



Installation Guide

6-68045-02, Rev. H

Quantum[®]

StorNext 6 Installation Guide, 6-68045-02, Rev. H, 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. ActiveScale, DXi, DXi Accent, FlexSync, FlexTier, iLayer, Lattus, Quantum, The Quantum Logo, QXS, Scalar, StorNext, SuperLoader, Vision, 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. Quantum specifications are subject to change.



Contents

| | |
|---|-----|
| Preface | vii |
| Chapter 1: Begin Here | 1 |
| Begin Here | 1 |
| Audience | 2 |
| Installation Scenarios | 2 |
| Chapter 2: StorNext Linux MDC Installation and Configuration | 5 |
| StorNext Linux MDC Installation and Configuration | 6 |
| System Requirements | 7 |
| Getting Ready to Install | 12 |
| Run the Pre-Installation Script | 15 |
| Run the StorNext Installation Script | 19 |
| Post-Installation Steps | 23 |
| Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs | 24 |
| Install Optional StorNext Add-on Software Features and Clients | 30 |
| Convert to a High Availability (HA) System | 30 |

| | |
|---|------------|
| Configure StorNext Using the StorNext GUI | 37 |
| Linux Device Mapper Multipath Support | 41 |
| Chapter 3: Linux SNFS-Only MDC Installation and Configuration | 44 |
| Linux SNFS-Only MDC Installation and Configuration | 44 |
| File System Requirements | 45 |
| Getting Ready to Install | 46 |
| Install a StorNext File System Only Using an RPM | 49 |
| Configure a StorNext RPM Server Using the Command Line | 51 |
| Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs | 56 |
| Obtain and Install StorNext Licenses without a GUI | 62 |
| Chapter 4: Windows SNFS-Only MDC Installation and Configuration .. | 68 |
| Windows SNFS-Only MDC Installation and Configuration | 68 |
| File System Requirements | 69 |
| Getting Ready to Install | 71 |
| Install the StorNext File System on a Windows MDC: Interactively | 73 |
| Configure a StorNext File System on Windows | 82 |
| Configure a High Availability System on Windows | 136 |
| Add or Remove StorNext Components | 138 |
| Restore a Previous Client Configuration | 140 |
| Upgrade or Reinstall StorNext on a Windows Server | 141 |
| Chapter 5: StorNext Client Installation and Configuration | 143 |
| StorNext Client Installation and Configuration | 143 |
| StorNext Client System Requirements | 144 |
| Download the StorNext Client Software | 146 |
| Install the StorNext Client Software on Linux | 149 |
| Install the StorNext Client Software on Windows | 162 |

| | |
|---|------------|
| Configure the StorNext Client Software on Windows | 169 |
| Chapter 6: Configure LAN Clients and Gateway Servers | 181 |
| Configure LAN Clients and Gateway Servers | 181 |
| Configure a StorNext LAN Client on Linux | 182 |
| Configure a Gateway Server on Linux | 183 |
| Linux LAN Clients or Gateway Servers in HA Environments | 186 |
| Configure LAN Clients or Gateway Servers on Windows | 186 |
| Configure Gateway Functionality on a Metadata Controller (MDC) | 187 |
| Appendix A: Uninstall StorNext | 191 |
| Uninstall StorNext | 191 |
| Uninstall StorNext on Linux | 191 |
| Uninstall StorNext on Windows | 195 |
| Appendix B: Configuration Tasks | 198 |
| Configuration Tasks | 198 |
| Switch to a Named Cluster Using the Command Line Interface | 199 |
| Share a Library Between Applications (Linux Only) | 200 |
| Manually Modify the Storage Manager Configuration | 201 |
| Connect an Xsan Client to a StorNext MDC | 203 |
| For Systems Running StorNext 6.x (or earlier) and macOS 11.x (or later) | 205 |
| For Systems Running StorNext 6.x (or earlier) and macOS 10.x (or earlier) | 217 |
| Special Considerations for Xsan Mount Restrictions | 219 |
| Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client (DLC) | 221 |
| View a Fibre Channel Device | 223 |
| StorNext Log Rolling Administration | 223 |
| Avoid Poor System Performance | 224 |
| Avoid Not Saving cvmkfs Coredump Files | 225 |
| Import a Secure Sockets Layer (SSL) Certificate | 225 |

| | |
|---|------------|
| ActiveScale™ P100/X100 Installation, Configuration and Tuning | 229 |
| Appendix C: StorNext Configuration Files | 231 |
| StorNext Files | 231 |
| Example FSM Configuration File | 231 |
| Cluster-Wide Central Control | 242 |
| Appendix D: Quality of Service Feature | 250 |
| Quality of Service | 250 |
| About Quality of Service | 251 |
| Configuration | 252 |
| Real-time I/O | 258 |
| Non-Realtime I/O Configuration | 261 |
| Callbacks | 262 |
| Monitoring | 266 |
| Sharing Disk Resources | 267 |



Preface

This manual provides installation and configuration procedures for StorNext software and contains the following chapters:

- [Chapter 1: Begin Here](#)
- [Chapter 2: StorNext Linux MDC Installation and Configuration](#)
- [Chapter 3: Linux SNFS-Only MDC Installation and Configuration](#)
- [Chapter 4: Windows SNFS-Only MDC Installation and Configuration](#)
- [Chapter 5: StorNext Client Installation and Configuration](#)
- [Chapter 6: Configure LAN Clients and Gateway Servers](#)
- [Appendix A: Uninstall StorNext](#)
- [Appendix B: Configuration Tasks](#)
- [Appendix C: StorNext Configuration Files](#)
- [Appendix D: Quality of Service Feature](#)

Audience

This manual is written for StorNext 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. | ./install |
| Computer output and command line examples are shown in monospace font. | ./install |
| User input variables are enclosed in angle brackets. | http://<ip_address>/cgi-bin/stats |
| For UNIX and Linux commands, the command prompt is implied. | ./install is the same as # ./install |
| 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:

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

Contacts

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

<https://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

Quantum University, 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 University 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:
<https://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:

<https://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:

<https://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:

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



Chapter 1: Begin Here

This chapter contains the following topics:

| | |
|------------------------------|---|
| Begin Here | 1 |
| Audience | 2 |
| Installation Scenarios | 2 |

Begin Here

StorNext provides high performance data sharing and intelligent archiving that is compatible with an industry-leading range of operating systems, server platforms, and storage devices.

StorNext includes two main components:

- StorNext File System (SNFS) is high performance data sharing software that lets applications on multiple operating systems access a single data set.
- StorNext Storage Manager (SNSM) is an intelligent, policy-based data mover that automatically moves and locates data across multiple storage tiers.

In addition, StorNext client software allows client operating systems to mount and work with StorNext and Xsan file systems.

Audience

This document is intended for service professionals and system administrators who install StorNext software. The installer should have strong familiarity with the operating system on which StorNext is installed.

Only service professionals, experienced StorNext users, or individuals who have successfully completed StorNext training should configure StorNext. Attempting to configure StorNext without proper experience or training can affect system performance and may result in data loss.

Installation Scenarios

The purpose of this topic is to help you better use this guide by providing a few of installation scenarios. These scenarios will help you navigate to the procedures that are appropriate to your environment.

Metadata Controller Installations

If you attempt to perform a full backup operation (for example, by using the command **snbackup**) immediately after upgrading, the backup operation will fail.

-
- 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.

Install a RedHat Linux MDC with Windows and Clients


- [StorNext Linux MDC Installation and Configuration](#) provides system requirements, pre-installation information, and instructions for installing and configuring StorNext File System and Storage Manger on a Linux server, called the Metadata Controller (MDC).
- [StorNext Client Installation and Configuration](#) provides system requirements, pre-installation information, and instructions for installing the StorNext client software on Linux and Windows machines.

Install a SNFS-only Windows MDC

- [Windows SNFS-Only MDC Installation and Configuration](#) provides system requirements, pre-installation information, and instructions for installing and configuring StorNext File System on Windows servers.

Install a SNFS-only Linux MDC

- [Linux SNFS-Only MDC Installation and Configuration](#) provides system requirements, pre-installation information, and instructions for installing and configuring RPM-only installations on Linux servers.

 **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.

Additional Resources

- [Configure LAN Clients and Gateway Servers on page 181](#) describes how to configure a distributed LAN client or server for both single-server and high-availability (HA) systems.
- [Uninstall StorNext on page 191](#) describes how to remove StorNext software from a server or client system.
- [Configuration Tasks on page 198](#) describes how to perform file system tasks that do not apply to a typical StorNext setup, including Connecting Xsan Clients to StorNext MDCs.
- [StorNext Files on page 231](#) describes the files used and modified to configure a file system, the cluster-wide control feature, and the file used to configure the cluster-wide security controls.
- [Quality of Service on page 250](#) describes how to configure and use the StorNext File System (SNFS) Quality of Service (QOS) feature.

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. |



Chapter 2: StorNext Linux MDC Installation and Configuration

This chapter contains the following topics:

| | |
|---|----|
| StorNext Linux MDC Installation and Configuration | 6 |
| System Requirements | 7 |
| Getting Ready to Install | 12 |
| Run the Pre-Installation Script | 15 |
| Run the StorNext Installation Script | 19 |
| Post-Installation Steps | 23 |
| Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs | 24 |
| Install Optional StorNext Add-on Software Features and Clients | 30 |
| Convert to a High Availability (HA) System | 30 |
| Configure StorNext Using the StorNext GUI | 37 |
| Linux Device Mapper Multipath Support | 41 |

StorNext Linux MDC Installation and Configuration

This topic describes how to install StorNext File System (SNFS) and StorNext Storage Manager (SNSM) together as a complete package on a Linux Metadata Controller (MDC). Install both SNFS and SNSM for storage systems that require policy-based data movement (for example, systems that use the Replication feature or include tape drives or libraries or to use the StorNext Linux GUI).

Quantum recommends performing a full StorNext install on a Linux MDC, with the StorNext GUI and StorNext Storage Manager. This provides a user-friendly interface, and the ability to use archiving.

If you do not need Storage Manager components or the StorNext GUI, you can perform a [Linux SNFS-Only MDC Installation and Configuration](#).

i Note: The Storage Manager component of StorNext can only be installed on Linux systems. For Windows installation, you can perform a [Windows SNFS-Only MDC Installation and Configuration](#).

Once the StorNext File System is installed, you can configure it using the StorNext Configuration Wizard. To configure StorNext, you will enter license information and create one or more file systems. In addition, on Metadata Controllers (MDCs) running StorNext Storage Manager, you can add storage devices and media, create storage policies, and set up e-mail notifications. To access the Configuration Wizard, use the Web-based GUI to access StorNext running on an MDC.

To ensure successful operation, do the following tasks before installing StorNext:

- Make sure the MDC meets all operating system and hardware requirements (see [System Requirements on the next page](#)).
- Make sure all storage devices are correctly configured and are visible to the MDC (see [Getting Ready to Install on page 71](#))
- **(Optional)** Run the pre-installation script to check for available disk space and view recommended locations for support directories (see [Run the Pre-Installation Script on page 15](#)).

When you are ready to install StorNext File System and Storage Manager on the MDC, run the installation script (see [Run the StorNext Installation Script on page 19](#)).

See the following for additional information.

- [System Requirements on the next page](#)
- [Getting Ready to Install on page 71](#)
- [Run the Pre-Installation Script on page 15](#)
- [Run the StorNext Installation Script on page 19](#)
- [Post-Installation Steps on page 23](#)
- [Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs on page 56](#)
- [Convert to a High Availability \(HA\) System on page 30](#)

- [Configure StorNext Using the StorNext GUI on page 37](#)
- [Linux Device Mapper Multipath Support on page 41](#)

System Requirements

To run StorNext, the Metadata Controller must meet all operating system, hardware, and LAN requirements.

-
- 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 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**.

Operating System Requirements

The operating systems, kernel versions, rpm package dependencies, and hardware platforms supported by StorNext File System are listed in the [StorNext 6 Compatibility Guide](#).

Make sure the system meets the operating system and platform requirements outlined in the guide, and, if necessary, update to a supported release or kernel version before installing StorNext.

-
- i Note:** For Linux clients, you must be disable C-States through kernel boot parameters. See [Modify the grub.conf File to Disable the Intel Sleep State below](#).

Modify the grub.conf File to Disable the Intel Sleep State

Newer versions of Linux include a driver that controls CPU C-states. It is strongly recommended that this driver be adjusted so that deep sleep states are disabled. Making this change may result in increased power consumption but improves stability and performance. This change should be applied to systems running RHEL 6.x and newer, SLES11 and newer, and all supported versions of Debian and Ubuntu:

1. Prior to installation, append the following text to the kernel line in the boot loader configuration file:

```
intel_idle.max_cstate=0 processor.max_cstate=1
```

i Note: Consult your distribution's documentation on how to modify the kernel boot loader.

- For **RHEL 6.x** and **Debian** systems, update:

```
/boot/grub/grub.conf
```

- For **RHEL 7.x** systems, update:

```
/boot/grub2/grub.cfg
```

- For **SLES** systems use `/sbin/yast` to change the boot loader options under **System > Boot Loader**. Save the configuration and verify that `/boot/grub/menu.lst` was updated.
- For **Ubuntu** systems update the `GRUB_CMDLINE_LINUX_DEFAULT` string in `/etc/default/grub` to include the above text, and then execute `update-grub`. Verify that the `/boot/grub/grub.cfg` file was updated.

2. Reboot the system for the change to take effect.

i Note: Disabling C-States in the system BIOS has no effect on Linux.

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.

Install StorNext from a DVD

Installation and upgrade media is shipped on DVDs. Both software and documentation are now stored on a single DVD, eliminating multiple discs for different operating systems. If you plan to install from media, you must have a DVD ROM drive to perform the installation or upgrade.

Refer to the following procedures:

1. [Install the StorNext Client Software on Linux on page 149](#)
2. [Install the StorNext Client Software on Windows on page 162](#)

Download a StorNext Media Kit (Create an ISO Image)

Another way to obtain a StorNext media kit is by downloading the software. This process consists of obtaining a download authorization certificate, creating a download account (or signing on if you already have an account), downloading a software ISO image onto DVD, and then installing the software as described in the following procedures.

Refer to the following procedures:

- [Install the StorNext Client Software on Linux on page 149](#)
- [Install the StorNext Client Software on Windows on page 162](#)

If you have problems or questions with this process, contact your Quantum StorNext support representative.

Install StorNext from a Web Download

A third StorNext delivery method is downloading from the Web. With this method you will download an executable program that will require you to accept the StorNext end user license agreement and will then extract the StorNext installer.

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

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

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

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

Refer to the following procedures:

- [Install the StorNext Client Software on Linux on page 149](#)
- [Install the StorNext Client Software on Windows on page 162](#)

Hardware Requirements

RAM and hard disk space requirements to run the StorNext File System are listed in the [StorNext Compatibility Guide](#).

Make sure the MDC meets the hardware requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

For improved performance and stability, it is strongly recommended that CPU C-States be disabled in the BIOS if they are supported by your system. Refer to system vendor documentation for additional information.

-
- i Note:** Gateway servers may require additional RAM depending on the number of file systems, LAN clients, and NICs used. See [Gateway Server/Client Network and Memory Tuning](#) for Gateway server/client memory tuning guidelines.

Partition Local Hard Disks

StorNext can be installed on any local file system (including the root file system) on the MDC. However, for optimal performance, as well as to aid disaster recovery, follow these recommendations:

- Avoid installing StorNext on the root file system.
- Partition local hard disks so that the MDC has four available local file systems (other than the root file system) located on four separate hard drives. This is so you can install the support directories in four separate locations (for more information, see [Run the Pre-Installation Script on page 15](#)).

i Note: You can run the pre-installation script to help determine the estimated size of and optimal location for StorNext support directories. For more information, see [Run the Pre-Installation Script on page 15](#).

Additional Memory and Disk Requirements for Deduplication and Replication

Additional requirements that must be met in order to use the data deduplication and replication features on the StorNext MDC are listed in the “StorNext MDC Additional RAM and Disk Requirements for Deduplication and Replication” section in the [StorNext Compatibility Guide](#).

Make sure the MDC meets the requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

LAN Requirements

StorNext LAN requirements are listed in the [StorNext 6 Compatibility Guide](#).

i Note: Make sure the MDC meets the network requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

Other Requirements

Additional requirements that must be met for the StorNext system are listed in the [StorNext 6 Compatibility Guide](#).

i Note: Make sure the system meets the additional requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

System Clock Synchronization

Quantum recommends that system clocks are synchronized using [NTP](#) for easier debugging. For HA installations, the system clocks must be synchronized for proper operation.

Nobody User and Group ID Values

The default UID/GID values for the **Nobody** user differ across platforms. Please ensure that your environment is set up with a common value for the **Nobody** user and group to ensure proper file access. Also note the default values for **Nobody** when files are generated by Windows clients, as defined in each file system's configuration.

Windows Recycle Bin Requirements

If you are using the StorNext client software with Windows Server 2003, Windows Server 2008, Windows XP, Windows Vista or Windows 7, disable the **Recycle Bin** in the StorNext file systems mapped on the Windows machine.

You must disable the **Recycle Bin** for the drive on which a StorNext file system is mounted. Also, each occurrence of file system remapping (unmounting/mounting) will require disabling the **Recycle Bin**. For example, if you mount a file system on E: (and disable the **Recycle Bin** for that drive) and then remap the file system to F:, you must then disable the **Recycle Bin** on the F: drive.

StorNext supports mounting file systems to a directory. For Windows Server 2003 and Windows XP you must disable the **Recycle Bin** for the root drive letter of the directory-mounted file system. For example: for **C:\MOUNT\File_System** you would disable the **Recycle Bin** for the **C:** drive.

For Windows Server 2003 or Windows XP

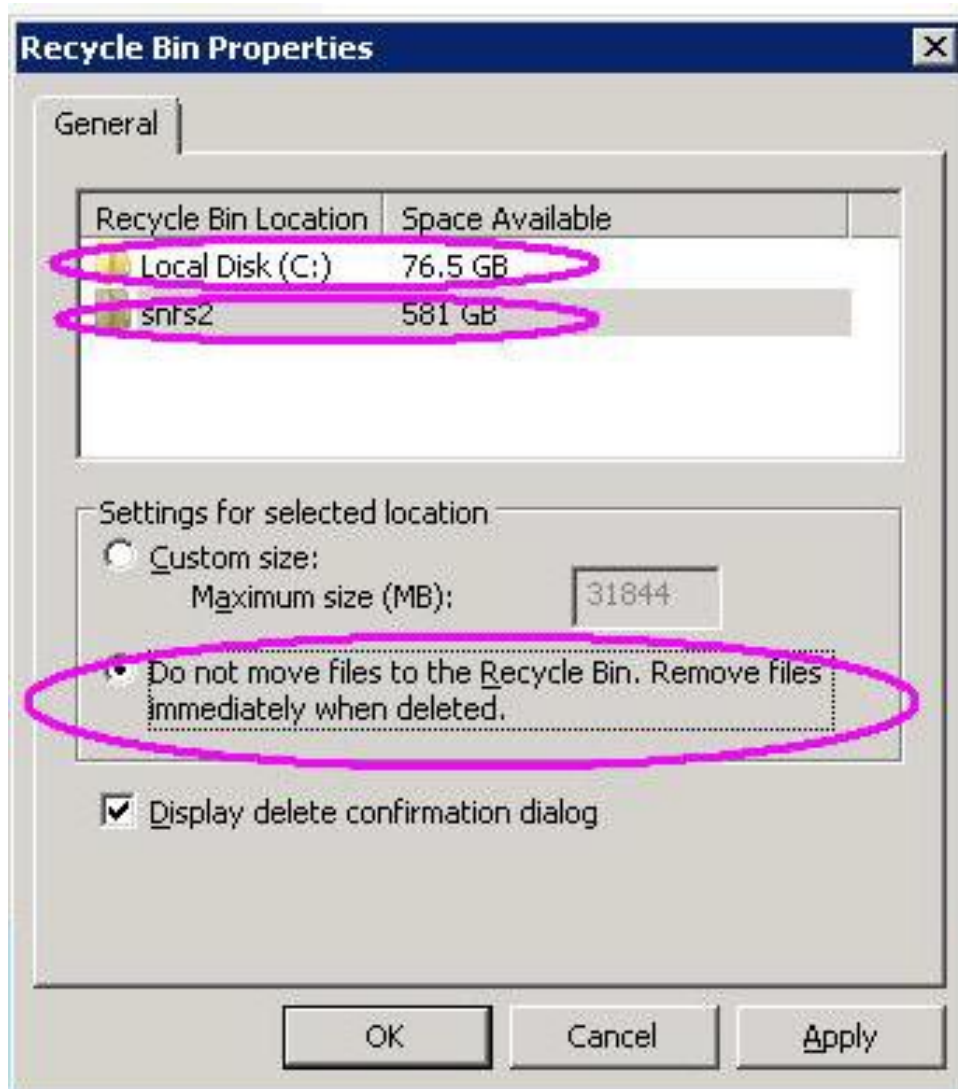
1. On the Windows client machine, right-click the **Recycle Bin** icon on the desktop and then click **Properties**.
2. Click **Global**.
3. Click **Configure drives independently**.
4. Click the **Local Disk** tab that corresponds to the mapped or directory-mounted file system.
5. Click the checkbox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

For Windows Server 2008, Windows Vista and Windows 7 Systems

You must disable the **Recycle Bin** on C: and the File system name.

1. On the Windows client machine, right-click the **Recycle Bin** icon on the desktop and then click **Properties**.
2. Click the **General** tab.
3. Select the mapped drive that corresponds to the StorNext mapped file system. For directory-mounted file systems, select the file system from the list.
4. Choose the option **Do not move files to the Recycle Bin. Remove files immediately when deleted**.
5. Click **Apply**.

- Repeat **Step 3** through **Step 5** for each remaining directory-mounted file system.
- When finished, click **OK**.



Getting Ready to Install

Before installing StorNext File System, complete the following tasks to ensure successful installation:

- Make sure the Metadata Controller (MDC) meets all operating system and hardware requirements (see [File System Requirements on page 69](#)).
- Correctly configure all storage devices and verify they are visible to the MDC (see [Configure a Storage Device on the next page](#)).

- If you plan to use LUNs larger than 1 TB, decide on a label type and install any necessary operating system patches (see [Plan for LUNs Larger than 1 TB below](#)).
- Review and install any package dependencies (see [Package Dependencies on the next page](#)).

Configure a Storage Device

Before installing StorNext File System, make sure that all LUNs are visible to the MDC.

i Note: A LUN, or logical unit number, is a logical device that corresponds to one or more disks, drives, or storage devices.

If there are any connection issues, resolve them before installing StorNext. For assistance in configuring storage devices, refer to the documentation that came with the storage device, or contact the manufacturer.

i Note: LUNs that you plan to use in the same stripe group must be the same size. Consider this when configuring storage devices. (For more information about stripe groups, see the *StorNext User's Guide*.)

! Caution: StorNext does not support the connection of multiple devices through fibre channel hubs. Multiple devices must be connected using fibre channel switches.

! Caution: Quantum recommends that only the storage devices required for an environment are visible to each system. Unnecessary disk visibility will affect system start-up performance and increase the risk of inadvertent reuse of resources.

! Caution: Additionally, it is recommended that sites never perform operating (OS) installations while hardware is attached to a StorNext storage network. Numerous instances of OS installations accidentally overwriting StorNext volumes have been observed in the past.

Plan for LUNs Larger than 1 TB

StorNext supports LUNs greater than 1 TB in size if they are allowed by the operating system. To enable support for 1 TB or larger LUNs, all StorNext LUNs must be correctly labeled according to the requirements of the operating system running on the MDC as well as the operating system running on all connected clients. Disk LUNs can be labeled as EFI. Before installing StorNext, decide what label type will be used.

Required disk LUN label settings based on operating system and LUN size are presented in [Table 1 on the next page](#). Before installing StorNext, install any required operating system patches or updates (for both MDC and client) as indicated in the notes for [Table 1 on the next page](#).

i Note: After labeling a disk LUN, you must reboot systems running Solaris before they can access the disk LUN.

Table 1: Required Disk LUN Labels

| Operating System | Label for LUNs < 1 TB | Label for LUNs 1–2 TB | Label for LUNs > 2 TB |
|--|-----------------------|-----------------------|-----------------------|
| Windows ^a | EFI | EFI | EFI |
| Linux 2.6 Kernel | EFI | EFI | EFI |
| Solaris 10 ^b | EFI | EFI | EFI |
| HP-UX ^c | EFI | EFI | EFI |
| IBM AIX ^d | EFI | EFI | EFI |
| SGI IRIX ^e | Not supported | Not supported | Not supported |
| Apple Xsan 2.1.1 (macOS X 10.5 Leopard) | EFI | EFI | EFI |
| Apple Xsan 2.2 (macOS X 10.6 Snow Leopard) | EFI | EFI | EFI |

- a. For Windows XP 32-bit systems only, support for LUNs larger than 2 TB is limited to the first 2 TB. All other supported Windows platforms (including Windows XP 64-bit) support LUNs greater than 2 TB.
- b. Support for 2 TB or larger LUNs requires Solaris 10 Update 2 (06/06).
- c. Supports only the first 2 TB of LUNs larger than 2 TB.
- d. Supports only LUNs of 2.2 TB or less.
- e. IRIX is not supported starting with StorNext 4.0. However, StorNext 3.5 IRIX clients can be used with StorNext MDCs.

i Note: For best results, Quantum recommends using persistent binding for disk LUNs. For more information, contact the vendor of your HBA (host bus adapter).

Package Dependencies

- For management servers running Red Hat Enterprise Linux version 6 or 7, before installing SNFS and SNSM you must first install the kernel header files (shipped as the kernel-devel or kernel-devel-smp RPM, depending on your Linux distribution).
- For servers running SUSE Linux Enterprise Server, you must install the first kernel source code (shipped as the kernel-source RPM). StorNext will not operate correctly if these packages are not installed. You can install the kernel header files or kernel source RPMs by using the installation disks for your operating system.
- In addition, for MDC installs on Linux the PERL package is required.
- For customer supplied metadata-controllers and hardware, install the fuse-libs package on the system before installing the stornext_full or snfs_client_plus_mover packages.

Verify Hostname Length

The maximum hostname length for a StorNext server is limited to 255 characters. Before you begin the installation, verify that the destination hostname is not longer than 255 characters.

Verify You Have Enough Space to Extract the Installation Software

The installation software files are extracted into the **/tmp/stornext** directory by default. The software, when extracted, requires approximately 40 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the **-d** option.

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

```
# ./<snfs_full.bin> -d <dir>
```

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

Run 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.

Support Directories

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

The StorNext support directories are described below.

| Support Directory* | Description |
|---|--|
| Database /usr/adic/mysql/db | Records information about where and how data files are stored. |
| Journal /usr/adic/mysql/journal | Records changes made to the database. |

| Support Directory* | Description |
|--|--|
| 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.

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?
- Which 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 the Pre-Installation Script

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. When a CD-ROM is mounted on Linux systems, it may have noexec (non-executable) permissions. If you encounter this condition, the workaround to this Linux security issue (not StorNext) is to remount the CD-ROM with an executable option, as follows:
 - a. At command prompt, enter: **mount -t iso9660**. The system responds with:

```
/dev/cdrom on /mnt/cdrom type iso9660 (ro,noexec,nosuid,nodev)
```

- b. Use this information to mount the cdrom:

```
mount -o remount,exec /mnt/cdrom
```

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 the following:

```
./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 Pre-Installation Script 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.

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

| 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.

i Note: Do not change the location of support directories manually. Instead, use the installation script to specify the location for support directories. For more information, see [Run the StorNext Installation Script below](#).

Run the StorNext Installation Script

When you are ready to install StorNext, use the StorNext installation script to install StorNext File System and Storage Manager on the Metadata Controller. The installation script also installs the client software on the MDC.

i Note: The installation script runs dependency checks to ensure the server has all the required RPM packages pre-installed. Refer to the **StorNext Compatibility Guide** for a list of required RPM packages for the operating system you are using.

Upgrade from a StorNext File System-Only Configuration

If you have StorNext File System only and are adding StorNext Storage Manager, some components remain from the File System installation and the new installation could fail. To prevent installation failure, you should first remove the old components.

Remove StorNext Components

The following procedure describes how to remove the old components before starting the new StorNext installation. Follow the steps below **only** if you are upgrading from a StorNext File System-only configuration to StorNext File System and Storage Manager.

1. Log on to the MDC as **root**.
2. Save **/etc/fstab** and **/usr/cvfs/config**
3. Remove the existing StorNext File System RPMs.
4. Install StorNext as described in [Launch the StorNext Installation Script below](#).
5. Restore **/etc/fstab** and **/usr/cvfs/config**

Launch the StorNext Installation Script

The StorNext installation script lets you choose installation configuration options and install StorNext. To launch the script, you will need to use the correct file on the StorNext installation DVD for your operating

system.

i Note: localhost cannot map to multiple IP addresses. Before running the StorNext installation script, make sure localhost in the /etc/hosts file is mapped to a single IPv4 address, and not an IPv6 address.

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 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**

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. For new StorNext installations, at the command line, type:

```
./install.stornext
```

Running the installation script displays:

```
Stornext Install MAIN MENU  
1) Installation Configuration
```

```
2) Install stornext
3) Show Status
4) Quit
Enter your choice <2>:
```

The `install.stornext` by default uses `snPreInstall` to automatically detect the available free space on the system for installing StorNext components. Components favor the location `/usr/adic`, however if there are file systems available that have more space, it will choose those file systems for installing components, creating symbolic links in `/usr/adic` to the installation location of the installed component as needed. Newly installed components may be changed interactively at install time (see [Change Installation Configuration Options below](#)). You may also choose to override the locations chosen by `snPreInstall` by passing in the `-installdir` and `-dbjrn1` to `install.stornext`. To specify the location to install StorNext components, execute the command:

```
./install.stornext -installdir <dir>
```

This causes `install.stornext` to install components at `<dir>/adic`, and will create symbolic links in `/usr/adic/<component>` pointing to `<dir>/adic/<component>`.

i Note: Once elements for a component have been installed the location of the component may not be changed by `install.stornext`.

8. Type the number corresponding to a menu option and press **Enter**. Installation menu options are described below.

| Menu Option | Description |
|-------------|---|
| 1 | Change installation options, including the location of application directories and the default media type (see Change Installation Configuration Options below). |
| 2 | Install StorNext using the installation options specified on the Configuration Menu (see Perform an Installation Operation on page 23). |
| 3 | View the status of an installation operation that is currently in progress. |
| 4 | Exit the StorNext installation script. |

Change Installation Configuration Options

Use the **Configuration Menu** to change StorNext installation options. You can specify the location of application and support directories and change the default media type for storage devices.

On the Main Menu, type **1** and press **<Enter>**. The **Configuration Menu** appears.


```
Stornext 5.0.0(38610) Configuration Menu  
[U.S. Pat. No. 5,990,810]
```

| <u>VALUE</u> | <u>DESCRIPTION</u> |
|-------------------------|---|
| 1) /usr/adic/perl | Perl installation directory |
| 2) /usr/adic/docs | Docs installation directory |
| 3) /adic/mysql | Mysql installation directory |
| 4) /usr/adic/database | Dbm installation directory |
| 5) /usr/adic/SRVCLLOG | Srvcllog installation directory |
| 6) /usr/adic/PSE | Pse installation directory |
| 7) /usr/adic/MSM | Media Manager installation directory |
| 8) /usr/adic/TSM | Tertiary Manager installation directory |
| 9) /usr/adic/java | Java5 installation directory |
| 10) /usr/adic/tomcat | Tomcat installation directory |
| 11) /usr/adic/gui | Sngui installation directory |
| /usr/cvfs | Disk Manager installation directory |
| 13) /adic/database_meta | Metadumps directory |
| 14) /adic/mapping_dir | TSM mapping directory |
| 15) /adic/mysql_jnl | MySQL journal directory |
| 16) LTO | Default media type |

Enter number to modify or hit <RETURN> to continue:

Change the Application Directory Location

When you first enter the **Configuration Menu**, it displays recommended locations for StorNext application directories based on the available file systems detected by the installation script. To accept the recommended directory locations and return to the **Main Menu**, press **<Enter>**.

Alternately, customize the installation by specifying the file systems where application and support files are stored.

1. On the **Configuration Menu**, type the number that corresponds to the directory location you want to customize (**1–16**), and then press **<Enter>**. A list of available local file systems is shown.
2. To change the file system where the directory is located, type the name of the file system and press **<Enter>**.
3. To confirm the change, type **yes** and press **<Enter>**. To customize additional directory locations, repeat steps 1–3.
4. When you are done customizing the installation, press **<Enter>** to return to the Main Menu.

i Note: To help determine the optimal location for StorNext support directories, use the StorNext pre-installation script. For more information, see [Run the Pre-Installation Script on page 15](#).

Change the Default Media Type

If you do not specify a different media type, the StorNext installation script selects LTO as the default media type for storage devices. If storage devices in your system use a different media type, change the default media type before installing StorNext.

1. On the Configuration Menu, type **16** and press **<Enter>**. A list of valid default media types is shown. The valid media types are:

- LTO
 - T10K
 - 3592
 - Various Object Storage
 - SDISK
2. Type the correct default media type for storage devices in your system and press **<Enter>**.
 3. To confirm the change, type **yes** and press **<Enter>**.
 4. When you are done customizing the installation, press **<Enter>** to return to the Main Menu.

Perform an Installation Operation

To install StorNext using the options specified on the **Configuration Menu**, perform an installation operation.

i Note: Storage Manager requires that the **rpcbind** service is running. Make sure that the **rpcbind** service is installed and running before installing StorNext or starting Storage Manager. Consult your distribution's documentation for more information on how to install and start the **rpcbind** service.

1. On the **Main Menu**, type **2** and press **<Enter>**. The StorNext installation script runs. The script displays the current progress of the installation.
2. When the installation is complete, press **<Enter>** to return to the **Main Menu**.
3. To exit the installation script, type **4** and press **<Enter>**.

i Note: The script displays the URL at which the MDC can be accessed. Make a note of this information.

4. After running the installer, source the profile on the MDC by running the command:

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

You can now access StorNext File System and Storage Manager, and run the Configuration Wizard. For more information, see [Configure StorNext Using the StorNext GUI on page 37](#).

Post-Installation Steps

After StorNext is installed, you are ready to configure your system. Quantum recommends using the StorNext Configuration Wizard, which walks you through the steps of configuring your system for the first time. For more information about logging in to StorNext and using the Configuration Wizard, see [Configure StorNext Using the StorNext GUI on page 37](#).

Log In to StorNext

For instructions, see [Access the StorNext GUI on page 37](#).

SSL Certificate Exception

When you log into StorNext for the first time, you might see a message warning you about a security certificate.

For a permanent workaround to this issue, see the Quantum Knowledge Base article titled *StorNext GUI: Installing a SSL Certificate from a Certificate Authority* available online at <https://qsupport.quantum.com/kb/node/4489>.

See also [Import a Secure Sockets Layer \(SSL\) Certificate on page 225](#).

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

Start the StorNext File System and StorNext Storage Manager

When you log in to StorNext for the first time, the **System Control** page appears. On this screen you can start StorNext File System and Storage Manager by clicking **Start**.

Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs

Non-Expiring vs. Subscription License Models

StorNext 7 (or later) with Subscription Licenses:

If your system is running StorNext 7 (or later), contact your [Quantum Sales](#) representative to determine which non-expiring or subscription licenses are best suitable for your environment. See [StorNext Licenses](#) for more information.

StorNext 7 (or later) and StorNext 6.x or earlier with Non-Expiring Licenses:

StorNext and some StorNext non-expiring license features are pre-installed on your system, and others must be purchased and then enabled by installing a license.dat file on the system. If your system is using a 30-day auto-generated license, you should upload a license file containing permanent licenses as soon as it is convenient, so that the **Maintenance** expiration date and the (**FlexTier™**) subscription date for the license is set correctly on the system.

- i Note:** For non-expiring licenses, you can determine the current license limits for many StorNext features, such as SAN clients, LAN clients, DDM, and Storage Manager licenses, on the **Configuration > Licenses** page of the StorNext GUI and on the Quantum **License Key Management** page available online at <https://www.quantum.com/licensekeys>.

Obtain StorNext Licenses

Prerequisites

In order to receive a license file (`license.dat`) to enable StorNext features or capacity, you will need the following:

- The **System Serial Number** or the **StorNext Software Serial Number**.

System serial numbers and StorNext software serial numbers are alphanumeric. The table below provides examples of serial numbers.

| Type of Serial Number | Example(s) of Serial Number(s) |
|--|--|
| System Serial Number | AV1607CKH12345 CX1306CKB12345 EW1450CKD12345 |
| StorNext Software Serial Number | SN12345 |

- The `cvfsid` string for each server node. See [How to Generate a cvfsid for Linux and Unix Configurations and Quantum Appliances on page 27](#), [How to Generate a cvfsid for Windows Configurations on page 27](#), or [How to Generate a cvfsid for Xsan Configurations on page 28](#), depending on your configuration.

Example: `ECF4BCEECC0E linux 0 xcellis13`

The table below provides information regarding which configurations require one `cvfsid` string or two `cvfsid` strings.

| Configurations Requiring One cvfsid String | Configurations Requiring Two cvfsid Strings |
|--|--|
| <ul style="list-style-type: none"> • Customer-configured MDC without an HA license • Xcellis Workflow Extender • G300 • Xcellis Workflow Director (single-node systems only) • Artico (single-node systems only) • Xcellis Foundation (single-node systems only) • aiWARE for Xcellis Standard and High Performance | <ul style="list-style-type: none"> • Customer-configured MDC with an HA license • Xcellis Workflow Director (dual-node systems only) • Xcellis Foundation (dual-node systems only) • Artico (dual-node systems only) • aiWARE for Xcellis Embedded (DAE) • Pro Solutions • M660 • M440 • M330 |

How to Obtain a New StorNext license.dat File

1. Visit the Quantum **License Key Management** page at <https://www.quantum.com/licensekeys>. If you cannot access the web page or need additional help, contact Quantum Support at <https://www.quantum.com/ServiceandSupport>.
2. Enter your serial number in the **Serial Number** field.
3. Click **Submit**. The license(s) associated with the serial number you entered in **Step 2** are displayed in the **License Details** section.

i Note: If you have not previously entered a `cvfsid` for the serial number, enter the strings in the **CVFSID 1** or **CVFSID 2** fields and then click **Set CVFSID**. See [How to Generate a cvfsid for Linux and Unix Configurations and Quantum Appliances on the next page](#), [How to Generate a cvfsid for Windows Configurations on the next page](#), or [How to Generate a cvfsid for Xsan Configurations on page 28](#), depending on your configuration. The **CVFSIDs** sub-section under **License Details** is updated with the `cvfsid` value(s) for the serial number.

⚠ Caution: Ensure your `cvfsid` value(s) are correct. Once entered and submitted, they cannot be changed.

i Note: At any time after you have submitted your serial number, you may enter another serial number by clicking **Choose Another**.

4. Click **Generate License Key File** once the `cvfsid` value(s) have been entered.

i Note: There may be some cases where the Quantum **License Key Management** site will not be able to generate a license file. In this case, you are re-directed to an alternate web page to continue the process.

5. Save the `license.dat` file to a temporary location on your computer. You will upload this file on the **Licenses** page of the StorNext GUI. To install the licenses in the `license.dat` file, use the [Install a StorNext License from the GUI for non-HA Installations](#) procedure if you have a single-node StorNext MDC, or use the [Install/Update Licenses from the GUI for HA Installations](#) procedure if you have a High Availability (HA) StorNext MDC pair.

How to Generate a `cvfsid` for Linux and Unix Configurations and Quantum Appliances

For Quantum appliances, open an SSH connection to the server

Using the IP address assigned to the server on the StorNext management/public network.

1. Log in to the command line using the following credentials:
 - User name: `stornext`
 - Password: `<stornext user account password>`

i Note: "password" is the default password for the `stornext` user account. If the password has been changed, use the current password. Beginning with StorNext 7.0.2, you must change the default password for the `stornext` user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

2. Enter `sudo rootsh` to gain root user access.
3. Enter the password for the `stornext` user account again.

Generate the `cvfsid`

1. Enter the `cvfsid` command.

The system returns the following information:

`<server_MAC_address> <server_OS> <server_name>`

Example:

```
ECF4BCDECC0E linux 0 xcellis13
```

2. Close the SSH session for the node.
3. **(Optional)** If your system is configured for HA, repeat for the other node.

How to Generate a `cvfsid` for Windows Configurations

Follow the procedure below to generate the `cvfsid` on the command line of the MDC node:

1. On the **Start** menu, click **Programs**, click **StorNext File System**, and then click **License Identifier**.
2. The `cvfsid` string appears in a popup.

How to Generate a cvfsid for Xsan Configurations


Follow the procedure below to generate the cvfsid on the command line of the MDC node:

1. Navigate to `/Library/Filesystems/Xsan/bin`.
2. Execute the command `sudo cvfsid`.


Install/Update StorNext Licenses From the GUI for Dual-node (HA) Systems

StorNext 7 (or later) with Subscription Licenses:

If your system is running StorNext 7 (or later), contact your [Quantum Sales](#) representative to determine which non-expiring or subscription licenses are best suitable for your environment. See [StorNext Licenses](#) for more information.

 **Caution:** At the end of this procedure, exiting Config Mode will restart StorNext services (on both server nodes for dual-server node systems). This will cause a disruption in client operations. Quantum strongly recommends halting I/O operations for connected StorNext clients before beginning the process.

1. If you are logging into StorNext for the first time, read the end-user license agreement (EULA) carefully, and then click **Accept** (you will not see the EULA page if you have previously logged into StorNext).
2. For new installations, the **Configuration > Licenses Entry** screen appears. Skip to step 4.
For operational systems, navigate to **Tools > High Availability > Manage**.
3. **(Optional)** If your **cvfsid** has changed, then click **Enter Config Mode** to halt StorNext client I/O operations. A dialog appears to confirm you want to lock the HA cluster configuration. Click **Yes**.
(See the **High Availability Systems** section in the appendix of the current *StorNext User's Guide* for further information about HA systems.)
4. Select **Configuration > Licenses**.
5. Scroll down to the bottom of the page, and click **Import...** to install and/or update StorNext licenses.
6. Click **Choose File**, select the `license.dat` file provided for your system from the temporary location where you downloaded it, and click **Open**.

 **Caution:** Importing a license file overwrites the existing license file and all existing licenses. Licenses not present in the imported license file are removed from StorNext.

You will know that the licenses have been applied to your system when you see the confirmation “**License file license.dat imported successfully.**” appear at the top of the page. The status indicator for the licensed features is displayed as a green check mark icon.

7. If the system is in config mode, navigate to **Tools > High Availability > Manage** and click **Exit Config Mode**. You will be prompted for confirmation that you want to unlock the cluster. Click **Yes** to exit config mode.

The procedure is complete.

Install StorNext Licenses From the GUI for Single-node (non-HA) Systems

-
- i Note:** Use these steps to install the license keys on a single-node system, or one that has not been configured as HA. For HA systems, see [Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs on page 24](#).

1. If you are logging into StorNext for the first time, read the end-user license agreement (EULA) carefully, and then click **Accept** (you will not see the EULA page if you have previously logged into StorNext).

-
- i Note:** To comply with the California Senate Bill No. 327 (CA SB-327), you must change your password for your **admin** and **service** user accounts on your first log in attempt.

2. For new installations, the **Configuration > Licenses Entry** page appears. Skip to step 5.
3. For operational systems, halt StorNext client I/O operations. Stop the StorNext filesystem by navigating to **Tools > System Control**, and selecting **Stop** in the **File System** section of the page.
4. Select **Configuration > Licenses**.
5. Scroll down to the bottom of the page, and click **Import...** to install and/or update StorNext licenses.
6. Click **Choose File**, select the **license.dat** file provided for your system from the temporary location where you downloaded it, and click **Open**.

-
- ! Caution:** Importing a license file overwrites the existing license file and all existing licenses. Licenses not present in the imported license file are removed from StorNext.

You will know that the licenses have been applied to your system when you see the confirmation “**License file license.dat imported successfully.**” appear at the top of the page. The status indicator for the licensed features is displayed as a green check mark icon.

7. After you install your licenses, restart the StorNext filesystem by navigating to **Tools > System Control**, and selecting **Start** in the **File System** section of the page. StorNext may not recognize some feature licenses until you restart services.

The procedure is complete.

Install a StorNext License on a Windows MDC

For help installing Windows MDC licenses, refer to the StorNext help, which can be found at: **Windows Start menu > All Programs > StorNext File System > StorNext Help**.

Install Optional StorNext Add-on Software Features and Clients

Download client software installers and optional add-on software features from the StorNext GUI on the **Tools > Client Download** page.

Convert to a High Availability (HA) System

i Note: The HA conversion process is required for initial installations.

If you are upgrading from StorNext 4.7.x (or later), and have previously converted your system to HA configuration, the conversion process is unnecessary. Previously-converted systems will not be taken out of HA configuration, so the option to convert to HA is unavailable. However, for StorNext releases prior to StorNext 4.7.x, an HA conversion is necessary after StorNext software upgrades.

This section describes the configuration steps necessary to convert two StorNext MDC servers into a High Availability MDC pair connected to a shared file system. Converting to HA consists of selecting the dedicated unmanaged StorNext file system for use as the controlling shared file system, and then instructing StorNext to convert each MDC server to operate as one MDC server of the HA pair.

The following note and bullet items apply **only** to customer-supplied MDCs.

i Note: The **Convert** menu option will be unavailable (grayed out) on the **Tools** menu if you have not specified a secondary system. If you have not already done so, specify a secondary system by using the Name Servers function. For more information, see [Step 2: Configure a File System Name Server on page 85](#).

The following are considerations you should be aware of concerning the HA conversion process:

- The HA shared file system **MUST** be configured as an unmanaged file system. The file system should not have quotas enabled. Enabling quotas on this file system can interfere with the proper functionality of Storage Manager and the HA infrastructure.
- The conversion process converts one MDC server at a time. The second MDC should be converted as soon as possible after the first server.
- StorNext operating files will be moved to the HaShared file system, and this move cannot easily be reversed.
- Following conversion, the primary MDC is identified by the vIP for Replication/Deduplication.
- Replication/Deduplication policies must be changed to use the vIP:
 - The global policy for each file system must use it as the “Address for Replication and Deduplication”
 - Replication policies must use it as the Target for Replication Schedules
 - If multilink is configured, the vIP address should be used.

Considerations for All MDCs

The following applies to all MDCs used in an HA configuration:

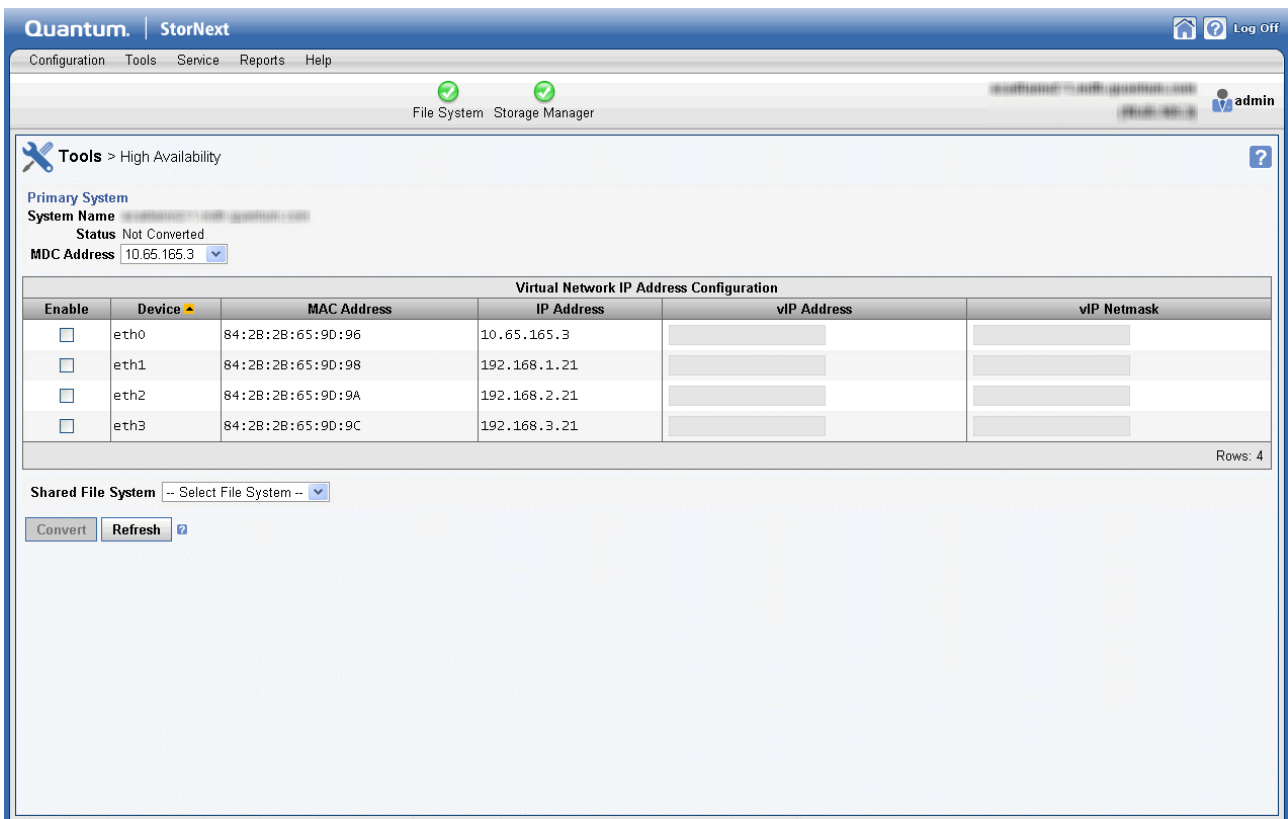
- The UIDs for the **quantumdb** and **tdlm** users and the **adic** group **must** be identical on both MDC server servers of an HA pair. If the UIDs are not identical, the MySQL database will not start (due to file permission errors), which in turn prevents storage manager from fully starting up. Quantum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

HA Conversion Procedure

Follow these steps to configure HA:

1. Using the StorNext GUI, on the **Tools** menu, click **High Availability**, and then click **Convert**. The **High Availability** page appears.

Note: If the system has previously been converted to HA mode, continue to [Step 6](#).



The screenshot shows the StorNext GUI interface for configuring High Availability. The main content area is titled "Tools > High Availability". Below this, there is a "Primary System" section with fields for "System Name", "Status" (set to "Not Converted"), and "MDC Address" (set to "10.65.165.3").

The central part of the interface is a table titled "Virtual Network IP Address Configuration". The table has the following columns: "Enable", "Device", "MAC Address", "IP Address", "vIP Address", and "vIP Netmask".

| Enable | Device | MAC Address | IP Address | vIP Address | vIP Netmask |
|--------------------------|--------|-------------------|--------------|-------------|-------------|
| <input type="checkbox"/> | eth0 | 84:2B:2B:65:9D:96 | 10.65.165.3 | | |
| <input type="checkbox"/> | eth1 | 84:2B:2B:65:9D:98 | 192.168.1.21 | | |
| <input type="checkbox"/> | eth2 | 84:2B:2B:65:9D:9A | 192.168.2.21 | | |
| <input type="checkbox"/> | eth3 | 84:2B:2B:65:9D:9C | 192.168.3.21 | | |

Below the table, there is a "Shared File System" dropdown menu with the text "-- Select File System --". At the bottom of the configuration area, there are two buttons: "Convert" and "Refresh".

2. On the **Shared File System** option, select the shared file system you want to use for the HA pair.

Caution: Once you convert a file system to HA you cannot undo the operation, so choose carefully when selecting the file system.

i Note: The file system should not have quotas enabled. Enabling quotas on the file system can interfere with the proper functionality of Storage Manager and the HA infrastructure.

! Caution: On the **Tools > High Availability** page, it is critical that you assign the Metadata network IP address in the MDC **Address** field and not the Management network IP address.

! Caution: Double-check that the **Shared File System** option is set to the shared file system on the metadata array for the metadata appliance/HA pair and **NO OTHER** shared file system on the list before applying changes. The file system selection on this page may only be done once, so make sure you select the shared file system on the metadata array.

i Note: A virtual IP (VIP) may also be set at this time. Use the IP addresses assigned to the port used for the Management network on each MDC server for the VIP values. Do not use the IP addresses assigned for the ports used for the Metadata network.

3. For the MDC Address, select one IP address to be placed in the `ha_peer` file for use in administrative operations between the MDC servers in the HA pair.
4. **For non-Lattus systems ONLY:** If your HA cluster also runs the blockpool, select **Enable** and then enter the virtual IP address and virtual netmask. Ask your network administrator for the vIP address and netmask.
5. Click **Convert** to convert the primary MDC server to HA.

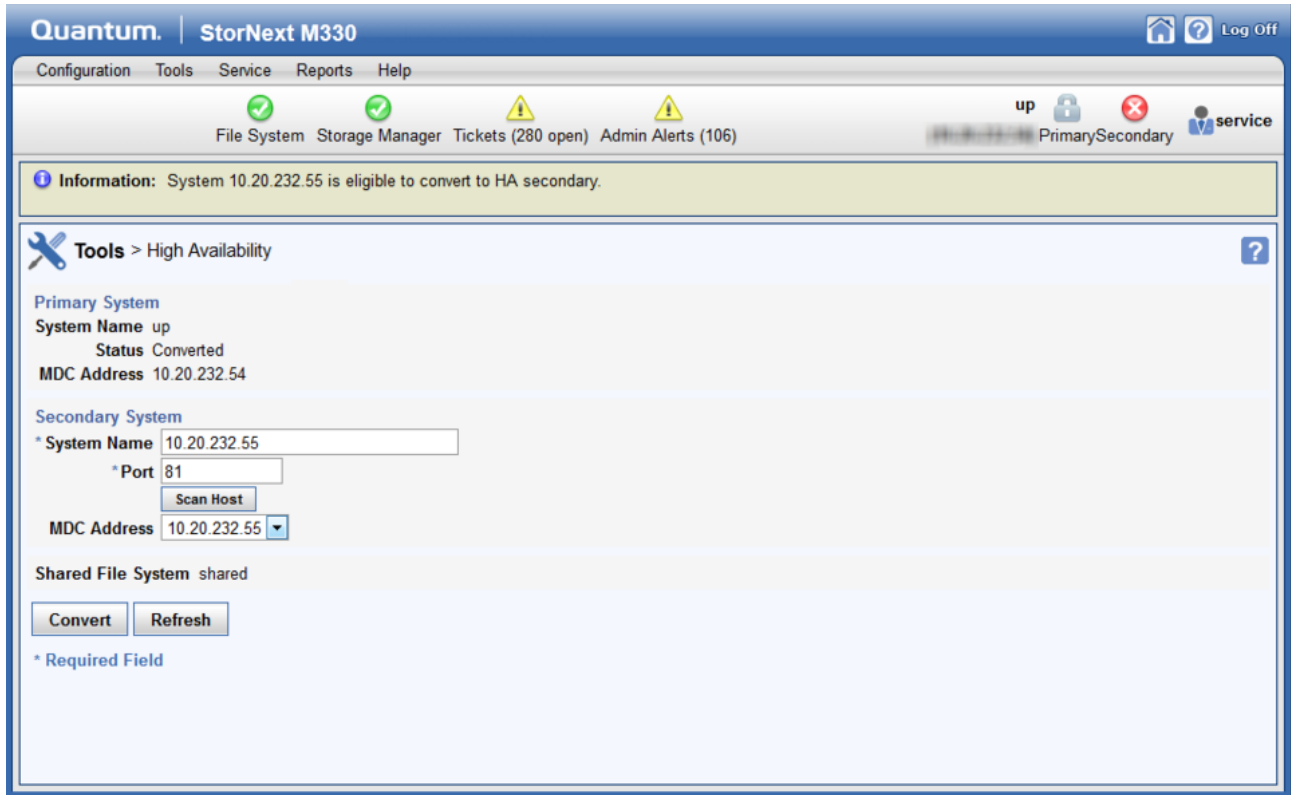
i Note: Once the primary MDC server has been converted, the **Status** displays **Converted**.

6. Enter the IP address of the secondary MDC server in the **System Name** field along with a port number on the same LAN as the primary MDC server.
7. (If the system is configured with a DNS entry, you may choose to use that name in place of the IP address in this field.) Click **Scan Host**. The system should resolve the secondary MDC server - the MDC Address will auto-fill with a value. If you do not already have licenses for the secondary system in the license file, you will be required to switch to the license page to import them before continuing. The information comes from the **license.dat** file

i Note: .Until you have performed the scan, you cannot import the license file for the secondary system using the StorNext import function. After you have performed the scan you can import licenses for the secondary. Following the conversion to HA, the license file contains licenses and is present on the servers.

8. Click **Convert** to convert the secondary MDC server.

i Note: Both MDC servers will reboot, which can take 30 minutes or more per server to complete.



9. Storage Manager may need to be started following the HA conversion if the system was in config mode at the time that HA conversion was initiated. To restart the Storage Manager components, click **Start** button in the Storage Manager panel of the **Tools > System Control** page.

GUI Feedback During HA Conversion

There are some indications within the GUI that the system is being upgraded. Here are some notes about this visual feedback:

- After the StorNext upgrade has completed, and the HA configuration has been done, the GUI for the secondary MDC server provides a message stating it is not the primary MDC server and a link to launch the primary MDC server.

Caution: Do not log in to the GUI of the secondary MDC server at any point during the upgrade/HA conversion process. System configuration and licensing for the system could be compromised.

- When you are able to log into the primary system/server, after accepting the EULA, the system will automatically display the **Tools > System Control** page. Click **Start** to restart the Storage Manager components.
- Wait until the system icons for both MDC servers of the system as well as **File System** and **Storage Manager** are green, which indicates normal operation.

Post-Conversion Steps

If you are using the DDM feature, perform the following:

1. If you use the secondary MDC server as a DDM mover, make sure the file systems are mounted.
2. Edit the **fs_sysparm** or the **fs_sysparm_override** files to use your preferred DDM mode: **All** or **Threshold**. Use the **adic_control restart TSM** command to put this change into effect.

Verify HA Failover

After an initial conversion of your system to HA, you may want to verify HA failover. To initiate the failover of a system that has just been converted to HA:


1. Open an SSH connection to the MDC server operating as the **primary**.
2. Log in to the command line of the **primary** MDC server, using the following credentials:
User: **stornext**
Password: *(You (the customer) should have the password for the StorNext account)*
3. Type **sudo rootsh** to gain root user access. Enter the password again.
4. Confirm that the MDC server is operating as the primary by entering the following at the command prompt:

