs_client_Windows_x86_64.zip** file (for 64-bit systems only) 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.
3. Unzip the file.
4. Open the unzipped folder and double-click the client software installer file. This file is named **SnfsSetup.exe**. The StorNext **Installation** screen appears.

i Note: If you are not logged in with Administrator privileges, right-click the file and choose **Run as Administrator**.



5. Click **Upgrade/Reinstall StorNext**. The StorNext **Component Selection** window appears.

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



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 may 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, etc. prohibit the automatic restart feature. If the installation is not automatically completed (for example, 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.



Chapter 6: Upgrade an HA System

This chapter contains the following topics:

Upgrade an HA System	26
----------------------------	----

Upgrade an HA System

This section describes how to upgrade a High Availability (HA) system using primary and secondary MDC nodes.

i Note: This section is not intended for single-node MDCs. If you are performing an upgrade on a single MDC node system using RPM files, see **Installing StorNext File System-Only from RPMs** in the current *StorNext Installation Guide*.

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 running commands from the command line. You should also have a solid understanding of the HA feature in StorNext and a general familiarity with the StorNext GUI.

-
- Note:** 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.
-
- Note:** If you use scripts in your environment to upgrade - removing StorNext from the secondary node first, then upgrading the primary node, they can still be used to perform the HA upgrade.
-
- Note:** 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.
-
- Caution:** 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.

Upgrade a StorNext HA System on the Linux OS

This procedure describes how to upgrade a StorNext HA system on the Linux operating system. See the [StorNext 6 Compatibility Guide](#) for more information about valid upgrade paths.

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 21](#) if you are not sure you have the necessary licenses for StorNext.

Pre-Upgrade Steps

1. It is recommended that you run and complete a StorNext backup prior to starting the upgrade. If you have not completed a backup, you can initiate one from the StorNext GUI.
2. Make sure all store and retrieve requests have finished.
3. If you are using DDM, disable it on all the hosts by using the StorNext GUI or the `fsddmconfig` command. Storage Manager requires that the MDC and DDM clients are running the same version of the software. Refer to the section on **Distributed Data Mover** in the *StorNext User's Guide* for information on DDM.

Upgrade the Secondary MDC Node

1. Login to command line of the secondary MDC node.
2. Identify the correct installation directory for your operating system and hardware platform, and then change to that directory. For example, for Red Hat Linux 7 running on an x86 64-bit platform, change to the `RedHat7_x86_64` directory.
3. Upgrade StorNext on the secondary MDC node by running the following command:

```
./install.stornext
```

Upgrade the Primary MDC Node

1. After the upgrade on the secondary node completes, log in to command line of the primary MDC node.
2. Upgrade StorNext on the primary MDC node by running the following command:

```
./install.stornext
```

-
- i Note:** Upgrading the primary MDC node triggers a failover of control to the secondary MDC node. At this point, the secondary node becomes the new primary node. Client connectivity can be checked to ensure that the new primary node is functioning properly.
-
- i Note:** The new primary MDC node converts the old StorNext file system to the new StorNext file system in the background. During this process, it is possible the old primary MDC running the older version of StorNext will still try to access the new StorNext file system. Because the file system is converting, the old primary StorNext log may show **ICB** error messages in the log.
-
- i Note:** Also, while the file system is being converted, stores and retrieves will not be available until the conversion is completed. It is possible that this process could take hours.

Initiate a Graceful System Failover and Verification

During the upgrade of the node originally operating as the primary, the node originally operating as the secondary node becomes the primary. After the database updates are complete and the upgrade of the original primary completes, a failover of the node currently operating as the primary should be triggered so it once again is operating as the secondary, and to reinstate the original primary node to operate as primary again. To initiate the graceful failover of an HA pair, perform the following procedure.

1. Open an SSH connection to the node operating as the **primary**.
2. Log in to the command line of the **primary** node as root.
3. Confirm that the node is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

4. Verify the output is (bold used for clarification):

```
:default:primary:default:running:
```

5. On the node operating as the **primary**, initiate an HA failover to the node operating as the **secondary**.

```
service cvfs stop
```

6. Wait until the **secondary** node becomes the **primary**, and leave your SSH connection to this node open. (Time may vary.)
7. Open an SSH connection to the node now operating as the **primary**.
8. Log in to the command line of the **primary** node as root.
9. Confirm that the node is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

10. Verify the output is:

```
:default:primary:default:stopped:
```

11. From the SSH connection to the node now operating as the **secondary**, enter the following:

```
service cvfs start
```

12. Confirm that the node is operating as the **secondary** by entering the following at the command line:

```
snhamgr -m status
```

13. Verify the output is:

```
:default:running:default:primary:
```

14. Repeat if desired to fail over to the original system operating as the **primary**.
15. Verify that all clients have full access.
16. Test access to all file systems.