```
snhamgr -m status
```

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

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

6. Initiate an HA failover on the **primary** MDC server to the **secondary** MDC server.

 **Caution:** Only use the **snhamgr force smith** or **cvadmin -e hamon smith** commands to initiate a system failover the first time the system is converted to HA and should only be done when the system is operating in a non-production environment. These commands cause a HARD RESET of the MDC server, and will impact StorNext operations. For initiating a failover any other time, see [Initiate a Graceful System Failover on the next page](#) for the graceful HA failover procedure.

- a. Issue the following command to stop StorNext services on the primary MDC server:

```
adic_control stop
```

- b. Then, issue the following command to initiate the failover:

```
snhamgr force smith
```

i Note: Restarting StorNext more than once every three minutes can trigger a startup delay of three minutes. The delay slows down looping HA resets so that an administrator can correct any configuration problems that are causing the resets.

7. Wait 3 minutes until the MDC server previously operating as the **secondary** becomes the **primary**.
8. Open an SSH connection to the MDC server now operating as the **primary**.
9. Log in to the command line of the **primary** MDC server.
10. Confirm that the MDC server is operating as the **primary** by entering the following at the command prompt:

```
snhamgr -m status
```

11. Verify the output is:

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

12. Repeat if desired to fail over to the original system operating as the **primary**.

i Note: When the command **force smith** is used on an MDC server, the system will reboot, and it may take a significant amount of time for the MDC server to come back online, so plan for this additional delay. Wait until the MDC server reboot has completed before initiating another fail over to the MDC server originally operating as the **primary**.

13. Once the failover has completed, restart SNFS services on any clients that were stopped earlier.
14. Mount the SNFS file systems on each client machine, if needed.
15. Verify that all clients have full access.
16. Verify access to all file systems and move files to/from disk and tape.

Initiate a Graceful System Failover

During the upgrade of the server originally operating as the primary, the server originally operating as the secondary server becomes the primary. After the database updates are complete and the upgrade of the original primary completes, a failover of the server currently operating as the primary should be triggered so it once again is operating as the secondary, and to reinstate the original primary server to operate as primary again. During StorNext 5.x or later upgrades, the Metadata Appliance fails over once. Because of this single failover, the MDC server originally operating as primary is set as the secondary after the

upgrade. This procedure provides the steps necessary to set the MDC servers back to the system state prior to the upgrade for this scenario.

To initiate the graceful failover of an HA pair:

1. Open an SSH connection to the MDC server operating as the **primary**.
2. Log in to the command line of the **primary** MDC server, using the following credentials:

User: stornext

Password: *(You (the customer) should have the password for the StorNext account)*

3. Type **sudo rootsh** to gain root user access.
4. Enter the password a second time.
5. Confirm that the MDC server is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

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

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

7. On the MDC server operating as the **primary**, initiate an HA failover to the MDC server operating as the **secondary**:

```
service cvfs stop
```

8. Wait until the **secondary** MDC server becomes the **primary**, and leave your SSH connection to this server open. The time may vary.
9. Open an SSH connection to the MDC server now operating as the **primary**.
10. Log in to the command line of the **primary** MDC server, using the following credentials:

User: stornext

Password: *(You (the customer) should have the password for the StorNext account)*

11. Type **sudo rootsh** to gain root user access.
12. Enter the password a second time.
13. Confirm that the MDC server is operating as the **primary** by entering the following at the command line:

```
snhamgr -m status
```

14. Verify the output is:

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

15. From the SSH connection to the MDC server now operating as the **secondary**, enter the following:

```
service cvfs start
```

16. Confirm that the MDC server is operating as the **secondary** by entering the following at the command line:

```
snhamgr -m status
```

17. Verify the output is:

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

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

Configure StorNext Using the StorNext GUI

The StorNext GUI provides access to StorNext features and displays system status information. Use the StorNext GUI to create and work with file systems. In addition, on Metadata Controllers running Storage Manager, you can use the StorNext GUI to configure storage devices and media, and to set up storage policies.

Access the StorNext GUI

The StorNext GUI is browser-based and can be remotely accessed from any machine with access to the StorNext server. Perform the following procedure to access the StorNext GUI.

1. Open a Web browser.


i Note: For a list of browsers that have been tested to work with StorNext, refer to the *StorNext Compatibility Guide* posted online at <https://www.quantum.com/snsdocs>.

2. In the browser's **Address** field, type the full address of the machine, and then press **Enter**. For example:

```
https://<machine name>:<port number>
```


Use the name of the machine you copied when you installed the StorNext software. After you enter the machine name and port number, the **Login** page appears.

i Note: Typically, the port number is 443. If port 443 is in use, use the next unused port number (for example, 444, 445, and so on).



3. In the **User ID** field, type **admin**.
4. In the **Password** field, type **password**.
5. Click **Login**. The initial StorNext **System Control** page appears.

Note: When you log into StorNext for the first time, you might see a message warning you about a security certificate. Either click the option that allows you to “continue to this site,” or, for a temporary solution, create a certificate exception that will allow you to log into StorNext without seeing the warning message during subsequent logins attempts.



6. On this screen you can determine if the StorNext File System and Storage Manager components are currently started. If not, click **Start** for each component to start them.

Note: If you are logging in for the first time, it is likely that Storage Manager will be stopped and require starting.

7. Click the home (house) icon in the upper right corner to go to the **StorNext Home Page**.

Use the StorNext Configuration Wizard

The first time you log onto the StorNext GUI, the Configuration Wizard appears. The wizard guides you step-by-step through the process of configuring StorNext.

Note: Using the Configuration Wizard to configure StorNext is **optional**. All configuration tasks can be accessed at any time using the StorNext Setup menu.

Chapter 2: StorNext Linux MDC Installation and Configuration

Configure StorNext Using the StorNext GUI

StorNext

Configuration Tools Service Reports Help

File System Storage Manager Tickets (1 open) Primary Secondary admin

Configuration Wizard

Configuration -> Welcome

Welcome to the StorNext Configuration Wizard. Below you will see the disks and libraries accessible to the application. Click Next to configure StorNext.

| Serial Number | Type | Label | Size | Status | Used | File System |
|------------------|-------------------|-----------------------|------------|--------|------|-------------|
| /dev/sda | GENERIC_67090399 | | 31.99 GiB | up | true | reserved |
| 3030303030303030 | GENERIC_134201083 | CvfsDisk_managed_data | 511.94 GiB | up | true | managedfs |
| 3030303030303031 | GENERIC_8371963 | CvfsDisk_managed_meta | 31.94 GiB | up | true | managedfs |
| 3030303030303032 | GENERIC_134199263 | CvfsDisk_shared | 63.99 GiB | up | true | shared |

Rows: 4

| Serial Number | Product ID | Device Type | Device Path |
|------------------|------------|-------------|-------------|
| (Table is empty) | | | |

Rows: 0

Next

Refresh

The Configuration Wizard consists of multiple steps. The wizard lets you navigate between steps and tracks your progress as you complete each step.

- To go to the next step, click **Next**.
- To return to a previous step, click **Back** (or click the step in the list).

These are the steps you'll complete in the Configuration Wizard:

- **Welcome:** View disks and libraries currently available for StorNext usage.
- **Licenses:** Enter StorNext License Information.
- **System:** Set up the networking configuration for StorNext Metadata Appliances (M660, M440 and M330). This step only appears on the Configuration Wizard of StorNext Metadata Appliances.
- **Name Servers:** Specify and order the machines acting as StorNext name servers.
- **File Systems:** Add a StorNext file system.
- **Storage Destinations:** Add a library, storage disks, and other storage destinations.
- **Storage Policies:** Add a Storage Manager or replication storage policy.
- **Email Server:** Specify an email server to handle StorNext notifications.
- **Email Notification:** Add e-mail notification recipients.

- **Done:** Signify that you are finished using the Configuration Wizard. You can also convert to a high availability (HA) system.

To see more detailed instructions for completing Configuration Wizard tasks, see [Configuration Wizard](#).

Linux Device Mapper Multipath Support

StorNext supports the Linux Device Mapper (DM) Multipath driver. This driver provides redundancy and improved I/O performance by taking advantage of multiple paths to storage.

Considerations for Using Linux DM Multipath with StorNext

- The StorNext File System does not currently support disks managed through the Linux Device Mapper tools other than those which are managed through **multipathd**. Disk paths showing up under the `/dev/mapper` path that are not **multipathd** devices will not be usable by StorNext due to Device Mapper taking exclusive control of the disk devices. See the Linux **dmsetup** documentation for details.

Device Mapper will release exclusive control of a disk device if all of the Device Mapper partition entries are removed. Some customers have successfully used Device Mapper devices by using **dmsetup** to remove the partition entries for a particular device before attempting to mount StorNext File Systems. See the Linux **dmsetup** documentation for details.

Note: Device Mapper control on specific devices may need to be removed after each reboot or after a device is first presented to a system.

- Not all RAID configurations work with the DM Multipath Driver. Check with your storage vendor for compatibility.
- For detailed instructions on installing and configuring the DM Multipath Driver, refer to the Red Hat, SUSE, Debian, or Ubuntu documentation provided with your version of Linux.
- For StorNext to use Linux Device Mapper Multipath devices, you must make two changes to the `/etc/multipath.conf` file.
 1. Set `user_friendly_names` to `yes`.
 2. Add the following entry to the `multipath.conf` blacklist section to prevent **multipathd** from managing the SNFS volume devices:

```
devnode "cvfsctl*"
```

In addition, the use of the `alias` attribute in a `multipath` subsection of the `multipath.conf` file is not currently supported for devices used by StorNext. Its use can lead to mount failures.

- Using the `cvpaths` file and `udev` rules configuration files is typically unnecessary with Linux Device Mapper with StorNext.

- Current versions of the Linux DM Multipath driver assign a default value of 1000 for `rr_min_io` which is too high for most configurations having multiple active paths to storage. Using a smaller value such as 32 will typically result in significantly improved performance. Refer to the Red Hat, SUSE, Debian, or Ubuntu documentation provided with your version of Linux for details on how to apply this setting.

i Note: Experimentation may be required to determine the optimal value.

- **(On SUSE Linux Systems)** In order to use Linux Device Mapper Multipath with Stornext, `/etc/multipath.conf` must be used because SUSE Linux does not install a `multipath.conf`, and Novell recommends against using it. Although SUSE Linux does not install a `multipath.conf` file by default, an example file, located at:

```
/usr/share/doc/packages/multipath-tools/ multipath.conf.synthetic
```

can be copied to:

```
/etc/multipath.conf
```

- **(On Ubuntu Linux Systems)** There is no `multipath.conf` file installed by default. However, an example file, located at:

```
/usr/share/doc/multipath-tools/examples/multipath.conf.synthetic
```

can be copied to:

```
/etc/multipath.conf
```

- **(On Debian Linux Systems)** There is no `multipath.conf` file installed by default and no example file is provided. Therefore, the `multipath.conf` file must be hand-crafted. Refer to your Debian documentation for additional information.
- **(On Red Hat Linux Systems)** Red Hat does install a `multipath.conf` file. By default, Red Hat `multipath.conf` file blacklists all multipath-capable targets. This means the following must be commented out:

```
blacklist { devnode "*" }
```

i Note: After installing and/or updating the `/etc/multipath.conf` file, **multipath** services must be restarted for changes to take effect. For details, refer the documentation for your version of Linux.

Third-Party Multipath Support

If you are using a third-party multipath driver, you may need to use a **cvpaths** file. The **cvpaths** file directs StorNext to the multipath devices for the entire lun.

By default, below are the devices StorNext uses:

- **Linux**

```
/dev/sd*
```

- **AIX**

```
/dev/rhdisk*
```

- **Solaris**

```
/dev/rdisk*
```

For example, the EMC Powerpath driver creates devices named `/dev/emcpower*` on Linux. If you do not tell StorNext to use the `/dev/emcpower*` devices, it will instead use the `/dev/sd*` devices. In the example below, a **cvpaths** file must be created to specify the following:

```
wildcard=/dev/emcpower*[^1-9]
```

Other third-party multipath devices may use other path names. If you run the `cvlabel -l` command, the output will show which devices are being used by StorNext.

For additional information on configuring the **cvpaths** file, refer to the **cvpaths** man page from the command line, or the **cvpaths** entry in the **MAN Pages Reference Guide**.



Chapter 3: Linux SNFS-Only MDC Installation and Configuration

This chapter contains the following topics:

| | |
|---|----|
| Linux SNFS-Only MDC Installation and Configuration | 44 |
| File System Requirements | 45 |
| Getting Ready to Install | 46 |
| Install a StorNext File System Only Using an RPM | 49 |
| Configure a StorNext RPM Server Using the Command Line | 51 |
| Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs | 56 |
| Obtain and Install StorNext Licenses without a GUI | 62 |

Linux SNFS-Only MDC Installation and Configuration

The RPM installation includes only StorNext File System (not StorNext Storage Manager) and does not include a GUI or components required to use the StorNext Replication feature. This topic describes how to install and configure File System-only on Linux MDCs. It also describes how to configure an RPM high-availability (HA) system.

This section covers:

- [File System Requirements on page 69](#)
- [Getting Ready to Install on page 71](#)
- [Install a StorNext File System Only Using an RPM on page 49](#)
- [Configure a StorNext RPM Server Using the Command Line on page 51](#)
- [Obtain and Install StorNext Licenses without a GUI on page 62](#)

For further information not covered in this topic, refer to the man pages that reside in the directory `/usr/cvfs/man`.

File System Requirements

To run StorNext File System, the machine onto which you are installing SNFS must meet all operating system, hardware, and LAN requirements.

Operating System Requirements

The operating systems, kernel versions, rpm package dependencies, and hardware platforms supported by StorNext File System are listed in the [StorNext 6 Compatibility Guide](#).

Make sure the system meets the operating system and platform requirements outlined in the guide, and, if necessary, update to a supported release or kernel version before installing StorNext.

i Note: For Linux clients, you must be disable C-States through kernel boot parameters. See [Modify the grub.conf File to Disable the Intel Sleep State on page 7](#).

Hardware Requirements

RAM and hard disk space requirements to run the StorNext File System are listed in the [StorNext Compatibility Guide](#).

Make sure the MDC meets the hardware requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

For improved performance and stability, it is strongly recommended that CPU C-States be disabled in the BIOS if they are supported by your system. Refer to system vendor documentation for additional information.

i Note: Gateway servers may require additional RAM depending on the number of file systems, LAN clients, and NICs used. See [Gateway Server/Client Network and Memory Tuning](#) for Gateway server/client memory tuning guidelines.

StorNext Client Hardware Requirements

RAM, hard disk, and CPU requirements to run the StorNext clients are listed in the [StorNext 6 Compatibility Guide](#).

-
- i Note:** Make sure the system to be used as a StorNext client meets the hardware requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

For improved performance and stability, on Windows clients it is strongly recommended that CPU C-States be disabled in the BIOS if they are supported by your system. Refer to system vendor documentation for additional information.

Windows Server 2012 SAN Disk Requirements

Windows disks on a Storage Area Network that will be used for StorNext, must be initialized before a StorNext volume label be written to the device. This involves writing a master boot record (MBR) to the disk.

This only needs to be done the first time new devices are connected to Windows.

LAN Requirements

StorNext LAN requirements are listed in the [StorNext 6 Compatibility Guide](#).

-
- i Note:** Make sure the MDC meets the network requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

Other Requirements

Additional requirements that must be met for the StorNext system are listed in the [StorNext 6 Compatibility Guide](#).

-
- i Note:** Make sure the system meets the additional requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

Getting Ready to Install

Before installing StorNext File System, complete the following tasks to ensure successful installation:

- Make sure the Metadata Controller (MDC) meets all operating system and hardware requirements (see [File System Requirements on page 69](#)).
- Correctly configure all storage devices and are verify they are visible to the MDC (see [Configure a Storage Device on the next page](#)).

- If you plan to use LUNs larger than 1 TB, decide on a label type and install any necessary operating system patches (see [Plan for LUNs Larger than 1 TB below](#)).
- Review and install any package dependencies (see [Package Dependencies on the next page](#)).

Configure a Storage Device

Before installing StorNext File System, make sure that all LUNs are visible to the MDC.

i Note: A LUN, or logical unit number, is a logical device that corresponds to one or more disks, drives, or storage devices.

If there are any connection issues, resolve them before installing StorNext. For assistance in configuring storage devices, refer to the documentation that came with the storage device, or contact the manufacturer.

i Note: LUNs that you plan to use in the same stripe group must be the same size. Consider this when configuring storage devices. (For more information about stripe groups, see the *StorNext User's Guide*.)

! Caution: StorNext does not support the connection of multiple devices through fibre channel hubs. Multiple devices must be connected using fibre channel switches.

! Caution: Quantum recommends that only the storage devices required for an environment are visible to each system. Unnecessary disk visibility will affect system start-up performance and increase the risk of inadvertent reuse of resources.

! Caution: Additionally, it is recommended that sites never perform operating (OS) installations while hardware is attached to a StorNext storage network. Numerous instances of OS installations accidentally overwriting StorNext volumes have been observed in the past.

Plan for LUNs Larger than 1 TB

StorNext supports LUNs greater than 1 TB in size if they are allowed by the operating system. To enable support for 1 TB or larger LUNs, all StorNext LUNs must be correctly labeled according to the requirements of the operating system running on the MDC as well as the operating system running on all connected clients. Disk LUNs can be labeled as EFI. Before installing StorNext, decide what label type will be used.

Required disk LUN label settings based on operating system and LUN size are presented in [Table 1 on the next page](#). Before installing StorNext, install any required operating system patches or updates (for both MDC and client) as indicated in the notes for [Table 1 on the next page](#).

i Note: After labeling a disk LUN, you must reboot systems running Solaris before they can access the disk LUN.

Table 1: Required Disk LUN Labels

| Operating System | Label for LUNs < 1 TB | Label for LUNs 1–2 TB | Label for LUNs > 2 TB |
|---|-----------------------|-----------------------|-----------------------|
| Windows ^a | EFI | EFI | EFI |
| Linux 2.6 Kernel | EFI | EFI | EFI |
| Solaris 10 ^b | EFI | EFI | EFI |
| HP-UX ^c | EFI | EFI | EFI |
| IBM AIX ^d | EFI | EFI | EFI |
| SGI IRIX ^e | Not supported | Not supported | Not supported |
| Apple Xsan 2.1.1 (macOS X 10.5 Leopard) | EFI | EFI | EFI |
| Apple Xsan 2.2 (macOS X 10.6 Snow Leopard) | EFI | EFI | EFI |

- a. For Windows XP 32-bit systems only, support for LUNs larger than 2 TB is limited to the first 2 TB. All other supported Windows platforms (including Windows XP 64-bit) support LUNs greater than 2 TB.
- b. Support for 2 TB or larger LUNs requires Solaris 10 Update 2 (06/06).
- c. Supports only the first 2 TB of LUNs larger than 2 TB.
- d. Supports only LUNs of 2.2 TB or less.
- e. IRIX is not supported starting with StorNext 4.0. However, StorNext 3.5 IRIX clients can be used with StorNext MDCs.

i Note: For best results, Quantum recommends using persistent binding for disk LUNs. For more information, contact the vendor of your HBA (host bus adapter).

Package Dependencies

- For management servers running Red Hat Enterprise Linux version 6 or 7, before installing SNFS and SNSM you must first install the kernel header files (shipped as the kernel-devel or kernel-devel-smp RPM, depending on your Linux distribution).
- For servers running SUSE Linux Enterprise Server, you must install the first kernel source code (shipped as the kernel-source RPM). StorNext will not operate correctly if these packages are not installed. You can install the kernel header files or kernel source RPMs by using the installation disks for your operating system.
- In addition, for MDC installs on Linux the PERL package is required.
- For customer supplied metadata-controllers and hardware, install the fuse-libs package on the system before installing the stornext_full or snfs_client_plus_mover packages.

Verify Hostname Length

The maximum hostname length for a StorNext server is limited to 255 characters. Before you begin the installation, verify that the destination hostname is not longer than 255 characters.

Verify You Have Enough Space to Extract the Installation Software

The installation software files are extracted into the `/tmp/stornext` directory by default. The software, when extracted, requires approximately 40 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

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

```
# ./<snfs_full.bin> -d <dir>
```

where `<snfs_full.bin>` is the name of the Linux RPM installation file, and `<dir>` is the directory into which you want to extract the installation files.

Install a StorNext File System Only Using an RPM

This section describes how to install StorNext using an RPM-only installation. To perform the install, you will need to use the correct file on the StorNext installation DVD or download for your operating system.

Install a StorNext File System Only Using an RPM

1. Log on to the MDC as **root**. For DVD installations, continue to next step. For installations from downloaded files, change to the directory containing the `.bin` file and continue to **Step 7**. If you only have the rpm files, continue to **Step 10**.
2. 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 ...
```

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

```
ls -l
```

4. Change to the **fs_only** directory.

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

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

```
ls -l
```

6. Identify the correct .bin file for your operating system and hardware platform. For example, for Red Hat Linux 5 running on an x86 64-bit platform, the file you will use is **snfs_full_RedHat50AS_26x86_64.bin**.
7. Launch the install script, where <file> is the name of the .bin file for your installation. At the command prompt, type:

```
./<file>
```

The End User License Agreement (EULA) displays.

8. Read the EULA. When finished you are asked if you accept.
9. Type **y** to accept the EULA and continue.

After you accept the EULA, the client software is extracted to **/tmp/stornext** (or to another directory if you specified an alternate location; see [Verify You Have Enough Space to Extract the Installation Software on page 73](#)).

i Note: If the directory into which the files are being extracted does already exist or 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, where <snfs_full.bin> is the name of the Linux server installation .bin file, and <dir> is the directory into which you want to extract the installation files. :

```
# ./<snfs_full.bin> -d <dir>
```

Quantum recommends that you do not force using an existing directory (by using the **-f** command) because this could lead to additional problems installing the .rpm files.

10. Change to the directory into which the files were extracted. The following command uses the default extract location for the files prompt, type:

```
cd /tmp/stornext
```

11. Display a list of the installed files. At the prompt, type:

```
ls -l
```

12. Install the required .rpm files in the following order. Use the command `rpm -ivh` before each package. For example, `rpm -ivh 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

Configure a StorNext RPM Server Using the Command Line

Below are procedures to configure a StorNext server in an RPM-only installation.

i Note: For RPM installations, specifically, when there is no “HA shared” file system, the HA manager program `/usr/cvfs/bin/snhamgr` is not functional.

Edit the `/etc/hosts` file in an HA Configuration

SUSE Linux distributions automatically associate the FQDN of the local machine with the address 127.0.0.2 in the `/etc/hosts` file. There is no benefit from doing this when the machine is connected to a

network that can resolve its name to an IP address. However, the existence of this entry can sometimes cause a failure of configuration synchronization within and between the server computers in an HA configuration. For this reason, the 127.0.0.2 entry should be deleted from the `/etc/hosts` file.

Configure a Single StorNext RPM Server or the Primary Server in an HA Configuration

The steps for setting up a single server or the primary server in an HA configuration are nearly identical. The HA configuration requires a few extra steps, which are noted in the procedure below.

i Note: If you do not have an HA configuration, ignore the steps marked **HA Only**.

Run the following commands on the **Primary** node if installing an HA system:

1. Install RPMs using the procedure outlined in [Install a StorNext File System Only Using an RPM on page 49](#).
2. Obtain and install licenses. See [Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs on page 56](#).
3. Populate the `fsnameservers` file with the IP addresses of the `fsnameserver` systems by running this command:

```
# vi /usr/cvfs/config/fsnameservers
```

i Note: For additional information, see the `fsnameservers` man page. If you add the host name instead of the IP address, the conversion to HA could fail. Use **only** the IP address when populating the `fsnameservers` file.

4. (**HA Only**) Create the `ha_peer` file by running the command below. This file should contain the IP address of the **Secondary** in the HA pair.

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

i Note: For additional information, see the `ha_peer` man page.

5. Start the **fsmpm** by running the command:

```
# /etc/init.d/cvfs start
```

6. Create the `.cfgx` file by running the command:

```
# /usr/cvfs/bin/sncfgedit -n snfs1
```

i Note: For additional information, see the `snfs.cfgx` and `sncfgedit` man pages.

After creating the file, do the following:

- a. Delete first two lines at top that start with `****`.
- b. **(HA Only)** Change the `haFsType` entry to `HaUnmanaged`. That is, change this:

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

to this:

```
<haFsType>HaUnmanaged</haFsType>
```

- c. Fix the `<diskTypes>` and `<stripeGroups>` sections at the bottom.

7. Make the file system by running the command:

```
# /usr/cvfs/bin/cvmkfs snfs1
```

8. Create the mount point by running the command:

```
# mkdir /stornext/snfs1
```

9. Edit the file `/usr/cvfs/config/fsmlist` and add your file system name(s) that the file system services will start. To do this:

- a. Open the file by running this command:

```
# vi /usr/cvfs/config/fsmlist
```

- b. Add the name of your file system to launch:

```
snfs1
```


10. Add file system entry to **/etc/fstab** file:
 - a. Edit the **/etc/fstab** file by running the command:

```
# vi /etc/fstab
```

- b. Add the file system name and mount options to the end of the list:

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

Important

There are six fields in the fstab file. Spaces are used between the entries to delineate one field from the other. The fields are:

- **First** field: Name of the file system to be mounted.
- **Second** field: Mount point for the file system.
- **Third** field: Type of file system.
- **Fourth** field: Mount options.

i Note: You must always specify the **noauto** option for the **Mount** options.

- **Fifth** field: Flag that tells Linux if this file system need to be dumped.
- **Sixth** field: Determines the order in which file system checks are done.

11. Start and mount the file system by running the command:

```
# /usr/cvfs/bin/cvadmin -e 'start snfs1'  
# mount /stornext/snfs1
```

Configure the Secondary StorNext RPM Server in an HA Configuration

Perform this procedure only if you have an HA configuration. Run the following commands on the **Secondary** node:

1. Install RPMs using the procedure outlined in [Install a StorNext File System Only Using an RPM on page 49](#).
2. Create the `fsnameservers` file that contains the IP address of the system acting as the name servers by doing the following:

- a. Open the file by running the command:

```
# vi /usr/cvfs/config/fsnameservers
```

- b. Add the IP address of the name servers.

i Note: The `fsnameservers` file on the secondary must be the same as the `fsnameservers` file on the primary. Remember, if you add the host name instead of the IP address, the conversion to HA could fail. Use **only** the IP address when populating the `fsnameservers` file.

3. Create the `ha_peer` file and add the IP address of the **Primary** in the HA pair by doing the following:

- a. Open the file by running the command:

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

- b. Add the IP address of the **Primary**.

4. Start up the `fsm` by running the command:

```
# /etc/init.d/cvfs start
```

5. Copy the `.cfgx` file from the Primary to the `/usr/cvfs/config` directory on the **Secondary**.

6. Edit the file `/usr/cvfs/config/fsmlist` and add your file system name(s) that the file system services will start by doing the following:

- a. Run the command:

```
# vi /usr/cvfs/config/fsmlist
```

- b. Add the name of your file system to launch:

```
snfs1
```

7. Add file system entry to **/etc/fstab** file:
 - a. Edit the **/etc/fstab** file by running the command:

```
# vi /etc/fstab
```

- b. Add the file system name and mount options to the end of the list:

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

8. Start and mount the file system by running the command:

```
# /usr/cvfs/bin/cvadmin -e 'start snfs1'  
# mount /stornext/snfs1
```

Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs

Non-Expiring vs. Subscription License Models

StorNext 7 (or later) with Subscription Licenses:

If your system is running StorNext 7 (or later), contact your [Quantum Sales](#) representative to determine which non-expiring or subscription licenses are best suitable for your environment. See [StorNext Licenses](#) for more information.

StorNext 7 (or later) and StorNext 6.x or earlier with Non-Expiring Licenses:

StorNext and some StorNext non-expiring license features are pre-installed on your system, and others must be purchased and then enabled by installing a license.dat file on the system. If your system is using a 30-day auto-generated license, you should upload a license file containing permanent licenses as soon as it is convenient, so that the **Maintenance** expiration date and the **(FlexTier™)** subscription date for the license is set correctly on the system.

- i Note:** For non-expiring licenses, you can determine the current license limits for many StorNext features, such as SAN clients, LAN clients, DDM, and Storage Manager licenses, on the **Configuration > Licenses** page of the StorNext GUI and on the Quantum **License Key Management** page available online at <https://www.quantum.com/licensekeys>.

Obtain StorNext Licenses

Prerequisites

In order to receive a license file (`license.dat`) to enable StorNext features or capacity, you will need the following:

- The **System Serial Number** or the **StorNext Software Serial Number**.

System serial numbers and StorNext software serial numbers are alphanumeric. The table below provides examples of serial numbers.

| Type of Serial Number | Example(s) of Serial Number(s) |
|--|--|
| System Serial Number | AV1607CKH12345 CX1306CKB12345 EW1450CKD12345 |
| StorNext Software Serial Number | SN12345 |

- The `cvfsid` string for each server node. See [How to Generate a cvfsid for Linux and Unix Configurations and Quantum Appliances on page 59](#), [How to Generate a cvfsid for Windows Configurations on page 59](#), or [How to Generate a cvfsid for Xsan Configurations on page 60](#), depending on your configuration.

Example: `ECF4BCEECC0E linux 0 xcellis13`

The table below provides information regarding which configurations require one `cvfsid` string or two `cvfsid` strings.

| Configurations Requiring One cvfsid String | Configurations Requiring Two cvfsid Strings |
|--|--|
| <ul style="list-style-type: none"> • Customer-configured MDC without an HA license • Xcellis Workflow Extender • G300 • Xcellis Workflow Director (single-node systems only) • Artico (single-node systems only) • Xcellis Foundation (single-node systems only) • aiWARE for Xcellis Standard and High Performance | <ul style="list-style-type: none"> • Customer-configured MDC with an HA license • Xcellis Workflow Director (dual-node systems only) • Xcellis Foundation (dual-node systems only) • Artico (dual-node systems only) • aiWARE for Xcellis Embedded (DAE) • Pro Solutions • M660 • M440 • M330 |

How to Obtain a New StorNext license.dat File

1. Visit the Quantum **License Key Management** page at <https://www.quantum.com/licensekeys>. If you cannot access the web page or need additional help, contact Quantum Support at <https://www.quantum.com/ServiceandSupport>.
2. Enter your serial number in the **Serial Number** field.
3. Click **Submit**. The license(s) associated with the serial number you entered in **Step 2** are displayed in the **License Details** section.

i Note: If you have not previously entered a `cvfsid` for the serial number, enter the strings in the **CVFSID 1** or **CVFSID 2** fields and then click **Set CVFSID**. See [How to Generate a cvfsid for Linux and Unix Configurations and Quantum Appliances on the next page](#), [How to Generate a cvfsid for Windows Configurations on the next page](#), or [How to Generate a cvfsid for Xsan Configurations on page 60](#), depending on your configuration. The **CVFSIDs** sub-section under **License Details** is updated with the `cvfsid` value(s) for the serial number.

⚠ Caution: Ensure your `cvfsid` value(s) are correct. Once entered and submitted, they cannot be changed.

i Note: At any time after you have submitted your serial number, you may enter another serial number by clicking **Choose Another**.

4. Click **Generate License Key File** once the `cvfsid` value(s) have been entered.

i Note: There may be some cases where the Quantum **License Key Management** site will not be able to generate a license file. In this case, you are re-directed to an alternate web page to continue the process.

5. Save the `license.dat` file to a temporary location on your computer. You will upload this file on the **Licenses** page of the StorNext GUI. To install the licenses in the `license.dat` file, use the [Install a StorNext License from the GUI for non-HA Installations](#) procedure if you have a single-node StorNext MDC, or use the [Install/Update Licenses from the GUI for HA Installations](#) procedure if you have a High Availability (HA) StorNext MDC pair.

How to Generate a `cvfsid` for Linux and Unix Configurations and Quantum Appliances

For Quantum appliances, open an SSH connection to the server

Using the IP address assigned to the server on the StorNext management/public network.

1. Log in to the command line using the following credentials:
 - User name: `stornext`
 - Password: `<stornext user account password>`

i Note: "password" is the default password for the `stornext` user account. If the password has been changed, use the current password. Beginning with StorNext 7.0.2, you must change the default password for the `stornext` user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

2. Enter `sudo rootsh` to gain root user access.
3. Enter the password for the `stornext` user account again.

Generate the `cvfsid`

1. Enter the `cvfsid` command.

The system returns the following information:

`<server_MAC_address> <server_OS> <server_name>`

Example:

```
ECF4BCDECC0E linux 0 xcellis13
```

2. Close the SSH session for the node.
3. **(Optional)** If your system is configured for HA, repeat for the other node.

How to Generate a `cvfsid` for Windows Configurations

Follow the procedure below to generate the `cvfsid` on the command line of the MDC node:

1. On the **Start** menu, click **Programs**, click **StorNext File System**, and then click **License Identifier**.
2. The `cvfsid` string appears in a popup.

How to Generate a cvfsid for Xsan Configurations


Follow the procedure below to generate the cvfsid on the command line of the MDC node:

1. Navigate to `/Library/Filesystems/Xsan/bin`.
2. Execute the command `sudo cvfsid`.


Install/Update StorNext Licenses From the GUI for Dual-node (HA) Systems

StorNext 7 (or later) with Subscription Licenses:

If your system is running StorNext 7 (or later), contact your [Quantum Sales](#) representative to determine which non-expiring or subscription licenses are best suitable for your environment. See [StorNext Licenses](#) for more information.

 **Caution:** At the end of this procedure, exiting Config Mode will restart StorNext services (on both server nodes for dual-server node systems). This will cause a disruption in client operations. Quantum strongly recommends halting I/O operations for connected StorNext clients before beginning the process.

1. If you are logging into StorNext for the first time, read the end-user license agreement (EULA) carefully, and then click **Accept** (you will not see the EULA page if you have previously logged into StorNext).
2. For new installations, the **Configuration > Licenses Entry** screen appears. Skip to step 4.
For operational systems, navigate to **Tools > High Availability > Manage**.
3. **(Optional)** If your **cvfsid** has changed, then click **Enter Config Mode** to halt StorNext client I/O operations. A dialog appears to confirm you want to lock the HA cluster configuration. Click **Yes**.
(See the **High Availability Systems** section in the appendix of the current *StorNext User's Guide* for further information about HA systems.)
4. Select **Configuration > Licenses**.
5. Scroll down to the bottom of the page, and click **Import...** to install and/or update StorNext licenses.
6. Click **Choose File**, select the `license.dat` file provided for your system from the temporary location where you downloaded it, and click **Open**.

 **Caution:** Importing a license file overwrites the existing license file and all existing licenses. Licenses not present in the imported license file are removed from StorNext.

You will know that the licenses have been applied to your system when you see the confirmation “**License file license.dat imported successfully.**” appear at the top of the page. The status indicator for the licensed features is displayed as a green check mark icon.

7. If the system is in config mode, navigate to **Tools > High Availability > Manage** and click **Exit Config Mode**. You will be prompted for confirmation that you want to unlock the cluster. Click **Yes** to exit config mode.

The procedure is complete.

Install StorNext Licenses From the GUI for Single-node (non-HA) Systems

-
- i Note:** Use these steps to install the license keys on a single-node system, or one that has not been configured as HA. For HA systems, see [Obtain and Install StorNext Non-Expiring Licenses – GUI-based Appliances and Customer-supplied MDCs on page 56](#).

1. If you are logging into StorNext for the first time, read the end-user license agreement (EULA) carefully, and then click **Accept** (you will not see the EULA page if you have previously logged into StorNext).

-
- i Note:** To comply with the California Senate Bill No. 327 (CA SB-327), you must change your password for your **admin** and **service** user accounts on your first log in attempt.

2. For new installations, the **Configuration > Licenses Entry** page appears. Skip to step 5.
3. For operational systems, halt StorNext client I/O operations. Stop the StorNext filesystem by navigating to **Tools > System Control**, and selecting **Stop** in the **File System** section of the page.
4. Select **Configuration > Licenses**.
5. Scroll down to the bottom of the page, and click **Import...** to install and/or update StorNext licenses.
6. Click **Choose File**, select the **license.dat** file provided for your system from the temporary location where you downloaded it, and click **Open**.

-
- ! Caution:** Importing a license file overwrites the existing license file and all existing licenses. Licenses not present in the imported license file are removed from StorNext.

You will know that the licenses have been applied to your system when you see the confirmation “**License file license.dat imported successfully.**” appear at the top of the page. The status indicator for the licensed features is displayed as a green check mark icon.

7. After you install your licenses, restart the StorNext filesystem by navigating to **Tools > System Control**, and selecting **Start** in the **File System** section of the page. StorNext may not recognize some feature licenses until you restart services.

The procedure is complete.

Install a StorNext License on a Windows MDC

For help installing Windows MDC licenses, refer to the StorNext help, which can be found at: **Windows Start menu > All Programs > StorNext File System > StorNext Help**.

Obtain and Install StorNext Licenses without a GUI

For File System installations, or in instances where there is no available StorNext GUI, StorNext license keys for the system can be installed manually.

Obtain StorNext Licenses

Prerequisites

In order to receive a license file (`license.dat`) to enable StorNext features or capacity, you will need the following:

- The **System Serial Number** or the **StorNext Software Serial Number**.

System serial numbers and StorNext software serial numbers are alphanumeric. The table below provides examples of serial numbers.

| Type of Serial Number | Example(s) of Serial Number(s) |
|--|--|
| System Serial Number | AV1607CKH12345 CX1306CKB12345 EW1450CKD12345 |
| StorNext Software Serial Number | SN12345 |

- The `cvfsid` string for each server node. See [How to Generate a cvfsid for Linux and Unix Configurations and Quantum Appliances on page 64](#), [How to Generate a cvfsid for Windows Configurations on page 64](#), or [How to Generate a cvfsid for Xsan Configurations on page 65](#), depending on your configuration.

Example: `ECF4BCEECC0E linux 0 xcellis13`

The table below provides information regarding which configurations require one `cvfsid` string or two `cvfsid` strings.

| Configurations Requiring One cvfsid String | Configurations Requiring Two cvfsid Strings |
|--|--|
| <ul style="list-style-type: none">• Customer-configured MDC without an HA license• Xcellis Workflow Extender• G300• Xcellis Workflow Director (single-node systems only)• Artico (single-node systems only)• Xcellis Foundation (single-node systems only)• aiWARE for Xcellis Standard and High Performance | <ul style="list-style-type: none">• Customer-configured MDC with an HA license• Xcellis Workflow Director (dual-node systems only)• Xcellis Foundation (dual-node systems only)• Artico (dual-node systems only)• aiWARE for Xcellis Embedded (DAE)• Pro Solutions• M660• M440• M330 |

How to Obtain a New StorNext license.dat File

1. Visit the Quantum **License Key Management** page at <https://www.quantum.com/licensekeys>. If you cannot access the web page or need additional help, contact Quantum Support at <https://www.quantum.com/ServiceandSupport>.
2. Enter your serial number in the **Serial Number** field.
3. Click **Submit**. The license(s) associated with the serial number you entered in **Step 2** are displayed in the **License Details** section.

i Note: If you have not previously entered a `cvfsid` for the serial number, enter the strings in the **CVFSID 1** or **CVFSID 2** fields and then click **Set CVFSID**. See [How to Generate a cvfsid for Linux and Unix Configurations and Quantum Appliances on the next page](#), [How to Generate a cvfsid for Windows Configurations on the next page](#), or [How to Generate a cvfsid for Xsan Configurations on page 65](#), depending on your configuration. The **CVFSIDs** sub-section under **License Details** is updated with the `cvfsid` value(s) for the serial number.

⚠ Caution: Ensure your `cvfsid` value(s) are correct. Once entered and submitted, they cannot be changed.

i Note: At any time after you have submitted your serial number, you may enter another serial number by clicking **Choose Another**.

4. Click **Generate License Key File** once the `cvfsid` value(s) have been entered.

i Note: There may be some cases where the Quantum **License Key Management** site will not be able to generate a license file. In this case, you are re-directed to an alternate web page to continue the process.

5. Save the `license.dat` file to a temporary location on your computer. You will upload this file on the **Licenses** page of the StorNext GUI. To install the licenses in the `license.dat` file, use the [Install a StorNext License from the GUI for non-HA Installations](#) procedure if you have a single-node StorNext MDC, or use the [Install/Update Licenses from the GUI for HA Installations](#) procedure if you have a High Availability (HA) StorNext MDC pair.

How to Generate a `cvfsid` for Linux and Unix Configurations and Quantum Appliances

For Quantum appliances, open an SSH connection to the server

Using the IP address assigned to the server on the StorNext management/public network.

1. Log in to the command line using the following credentials:
 - User name: `stornext`
 - Password: `<stornext user account password>`

i Note: "password" is the default password for the `stornext` user account. If the password has been changed, use the current password. Beginning with StorNext 7.0.2, you must change the default password for the `stornext` user account. As of StorNext 7.0.2, the first time you log in, you will be prompted to change the password to a different one.

2. Enter `sudo rootsh` to gain root user access.
3. Enter the password for the `stornext` user account again.

Generate the `cvfsid`

1. Enter the `cvfsid` command.

The system returns the following information:

`<server_MAC_address> <server_OS> <server_name>`

Example:

```
ECF4BCDECC0E linux 0 xcellis13
```

2. Close the SSH session for the node.
3. **(Optional)** If your system is configured for HA, repeat for the other node.

How to Generate a `cvfsid` for Windows Configurations

Follow the procedure below to generate the `cvfsid` on the command line of the MDC node:

1. On the **Start** menu, click **Programs**, click **StorNext File System**, and then click **License Identifier**.
2. The `cvfsid` string appears in a popup.

How to Generate a cvfsid for Xsan Configurations

Follow the procedure below to generate the cvfsid on the command line of the MDC node:

1. Navigate to `/Library/Filesystems/Xsan/bin`.
2. Execute the command `sudo cvfsid`.

Install the StorNext License File on Xcellis Workflow Extender and G300 Systems

1. When you receive the new `license.dat` file from Quantum, copy the file to the service laptop you will use to remotely access the server.
2. Initiate an SSH session to the command line of the server.
3. Stop StorNext services:

```
systemctl stop cvfs
```

4. Copy the license file from the service laptop to the following directory on the system:

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

If you don't know how to do this, open the following section. (You will be unable to apply the `license.dat` file if it has not first been transferred to the system.):

Prerequisites before you move the license.dat file

An scp client must be installed on the system you plan to use for remote access and configured the software for binary scp file transfers.

Examples:

- **WinSCP** – see Transfer the license.dat file to the system with an SCP client
- **PuTTY pscp.exe** (Secure Copy client) – see Transfer the license.dat file to the system with the PuTTY SCP client

Transfer the license.dat file to the system via an SCP client with GUI:

For this example, we are using a Windows SCP client. The procedures may vary from the client you will use.

- a. Open an SCP session to the server, using the server IP on the metadata network or by connecting to the service port and using its IP address, if you are directly connected to the server.
- b. Login with the default **stornext** credentials:

User name: stornext

password: password

- c. Drag and drop the license.dat file from the directory on the service laptop where you moved the file after downloading to the /home/stornext directory on the server.

OR Transfer the license.dat file to the system with a command-line SCP client:

For this example, we are using the Windows command-line PuTTY SCP client (client software must be installed before doing the procedure). These procedures will vary from the client you will use.

- a. Open a command prompt (from the Windows Start Menu, type cmd in the search field and click Enter).
- b. Enter the following (the default installation location for PuTTY pscp.exe is used in the first argument):

```
"C:\<path to pscp.exe>\pscp.exe" C:\<local directory to copy>\license.dat  
stornext@host server>:/home/stornext
```

Example:

```
"C:\Program Files (x86)\PuTTY\pscp.exe" C:\<local directory to  
copy>\license.dat stornext@host server>:/home/stornext
```

- c. When prompted, enter **stornext** for the user account.
- d. When prompted, enter the password for the **stornext** user account on the system.
- e. The license.dat file is transferred to the system.

Transfer the file to the usr/cvfs/config directory from the server command-line:

The /usr/cvfs/config directory on the system, which is used by the Service Menu to apply the license.

- a. Open an SSH connection to the system using IP address 10.17.21.1 StorNext public (management), Metadata or LAN client network. Use the IP address assigned if different from the default used here
- b. Login with the default **stornext** credentials:

User name: stornext

password: password

- c. At the command prompt enter the following to gain root user access:

```
sudo rootsh
```

- d. Enter the password for the **stornext** user account again when prompted.

- e. Copy the license.dat file from /home/stornext to /usr/cvfs/config.

Example:

```
cp /home/stornext/license.dat /usr/cvfs/config/
```

- f. Verify the license.dat file is located in the installation directory:

```
ls /usr/cvfs/config
```

- g. Output should display a license.dat file:

Example:

```
deviceparams fsnameservers pseudodevs sngateway_configfilelist license.dat  
qustat.conf
```

5. Start StorNext services:

```
systemctl start cvfs
```

6. The system will restart the services and use the new license file.
7. Close your SSH session and disconnect from the server.



Chapter 4: Windows SNFS-Only MDC Installation and Configuration

This chapter contains the following topics:

| | |
|--|-----|
| Windows SNFS-Only MDC Installation and Configuration | 68 |
| File System Requirements | 69 |
| Getting Ready to Install | 71 |
| Install the StorNext File System on a Windows MDC: Interactively | 73 |
| Configure a StorNext File System on Windows | 82 |
| Configure a High Availability System on Windows | 136 |
| Add or Remove StorNext Components | 138 |
| Restore a Previous Client Configuration | 140 |
| Upgrade or Reinstall StorNext on a Windows Server | 141 |

Windows SNFS-Only MDC Installation and Configuration

The following sections describe how to install and configure StorNext File System (SNFS) on a Windows Metadata Controller (MDC). Since Windows systems do not support StorNext Storage Manager, the

Windows installation does not include procedures for installing the StorNext Storage Manager.

Once StorNext File System is installed, you can configure it using the Windows-based configuration utilities to set up server and client features. To configure StorNext, you enter license information and create one or more file systems.

- [File System Requirements below](#)
- [Getting Ready to Install on page 71](#)
- [Install the StorNext File System on a Windows MDC: Interactively on page 73](#)
- [Configure a StorNext File System on Windows on page 82](#)
- [Configure a High Availability System on Windows on page 136](#)
- [Add or Remove StorNext Components on page 138](#)
- [Restore a Previous StorNext Client Configuration on page 174](#)
- [Upgrade or Reinstall StorNext on a Windows Server on page 141](#)

File System Requirements

To run StorNext File System, the machine onto which you are installing SNFS must meet all operating system, hardware, and LAN requirements.

Operating System Requirements

The operating systems, kernel versions, rpm package dependencies, and hardware platforms supported by StorNext File System are listed in the [StorNext 6 Compatibility Guide](#).

Make sure the system meets the operating system and platform requirements outlined in the guide, and, if necessary, update to a supported release or kernel version before installing StorNext.

i Note: For Linux clients, you must be disable C-States through kernel boot parameters. See [Modify the grub.conf File to Disable the Intel Sleep State on page 7](#).

Hardware Requirements

RAM and hard disk space requirements to run the StorNext File System are listed in the [StorNext Compatibility Guide](#).

Make sure the MDC meets the hardware requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

For improved performance and stability, it is strongly recommended that CPU C-States be disabled in the BIOS if they are supported by your system. Refer to system vendor documentation for additional information.

-
- i Note:** Gateway servers may require additional RAM depending on the number of file systems, LAN clients, and NICs used. See [Gateway Server/Client Network and Memory Tuning](#) for Gateway server/client memory tuning guidelines.

StorNext Client Hardware Requirements

RAM, hard disk, and CPU requirements to run the StorNext clients are listed in the [StorNext 6 Compatibility Guide](#).

-
- i Note:** Make sure the system to be used as a StorNext client meets the hardware requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

For improved performance and stability, on Windows clients it is strongly recommended that CPU C-States be disabled in the BIOS if they are supported by your system. Refer to system vendor documentation for additional information.

Windows Server 2012 SAN Disk Requirements

Windows disks on a Storage Area Network that will be used for StorNext, must be initialized before a StorNext volume label be written to the device. This involves writing a master boot record (MBR) to the disk.

This only needs to be done the first time new devices are connected to Windows.

LAN Requirements

StorNext LAN requirements are listed in the [StorNext 6 Compatibility Guide](#).

-
- i Note:** Make sure the MDC meets the network requirements outlined in the guide, and, if necessary, update the MDC before installing StorNext.

Other Requirements

Additional requirements that must be met for the StorNext system are listed in the [StorNext 6 Compatibility Guide](#).

-
- i Note:** Make sure the system meets the additional requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

Getting Ready to Install

Before installing StorNext File System, complete the following tasks to ensure successful installation:

- Make sure the Metadata Controller (MDC) meets all operating system and hardware requirements (see [File System Requirements on page 69](#)).
- Correctly configure all storage devices and are verify they are visible to the MDC (see [Configure a Storage Device below](#)).
- If you plan to use LUNs larger than 1 TB, decide on a label type and install any necessary operating system patches (see [Plan for LUNs Larger than 1 TB below](#)).
- Review and install any package dependencies (see [Package Dependencies on the next page](#)).

Configure a Storage Device

Before installing StorNext File System, make sure that all LUNs are visible to the MDC.

i Note: A LUN, or logical unit number, is a logical device that corresponds to one or more disks, drives, or storage devices.

If there are any connection issues, resolve them before installing StorNext. For assistance in configuring storage devices, refer to the documentation that came with the storage device, or contact the manufacturer.

i Note: LUNs that you plan to use in the same stripe group must be the same size. Consider this when configuring storage devices. (For more information about stripe groups, see the *StorNext User's Guide*.)

! Caution: StorNext does not support the connection of multiple devices through fibre channel hubs. Multiple devices must be connected using fibre channel switches.

! Caution: Quantum recommends that only the storage devices required for an environment are visible to each system. Unnecessary disk visibility will affect system start-up performance and increase the risk of inadvertent reuse of resources.

! Caution: Additionally, it is recommended that sites never perform operating (OS) installations while hardware is attached to a StorNext storage network. Numerous instances of OS installations accidentally overwriting StorNext volumes have been observed in the past.

Plan for LUNs Larger than 1 TB

StorNext supports LUNs greater than 1 TB in size if they are allowed by the operating system. To enable support for 1 TB or larger LUNs, all StorNext LUNs must be correctly labeled according to the requirements of the operating system running on the MDC as well as the operating system running on all connected clients. Disk LUNs can be labeled as EFI. Before installing StorNext, decide what label type

will be used.

Required disk LUN label settings based on operating system and LUN size are presented in [Table 1 below](#). Before installing StorNext, install any required operating system patches or updates (for both MDC and client) as indicated in the notes for [Table 1 below](#).

i Note: After labeling a disk LUN, you must reboot systems running Solaris before they can access the disk LUN.

Table 1: Required Disk LUN Labels

| Operating System | Label for LUNs < 1 TB | Label for LUNs 1–2 TB | Label for LUNs > 2 TB |
|---|-----------------------|-----------------------|-----------------------|
| Windows ^a | EFI | EFI | EFI |
| Linux 2.6 Kernel | EFI | EFI | EFI |
| Solaris 10 ^b | EFI | EFI | EFI |
| HP-UX ^c | EFI | EFI | EFI |
| IBM AIX ^d | EFI | EFI | EFI |
| SGI IRIX ^e | Not supported | Not supported | Not supported |
| Apple Xsan 2.1.1 (macOS X 10.5 Leopard) | EFI | EFI | EFI |
| Apple Xsan 2.2 (macOS X 10.6 Snow Leopard) | EFI | EFI | EFI |

a. For Windows XP 32-bit systems only, support for LUNs larger than 2 TB is limited to the first 2 TB. All other supported Windows platforms (including Windows XP 64-bit) support LUNs greater than 2 TB.

b. Support for 2 TB or larger LUNs requires Solaris 10 Update 2 (06/06).

c. Supports only the first 2 TB of LUNs larger than 2 TB.

d. Supports only LUNs of 2.2 TB or less.

e. IRIX is not supported starting with StorNext 4.0. However, StorNext 3.5 IRIX clients can be used with StorNext MDCs.

i Note: For best results, Quantum recommends using persistent binding for disk LUNs. For more information, contact the vendor of your HBA (host bus adapter).

Package Dependencies

- For management servers running Red Hat Enterprise Linux version 6 or 7, before installing SNFS and SNSM you must first install the kernel header files (shipped as the kernel-devel or kernel-devel-smp RPM, depending on your Linux distribution).
- For servers running SUSE Linux Enterprise Server, you must install the first kernel source code (shipped as the kernel-source RPM). StorNext will not operate correctly if these packages are not

installed. You can install the kernel header files or kernel source RPMs by using the installation disks for your operating system.

- In addition, for MDC installs on Linux the PERL package is required.
- For customer supplied metadata-controllers and hardware, install the `fuse-libs` package on the system before installing the `stornext_full` or `snfs_client_plus_mover` packages.

Verify Hostname Length

The maximum hostname length for a StorNext server is limited to 255 characters. Before you begin the installation, verify that the destination hostname is not longer than 255 characters.

Verify You Have Enough Space to Extract the Installation Software

The installation software files are extracted into the `/tmp/stornext` directory by default. The software, when extracted, requires approximately 40 MB of space. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

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

```
# ./<snfs_full.bin> -d <dir>
```

where `<snfs_full.bin>` is the name of the Linux RPM installation file, and `<dir>` is the directory into which you want to extract the installation files.

Install the StorNext File System on a Windows MDC: Interactively

The StorNext setup wizard guides you through the process of installing the StorNext File System on Windows machines.

When you are ready, use the setup wizard to install StorNext (see [Install a StorNext File System on the next page](#)).

If you did not install all StorNext components when you initially installed StorNext or wish to remove them later, you can add or remove components at any time (see [Add or Remove StorNext Components on page 138](#)).

(Optional) After installation, restore the previous client configuration (see [Restore a Previous StorNext Client Configuration on page 174](#)). This is only necessary if you have performed a complete removal of StorNext. The StorNext client configurations are fully maintained on upgrades from version 3.0 and later.

As new versions of StorNext are released, you can upgrade to a new version or reinstall (see [Upgrade or Reinstall StorNext on a Windows Server on page 141](#)).

-
- i Note:** You must log on as an Administrator to install StorNext. If you are installing on Windows Vista, answer Yes to any messages asking if you want to run the installation process with administrative privileges.

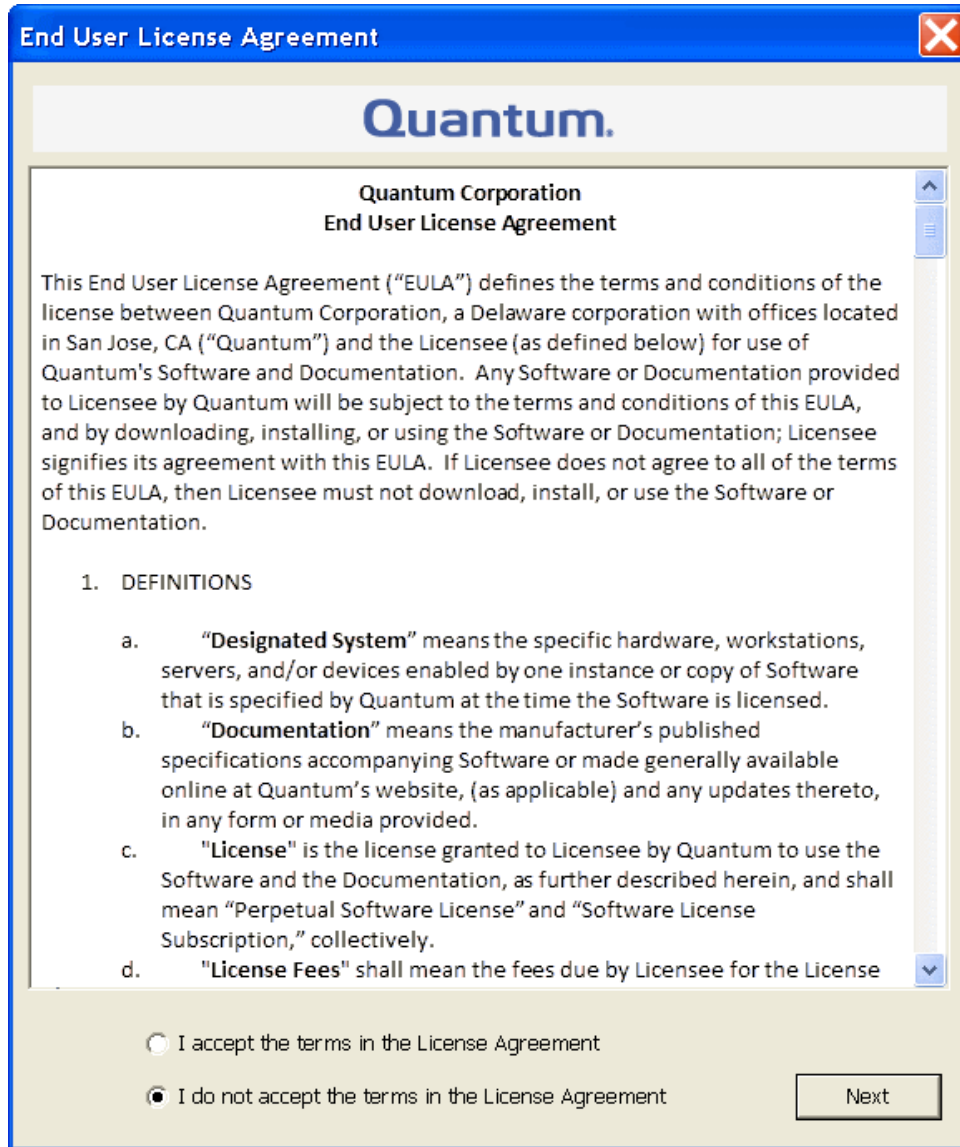
Install a StorNext File System

1. Copy the file named **/fs_only/snfs_full_Windows_x86_64.zip** from the installation DVD and place it on the system onto which you want to install it. For the SNFS-only installation, be sure to select the file from the **fs_only** directory and not the **stornext_full** directory.
2. Unzip the file.
3. Double-click the file named **Snfssetup.exe** to launch the installer. The **StorNext Installation** window appears.


-
- i Note:** For existing installations, do not use this procedure. Refer to the **StorNext Upgrade Guide** for upgrading existing StorNext.



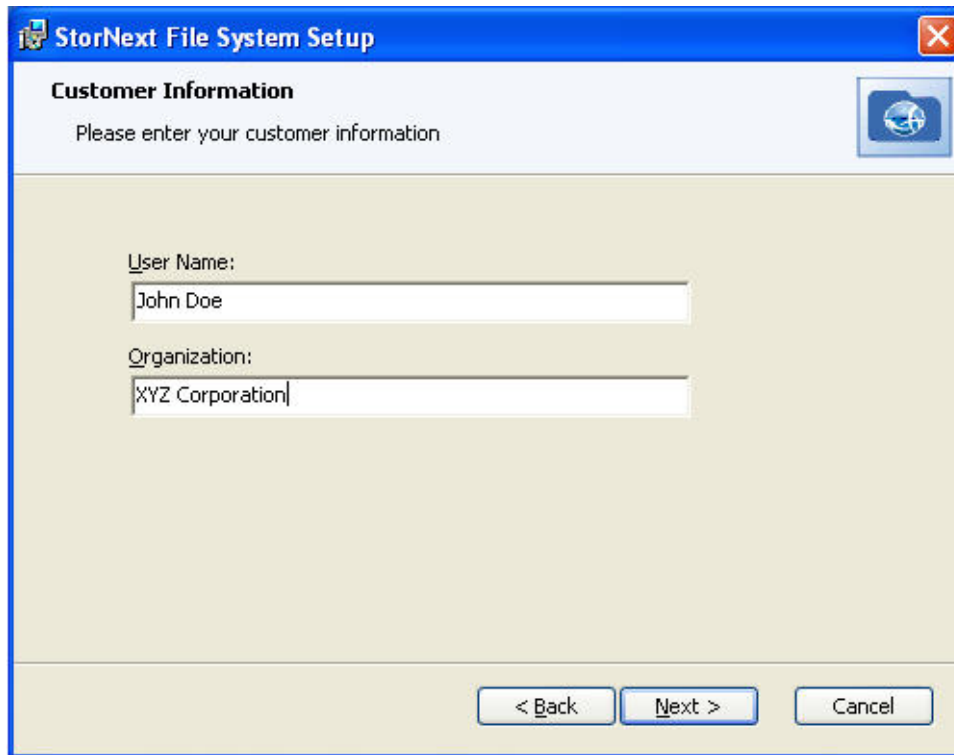
4. Click **Install StorNext** to begin installation. The **StorNext File System License Agreement** window appears.



5. Read the license. If you agree to the terms, select "I accept the terms in the License Agreement" and click **Next** to continue the install. The StorNext setup wizard launches.




6. Click **Next** to continue. The **Customer Information** window appears.



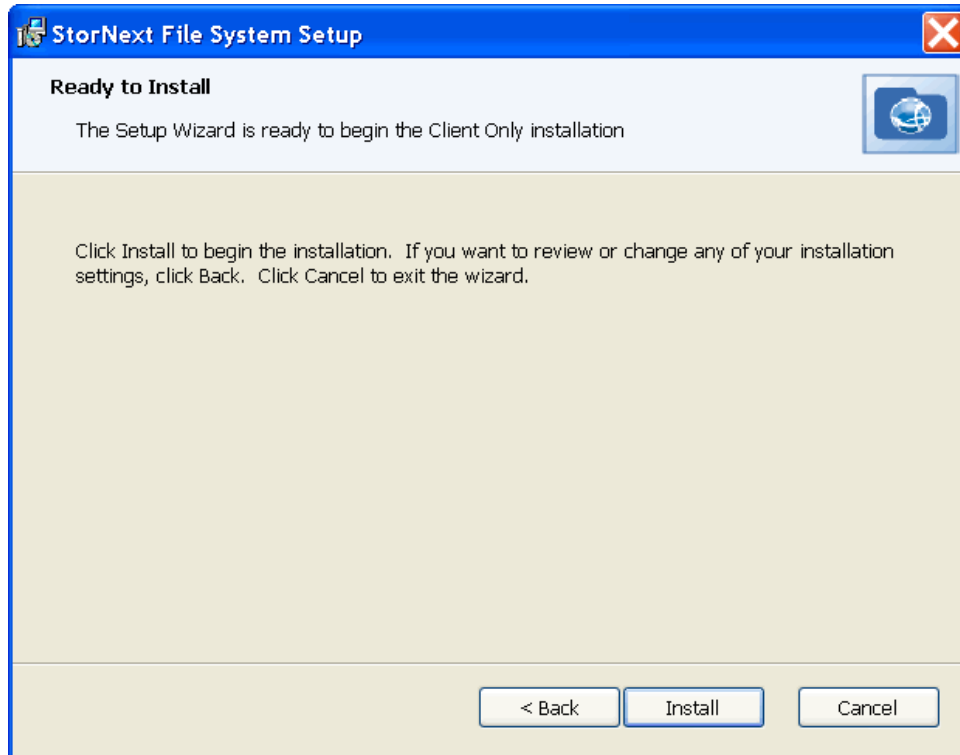
The screenshot shows a Windows-style dialog box titled "StorNext File System Setup". The main heading is "Customer Information" with a sub-instruction "Please enter your customer information". There are two text input fields: "User Name:" containing "John Doe" and "Organization:" containing "XYZ Corporation". At the bottom, there are three buttons: "< Back", "Next >", and "Cancel". A small globe icon is visible in the top right corner of the dialog's content area.

7. Type your name and the name of your company or organization in the boxes, and then click **Next** to continue. The **Choose Setup Type** window appears.




8. Click one of the following installation options (to install StorNext File System you would click **Custom** or **Complete**):
- **Client Only:** Installs only the StorNext client software and help files in the default location (**C:\Program Files\StorNext**).
 - **Custom:** Lets you choose which components to install and specify an installation location. When ready, click **Next**.
 - **Complete:** All installable components will be installed in the default location (**C:\Program Files\StorNext**). If the system is capable of running the StorNext server, the StorNext server will also be installed.

The **Ready to Install** window appears.



9. Click **Install** to continue. Wait while the setup wizard installs StorNext. When installation is complete, the **Finish** window appears.

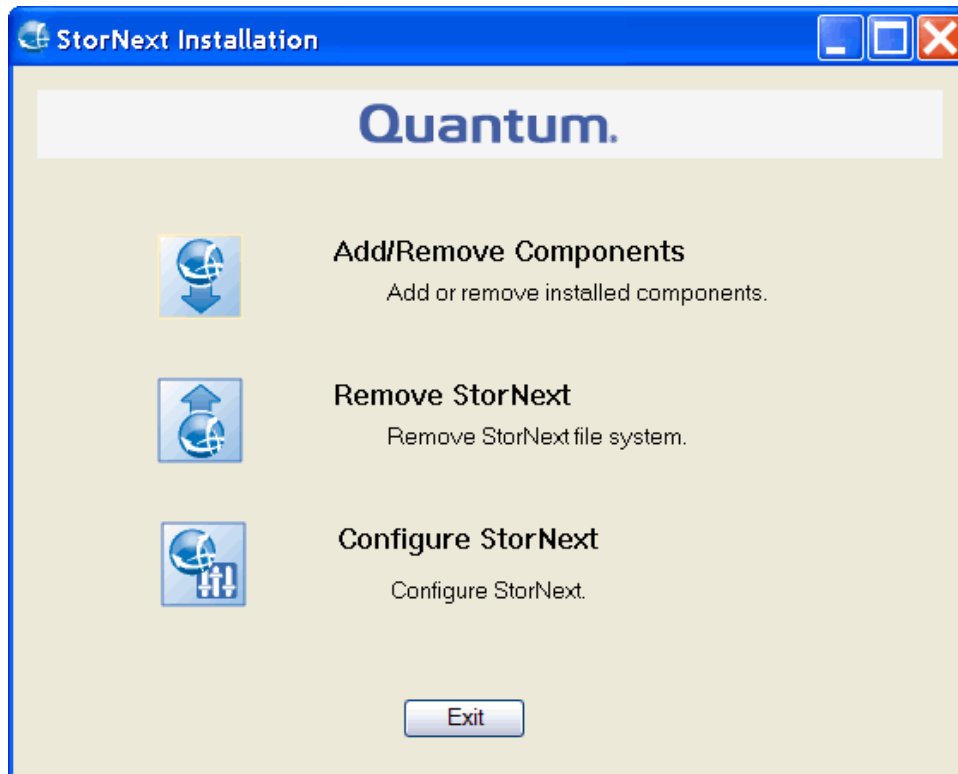
i Note: It may take several minutes for the installation to finish.



10. Click **Finish** to close the setup wizard. A dialog box opens letting you know that 30-day temporary licenses are installed and that you should install permanent licenses now. If you choose not to install them now, you can install them later using the Windows **Start** button (**Start > Licenses** or **Start > All Programs > StorNext Licenses**) or by following the instructions in [Add or Remove StorNext Components on page 138](#).



11. Click **Yes** to install permanent licenses now, or click **No** to install licenses later. If you choose to install now, complete the license installation screens that follow (for details, see [Configure a StorNext File System on Windows on the next page](#)). The **StorNext Installation** screen displays again with choices to **Add/Remove Components**, **Remove StorNext**, or **Configure StorNext**.



12. The installation is complete. Click one of the following options:

- **Add/Remove Components:** The components you selected during the install are already added, but if you changed your mind and want to select different components, you can do so now. See [Add or Remove StorNext Components on page 138](#) for details.
- **Remove StorNext:** Removes StorNext file system. For more information, see [Uninstall StorNext on Windows on page 195](#).
- **Configure StorNext:** Configures StorNext. For more information, see [Configure a StorNext File System on Windows below](#).
- **Exit:** Exits without making any further changes at this time (you can always perform configuration later).

Configure a StorNext File System on Windows

To configure a StorNext File System or the StorNext client on Windows operating systems, use the Windows configuration utilities. The configuration utilities let you set up a Windows-based Metadata Controller, configure a StorNext client, and work with StorNext file systems.

This section contains the required steps to configure StorNext, and also provides a description of all the other configuration utilities that StorNext provides.

Once you have installed StorNext, follow the steps below to configure StorNext.


Step 1: Obtain and Install a StorNext License

To obtain StorNext component and feature licenses, see [Obtain and Install StorNext Licenses – GUI-based Appliances and Customer-supplied MDCs](#). After following those steps, you will have a **license.dat** file, and then continue with the following procedure.


1. Copy the file to the **C:\Program Files\StorNext\config** directory. If there is a temporary license file, rename the file or move it to a backup location.

Note: To prevent accidentally losing a valid license, be sure to back up or rename any existing **license.dat** files.


2. Install StorNext File System (see [Install the StorNext File System on a Windows MDC: Interactively on page 73](#)). Licenses can be installed as part of the installation process or later if desired.
3. If installing licenses after the initial install, click **Start > All Programs > StorNext File System > Licenses**. The **StorNext License** screen displays.



4. Click **Import License File**. The **License Import File** window opens.



5. Browse to the license file location and click **Open**. The licenses are installed and a confirmation dialog box displays.
6. Click **OK**. The **StorNext License** screen displays again.
7. Close the window by clicking the X in the upper right corner. The **Installation Complete** dialog box displays.



8. Click **OK** in the **Installation Complete** dialog box.
9. Reboot the system if instructed to do so.

Step 2: Configure a File System Name Server

The purpose of a File System Name Server is to track the IP addresses of StorNext component servers, such as file system servers and gateways. StorNext SAN clients and LAN clients query name servers to locate appropriate services. The name server is a critical component for handling failover.

Every computer running StorNext components must have an `fsnameservers` file. The default location is **Program Files\StorNext\config\fsnameservers**.

At least one, but preferably two of the systems that have StorNext Server installed must be designated as a StorNext Name Server. The Name Server and Metadata Controller (MDC) can be the same machine.


-
- i Note:** The Name Servers configuration (including the ordering) must be identical on every computer running StorNext.

StorNext name servers can be added as either IP addresses or computer names, but when possible you should use IP addresses to avoid ambiguity.

Follow these steps to add a name server:

1. Access the name servers utility from the Windows Start menu (**Start > All Programs > StorNext File System > Name Servers**). The **StorNext Name Servers** window appears.

-
- i Note:** You may also access the name servers utility from the StorNext configuration window (see [Use the StorNext Configuration Window on page 114](#)).



2. Enter the IP address, and then click **Add**. The IP address will be added to the list on the left side of the window.
3. If desired, manipulate the order by selecting a name server and then clicking **Move Up** or **Move Down**. Name servers will be used in the order specified.

i Note: The Name Servers configuration (including the ordering) must be identical on all computers running StorNext.

4. To remove a name server, click the name server you want to remove and then click **Remove**. Confirm the deletion when prompted.
5. When you are finished adding name servers, click **OK**.

Step 3: Label a Disk Drive

i Note: This feature is not available for client-only configurations.

Use the **Disk Labeler** to configure StorNext File System storage area network (SAN) disks. Each disk that is to be used in the file system must have a StorNext File System Label. The Disk Labeler formats the disk so it can be used by StorNext. Use the Disk Labeler to create a list of disk labels, associated device names, and optionally the sectors to use.

Run the Disk Device Label Tool on a system that has visibility to all disk devices on the SAN.

Caution: Use extreme caution when labeling your storage area network disk devices. Inadvertently labeling a Windows system disk will cause irreparable harm to your system and/or loss of data! Only label disk devices that are intended for use with the StorNext File System SAN.

The file system uses the volume labels to determine which disk drives to use. The label name written to a disk device must match the disk name specified in the Server Configuration utility (see [Step 4: Configure a Server \(Advanced\) on page 89](#)).

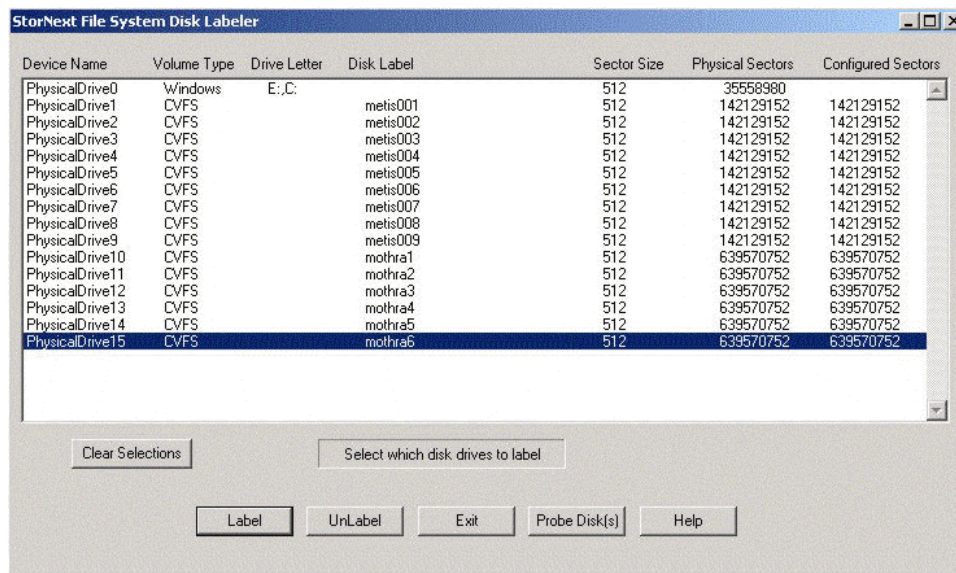
Follow the steps below to label a disk drive:

1. Access the disk labeler from the Windows **Start** menu (**Start > All Programs > StorNext File System > Disk Labeler**).

Note: You may also access the disk labeler from the StorNext configuration window (see [Use the StorNext Configuration Window on page 114](#)).

The **StorNext File System Disk Labeler** window appears. All disk devices are scanned for labels, and those devices with labels are displayed.

Note: The label type **NotVerified** is used for disks that are unusable or represent a secondary (passive) path to a disk that appears elsewhere in the list. The label type **TooLarge** is used for disks that have valid StorNext labels, but are not usable on this system without re-labeling, because the system SCSI driver cannot access disk blocks beyond block number 4294967295.




2. Do one of the following:
 - Click **Label** to write a volume label to the selected device(s) (see [Label a Disk below](#)).
 - Click **UnLabel** to remove the existing volume label from the selected device(s).
 - Click **Exit** to quit the Disk Device Labeler utility.
 - Click **Probe Disk(s)** to pulse the activity light of the selected device(s). Use this feature to correctly identify disks before labeling them.
 - Click **Clear Selections** to deselect all devices in the list.

Label a Disk

Follow the steps below to label a disk.

1. Select the disks you wish to label. If desired, verify that the disks selected are the correct ones by pressing the **Probe Disk(s)** button. The activity light on selected disks will pulse, allowing visual verification that the correct disk drives are selected. This procedure will continue until you click the **Stop** button.
2. After you have verified that the correct disk devices are selected, press the **Label** button. A confirmation screen appears asking if you are sure you want to proceed.
3. Click **OK** to continue.
4. A properties screen for each disk appears, one at a time. This screen allows you to name the disk and assign the configured number of sectors to use on the disk.



5. Enter the following fields:

- **New Disk Label:** The label you want to give the disk.
- **New Sectors:** The number in 512-byte sectors you will use in the **DiskType** configuration during file system configuration.

Normally you will use the default number of new sectors provided by the tool. However, some disks must be configured smaller than their actual size. For example, MPIRE video disks must be under-configured to eliminate using the last zone of the disk.

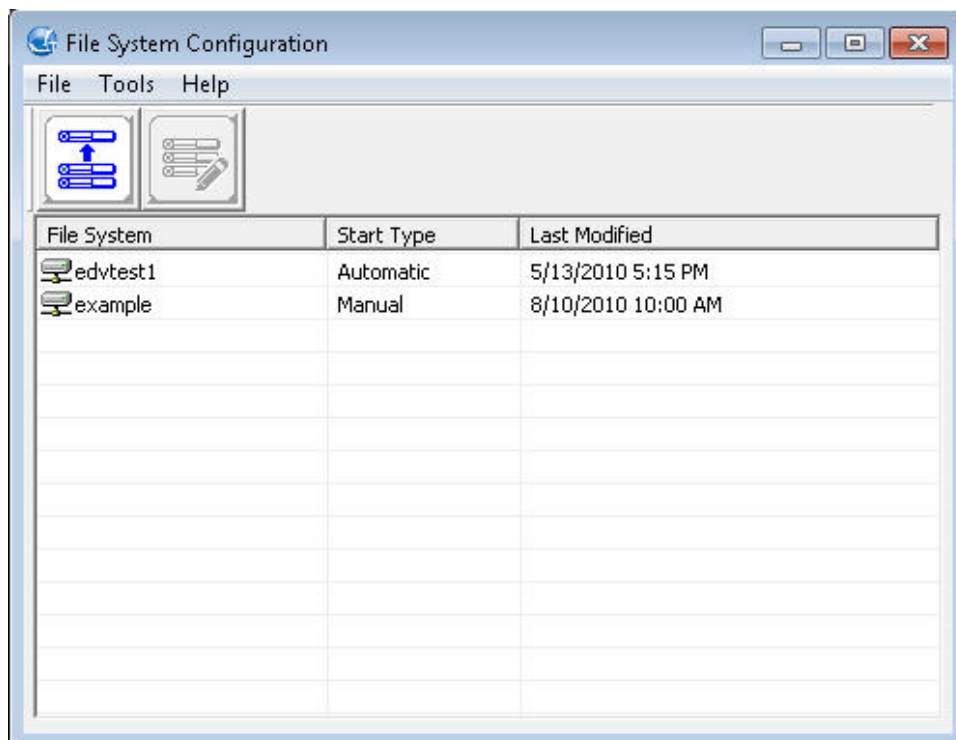
6. Click **Create Label** to commit the label to the disk device. If more than one disk was selected, a properties screen for the next device will appear until all of the selected devices are completed. If you click **Skip Disk**, then you proceed to the next selected disk.

Step 4: Configure a Server (Advanced)

Note: Not available for client-only configurations.

The **Advanced Server Configuration** utility lets you view and modify properties for an MDC. Use this utility to create a new server configuration or modify an existing configuration. This tool provides access to advanced features of StorNext File Systems. It is assumed that you are familiar with individual features available in the File System Cfg File (.cfg).

1. Access the advanced server configuration utility from the Windows Start menu **Start > All Programs > StorNext File System > File System Cfg (Advanced)**. The File System Configuration window appears.



2. Do one of the following:

- Click **Create New Configuration** to create a new configuration file.
- Select an existing file and click **Open Existing Configuration** to edit an existing configuration file.

The **Properties** window appears. This window includes four tabs:

Global Settings

Use the **Global Settings** parameters to specify general file system properties. The table below describes the fields on the **Global Settings** tab.

The screenshot shows the 'Properties' dialog box with the 'Global Settings' tab selected. The dialog is divided into several sections:

- Disk:** Allocation Strategy (Round), Inode Stripe Width (0 MBs), Allocation Session Reservation Size (0 MBs), Journal Size (64 MBs), Max Debug Log Size (4 MBs).
- Memory:** Inode Cache Size (131072 Entries), Buffer Cache Size (256 MBs).
- Performance:** Cluster File Locking (unchecked).
- Security:** Global Superuser (checked), Windows Security (checked), Use Active Directory SFU (checked), Global ShareMode (unchecked).
- XSAN Client Behavior:** Xsan Named Streams (unchecked), Enforce ACLs (unchecked).
- Default Unix User Mapping:** "Nobody" UID (60001), "Nobody" GID (60001), Default File Mode (0644), Default Directory Mode (0755), Use Fabricated IDs (unchecked).
- Quotas:** Use Quotas (unchecked), Duration of Quota Logs (7 Days).
- File Size Dynamic Expansion:** Minimum (0 Blocks), Increment (0 Blocks), Maximum (0 Blocks). A warning message states: "WARNING: 'Dynamic Expansion' parameters are deprecated and will be removed from a future release".

Buttons at the bottom: OK, Cancel, Apply, Help.

| Field / Button | Description |
|---|---|
| <p>Disk</p> <p>Allocation Strategy</p> | <p>Select the method of allocation for new disk blocks:</p> <ul style="list-style-type: none"> • Round: Alternate between similar stripe groups. (Default) • Fill: Use the stripe group with the least amount of free space (that can accommodate the file) before using the stripe group with the next smallest amount of free space. • Balance: Use the stripe group with the most free space. |
| <p>Inode Stripe Width</p> | <p>(Optional) Specify a value in MB. This feature defines how a file is striped across the file system's data stripe groups. See the Windows on-line help for additional information.</p> |
| <p>Allocation Session Reservation Size</p> | <p>(Optional) Specify the size (in MB) this feature should use when allocating segments for a session. This feature allows a file system to benefit from optimized allocation behavior for certain rich media streaming applications, and potentially other workloads. The feature also focuses on reducing free space fragmentation. This feature is disabled by default; a value of 0 is the default value, which means the feature is turned off. When enabled, the value can range from 128 MB (134217728) to 1 TB (1099511627776); the largest value would indicate segments are 1 TB in size, which is extremely large.</p> |
| <p>Journal Size</p> | <p>Type the maximum size (in MB) for the journal.</p> |
| <p>Max Debug Log Size</p> | <p>Type the maximum size (in MB) for the debug log.</p> |


| Field / Button | Description |
|---|--|
| Memory Inode Cache Size Buffer Cache Size | Type the number of entries in the inode cache. Type the amount of memory (in MB) used for general metadata information caching. |
| Performance Cluster File Locking | (Optional) Select to send a warning when a user attempts to access a file that is locked or already in use. |
| Security Global Superuser Windows Security Use Active Directory SFU Global ShareMode | (Optional) Select to allow a user with superuser privileges to have global access rights on the file system. (Optional) Select to use Windows Security Reference Monitor (ACLs) on Windows clients. (Optional) Select to enable or disable the use of Microsoft's Active Directory Services for UNIX (SFU) on Windows based SNFS clients. Always use Active Directory SFU in a mixed Windows/Unix environment, or if there is a possibility in the future of moving to a mixed environment. If Use Active Directory SFU is disabled, files created on Windows based SNFS clients will always have their UID and GID set to NOBODY with default permissions. (Optional) Select to enable or disable the enforcement of Windows Share Modes across StorNext clients. This feature is limited to StorNext clients running on Microsoft Windows platforms. See the Windows CreateFile documentation for the details on the behavior of share modes. When enabled, sharing violations will be detected between processes on different StorNext clients accessing the same file. Otherwise, sharing violations will only be detected between processes on the same system. |

| Field / Button | | Description |
|----------------------------------|------------------------|---|
| XSAN Client Behavior | Xsan Named Streams | <p>(Optional) Select to enable or disable support for Apple Named Streams.</p> <p>Named Streams are utilized by Apple Xsan clients. Enabling Named Streams support on a file system is a permanent change. It cannot be disabled once enabled.</p> <p>Only Apple Xsan clients should be used with named streams enabled file systems. Use of clients other than Apple Xsan may result in loss of named streams data.</p> |
| | Enforce ACLs | <p>(Optional) Select to enable Access Control List enforcement on XSan clients. On non-XSan MDCs, Windows Security should also be enabled for this feature to work with XSan clients.</p> |
| Default Unix User Mapping | "Nobody" UID | Type the Windows user ID to use when no user ID can be found using Active Directory. |
| | "Nobody" GID | Type the Windows group ID to use when no group ID can be found using Active Directory. |
| | Default File Mode | Type the permission mode to use when creating files. |
| | Default Directory Mode | Type the permission mode to use when creating directories. |
| | Use Fabricated IDs | Select if the Metadata Controller is running on Apple macOS X. |
| Quotas | Use Quotas | <p>(Optional) Select to enforce file system quotas and track storage usage for individual users and groups.</p> |
| | Duration of Quota Logs | Type the duration (in days) of the system quota logs. |

| Field / Button | | Description |
|--|-----------|---|
| File Size Dynamic Expansion i Note: File Size Dynamic Expansion is calculated on a file-by-file basis for better performance. If this is not working optimally, you can use these settings to override the file-by-file calculations. | Minimum | (Deprecated) Type the minimum number (in blocks) to allocate when a file is created. |
| | Increment | (Deprecated) Type the initial increment (in blocks) by which a file is expanded. |
| | Maximum | (Deprecated) Type the maximum increment (in blocks) by which a file is expanded. |

Disk Types

The **Disk Types** parameters define disk types used in the file system.

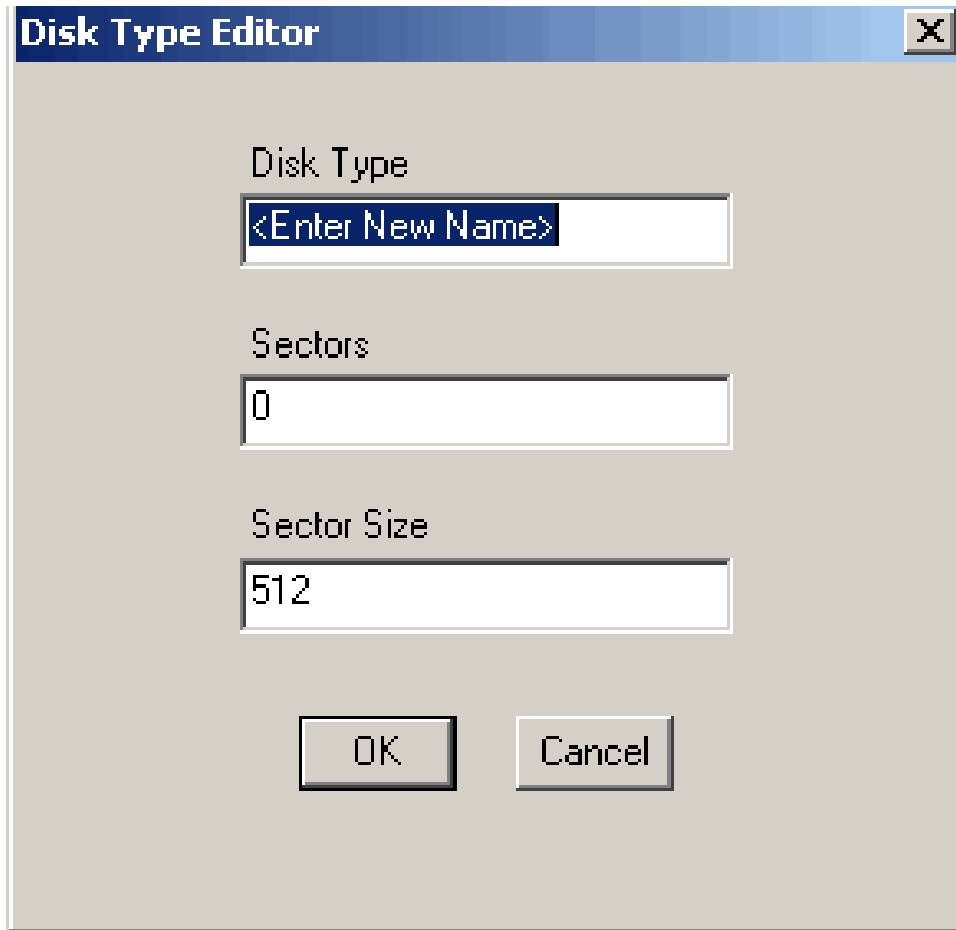


On the **Disk Types** tab, do one of the following:

- Click **New** to add a new disk type.
- Select a disk type in the list and click **Edit** to modify it.
- Select a disk type in the list and click **Delete** to delete it.

When you add or modify a disk type, the **Disk Type Editor** dialog box appears. The table below describes the fields on the **Enter New Disk Type** dialog box.

i Note: The **Sectors** and **Sector Size** fields are populated with values from the Disk Labeler utility. For more information, see [Step 3: Label a Disk Drive on page 86](#).



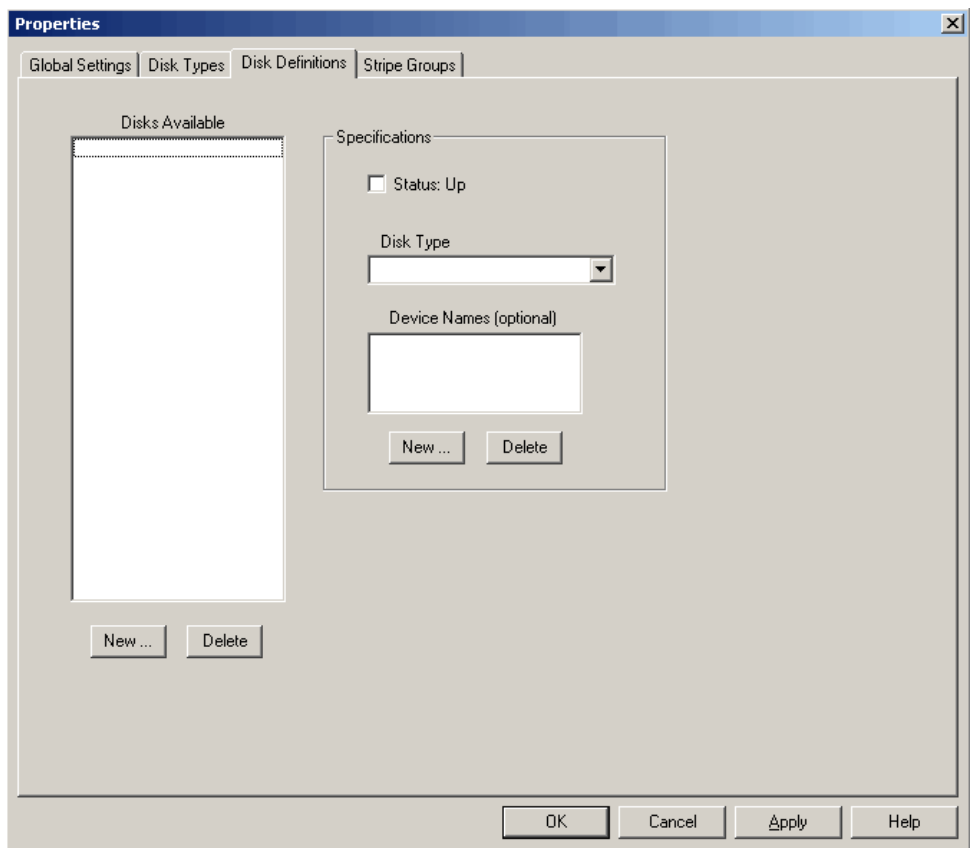
The screenshot shows a Windows-style dialog box titled "Disk Type Editor". It has a close button (X) in the top right corner. The dialog contains three text input fields. The first field is labeled "Disk Type" and contains the text "<Enter New Name>". The second field is labeled "Sectors" and contains the value "0". The third field is labeled "Sector Size" and contains the value "512". At the bottom of the dialog are two buttons: "OK" and "Cancel".

| Field / Button | Description |
|------------------|---------------------------------------|
| Disk Type | Type a descriptive name for the disk. |

| Field / Button | Description |
|--------------------|---|
| Sectors | Type the number of physical sectors on the disk. |
| Sector Size | Type the size of sectors on the disk. |
| OK | Click to create the new disk type with the specified properties. |
| Cancel | Click to close the Disk Type Editor dialog box without saving changes. |

Disk Definitions

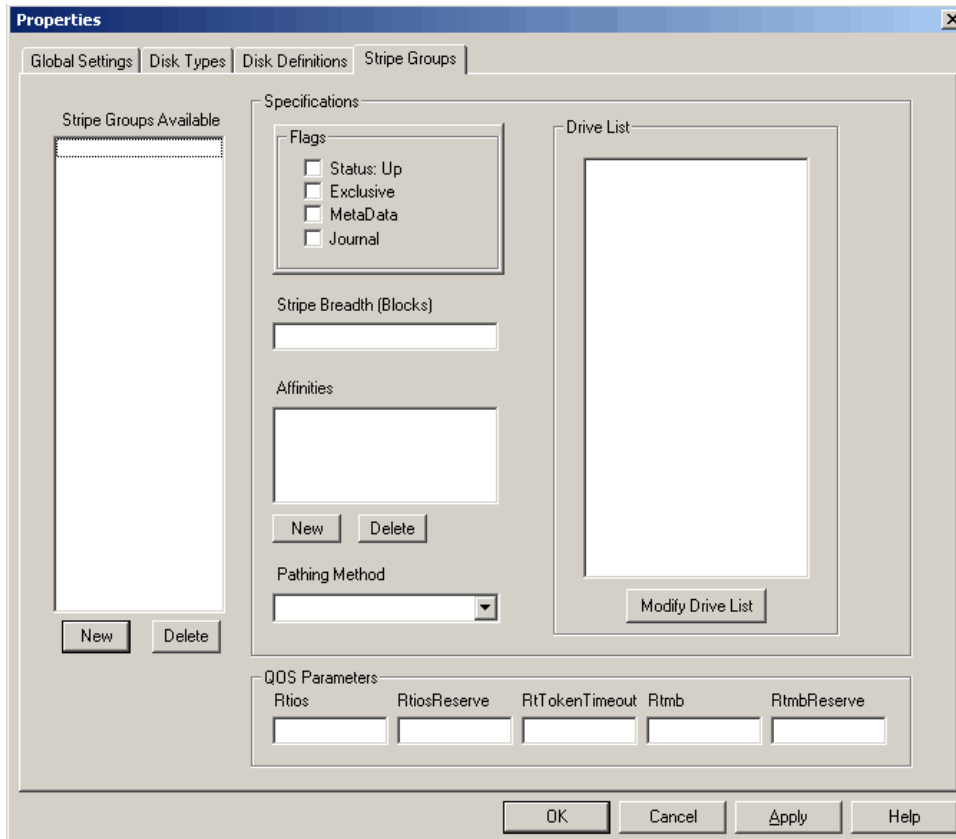
The **Disk Definitions** parameters create disk definitions and modify disk specifications. The table below describes the fields on the **Disk Definitions** tab.




| Field / Button | Description |
|---------------------------------|---|
| Disks Available - New | Click to create a new disk definition. Type a name for the disk and click OK . |
| Disks Available - Delete | Click to delete the selected disk definition. |
| Status: Up | Select to enable the selected disk definition. When enabled, the defined disk is available to the configured file system for I/O. |
| Disk Type | Select a disk type for the selected disk definition. (Disk types are defined on the Disk Types tab.) |
| Device Names - New | Click to add a device name to the disk definition. Type a device name and click OK . |
| Device Names - Delete | Click to delete the selected device name. |


Stripe Groups

The **Stripe Groups** define stripe groups. A stripe group is a logical storage unit consisting of one or more disk drives. The table below describes the fields on the **Stripe Groups** tab.




| Field / Button | Description |
|---|--|
| Stripe Groups Available - New | Click to create a new stripe group. Type a name for the stripe group and click OK . |
| Stripe Groups Available - Delete | Click to delete the selected stripe group. |
| Status: Up | Select to enable the selected stripe group. When enabled, the stripe group is available to the configured file system for I/O. |
| Exclusive | Select to allow only file types associated with the affinities specified under Affinities to be stored on the stripe group. |
| MetaData | Select to allow file system metadata to be stored on the stripe group. |
| Journal | Select to allow the file system logging journal to be stored on the stripe group. |

| Field / Button | Description |
|--------------------------------|---|
| Stripe Breadth (Blocks) | Describes the number of file system blocks that are appended to a file before switching to the next disk in the stripe group. |
| Affinities - New | Click to add a new affinity to the stripe group. Type the name of the new affinity and click OK .  Note: Affinity names cannot be longer than eight characters. |
| Affinities - Delete | Click to delete the selected affinity. |
| Pathing Method | Select a method for how the file system accesses the disk: <ul style="list-style-type: none"> • Rotate: Chooses an I/O path based on load. • Static: Always uses the same I/O path. • Sticky: Chooses an I/O path based on the identity of the target file. |
| Rtios | Type the maximum number of disk I/O operations per second available to real-time applications using the Quality of Service (QOS) API. |
| RtiosReserve | Type the minimum number of disk I/O operations per second reserved for non-real-time applications. |
| RtTokenTimeout | Type the number of seconds to wait for clients to respond to a token callback. |
| Rtmb | Type the maximum number of MBs per second available to real-time applications using the Quality of Service (QOS) API. |
| RtmbReserve | Type the minimum number of MBs per second reserved for non-real-time applications. |
| Modify Drive List | Click to add or remove disks for the selected stripe group. Use the left and right arrows to add or remove disks in the Primary Disks list, and then click OK . |

 **Caution:** When adding a disk to an existing stripe group, make sure the disk is added to the bottom of the **Drive List**. If the disk is inserted in the middle of the list, the file system behaves as if files on the file system are corrupted.