Post-Upgrade Steps

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

1. Upgrade the DDM hosts to same version of StorNext as the MDC.
2. Enable the DDM host using the **fsddmconfig** command or the StorNext GUI.

Upgrade a Windows 4.3.x (or later) HA System

To upgrade a Windows HA server to the latest version of StorNext, perform the following procedure:

1. On the secondary MDC, navigate to **Start Menu > Programs > StorNext File System > Stop and Remove** to stop the File System Managers (FSMs).
2. Upgrade the secondary MDC 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.

```
cvadminfail <fs_name>
```

5. Now the upgraded MDC should now be the primary MDC.
6. Repeat **Step 1** through **Step 3** on the other (now secondary) MDC.
7. You can initiate a failover of the MDCs to verify HA is correctly failing over if desired.

How to Configure the fsCapacityThreshold Configuration File Parameter

Beginning with StorNext 6.3, StorNext automatically monitors the used capacity of the HA shared file system and generates a RAS event when the file system reaches a capacity of 85% or higher.

i Note: StorNext continues to generate a RAS event once an hour until the used capacity level falls below 85%.

The percent at which a RAS event is generated is controlled by the HA shared file system's configuration file parameter **fsCapacityThreshold**.

i Note: If the **fsCapacityThreshold** parameter is not present in the configuration file or if it is set to zero (0), then the default of 85% is used.

To effectively disable this RAS event for the HA shared file system, set the **fsCapacityThreshold** parameter to **100**.

Do the following to change the **fsCapacityThreshold** parameter for the HA shared file system:

Step 1: On the secondary MDC

Run the following command to set the secondary MDC to the **locked** state; this shuts down the cvfs services on the MDC:

```
snhamgr mode=locked
```

Step 2: On the primary MDC

1. Run the following command to set the primary MDC to the **config** state:

```
snhamgr mode=config
```

2. Edit the configuration file to change the **fsCapacityThreshold** parameter to a value between **0** and **100** inclusive:

```
sncfgedit /usr/cvfs/config/<ha_shared_fs_name>.cfgx
```

3. Run the following command to return the primary MDC to the **default** state; this shuts down the cvfs services on the MDC:

```
snhamgr mode=default
```

4. Run the following command to restart the cvfs services:

```
service cvfs start
```

5. Run the following command and wait for the MDC status to become **primary**; this is displayed by the **LocalStatus** parameter.

```
snhamgr status
```

Below is an example of the output; note the **LocalStatus=primary** line.

```
LocalMode=default  
LocalStatus=primary  
RemoteMode=locked  
RemoteStatus=stopped
```


Step 3: On the secondary MDC

1. Run the following command to return the secondary MDC to the **default** state:

```
snhamgr mode=default
```

2. Run the following command to restart the cvfs services:

```
service cvfs start
```

After you complete the procedure, StorNext begins to monitor the used capacity of the HA shared file system using the updated **fsCapacityThreshold** value.



Chapter 7: Upgrade an HA System Using RPM Files

This chapter contains the following topics:

Upgrade an HA System Using RPM Files	33
--	----

Upgrade an HA System Using RPM Files

This section describes how to upgrade to a High Availability (HA) system with primary and secondary MDC nodes using RPM files.

i Note: This section is not intended for single-node MDCs. If you are performing an upgrade on a single MDC node system using RPM files, see **Installing StorNext File System-Only from RPMs** in the current *StorNext Installation Guide*.

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

Upgrade a StorNext 4.1 (or later) HA System using RPM files

You can upgrade directly from StorNext 4.1 if you are using RPM files to upgrade the MDC. Acquire the StorNext client, server and core StorNext RPMs for your server OS architecture from the installation DVD, or by downloading the tarball from the Quantum web site. For example:

- a. quantum_jansson-<version>.rpm
- b. quantum_libevent-<version>.rpm

- c. quantum_protobuf-<version>.rpm
- d. quantum_zeromq-<version>.rpm
- e. snfs-common-<version>.rpm
- f. snfs-client-<version>.rpm
- g. snfs-server-<version>.rpm
- h. snfs-extapi-<version>.rpm
- i. snupdate-<version>.rpm

i Note: Make sure you have obtained and installed the proper licenses for both nodes for the targeted build (such as StorNext). Separate licenses are required for different StorNext components (such as File System and Storage Manager) and features (such as Replication, HA and Distributed Data Mover). For a list of all the StorNext licenses, refer to the current *StorNext Licensing Guide* available online at <http://www.quantum.com/snsdocs>. Disk licensing has been added to StorNext in order to help quantify customers' configurations. Disk licenses are required on all StorNext Metadata Controllers (MDCs) but may be updated after upgrade. For additional information on disk licensing, see the [StorNext 6 Licensing Guide](#).