3. Click a tab to view or modify StorNext File System properties. After making changes on one or more tabs, click **OK** to save the changes. Specify a name and a location for the *.cfg file, and then click **Save**.

 **Note:** Before configuring a file system, you should label disk devices. See [Step 3: Label a Disk Drive on page 86](#).

4. **Import and Export Client Settings.** This feature is described in [Restore a Previous StorNext Client Configuration on page 174](#).

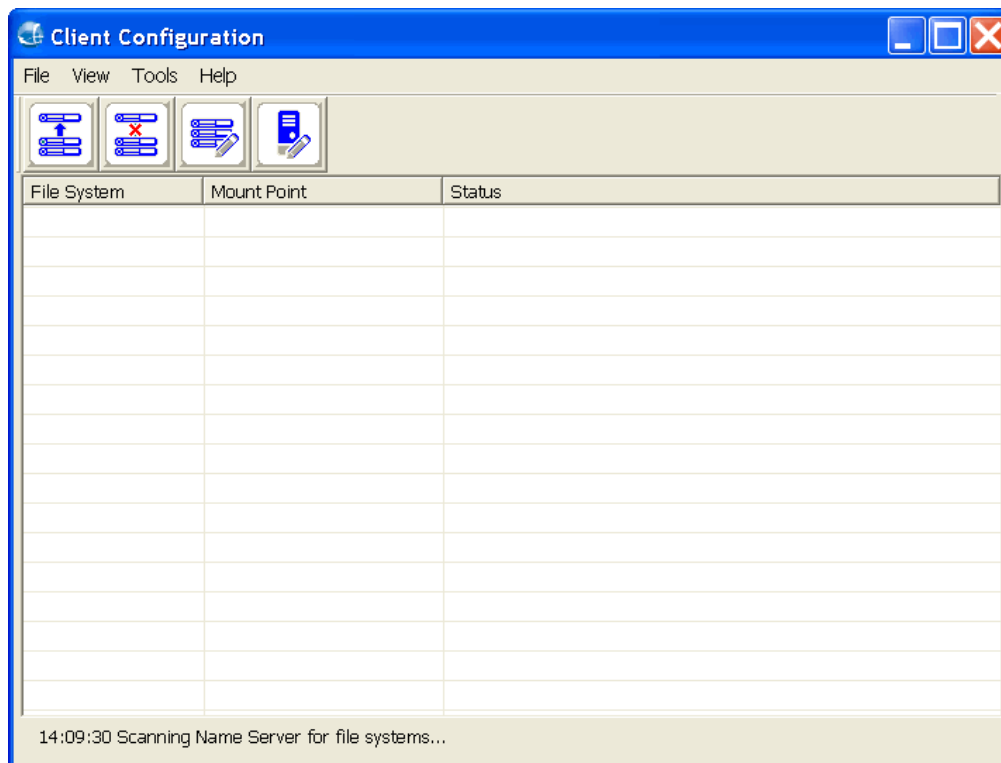
Step 5: Configure a Client

In order for a client to access a StorNext file system, the file system must be mapped to a folder or drive letter. The **Client Configuration Tool** maps a StorNext file system to a Windows drive letter or folder and controls various configuration settings.

The values for a Client Configuration are stored in the Windows Registry and are referenced by StorNext when the file system is mounted.

Access the client configuration utility from the Windows Start menu **Start > All Programs > StorNext File System > Client Configuration**. The **Client Configuration** dialog box appears. Displayed under the **File System** column heading are all file systems advertised by the name server, and currently mapped StorNext file systems found in the Windows registry.

i Note: You may also access the client configuration utility from the StorNext configuration window (see [Use the StorNext Configuration Window on page 114](#)).



Review the sections below to configure a client. The **Client Configuration** dialog box includes the following functions.

Map a Drive

To map a StorNext file system to a folder or drive letter, double-click the desired file system (or select it and click the **Map Drive** icon or **Tools >Map Drive**). The **Client Properties** screen appears.

Use the **Mount Point** tab to specify drive mappings. The table below describes the fields on the **Mount Point** tab.

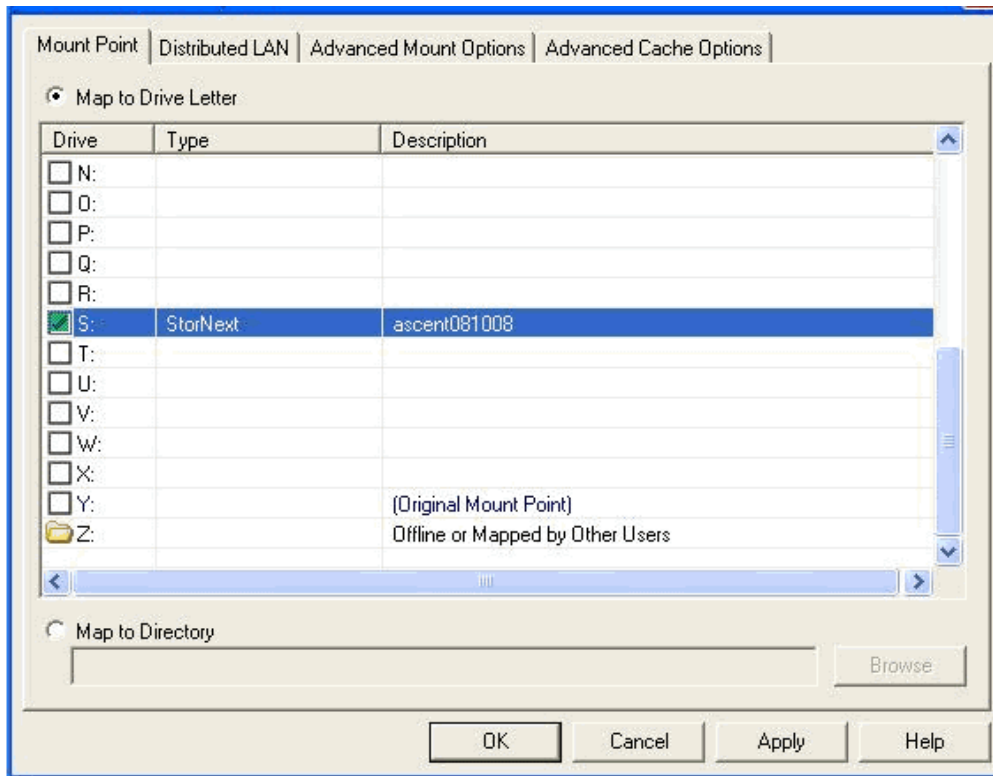
If your file systems are not displayed, check the status bar at the bottom. The most common problems are:

- An incorrect Name Server Configuration.
- Network problems.
- The file system is not running.
- The StorNext Services are not running.

StorNext file systems can be mapped to a drive letter or to a folder. When mapping to a folder, the following conditions must be met:

- The folder must be empty or non-existent.
- The folder must exist on an NTFS file system volume, for example: C : \Mount\MySnfs.
- The NTFS volume must be a local volume. In other words, StorNext cannot be mounted in a folder on a Cluster Shared Volume.

i Note: Verify the **fsnameservers** configuration is the same for all clients on the SAN. An incorrect or incomplete **fsnameservers** configuration may prevent the client from connecting to the file system.



| Field / Button | Description |
|----------------------------|--|
| Map to Drive Letter | Select Map to Drive Letter to map the drive to a specific drive letter. |
| Drive (Letter) | Select the desired drive letter. |
| Map to Directory | Instead of selecting Map to Drive Letter , select Map to Directory . Enter a directory path, or click Browse to navigate to the directory path. |

i Note: If this is a LAN client or gateway, be sure to enter the parameters on the **Distributed LAN** tab (see [Configure a LAN Client and a Gateway Server on the next page](#)).

Unmap a Drive

Select the file system you want to unmap and then click the **Unmap Drive** icon or select **Unmap Drive** from the **Tools** menu. You must close all applications (included Windows Explorer) which might be using the file system. If the file system is busy, you will not be allowed to unmap the drive.

i Note: In order to unmap a drive, you must have Administrator rights on the local computer, and you must have read/write access to root of the StorNext file system.

Edit a Drive Map

Once a file system is mapped, double-click the file system to change the map. Alternatively, select the file system and click the **Properties** button, or select the file system and click **Tools > Edit Drive Mapping**. This opens the **Properties** window where you can change drive mapping. See [Map a Drive on page 100](#).

Configure a File System Name Server


Click **Edit Name Servers** icon or choose **Configure Name Servers** from the **Tools** menu to edit, add or remove a file system name server. See [Step 2: Configure a File System Name Server on page 85](#).

Configure a LAN Client and a Gateway Server

To configure a StorNext LAN client and a gateway server, double-click the desired file system (or select it and click the **Properties** icon). The **Properties** dialog box appears.

i Note: Do not enable gateway functionality for MDCs, connected to Lattus systems.

Use the **Distributed LAN** tab to enable and configure StorNext LAN clients and gateway servers for a file system.



The **Enable Client** and **Enable Server** radio buttons determine whether the system is a LAN client or gateway Server. Note that separate licensing is required for LAN clients.

StorNext LAN clients provide the same functionality as StorNext SAN clients but do not need to be connected directly to the SAN. They instead send I/O via the TCP/IP network through a gateway server - a StorNext SAN client enabled as a gateway server. Metadata is sent over the TCP/IP network directly to the File System Metadata Controller (Server/MDC).

LAN Clients request the location of gateway servers via the Name Server.

- **Gateway Server:** A gateway server has direct access to data disks (using fibre channel) and makes the shared storage pool available to distributed LAN clients over a network.
- **LAN Client:** A StorNext LAN client does not have direct access to data disks and instead accesses the shared storage pool by connecting to a gateway server over a network. (LAN clients have full StorNext client functionality.)


i Note: A StorNext SAN client configured as a gateway server requires a dedicated system. Do not run other applications or services on a system configured as a gateway server.

The table below describes the fields on the **Distributed LAN** tab.

| Field / Button | Description |
|-----------------------------------|--|
| Disable | Select this option to disable the StorNext client as a StorNext LAN client or gateway server. |
| Enable Client | Select to configure the StorNext client as a LAN client. The StorNext client will connect to all available gateway servers. (If you select this check box, all other fields on the tab become unavailable.) |
| Enable Server | Select to configure the StorNext SAN client as a gateway server. The StorNext client will make the disks for this file system available to LAN clients over the specified LAN. |
| Read Timeout (sec) | Specify how long the LAN client will wait for a read operation to complete before retrying it on a different gateway server. |
| Write Timeout (sec) | Specify how long the LAN client will wait for a write operation to complete before retrying it on a different gateway server. |
| Number of Transfer Buffers | Specify the number of socket transfer buffers used for LAN client I/O. |
| Transfer Buffer Size (kb) | Specify the size in kilobytes of the socket transfer buffers used for LAN client I/O. |
| TCP Window Size (kb) | Specify the size in kilobytes of the TCP window used for LAN client I/O connections. |
| Max Threads | Specify the number of kernel threads that will be used to service LAN client I/O requests. |
| Server Interfaces | To complete gateway server configuration, select one or more network interfaces to be used for LAN client traffic. |

Configure Logging Options

From the **Client Configuration** dialog box, click **Tools Global Options**. The **Global Client Options** window appears.



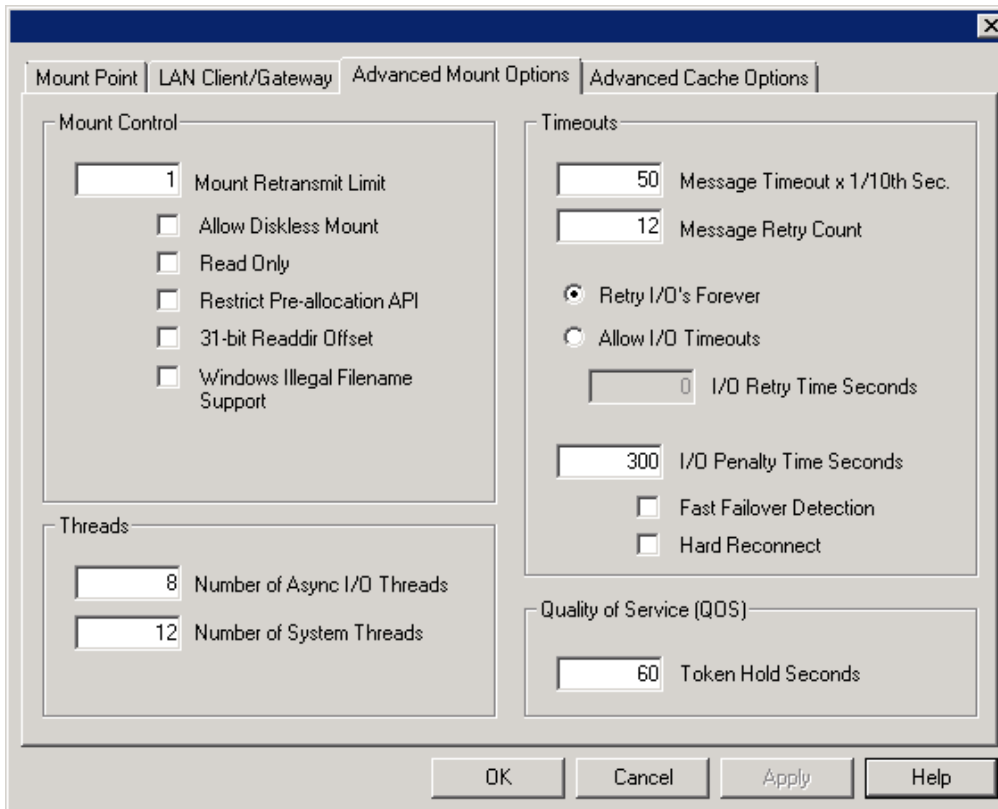
Select the level of logging you want to receive and whether you want to display critical events on your desktop. Log information is available in the **Windows Event Viewer**.

| Field / Button | Description |
|--|--|
| Event Log Level | Controls how many messages are inserted into the Windows Event Viewer. Critical provides the fewest messages; Debug provides the most. Each level includes all the messages in the levels listed above it. |
| Display Critical Events to Desk Top | Enables a notification of critical events to be sent to work station's desktop in a message box. If this is not desired, clear the check box. |

Advanced Mount Options

The **Mount Options** tab displays the mount options for the drive currently selected on the **Drive Mappings** tab. The table below describes the fields on the **Distributed LAN** tab.

Caution: Changing the values on the **Mount Options** tab can affect system performance and stability. Do not change mount options unless instructed to do so by the Quantum Technical Assistance Center.



| Field / Button | Description | |
|----------------------|---|---|
| Mount Control | Mount Retransmit Limit | Specify the number of times the driver will re-transmit a message to the FSM. After a successful connection, this value is no longer used. The default is 1. |
| | Allow Diskless Mount | When this box is checked, the file system can be mounted and accessed without all the disks being accessible in the file system stripe groups. In this mode, file metadata can be viewed and changed, but no data can be accessed. The default is to not allow a file system to be mounted unless all disks are accessible. |
| | Read Only | When this box is checked, the file system is mounted in read-only mode. The default is off or not checked. |
| | Restrict Pre-allocation API | When set to yes, non-administrator users are unable to use the preallocation ioctl. |
| | 31-bit Readdir Offset | When this option is enabled, you force the directory offsets to fit into 31 bits. i Note: Enable this option only when a problem has been identified with using the full size of the struct <code>d_off</code> field from <code>readdir(2)</code> or related system calls. |
| | Windows Illegal Filename Support (Previously: Apple SFM Filename Support) | If this box is checked, filenames that Windows considers <i>illegal</i> due to having certain reserved characters or having a space or period at the end will have the offending characters mapped to Unicode Private Use Area code points with the method used by the Microsoft Windows feature Apple Services For Macintosh®. i Note: Before you enable the Windows Illegal Filename Support option, read the section StorNext File Name Considerations in the StorNext 6 User's Guide to understand the system wide ramification of using this option. |
| Threads | Number of Async I/O Threads | Specify the number of threads created to be used as asynchronous I/O threads for user applications. The default is 8. The allowed range is 4 to 32. |
| | Number of System Threads | Specify the number of threads created for use by the file system. The default value is 16. |


| Field / Button | Description |
|---|--|
| Timeouts Message Timeout x 1/10th Sec. | Specify (in tenths of a second) the amount of time before a message to the FSM is timed out. The default is 50 (five seconds). |
| Message Retry Count | Specify the number of attempts to be made for sending a message to the File System Service. The default is 12. |
| Retry I/Os Forever or Allow I/O Timeouts | Choose either Retry I/Os Forever or Allow I/O Timeouts, depending on whether you want to allow timeouts or keep trying I/O. The default value is to Retry I/Os Forever. |
| I/O Retry Time Seconds | <p>When you choose Allow I/O Timouts, this option controls the I/O retry behavior. This parameter establishes the amount of time that may elapse during an I/O retry sequence.</p> <p>An I/O retry sequence consists of the following:</p> <ul style="list-style-type: none"> • Retry an I/O request across all available paths that are currently present. • Compare the current time to the Instantiation time of the I/O request, if at least <n> seconds have elapsed, return the I/O request in error, otherwise reset the paths used, and retry again. |
| I/O Penalty Time Seconds | <p>This option controls the Multi-Path I/O penalty value, where <n> is expressed in seconds with a minimum value of 1 and a maximum value of 15.</p> <p>This parameter establishes the amount of time a Path_In_Error will be bypassed in favor of an Operational Path during a Multi-Path selection sequence. If all paths are currently in the Path_In_Error state, the first available path will be selected, independent of the Path_In_Error state. The default value is 5 seconds.</p> |
| Fast Failover Detection | <p>When this box is checked, the client uses a more aggressive approach to detecting if the FSM has failed. With this option on, the client monitors the FSM heartbeat. If no heartbeat is detected within three (3) seconds, a failover request is initiated.</p> <p>This option is desirable for near realtime activities that cannot sustain more than 10 seconds of access loss to the server. Quantum does not recommend enabling this option for normal file system operations.</p> |

| Field / Button | | Description |
|---------------------------------|--------------------|---|
| | Hard Reconnect | When this box is checked, the file system attempts to reconnect to the FSM forever. If this option is not enabled, the file system attempts to reconnect to the FSM for the number of times specified at the Mount Retransmit Limit field before failing the request. The default value is off. |
| Quality of Service (QOS) | Token Hold Seconds | The QOS Token Hold Time parameter is applicable only when using the StorNext Quality of Service (QOS) feature for real-time IO. This parameter determines the number of seconds a client stripe group holds on to a non-realtime I/O token during periods of inactivity. If no I/O is performed on a stripe group within the specified number of seconds, the token is released back to the FSM. The default value is 60 seconds. The parameter should be specified in five-second increments; if the parameter is not a multiple of five, it will be rounded up automatically. |

Advanced Cache Options

The **Data Buffer Cache** keeps the file system data in memory on the local computer to speed up performance for small I/O operations.

The **Advanced Cache Options** tab displays performance values that control how many file system lookup names are kept in memory.

 **Caution:** Changing the values on the **Advanced Cache Options** tab can affect system performance and stability. Do not change cache parameters unless instructed to do so by the Quantum Technical Support.

Mount Point | Distributed LAN | Advanced Mount Options | **Advanced Cache Options**

Data Buffer Cache

Enable Data Buffer Cache

Individual Buffer Size

Shared by drives with the same Buffer Size:

Minimum Total Cache Size

Auto-DMA Read Size

Auto-DMA Write Size

Number of Read-ahead Buffers

Number of Threads

Attribute Flush Time

Seconds for Non-Shared

Seconds for Shared

Delay Atime Updates

Metadata/Inode Cache

Low Water Mark

High Water Mark

Max. Entries

Purge Period (Seconds)

Directory Cache

Shared by all drives

Minimum Directory Cache Size

OK Cancel Apply Help

| Field / Button | Description |
|--|--|
| Data Buffer Cache Enable Data Buffer Cache | When this box is not checked, the file system will not buffer any data. All files will be read directly into the application's memory using DMA. Requests that are not properly aligned will be read into temporary buffers, then copied to the application's buffer. If Buffer Cache is disabled, the remaining Data Buffer Cache options are ignored. |
| Individual Buffer Size | This option sets the size of each cache buffer. This determines the I/O transfer size used for the buffer cache. To determine optimal performance, try different sizes or contact your RAID vendor. This size must be divisible by the file system block size. The maximum value allowed is 512K and the minimum value allowed is the file system block size. The value default is 64K. |
| Minimum Total Cache Size | This value controls the amount of memory used to cache data from the file system. This parameter is shared by all file systems with the same block size. |
| Auto-DMA Read/Write Sizes | Buffer Cache Auto DMA Read Length and Buffer Cache Auto DMA Write Length control the minimum I/O size for DMA transfers. That is, any smaller transfer will always go through the buffer cache. The default value is 1MB. |
| Number of Read-ahead Buffers | This option controls the size of the read-ahead window used by the buffer cache. The default value is 16 buffers with a maximum of 100. Using a value of 0 disables read-ahead. |
| Number of Threads | This option controls the number of dedicated background threads used by the buffer cache for performing I/O. The default value is 8, with a minimum of 1 and a maximum of 100. Using a larger value can increase I/O concurrency in the buffer cache and therefore improve performance. |

| Field / Button | | Description |
|-----------------------------|-----------------------------------|--|
| Attribute Flush Time | Seconds for Shared/ Non-Shared | <p>The Attribute Flush Time parameters control the length of time attributes are held on the client before they are flushed back to the FSM (unless there is an explicit attribute request from the FSM).</p> <p>The default values are 30 seconds for non-shared files (Attribute Flush Time, non shared), and 2 seconds for shared files (Attribute Flush Time, shared).</p> <p>Setting these parameters lower will greatly increase the amount of metadata traffic between clients and the FSM. However, in circumstances where files are frequently shared on the SAN, setting the Attribute Flush Time Shared can result in other clients seeing size updates more frequently if a file is being written on one client and read on another.</p> <p>A value of zero is invalid, and will result in using the default setting.</p> |
| | Delay Atime Updates | <p>When this box is checked, the file system delays Atime (access time) updates when reading a file to the time when the file is closed. This cuts down on FSM metadata updates at the expense of coherency.</p> |
| Metadata/Inode Cache | Low Water Mark | <p>A metadata/inode cache contains non-data information such as file names, sizes, update times, and so on.</p> |
| | High Water Mark | |
| | Max. Entries | <p>Tuning the low and high water marks and the frequency of purge passes can help certain large mix applications.</p> |
| | Purge Period (Seconds) | |
| Directory Cache | Minimum Directory Cache Size | <p>This option sets the size of the directory cache. Directory entries are cached on the client to reduce client-FSM communications during directory reads.</p> <p>The default value is 10 MB.</p> |

Additional Windows configuration utilities are available. See below for details.

Access the Windows Configuration Utilities

There are two ways to access the configuration utilities.

Use the Windows Start Menu

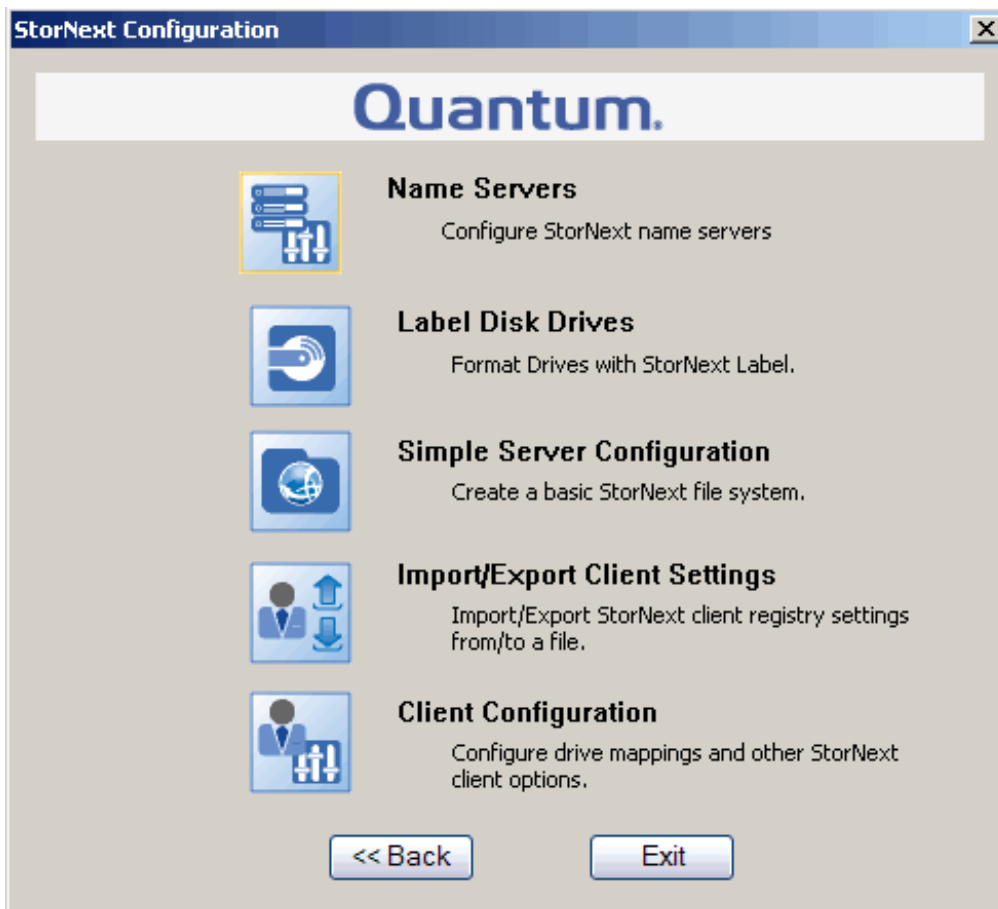
You can access all of the configuration utilities, along with several others, from the Windows **Start** menu.

To run a configuration utility from the Windows **Start** menu, click **Start > All Programs > StorNext File System**, and then click one of the available configuration utilities (see [Available Configuration Utilities on the next page](#)).

Use the StorNext Configuration Window

As soon as you finish installing StorNext, the StorNext Installation window remains open. If you wish to begin configuring immediately, click the **Configure StorNext** button to open the StorNext Configuration window.

You can return to the StorNext Configuration window at any time by double-clicking the file named **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems) and then clicking the **Configure StorNext** button.



The StorNext Configuration window provides buttons to access several of the most common configuration utilities:

- Name Servers, see [Step 2: Configure a File System Name Server on page 85](#).
- Label Disk Drives, see [Step 3: Label a Disk Drive on page 86](#).

- Import/Export Client Settings, see [Configure the StorNext Client Software on Windows on page 169](#).
- Client Configuration, see [Step 5: Configure a Client on page 100](#).

Available Configuration Utilities

The following configuration utilities are accessible from the Windows **Start** menu **Start > All Programs > StorNext File System**.

For more information on using these utilities, see the online help accessible from **Start > All Programs > StorNext File System > StorNext Help**.

Configure a Client

Use this utility to view and modify properties for the StorNext client software. See [Step 5: Configure a Client on page 100](#).

Label a Disk Drive

Formats a hard disk so it can be used by StorNext. See [Step 3: Label a Disk Drive on page 86](#).

Configure a Server (Advanced)

Allows you to configure StorNext file systems. This tool provides access to advanced features of StorNext File Systems. It is assumed that you are familiar with individual features available in the File System Cfg File (.cfg). See [Step 4: Configure a Server \(Advanced\) on page 89](#).

The File System Startup List Utility

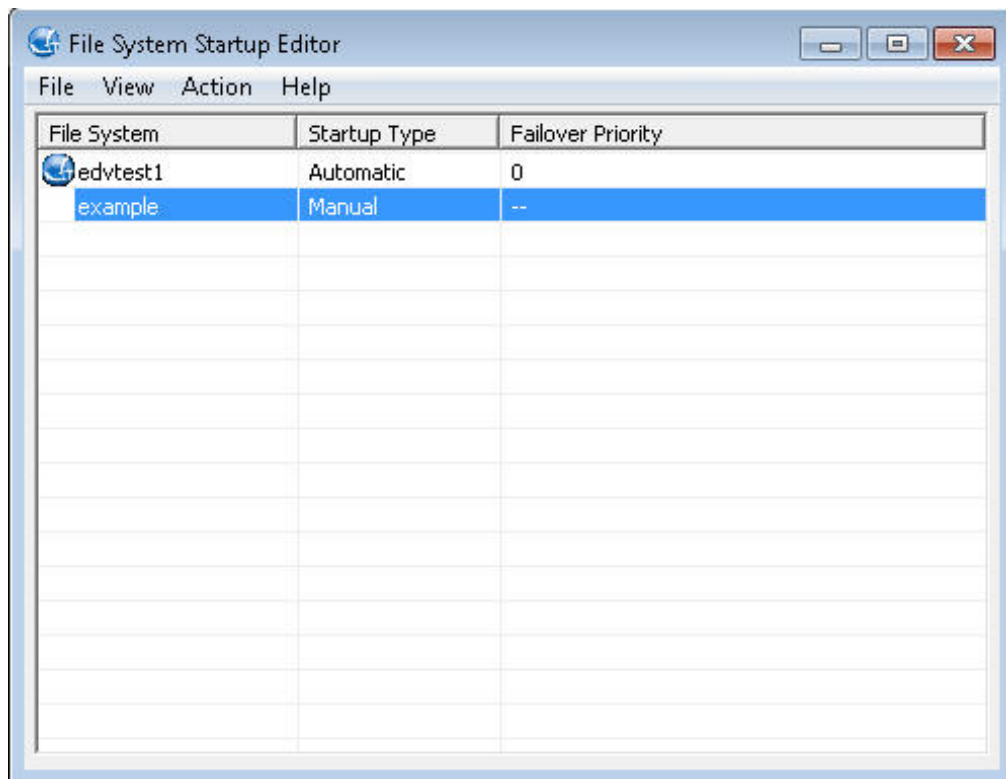
The **File System Startup List** utility lets you modify the File System Manager (FSM) services list and set file system priority. You can check the StorNext file system health status while the file system is active. You can note changes but not commit them in read-only mode. You can also list all of the existing files and their pertinent statistics, such as inode number, size, file type and location in the file system

The File System Manager is a process that manages the name space, allocations, and metadata coherency for a file system. Each file system uses its own FSM process. When there are multiple file systems (and therefore multiple FSM processes), the FSM services list controls which FSM processes are run when the server starts up, and also sets the priority for each file system (for failover configurations).

To run the File System Startup List utility, on the Windows **Start** menu, click **All Programs > StorNext File System > File System Startup List**. The **File System Startup Editor** appears.

Select a file system do any of the following. When finished making changes, select **File > Save** to save your changes.

| Action | Description |
|-----------------------------------|---|
| Action Toggle Startup Type | Allows you to choose if the file system will start up automatically upon server startup. This action toggles the startup type between automatic and manual. |
| Action Priority | (Optional) Select the priority for the selected file system. This priority is used for failover configurations. 0 is highest priority and 9 is lowest priority. |



The File System Check (Repair) Utility

The **File System Check (Repair)** utility allows you to repair corrupted metadata on a file system. You can also list all of the existing files and their pertinent statistics, such as inode number, size, file type and location in the file system. Repair a file system if errors were identified when checking the file system (see [The File System Check \(Read-Only\) Utility on the next page](#)).

The file system must be inactive in order to be repaired. To stop a file system, see [The Stop File System Services Utility on page 121](#).

To repair a file system, on the Windows **Start** menu, click **All Programs > StorNext File System > File System Check (Repair)**. Type the number that corresponds to the file system you want to repair, and then press **Enter**.

i Note: Multiple repair passes may be needed to fully fix a damaged file system. Run the repair utility repeatedly until no problems are reported by the utility.

The File System Check (Read-Only) Utility

The **File System Check (Read-Only)** utility allows you to check a StorNext file system for metadata corruption (due to a system crash, bad disk, or other failure). Run the utility on an MDC that contains the file system you want to check.

To check a file system, on the Windows **Start** menu, click **All Programs > StorNext File System > File System Check (Read-Only)**. Type the number that corresponds to the file system you want to check, and then press **Enter**.

Because the file system check is run in read-only mode, any problems that exist are not repaired. If the utility identifies errors or corruption in metadata, you must repair the file system (see [The File System Check \(Repair\) Utility on the previous page](#)).

The Re-initialize a File System Utility

The Re-initialize a File System utility allows you to re-initialize an existing file system or initialize a file system for the first time. Initializing a file system prepares it for use.

⚠ Caution: Use extreme caution when using this tool. This will destroy all data on the selected file system.

To initialize a file system, on the Windows **Start** menu, click **All Programs > StorNext File System > Re-initialize a File System**. Type the number that corresponds to the file system you want to re-initialize, and then press **Enter**.

The License Identifier Utility

Use the **License Identifier** utility to display the host license identifier. The host license identifier is required to obtain a permanent license for StorNext.

To run the License Identifier utility, on the Windows **Start** menu, click **All Programs > StorNext File System > License Identifier**. A dialog box displays the host license identifier. Record this information.

Obtain and Install a StorNext License

Allows you to import StorNext licenses. See [Step 1: Obtain and Install a StorNext License on page 83](#).

Edit a Drive Map


Specifies the name server(s) used by StorNext File System services. See [Edit a Drive Map on page 103](#).

The Server Administration Utility

The **Server Administration** utility allows you to look at the current status of file systems and to start, stop and perform basic file system administration tasks (similar to the CLI command, **cvadmin**).

i Note: A stripe group is a logical storage unit made up of one or more disks. A quota is a space limit that is set for specified users or groups.

To run the Server Administration utility, on the Windows Start menu, click **All Programs > StorNext File System > Server Administration**. The **StorNext Administration** window appears.



In the **StorNext Administration** window, you can perform the following procedures.

Connect to an MDC

1. Click **Connections**, and then click **Actions > New**.
2. Type the host name and click **OK**. **Disks** and **File Systems** are displayed in the left pane.

Start a File System

1. Right-click **File Systems**, and then click **Start**.
2. Enter file system information. For additional details, see [Start a File System on page 122](#).

Activate a File System

1. Click **File Systems** in the left pane.
2. Right-click the name of the file system under **File Systems**, and then click **Activate**.

Stop a File System

1. Click **File Systems** in the left pane.
2. Right-click the name of the file system under **File Systems**, and then click **Stop**.


Update the List of File Systems

Press **Ctrl+R**, or click **View**, and then click **Refresh**.

Create a Connection


You can create a connection to any StorNext client or server.

1. To create a new connection, click **Actions**, and then click **New**. The **New Connection** dialog box appears.



2. In the **New Connection** dialog, enter the host name or IP address of the computer to add, and then click **OK**. The new connection information appears. The left pane displays file systems running on the connected MDC.


Note: Expand a file system to see stripe groups, quotas, and additional file system properties.



Close a Connection

You can close a connection to any StorNext client or server.

To close an existing connection, right-click the connection, and then click **Close**.



Save a Connection

You can save the current connections. The next time you open the StorNext Administration utility, the utility will query the saved list to automatically open connections.

To save the list, click **Connections** (in the left pane.), click **Actions** on the menu, and then click **Save**.

The Start File System Services Utility

The **Start File System Services** utility starts all StorNext services on an MDC or StorNext client.

The StorNext services must be running on the MDC for file systems to be active and available. In addition, the StorNext services must be running to use the StorNext configuration utilities and to mount file systems using the client software.

To start StorNext File System services, on the Windows **Start** menu, click **All Programs > StorNext File System > Services Start**.

The Stop File System Services Utility

The **Stop File System Services** utility stops all StorNext services on an MDC or StorNext client.

To stop StorNext File System services, on the Windows **Start** menu, click **All Programs StorNext File System Services Stop**.

To start the StorNext services again, reboot the system or use the Start File System Services utility. For more information, see [The Start File System Services Utility above](#).

The Stop and Remove File System Services Utility

The **Stop and Remove File System Services** utility stops all StorNext services on an MDC or StorNext client, and also removes registry entries that automatically start the services on bootup.

To stop and remove StorNext File System services, on the Windows **Start** menu, click **All Programs > StorNext File System > Services Stop and Remove**.

To start the StorNext services again, you must use the Start File System Services utility. Rebooting the system will not restart services. For more information, see [The Start File System Services Utility above](#).

StorNext Help

View StorNext help documentation for utilities and other commands.

The Version Information Utility

The **Version Information** utility displays information about the currently installed version of the StorNext server and/or client software, such as the version number, build number, and platform.

To run the Version Information utility, on the Windows **Start** menu, click **All Programs > StorNext File System > Version Information**. A dialog box displays version information for the StorNext server and/or client software installed on your system.

StorNext Administration's File System Control


The StorNext Administration's File System Control dialog is used to Start, Stop, Failover and Activate a StorNext file systems. The actions in the File System Control dialog perform actions equivalent to the **cvadmin** commands.

- i Note:** To prevent long delays, the StorNext Administration tool does not automatically refresh. After issuing commands from the File System Control dialog, you may need to issue a **Refresh** before the changes will appear.

Start a File System

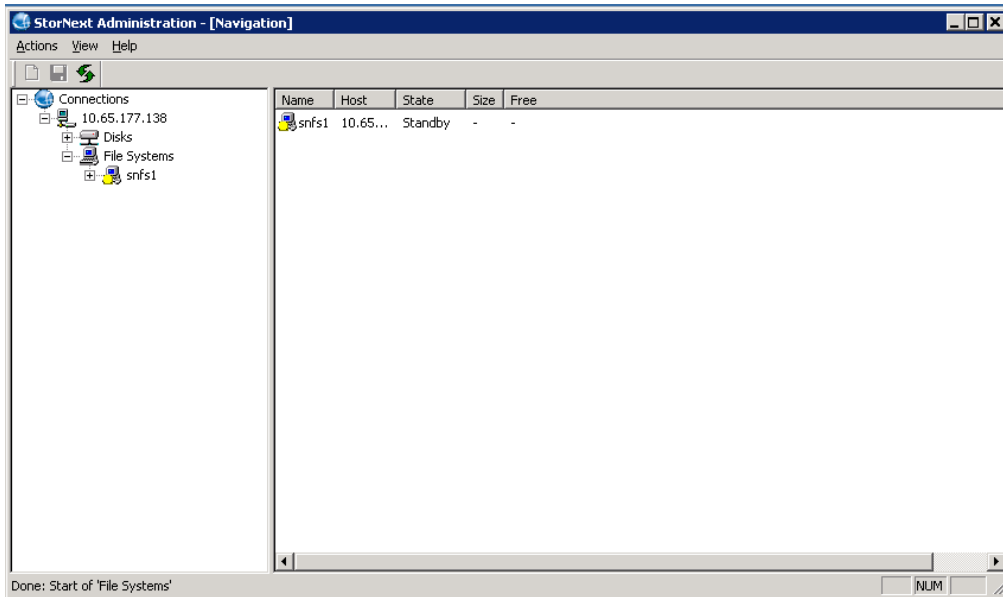
To start a file system, perform the following procedure.

1. Right-click **File Systems**, and then click **Start**. The **File System Control** dialog displayed.
2. Enter the file system name and the IP address of the MDC.
 - The **File System Name** is the name of a file system configuration file without the .cfg/.cfgx file extension.
 - The **Meta-data Controller (MDC) Host or IP** is the hostname or IP address of the MDC managing the specified file system.
 - If **Start a Single File System Manager** is checked, this will start a single FSM on the specified host. If it is unchecked, it will start FSMs on both HA (High Availability) MDCs if they are an HA pair. Optionally, you may have the dialog wait for the file system to start before returning. Otherwise the request will be issued and return immediately.



Once the file system is started, the file system status is displayed.

- A **yellow dot** indicates the file system is in the **standby** state.
- A **green arrow** indicates the file system is in the **active** state.



Activate a File System



Once the file system is started, perform the following procedure to activate the file system.

1. Right-click the file system, and then click **Activate**. The **File System Control** dialog is displayed with **Initiate Activation Vote** checked.


i Note: **Initiate Activation Vote** causes a client "vote" to commence and might cause a file system to activate. If an FSM is already active, no action is taken.

2. Click **OK** to activate.

Chapter 4: Windows SNFS-Only MDC Installation and Configuration
Configure a StorNext File System on Windows



If the file system activates successfully, the **File System Control** status is updated to display the following.




Other radio buttons in the **File System Control** specify the following parameters:

- The **Stop File System** radio button is used to stop a single FSM or to stop both FSMs, if they are HA pairs. As with the **Start File System** option, you can have the dialog wait for the command to complete before return control back to you.
- The **Fialover** option can be used on a HA MDC. This radio button will cause the file system to fail-over from the active FSM to the standby FSM (regardless of which MDC host is specified).

File System Displays

You can select a file system to display clients, debug options, quality of service, quotas and stripe groups. The table below provides the description of the various fields in the **File System Properties** display.

Chapter 4: Windows SNFS-Only MDC Installation and Configuration
 Configure a StorNext File System on Windows




| Name | Value | Type |
|---------------------------|--|---|
| Clients | | Expandable list of connects to the file system (clients with file system mounts and control connections) and the client properties. |
| Debug | None | Click to expand to a clickable list of debug flags. |
| Quality Of Service | None | Click to display QOS properties (if QOS is enabled on file system.) |
| Quotas | None | Click to display Quotas properties (if Quotas are enabled on file system.) |
| Stripe Groups | None | Expandable list of stripe groups and stripe group properties. |
| Host | IP address of the MDC running the file system. | |

| Name | Value | Type |
|---------------------|--|------|
| Percent Free | Percentage of free space on the file system. | |
| Size | Size of the file system | |
| State | Current state of the file system. | |


Client Information

To view information about a client connected to the file system, expand a file system in the left pane, and then click **Clients**. Information about the clients currently connected appears in the right pane.



Click on a specific client to see the client's properties.

Chapter 4: Windows SNFS-Only MDC Installation and Configuration
 Configure a StorNext File System on Windows



Select a specific **Client** to display the following information.

| Name | Value |
|------------------------|---|
| Build | StorNext version |
| DLAN Client | Is it a Distributed LAN Client? |
| DLAN Server | Is it a Distributed LAN Gateway? |
| Expires | When the license expires |
| Gateway License | Distributed LAN Gateway license installed |
| Hashed ID | SNFS Identifier hash value |
| ID | ID or IP address |
| IO Test Last | Last IO Test result |
| IO Test Rate | IO Test rate |
| IP Address | Address of connection |
| Local Address | Alternate local address |
| Nbr CPUs | CPU count on Client |


| Name | Value |
|---------------|-------------------------|
| Ordinal | FSM Connection number |
| Reconnect Cnt | Number of reconnections |
| Type | Type of connection |
| Uptime | Connection time |

Debug Information

Click the **Debug** icon to display the current active debug flags used by the File System Manager (FSM) process. If any of the debug flags are turned on, the debug output is printed to the file:

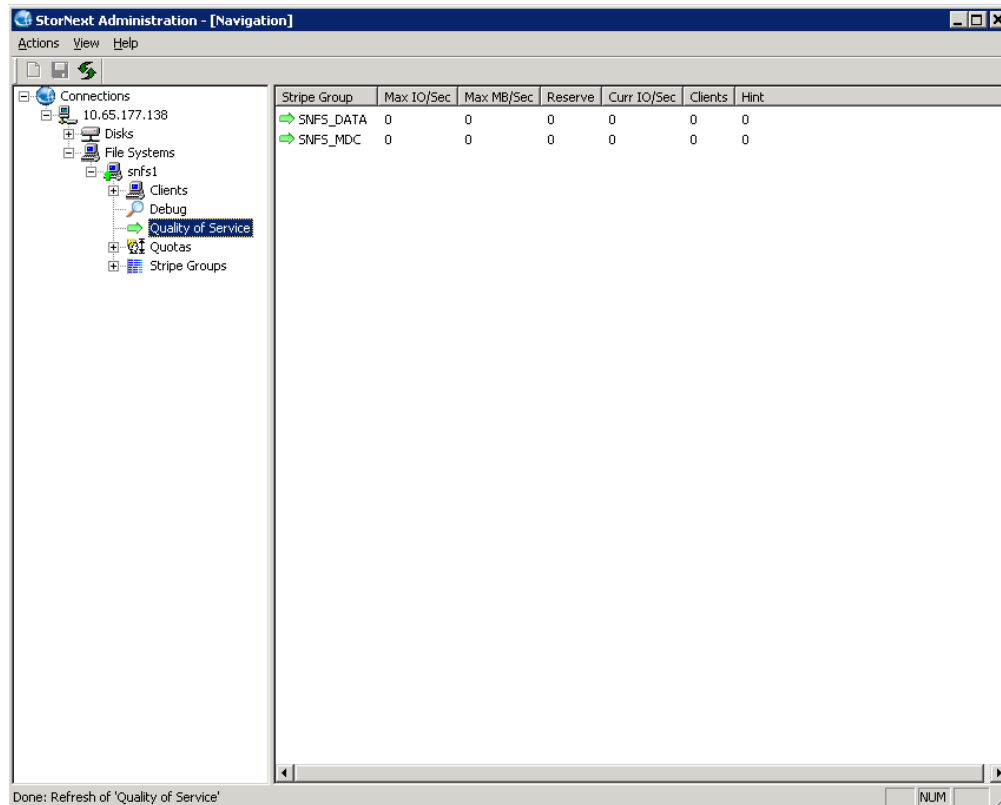
```
<StorNext Install Directory>/data/<file system name>/log/cvlog.txt
```

Caution: If you enable debug flags, this might severely impact file system performance. Do not enable debug flags unless recommended by Quantum.



Quality of Service Information

The **Quality of Service** (QOS) pane displays the settings and current metrics. Quality of Service must be enabled on the file system for it to be active. See [Quality of Service on page 250](#) for additional details.



Quota Properties


Quotas define the amount of storage space a specified user or group is allowed to own. You can set hard or soft quota limits. Quotas must be enabled in the File System Configuration file to be able to set and change quotas.

- To view quota properties, expand a file system in the left pane, expand **Quotas** and then click **Directories** or **Users & Groups**. The **User/Group Quota** dialog box appears
- To change quotas, right-click on **Directories** or **Users & Groups**. The **User/Group Quota** dialog box appears

There are three types of quotas you can configure:

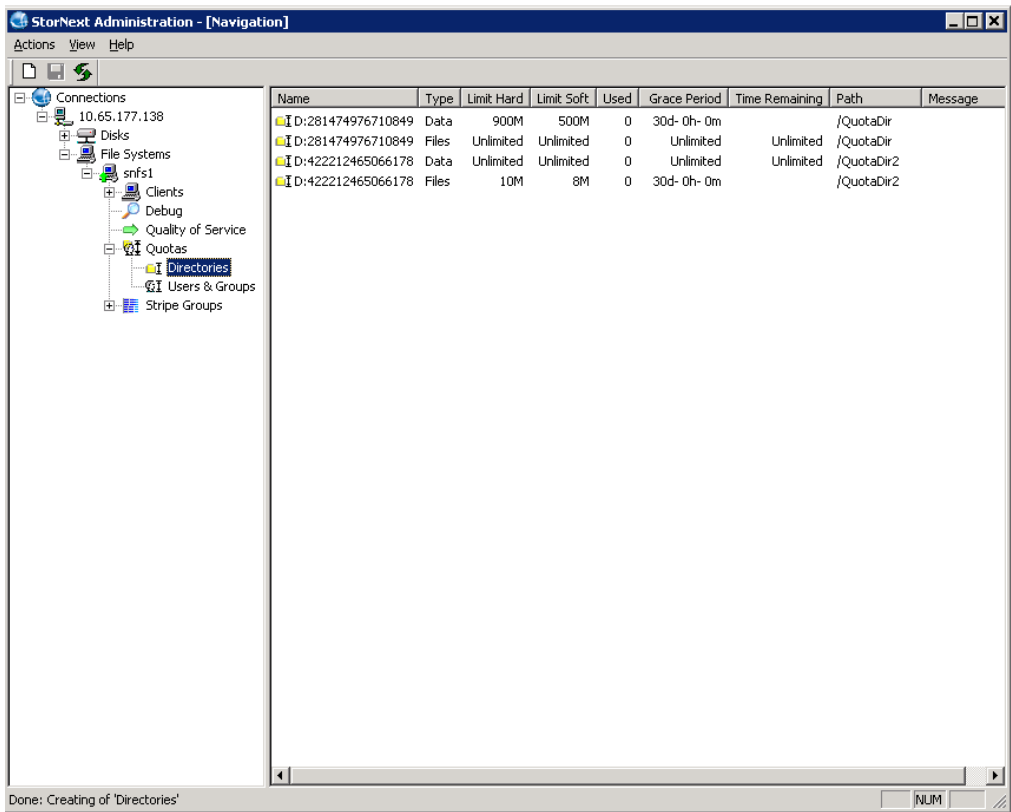
- Set quotas for users on the **User Quotas** button.
- Set quotas for groups on the **Group Quotas** button.
- Set data or file quotas for directories under the **Directory Quotas** heading.

Chapter 4: Windows SNFS-Only MDC Installation and Configuration
 Configure a StorNext File System on Windows



| Field / Button | Description | |
|-------------------|---|---|
| Quota Type | User Quota Type | Type the name of the user for which you want to set a quota. |
| | Group Quota | Type the name of the group for which you want to set a quota. |
| Directory Quotas | <ul style="list-style-type: none"> Data Limit: Set a data size limit on the directory. File Count Limit: Set a limit on the number of files under a directory. Directory: Specify the directory to apply the quotas to. | |

| Field / Button | Description | |
|-------------------|---|--|
| Data Limit | Hard Limit | Specify an amount in bytes or with the extensions M (megabytes), G (gigabytes) or T (terabytes). This is the maximum amount of space the specified user or group can use. |
| | Soft Limit | Specify an amount in bytes or with the extensions M (megabytes), G (gigabytes) or T (terabytes). Once the user or group exceeds this limit, a warning is sent. Note: Typically, this is set at 80% of the hard limit. |
| Time Limit | Specify the amount of time it takes for the soft limit to turn into a hard limit. | |

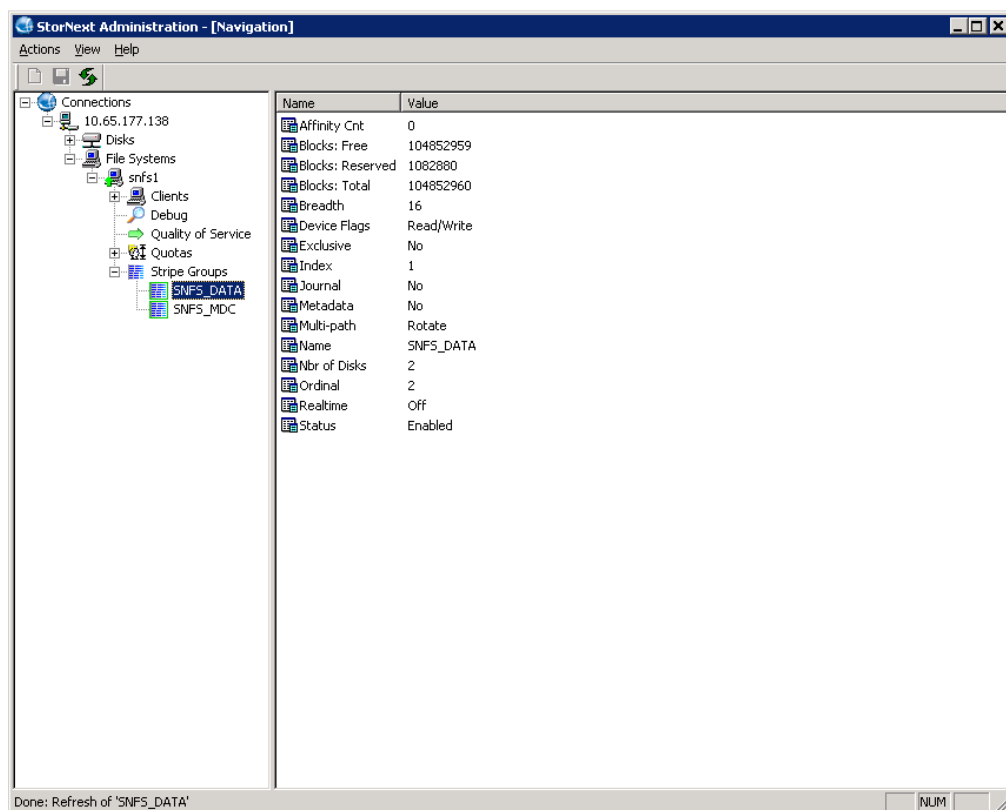


| Field / Button | Description |
|----------------|------------------------------|
| Name | The Linux user or group name |

| Field / Button | Description |
|-----------------------|--|
| Type | The type of quota (User or Group) |
| Limit Hard | The absolute maximum amount of data the user or group can consume |
| Limit Soft | The maximum amount of data the user or group can consume. They can exceed this limit for Grace Period amount of time, but cannot exceed the Hard Limit. |
| Used | The amount of data currently consumed. |
| Grace Period | The amount of time the user or group is allowed to exceed the Soft Limit |
| Time Remaining | Amount of time remaining before the user or group cannot allocate any more space in the file system. If the value is negative, the user can no longer allocate any more space. |

Stripe Group Properties

To view stripe group properties, expand a file system in the left pane, click **Stripe Groups**, and then click the stripe group name in the left pane. The stripe group's property appears in the right pane. The table below describes the fields of the stripe group property display.