Stop the StorNext File System processes on each client machine by running `/etc/init.d/cvfs fullstop`.

On the Secondary MDC

1. Stop the StorNext services by running the following command:

```
service cvfs fullstop
```

2. Upgrade the required .rpm files in the following order. Use the command `rpm -Uvh` before each package. For example, `rpm -Uvh quantum_jansson-<version>.rpm`. The files are installed in the directory: `/usr/cvfs`.
 - a. quantum_jansson-<version>.rpm
 - b. quantum_libevent-<version>.rpm
 - c. quantum_protobuf-<version>.rpm
 - d. quantum_zeromq-<version>.rpm
 - e. snfs-common-<version>.rpm
 - f. snfs-client-<version>.rpm
 - g. snfs-server-<version>.rpm
 - h. snfs-extapi-<version>.rpm
 - i. snupdate-<version>.rpm

3. Start cvfs by running the following command:

```
service cvfs start
```

4. Verify that all SNFS mounts completed successfully.

On the Primary MDC

5. Stop StorNext services by running the following command:

```
service cvfs fullstop
```

6. Verify that the Secondary MDC is now serving the configured file systems by running the following command:

```
cvadmin -e select
```

7. Upgrade the required .rpm files in the following order. Use the command `rpm -Uvh` before each package. For example, `rpm -Uvh quantum_jansson-<version>.rpm`. The files are installed in the directory: `/usr/cvfs`.

- a. `quantum_jansson-<version>.rpm`
- b. `quantum_libevent-<version>.rpm`
- c. `quantum_protobuf-<version>.rpm`
- d. `quantum_zeromq-<version>.rpm`
- e. `snfs-common-<version>.rpm`
- f. `snfs-client-<version>.rpm`
- g. `snfs-server-<version>.rpm`
- h. `snfs-extapi-<version>.rpm`
- i. `snupdate-<version>.rpm`

8. Start cvfs by running the following command:

```
service cvfs start
```

9. Verify that all SNFS mounts completed successfully.

Back on the Secondary MDC

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

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

11. 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>"
```



Chapter 8: Upgrade Clients

This chapter contains the following topics:

Upgrade a Client	37
------------------------	----

Upgrade a Client

This section describes how to upgrade a client. Refer to the instructions pertaining to your particular operating system.

Procedures for how to upgrade a Windows client and server for non-HA systems are described in [Upgrade a Windows Client or Non-HA Server on page 23](#).

StorNext Software Requirements

Upgrading a Debian Linux client requires installing these quantum-supplied packages:

- zeromq
- protobuf

These packages are used by StorNext and are bundled with the StorNext .bin file.

i Note: When upgrading RedHat and SUSE Linux clients, the **snupdate** service installs these packages.

Upgrading Redhat and SUSE mover clients requires installing these additional quantum-supplied DDM packages:

- quantum_curl
- quantum_openssl
- snfs-mover
- quantum_unixODBC
- snlfs

i Note: Upgrading these packages might require using the `--force` option.

Upgrade Linux Clients

The **snupdate** service allows you to manage StorNext software packages by providing the ability to install, upgrade, remove, and rollback StorNext software packages. The **snupdate** service works with supported clients listed in the [StorNext 6 Compatibility Guide](#).

Packages can be applied immediately or they can be staged and updated at a later time. The snupdate service is a boot time script that starts before cvfs. It checks to see if there are staged updates available. If an update has been staged, the snupdate service will apply that update.

The snupdate service, located in the `/etc/init.d/` directory, calls the snupdate utility (`/opt/quantum/snupdate/bin/snupdate`) which performs the actions and manages the snupdate repository (`/opt/quantum/snupdate/var`).

The snupdate repository consists of a set of packages. These packages are the contents of the extracted StorNext **.bin** archives. The snupdate has the ability to install, update, and remove a package that is copied into the repository. It also has the ability to rollback a package to a previous version, provided that the package supports the ability to downgrade its version.

The snupdate utility can stage an install, upgrade, or rollback (if the package supports rollback) by marking in the repository the next version that the package should be updated to. The staged action is applied at a later time by the snupdate apply command.

Multiple versions of a package can be copied into the snupdate repository. The snupdate utility will determine which version of a package is currently installed and keep track of the previously installed version for the purpose of rollback. It can also detect if there is a newer version in the repository for upgrades.

i Note: You must use the latest version of **snupdate** when you install or upgrade StorNext.

Get Help

The snupdate utility supports an interactive help system. Running `/opt/quantum/snupdate/bin/snupdate help` will provide a list of available snupdate commands. The table below lists the available commands.

The help system syntax is:

```
/opt/quantum/snupdate/bin/snupdate help [COMMAND]
```

To get additional information on a command, type:

```
/opt/quantum/snupdate/bin/snupdate help import
```

Command	Description
apply	Apply a staged action.
cancel	Cancel a staged action.
clean	Delete a package from the repository.
help	Print extended help on a command.
import	Import packages into the repository.
install	Install package on the system.
remove	Remove package from the system.
rollback	Rollback package.
show	Display package and version information.
stage	Stage an upgrade.
status	Status for staged actions.
upgrade	Upgrade package on the system.

How to Install or Upgrade the snupdate Service

Note: You must use the latest available version of **snupdate** when you install StorNext.

To install or upgrade the **snupdate** service, complete the following:

1. Run the StorNext **.bin** command to extract a StorNext package. The snupdate package is bundled with the StorNext **.bin** file.

Note: The StorNext package will extract into `/tmp/stornext` directory by default. If this directory exists, the installation will fail and you will receive an error message. You will need to identify an alternate location into which to extract the files by entering the following command, for example:


```
# ./<archive name> -d <dir>
```

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

2. Change to the directory where you extracted the contents of the **.bin** file.
3. Install the snupdate service by running this command:
 - On RedHat or SUSE Linux execute:

```
# rpm -Uvh snupdate-<version>.rpm
```

- On Debian or Ubuntu Linux execute:

```
# dpkg --install snupdate-<version>.deb
```

4. Run the following command to check the installed version of **snupdate**:
 - On RedHat or SUSE Linux execute:

```
# rpm -q snupdate
```

- On Debian or Ubuntu Linux execute:

```
# dpkg -list snupdate
```

Import a StorNext Package Into a Repository

To import a StorNext package into an snupdate repository, complete the following:

1. Import the `/tmp/stornext` package into the snupdate repository. Snupdate will determine the package name and version of the package and copy the package into `/opt/quantum/snupdate/var/PACKAGE/VERSION`, for example:

```
# /opt/quantum/snupdate/bin/snupdate import /tmp/stornext/  
Checking for new packages in /tmp/stornext  
Copying 6.0.0-38570/ to /opt/quantum/snupdate/var/snfs_client/6.0.0-38570
```

If you extracted the .bin file into another directory, such as /tmp/stornext6/, then import the contents of the directory to the snupdate repository.

```
# /opt/quantum/snupdate/bin/snupdate import /tmp/stornext6/
Checking for new packages in /tmp/stornext6
Copying 6.0.0-38570/ to /opt/quantum/snupdate/var/snfs_client/6.0.0-38570
```

2. Run the show versions command to ensure the package was imported:

```
# /opt/quantum/snupdate/bin/snupdate show versions
package: snfs_client
version: 5.0.0-38570 (Current)
version: 6.0.0-38570
package: snfs_full
version: (none)
package: stornext_full
version: (none)
```

Upgrade a Package in the Repository

To immediately upgrade a package, complete the following:

1. Upgrade to the latest version of the **snfs_client** that is available in the /opt/quantum/snupdate/var repository:

```
#/opt/quantum/snupdate/bin/snupdate upgrade snfs_client
Upgrading snfs_client package /opt/quantum/snupdate/var/snfs_client/6.0.0-38570...
Preparing... ##### snfs-common
##### snfs-client
#####
CC [M] /usr/cvfs/install/linuxif_fs.o
CC [M] /usr/cvfs/install/linuxif_fs_common.o
CC [M] /usr/cvfs/install/linuxif_fsctl.o
LD [M] /usr/cvfs/install/cvfs.o
Building modules, stage 2.
MODPOST 1 modules
CC /usr/cvfs/install/cvfs.mod.o
```

```
LD [M] /usr/cvfs/install/cvfs.ko.unsigned  
NO SIGN [M] /usr/cvfs/install/cvfs.ko  
INSTALL /usr/cvfs/install/cvfs.ko  
DEPMOD 2.6.32-279.el6.x86_64
```

2. Run the **show versions** command to verify the installed version. The output will show a status of Current, as shown:

```
#/opt/quantum/snupdate/bin/snupdate show versions snfs_client  
package: snfs_client  
version: 5.0.0-38570 (Previous)  
version: 6.0.0-38570 (Current)
```

Stage an Upgrade

To stage an upgrade of the package to another version, complete the following:

1. Assuming the packages have been imported as shown in the previous procedure, list the available versions.

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client  
package: snfs_client  
version: 5.0.0-38570 (Current)  
version: 6.0.0-38570
```