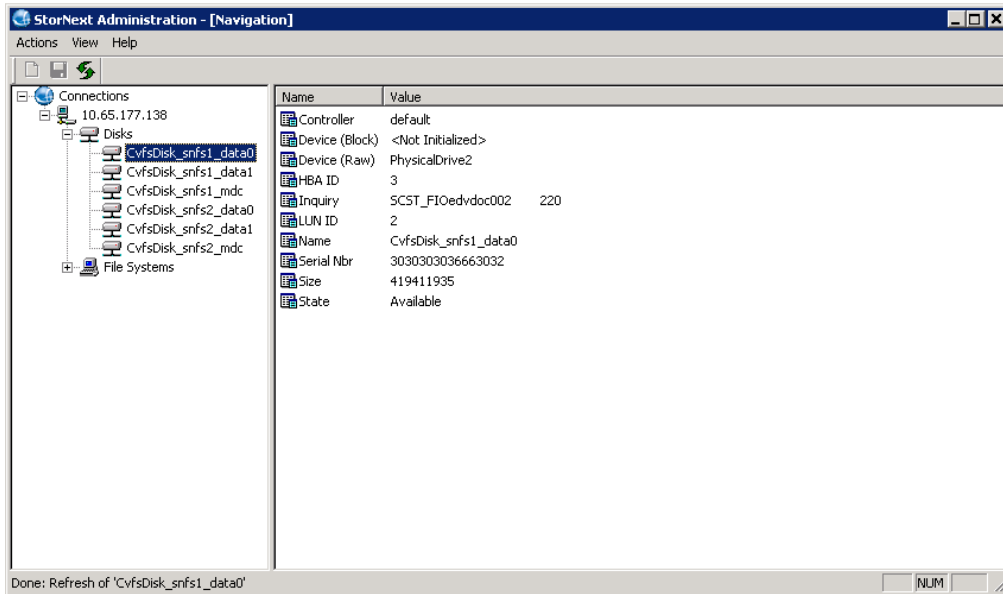


| Name | Value |
|-------------------------|---|
| Affinity Cnt | Number of affinities |
| Blocks: Free | Number of blocks free |
| Blocks: Reserved | Number of blocks reserved |
| Blocks: Total | Total number of blocks |
| Stripe Breadth | The number of file system blocks to write before switching to the next disk in the stripe group. |
| Stripe Depth | The number of disks in the stripe group. |
| Device Flags | Read and Write flags |
| Exclusive | Indicates if only specified file types (associated with the stripe group affinities) can be stored on the stripe group. |
| Index | Disk index in configuration |
| Journal | Indicates if the file system logging journal can be stored on the stripe group. |
| Metadata | Indicates if file system metadata can be stored on the stripe group. |
| Multi-path | Indicates the method the file system uses to access the disk: round, static, or sticky. |
| Name | Stripe group name |
| Nbr of Disks | Number of disks in group |
| Ordinal | Internal index |
| Realtime | Uses Real-time IO |
| Status | Current status |

Disk Properties

Click **Disks** to see all disks visible to the currently selected host. Click a specific disk to display the various attributes of that disk.

Chapter 4: Windows SNFS-Only MDC Installation and Configuration
 Configure a StorNext File System on Windows




| Name | Value |
|-----------------------|-----------------------------|
| Controller | Device driver |
| Device (Block) | The block device |
| Device (Raw) | The Raw Device |
| HBA ID | Host Bus Adapter identifier |
| Inquiry | Manufacturer string |
| LUN ID | Logical Unit Number |
| Name | Disk Label |
| Serial Nbr | Disk's Serial Number |
| Size | Size of the disk |
| State | State of the disk |

Configure a High Availability System on Windows

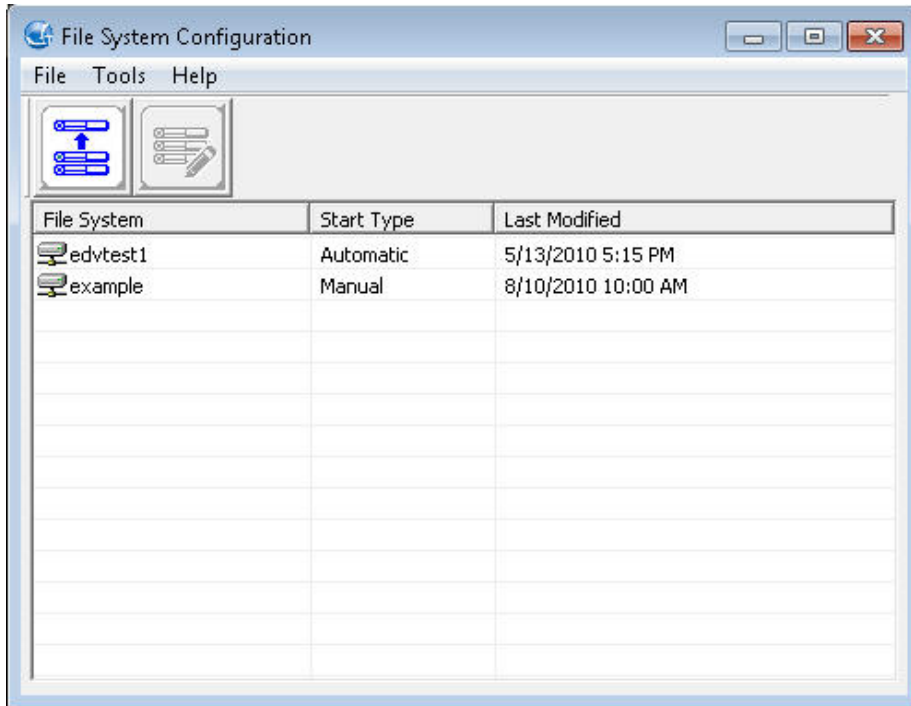
Configure a High Availability (HA) System on Windows

Follow these steps to configure a High Availability (HA) system on Windows:

1. Install StorNext if not already installed (see [Install the StorNext File System on a Windows MDC: Interactively on page 73](#)).
2. Obtain and install StorNext licenses, which must include at least Maintenance, Server and Failover licenses for both MDCs (see [Obtain and Install StorNext Licenses – GUI-based Appliances and Customer-supplied MDCs](#)).
3. If this is an existing configuration with clients, it is recommended that you stop StorNext services on all StorNext clients to avoid the clients hanging (see [Configure a StorNext File System on Windows on page 82](#)).
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.
 - d. If the optional `ha_smith_interval` file is configured, it must be identical on both computers (for more details, see [High Availability Systems](#)).
 - e. Copy and paste the *.cfg files and **fsmlist** files from the primary MDC to the secondary MDC before proceeding to **Step 5**.
5. Create the **ha_peer** file in the configuration directory of each MDC (**C:\Program Files\StorNext\config\ha_peer**). The **ha_peer** file on each MDC must contain the numerical IP address of its redundant-peer MDC, which enables intra-MDC communication for HA functionality.
6. Enable High Availability (HA) on both redundant-peer MDCs by doing the following:

 **Caution:** If you do not perform these steps on both MDCs your metadata could become corrupted.

- a. After installing StorNext, on each MDC run **Start > All Programs > StorNext > File System Cfg (Advanced)**. The **File System Configuration** window displays.



- b. On the **Tools** menu, click **HA Configuration**.
 - c. Click **Enable High Availability (HA)**. Note that you cannot enable HA if there are no file systems configured.
7. High Availability is now configured. You may now test the HA reset functionality by running the following command on either MDC:

```
cvadmin -e "hamon smith"
```

The MDC will immediately reboot. Clients for any StorNext file systems that were being served by active FSMs on the reset MDC will trigger elections that will result in the activations of corresponding FSMs on the redundant HA MDC. This can be observed by running the following command on the redundant MDC both before and after resetting its peer MDC:

```
cvadmin -e "select"
```

8. Restart services on any clients that you stopped earlier (see [Configure a StorNext File System on Windows on page 82](#)).

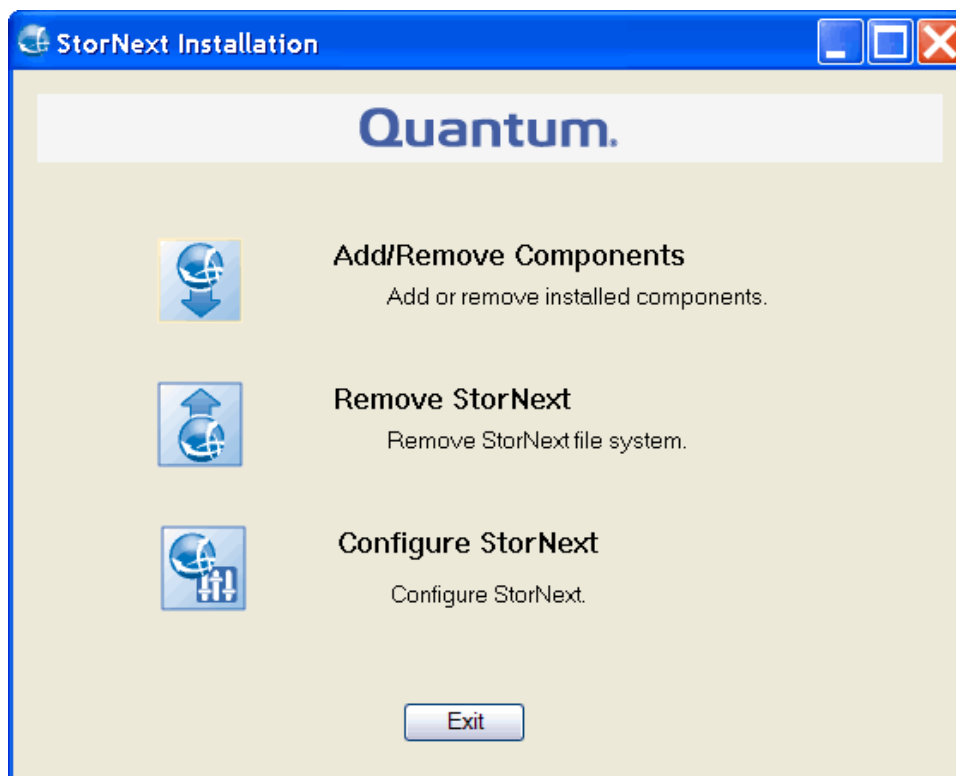
- i Note:** Restarting StorNext faster than once per three minutes can trigger a startup delay of three minutes. The delay slows down looping HA resets so that an administrator can correct any configuration problems that are causing the resets.

Add or Remove StorNext Components


After installing StorNext File System, you may wish to add or remove components. Specifically, you can add or remove the online help files, the StorNext client, and the StorNext server.

Add or Remove StorNext Components

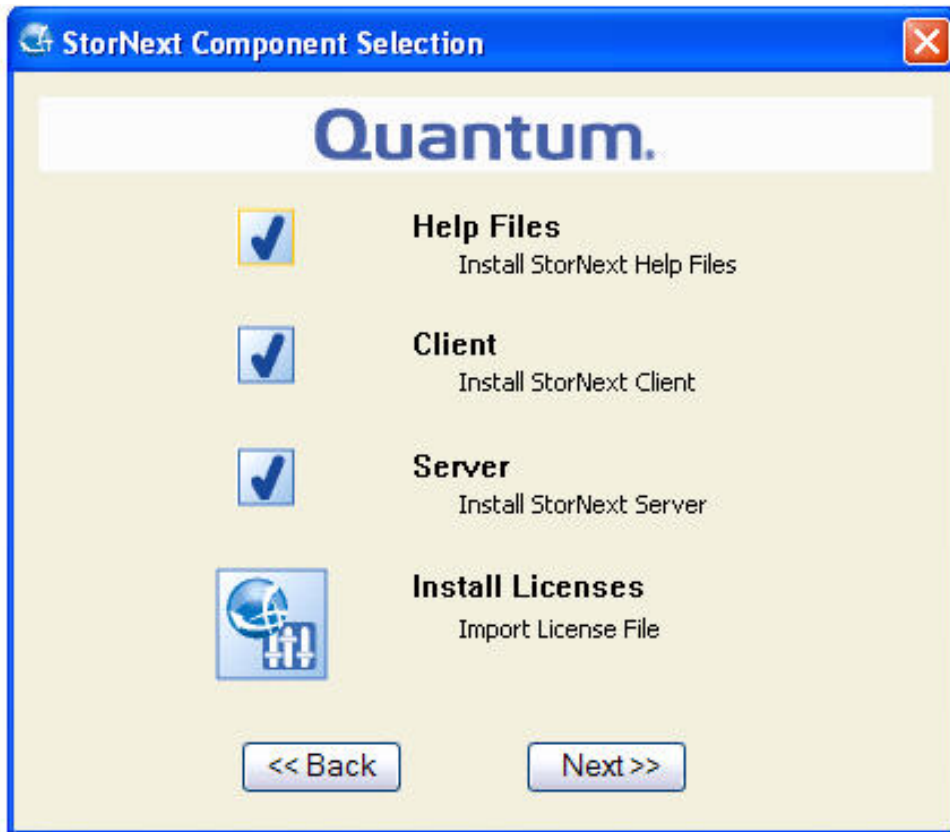
1. If you just completed installing StorNext, the **StorNext Installation** window displays. The components you selected during the install are already added, but if you changed your mind and want to select different components, you can do so now.



- i Note:** If you previously installed StorNext and would like to add or remove components, display the **StorNext Installation** window by double-clicking the file named **SnfsSetup64.exe** (64-bit systems). Note that the screen is slightly different from above; instead of **Add/Remove Components**, the button you must click is labeled **Upgrade/Reinstall StorNext**.



2. Click **Add/Remove Component** or **Upgrade/Reinstall StorNext**, depending on which screen you are looking at. The **Component Selection** screen appears.



3. Select the check boxes of the components you want to add. Clear the check boxes of the components you want to remove. Click **Next** to proceed.

i Note: The **Install Licenses** option allows you to import your StorNext licenses onto the StorNext server, if you have not already done so. You may also install licenses later using this screen or the Windows Start button (**Start > Licenses** or **Start > All Programs > StorNext > Licenses**). The StorNext End User License Agreement (EULA) appears.

4. Read and accept the EULA by selecting “I accept the terms in the License Agreement” and clicking **Next**. It may take several minutes for the selected components to finish being added or removed.
5. When a message informs you that the process is complete, click **OK**.
6. Click **Exit**.

Restore a Previous Client Configuration

See [Restore a Previous StorNext Client Configuration on page 174](#).

Upgrade or Reinstall StorNext on a Windows Server

The **Upgrade/Reinstall** button allows you to upgrade to a new version of StorNext or reinstall the current version. You might want to reinstall if you accidentally removed some components and need to get them back. During an upgrade/reinstall, the old software is removed and the new software is installed. Your previous configuration settings remain intact.

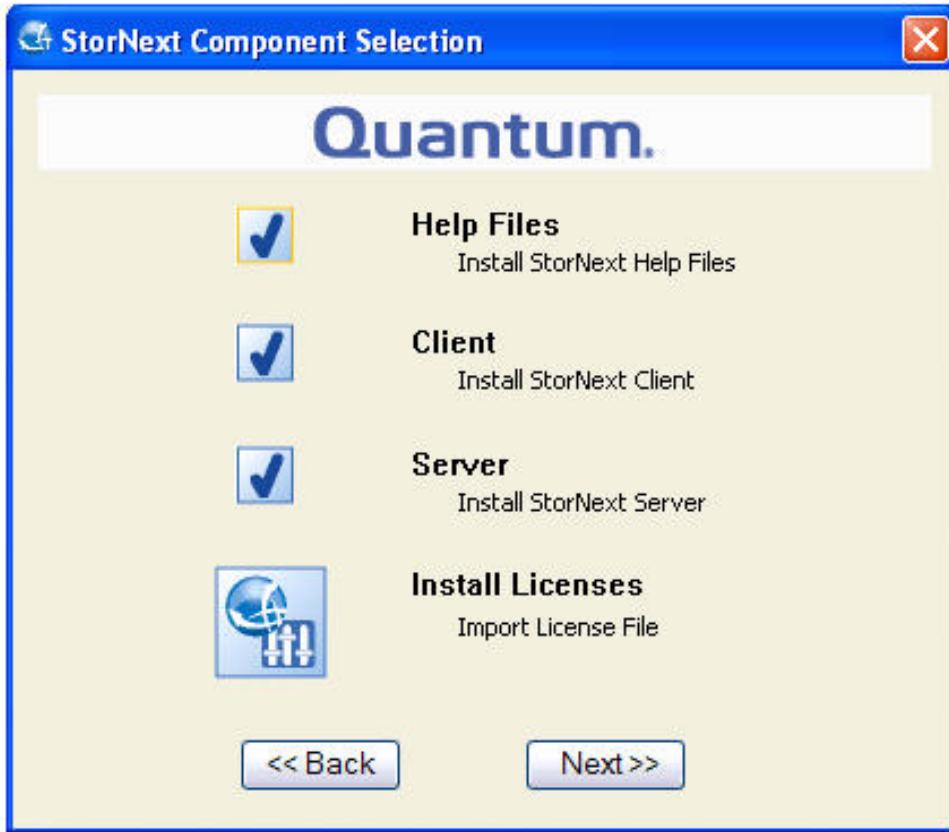
i Note: If you are upgrading a Windows 3.x high-availability system to StorNext, additional steps are required. Follow the instructions in the StorNext Upgrade Guide.

Upgrade or Reinstall StorNext on a Windows Server

1. Copy the file named `/fs_only/snfs_full_Windows_x86_64.zip` from the 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.
2. Unzip the file.
3. Double-click the file named **SnfsSetup64.exe** to launch the installer. The StorNext Installation screen appears.



4. Click **Upgrade/Reinstall StorNext** to begin installation. The **StorNext Component Selection** window appears.



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



Chapter 5: StorNext Client Installation and Configuration

This chapter contains the following topics:

| | |
|---|-----|
| StorNext Client Installation and Configuration | 143 |
| StorNext Client System Requirements | 144 |
| Download the StorNext Client Software | 146 |
| Install the StorNext Client Software on Linux | 149 |
| Install the StorNext Client Software on Windows | 162 |
| Configure the StorNext Client Software on Windows | 169 |

StorNext Client Installation and Configuration

This topic describes how to perform a StorNext client-only installation (without StorNext File System or StorNext Storage Manager) on supported platforms. StorNext client software lets you mount and work with StorNext file systems.

i Note: To ensure successful operation, make sure the StorNext client system meets all operating system and hardware requirements (see [StorNext Client System Requirements on the next page](#)).

i Note: Before installing the StorNext client software you must install and configure your StorNext or Xsan MDC, activate StorNext with a license key, and create necessary File Systems. For Xsan MDCs consult your Apple Xsan documentation for information on configuring the MDC.

See the sections below for additional information:

- [StorNext Client System Requirements below](#)
- [Download the StorNext Client Software on page 146](#)
- [Install the StorNext Client Software on Linux on page 149](#)
- [Install the StorNext Client Software on Windows on page 162](#)
- [Configure the StorNext Client Software on Windows on page 169](#)

StorNext Client System Requirements

To run the StorNext client software, the client system must meet all operating system and hardware requirements.

StorNext Software Requirements

Installing StorNext on all UNIX-based (including Linux) platforms requires installing these quantum-supplied packages:

- **jansson**
- **libevent**
- **protobuf**
- **zeromq**

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

Operating System Requirements

The operating systems, kernel versions, rpm package dependencies, and hardware platforms supported by StorNext File System are listed in the [StorNext 6 Compatibility Guide](#).

Make sure the system meets the operating system and platform requirements outlined in the guide, and, if necessary, update to a supported release or kernel version before installing StorNext.

i Note: For Linux clients, you must be able to disable C-States through kernel boot parameters. See [Modify the grub.conf File to Disable the Intel Sleep State on page 7](#).

Hardware Requirements

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

- For SAN (FC-attached) clients or for LAN clients:
 - 1 GB RAM
 - 500 MB available hard disk space
- For SAN clients acting as a Gateway server:
 - 2 GB RAM
 - 500 MB available hard disk space

i Note: Gateway servers may require additional RAM depending on the number of file systems, LAN clients, and NICs used. See [Gateway Server/Client Network and Memory Tuning](#) for Gateway server/client memory tuning guidelines.

StorNext Client Hardware Requirements

RAM, hard disk, and CPU requirements to run the StorNext clients are listed in the [StorNext 6 Compatibility Guide](#).

i Note: Make sure the system to be used as a StorNext client meets the hardware requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

For improved performance and stability, on Windows clients it is strongly recommended that CPU C-States be disabled in the BIOS if they are supported by your system. Refer to system vendor documentation for additional information.

StorNext Client File System Buffer Cache Requirements

File System Buffer cache requirements for the StorNext clients are listed in the [StorNext 6 Compatibility Guide](#).

i Note: Make sure the system to be used as a StorNext client meets the hardware requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

Verify You Have Enough Space to Extract the Installation Software

On Linux systems, the StorNext client installation software files are extracted into the `/tmp/stornext` directory by default. The software, when extracted, requires approximately 60 MB of space, and requires approximately 220 MB of space when installed. Make sure there is enough free space in the default directory to extract the files. If there is not sufficient space, you may need to specify an alternative directory with the `-d` option.

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

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

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

Other Requirements

Additional requirements that must be met for the StorNext system are listed in the [StorNext 6 Compatibility Guide](#).

-
- Note:** Make sure the system meets the additional requirements outlined in the guide, and, if necessary, update the system before installing StorNext.

StorNext File System Buffer Cache

Multiple mounted file systems typically share a single buffer cache. A StorNext Client buffer cache is created for each different **cachebufsize**. By default, all file systems have the same **cachebufsize** of 64K, so they all share the same buffer cache.

-
- Note:** These settings do not apply to Apple Xsan Clients, which do not use the StorNext buffer cache.

The amount of memory consumed by default for each **cachebufsize** depends on the platform type and the amount of memory in the system. The table below displays the default amount of memory consumed by **cachebufsize**.

| Platform Type | <=2GB Memory | 2GB Memory |
|---------------------|--------------|------------|
| 32-bit Windows | 32MB | 64MB |
| All other platforms | 64MB | 256MB |

To see information about the buffer cache after mounting file systems, use the **cvdb(1)** command with the **-b** option. To change the amount of memory used by the buffer cache at mount time, use the **buffercachecap** parameter.

On Windows, the non-paged pool is used for buffer cache memory until it consumes up to 64 megabytes (32-bit systems) or 64 gigabytes (64-bit systems). Any additional buffer cache memory comes from the paged pool.

Download the StorNext Client Software

You can download the StorNext client software from the Quantum web site, or from a StorNext Metadata Controller (MDC) running the StorNext GUI.

There are three types of Client packages available, depending on the Operating System:

- **StorNext Client** (snfs-client)
- **StorNext Client with Mover** (snfs_client_plus_mover) [DDM]
- **StorNext Gateway Agent** (snfs_gateway-agent)

The table below provides information about the disk space required to install each package.

| Package Type | Description | Required Disk Space for Install |
|-----------------------------------|---|---------------------------------|
| StorNext Client | Base StorNext file system client capability | 280 MB |
| StorNext Client with Mover | Base StorNext file system client capability with packages for supporting Distributed Data Mover (DDM) | 390 MB |
| StorNext Gateway Agent | Gateway metrics reporting agent add-on | 5 MB |

Note: You can only download StorNext client software from a StorNext MDC if you have SNFS and SNSM installed and are running the StorNext GUI.

Note: StorNext client downloads are not supported from an Xsan MDC.

Locate the StorNext Client Software on the Quantum Web Site

Follow the instructions below to download the client software from the Quantum web site.

For New Customers

1. Go to <https://swdownload.quantum.com>.
2. Click **Register Now**. Complete the registration form.
3. Click **Submit**. The system will create an account and then send you an email with the password within five minutes. If you do not receive the email after five minutes, go to login page, enter your email and click **Recover Password**. You will receive an encrypted link to reset your password. Once your email has been registered and verified, you will be able to enter and manage your **Download Authorization Code**.
4. Go to the login page and sign in using your **User ID** and **Password**.
5. Register the **Download Authorization Code** and click **Submit**.
6. Click the **Product Name** link to begin downloading the software.

For Existing Customers


1. Go to <https://swdownload.quantum.com> and log in.
2. Click **Register Additional Download Authorization Codes**.
3. Enter your **Download Authorization Code** and click **Register**.

4. The system will register the download authorization code and make the download available.
5. Click the **Product Name** link to begin downloading the software.


Download the StorNext Software from the StorNext GUI

Follow the instructions below to download the client software from the StorNext GUI.

- i Note:** To download the client software from the StorNext MDC GUI, the client system must have network access to the MDC.
1. From a web browser, access the URL (host name and port number) of the StorNext MDC (for example, **http://servername:81**). Use one of the web browsers supported by StorNext to access the MDC. For a list of currently-supported browsers, see the [StorNext Compatibility Guide](#).
 2. When prompted, type the username and password for the MDC, and then click **OK**. The default username is **admin**, and the default password is **password**. The StorNext home page appears.
 3. On the **Tools** menu, click **Client Download**. The **Tools > Client Download** page appears.



4. Select from the **Platform** list the desired operating system.
5. Select from the **OS Version** list the desired operating system version corresponding to the platform you selected.
6. When a window appears containing a link to the client software download location, click the link to begin downloading.



7. Click **Download** to begin the process.
8. When prompted, choose the **Save to Disk** option, and then click **OK**.
9. Browse to the location where you want to save the file, and then click **Save**.
10. After the client package has been saved, click **Done**.
11. Continue with the installation procedure for your operating system. Depending on your operating system, proceed to one of the following sections to install the software:
 - [Install the StorNext Client Software on Linux below](#)
 - [Install the StorNext Client Software on Windows on page 162](#); after you have installed the software, see [Configure the StorNext Client Software on Windows on page 169](#) to configure the software.

Install the StorNext Client Software on Linux

After downloading the StorNext client software from the MDC or installation DVD, first make sure you have enough space to extract the software files by following the instructions in [Verify You Have Enough Space to Extract the Installation Software on page 145](#).

⚠ Caution: Firewalls can interfere with the proper functioning of StorNext and result in unexpected errors unless specifically configured for use with StorNext. Quantum recommends you disable the firewall on the metadata network. See [Ports Used By StorNext](#) for additional information.

Install the software on the StorNext client system using one of the following procedures for your operating system:

- [Install the StorNext Client Software on Red Hat or SUSE Linux on the next page](#)
- [Install the StorNext Client Software on Debian or Ubuntu Linux on page 153](#)
- [Install the StorNext Client Software on Sun Solaris on page 154](#)
- [Install the StorNext Client Software on HP-UX on page 157](#)
- [Install the StorNext Client Software on IBM AIX on page 159](#)

i Note: Regardless of the operating system platform onto which you are installing, you must view and accept the StorNext end user license agreement (EULA) before you can install StorNext.

Install the StorNext Client Software on Red Hat or SUSE Linux

To run the StorNext client software on Red Hat Linux or SUSE Linux Enterprise, first install the client software package, and then configure the client.

⚠ Caution: Before you install the StorNext client software, you must install the kernel headers for the running kernel required for building kernel modules. Obtain the kernel headers package from your Operating System vendor.

i Note: Client software is extracted by default to the directory `/tmp/stornext`.

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. Run the client software binary. At the command prompt, where `<archive name>` is the name of the client software archive file you downloaded, type:

```
./<archive name>
```

For example, the file for **Red Hat 7** is `snfs_client_RedHat7_x86_64`.

4. When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA enter **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; see [StorNext Client System Requirements on page 144](#)).

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, 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:

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

⚠ 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, 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:

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

Quantum recommends that you do not force using an existing directory (by using the `-f` command) because this could lead to additional problems installing the `.rpm` files.

5. Change directories to `/tmp/stornext` (or to the alternate directory if you specified one in the step above).
6. List the packages extracted from the software archive file. At the command prompt, type:

```
ls -l
```

Identify the correct package to install. The correct package begins with **snfs-client** and ends with the **.rpm** file name extension.

i Note: The file that ends with **.rpm.md5sum** is a checksum file, not the client software package.

7. Install the required `.rpm` files in the following order. Use the command `rpm -ivh` before each package. For example, `rpm -ivh 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-extapi-<version>.rpm`
 - h. `snupdate-<version>.rpm`

The following additional `.rpm` files are required when installing the **snfs_client_plus_mover** packages:

i Note: For StorNext 6.2, only RHEL 6 and RHEL 7 mover packages are supported.

- i. `quantum_curl-<version>.rpm`
- j. `quantum_libarchive-<version>.rpm`
- k. `quantum_mariadb-connector-odbc-<version>.rpm`
- l. `quantum_unixODBC-<version>.rpm`
- m. `quantum_openssl-util-<version>.rpm`
- n. `quantum_sndbapi-<version>.rpm`
- o. `quantum_sndbapicxx-<version>.rpm`

- p. `snltfs-<version>.rpm`
- q. `snfs-mover-<version>.rpm`

The following additional `.rpm` files are required when installing the **snfs_gateway-agent** packages:

i Note: The gateway packages are only supported on RHEL 6, RHEL 7, SLES 11, and SLES 12.

- r. `quantum_curl-<version>.rpm`
- s. `snfs-gateway-agent-<version>.rpm`

8. Configure the boot order for system services.

i Note: Skip this step for Red Hat Enterprise Linux 7.x and SUSE Linux Enterprise Server 12.x clients. The boot order is determined by `systemd.unit(5)` unit file dependencies.

Perform one of the following:

- For Red Hat Linux, at the command prompt, type:

```
chkconfig --level 345 cvfs on
```

- For SUSE Linux Enterprise, at the command prompt, type:

```
chkconfig -s raw 235  
chkconfig -s cvfs 345
```

9. 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.

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

10. Create a mount point for the file system. At the command prompt, where `<mount point>` is the directory path where you want the file system to be mounted, type:

```
mkdir -p <mount point>  
chmod 777 <mount point>
```

For example: `/stornext/snfs1`.

11. Add file system entry to `/etc/fstab` file. To do this, edit the `/etc/fstab` file so that it contains the following line, where `<file system>` is the name of the StorNext file system and `<mount point>` is

the directory path created in the previous step:

```
<file system> <mount point> cvfs rw,noauto 0 0
```

12. Reboot the StorNext client system. After reboot, the StorNext file system is mounted at the mount point you specified.

i Note: To manually mount a file system, at the command prompt, where `<file system>` is the name of the StorNext file system and `<volume>` is the name of the Xsan volume you will use, and `<mount point>` is the directory path where you want the file system to be mounted, type:

```
mount -t cvfs <file system/volume> <mount point>
```

Install the StorNext Client Software on Debian or Ubuntu Linux

To run the StorNext client software on Debian or Ubuntu Linux, first install the client software package, and then configure the StorNext client. To install the client software package:

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. Execute the **snfs_client_Debian60_26x86_64.bin** or **snfs_client_Ubuntu14_x86_64.bin** package.
4. When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA enter **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; see [StorNext Client System Requirements on page 144](#)).

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, 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:

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

⚠ 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, 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:

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

Quantum recommends that you do not force using an existing directory (by using the `-f` command) because this could lead to additional problems installing the `.rpm` files.

- From the `/tmp/stornext` directory, install all the `.deb` packages packages by typing:

```
# dpkg -i *.deb
```

- Copy the **fsnameservers** file from MDC to `/usr/cvfs/config/fsnameservers`.

i Note: Note: The **fsnameservers** file on the client must be exactly the same as on the MDC.

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

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

- Create a mount point, where **<snfs1>** is the name of the StorNext File System/Xsan volume you are using on your system. For example:

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

- Edit the `/etc/fstab` to add the client mount options, where **<snfs1>** is the name of the StorNext File System/Xsan volume you are using on your system. For example:

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

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

```
service cvfs start
```

- Verify file systems are mounted.

Install the StorNext Client Software on Sun Solaris

To run the StorNext Client software on Sun Solaris, first install the client software package, and then configure the 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. Run the client software binary. At the command prompt, where `<archive name>` is the name of the client software archive file you downloaded, type:

```
./<archive name>
```

For example, the file for RH5 is `snfs_client_RedHat50AS_26x86_64.bin`.

4. When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA enter **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; see [StorNext Client System Requirements on page 144](#)).

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, 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:

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

! 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, 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:

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

Quantum recommends that you do not force using an existing directory (by using the `-f` command) because this could lead to additional problems installing the `.rpm` files.

! Caution: If you have other Quantum products that share the **zeromq** and **protobuf** packages, it is recommend you use the newer versions. Consult the other product's documentation on how to remove the older versions.

5. Change to the `/tmp/stornext` directory.
6. Install the required SNFS packages first, and then install the client software package. At the command prompt, type:

```
pkgadd -d . QUANTUMjansson
pkgadd -d . QUANTUMprotobuf
pkgadd -d . QUANTUMzeromq
pkgadd -d . ADICsnfs
pkgadd -d . QUANTUMsnupdate
```

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <ADICsnfs> [y,n,?]

7. Type y to confirm installation of the ADICsnfs package.
8. Edit the `/usr/cvfs/config/fsnameservers` text file to contain the IP address of the MDC to which the 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.

i Note: The `fsnameservers` file needs to contain only the IP address of the MDC. For example, if the IP address for the MDC is **192.168.143.1**, then the `fsnameservers` file should contain a single line with **192.168.143.1**

9. Create a mount point for the file system. At the command prompt, where **<mount point>** is the directory path where you want the file system to be mounted. For example: `/stornext/snfs1`, type:

```
mkdir -p <mount point>
chmod 777 <mount point>
```

10. Configure the file system to automatically mount after a reboot. To do this, edit the `/etc/vfstab` file so that it contains the following line, where **<file system>** is the name of the StorNext file system and **<mount point>** is the directory path created in **Step 9**:

```
<file system> - <mount point> cvfs 0 auto rw
```

While a reboot is not required at this time, in the future, a reboot of the client will mount the StorNext file system at the mount point you specified.

11. Restart StorNext services on the solaris client. To restart the services, type:

```
/usr/cvfs/bin/cvfs stop
/usr/cvfs/bin/cvfs start
```

To manually mount a file system, at the command prompt, where **<file system/volume>** is the name of the StorNext file system/Xsan volume you will use, and **<mount point>** is the directory path where you want the file system to be mounted, type:

```
mount -F cvfs <file system/volume> <mount point>
```


Install the StorNext Client Software on HP-UX

To run the StorNext Client software on HP-UX, first install the client software package, and then configure the 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. Run the client software binary. At the command prompt, where **<archive name>** is the name of the software archive file you downloaded from the MDC, type:

```
./<archive name>
```

For example, the file for RH5 is **snfs_client_RedHat50AS_26x86_64.bin**.

4. When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA enter **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; see [StorNext Client System Requirements on page 144](#)).
5.  **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, 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:

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

6.  **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, 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:

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


Quantum recommends that you do not force using an existing directory (by using the `-f` command) because this could lead to additional problems installing the `.rpm` files.

7. Change directories to `/tmp/stornext`.
8. List the packages extracted from the software archive file. At the command prompt, type:

```
ls -l
```

Identify the correct package to install. The correct package begins with **snfs-client** and ends with the **.depot** file name extension.

9. Install the required SNFS packages first, and then install the client software package. At the command prompt, type:

```
swinstall -s /tmp/stornext/quantum_zeromq-<version>.depot -x mount_all_
filesystems=false \*
swinstall -s /tmp/stornext/quantum_protobuf-<version>.depot-xmount_all_
filesystems=false \*
swinstall -s /tmp/stornext/snfs-client_hpux-<version>.depot -x mount_all_
filesystems=false \*
```

10. Edit the `/usr/cvfs/config/fsnameservers` text file to contain the IP address of the MDC to which the 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.

i Note: The **fsnameservers** file needs to contain only the IP address(es) of the MDC(s).

11. Create a mount point for the file system. At the command prompt, where **<mount point>** is the directory path where you want the file system to be mounted. For example: `/stornext/snfs1`, type:

```
mkdir -p <mount point>
chmod 777 <mount point>
```

12. Add file system entry to `/etc/fstab` file. To do this, edit the `/etc/fstab` file so that it contains the following line:

```
/stornext/snfs1 /stornext/snfs1 cvfs rw,noauto,<fsname>=RH5-DATA 0 0
```

- *fsname* defines the StorNext file system name (for example, **RH5-DATA**).
- `/stornext/snfs1` is the StorNext mount point.

i Note: Regarding the two file system mount points in the `/etc/fstab` file, the first mount point entry is ignored and the second mount point entry is the StorNext mount point you want to have the file system mounted on.

13. Reboot the client system. After reboot, the StorNext file system is mounted at the mount point you specified. To manually mount a file system, where **<mount point>** is the directory path where you want the file system to be mounted, at the command prompt type:

```
mount -F cvfs <mount point> <mount point>
```

Install the StorNext Client Software on IBM AIX

To run the StorNext Client software on IBM AIX, first install the client software package, and then configure the 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. Run the client software binary. At the command prompt, where **<archive name>** is the name of the software archive file you downloaded, type:

```
./<archive name>
```

For example, the file for RH5 is **snfs_client_RedHat50AS_26x86_64.bin**.

4. When you are presented with the StorNext End User License Agreement (EULA), press **Enter** to read the EULA. After reading the EULA enter **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; see [StorNext Client System Requirements on page 144](#)).

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, 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:

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

⚠ 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, 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:

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

Quantum recommends that you do not force using an existing directory (by using the `-f` command) because this could lead to additional problems installing the `.rpm` files.

5. Change directories to `/tmp/stornext`.
6. List the packages extracted from the software archive file. At the command prompt, type:

```
ls -l
```

Identify the correct package to install. The correct package begins with **snfs** and ends with the **.bff** file name extension.

7. Install the required SNFS packages, and then install the client software package. At the command prompt, type:

```
installp -ac -d . all
```

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

i Note: The **fsnameservers** file needs to contain only the IP address(es) of the MDC(s).

9. Create a mount point for the file system. At the command prompt, where **<mount point>** is the directory path where you want the file system to be mounted, type:

```
mkdir -p <mount point>  
chmod 777 <mount point>
```

For example: `/stornext/snfs1`.

10. Configure the file system to automatically mount. For default mount options, at the command prompt, where **<file system>** is the name of the StorNext file system and **<mount point>** is the directory path created in **Step 9**, type:

```
crfs -v cvfs -d <file system> -a verbose=yes -a type=cvfs -A yes -m <mount  
point>
```

Or, if any mount options are required, where `<mount options>` is a comma separated list. StorNext mount options are documented in the `mount_cvfs(8)` man-page, type:

```
crfs -v cvfs -d <file system> -a verbose=yes -a options=<mount options> -a  
type=cvfs -A yes -m <mount point>
```

The StorNext file system is mounted at the mount point you specified.

i Note: To manually mount a file system, at the command prompt, where `<mount point>` is the directory path where you want the file system to be mounted, type:

```
mount <mount point>
```

11. Start the new StorNext client software. At the command prompt, type:

```
/etc/rc.cvfs start
```

Alternatively, you can reboot the system.

Additional Considerations

In Red Hat Enterprise Linux 7.x and SUSE Linux Enterprise Server 12.x clients, the `init.d(7)/service(8)` system service utilities have been replaced by the `systemd(1)/systemctl(1)` utilities.

The most notable difference is that the `systemctl(1)` utility does not write status output to the controlling TTY during start and stop operations. The user must run `systemctl status cvfs` after the start or stop request to view the status information.

Red Hat Enterprise Linux 7.x and SUSE Linux Enterprise Server 12.x clients provide a script in `/sbin/service` for convenience and compatibility that is simply a wrapper that calls the `systemctl(1)` command with appropriate arguments. Like `systemctl(1)`, the `service(8)` script will not write status output to the TTY during start and stop operations. The user must execute `/sbin/service cvfs status` to view the status information.

Additional kernel module verification is enabled in Red Hat Enterprise Linux 7.x and SUSE Linux Enterprise Server 12.x clients. Since StorNext is not shipped with a signed `cvfs` kernel module, an additional warning message will be seen in the `syslog` when the `cvfs` kernel module is loaded:

```
[44.915641] cvfs: module verification failed: signature and/or required key  
missing - tainting kernel
```

You can safely ignore the warning message.

Install the StorNext Client Software on Windows

The StorNext setup wizard guides you through the process of installing the StorNext client software on Windows machines.

If you did not install all StorNext components when you initially installed StorNext or wish to remove them later, you can add or remove components at any time (see [Add or Remove StorNext Components on page 172](#)).

i Note: After installation, it is completely optional for you to restore the previous client configuration (see [Restore a Previous StorNext Client Configuration on page 174](#)). This is only necessary if you have performed a complete removal of StorNext. The StorNext client configurations are fully maintained on upgrades from version 3.0 (and later).

As new versions of StorNext are released, you can upgrade to a new version or reinstall (see [Upgrade or Reinstall StorNext Clients on page 178](#)).

i Note: You must log on as an **Administrator** to install StorNext. If you are installing on Windows Vista, answer **Yes** to any messages asking if you want to run the installation process with administrative privileges.

Install the StorNext Client Software on Windows

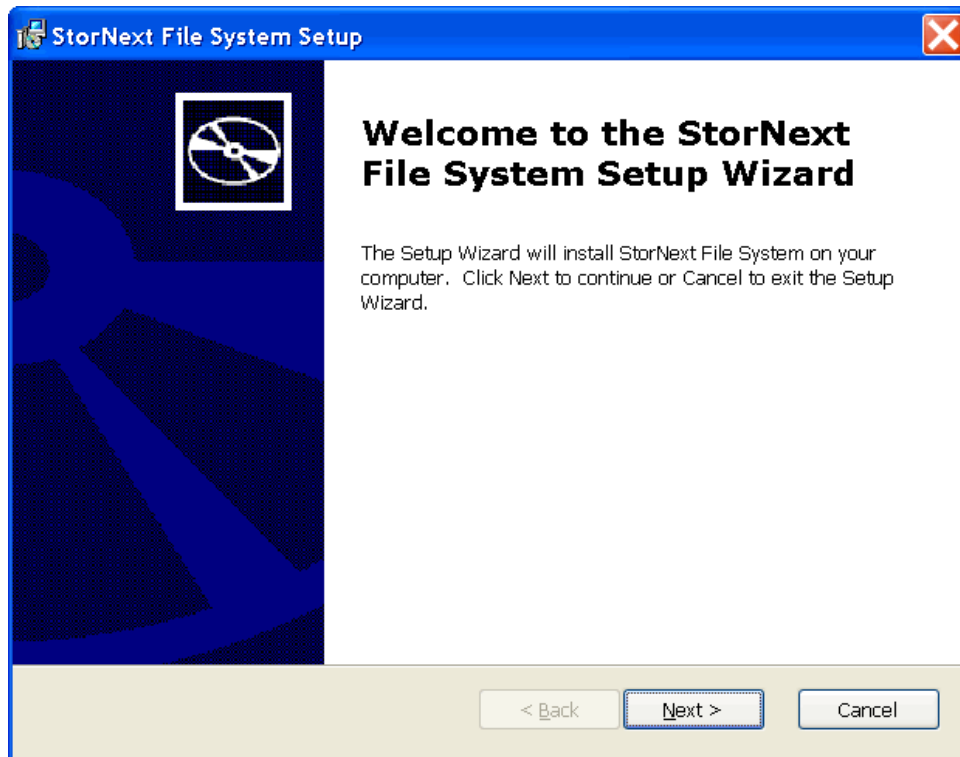
1. Download the StorNext client software from the MDC as described in [Download the StorNext Client Software on page 146](#).

i Note: If you are running SNFS-only on a Windows machine or have an Xsan MDC, you cannot download the StorNext client software from the MDC. Instead, you must obtain the StorNext client software from the StorNext installation DVD. Copy one of the following files from the DVD and place it in an accessible location on the client system, then proceed to the next step.


- **fs_only/snfs_client_Windows_i386.zip** (for 32-bit systems)
 - **fs_only/snfs_client_Windows_x86_64.zip** (for 64-bit systems)
2. If necessary, unzip the client software file you downloaded.
 3. Open the unzipped folder and double-click the client software installer file. This file is named **SnfsSetup.exe**. The **StorNext Installation** window appears.



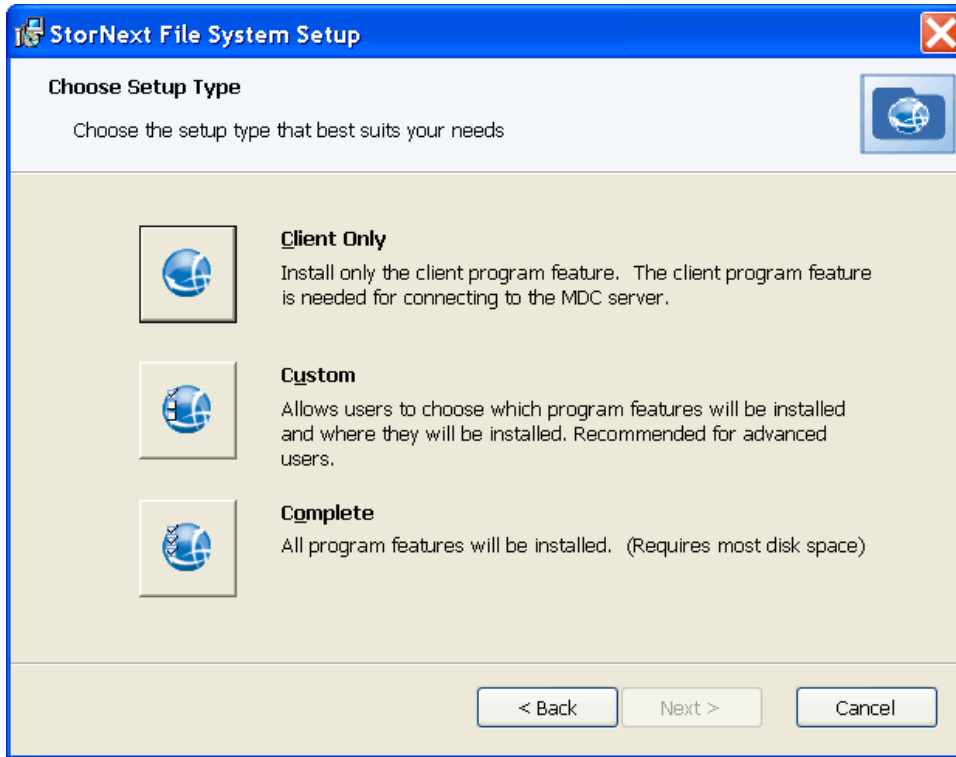
4. Click **Install StorNext** to begin installation. The **End User License Agreement** dialog appears.
5. Read the license. If you agree to the terms, click **I accept the terms in the License Agreement** and click **Next** to continue the installation process. The StorNext File System setup wizard launches.



6. Click **Next** to continue. The **Customer Information** window appears.



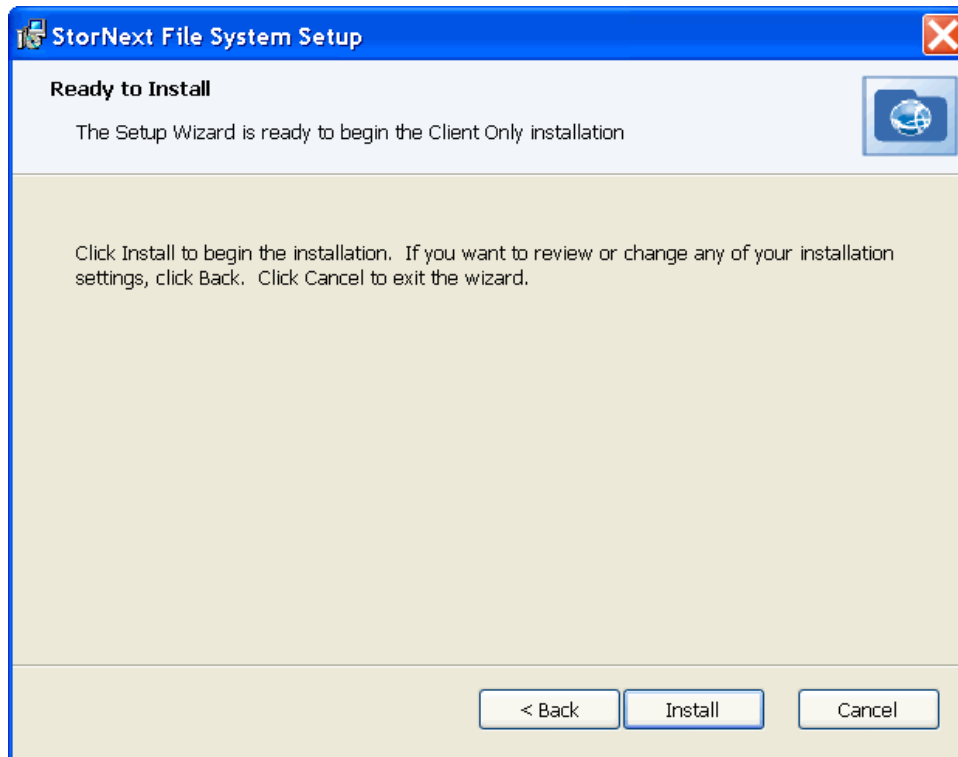
7. Type your name and the name of your company or organization in the boxes, and then click **Next** to continue. The **Choose Setup Type** window appears.




8. Click one of the following installation options. Since this is a client-only install, you should select **Client Only** or **Custom**.

- **Client Only:** Installs only the client program feature and help files in the default location.
- **Custom:** Allows you to select which program features are installed and where to install the program features.
- **Complete:** Installs all program features.


The **Ready to Install** window appears.