2. Verify that the package update is available:

```
# /opt/quantum/snupdate/bin/snupdate status  
Upgrade available for package: snfs_client (6.0.0-38570)
```

3. Stage the upgrade using the **snupdate stage** command:

```
#/opt/quantum/snupdate/bin/snupdate stage upgrade snfs_client  
Staging upgrade of snfs_client to version 6.0.0-38570
```

4. Check on the status of snfs_client package to see that it has been marked to be updated:

```
#/opt/quantum/snupdate/bin/snupdate status  
Staged upgrade scheduled for package: snfs_client (6.0.0-38570)
```

You could also verify the updates by running the **snupdate service** command:

```
# snupdate service status  
Staged upgrade scheduled for package: snfs_client (6.0.0-38570)
```

5. Verify the new update has been staged by running the **show versions** command:

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client package: snfs_  
client  
version: 5.0.0-38570 (Current)  
version: 6.0.0-38570 (StagedUpgrade)
```

Apply a Staged Upgrade

1. Verify that a staged update is available.

```
# /opt/quantum/snupdate/bin/snupdate status  
Staged upgrade scheduled for package: snfs_client (6.0.0-38570)
```

2. The cvfs service will need to be stopped and unloaded prior to starting the snupdate service.
 - On Linux execute:

```
/usr/cvfs/bin/DSM_control fullstop
```

3. Staged updates are applied on the next reboot. Alternatively, you can apply a staged update manually.

```
# /opt/quantum/snupdate/bin/snupdate apply  
Checking for new packages in /tmp/stornext package  
/opt/quantum/snupdate/var/snfs_client/6.0.0-38570 is already up to date  
Staged upgrade scheduled for package: snfs_client (6.0.0-38570)  
Upgrading snfs_client package /opt/quantum/snupdate/var/snfs_client/5.0.0-  
38570...
```

```
Preparing... ##### quantum_zeromq
##### quantum_protobuf
#####
snfs-common ##### snfs-client
#####
CLEAN /usr/cvfs/install/.tmp_versions
CLEAN /usr/cvfs/install/Module.symvers /usr/cvfs/install/modules.order
CC [M] /usr/cvfs/install/linuxif_fs.o
CC [M] /usr/cvfs/install/linuxif_fs_common.o
CC [M] /usr/cvfs/install/linuxif_fsctl.o
LD [M] /usr/cvfs/install/cvfs.o
Building modules, stage 2.
MODPOST 1 modules
CC /usr/cvfs/install/cvfs.mod.o
LD [M] /usr/cvfs/install/cvfs.ko.unsigned
NO SIGN [M] /usr/cvfs/install/cvfs.ko
INSTALL /usr/cvfs/install/cvfs.ko
DEPMOD 2.6.32-131.0.15.el6.x86_64
No staged action defined for snfs_full
No staged action defined for stornext_full
```

4. After running **snupdate** to upgrade, rollback, or apply a staged operation, the cvfs service will need to be restarted:
 - On Linux execute:

```
/usr/cvfs/bin/DSM_control start
```

Rollback to a Previous Version

To rollback to the previously installed version, complete the following:

1. Identify the versions currently in the repository using the **show versions** option, for example:

```
#!/opt/quantum/snupdate/bin/snupdate show versions snfs_client
package: snfs_client
version: 5.0.0-38570 (Previous)
version: 6.0.0-38570 (Current)
```

2. Stage a rollback to a previously installed version by using the **stage rollback** option. You can also use

a specific version to rollback to a different version instead of the last installed version:

```
#!/opt/quantum/snupdate/bin/snupdate stage rollback \ snfs_client 5.0.0-38570  
Staging rollback of snfs_client to version 5.0.0-38570
```

3. Use the option, **show versions**, to list the current versions of a package and the intended states.

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client  
package: snfs_client  
version: 5.0.0-38570 (Previous,StagedRollback)  
version: 6.0.0-38570 (Current)
```

4. Use the **snupdate apply** command or the **service snupdate start** command to make the StorNext software update take effect.

```
# /opt/quantum/snupdate/bin/snupdate apply snfs_client  
Rolling back snfs_client to /opt/quantum/snupdate/var/snfs_client/5.0.0-  
38570...  
CLEAN /usr/cvfs/install/.tmp_versions  
CLEAN /usr/cvfs/install/Module.symvers /usr/cvfs/install/modules.order  
Preparing... #####  
snfs-common ##### snfs-client  
#####  
CC [M] /usr/cvfs/install/linuxif_fs.o  
CC [M] /usr/cvfs/install/linuxif_fs_common.o  
CC [M] /usr/cvfs/install/linuxif_fsctl.o  
LD [M] /usr/cvfs/install/cvfs.o  
Building modules, stage 2.  
MODPOST 1 modules  
CC /usr/cvfs/install/cvfs.mod.o  
LD [M] /usr/cvfs/install/cvfs.ko.unsigned  
NO SIGN [M] /usr/cvfs/install/cvfs.ko
```

```
INSTALL /usr/cvfs/install/cvfs.ko  
DEPMOD 2.6.32-279.el6.x86_64
```

5. Use the **show versions** to show that rollback succeeded.

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client  
package: snfs_client  
version: 5.0.0-38570 (Current)  
version: 6.0.0-38570 (Previous)
```

Cancel a Staged Update

A staged update operation may be canceled using the **snupdate cancel** command. To cancel a staged update, complete the following:

1. Use the **snupdate show versions** command to list the available versions:

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client  
package: snfs_client  
version: 5.0.0-38570 (Previous,StagedRollback)  
version: 6.0.0-38570 (Current)
```

2. Use the **snupdate cancel** command to cancel the pending staged update for a package:

```
# /opt/quantum/snupdate/bin/snupdate cancel snfs_client  
package: snfs_client  
Canceling staged rollback for package snfs_client
```

3. Verify that the staged update has been canceled using the **show versions** command:

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client  
package: snfs_client  
version: 5.0.0-38570(Previous)  
version: 6.0.0-38570 (Current)
```

Clean the snupdate Repository

Use the **snupdate clean** command to clear stale versions from the snupdate repository. Complete the following:

1. List the current packages using the **show versions** command.

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_clientpackage: snfs_client
version: 5.0.0-38570 (Previous)
version: 6.0.0-38570 (Current)
```

2. Delete the stale StorNext packages from the repository in the **/opt/quantum/snupdate/var** directory using the **clean** command. The **clean** command requires both a package name and a version:

```
# /opt/quantum/snupdate/bin/snupdate clean snfs_client 5.0.0-38570
Removing /opt/quantum/snupdate/var/snfs_client/5.0.0-38570 from repository
```

3. Verify the package has been removed from the repository using the **show versions** command:

```
# /opt/quantum/snupdate/bin/snupdate show versions snfs_client
package: snfs_client
version: 6.0.0-38570 (Current)
```

Upgrade a Debian Linux Client


Follow these steps to upgrade a Debian Linux client.

1. Log on to the StorNext client system as **root**.
2. Change to the directory where the client software archive file you downloaded is located.
3. Extract the Debian client package (for example, **snfs_client_Debian8_x86_64.bin**). When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA, type **y** to accept the EULA, or enter any other key to decline. After you accept the EULA, the client software is extracted to **/tmp/stornext** (or to another directory if you specified an alternate location).

i Note: If the directory into which the files are being extracted does not have enough space, the installation will fail and you will receive an error message. You will need to identify an alternate location into which to extract the files by entering the following command:


```
# ./<archive name> -d <dir>
```

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

 **Caution:** If the directory into which the files are being extracted already exists, you will receive a warning message. If this occurs, either remove the directory and try the installation again, or change the installation directory by entering the following command:


```
# ./<archive name> -d <dir>
```

where **<archive name>** is the name of the client software archive file you downloaded, and **<dir>** is the directory into which you want to extract the installation files. It is NOT recommended that you force using an existing directory (by using the `-f` command) because this could lead to additional problems installing the `.rpm` files.

4. From the **/tmp/stornext** directory, install the required SNFS packages first, and then install the client and common software packages. For example:

```
dpkg -i quantum-protobuf_<version>.deb  
dpkg -i quantum-zeromq_<version>.deb  
dpkg -i snfs-common-<version>.deb  
dpkg -i snfs-client-<version>.deb
```

5. Edit the **/usr/cvfs/config/fsnameservers** text file to contain the IP address of the MDC to which the StorNext client will connect. The **fsnameservers** file on the client must be exactly the same as on the MDC. If the **fsnameservers** file does not exist, use a text editor to create it.

 **Note:** The **fsnameservers** file needs to contain only the IP address of the MDC.

6. If this is a SAN client, check to make sure LUNs are visible. At the command prompt, type:

```
/usr/cvfs/bin/cvlabel -l
```

7. Create a mount point. For example:

```
/stornext/<snfs1>
```

Where **<snfs1>** is the name of the StorNext File System/Xsan volume you are using on your system.

8. Edit the **/etc/fstab** to add the client mount options. For example:

```
<snfs1> /stornext/snfs1 cvfs rw,noauto 0 0
```

Where **<snfs1>** is the name of the StorNext File System/Xsan volume you are using on your system.