9. Click **Install** to continue. Wait while the setup wizard installs StorNext. When installation is complete, the **Finish** window appears. It may take several minutes for the installation to finish.



10. Click **Finish** to close the setup wizard. The **Installation Complete** dialog box opens.



11. Click **OK** in the **Installation Complete** dialog box. The **StorNext Installation** dialog appears with options to **Add/Remove Components**, **Remove StorNext**, or **Configure StorNext**.



12. The installation is complete. Click one of the following options:

- **Add/Remove Components:** The components you selected during the install are already added, but if you changed your mind and want to select different components, you can do so now. See [Configure the StorNext Client Software on Windows below](#) for details.
- **Remove StorNext:** Removes StorNext file system. For more information, see [Uninstall StorNext on Windows on page 195](#).
- **Configure StorNext:** Configures the StorNext client. For more information, see [Configure the StorNext Client Software on Windows below](#).
- **Exit:** Exits without making any further changes at this time (you can always perform configuration later).

Configure the StorNext Client Software on Windows

Configuring StorNext clients on Windows is similar to configuring a StorNext File System, but with more limited configuration options.

To configure the StorNext client software on Windows operating systems, use the Windows configuration utilities.

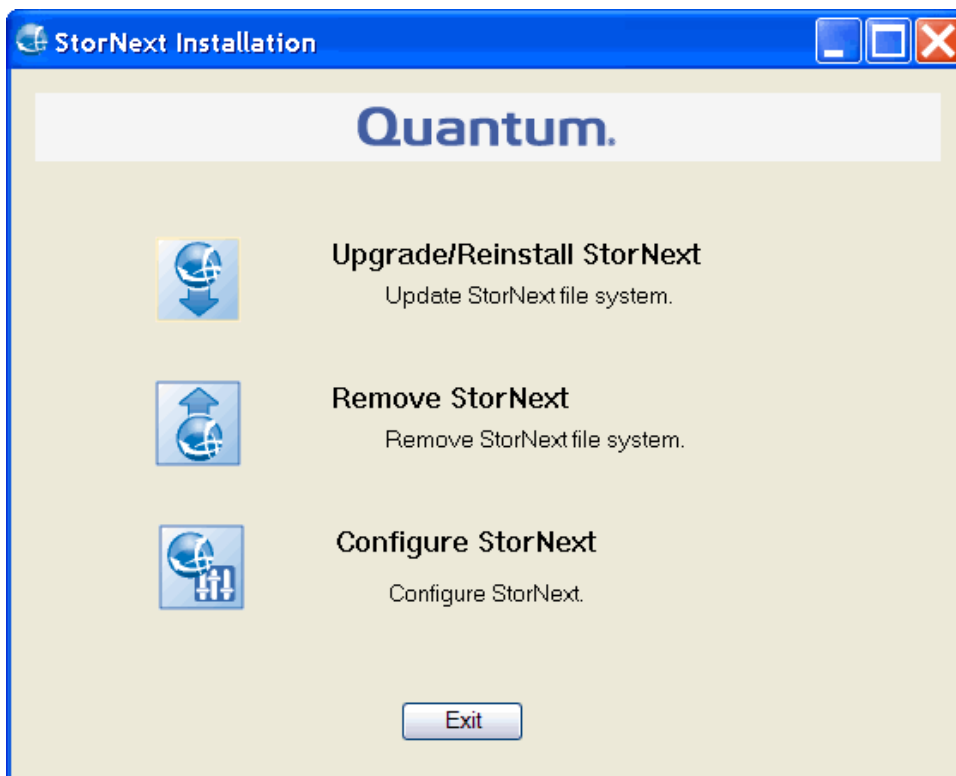
The following tools are provided with the StorNext Configuration utility, which is accessible from the StorNext Installation program:

- **Name Servers:** Configure StorNext name servers.
- **Import/Export Client Settings:** Import or export StorNext registry settings from or to a file.
- **Client Configuration:** Configure drive mappings and other StorNext client options.

Additional tools are accessible from the Windows **Start** menu (see [Additional Windows Configuration Utilities on the next page](#)).

Access the Configuration Utility

1. Open the **StorNext Installation** window by double-clicking the file named **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems). The **StorNext Installation** window displays appears.



2. Click **Configure StorNext**. The **StorNext Configuration** screen appears.



3. Click one of the following options:

- **Name Servers:** This feature is described in [Step 5: Configure a Client on page 100](#).
- **Import/Export Client Settings:** This feature is described in [Restore a Previous StorNext Client Configuration on page 174](#).
- **Client Configuration:** This feature is described in [Configure a StorNext File System on Windows on page 82](#).

Additional Windows Configuration Utilities

The following additional configuration utilities are accessible from the Windows **Start** menu (**Start > All Programs > StorNext File System**).

For instructions on using any of these utilities, see the online help accessible from **Start > All Programs > StorNext File System > StorNext Help**.

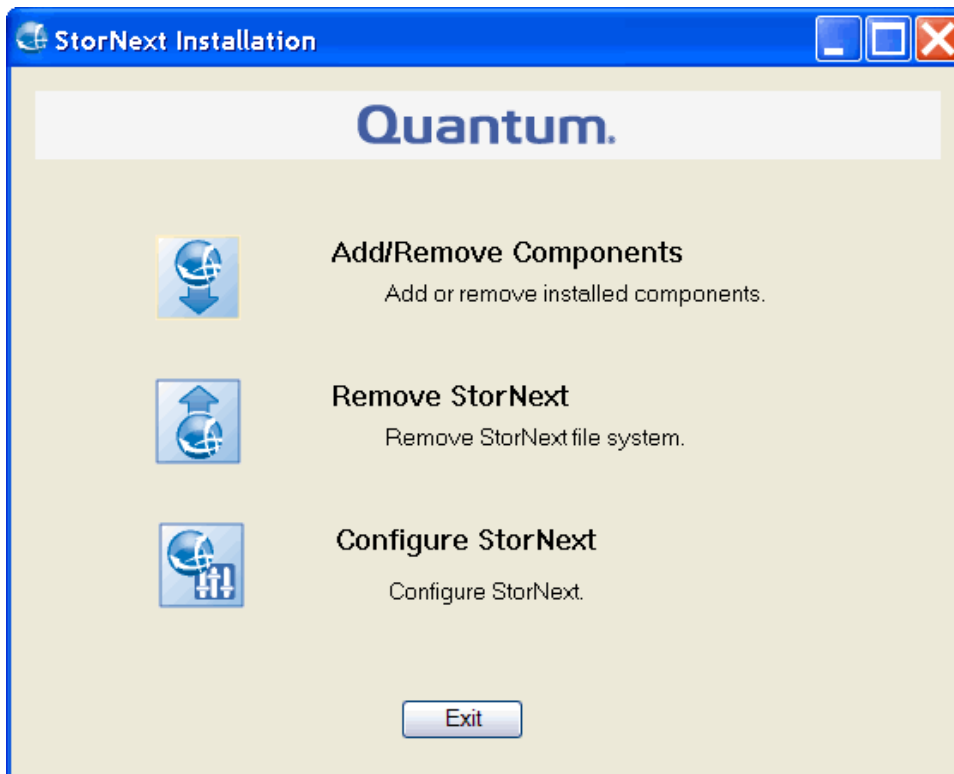
- **License Identifier:** Use this utility to display the host license identifier. The host license identifier is required to obtain a permanent license for StorNext. For details, see [The License Identifier Utility on page 117](#).

- **Services Start:** Use this utility to start all StorNext services on an MDC or StorNext client. For details, see [The Start File System Services Utility on page 121](#).
- **Services Stop:** Use this utility to stop all StorNext services on an MDC or StorNext client. For details, see [The Stop File System Services Utility on page 121](#).
- **Services Stop and Remove:** Use this utility to stop all StorNext services on an MDC or StorNext client, and also remove registry entries that automatically start the services on bootup. For details, see [The Stop and Remove File System Services Utility on page 121](#).
- **StorNext Help:** View StorNext help for these utilities and other commands.
- **Version Information:** View information about the currently installed version of the StorNext server and/or client software, such as the version number, build number, and platform. For details, see [The Version Information Utility on page 121](#).

Add or Remove StorNext Components


After installing the StorNext client, you may wish to add or remove components. If you installed the client only, you can add or remove the online help files or the StorNext client.

1. If you just completed installing the client, the **StorNext Installation** window displays. The components you selected during the install are already added, but if you changed your mind and want to select different components, you can do so now.



If you previously installed StorNext and would like to add or remove components, display the **StorNext Installation** window by double-clicking the file named **SnfsSetup32.exe** (32-bit systems)

or **SnfsSetup64.exe** (64-bit systems). Note that the screen is slightly different from the image above — instead of **Add/Remove Components**, the button you need to click is called **Upgrade/Reinstall StorNext**.



2. Click **Add/Remove Component** or **Upgrade/Reinstall StorNext**, depending on which screen you are looking at. The **Component Selection** screen appears.




3. Select the check boxes of the components you want to add. Clear the check boxes of the components you want to remove. Click **Next** to proceed. The StorNext EULA appears.
4. Read and accept the EULA by selecting “I accept the terms in the License Agreement” and clicking **Next**. It may take several minutes for the selected components to finish being added or removed.
5. When a message informs you that the process is complete, click **OK**.
6. Click **Exit**.

Restore a Previous StorNext Client Configuration


StorNext client configuration files are automatically saved when you remove StorNext, or can be manually saved for backup purposes.

If you want to restore a saved client configuration file, you can import it using the StorNext Installation tool. This configures StorNext using the settings from a previous installation.

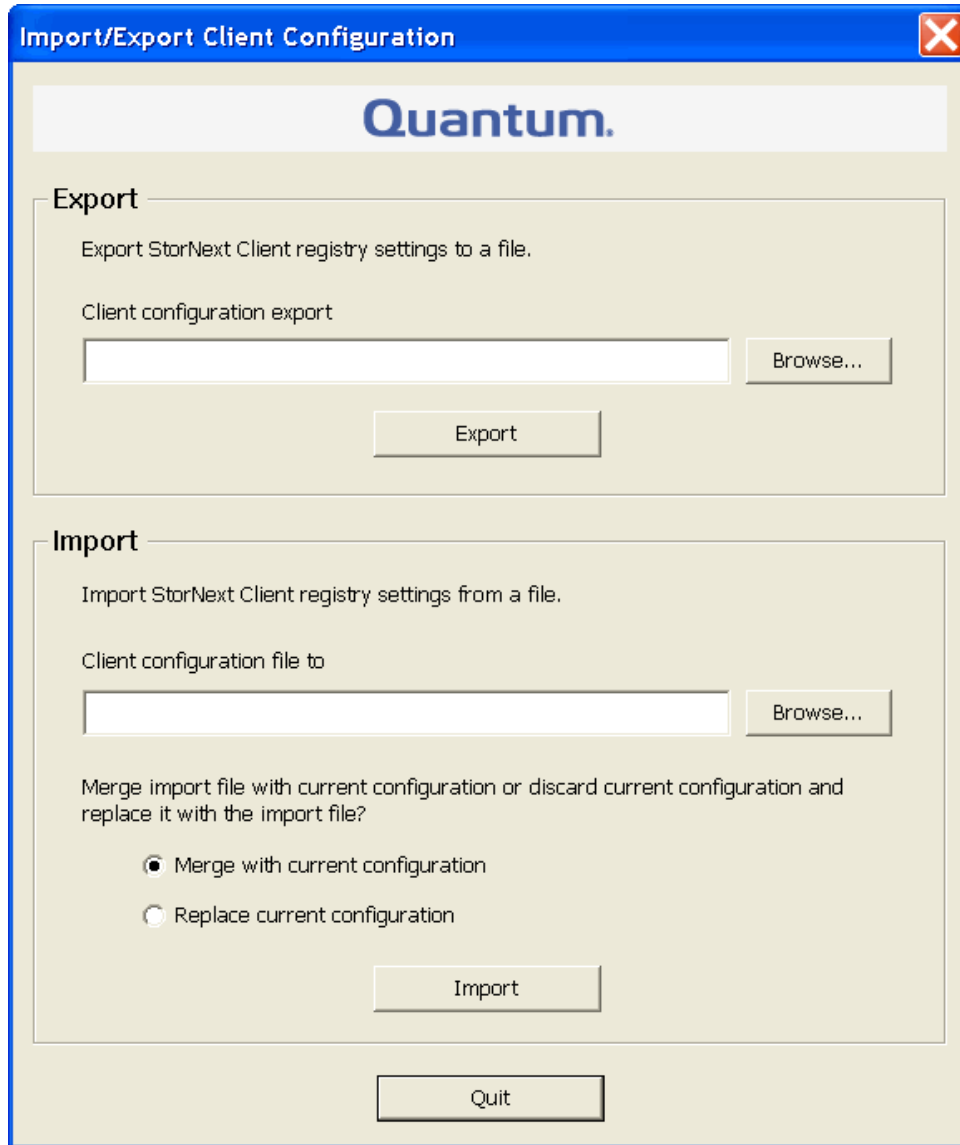
1. Start the StorNext installation tool if it is not already running by double-clicking the file named **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems). See **Step 1** and **Step 2** of [Install the StorNext Client Software on Windows on page 162](#). The **StorNext Installation** window appears.



2. Click **Configure StorNext**. The **StorNext Configuration** window appears.



3. Click **Import/Export Client Settings**. The **Import/Export Client Configuration** window appears.



4. Under **Import**, click **Browse**. Locate the client configuration (*.reg) file to import, and then click **Open**. Client configuration files saved during removal of a previous version of StorNext are located in the following directory by default:

```
C:\Program Files\StorNext\config\
```

5. Click an option for handling current configuration settings:
 - **Merge with current configuration:** The imported configuration is merged with the current configuration settings. Imported mount information is added to, or overwrites, existing information.

- **Merge with current configuration:** The imported configuration is merged with the current configuration settings. Imported mount information is added to, or overwrites, existing information.
6. Click **Import**, and then click **Yes** to confirm. A message appears informing you the configuration settings were successfully added to the registry.
 7. Click **OK**, and then click **Quit** to close the **Import/Export Client Configuration** window.

Upgrade or Reinstall StorNext Clients

The **Upgrade/Reinstall** button allows you to upgrade to a new version of StorNext client software, or reinstall the current version. You might want to reinstall if you accidentally removed some components and need to get them back. During an upgrade/reinstall, the old software is removed and the new software is installed. Your previous configuration settings remain intact.


1. To upgrade or reinstall a StorNext client after the initial installation, download the new client software from the MDC as described in [Download the StorNext Client Software on page 146](#).

Exception

If you are running SNFS-only on a Windows machine, you cannot download the client software from the MDC. Instead, you will need to obtain the client software from the StorNext installation DVD. Copy one of the following files from the DVD and place it in an accessible location on the client system, then proceed to the next step.

- **fs_only/snfs_client_Windows_i386.zip** (for 32-bit systems)
- **fs_only/snfs_client_Windows_x86_64.zip** (for 64-bit systems)

2. If necessary, unzip the client software file you downloaded.
3. Open the unzipped folder and double-click the client software installer file. This file is named **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems). The **StorNext Installation** screen appears.



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



5. 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.

6. The StorNext EULA appears.
7. Read and accept the EULA by selecting “I accept the terms in the License Agreement” and then clicking **Next**.
8. It may take several minutes for the selected components to finish being added or removed.
9. When a message informs you that the installation is finished, click **OK**.



Chapter 6: Configure LAN Clients and Gateway Servers

This chapter contains the following topics:

| | |
|--|-----|
| Configure LAN Clients and Gateway Servers | 181 |
| Configure a StorNext LAN Client on Linux | 182 |
| Configure a Gateway Server on Linux | 183 |
| Linux LAN Clients or Gateway Servers in HA Environments | 186 |
| Configure LAN Clients or Gateway Servers on Windows | 186 |
| Configure Gateway Functionality on a Metadata Controller (MDC) | 187 |

Configure LAN Clients and Gateway Servers

Linux or Windows StorNext clients can be configured as LAN (local area network) clients or gateway servers (including the StorNext G300 Gateway Appliance and the M660 Metadata Appliance). The StorNext G300 is a dedicated gateway server. The M660 is a meta-data controller (MDC) that can be configured as a gateway server (in non-Lattus systems only).

Traditional StorNext SAN clients access data storage (hard disk drives) in the SAN using a high throughput, low-latency fibre channel connection. In addition to this fibre channel connection, StorNext also provides access for LAN clients through a StorNext G300 or M660, or SAN client configured as a gateway for those LAN clients. In this environment, StorNext LAN clients access data disks over a network instead of directly accessing disks on the SAN.

A StorNext client running on Linux or Windows can be configured as a LAN client or a gateway. This is a licensed feature.

-
- i Note:** A macOS X MDC is not supported as a proxy server for an FX LAN client.
 - **Gateway:** A gateway is a StorNext client which has direct access to data disks (using fibre channel) on the SAN and makes the shared storage pool available to LAN clients over a LAN connected to that gateway.
 - **LAN Client:** A LAN client cannot have direct access to data disks and instead accesses the shared storage pool by connecting to a gateway over a network (LAN clients have full StorNext client functionality).
-
- i Note:** If the **diskproxy** option is set to **client**, then the mount may use a Proxy Server to do its data I/O. If the host has SAN connectivity to some or all of the disks in the file system, then those disks will be accessed via the SAN connection, not the network. This client is then referred to as a disk proxy hybrid client. When SAN connectivity is used, the server license on the MDC will be charged for this mount. If it is desired that this client use the network for the mount, then the disks should be made unavailable to this host or the **cvpaths** file should be configured to prevent StorNext from using the directly attached disks. The **who** subcommand of **cvadmin** shows the type of proxy mount.
-
- i Note:** A StorNext client configured as a gateway requires a dedicated system. Do not run other applications or services on a system configured as a gateway.
-
- i Note:** You can enable and configure gateway functionality for an MDC connected to a Lattus system. StorNext only supports a single DLC network for each file system. A Lattus system and an MDC are normally divided into two separate segments which makes it a challenge to enable a gateway; a challenge for a single network segment. Additionally, you can experience slower than normal throughput if you share both the 10G Lattus and gateway network segments. You must evaluate your scenario depending on the S3 configuration relating to the object segment size and your environment workload.

Review the sections below for additional information.

- [Configure a StorNext LAN Client on Linux below](#)
- [Configure a Gateway Server on Linux on the next page](#)
- [Linux LAN Clients or Gateway Servers in HA Environments on page 186](#)
- [Configure LAN Clients or Gateway Servers on Windows on page 186](#)
- [Configure Gateway Functionality on a Metadata Controller \(MDC\) on page 187](#)

Configure a StorNext LAN Client on Linux

To configure a StorNext client as a StorNext LAN client, edit mount options in the **/etc/fstab** file.

1. Stop the StorNext client. At the command prompt, type:

```
/etc/init.d/cvfs stop
```

2. Configure the client to mount a file system as a LAN client. To do this, edit the **/etc/fstab** file so that it contains the following line:

```
<file_system> <mount_point> cvfs rw,noauto,diskproxy=client 0 0
```

where **<file_system>** is the name of the StorNext file system and **<mount_point>** is the directory path where the file system is mounted.

3. Start the StorNext client. At the command prompt, type:

```
/etc/init.d/cvfs start
```

Configure a Gateway Server on Linux

To configure a StorNext SAN client as a gateway server, edit mount options in the **/etc/fstab** file, and then configure gateway server options.

i Note: The StorNext G300 and M660 come pre-configured as gateways. Do not use these commands to configure those systems.

Configure the Gateway Server

1. Stop the StorNext SAN client. At the command prompt, type:

```
/etc/init.d/cvfs stop
```

2. Configure the client to mount a file system as a gateway. To do this, edit the **/etc/fstab** file so that it contains the following line:

```
<file_system> <mount_point> cvfs rw,noauto,diskproxy=server 0 0
```

where <file_system> is the name of the StorNext file system and <mount_point> is the directory path where the file system is mounted.

3. Change to the `/usr/cvfs/bin/` directory.
4. Launch the StorNext gateway configuration utility. At the command prompt, type:

```
./sndpscfcfg -e
```

i Note: To get help for the `sndpscfcfg` utility, at the command prompt, type: `./sndpscfcfg -h`

The gateway configuration utility appears.

```
# Disk Proxy Server Configuration
# -----

# =====
# Tunable Parameters
# =====

# Un-comment and modify any of the following lines to override the default
# values for tunable parameters.
#tcp_window_size_kb 64
#transfer_buffer_size_kb 256
#transfer_buffer_count 16
#server_buffer_count 8
#daemon_threads 8

# =====
# Interface Configuration
# =====

# One or more interfaces must be configured for Disk Proxy Server I/O.

# Un-comment one (and only one) of the 2 lines below to allow Disk Proxy
# Server I/O on interface 'eth0' and to select the address to listen for
# connections on.
#interface eth0
interface eth0 address fec0::215:f2ff:fe68:afa2
```

5. Edit the default gateway configuration:
 - **(Required)** Uncomment (delete the `#` next to) the line that corresponds to the network interface used by the gateway server when connecting to LAN clients. (For best performance, do not choose the network used for metadata.)
 - **(Optional)** To change a tunable parameter from the default value, uncomment the parameter, and then specify a new value. The table below describes the tunable parameters.

i Note: The gateway configuration utility uses the default text editor. If a default text editor has not been specified, the utility uses the `vi` editor.

| Field / Button | Description |
|--------------------------------------|---|
| <code>tcp_window_size_kb</code> | The TCP window size (in KB) used by the gateway server. (Default: 64) |
| <code>transfer_buffer_size_kb</code> | The transfer buffer size (in KB) used by the gateway server. A larger buffer may increase performance for larger files. (Default: 256) |
| <code>transfer_buffer_count</code> | The number of transfer buffers used by the gateway server. This parameter is used only by Windows servers and clients. Linux servers pass the value of this parameter to Windows clients. (Default: 16) |
| <code>server_buffer_count</code> | The number of I/O buffers allocated for each network interface on the gateway server. This parameter is used only by Linux servers. (Default: 8) |
| <code>daemon_threads</code> | The maximum number of daemon threads used by the gateway server. (Default: 8) |

- When you are done making edits, type `:x` to save and close the configuration file. The configuration is saved in the file `dpserver` in the `/usr/cvfs/config/` directory.

i Note: To edit this file at a later date, you can run the `sndpscfcg` utility again or edit the `dpserver` file directly.

- Start the StorNext SAN client, which is now functioning as a gateway server. At the command prompt, type:

```
/etc/init.d/cvfs start
```

Create an Override Configuration

By default, the settings specified in the `dpserver` configuration file are used for all file systems mounted by the gateway server. To use different settings for a file system, create an override configuration file for that file system. At the command prompt, type:

```
/usr/cvfs/bin/sndpscfcg -E <file_system>
```

where `<file_system>` is the name of the file system to which the override settings apply. Edit the configuration file, and then save and exit. A new configuration file with the name `dpserver. <file_system>` is created in the `/user/cvfs/config/` directory.

Use the LAN Client and Server Commands

Use the **cvadmin** commands described below to see information about gateway servers and LAN clients. For more information about using **cvadmin** commands, see the *MAN Pages Reference Guide*.

| Command | Description |
|---|---|
| proxy | Lists all gateway servers for the currently selected file system. |
| proxy long | Lists all gateway servers for the currently selected file system, as well as tunable parameters and attached disks. |
| proxy who <LAN_client> | Lists gateway servers connected to the specified LAN client, as well as current I/O statistics. |
| proxy who <gateway_server> | Lists LAN clients connected to the specified gateway server, as well as current I/O statistics. |

Linux LAN Clients or Gateway Servers in HA Environments

When a StorNext High Availability (HA) configuration is used and the StorNext LAN Client/Gateway Server feature is configured on the HA server nodes, each node in the HA pair should be configured separately and independently. This allows for the best and highest availability of the LAN client capability.

Ideally, each node in the HA pair should have the same number of network interfaces and be on the same networks. The dpserver configuration should be done before converting an MDC pair into an HA pair.

When deduplication/replication is enabled, Virtual IPs (VIPs) are configured and applied to the active blockpool server. In such a setting, the network interface(s) with VIP configured can potentially have more than one IP addresses (physical and virtual IPs). If such an interface is also used for a LAN client, use the interface with the physical IP address to specify the network interface.

Configure LAN Clients or Gateway Servers on Windows

To configure a Windows client as a gateway server or a LAN client, use the Windows Client Configuration utility (see [Configure a StorNext File System on Windows on page 82](#)).

Configure Gateway Functionality on a Metadata Controller (MDC)

The StorNext M660 Metadata Appliance and appropriately provisioned MDCs can be configured to also operate as a LAN gateway. The following instructions describe the steps to configure gateway functionality on MDCs.

- i Note:** Do not enable gateway functionality for M662s connected to Lattus systems. While there are no functional issues in enabling the LAN gateway functionality within a Lattus environment, observed throughput when moving data to Lattus through the 10GbE ports will be slower than expected if this feature is enabled.

Enable Gateway Functionality on an MDC

To enable gateway functionality:

1. Log on to the primary and, if applicable, secondary nodes using as the “**stornext**” user with the default password of **password**.
2. Run the **sudo rootsh** command to gain root access.
3. At the command prompt, type:

```
/usr/cvfs/bin/sndpscfcfg -e
```

4. Select an appropriate interface and set tuning parameters. This creates or updates **/usr/cvfs/config/dpserver**.

- i Note:** This must be done on both the primary and secondary nodes, if applicable. This file is not synchronized from the primary to the secondary node.

5. For HA pairs, update the **/usr/cvfs/config/sngateway_config** file so that both files include the IP address of both MDC nodes.
6. On the primary node, edit the **/etc/fstab** file. Add the following to the mount option on the file system to be used as the gateway:

```
diskproxy=server
```

For example, the file system mount option:

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

Becomes:

```
snfs1 /stornext/snfs1 cvfs rw,noauto,diskproxy=server 0 0
```

i Note: If the file system being configured is a managed file system, stop the TSM processes before unmounting the file system. At the command prompt, type:

```
TSM_control stop
```

7. On the primary node, un-mount and re-mount the file system. At the command prompt, type:

Important Consideration for NAS

If your file system contains a NAS share or is a NAS cluster mount point, then you might have to disable your NAS service before you un-mount your file system. At the prompt, run the following command to **disable** your NAS service:

```
/usr/cvfs/lib/snnas_control stop
```

If you disable your NAS service, then you must enable your NAS service after you re-mount your file system. At the prompt, run the following command to **enable** your NAS service:

```
/usr/cvfs/lib/snnas_control start
```

```
umount /stornext/snfs1
```

```
mount snfs1
```

8. If Gateway functionality is being configured on a single-node MDC continue to **Step 12**. Otherwise continue on with the next steps to complete the configuration on the secondary MDC node.
9. On the secondary node, wait for the **/etc/fstab** change to synchronize with the primary node. It should take less than two minutes for the change to appear in this file.
10. When you see **diskproxy=server** change in the **fstab** file on the secondary, un-mount and re-mount the file system on the secondary.
11. Run the mount command with no parameters on both the primary and secondary nodes. Verify that the file system is mounted and that the **diskproxy=server** mount option is present. Example output:

```
/dev/cvfsctl2_snfs1 on /stornext/snfs1 type cvfs  
(rw,diskproxy=server,sparse=yes)
```

12. If the TSM processes were stopped, restart the TSM process from the primary MDC. At the command prompt, type:

```
TSM_control start
```

Disable Gateway Functionality on an MDC

To disable gateway functionality on these devices:

1. Log into the primary MDC node, and, if applicable, the secondary MDC node as root.
2. On the primary node, edit the `/etc/fstab` file. Remove the following from the mount option on the file system to be used as the gateway:

```
diskproxy=server
```

For example, the file system mount option:

```
snfs1 /stornext/snfs1 cvfs rw,noauto,diskproxy=server 0 0
```

Becomes:

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

i Note: If the file system being configured is a managed file system, stop the TSM processes before unmounting the file system. At the command prompt, type:

```
TSM_control stop
```

3. On the primary node, un-mount and re-mount the file system. At the command prompt, type:

Important Consideration for NAS

If your file system contains a NAS share or is a NAS cluster mount point, then you might have to disable your NAS service before you un-mount your file system. At the prompt, run the following command to **disable** your NAS service:

```
/usr/cvfs/lib/snnas_control stop
```

If you disable your NAS service, then you must enable your NAS service after you re-mount your file system. At the prompt, run the following command to **enable** your NAS service:

```
/usr/cvfs/lib/snnas_control start
```

```
umount /stornext/snfs1
```

```
mount snfs1
```

4. If Gateway functionality is being removed from a single-node MDC continue to **Step 8**. Otherwise continue on with the next steps to complete the configuration on the secondary MDC node.
5. On the secondary node, wait for the **/etc/fstab** change to synchronize with the primary node. It should take less than two minutes for the change to appear in this file.
6. When you see **diskproxy=server** removed from the **fstab** file on the secondary, un-mount and re-mount the file system on the secondary.
7. Run the mount command with no parameters on both the primary and secondary nodes. Verify that the file system is mounted and that the **diskproxy=server** mount option has been removed. Example output:

```
/dev/cvfsct12_snfs1 on /stornext/snfs1 type cvfs (rw,sparse=yes)
```

8. If the TSM processes were stopped, restart the TSM process from the primary MDC. At the command prompt, type:

```
TSM_control start
```



Appendix A: Uninstall StorNext

This appendix contains the following topics:

Uninstall StorNext

If you no longer need to use StorNext, uninstall the software to remove it from your system. To uninstall StorNext, use the appropriate method for your operating system:

- [Uninstall StorNext on Linux below](#)
- [Uninstall StorNext on Windows on page 195](#)

Uninstall StorNext on Linux

To uninstall StorNext on a Metadata Controller (MDC), use the installation script (see [Uninstall StorNext Storage Manager and StorNext File System on the next page](#)).

To uninstall the StorNext client software, use the appropriate method for your operating system (see [Uninstall the StorNext Client Software on page 193](#)).

Uninstall StorNext Storage Manager and StorNext File System

To uninstall StorNext File System and Storage Manager on an MDC running Unix or Linux, run the installation script with the **-remove** or **-removeall** option. To launch the script, access the folder for your operating system on 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. 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.
5. Do one of the following steps:
 - **(Recommended)** Completely uninstall StorNext and return the system to a pristine state. At the command prompt, type:

```
./install.stornext -removeall
```

- Uninstall StorNext binaries but leave files that are useful if you are going to perform an HA upgrade (such as license files, some configuration files, and log files). At the command prompt, type:

```
./install.stornext -remove
```

The installation script runs.

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

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

Enter your choice <4>: 
```

6. Type **2** and press **Enter**.

Uninstall the StorNext Client Software

To uninstall the StorNext client software, unmount all file systems and stop StorNext software. Then remove the client software package using the appropriate command for your operating system.

i Note: This procedure removes the client software only. To uninstall StorNext File System and Storage Manager, see [Uninstall StorNext Storage Manager and StorNext File System on the previous page](#).

1. Unmount all file systems. To unmount a file system, at the command prompt, type:

```
umount <mount_point>
```

where **<mount_point>** is the location where the file system is mounted.

2. Stop the StorNext client software. At the command prompt, type:

- **HP-UX**

```
/sbin/init.d/cvfs fullstop
```

- **IBM AIX**

```
/etc/rc.cvfs fullstop
```

- **Solaris**

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

- **Others**

```
/etc/init.d/cvfs fullstop
```

3. Remove the StorNext software packages using the correct commands for your operating system.

- **Red Hat and SUSE Linux:**

```
rpm -e <package_names>
```

where <package_names> are the names of the client software packages that were installed.

- **Debian and Ubuntu Linux:**

```
dpkg -P <package_names>
```

where <package_names> are the names of the client software packages that were installed.

- **Solaris:**

```
pkgrm ADICsnfs
```

- **HP-UX:**

```
swremove -x mount_all_filesystems=false snfsclient
```

- **IBM AIX:**

```
installp -u <package_name>
```

where <package_name> is the name of the client software package that was installed.

- **SGI IRIX:**

```
Inst  
remove cv_base  
remove cv_clnt  
go
```

i Note: If you do not know the package name, you can download the client software for your operating system and check the package name. See [Download the StorNext Client Software on page 146](#).

Uninstall a StorNext RPM Server (SNFS server-only)

To uninstall the SNFS-only server install, unmount all file systems and stop the SNFS software. Then remove the SNFS software package using the rpm command.

i Note: This procedure only removes the RPM software. To uninstall StorNext File System with Storage Manager, see [Uninstall StorNext Storage Manager and StorNext File System on page 192](#).

1. Unmount all SNFS file systems using the following at the command prompt, where `<mount_point>` is the location where the file system is mounted:

```
umount <mount_point>
```

2. Stop the StorNext client software. At the command prompt, type:

```
/etc/init.d/cvfs fullstop
```

3. Remove the StorNext software package using the following commands, where `<package_names>` are the names of the client software packages that were installed:

```
rpm -e <package_names>
```

Uninstall StorNext on Windows

There are two methods to uninstall StorNext (server or client) on Windows:

- [Uninstall StorNext Using the StorNext Software on the next page](#)
- [Uninstall StorNext Using the Windows Control Panel on page 197](#)

Either method removes all installed StorNext components from your system, but will save the previous client configurations, debug file, and RAS files. All of the saved files are located in the following directory:

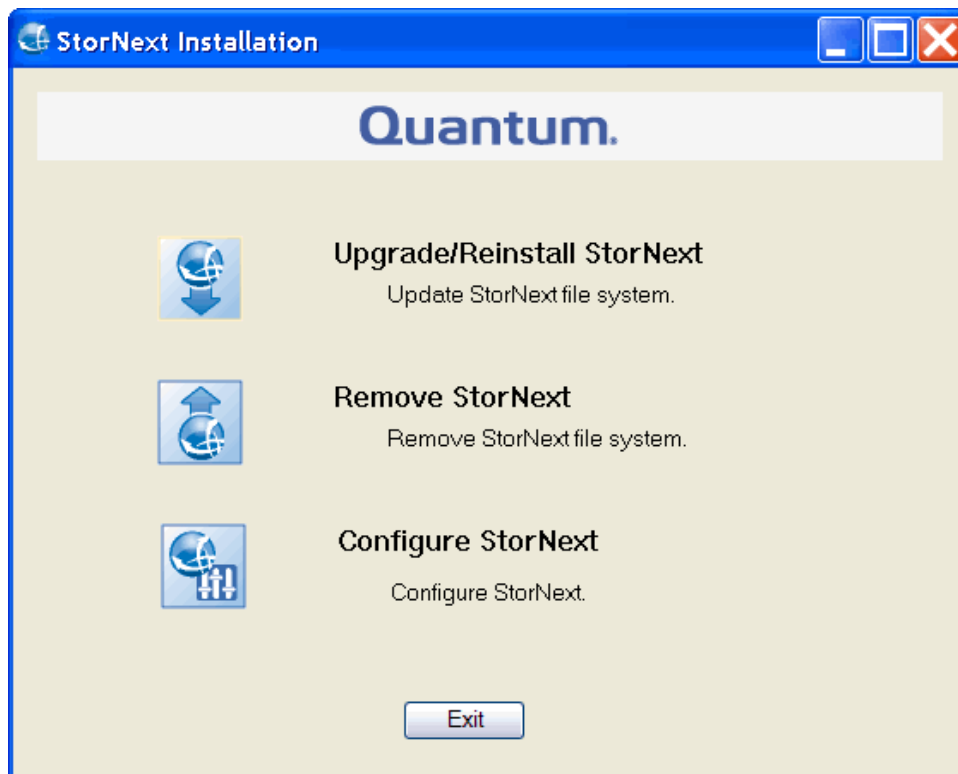
```
C:\Program Files\StorNext\config\
```

Note: Systems prior to StorNext 3.0 used the default directory **C:\SNFS\config**.

Uninstall StorNext Using the StorNext Software

Perform the following procedure to remove StorNext using the StorNext software.

1. Start the StorNext installation tool if it is not already running by double-clicking the file named **SnfsSetup32.exe** (32-bit systems) or **SnfsSetup64.exe** (64-bit systems). The **StorNext Installation** window appears.



2. Click **Remove StorNext**.
3. Click **OK** to confirm you want to remove StorNext and all of its components. A dialog box informs you that the current client configuration has been saved and lists the location. You can use this to restore the client configuration later if you wish.
4. Note the name and location of the saved configuration file, and then click **OK**.
5. If prompted, click **Yes** to continue, and then click **OK** to finish the removal.

6. When the removal is complete, you will be notified that you must reboot for the change to take effect. Click **Yes** to reboot the system immediately, or **No** to reboot later.

Uninstall StorNext Using the Windows Control Panel

To uninstall StorNext (server or client) on a Windows system, use the **Add or Remove Programs** feature located on the **Windows Control Panel**.

i Note: On Windows Vista, Windows Server 2008 and later, the **Control Panel** feature is called **Programs and Features**.

1. Open the **Windows Control Panel**.
2. Double-click **Add or Remove Programs** (or **Programs and Features** on Windows Vista or Windows Server 2008 and later). The **Add or Remove Programs** (or **Programs and Features**) window appears.
3. In the list of currently installed programs, click **StorNext File System**, and then click **Remove**.
4. Click **Yes** to confirm the removal.
5. After the uninstall process is complete, reboot the system.



Appendix B: Configuration Tasks

This appendix contains the following topics:

Configuration Tasks

Use the following information to perform specific file system configuration tasks that do not apply to a typical StorNext set up.

- [Switch to a Named Cluster Using the Command Line Interface on the next page](#)
- [Share a Library Between Applications \(Linux Only\) on page 200](#)
- [Manually Modify the Storage Manager Configuration on page 201](#)
- [Connect an Xsan Client to a StorNext MDC on page 203](#)
- [Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client \(DLC\) on page 221](#)
- [View a Fibre Channel Device on page 223](#)
- [StorNext Log Rolling Administration on page 223](#)
- [Avoid Poor System Performance on page 224](#)
- [Avoid Not Saving cvmkfs Coredump Files on page 225](#)
- [Import a Secure Sockets Layer \(SSL\) Certificate on page 225](#)

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.

Share a Library Between Applications (Linux Only)

If two or more software applications (on one or multiple machines), are configured to use the same library, then you must use the following procedure to enable the library to be shared and to prevent media being overwritten when these applications are used. This task must be performed when the Configuration Wizard (CW) is run during initial StorNext setup and configuration. For more information on using the CW, see [Configure StorNext Using the StorNext GUI on page 37](#).

⚠ Caution: Quantum strongly recommends using library partitioning features to present each application using the library with its own virtual library. Unless partitioning is used, two applications sharing the robot arm, drives and slots may over-write each others data.

Use the following instructions *only* if library partitioning is not an option.

i Note: For library sharing to be effective, the steps in the following procedure must be performed in the order in which they are presented.

Enable the Library to be Shared

i Note: You must have root privileges to perform this task.

1. When you match hardware devices to their correct slots on the **Matched Devices With Slots** screen (you will see this screen during the execution of Step 4 of the Configuration Wizard, **Add Tape Drive**), only configure the tape drives that StorNext will use. You do not need to configure all tape drives listed as part of the library.
2. Once the tape drives are added and configured in Step 4 of the CW (**Add Tape Drive**), but before continuing to Step 5 of the CW (**Add Media**), go to a Metadata Controller and perform these steps:
 - a. Change directory to the **config** directory. Type:

```
cd /usr/adic/MSM/internal/config
```

- b. List all files in the config directory. Type:

```
ls -l
```

- c. Using a text editor, edit the **media_file_<library_name>** file. For example:

```
vi media_file_<library_name>
```

The **media_file** contains a filter for barcode labels that are read by StorNext. In the example above, *<library_name>* is the library that was configured in Step 3, **Add Library**, of the CW.

- d. Set up the pattern matching of the media file.

Multiple entries may exist in the media file. Read the examples in the file to learn how to set up pattern matching.


These are examples of pattern matching:

scsi_archive 1111.. (brings in all media with the 1111 prefix)

scsi_archive AB9... (brings in all six-character media labels with the AB9 prefix)

scsi_archive CLN1.. (brings in all media that have the CLN1 prefix)

3. Continue to the **Add Media** section of the **Configuration Wizard**.

 **Caution:** After this procedure is done, StorNext will not use other media, but other applications may try to overwrite StorNext media. To avoid data loss, be sure you configure the other application to prevent overwriting StorNext data.

Manually Modify the Storage Manager Configuration

StorNext employs global parameters that define the Storage Manager configuration. Although the default values in these files are sufficient for most installations, you can edit the files to customize the StorNext software for your specific installation. This section provides a procedure that can be used to modify the following StorNext configuration files:

- /usr/adic/TSM/config/fs_sysparm_override
- /usr/adic/TSM/config/filesize.config
- /usr/adic/TSM/logs/log_params

The `fs_sysparm_override` Configuration File

The `fs_sysparm_override` configuration file is used to tune items related to system administration, drives and media, media manager interface, and other miscellaneous items. For more information about the parameters that can be placed into the `fs_sysparm_override` configuration file, refer to the `fs_sysparm.README` file.

⚠ Caution: Incorrect modification of the `fs_sysparm_override` configuration file can cause a serious, adverse effect on StorNext functionality. Before modifying this file, Quantum recommends that you contact Quantum Global Services.

⚠ Caution: If you have files larger than 100GB and are using LTO2 media, the `MED_SEG_OVER_LTO` parameter may be modified to a value less than or equal to 190G to reduce file fragmentation. This has the adverse effect of reducing the potential for parallel I/O for multiple file segments. Setting the `MED_SEG_OVER_LTO` parameter to a value larger than 190GB may result in allocation failures that prevent file movement to tape.

The `filesize.config` Configuration File

The `filesize.config` configuration file is used to control the file steering feature and has these characteristics:

- Allows the placement of files on different media types, based on the size of the files
- Specifies which drive pool a policy class should use when storing data
- Contains specific syntax and usage information
- Enables the system administrator to make changes without recycling the Tertiary Manager software
- Read the instructions in the `filesize.config.default` file for usage and syntax information.

The `log_params` Configuration File


The `log_params` configuration file is used to control various levels of trace logging. The file specifies each log level and how to enable and/or disable it.

Use the following procedure to modify the `fs_sysparm_override`, `filesize.config`, or `log_params` configuration files.

1. Use a text editor to open one of the following configuration files:
 - `/usr/adic/TSM/config/fs_sysparm_override`
 - `/usr/adic/TSM/config/filesize.config`
 - `/usr/adic/TSM/logs/log_params`
2. Locate the parameter you want to modify and replace the setting with a new, valid value. When editing a file, be sure to follow the format used by entries in the file. For example, in the `fs_sysparm_override`, or `log_params` files, all entries must be in the format: **name=value;**

3. Restart the Tertiary Manager software.
 - a. Stop the software by typing **TSM_control stop**.
 - b. Restart the software by typing **TSM_control start**.

Connect an Xsan Client to a StorNext MDC



Requirements

In order to connect Xsan Clients to StorNext MDCs, verify you are meeting the following requirements:

- StorNext MDCs and Xsan clients must be connected to the same SAN and be able to see the same LUNs
- StorNext MDCs must have **license.dat** file in the configuration directory of the system.
 - For Linux MDCs, the configuration directory is:

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

- For Windows MDCs, the configuration directory is:

```
C:\Program Files\StorNext\config\
```

- If the system configuration includes both Xsan and StorNext MDCs, and Xsan clients, **ALL** machines in the configuration must have the identical **.auth_secret** file. However, if the only MDCs in use are StorNext MDCs, the **.auth_secret** files should be removed from the Xsan clients, the StorNext clients, and the MDCs.

Important Considerations Regarding the .auth_secret File

If you add or remove the **.auth_secret** file, then all clients and MDCs fail to mount file systems until the following services are restarted.

For Linux clients:

1. On the **secondary** node, run the following command:

```
service cvfs stop
```

2. On the **primary** node, run the following command:

```
service cvfs stop
```

3. On the **primary** node, run the following command:

```
service cvfs start
```

4. On the **secondary** node, run the following command:

```
service cvfs start
```

For macOS clients:

- On a macOS client, run the following command:

```
sudo launchctl unload /System/Library/LaunchDaemons/com.apple.xsan.plist
```

Followed by the following command:

```
sudo launchctl load /System/Library/LaunchDaemons/com.apple.xsan.plist
```

- The **fsnameservers** file must be identical on all StorNext MDCs, StorNext clients, and Xsan clients.

For Systems Running StorNext 6.x (or earlier) and macOS 11.x (or later)

Mount the StorNext File System

Perform the following procedure to connect an Xsan 5.0 (or later) client to a StorNext MDC.

1. Generate a configuration profile:

```
mdc# cd /usr/cvfs/config  
mdc# /usr/cvfs/lib/make_profile mysan
```

2. Rename the resultant profile file to have a **.mobileconfig** extension instead of a **.configprofile** extension:

```
mdc# mv mysan.configprofile mysan.mobileconfig
```

3. Copy the profile to the Xsan client:

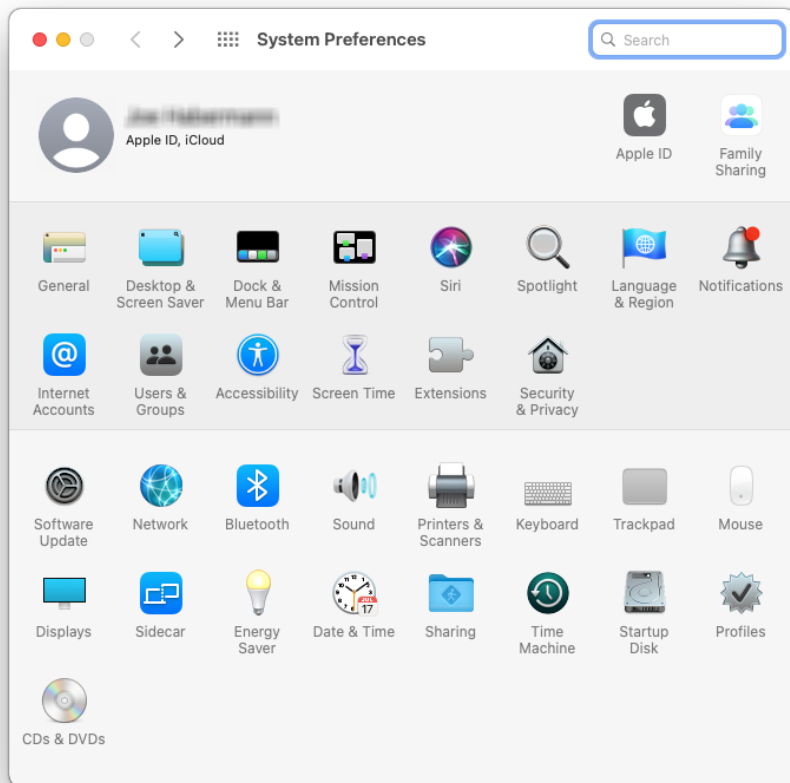
```
mdc# scp mysan.mobileconfig root@xsan50client:/tmp
```

4. On your macOS client (desktop), remove any existing profiles as follows:

i Note: If your system is a fresh install and does not contain an existing profile, then proceed to the next step to install the profile.


Appendix B: Configuration Tasks
Connect an Xsan Client to a StorNext MDC

a. Click **System Preferences** to display its dialog:



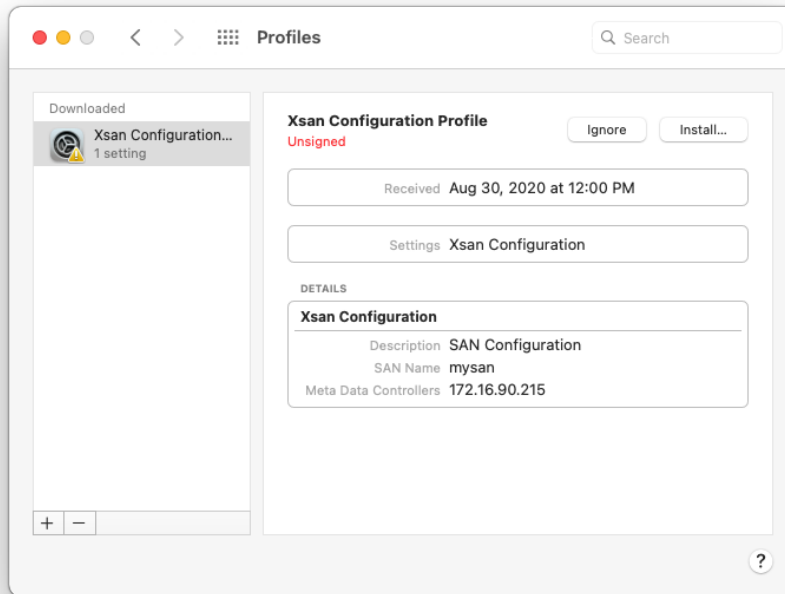
b. In the **System Preferences** dialog, click **Profiles** to display its dialog:

Appendix B: Configuration Tasks
Connect an Xsan Client to a StorNext MDC

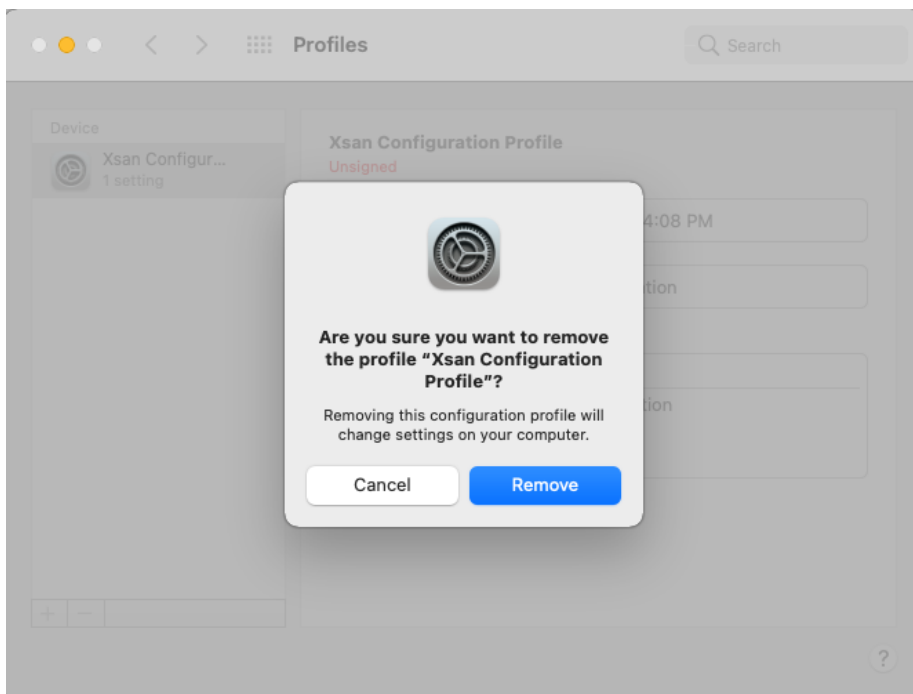


- c. On the left pane of the **Profiles** dialog, click the device profile you want to remove and then click - on the left pane:

Appendix B: Configuration Tasks
Connect an Xsan Client to a StorNext MDC

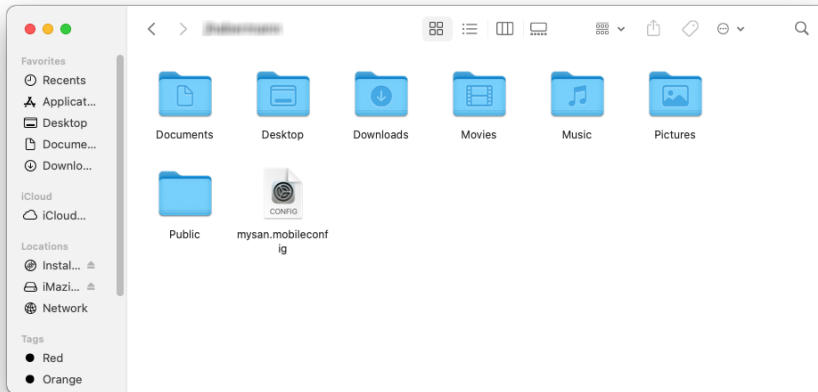


d. You are prompted to verify you want to remove the device profile. Click **Remove**:

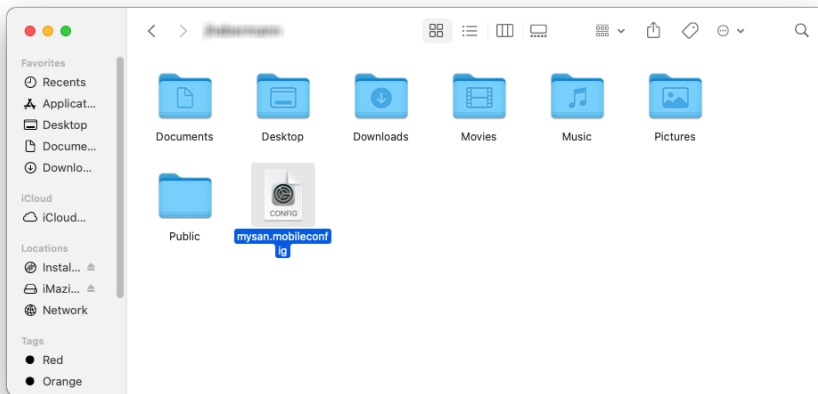


5. On your macOS client (desktop), install the new profile as follows:


- a. Click **Finder** and navigate to the **/tmp** directory where the profile is located:



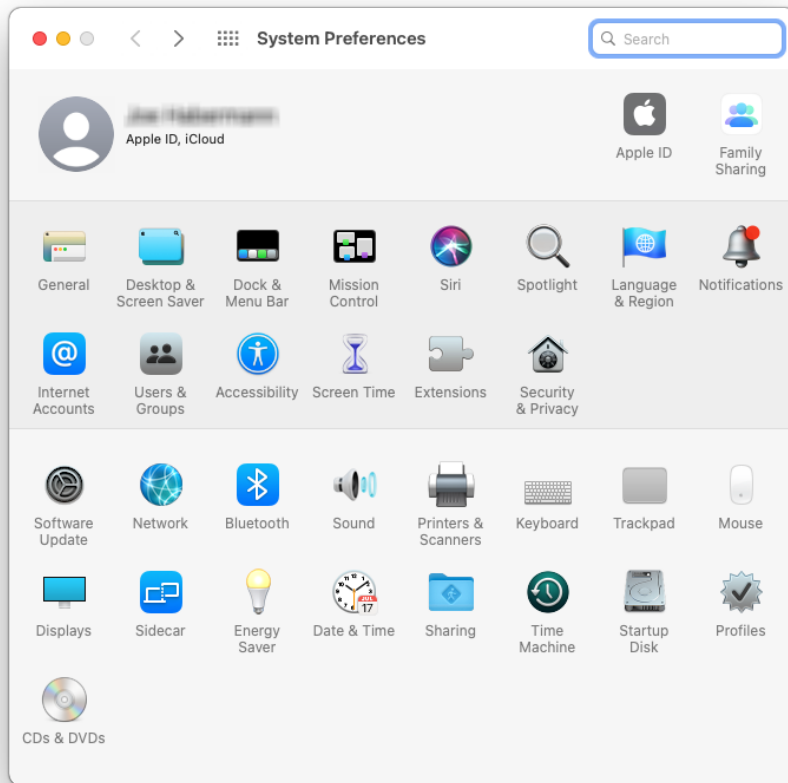
- b. Double-click the **.mobileconfig** profile:



- c. The **Profiles** dialog appears, prompting you to review the profile in **System Preferences** if you want to install it:




d. Click **System Preferences** to display its dialog:




e. In the **System Preferences** dialog, click **Profiles** to display its dialog:

Appendix B: Configuration Tasks
Connect an Xsan Client to a StorNext MDC

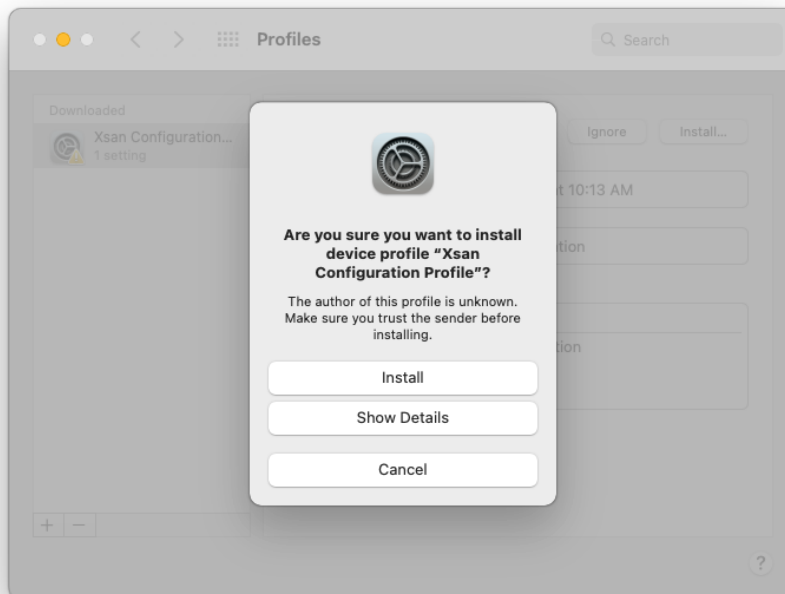


- f. On the left pane of the **Profiles** dialog, click the device profile you want to install and then click **Install...** on the right pane:

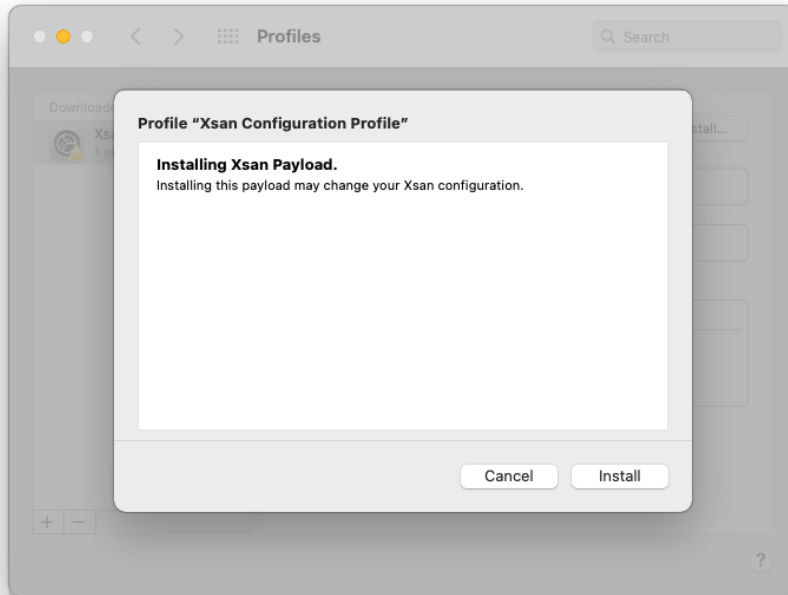
Appendix B: Configuration Tasks
Connect an Xsan Client to a StorNext MDC



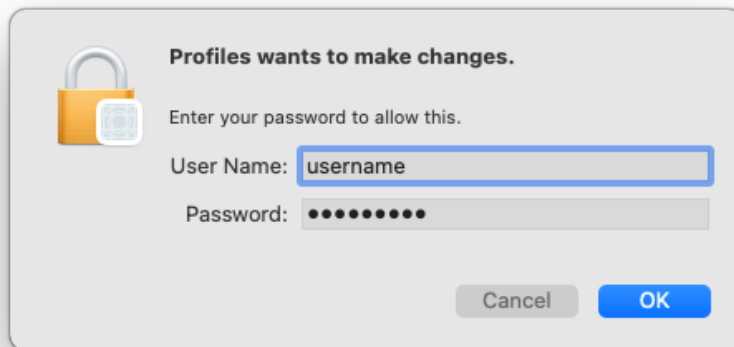
g. You are prompted to verify you want to install the device profile. Click **Install**:



- h. You are prompted to verify you want to install the Xsan payload. Click **Install**:

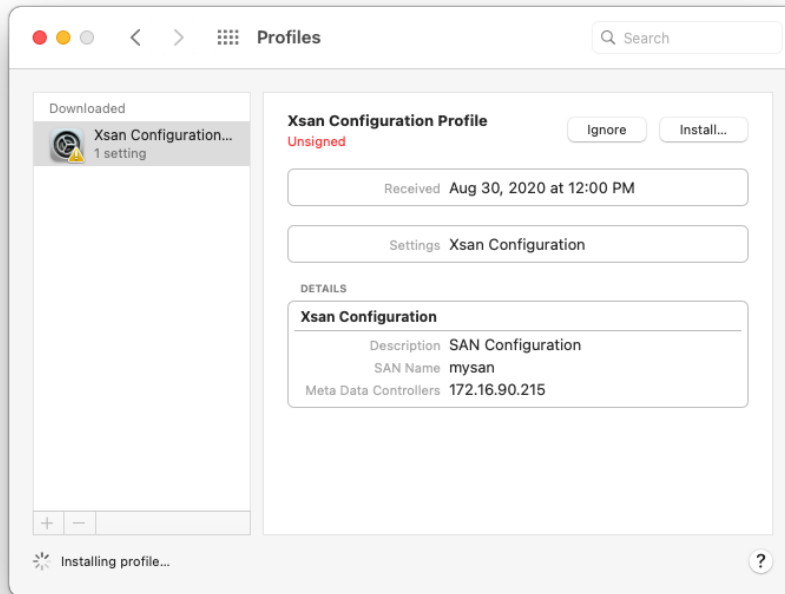


- i. Enter your credentials in the **User Name** and **Password** fields:

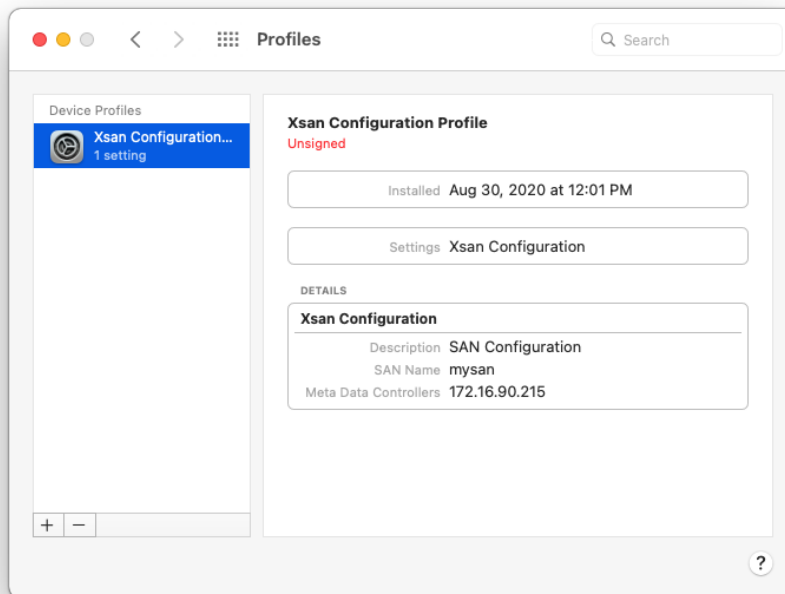


- j. The **Profiles** dialog displays the device profile installation progress in the lower left of the dialog box:

Appendix B: Configuration Tasks
Connect an Xsan Client to a StorNext MDC



k. After the installation is complete, the **Profiles** dialog displays the installed device profile:



6. Restart Xsan:

i Note: On a new Xsan 5.0 (or later) driver installation, an Xsan client may not have been loaded and an error will be displayed when you execute `launchctl unload`. Quantum recommends you verify a driver is loaded and skip the unload step if the driver is not loaded.

When the Xsan client is loaded using the command `launchctl load`, the client will need to be rebooted before the Xsan client is loaded. Quantum recommends you reboot the client.

If an upgrade to Xsan 5.0 (or later) has been performed, Xsan is disabled after the upgrade. Re-enable Xsan using the commands below.

```
xsan50client$ sudo launchctl unload
/System/Library/LaunchDaemons/com.apple.xsan.plist
xsan50client$ sudo launchctl load -w
/System/Library/LaunchDaemons/com.apple.xsan.plist
```

7. Execute the command `cvadmin` to verify the file systems are present.
8. Execute the command `cvadmin -e disks` to verify the all the LUNs used for file systems are displayed.
9. In the Terminal type `xsanctl mount <fsname>`, where `<fsname>` is the file system name displayed in `cvadmin`. See [Special Considerations for Xsan Mount Restrictions on page 219](#).
10. Verify the StorNext file system is mounted on the macOS X computer.

Mount an Xsan Server from a StorNext Client

Perform the following procedure to mount the StorNext file system on a macOS computer.

1. Make sure the Xsan MDC has a valid license.dat file.
2. Create file systems/Volumes on Xsan MDC
3. Install StorNext client build on your desired systems.
4. Verify the StorNext clients can see the same disks/LUNs as the Xsan MDC (use "cvlabel -L" or "cvadmin -e disks" to verify this)
5. Copy the `.auth_secret` and `fsnameservers` files from the Apple Xsan MDC to each StorNext client (may need to change permissions on both files).
 - a. **For Linux Clients:**
 - Place both files under the `/usr/cvfs/config` directory.
 - b. **For Windows Clients:**
 - Place both files under the `C:\Program Files\StorNext\config` directory.
6. Restart the StorNext services on the client

Verify Access to the Xsan File Systems

Depending on your operating system, use one of the following procedures to verify access to the Xsan file systems.

- **For Linux Clients**

1. From the command line, type:

```
/usr/cvfs/bin/cvadmin -e select
```

2. Output similar to the following is shown:

```
Select FSM "none"  
List FSS  
File System Services (* indicates service is in control of FS):  
1>*Xsan1[0] located on pan:33355 (pid 7484)  
2>*Xsan2[0] located on pan:33355 (pid 7485)
```

- **For Windows Clients**

1. Click **Start > Programs > StorNext > Client Configuration**.
2. Verify file system are displayed here and they can be mounted.

Mount the File System on a Linux Client

1. Create mount points to mount the file systems (for example, **mkdir -p /stornext/Xsan1**). Add mount entries in the **/etc/fstab** file similar to:

```
Xsan1 /stornext/Xsan1 cvfs rw,noauto 0 0
```

2. Mount the Xsan File System on Linux client by issuing below mount command:

```
mount -t cvfs Xsan1 /stornext/Xsan1
```

For Systems Running StorNext 6.x (or earlier) and macOS 10.x (or earlier)

Mount the StorNext File System

Perform the following procedure to connect an Xsan 5.0 (or later) client to a StorNext MDC.

1. Generate a configuration profile:

```
mdc# cd /usr/cvfs/config
mdc# /usr/cvfs/lib/make_profile mysan
```

2. Copy the resultant profile to the Xsan client:

```
mdc# scp mysan.configprofile root@xsan50client:/tmp
```

3. Remove any existing Xsan profiles from the client:

```
xsan50client$ sudo profiles -L | grep com.apple.xsan | awk '{print $NF}' > /tmp/oldprofiles
xsan50client$ for i in `cat /tmp/oldprofiles`; do sudo profiles -R -p $i; done
```

4. Install the new profile:

```
xsan50client$ cd /tmp
xsan50client$ sudo profiles -I -F mysan.configprofile
```

5. Restart Xsan:

i Note: On a new Xsan 5.0 (or later) driver installation, an Xsan client may not have been loaded and an error will be displayed when you execute **launchctl unload**. Quantum recommends you verify a driver is loaded and skip the unload step if the driver is not loaded.

When the Xsan client is loaded using the command **launchctl load**, the client will need to be rebooted before the Xsan client is loaded. Quantum recommends you reboot the client.

If an upgrade to Xsan 5.0 (or later) has been performed, Xsan is disabled after the upgrade. Re-enable Xsan using the commands below.

```
xsan50client$ sudo launchctl unload  
/System/Library/LaunchDaemons/com.apple.xsan.plist  
xsan50client$ sudo launchctl load -w  
/System/Library/LaunchDaemons/com.apple.xsan.plist
```

6. Execute the command **cvadmin** to verify the file systems are present.
7. Execute the command **cvadmin -e disks** to verify the all the LUNs used for file systems are displayed.
8. In the Terminal type **xsanctl mount <fsname>**, where **<fsname>** is the file system name displayed in **cvadmin**. See [Special Considerations for Xsan Mount Restrictions on the next page](#).
9. Verify the StorNext file system is mounted on the macOS X computer.

Mount an Xsan Server from a StorNext Client

Perform the following procedure to mount the StorNext file system on a macOS computer.

1. Make sure the Xsan MDC has a valid **license.dat** file.
2. Create file systems/Volumes on Xsan MDC
3. Install StorNext client build on your desired systems.
4. Verify the StorNext clients can see the same disks/LUNs as the Xsan MDC (use "cvlabel -L" or "cvadmin -e disks" to verify this)
5. Copy the **.auth_secret** and **fsnameservers** files from the Apple Xsan MDC to each StorNext client (may need to change permissions on both files).
 - a. **For Linux Clients:**
 - Place both files under the **/usr/cvfs/config** directory.
 - b. **For Windows Clients:**
 - Place both files under the **C:\Program Files\StorNext\config** directory.
6. Restart the StorNext services on the client

Verify Access to the Xsan File Systems

Depending on your operating system, use one of the following procedures to verify access to the Xsan file systems.

- **For Linux Clients**

1. From the command line, type:

```
/usr/cvfs/bin/cvadmin -e select
```

2. Output similar to the following is shown:

```
Select FSM "none"
List FSS
File System Services (* indicates service is in control of FS):
1>*Xsan1[0] located on pan:33355 (pid 7484)
2>*Xsan2[0] located on pan:33355 (pid 7485)
```

- **For Windows Clients**

1. Click **Start > Programs > StorNext > Client Configuration**.
2. Verify file system are displayed here and they can be mounted.

Mount the File System on a Linux Client

1. Create mount points to mount the file systems (for example, `mkdir -p /stornext/Xsan1`). Add mount entries in the `/etc/fstab` file similar to:

```
Xsan1 /stornext/Xsan1 cvfs rw,noauto 0 0
```

2. Mount the Xsan File System on Linux client by issuing below mount command:

```
mount -t cvfs Xsan1 /stornext/Xsan1
```

Special Considerations for Xsan Mount Restrictions

- The client attempts to mount all of the StorNext file systems automatically, including the **High Availability (HA)** shared file system.
- By default, the **Distributed LAN Client (DLC)** is enabled for all file systems.
- If the `automount.plist` file exists, the client ignores the `/Library/Preferences/Xsan/automount.plist` file. On a fresh installation, the `automount.plist` file does **not** exist.
- To **only** mount a given file system, execute the command similar to the example below. The command enforces mounts on `snfs1` and `snfs2` **only**:

```
sudo defaults write /Library/Preferences/com.apple.xsan onlyMount -array snfs1 snfs2
```

- To deny automatic mounting of a given file system, execute the command similar to the example below. The command enforces not to automatically mount the HA shared file system **shared-AV1234ABC00567**:

```
sudo defaults write /Library/Preferences/com.apple.xsan denyMount -array shared-AV1234ABC00567
```

- To disable **DLC** for a list of file systems, execute the command similar to the example below. The command disables **DLC** on snfs1 and snfs2 file systems **only**:

```
sudo defaults write /Library/Preferences/com.apple.xsan denyDLC -array snfs1 snfs2
```

- At system boot, the Xsan daemon might attempt to use DLC to mount a volume if the mount event occurs before the Fibre Channel interface is visible to the system, if the Xsan client is a SAN only client disable DLC for all file systems:

```
sudo defaults write /Library/Preferences/com.apple.xsan useDLC -bool NO
```

- Confirm what Xsan Preferences are set on the system, execute the command below and review the output that is similar to the below:

```
defaults read /Library/Preferences/com.apple.xsan.plist
{
    denyMount = (
        "shared-AV1234ABC00567"
    );
    onlyMount = (
        snfs1,
        snfs2
    );
    denyDLC = (
        snfs1,
        snfs2
    );
    useDLC = 0;
}
```

i Note: For additional information, see the Apple support article available online at <https://support.apple.com/en-us/HT205333>.

Conditions for Mounting an Xsan Volume as a StorNext Distributed LAN Client (DLC)

Beginning with Apple's macOS El Capitan (10.11), support is provided to mount Xsan volumes as StorNext DLC clients. This section provides additional information which should clarify the conditions necessary for this mount to succeed.

When the StorNext services are loaded, the environment is scanned to determine the default mount type - storage area network (SAN) with access to data disks, or DLC with proxy access to data disks through a StorNext gateway node.

If a fibre channel card is detected or internet Small Computer Systems Interface (iSCSI) logical unit numbers (LUNs) are present, the default mode is SAN. If neither of these conditions are true, the default mode is DLC.

In the case that a SAN mount is attempted and not all data LUNs for the file system are present, the mount fails and no error message is logged. If a DLC mount is desired, there are several ways to accomplish this.

First, if the presence of iSCSI LUNs is causing the SAN mount attempt, and these LUNs are not needed, the iSCSI targets can be disconnected. If this is done, the StorNext service must be unloaded and loaded to detect the change.

Example of unload and load:

```
sh-3.2# launchctl unload /System/Library/LaunchDaemons/com.apple.xsan.plist
sh-3.2# launchctl load -w /System/Library/LaunchDaemons/com.apple.xsan.plist
```

If iSCSI LUNs or a fibre channel card are present, the DLC mount is accomplished by indicating this to the driver. There are two ways to achieve this:

- [Set the Xsan Preferences Payload "preferDLC" for the Volume that is to be DLC Mounted below](#)
- [Set the diskproxy=client Mount Options in the /etc/fstab File on the next page](#)

Set the Xsan Preferences Payload "preferDLC" for the Volume that is to be DLC Mounted

Example of setting preferDLC payload and verifying:


```
sh-3.2# defaults read /Library/Preferences/com.apple.xsan
{
    enableSpotlightServer = 1;
}
sh-3.2# defaults write /Library/Preferences/com.apple.xsan preferDLC '(snfs1)'
sh-3.2# defaults read /Library/Preferences/com.apple.xsan
{
    enableSpotlightServer = 1;
    preferDLC = (
        snfs1
    );
}
```

For more information on the preferDLC payload and Xsan configuration profiles, see <https://support.apple.com/en-us/HT205333>.

Set the diskproxy=client Mount Options in the /etc/fstab File

An alternative way to indicate to the driver that a DLC mount is desired is to set the `diskproxy=client` mount options in the `/etc/fstab` file.

Example of the `diskproxy=client` mount options in `/etc/fstab` file:

```
sh-3.2# cat /etc/fstab
LABEL=snfs1    none    acfs    rw,diskproxy=client
sh-3.2#
```

In either of the above scenarios, a driver unload and load must be performed to enact the change. Once the Volume is mounted, its status as a DLC client can be verified by using the sub-command `who` of the `cvadmin` CLI. First select the file system, then enter `who`. Look for the `C` after the `CLI` in the client entry.

Example of the `cvadmin who` sub-command:

```
Xsanadmin (snfs1) > select snfs1
Xsanadmin (snfs1) > who
Who (File System "snfs1")
```

| # | acfs | I.D. | Type | Location | Up Time | License Expires |
|----------|------------|------|-------|-------------------------------|------------|-----------------|
| 0> | | | FSM | smw-rhel64.mdh.quantum.com | 4d 23h 33m | N/A |
| 4> | 50569B3B7F | | CLI S | smw-rhel64.mdh.quantum.com | 4d 23h 33m | *10-27-2016 |
| 23:59:59 | | | | | | |
| 9> | | | ADM | smw-elcapitan.mdh.quantum.com | 0d 0h 0m | N/A |
| 17> | | | CLI C | smw-elcapitan.mdh.quantum.com | 0d 0h 6m | N/A |

```

1 active connection out of 65535 licensed connections
0 active connections out of 65535 licensed proxy connections
1 active gateway connection
1 active client-licensed connection

Xsanadmin (snfs1) >
```

View a Fibre Channel Device

Use the procedure below to add a fibre channel (FC) device. Before adding a FC device, first configure the Host Bus Adapter (HBA) card so you can view the device. Use the **fs_scsi -p** command to make sure you can view your devices over FC. FC devices include tape libraries, individual drives, or RAID disk.

- Type the **/usr/cvfs/bin/fs_scsi -p** command to make sure that the necessary devices can be viewed over FC.
 - To view tape libraries, access the StorNext GUI to view your tape libraries.
 - To view individual drives or RAID LUNs, access either the StorNext or SNFS GUI.

To add a disk, see [File System Tasks](#) in the *StorNext User's Guide*. To add a tape drive to a configured library, see [Setting Up Libraries](#) in the *StorNext User's Guide*.

StorNext Log Rolling Administration

Use the following procedures to configure StorNext log rolling options.

Change the Log Rolling Times

Use this procedure to change the frequency of rolling the StorNext logs. This process requires that you edit the **tdlm crontab** to set log times. Use this procedure to edit **tdlm crontab**.

1. Log on as **root**.
2. Edit the **tdlm crontab** and update the **sn_log_update script**. Below is an example **crontab**, where **0 1,7,13,19 * * *** designates the times when logs run.

```
0 1,7,13,19 * * * /usr/adic/guib/bin/cmdwrap /usr/adic/util/sn_log_update /usr/adic
```

Log Rolling Options

You can change these options to optimize log rolling.

- **-s**: This option sets the directory where logs are saved (copied) to as they are rolled. This directory is typically a managed directory, where **<dir>** is the directory where you want the logs to reside. For example:

```
sn_log_update [-s <dir>]
```

- **\$DEF_SIZE = 2000000**: This is the default size at which logs are rolled. Edit this entry in the **/usr/adic/util/sn_log_update.cfg** file if you want the log sizes to be larger or smaller.
- **\$DEF_LOGS = 28**: This is the default number of logs that are saved before they are deleted. Edit this entry in the **/usr/adic/util/sn_log_update.cfg** file if you want to save less than 28 logs or are saving the logs to a managed directory.

Avoid Poor System Performance

Most systems run administrator scripts to rebuild namespace databases and/or to remove certain files. Unfortunately, running these scripts can degrade StorNext performance. It is recommended that you review administrator scripts to determine if they may adversely affect your StorNext installation and its performance.

For example, many UNIX systems have a **find** command in the **root crontab** that traverses the entire directory hierarchy while searching for old temporary and core files. Typically, the **find** commands are configured only to run on the local file system. However by default, these commands do not skip StorNext file systems because they are considered to be local, even though the file system manager may be running remotely. StorNext performance can be dramatically degraded when these commands visit a StorNext file system if the system contains a large number of files.

If StorNext performance is degraded to an unacceptable level due to the operation of administrator scripts, you can modify them to skip StorNext file systems. A sample command (original and modified is shown below):

On IRIX, by default, the **root crontab** contains the following find command:

```
find / -local -type f '(' -name core -o -name dead.letter ')' -atime +7 -mtime +7 -exec rm -f '{} ' ;'
```

To prevent the find command from searching in StorNext file systems, modify it as shown:

```
'(' -type d -fstype cvfs -prune ')' -o
```

The modified find command looks like this:

```
find / -local '(' -type d -fstype cvfs ')' -o -type f '(' -name core -o -name dead.letter ')' -atime +7 -mtime +7 -exec rm -f '{} ' ;'
```

Avoid Not Saving cvmkfs Coredump Files

Both RHEL and CentOS systems enable a service called **abrt** by default. This service intercepts StorNext core files and tries to handle them. To avoid this situation, implement either one of these options:

- Add `/usr/adic/*/bin/*` and `/usr/cvfs/bin/*` to the **abrt** blacklist path.
- Disable the **abrt** service using the **chkconfig** command:

```
# chkconfig abrt off
```

Import a Secure Sockets Layer (SSL) Certificate

This section provides procedures to import a SSL certificate.

i Note: Quantum recommends you review your specific SSL Certificate Authority vendor requirements to upload/submit your Certificate Signed Request to purchase a certificate.

Important

The procedures must be executed on each node of your system where StorNext runs an Apache Tomcat server.

Sign the Tomcat Operating Certificate with a Custom Certification Authority (CA)

Keytool Usage:

For all uses of the **keytool** utility, you will be prompted to enter a password, which is stored in the following file as the value of **keystorePass**. By default, the **keytool** utility password is **changeit**.

```
grep Pass /usr/adic/tomcat/conf/server.xml  
keystorePass="changeit"
```

The **keytool** utility is located in the following directory:

```
/usr/adic/java/jre/bin
```

Step 1: Generate a Certificate Signed Request

You must generate a .csr file to submit for self-signing or by a trusted Certificate Authority.

Execute the command below to generate the file, **tomcat.csr**.

```
/usr/adic/java/jre/bin/keytool -certreq -keyalg RSA -keystore  
/usr/adic/gui/config/.keystore -alias tomcat -file tomcat.csr  
Enter keystore password:
```

You can expect the following JKS format warning:

```
Warning: The JKS keystore uses a proprietary format. It is recommended to  
migrate to PKCS12 which is an industry standard format using "keytool -  
importkeystore -srckeystore /usr/adic/gui/config/.keystore -destkeystore  
/usr/adic/gui/config/.keystore -deststoretype pkcs12".
```

Execute the following command to view the .csr file you generated above:

```
openssl req -in tomcat.csr -noout -text
```

```
Certificate Request:
  Data:
    Version: 0 (0x0)
    Subject: C=US, ST=California, L=San Jose, O=StorNext Software,
    OU=Quantum Corp., CN=node-1.node-1
    Subject Public Key Info:
      Public Key Algorithm: rsaEncryption
      Public-Key: (2048 bit)
      Modulus:
```

Step 2: Sign the Certificate

If you have a public domain open on port 80 you can submit the .csr file to be signed by a Trusted CA. There are both paid and open source options. If you sign a .csr in this manner, then it does not work for a domain that cannot be reached publicly on port 80 and is not a registered domain name. If you have a Trusted CA, sign the .csr and then proceed to [Step 3: Combine the Two Certificates to Create a Chain on the next page](#).

When you self-signing a certificate, you create your own .CRT file.

You must provide the **tomcat.cnf** file with a **subjectAltName** that reflects the IP address to Hostname resolution of the system that users connect to the Stornext GUI with:

```
echo "subjectAltName=DNS:spsxcellis2.sps.lab,DNS:sps.lab,IP:10.20.232.6" >
/etc/pki/tls/tomcat.cnf

[root@sps-xcellis2 jon]# nslookup sps-xcellis2
Server: 10.20.232.95
Address: 10.20.232.95#53
Name: sps-xcellis2.sps.lab
Address: 10.20.232.6
```

Create the .CRT file:

```
[root@sps-xcellis2 cert]# openssl x509 -req -in tomcat.csr -out tomcat.crt -CA
/usr/cvfs/config/certs/server.crt -CAkey /usr/cvfs/config/certs/server.key -
sha256 -days 365 -CAcreateserial -extfile /etc/pki/tls/tomcat.cnf
Signature ok
```

```
subject=/C=US/ST=California/L=San Jose/O=StorNext Software/OU=Quantum  
Corp./CN=node-1.node-1  
Getting CA Private Key  
[root@sps-xcellis2 cert]# ls  
tomcat.crt tomcat.csr
```

Step 3: Combine the Two Certificates to Create a Chain

The following procedure creates the file, **myTomcat.crt**. Execute the following command:

```
[root@sps-xcellis2 cert]# cat tomcat.crt /usr/cvfs/config/certs/server.crt >  
myTomcat.crt
```

Step 4: Import Your Intermediate CA Certificate and Your New Tomcat Certificate into Tomcat's keystore

 **Caution:** You must first import the CA certificate.

Execute the following commands:

```
[root@sps-xcellis2 cert]# /usr/adic/java/jre/bin/keytool -keystore  
/usr/adic/gui/config/.keystore -import -alias xmed -file ./myCA.crt  
Enter keystore password:  
Certificate reply was installed in keystore
```

```
[root@sps-xcellis2 cert]# /usr/adic/java/jre/bin/keytool -keystore  
/usr/adic/gui/config/.keystore -import -alias tomcat -file ./myTomcat.crt  
Enter keystore password:  
Certificate reply was installed in keystore
```

Step 5: Restart Tomcat

Execute the following command:

```
/etc/init.d/stornext_web restart
```

-
- i Note:** The procedure to add the root certificate into the browser's list of recognized Trusted Certificate Authorities is specific to the browser you are using; for additional information, refer to the browser documentation.

ActiveScale™ P100/X100 Installation, Configuration and Tuning

This section contains important information about performance configuration options for your ActiveScale P100/X100 integrated object storage solution system.

Prior to performing your installation and configuration of your ActiveScale P100/X100 system, you should review and consider the following performance configuration options.

-
- i Note:** Contact [Quantum Service and Support](#), or your ActiveScale Pre-Sales Service and Support representative to assess network, storage policies, and encryption requirements.

Configure the P100/X100 Network Settings

You must select one of the following 10G network configuration options on your P100/X100:

- **Bonded Round Robin:** This network setting provides reduced performance and simplified configuration setup.
- **Un-bonded Using Different Subnets:** This network setting provides acceptable performance that is Quantum tested and recommended.
- **Bonded Interface with LACP:** This network setting provides improved performance, but has not been specifically tested by Quantum. However, if you choose this option, you should obtain the settings for configuring LACP from the switch vendor, or its documentation.

-
- i Note:** Configuring the network settings on Xcellis is performed using the StorNext GUI. This must be completed during scheduled downtime since both servers/nodes of the HA pair are rebooted. See [System](#) for additional information.

-
- ⚠ Caution:** The **Network Settings** of your P100/X100 system cannot be changed after the installation and configuration process is complete.

Configure the P100/X100 Storage Policy

With P100/X100 single site installation, there are 2 types of storage policies available:

- **Balanced (13/4)**
- **Capacity Optimized (15/4)**

i Note: Quantum recommends you select the **Capacity Optimized (15/4)** policy for its highly efficient use of the P100/X100 capacity.

! Caution: The **Storage Policy** configuration cannot be changed after the P100/X100 installation and configuration process is complete.

Configure the P100/X100 Encryption Level

The P100/X100 supports server side encryption, which you can enable after the installation and configuration process is complete.

! Caution: Once you enable server side encryption on your P100/X100, you cannot disable it.

With the P100/X100, you have the following 3 options for its encryption level:

- **No Encryption:** The P100/X100 is not configured for encryption. StorNext server side encryption requests from StorNext Storage Manager are ignored silently.
- **System Level Encryption:** All the objects are encrypted, regardless if StorNext uses server side encryption request or not.
- **Object Level Encryption:** Only objects which contain requests for encryption are encrypted. This is set when creating or modifying the StorNext policy.

Configure the StorNext Storage Destination

For the P100/X100, Quantum recommends the default setting of 48 streams per controller, since it provides optimal performance.



Appendix C: StorNext Configuration Files

This appendix contains the following topics:

StorNext Files

StorNext creates a series of files that are used and modified to configure a file system. This section includes an expanded configuration file, a cluster-wide central control file, and a listing of the most commonly used StorNext files with descriptions and locations for each.

- [Example FSM Configuration File below](#)
- [Cluster-Wide Central Control on page 242](#)

Example FSM Configuration File

On Linux, the StorNext configuration file uses an XML format (.cfgx). On Windows, the configuration file uses a text format (.cfg). However, the values contained in both files are similar.

You can locate an example StorNext configuration file in the following directory:

- Linux: **/usr/cvfs/examples/example.cfgx**
- Windows: **C:\Program Files\Stornext\config\example.cfg**

If you installed StorNext in a location other than the default installation directory, the example configuration file is located in **C:\<install_directory>\config\example.cfg**

For additional information about the snfs.cfgx and snfs.cfg files, refer to the “snfs_config” entry in the *Man Pages Reference Guide*.

Linux Example Configuration File

Below are the contents of the StorNext example configuration file for Linux (**example.cfgx**):

```
<?xml version="1.0"?>
<snfs:configDoc xmlns:snfs="https://www.quantum.com/snfs" version="1.0">
  <snfs:config configVersion="0" name="example" fsBlockSize="4096"
journalSize="67108864">
    <snfs:globals>
      <snfs:affinityPreference>false</snfs:affinityPreference>
      <snfs:allocationStrategy>round</snfs:allocationStrategy>
      <snfs:haFsType>HaUnmonitored</snfs:haFsType>
      <snfs:bufferCacheSize>268435456</snfs:bufferCacheSize>
      <snfs:cvRootDir>/</snfs:cvRootDir>
      <snfs:storageManager>false</snfs:storageManager>
      <snfs:debug>00000000</snfs:debug>
      <snfs:dirWarp>true</snfs:dirWarp>
      <snfs:extentCountThreshold>49152</snfs:extentCountThreshold>
      <snfs:enableSpotlight>false</snfs:enableSpotlight>
      <snfs:spotlightUseProxy>false</snfs:spotlightUseProxy>
      <snfs:enforceAcls>false</snfs:enforceAcls>
      <snfs:fileLocks>false</snfs:fileLocks>
      <snfs:fileLockResyncTimeout>20</snfs:fileLockResyncTimeout>
      <snfs:forcePerfectFit>false</snfs:forcePerfectFit>
      <snfs:fsCapacityThreshold>0</snfs:fsCapacityThreshold>
      <snfs:globalSuperUser>true</snfs:globalSuperUser>
      <snfs:ignorePermissions>false</snfs:ignorePermissions>
      <snfs:inodeCacheSize>131072</snfs:inodeCacheSize>
      <snfs:inodeExpandMin>0</snfs:inodeExpandMin>
      <snfs:inodeExpandInc>0</snfs:inodeExpandInc>
      <snfs:inodeExpandMax>0</snfs:inodeExpandMax>
      <snfs:inodeDeleteMax>0</snfs:inodeDeleteMax>
      <snfs:inodeStripeWidth>4294967296</snfs:inodeStripeWidth>
      <snfs:ioTokens>true</snfs:ioTokens>
      <snfs:maintenanceMode>false</snfs:maintenanceMode>
      <snfs:maxLogs>4</snfs:maxLogs>
      <snfs:namedStreams>false</snfs:namedStreams>
      <snfs:remoteNotification>false</snfs:remoteNotification>
      <snfs:renameTracking>false</snfs:renameTracking>
      <snfs:reservedSpace>true</snfs:reservedSpace>
      <snfs:fsmRealTime>false</snfs:fsmRealTime>
```

Appendix C: StorNext Configuration Files
Example FSM Configuration File

```
<snfs: fsmMemLocked>false</snfs: fsmMemLocked>
<snfs: opHangLimitSecs>180</snfs: opHangLimitSecs>
<snfs: perfectFitSize>131072</snfs: perfectFitSize>
<snfs: quotas>false</snfs: quotas>
<snfs: quotaHistoryDays>7</snfs: quotaHistoryDays>
<snfs: metadataArchive>false</snfs: metadataArchive>
<snfs: metadataArchiveDir></snfs: metadataArchiveDir>
<snfs: metadataArchiveDays>0</snfs: metadataArchiveDays>
<snfs: metadataArchiveSearch>true</snfs: metadataArchiveSearch>
<snfs: metadataArchiveCache>2147483648</snfs: metadataArchiveCache>
<snfs: drCopy>false</snfs: drCopy>
<snfs: audit>false</snfs: audit>
<snfs: restAccess>privileged</snfs: restAccess>
<snfs: securityModel>legacy</snfs: securityModel>
<snfs: stripeAlignSize>-1</snfs: stripeAlignSize>
<snfs: trimOnClose>0</snfs: trimOnClose>
<snfs: useL2BufferCache>true</snfs: useL2BufferCache>

<snfs: unixDirectoryCreationModeOnWindows>755</snfs: unixDirectoryCreationM
odeOnWindows>

<snfs: unixIdFabricationOnWindows>false</snfs: unixIdFabricationOnWindows>

<snfs: unixFileCreationModeOnWindows>644</snfs: unixFileCreationModeOnWindo
ws>
  <snfs: unixNobodyUidOnWindows>60001</snfs: unixNobodyUidOnWindows>
  <snfs: unixNobodyGidOnWindows>60001</snfs: unixNobodyGidOnWindows>
  <snfs: windowsSecurity>true</snfs: windowsSecurity>
  <snfs: windowsIdMapping>ldap</snfs: windowsIdMapping>
  <snfs: unixIdMapping>none</snfs: unixIdMapping>
  <snfs: globalShareMode>false</snfs: globalShareMode>
  <snfs: useActiveDirectorySFU>true</snfs: useActiveDirectorySFU>
  <snfs: eventFiles>true</snfs: eventFiles>
  <snfs: eventFileDir></snfs: eventFileDir>

<snfs: allocSessionReservationSize>1073741824</snfs: allocSessionReservatio
nSize>
  </snfs: globals>
  <snfs: diskTypes>
    <snfs: diskType typeName="MetaDrive" sectors="99999999"
sectorSize="512"/>
    <snfs: diskType typeName="JournalDrive" sectors="99999999"
sectorSize="512"/>
    <snfs: diskType typeName="VideoDrive" sectors="99999999"
```

Appendix C: StorNext Configuration Files
Example FSM Configuration File

```
sectorSize="512"/>
  <snfs:diskType typeName="AudioDrive" sectors="99999999"
sectorSize="512"/>
  <snfs:diskType typeName="DataDrive" sectors="99999999"
sectorSize="512"/>
</snfs:diskTypes>
<snfs:autoAffinities/>
<snfs:stripeGroups>
  <snfs:stripeGroup index="0" status="up" name="MetaFiles"
stripeBreadth="262144" read="true" write="true" alloc="true"
metadata="true" journal="false" userdata="false" realTimeIOs="0"
realTimeIOsReserve="0" realTimeMB="0" realTimeMBReserve="0"
realTimeTokenTimeout="0" multipathMethod="rotate">
  <snfs:disk index="0" diskLabel="CvfsDisk0" diskType="MetaDrive"
ordinal="0"/>
  </snfs:stripeGroup>
  <snfs:stripeGroup index="1" status="up" name="JournFiles"
stripeBreadth="262144" read="true" write="true" alloc="true"
metadata="false" journal="true" userdata="false" realTimeIOs="0"
realTimeIOsReserve="0" realTimeMB="0" realTimeMBReserve="0"
realTimeTokenTimeout="0" multipathMethod="rotate">
  <snfs:disk index="0" diskLabel="CvfsDisk1" diskType="JournalDrive"
ordinal="1"/>
  </snfs:stripeGroup>
  <snfs:stripeGroup index="2" status="up" name="VideoFiles"
stripeBreadth="4194304" read="true" write="true" alloc="true"
metadata="false" journal="false" userdata="true" realTimeIOs="0"
realTimeIOsReserve="0" realTimeMB="0" realTimeMBReserve="0"
realTimeTokenTimeout="0" multipathMethod="rotate">
  <snfs:affinities exclusive="true">
    <snfs:affinity>Video</snfs:affinity>
  </snfs:affinities>
  <snfs:disk index="0" diskLabel="CvfsDisk2" diskType="VideoDrive"
ordinal="2"/>
  <snfs:disk index="1" diskLabel="CvfsDisk3" diskType="VideoDrive"
ordinal="3"/>
  <snfs:disk index="2" diskLabel="CvfsDisk4" diskType="VideoDrive"
ordinal="4"/>
  <snfs:disk index="3" diskLabel="CvfsDisk5" diskType="VideoDrive"
ordinal="5"/>
  <snfs:disk index="4" diskLabel="CvfsDisk6" diskType="VideoDrive"
ordinal="6"/>
  <snfs:disk index="5" diskLabel="CvfsDisk7" diskType="VideoDrive"
ordinal="7"/>
```

```
        <snfs:disk index="6" diskLabel="CvfsDisk8" diskType="VideoDrive"
ordinal="8"/>
        <snfs:disk index="7" diskLabel="CvfsDisk9" diskType="VideoDrive"
ordinal="9"/>
    </snfs:stripeGroup>
    <snfs:stripeGroup index="3" status="up" name="AudioFiles"
stripeBreadth="1048576" read="true" write="true" alloc="true"
metadata="false" journal="false" userdata="true" realTimeIOs="0"
realTimeIOsReserve="0" realTimeMB="0" realTimeMBReserve="0"
realTimeTokenTimeout="0" multipathMethod="rotate">
        <snfs:affinities exclusive="true">
            <snfs:affinity>Audio</snfs:affinity>
        </snfs:affinities>
        <snfs:disk index="0" diskLabel="CvfsDisk10" diskType="AudioDrive"
ordinal="10"/>
        <snfs:disk index="1" diskLabel="CvfsDisk11" diskType="AudioDrive"
ordinal="11"/>
        <snfs:disk index="2" diskLabel="CvfsDisk12" diskType="AudioDrive"
ordinal="12"/>
        <snfs:disk index="3" diskLabel="CvfsDisk13" diskType="AudioDrive"
ordinal="13"/>
    </snfs:stripeGroup>
    <snfs:stripeGroup index="4" status="up" name="RegularFiles"
stripeBreadth="262144" read="true" write="true" alloc="true"
metadata="false" journal="false" userdata="true" realTimeIOs="0"
realTimeIOsReserve="0" realTimeMB="0" realTimeMBReserve="0"
realTimeTokenTimeout="0" multipathMethod="rotate">
        <snfs:disk index="0" diskLabel="CvfsDisk14" diskType="DataDrive"
ordinal="14"/>
        <snfs:disk index="1" diskLabel="CvfsDisk15" diskType="DataDrive"
ordinal="15"/>
        <snfs:disk index="2" diskLabel="CvfsDisk16" diskType="DataDrive"
ordinal="16"/>
        <snfs:disk index="3" diskLabel="CvfsDisk17" diskType="DataDrive"
ordinal="17"/>
    </snfs:stripeGroup>
</snfs:stripeGroups>
<snfs:auxData/>
</snfs:config>
</snfs:configDoc>
```

The renameTracking Feature

Some Microsoft applications end up making temporary copies of operational files while they are being updated. For example while a file myDocument.doc is being updated with Word it may be renamed to a random sequence of characters like ABC123DYXab. When Word completes the updates it will save the

new myDocument.doc file and remove ABC123DYXab. When Storage Manager recognizes that file ABC123DYXab has been removed it makes it recoverable but with the current temporary name. This makes tracking old instances of myDocument.doc difficult if not impossible to accomplish. To get around this scenario the **renameTracking** feature is provided. You can enable this feature on in the file system (see the **snfs_config(5)** man page) and makes tracking of old instances of these temporary files possible.

If you enable the **renameTracking** feature, as a file is renamed it is made recoverable at that time. For example, in the scenario described above as Word renames myDocument.doc to ABC123DYXab, Storage Manager makes myDocument.doc recoverable with its original name before the rename proceeds; that renamed instance becomes inactive. Then after Word is done, the new myDocument.doc becomes the new instance of the file and the previous instance becomes available for recovery using **fsrecover(1)** with the correct name.

i Note: The **renameTracking** functionality is NOT a general purpose feature for keeping the Storage Manager database up to date with rename activity. If you enable the feature, then the files that are being renamed are re-stored. This is not, in general, a desirable behavior except in the specific case described.

Windows Example Configuration File

Below are the contents of the StorNext example configuration file for Windows (**example.cfg**):

```
# Globals

affinityPreference      No
allocationStrategy      Round
HaFsType                HaUnmonitored
fileLocks               No
brlResyncTimeout        20
bufferCacheSize         256M
cvRootDir               /
dataMigration           No
debug                   0x0
dirWarp                 Yes
ForcePerfectFit         No
FsBlockSize             4K
GlobalSuperUser         Yes
IgnorePermissions       No
IoTokens                Yes
InodeCacheSize          128K
InodeExpandMin          0
InodeExpandInc          0
InodeExpandMax          0
InodeDeleteMax          0
InodeStripeWidth        4G
JournalSize             64M
MaxLogs                 4
```

Appendix C: StorNext Configuration Files
Example FSM Configuration File

```
PerfectFitSize      128K
RemoteNotification  No
ReservedSpace      Yes
FSMRealtime        No
FSMMemlock         No
Quotas             No
QuotaHistoryDays   7
RestoreJournal     No
RestoreJournalMaxHours 0
RestoreJournalMaxMB 0
SecurityModel      Legacy
MetadataArchive    No
MetadataArchiveDays 0
DrCopy             No
audit             No
UnixIdFabricationOnWindows No
UnixIdMapping      none
UnixNobodyUidOnWindows 60001
UnixNobodyGidOnWindows 60001
SecurityModel      legacy
WindowsIdMapping   ldap
WindowsSecurity    Yes
```

Disk Types

```
[DiskType MetaDrive]
Sectors 99999999
SectorSize 512
[DiskType JournalDrive]
Sectors 99999999
SectorSize 512
[DiskType VideoDrive]
Sectors 99999999
SectorSize 512
[DiskType AudioDrive]
Sectors 99999999
SectorSize 512
[DiskType DataDrive]
Sectors 99999999
SectorSize 512
```

Disks

```
[Disk CvfsDisk0]
```


Appendix C: StorNext Configuration Files
Example FSM Configuration File

```
Type MetaDrive
Status UP
[Disk CvfsDisk1]
Type JournalDrive
Status UP
[Disk CvfsDisk2]
Type VideoDrive
Status UP
[Disk CvfsDisk3]
Type VideoDrive
Status UP
[Disk CvfsDisk4]
Type VideoDrive
Status UP
[Disk CvfsDisk5]
Type VideoDrive
Status UP
[Disk CvfsDisk6]
Type VideoDrive
Status UP
[Disk CvfsDisk7]
Type VideoDrive
Status UP
[Disk CvfsDisk8]
Type VideoDrive
Status UP
[Disk CvfsDisk9]
Type VideoDrive
Status UP
[Disk CvfsDisk10]
Type AudioDrive
Status UP
[Disk CvfsDisk11]
Type AudioDrive
Status UP
[Disk CvfsDisk12]
Type AudioDrive
Status UP
[Disk CvfsDisk13]
Type AudioDrive
Status UP
[Disk CvfsDisk14]
Type DataDrive
Status UP
```

Appendix C: StorNext Configuration Files

Example FSM Configuration File

```
[Disk CvfsDisk15]
Type DataDrive
Status UP
[Disk CvfsDisk16]
Type DataDrive
Status UP
[Disk CvfsDisk17]
Type DataDrive
Status UP

# Stripe Groups

[StripeGroup MetaFiles]
Status Up
StripeBreadth 64
Metadata Yes
Journal No
Exclusive Yes
Read Enabled
Write Enabled
Alloc Enabled
Rtmb 0
Rtios 0
RtmbReserve 0
RtiosReserve 0
RtTokenTimeout 0
MultiPathMethod Rotate
Node CvfsDisk0 0

[StripeGroup JournFiles]
Status Up
StripeBreadth 64
Metadata No
Journal Yes
Exclusive Yes
Read Enabled
Write Enabled
Alloc Enabled
Rtmb 0
Rtios 0
RtmbReserve 0
RtiosReserve 0
RtTokenTimeout 0
MultiPathMethod Rotate
```

Appendix C: StorNext Configuration Files
Example FSM Configuration File

```
Node CvfsDisk1 0

[StripeGroup VideoFiles]
Status Up
StripeBreadth 1024
Metadata No
Journal No
Exclusive No
Read Enabled
Write Enabled
Alloc Enabled
Rtmb 0
Rtios 0
RtmbReserve 0
RtiosReserve 0
RtTokenTimeout 0
MultiPathMethod Rotate
Node CvfsDisk2 0
Node CvfsDisk3 1
Node CvfsDisk4 2
Node CvfsDisk5 3
Node CvfsDisk6 4
Node CvfsDisk7 5
Node CvfsDisk8 6
Node CvfsDisk9 7
Affinity Video

[StripeGroup AudioFiles]
Status Up
StripeBreadth 256
Metadata No
Journal No
Exclusive No
Read Enabled
Write Enabled
Alloc Enabled
Rtmb 0
Rtios 0
RtmbReserve 0
RtiosReserve 0
RtTokenTimeout 0
MultiPathMethod Rotate
Node CvfsDisk10 0
Node CvfsDisk11 1
```

```
Node CvfsDisk12 2
Node CvfsDisk13 3
Affinity Audio

[StripeGroup RegularFiles]
Status Up
StripeBreadth 64
Metadata No
Journal No
Exclusive No
Read Enabled
Write Enabled
Alloc Enabled
Rtmb 0
Rtios 0
RtmbReserve 0
RtiosReserve 0
RtTokenTimeout 0
MultiPathMethod Rotate
Node CvfsDisk14 0
Node CvfsDisk15 1
Node CvfsDisk16 2
Node CvfsDisk17 3
```

The renameTracking Feature

Some Microsoft applications end up making temporary copies of operational files while they are being updated. For example while a file `myDocument.doc` is being updated with Word it may be renamed to a random sequence of characters like `ABC123DYXab`. When Word completes the updates it will save the new `myDocument.doc` file and remove `ABC123DYXab`. When Storage Manager recognizes that file `ABC123DYXab` has been removed it makes it recoverable but with the current temporary name. This makes tracking old instances of `myDocument.doc` difficult if not impossible to accomplish. To get around this scenario the **renameTracking** feature is provided. You can enable this feature on in the file system (see the **snfs_config(5)** man page) and makes tracking of old instances of these temporary files possible.

If you enable the **renameTracking** feature, as a file is renamed it is made recoverable at that time. For example, in the scenario described above as Word renames `myDocument.doc` to `ABC123DYXab`, Storage Manager makes `myDocument.doc` recoverable with its original name before the rename proceeds; that renamed instance becomes inactive. Then after Word is done, the new `myDocument.doc` becomes the new instance of the file and the previous instance becomes available for recovery using **fsrecover(1)** with the correct name.

- Note:** The **renameTracking** functionality is NOT a general purpose feature for keeping the Storage Manager database up to date with rename activity. If you enable the feature, then the files that are being renamed are re-stored. This is not, in general, a desirable behavior except in the specific case described.

Cluster-Wide Central Control

The purpose of this feature is to provide cluster-wide central control.

i Note: The central control file is supported on the Linux platform only.

A central control file called **nss_cctl.xml** provides a way to restrict the behavior of SNFS cluster nodes (fsm, file system client, cvadmin client) from a central place: an NSS server.

This feature currently supports the following controls that allow you to specify:

1. Whether a client is allowed to mount as a proxy client.
2. Whether a client is allowed to mount as read/write or read-only.
3. Whether a user (especially a local administrator on Windows clients,) is allowed to take ownership of a file or directory on a StorNext file system.
4. Whether **cvadmin** running on a certain client is allowed to have super admin privilege to run destructive commands such as starting/stopping the file system, refreshing disks, changing quota settings, and so on.
5. Whether **cvadmin** running on certain client is allowed to connect to other fsms via the **-H** option.
6. Whether binary executable files on the StorNext file system are allowed to be executed.
7. Whether the **setuid** bit of a file is allowed to take effect.

The control file is in xml format and has a hierarchical structure. The top level element, **snfsControl**, contains control elements with the **securityControl** label for certain file systems. If you have different controls for different file systems, each file system should have its own control definition. A special virtual file system, **#SNFS_ALL#**, is used as the default control for file systems not defined in this control file. It is also used to define the **cvadmin**-related control on clients.

i Note: You cannot have a file system named **#SNFS_ALL#**.

Each file system-related element (indicated by the label **securityControl**) has a list of **controlEntry** items. Each **controlEntry** item defines the client and the controls. The client type can be either **host** or **netgrp**. A host can be the IP address or the host name. (Both IP V4 and IP V6 are supported.) Netgrp specifies a group of consecutive IP addresses and has a network IP address (either IP V4 or V6,) and network mask bits. It is possible for there to be overlapping in IP address between an individual host and netgrp, but the individual host should be defined before the netgrp. If a client node has more than one IP address, then define the controls for each IP address.

The following controls are currently supported:

1. **mountReadOnly:** Control whether a client should mount as read-only. The default is read/write.
2. **mountDlanClient:** Control whether a client can mount as a proxy client. The default is not allowed.
3. **takeOwnership:** Control whether users on a Windows client are allowed to take ownership of a file or directory in a StorNext file system.
4. **snfsAdmin:** Controls whether **cvadmin** running on a host is allowed to have super admin privilege to run privileged commands such as start/stop fs.

5. **snfsAdminConnect**: Controls whether **cvadmin** running on a client is allowed to connect to other fsm via the **-H** option.
6. **exec**: Controls whether binary executable files on the file system are allowed to be executed. The default value is “true” (that is, the execution is allowed).
7. **suid**: Controls whether the **setuid** bit is allowed to take effect. (The default value is “true”.)

If no match is found for a given client's IP address, the client has no privileges. If a file system has been defined but the client is not defined in that file system's control section (**securityControl**), the client has no access privileges to the specified file system.

The **denyRetrieves** control prevents a client from triggering on-demand file retrieves. Set this to **true** to prevent the client from triggering retrieves. This control functionality runs on the MDC, so older clients will have the control enforced after you have configured the file **nss_cctl.xml**.

The client controls relate to SAN and LAN clients. SMB and NFS client control is routed through those client's NAS server and cannot be individually controlled. This is particularly important when configuring **denyRetrieves** for Offline File Manager.

- **SAN and DLC clients**: The addresses of the clients themselves.
- **NAS clients via a NAS server on an HA pair**: The VIP of the HA pair.
- **NAS clients via a NAS server on an MDC**: The address of the MDC.
- **NAS clients via a NAS server on a gateway**: The address of gateway.

i Note: If you set **denyRetrieves** to **true** for a NAS server, then this disables retrieves for that server and all NAS clients connecting using that server.

Cluster-Wide Control File Example

The control file called **nss_cctl.xml** provides a feature overview and describes how to configure the central control file. The values entered in the following sample file are for example purposes only and should not necessarily be entered in your own control files.