9. Start the file system service. At the command prompt, type:

```
service cvfs start
```

10. Verify file systems are mounted.



Chapter 9: Upgrade a StorNext Storage Manager Object Storage Configuration

This chapter contains the following topics:

Upgrade a StorNext Storage Manager Object Storage Configuration50

Upgrade a StorNext Storage Manager Object Storage Configuration

If you have existing systems configured with Object Storage, upon upgrading to StorNext 5 release 5.1 (or later), you will notice the following changes to your configuration:

Component	Description
Appliance Component	The Provider will be set to LATTUS .
Controller Component	The HTTP and HTTPS Protocol and ioport are no longer an attribute of the Controller. These attributes are now owned by the iopath component. During upgrade, these parameters are no longer set for the Controller Component.

Component	Description
Iopath Component	<p>The HTTP HTTPS Protocol, AXR S3 Object Storage Service API, PATH VHOST are attributes of the iopath. The host address is now defined as a TCP connection endpoint. This parameter now includes an IP address (or a DNS name) and a TCP/IP port number (for example, 123.123.123.123:8080). The port number is Optional. During an upgrade, the following changes are made:</p> <ul style="list-style-type: none"> • An iopath is created for each defined ioport. • The alias of the iopath is modified. • The Object Storage Service API is set to AXR. • The path style is set to PATH.
Namespace Component	<p>The namespace component now contains AXR or S3. During an upgrade, this parameter is set to AXR.</p>
DDM Device Component	<p>The DDM device configuration are also changed to reflect the new iopath configuration.</p>

Figure 1: Example of an Object Storage configuration before a StorNext 5 release 5.1 upgrade

```

Example of an Object Storage configuration before Stornext 5.1 upgrade

Alias      Type      Strms Protocol IP Address      Port(s)
-----
was1      APPLIANCE  0      https  10.65.166.80    443

mdh2-controller1
  iop1     NODE      48     http   192.168.20.1   8082,8081,8084,8085
  iop2     IOPATH    0      http   192.168.10.1

NOTE: There are 4 ioport and 2 iopaths per controller, in the above configuration
    
```

Figure 2: Example of a DDM Object Storage device configuration before a StorNext 5 release 5.1 upgrade

```

Example of a DDM Object Storage device configuration before Stornext 5.1 upgrade

Device Identifier Type      State      Host      Path
-----
iop1          LATTUS   Enabled   carlsbad2-1  192.168.20.1
              Enabled   carlsbad2-2  192.168.20.1
              Enabled   summit-9     192.168.20.1
              Enabled   summit-8     192.168.20.1
iop2          LATTUS   Enabled   carlsbad2-1  192.168.10.1
              Enabled   carlsbad2-2  192.168.10.1
              Enabled   summit-9     192.168.10.1
              Enabled   summit-8     192.168.10.1
    
```

Figure 3: Example of an Object Storage configuration after a StorNext 5 release 5.1 upgrade

Example of the Object Storage configuration after Stornext 5.1 upgrade

Alias	Type	Strms	Protocol	RESTAPI	URLSTYLE	IP Address	Port(s)	Provider
was1	APPLIANCE	0	https			10.65.166.80	443	LATTUS
mdh2-controller1	NODE	48						
iop1-1	IOPATH	0	http	AXR	PATH	192.168.20.1:8082		
iop1-2	IOPATH	0	http	AXR	PATH	192.168.20.1:8081		
iop1-3	IOPATH	0	http	AXR	PATH	192.168.20.1:8084		
iop1-4	IOPATH	0	http	AXR	PATH	192.168.20.1:8085		
iop2-1	IOPATH	0	http	AXR	PATH	192.168.10.1:8082		
iop2-2	IOPATH	0	http	AXR	PATH	192.168.10.1:8081		
iop2-3	IOPATH	0	http	AXR	PATH	192.168.10.1:8084		
iop2-4	IOPATH	0	http	AXR	PATH	192.168.10.1:8085		

NOTE: Before Stornext 5.1 upgrade, there were 4 ioports and 2 iopaths per controller in the configuration.
 After Stornext 5.1 upgrade, there are 8 iopaths per controller. The old iop1 now becomes iop1-1, iop1-2, iop1-3, iop1-4, and so on, depending on the number of configured ioports.

Figure 4: Example of a DDM Object Storage device configuration after a StorNext 5 release 5.1 upgrade

Example of the DDM Object Storage device configuration after Stornext 5.1 upgrade

Device Identifier	Type	State	Host	Path
iop1-1	LATTUS	Enabled	carlsbad2-1	192.168.20.1:8082
		Enabled	carlsbad2-2	192.168.20.1:8082
		Enabled	summit-9	192.168.20.1:8082
		Enabled	summit-8	192.168.20.1:8082
iop1-2	LATTUS	Enabled	carlsbad2-1	192.168.20.1:8081
		Enabled	carlsbad2-2	192.168.20.1:8081
		Enabled	summit-9	192.168.20.1:8081
		Enabled	summit-8	192.168.20.1:8081
iop1-3	LATTUS	Enabled	carlsbad2-1	192.168.20.1:8084
		Enabled	carlsbad2-2	192.168.20.1:8084
		Enabled	summit-9	192.168.20.1:8084
		Enabled	summit-8	192.168.20.1:8084
iop1-4	LATTUS	Enabled	carlsbad2-1	192.168.20.1:8085
		Enabled	carlsbad2-2	192.168.20.1:8085
		Enabled	summit-9	192.168.20.1:8085
		Enabled	summit-8	192.168.20.1:8085
iop2-1	LATTUS	Enabled	carlsbad2-1	192.168.10.1:8082
		Enabled	carlsbad2-2	192.168.10.1:8082
		Enabled	summit-9	192.168.10.1:8082
		Enabled	summit-8	192.168.10.1:8082
iop2-2	LATTUS	Enabled	carlsbad2-1	192.168.10.1:8081
		Enabled	carlsbad2-2	192.168.10.1:8081
		Enabled	summit-9	192.168.10.1:8081
		Enabled	summit-8	192.168.10.1:8081
iop2-3	LATTUS	Enabled	carlsbad2-1	192.168.10.1:8084
		Enabled	carlsbad2-2	192.168.10.1:8084
		Enabled	summit-9	192.168.10.1:8084
		Enabled	summit-8	192.168.10.1:8084
iop2-4	LATTUS	Enabled	carlsbad2-1	192.168.10.1:8085
		Enabled	carlsbad2-2	192.168.10.1:8085
		Enabled	summit-9	192.168.10.1:8085
		Enabled	summit-8	192.168.10.1:8085

Converting an AXR Namespace to an S3 Bucket

With Lattus 3.5.1, it is possible to convert AXR namespaces to S3 buckets and make them accessible via the S3 interface. StorNext provides the capability to convert the media type from AXR to S3.

To convert one or more AXR namespaces to S3 buckets, see [Converting an AXR Namespace to an S3 Bucket](#).



Chapter 10: Upgrade Firmware

This chapter contains the following topics:

Upgrade Firmware	54
Appliance Release Notes	54

Upgrade Firmware

For instructions on upgrading your firmware, see [Upgrade the System \(Upgrade Firmware\)](#) on the *Appliance InfoHub Documentation Center* (www.quantum.com/ApplianceInfoHub).

Appliance Release Notes

Refer to the respective Release Notes document for important information you should know about your system.

- [Xcellis Foundation](#)
- [aiWARE for Xcellis](#)
- [Xcellis Workflow Extender](#)
- [Xcellis Workflow Director](#)

- [Artico](#)
- [Pro Foundation](#)
- [G300](#)
- [M660](#)
- [M440](#)



Chapter 11: Compression and Encryption Usage Report

This chapter contains the following topics:

Compression and Encryption Usage Report	56
---	----

Compression and Encryption Usage Report

With compression and encryption for Q-Cloud devices, you can request to view compression and encryption usage information. The compression and encryption usage information is maintained in two new tables:

- `filecomp_obj`
- `classobj_info`

The `fsobjinfo` Command

Execute the command `fsobjinfo` to generate the compression and encryption usage report. The command `fsobjinfo` produces a summary usage report for object store media. Object store usage is summarized based on object store media ID and policy class ID. Reported usage can be limited to the optionally specified set of policy class IDs or object store media IDs.

In order to generate accurate reports, existing Q-Cloud Archive usage must be accounted for. The usage information in the `filecomp` table must be populated in the `filecomp_obj` and `classobj_info` tables. The command `qcloud_migrate.pl` provides the capability; execute the command `qcloud_migrate.pl` after you upgrade your system to StorNext 5 release 5.3.

The full path of the command is:

```
/usr/adic/TSM/util/install/qcloud_migrate.pl
```

i Note: Your system will operate normally without execution of the `qcloud_migrate.pl` command.

Considerations for the `qcloud_migrate.pl` Command

- You can execute the command anytime after an upgrade to StorNext 5 release 5.3.
- Execute the command `qcloud_migrate.pl` only once.
- The compression and encryption usage report may not be accurate if your system contains existing Q-Cloud Archive devices.
- If your system does not contain existing Q-Cloud Archive devices, do not execute the command `qcloud_migrate.pl`.