```
<?xml version="1.0" encoding="UTF-8"?>

<!-- Copyright 2016. Quantum Corporation. All Rights Reserved.      -
->
<!-- StorNext is either a trademark or registered trademark of      -
->
<!-- Quantum Corporation in the US and/or other countries.         -
->

<!--           Cluster-Wide Central Control File                   -->

<!-- The nss_cctl.xml file provides a way to restrict the behavior of -->
<!-- SNFS cluster nodes (fsm, file system client, cvadmin client) from --
```

Appendix C: StorNext Configuration Files

Cluster-Wide Central Control

```
>
<!-- a central place, i.e on nss server. As for SNFS 3.5, we support -->
<!-- the following controls: -->
<!-- 1. Whether a client is allowed to mount as a proxy client -->
<!-- 2. Whether a client is allowed to mount as read/write or read-only.--
>
<!-- 3. Whether a user especially local admin on windows is allowed to --
>
<!-- take ownership of file or directory on a Stornext file system. --
>
<!-- 4. Whether cvadmin running on certain client is allowed to have -->
<!-- super admin privilege to run destructive commands such as -->
<!-- start/stop file system, refresh disks, and change quota setting,--
>
<!-- etc. -->
<!-- 5. whether cvadmin running on certain client is allowed to connect --
>
<!-- to other fsm via "-H" option. -->
<!-- 6. whether an executable file on the file system can be executed. --
>
<!-- 7. whether to allow set-user-identifier bit to take effect. -->

<!-- The control file is in xml format and has hierarchical structure. --
>
<!-- The top level element is "snfsControl", it contains control element -
->
<!-- "securityControl" for certain file system. If you have different -->
<!-- controls for different file systems, then each file system should --
>
<!-- has its own control definition. A special virtual file system -->
<!-- "#SNFS_ALL#" is used as the default control for file systems not -->
<!-- defined in this control file. It is also the required file system --
>
<!-- name when configuring the snfsAdmin and snfsAdminControl options. --
>
<!-- Note: you cannot have a real file system named as "#SNFS_ALL#". --
>

<!-- Each file system related control element (securityControl) has a -->
<!-- list of "controlEntry", each "controlEntry" defines the client and --
>
<!-- the controls. The simplest and preferred way of defining a client -
->
```

Appendix C: StorNext Configuration Files
Cluster-Wide Central Control

```
<!-- is by specifying its IP address (or hostname) by itself, or followed-
->
<!-- by a netmask length separated by a slash (e.g "192.0.2.0/24")      -
->
<!-- if one would like to specify a subnet. Both IPv4 and IPv6 are      -
->
<!-- supported. For backwards compatibility, we support two other ways  -
->
<!-- of defining a client wherein we explicitly specify its type: "host" -
->
<!-- or "netgrp". A "host" can be an IP address or host name. "netgrp"  -
->
<!-- specifies a group of consecutive IP addresses. It has a network IP  -
->
<!-- address (either IPv4 or IPv6) and netmask length. In the case of   -
->
<!-- overlap between client IP addresses, the controls which correspond -
->
<!-- to the IP address with the longest netmask length will take       -
->
<!-- precedence.                                                       -
->

<!-- Currently there are eight controls supported:                       -->
<!-- 1. mountReadOnly: control whether a client should mount as        -->
<!--    readonly. The default is read/write.                            -->
<!-- 2. mountDlanClient: control whether a client can mount as proxy    -->
<!--    client, the default is "mount not allowed".                     -->
<!-- 3. takeOwnership: control whether users on a windows client is     -->
<!--    allowed to take ownership of file or directory of a stornext    -->
<!--    file system. The default is "take ownership not allowed".      -->
<!-- 4. snfsAdmin: whether cvadmin running on a host is allowed to have --
>
<!--    super admin privilege to run privileged commands such as      -->
<!--    start/stop fs. The default is that super admin privilege is not --
>
<!--    honored.                                                         -->
<!-- 5. snfsAdminConnect: whether cvadmin running on a client is allowed -
->
<!--    to connect to other fsm via "-H" option. The default is "-H" is --
>
<!--    not allowed.                                                     -->
<!-- 6. exec: whether binary files on the file system is allowed to    -->
```


Appendix C: StorNext Configuration Files

Cluster-Wide Central Control

```
<!--      be executed.                                -->
<!-- 7. suid: whether set-user-identifier bit is allowed to take effect. -
->
<!-- 8. denyRetrieves: whether the client is allowed to trigger      -->
<!--      dmap read events and retrieve offline files by reading then -->
<!--      default to false, set to true to deny retrieves. The client -->
<!--      will get permission denied errors when reading a truncated file. -
->

<!-- If no match is found for a given client's IP address, then the      -->
<!-- client has no privilege to access a SNFS cluster. If a file system --
>
<!-- has been defined but the client is not defined in that file      -->
<!-- system's control (securityControl), then the client has no access --
>
<!-- privilege to the specified file system.                            -->

<!-- The element "nonVotingCluster" can be included (on the same level as-
->
<!-- the "securityControl" element) to set the default client behavior -
->
<!-- (voting or non-voting) within the cluster during the election that -
->
<!-- will choose the host on which a specific file system manager will -
->
<!-- run. The cluster to which this control is applied will be the one -
->
<!-- specified in the filename. If no cluster is specified in the      -
->
<!-- filename, please refer to the beginning of the DESCRIPTION section -
->
<!-- of the nss_ctl man page for more information on which cluster this -
->
<!-- control will take effect.                                          -
->

<!-- NOTE: There always needs to be voting clients within the cluster so -
->
<!-- that a decision can be derived from the election. Therefore, when -
->
<!-- the "nonVotingCluster" element is set to true, it should be used in -
->
<!-- conjunction with the "votingClients" element (described in the -
```

```
->
<!-- following paragraphs) which allows one to specify an explicit list -
->
<!-- of voting clients. -
->

<!-- It is also possible to specify a group of non-voting clients within -
->
<!-- a cluster by creating a list of client addresses with the element -
->
<!-- "nonVotingClients" (also used on the same level as that of the -
->
<!-- "securityControl" element). The format of the client addresses -
->
<!-- within the "nonVotingClients" element is the same as that used in -
->
<!-- defining a client in the simplest and preferred way within a -
->
<!-- "controlEntry". And there must be at least one address in the list. -
->
<!-- To specify a group of voting clients, the same format is used but -
->
<!-- replacing "nonVotingClients" with "votingClients". -
->

<!-- All three elements (i.e. "nonVotingCluster", "nonVotingClients" and -
->
<!-- "votingClients") may be in the nss_ctl(4) man page at the same -
->
<!-- time. The "votingClients" and "nonVotingClients" elements will take -
->
<!-- precedence over the "nonVotingCluster" element. When a client IP -
->
<!-- address matches elements in both "nonVotingClients" and -
->
<!-- "votingClients", the element with the longest netmask will take -
->
<!-- precedence; if there is a tie, the "votingClients" element will be -
->
<!-- used. -
->

<!-- Currently only Linux platform is supported to be a nss server -->
```

```
<!-- capable of parsing this xml file. -->

<!-- The following is an example to define the nss_ctl.xml. It defines --
>
<!-- the control of file system "snfs", and also the special virtual -->
<!-- file system "#SNFS_ALL#". -->

<snfsControl xmlns="https://www.quantum.com/snfs/cctl/v1.0">
  <nonVotingCluster value="true"/>
  <votingclients>
    <address value="192.0.2.108/24"/>
    <address value="198.51.100.215"/>
  </votingClients>
  <securityControl fileSystem="snfs">
    <controlEntry>
      <client>
        <address value="192.0.2.108"/>
        <address value="198.51.100.215"/>
      </client>
      <controls>
        <mountReadOnly value="false"/>
        <mountDlanClient value="false"/>
        <takeOwnership value="false"/>
        <exec value="true"/>
        <suid value="false"/>
      </controls>
    </controlEntry>
    <controlEntry>
      <client type="host">
        <hostName value="192.0.2.132"/>
      </client>
      <controls>
        <mountReadOnly value="true"/>
        <mountDlanClient value="true"/>
        <takeOwnership value="false"/>
        <denyRetrieves value="true"/>
        <exec value="true"/>
        <suid value="false"/>
      </controls>
    </controlEntry>
    <controlEntry>
      <client type="netgrp">
        <network value="192.0.2.0"/>
      </client>
    </controlEntry>
  </securityControl>
</snfsControl>
```

```
        <maskbits value="24"/>
    </client>
    <controls>
        <takeOwnership value="true"/>
        <mountReadOnly value="true"/>
        <denyRetrieves value="true"/>
    </controls>
</controlEntry>
</securityControl>
<securityControl fileSystem="#SNFS_ALL#">
    <controlEntry>
        <client type="host">
            <hostName value="linux_ludev"/>
        </client>
        <controls>
            <snfsAdmin value="true"/>
            <snfsAdminConnect value="true"/>
        </controls>
    </controlEntry>
</securityControl>
</snfsControl>
```



Appendix D: Quality of Service Feature

This appendix contains the following topics:

Quality of Service

This topic describes how to configure and use the StorNext File System (SNFS) Quality of Service (QOS) feature. QOS allows real-time applications to reserve some amount of bandwidth on the storage system. This is known as real-time I/O (RTIO). SNFS gates (that is, throttles) non-real-time applications so their I/O accesses do not interfere with the real-time application.

QOS is a *passive* implementation in that it does not actively monitor a process' activity and then schedule the process so that it receives the bandwidth it has requested. It is up to real-time applications to gate their own I/O requests to the desired rate. SNFS QOS provides a “get out of the way” gating for non-real-time I/O requests so they do not hinder the real-time requests.

The remainder of this document explains the client and server configuration settings for QOS; describes the use of tokens and callbacks for gating non-real-time I/O; describes setting real-time I/O on a file; and discusses the performance monitoring tools that are available for diagnosis.

- [About Quality of Service on the next page](#)
- [Configuration on page 252](#)
- [Real-time I/O on page 258](#)
- [Non-Realtime I/O Configuration on page 261](#)
- [Callbacks on page 262](#)

- [Monitoring on page 266](#)
- [Sharing Disk Resources on page 267](#)

About Quality of Service

QOS is stripe-group centric; all configuration and operations act on a specific stripe group. This allows stripe groups with differing access characteristics to use different real-time settings. When a client requests a certain amount of real-time I/O, the FSM informs all connected clients of the amount remaining and available for non-real-time I/O accesses.

This information is encapsulated in a token. Think of the token as a capability that enables the client to perform I/O without contacting the FSM. When the amount of real-time I/O on a stripe group changes, the FSM informs all connected clients about the change via a callback.

In both the client and the server, QOS is implemented as a state machine. Each state has a set of events and allowable actions. The term *state* thusly refers to the internal state machine.

Active versus Passive

QOS is a passive, not active implementation of real-time I/O. In an active implementation, guaranteed rate I/O known as GRIO), the scheduler is tightly coupled with the I/O subsystem. The qualities of the disk subsystem are well known so the scheduler can guarantee that a process will be scheduled such that it will receive the required amount of bandwidth. Since SNFS is a cross-platform file system that does not have hooks in the operating system scheduler, it cannot provide such a guarantee.

In a passive implementation, a real-time process gates its I/O according to some outside metric (such as a frame rate for specific video formats). The file system then gates all other non-real-time I/O so they do not interfere.

It is very important that the differences between Active and Passive OOS operation are understood by the user. It is a misconception to think that QOS, despite its name, guarantees a specific amount of real-time I/O to a process.

Supported Platforms

QOS has been tested on Windows XP, Windows 2008, Linux, and Solaris. In Windows, an application gets a handle to a file to perform I/O, usually via the **Win32 CreateFile() API**. In UNIX, an application receives a file descriptor (**fd**) via the **open(2)** system call. The term “handle” is synonymous with **fd**.

Non-Realtime Operation

QoS reserves bandwidth for non-realtime I/O clients. Differentiating between realtime and non-realtime I/O clients should result in better bandwidth utilization and overall performance.

During the mounting process, there is now provision to specify the amount of bandwidth to be reserved. Applications that do not have awareness of the external API can leverage the bandwidth amount. Also, clients with different load characteristics can be differentiated.

Non-realtime I/O support is provided through a central configuration file on the FSM server side, which defines the bandwidth reservation for non-realtime I/O requests on certain clients. This bandwidth reservation (RVIO) has higher priority than normal non-realtime clients, but lower priority over the realtime application (RTIO). FSM will always attempt to first satisfy realtime requests, and then follow the bandwidth reservation request. The remaining bandwidth is evenly distributed to all active non-realtime clients (excluding RVIO reserved clients).

To summarize, bandwidth priority is allocated in the following order:

- Realtime applications (RTIO)
- Bandwidth reservation (RVIO)
- Non-realtime I/O requests

i Note: If there is not sufficient bandwidth to accommodate the bandwidth reservation request (RVIO), the client receives whatever bandwidth remains after RTIO requests have been granted. If there are multiple RVIO clients, then each client is allocated a RVIO bandwidth directly proportional to their requested bandwidth.

Configuration

The primary configuration for QOS is in the FSM configuration file. No client configuration is required, although there is a QOS tuning parameter that can be specified when the file system is mounted.

Unit of I/O

Real-time I/O is based on well-formed I/O. This means that for the purposes of determining bandwidth rates, well-formed I/O is characterized as being a stripe width in size. This makes the best utilization of the disks in the stripe group and maximizes the transfer rate. Internally, non-real-time I/O is tracked by number of I/O operations per second. An I/O operation is a minimum of a file system block size, and a maximum of the file system block size multiplied by the stripe breadth.

$(FsBlocksize \times StripeBreadth)$

$FsBlockSize \leq I/O \leq (FsBlockSize \times StripeBreadth)$

Converting MB/sec to I/O/sec

Typically, it is easier to qualify an I/O subsystem in terms of MB/sec that can be sustained. However, internally the file system tracks everything on an I/O/sec basis. Note that the file system tracks only non-

real-time I/O (that is, it gates only non-real-time I/O). An I/O is a minimum of the file system block size, and is typically the point at which the file system hands the request off to the disk driver (**IoCallDriver** in Windows, or a **strategy** call in UNIX).

The file system counts the number of I/Os that have taken place during a given second. If the number exceeds that which is allotted, the request is pended until I/O becomes available (typically in the next second). I/O is honored in FIFO fashion; no priority is assigned.

To convert between I/Os and MB/sec, SNFS uses a somewhat unique formula that quantifies I/O as well-formed. The rationale behind this is due to the way in which many video applications make real-time I/O requests. To optimize the disk subsystem, real-time I/Os are well-formed so they saturate the disks. In SNFS terminology, this would be an I/O that covers all of the disks in a stripe.

This can be expressed as follows:

$$\frac{IO}{sec} = \frac{\frac{Mb}{sec}}{StripeBreadth \times StripeDepth \times FsBlockSize}$$

For example, with a file system blocksize of 4k, a **StripeBreadth** of 384, and a **StripeDepth** of four, the equivalent number of I/Os/sec for each well-formed I/O would be 216 Mb/sec / (384 * 4 * 4k). This is equivalent to 221184 k/sec / 6144k = 36 I/O/sec.

Server Configuration

All storage subsystems are different, so users must qualify the I/O subsystem and determine the maximum amount of I/O bandwidth available. SNFS relies on the correct setting in the configuration file; if the storage system changes (for example, because of a new disk array,) the user must re-qualify the I/O subsystem to determine the amount of bandwidth available. This amount will be specified in the FSM configuration file. The user can also specify the minimum amount of bandwidth to be provided to non-real-time applications.

There are five keywords controlling QOS that can be specified in the stripe group section of the FSM configuration file. Not all keywords need be present. Typically, the user specifies the RTIO bandwidth in terms of either number of I/O operations per second (**rtios**) or megabytes per second (**rtmb**). Keywords are not case sensitive.

For a minimum configuration, only the real-time limit (either **rtios** or **rtmb**) need be specified. All other configuration variables default to reasonable values.

Table 1: QOS Configuration Keywords

| Name | Description | Default |
|--------------|--|------------------|
| Rtios | The maximum number of real-time I/Os allowed in a stripe group during any one-second period. | 0 (no real-time) |
| Rtmb | Maximum amount of real-time MB/sec allowed on the stripe group during any one-second period. | 0 (no real-time) |

| Name | Description | Default |
|-----------------------|--|---------------------------------|
| RtiosReserve | Amount of reserve in I/Os/sec from the maximum allowed for non-real-time I/Os. Must be greater than the equivalent to 1MB/sec or the amount that can be transferred to a single stripe line. | Equivalent to 1MB/sec |
| RtmbReserve | Amount to reserve in MB/sec from the maximum allowed for non-real-time I/O. | Must be greater than 1. 1MB/sec |
| RtTokenTimeout | Time in seconds to wait for clients to respond to a token callback. | 1.5 seconds |

RTIO Upper Bound

The limit will be specified in terms of I/Os per second (parameter **Rtios**) or in terms of MB/sec (parameter **Rtmb**). Case is not sensitive. Note that I/Os per second are I/Os of any size to the disk subsystem. Either or both may be specified. If both are specified, the lower limit is used to throttle I/O. If neither is specified, no real-time I/O is available on the stripe group. These parameters are applied to a stripe group definition.

Example (Linux)

```
<stripeGroup> index="1" name="MyStripeGroup" realTimeIOs="2048" realTimeMB="10"  
</stripeGroup>
```

Example (Windows)

```
[StripeGroup MyStripeGroup]  
Rtios 2048  
Rtmb 10
```

The above example specifies that the storage system can support a maximum of 2048 I/Os per second at any instant, aggregate among all the clients, or 10 MB/sec, whichever is lower.

Most real-time I/O requests will be a stripe line at a time to maximize performance. Non-real-time I/Os will be a minimum of a file system block size.

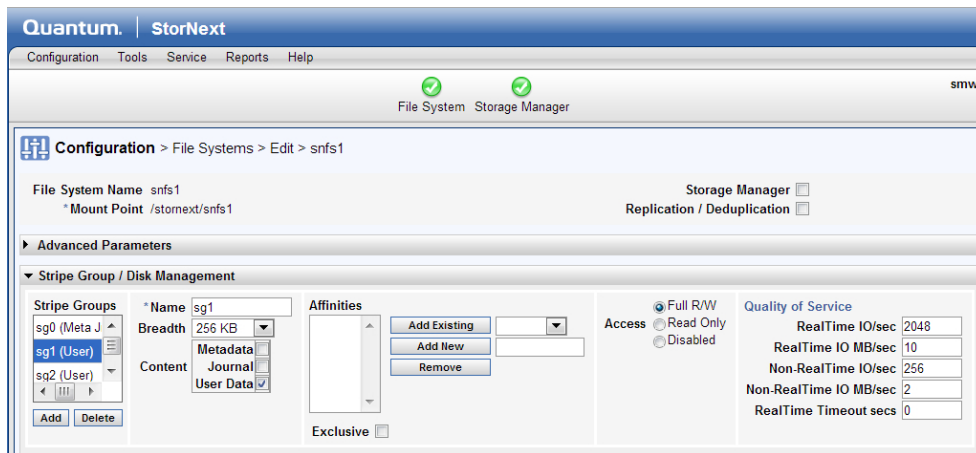
i Note: It is important to realize that the **rtios** and **rtmb** settings refer to the total amount of sustained bandwidth available on the disk subsystem. Any I/O, either real-time or non-real-time, will ultimately be deducted from this overall limit. The calculations of available real-time and non-real-time are discussed later.

Specifying **rtmb** in the FSM configuration file is only recommended if all I/Os are well formed (that is, a full stripe width). Otherwise, the conversion between MB/sec and I/Os/sec using the well-formed I/O calculation could lead to unexpected results.

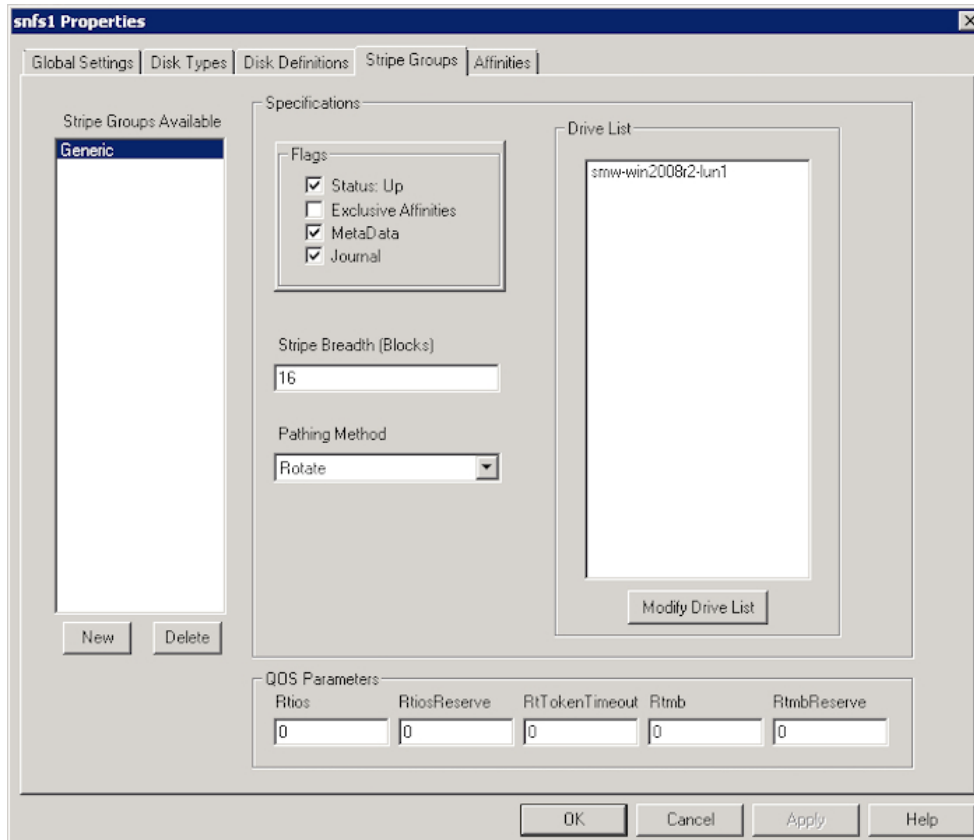
Modify

To change the RTIO parameters on the MDC, modify the `/usr/cvfs/config/<fname.cfgx>` file, and restart StorNext services. Refer to the examples above for the Linux and Windows RTIO parameter names. You can also use the StorNext GUI.

- For a file system on a Linux MDC, select a Stripe Group and modify the QOS, as shown:



- In a Windows MDC, you can modify these values with the **File System Cfg(Advanced)** tool. To change the QOS parameters, select **File System** and then click **Stripe Groups Tab**, as shown:



Reserve

To prevent deadlock, the QOS implementation never allows zero I/O/sec for non-real-time I/O. Otherwise, a system could block with many critical file system resources held waiting for I/O to become available. This is especially true via flush-on-close I/O via the buffer cache. It becomes extremely difficult to diagnose system hangs because no I/O is available. For this reason, QOS always reserves some amount of I/O for non-real-time I/O.

The minimum amount of non-real-time I/O reserved for non-real-time applications is one MB/sec. This can be changed via the stripe group section parameters (again, case is not sensitive). If both are specified, the lower of the two amounts is chosen. This amount is shared by all non-real-time applications on each client.

Example (Linux)

```
<stripeGroup> index="1" name="MyStripeGroup" realTimeIOsReserve="256"
realTimeMBReserve="2"
```

```
</stripeGroup>
```

Example (Windows)

```
[StripeGroup MyStripeGroup]  
RtiosReserve 256  
RtmbReserve 2
```

Token Timeouts

The **RtTokenTimeout** parameter controls the amount of time the FSM waits for clients to respond to callbacks. In most normal SANs, the default two-second setting is sufficient. This value may need to be changed for a SAN that has a mixture of client machine types (Linux, Windows NT, etc.) that all have different TCP/IP characteristics. Also, large numbers of clients (greater than 32) may also require increasing the parameter.

For example, if the FSM should ever fail, the clients will attempt to reconnect. When the FSM comes back online, the amount of time the clients take to re-establish their TCP/IP connection to the FSM can differ wildly. To avoid unnecessary timeouts, the **RtTokenTimeout** parameter can be increased, meaning the FSM waits longer for callback responses.

If a client times out on a token retraction, the original requestor receives an error from the FSM that includes the IP address of the offending client. This error is logged to **syslog**, and alternatively to the desktop on Windows clients. This can help in diagnosing reconnect failures, and in determining if the token time value should be increased.

Client Configuration

When a client obtains a non-real-time I/O token from the FSM, the token allows the client a specific amount of non-real-time I/O. If the client is inactive for a period of time, the token is relinquished and the non-real-time I/O released back to the FSM for distribution to other clients. The timeout period is controlled by the **nrtiotokenhold** mount option on UNIX platforms, and the **QOS Token Hold Time** parameter in the mount options tab of the SNFS control panel on Windows platforms. The default is sixty (60) seconds.

This means that after sixty seconds without non-real-time I/O on a stripe group, the non-real-time token for that stripe group is released. The parameter should be specified in five (5) second increments. If it is not, it will be silently rounded up to the next five-second boundary. If the **syslog** level is set to **debug**, the file system dumps out its mount parameters so the value can be seen.

Real-time I/O

A process requests real-time (ungated) I/O by using the SNFS External API **SetRtio** call (**F_SETRIO ioctl**). A library function is included in the External API sample source code that provides all the required cross-platform handling.

As an example, assume that a video playback application requires a constant rate of 186 MB/sec to correctly display images without dropping any frames. The application gates itself; that is, it requests I/O at a rate to satisfy the requirements of correctly displaying an image. QOS provides a mechanism so other I/O requests do not perturb the real-time display.

In the following example, assume the I/O subsystem has been qualified at 216 MB/sec. The file system block size is 4k. The disk subsystem is actually a large RAID array that internally maps many drives to a single LUN. There are four LUNs in the stripe group; each LUN is optimized for a 1.5 MB transfer. This corresponds to the following in the fsm configuration file:

Example (Linux)

```
<stripeGroup index="1" name="MyStripeGroup" stripeBreadth="384"
realTimeMB="216">
<disk index="0" diskLabel="CvfsDisk0" diskType="VideoDrive"/>
<disk index="0" diskLabel="CvfsDisk1" diskType="VideoDrive"/>
<disk index="0" diskLabel="CvfsDisk2" diskType="VideoDrive"/>
<disk index="0" diskLabel="CvfsDisk3" diskType="VideoDrive"/>
</stripeGroup>
```

Example (Windows)

```
[StripeGroup MyStripeGroup]
StripeBreadth 384
Node CvfsDisk0 0
Node CvfsDisk1 1
Node CvfsDisk2 2
Node CvfsDisk3 3
Rtmb 216
```

Also, assume there is only one stripe group for user data in the file system. As recommended by Quantum, there may be other stripe groups for metadata and journal that are not shown.

SetRtio

Initially, all stripe groups in the file system are in non-real-time mode. Clients make their requests directly to the I/O subsystem without any gating. In our example, the process requires 186 MB/sec and the system designers know there will never be a need to support more than one stream at 216 MB/sec.

The **SetRtio** request has a number of flags and parameters to control its operation. These are all documented in the [StorNext File System API Guide](#) that describes the external API in detail. For this example, set the handle for the indicated stripe group using the **RT_SET** parameter.

Oversubscription

In most cases, system designers ensure that the amount of **RTIO** is not oversubscribed. This means that processes will not ask for more **RTIO** than is specified in the configuration file. However, it is possible to request more **RTIO** than is configured. The API uses the **RT_MUST** flag to indicate that the call must succeed with the specified amount. If the flag is clear, the call allocates as much as it can. In both cases, the amount allocated is returned to the caller.

Handles


The **SetRtio** call accepts two different types of handles. The first is a handle to the **root** directory. In this mode the stripe group is put into real-time mode, but no specific file handle is tagged as being **ungated**. Real-time I/O continues on the stripe group until it is explicitly cleared with a **SetRtio** call on the **root** directory that specifies the **RT_CLEAR** flag; the file system is unmounted; or the system is rebooted. It is up to the application to make a subsequent call to **EnableRtio (F_ENABLERTIO)** on a specific handle.

If the handle in the **SetRtio** call refers to a regular file, it is the equivalent of a **SetRtio** call on the root directory followed by an **EnableRtio** call. The file handle will be ungated until it is closed, cleared (**RT_CLEAR** in a **SetRtio** call), or disabled (**DisableRtio**). When the handle is closed, the amount of real-time I/O is released back to the system. This causes the FSM to readjust the amount of bandwidth available to all clients by issuing a series of callbacks.

The client automatically issues a call to the FSM with the **RT_CLEAR** flag specifying the amount of real-time I/O set on the file. If multiple handles are open on the file—each with a different amount of real-time I/O—only the last file close triggers the releasing action; all aggregate **RTIO** are released.

This automatic clearing of real-time I/O is carried out in the context of the process that is closing the file. If the FSM cannot be reached for some reason, the request is queued on a daemon and the process closing the file is allowed to continue. In the background, the daemon attempts to inform the FSM that the real-time I/O has been released.

Different processes can share the same file in real-time and non-real-time mode. This is because the level of gating is at the handle level, not the file level. This allows a real-time process to perform ingest of material (video data) at the same time as non-real-time processes are performing other operations on the file.



In the illustration, Process A has ungated access to file **foo**. Processes B and C also are accessing file **foo**, but the client gates their I/O accesses. If multiple handles are open to the same file and all are in real-time mode, only the last close of the handle releases the real-time I/O back to the system. This is because on most platforms the file system is informed only on the last close of a file.

Ungated files

It is also possible to denote using the **RT_NOGATE** flag that a handle should not be gated without specifying any amount of real-time I/O. This is useful for infrequently accessed files (such as index files) that should not be counted against the non-real-time I/O. System designers typically allow for some amount of overage in their I/O subsystem to account for non-gated files.

Calculating Available RTIO

When the FSM receives a request for RTIO, it takes the amount reserved into consideration. The reserve amount functions as a soft limit beyond which the FSM will not triapse. The calculation for **RTIO** is as follows:

$$\text{available_rtio} = (\text{rtio_limit}) - (\text{rtio_current}) - (\text{rtio_reserve})$$

In the above calculation, **rtio_limit** is the stripe group's maximum number of IOs (the value of parameter **Rtios**), **rtio_current** is the total number of currently reserved real-time IOs, and **rtio_reserve** is the minimum IOs reserved for non-realtime IOs (the value of **RtiosReserve**).

All internal calculations are done in terms of I/O/sec.

Non-Realtime I/O Configuration

In order to reserve bandwidth for non-realtime I/O (RVIO) operations on certain clients, you must specify the amount in the QOS central configuration file. This file is located under **\$CVFSROOT/config/** (for example, **/usr/cvfs/config**). The configuration file name consists of the file system name and the suffix **_rvio.opt**. For example, the central configuration file name for a file system called **fs1** would be **fs1_rvio.opt**.

The configuration does not affect the normal RTIO requests; it affects only the bandwidth allocation to non-RTIO requests. Note the following characteristics of RVIO:

1. RTIO has higher priority and is always satisfied first. RVIO uses whatever bandwidth is left.
2. RVIO reservation is not guaranteed. Instead, it is adjusted dynamically. When there are new RVIO or RTIO requests, the bandwidth reserved for clients with previous RVIO requests is readjusted. Each RVIO client will get the RVIO bandwidth reserved proportional to its request amount.
3. The RVIO bandwidth allocated for a client is the total bandwidth for non-RTIO requests on the client. The bandwidth (token) will not return like a normal non-RTIO client does if it's been idle for a hold time. Non-RTIO requests on clients with an RVIO reservation is subject to the same throttling mechanism as non-RTIO requests on normal clients.
4. Other active non-RTIO clients equally share the bandwidth left over after the possible RTIO and RVIO allocations. If no bandwidth is left, the minimal reserved bandwidth (defined in the stripe group section of the file system configuration file) is used.

A central configuration file has multiple entries, and each entry defines the RVIO reservation for a client. Following are the fields of an entry:

`host bw-type sg=yy,[sg=yy]`

- **host** is the host name, IP address (either V4 or V6), or FQDN, etc. of the client. The host name should be able to be resolved (converted to a valid IP address) by the FSM server. If there are multiple IPs on the host, use the IP that is used to communicate with the FSM.
 - **bw-type** is the type of bandwidth to be specified. There are two bandwidth type options:
 - **qosmb** means the subsequent value for bandwidth reservation uses MB/s (megabytes per second) as the unit of measure.
 - **qosios** means the subsequent value for bandwidth reservation uses IOs/s (inputs/outputs per second) as the unit of measure.
 - **sg** stands for stripe group. You may specify a stripe group by entering the stripe group name, or by stripe group number starting from 0 for the first stripe group. If you'd like to specify the same RVIO bandwidth reservation for all eligible stripe groups, you can use the asterisk wildcard in the format ***=yy**.

Any line in the configuration file starting with the pound sign (#) is considered a comment line and is not executed.

For additional information about entering the configuration file for QOS, see the man page called **qos_config**.

Configuration File Examples

Below are some configuration file entry examples:

```
box1.xyz.com qosios 2=200,4=50,sg5=45  
10.10.1.8 qosmb *=60
```


Callbacks

The cornerstones of the communications between the FSM and the client are callbacks and tokens. A callback is an unsolicited message from the FSM to the client requesting that the client adjust its real-time I/O parameters. The callback contains a token that specifies the amount of non-real-time I/O available on a stripe group.

Initially, all stripe groups in a file system are in non-real-time (ungated) mode. When the FSM receives the initial request for real-time I/O, it first issues callbacks to all clients informing them that the stripe group is now in real-time mode. The token accompanying the message specifies no I/O is available for non-real-time I/O. Clients must now obtain a non-real-time token before they can do any non-real-time I/O.

After sending out all callbacks, the FSM sets a timer based on the **RtTokenTimeout** value, which by default is set to 1.5 seconds. If all clients respond to the callbacks within the timeout value the **RTIO** request succeeds, and a response is set to the requesting client.

Figure 1: Callback Flow for Initial RTIO Request



In the above diagram, a process on client A requests some amount of **RTIO** in Step 1. Since this is the first request, the FSM issues callbacks to all connected clients (Steps 2-5) informing them that the stripe group is now in real-time mode. The clients respond to the FSM in Steps 6-9. After all the clients have responded, the FSM responds to the original requesting client in Step 10.

If the timer expires and one or more clients have not responded, the FSM must retract the callbacks. It issues a response to the requesting client with the IP number of the first client that did not respond to the callback. This allows the requesting client to log the error with the IP number so system administrators have a chance of diagnosing the failure. It then sends out callbacks to all the clients to which it first sent

the callbacks, retracting them to the original state. In our example, it would set the stripe group back to non-real-time mode.


After sending out the callbacks, the FSM waits for a response using the **RtTokenTimeout** value as before. If a client again does not respond within the timeout value, the callbacks are retracted and sent out again. This repeats until all clients respond. During this time of token retractions, real-time requests cannot be honored and will only be queued.

Callback Failures

The FSM must handle a case where a client does not respond to a callback within the specified timeout period (**RtTokenTimeout**). If a client does not respond to a callback, the FSM must assume the worst: that it is a rogue that could wreak havoc on real-time I/O. It must retract the tokens it just issued and return to the previous state.

As mentioned earlier, the original requestor will receive an error (**EREMOTE**) and the IP address of the first client that did not respond to the callback. The FSM enters the token retraction state, and will not honor any real-time or token requests until it has received positive acknowledgement from all clients to which it originally sent the callbacks.

Figure 2: Callback Retraction Example



In [Figure 2 above](#), Client A requests some amount of RTIO as in [Figure 1 on the previous page](#). However, assume that Client C did not respond to the initial callback in time (step 7). The FSM will return a failure to Client A for the initial RTIO request, then send out callbacks to all clients indicating the stripe group is no longer real-time (steps 11-14). In the example, Client C responds to the second callback, so the FSM will not send out any more callbacks. The stripe group is back in non-real-time mode.

Note that this can have interesting repercussions with file systems that are soft mounted by default (such as Windows). When the caller times out because other clients are not responding and then gives up and returns an error to the application, if at some point the FSM is able to process the **RTIO** request it may result in the stripe group being put into real-time mode after the original caller has received an error code. Both the FSM and clients log their actions extensively to **syslog**, so if this situation arises it can be detected.

In [Figure 2 on the previous page](#), if the stripe group were already in real-time mode the FSM would only send out callbacks to those clients that already have tokens. Once all clients responded to the token callbacks, the stripe group would be back in its original state.

Tokens

A token grants a client some amount of non-real-time I/O for a stripe group. Tokens are encapsulated in callback messages from the FSM. Initially, no tokens are required to perform I/O. Once a stripe group is put into real-time mode, the FSM sends callbacks to all clients informing them that they will need a token to perform any non-real-time I/O. The first I/O after receiving the callback will then request a non-real-time I/O token from the FSM.

The FSM calculates the amount of non-real-time bandwidth using the following formula:

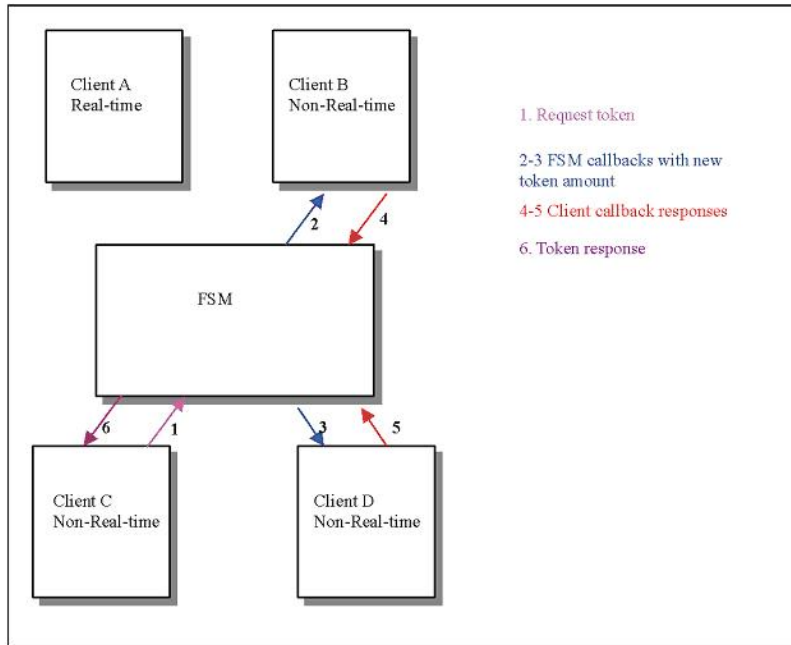
$$\text{available_nrtio} = \frac{(\text{rtio_limit}) - (\text{rtio_current}) - (\text{rvio_current})}{(\text{current_num_nrtio_clients}) + 1}$$

In the above calculation, **rvio_current** is the total bandwidth reserved by current RVIO clients. The amount of existing real-time I/O (**rtio_current**) has already been adjusted with the reserve parameter. As each client requests a non-real-time I/O token, the number of clients increases (**current_num_nonrtio_clients** in the above formula) and the amount of available non-real-time I/O decreases.

Each time there is a change in the amount of non-real-time I/O available, the FSM sends callbacks to the clients with tokens. It is important to note that unlike the initial set of callbacks where the FSM sent callbacks to all connected clients, it is now only necessary to send callbacks to those clients that have an existing token.

Once a client has a token, it can perform as much I/O per second as is allowed by that token. It does not need to contact the FSM on every I/O request. The FSM will inform the client whenever the token changes value.

Figure 3: Non-Real-time Token Adjustments



In [Figure 3 above](#), assume the stripe group is already in real-time mode as a result of an **RTIO** request from client A. Clients B and D are doing non-real-time I/O to the stripe group and have a token that specifies the amount of non-real-time I/O available. Client C then requests a non-real-time I/O token in Step 1. The FSM calls back to Clients B and D and specifies the new token amount in Steps 2-3. The clients respond in Steps 4-5, acknowledging the new token amount. The FSM then responds to Client C with the new token in Step 6.

Failure Semantics

There are two major failures that affect QOS: FSM crashes and client crashes. These can also be loss of communication (network outages). For client and server failures, the system attempts to readjust itself to the pre-failure state without any manual intervention.

FSM Failures

If the FSM crashes or is stopped, there is no immediate affect on real-time (ungated) I/O. As long as the I/O does not need to contact the FSM for some reason (attribute update, extent request, etc.), the I/O will continue. From the standpoint of QOS, the FSM being unavailable has no affect.

Non-real-time I/O will be pended until the FSM is re-connected. The rationale for this is that since the stripe group is in real-time mode, there is no way to know if the parameters have changed while the FSM is disconnected. The conservative design approach was taken to hold off all non-real-time I/O until the FSM is reconnected.

Once the client reconnects to the FSM, the client must re-request any real-time I/O it had previously requested. The FSM does not keep track of QOS parameters across crashes; that is, the information is

not logged and is not persistent. Therefore, it is up to the clients to inform the FSM of the amount of required **RTIO** and to put the FSM back into the same state as it was before the failure.

In most cases, this results in the amount of real-time and non-real-time I/O being exactly the same as it was before the crash. The only time this would be different is if the stripe group is oversubscribed. In this case, since more **RTIO** had been requested than was actually available, and the FSM had adjusted the request amounts, it is not deterministically possible to re-create the picture exactly as it was before. Therefore, if a deterministic picture is required across reboots, it is advisable to not over-subscribe the amount of real-time I/O.

The process of each client re-requesting **RTIO** is exactly the same as it was initially; once each client has reestablished its **RTIO** parameters, the non-real-time I/O is allowed to proceed to request a non-real-time token. It may take several seconds for the SAN to settle back to its previous state. It may be necessary to adjust the **RtTokenTimeout** parameter on the FSM to account for clients that are slow in reconnecting to the FSM.

Client Failures

When a client disconnects either abruptly (via a crash or a network partition,) or in a controlled manner (via an unmount), the FSM releases the client's resources back to the SAN. If the client had real-time I/O on the stripe group, that amount of real-time I/O is released back to the system. This causes a series of callbacks to the clients (all clients if the stripe group is transitioning from real-time to non-real-time,) informing them of the new amount of non-real-time I/O available.

If the client had a non-real-time I/O token, the token is released and the amount of non-real-time I/O available is recalculated. Callbacks are sent to all clients that have tokens informing them of the new amount of non-real-time I/O available.

Client Token Releases

While it is not a failure case, the handling of a client token release is exactly the same as in the case where the client disconnected. All clients retain non-real-time tokens for a fixed amount of time. The default is 60 seconds. This can be controlled via the **nrtiotokentimeout** mount option. After the specified period of inactivity (i.e., no non-real-time I/O on the stripe group), the client will release the token back to the FSM. The FSM will re-calculate the amount of non-real-time bandwidth available, and send out callbacks to other clients.

Therefore, if a situation exists where a periodic I/O operation occurs every 70 seconds, it would be beneficial to set the **nrtiotokentime** mount option to something greater than or equal to 70 seconds to cut down on system and SAN overhead.

Monitoring

The current real-time statistics are available via the **cvadmin** utility. The **show long** command has been enhanced to provide information as to the current limit, the minimum amount reserved for non-real-time I/O, the number of active clients, the amount currently committed, and the amount a non-real-time

application could hope to get when requesting I/O.

Whenever the stripe group status changes (such as from non-real-time to real-time mode), an event is logged to **syslog** (system event log on Windows platforms).

On the Windows platform, real-time performance is also available via the **perfmon** utility. There are counters for both the client (SNFS Client) and FSM (SNFS File System Server [FSS]). In the client, a number of **rtio_xxx** counters are available to track the number of real-time I/Os/sec, number of non-real-time I/O requests, non-real-time I/O starvation, and other counters. A full explanation of each counter is provided with the **perfmon** utility by clicking Explain. In the FSM, information about the number of outstanding non-real-time clients, available **RTIO**, and other QOS information is available.

Performance counters are not enabled by default in the client. To enable them, you must execute **cvdb -P**; this toggles the state of collection. To turn off the **perfmon** counters, do another **cvdb -P**. To re-enable, again do **cvdb -P**, and so on. To clear all client counters to zero, disable, then re-enable the counters with **cvdb -P**.

Sharing Disk Resources

As a final example, consider the case of a full-featured disk subsystem (Redundant Array of Very Expensive Disks, or RAVED) using a very high speed interconnect. Many clients can share these disk arrays, but it is sometimes desirable to limit a client's access to the array. QOS provides a mechanism for political bandwidth management so that no one client can consume all the bandwidth of the channel interconnects.

In such a scenario, the stripe group would always be in real-time mode. Each client would have a token specifying the number of I/Os/sec permissible. If there is need to assign different reserved bandwidth for non-real-time clients, specify the client's bandwidth reservation in the RVIO config file.

The foundation of such an approach is a simple program that puts the stripe group into real-time mode as soon as the FSM is up and servicing requests. An example of such a program is included in the source code for the External API.

Once the stripe group is in real-time mode, the bandwidth as specified in the FSM configuration file is shared by all clients. In cases such as this, the real-time limit (**rtios** or **rtmb**) is calculated to be the total bandwidth desired to each client times the number of possible clients.

As each client attempted to access the disk subsystem, it would obtain a token. The FSM would send out callbacks adjusting down the amount of bandwidth available. No one client would be allowed to exceed the threshold specified in the non-realtime token. This assures fairness among all the clients.