SGI® InfiniteStorage NAS System Administration Guide

007-5595-001

COPYRIGHT

LIMITED RIGHTS LEGEND

TRADEMARKS AND ATTRIBUTIONS

Record of Revision

Version	Description
001	October 2009. Original publication.

Table of Contents

1 About This Guide

2 Overview of SGI InfiniteStorage NAS Storage Systems

3 Quick System Configuration

4 Optional Configuration Steps

5 Network Configuration

6 Storage Management

7 File Services

8 Data Protection

9 Scalability and Clustering

10 Status and Monitoring

11 Maintenance Tasks

A Using Storage Management Applications

1 About This Guide

This guide provides an overview of the architecture, general operation, and descriptions of the major components in the SGI® InfiniteStorage NAS Server.

The SGI InfiniteStorage NAS Server (IS-NAS Server) and the Titan Server are enterprise-class network storage servers that provide high-performance read/write access to data through multiple protocols, such as CIFS, NFS, iSCSI, and FTP. These systems support:

- Stand-alone servers, or clusters of up to eight nodes
- Tiered storage
- A single name space with global access
- Virtualized Storage Pools, servers, file systems, and volumes
- Read caching
- Global symlinks
- Data protection through NVRAM buffering, snapshots, NDMP-based backup, virus scanner integration, automatic data migration, policy-based data replication, data relocation (with transfer of primary access), quorum devices
- User and group quotas
- Complete tools for managing the system and monitoring system status, including Web Manager (a browser-based graphical user interface) and a command line interface.

License keys are used to control the availability of some system features and functionality. For more information on licenses, see Managing License Keys, on page 535 or contact SGI Global Services.

Audience

This guide is written for owners and system administrators of SGI InfiniteStorage NAS Server systems. It is written with the assumption that the reader has a good working knowledge of networking concepts and practices, computers, and computer systems.

Chapter Descriptions

The following topics are covered in this document:

- Chapter 1: "About This Guide"

  Provides an introduction to this guide, lists other documentation resources available for this product, and explains the conventions used in this document.

- Chapter 2: "Overview of SGI InfiniteStorage NAS Storage Systems"

  Provides an overview of the NAS storage server system, including the parts of the system, management roles and management interfaces.

- Chapter 3: "Quick System Configuration"

  Explains how to use a wizard to set up an external SMU, selecting managed systems, and setting up an SGI InfiniteStorage NAS Server using a wizard.

- Chapter 4: "Optional Configuration Steps"

  Provides information on configuring administrative user accounts, SMU security, SMU access to managed servers, the private management network, and receiving SNMP traps.

- Chapter 5: "Network Configuration"

  Provides information about all aspects of the usage and configuration of the networks to which the SGI InfiniteStorage NAS Server system is connected. This includes network interfaces, IP addressing concerns, network statistics, name services, and directory services.

- Chapter 6: "Storage Management"

  Provides information about the physical and logical elements of the storage attached to the SGI InfiniteStorage NAS Server system. This chapter also explains how to use features of the NAS serve system and manage the various storage elements making up the complete storage system.

- Chapter 7: "File Services"

  Provides information about the file services supported, and explains how to configure and manage those services.

- Chapter 8: "Data Protection"

  Provides information about how to protect your data, including feature explanations, and the configuration, management, and usage of these features.

- Chapter 9: "Scalability and Clustering"

  Provides information about enlarging the NAS storage system, creating clusters, and managing large installations.

- Chapter 10: "Status and Monitoring"

  Provides information about how status information and statistics are provided, what that information means, and how to configure the system

to provide and display the information necessary for managing the NAS server system.

- Chapter 11: "Maintenance Tasks"

  Provides information and instructions on routine maintenance procedures and requirements.

- Appendix A: "Using Storage Management Applications"

  Provides information about the NDMP application variables supported by the SGI InfiniteStorage NAS Server.

Related Publications

The following documents are relevant to the SGI InfiniteStorage NAS Server:

- *InfiniteStorage NAS Server Storage Subsystem Guide:* In PDF format, this guide provides information about using the storage subsystems attached to the storage server/cluster.
- *InfiniteStorage NAS Server Software Installation Guide:* In PDF format, this guide provides information about installing software and firmware, including instructions on how to upgrade and downgrade the storage server and the SMU.
- *InfiniteStorage NAS Server Hardware Installation Guide:* In PDF format, this guide provides information about installing the storage server and connecting it to your network.
- *InfiniteStorage NAS Server Hardware Reference:* This guide (in PDF format) provides an overview of the InfiniteStorage NAS Server hardware, describes how to resolve any problems, and shows how to replace faulty components.
- *Titan Server Hardware Reference:* This guide (in PDF format) provides an overview of the Titan Server hardware, describes how to resolve any problems, and shows how to replace faulty components.
- *InfiniteStorage NAS Server Command Line Reference:* This guide (in HTML format) describes how to administer the system by typing commands at a command prompt.
- *InfiniteStorage NAS Server Release Notes:* This document gives late-breaking news about the system.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
`Command`	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
variable	The italic typeface denotes variable entries and words or concepts being defined. Italic typeface is also used for book titles.
user input	This bold fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[and]	Brackets enclose optional portions of a command or directive line.
…	Ellipses indicate that a preceding element can be repeated.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.
	**Tip**: A tip contains supplementary information that is useful in completing a task.
	**Note**: A note contains information that helps to install or operate the system effectively.
	**Caution**: A caution indicates the possibility of damage to data or equipment. Do not proceed beyond a caution message until the requirements are fully understood.

Browser Support

Any of the following browsers can be used to run Web Manager, the System Management Unit (SMU) web-based graphical user interface.

- Microsoft Internet Explorer: version 7.0 or later.
- Mozilla Firefox: version 1.5 or later.

The following Java Runtime Environment is required to enable some advanced Web Manager functionality: Sun Microsystems Java Runtime Environment: version 5.0, update 6, or later.

Some product documentation is included for download or viewing through Web Manager. The following software is required to view this documentation: Adobe Acrobat: version 7.0.5 or later.

Product Support

SGI provides a comprehensive product support and maintenance program for its products. SGI also offers services to implement and integrate Linux applications in your environment.

- Refer to http://www.sgi.com/support/.
- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

Send e-mail to the following address: techpubs@sgi.com.

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

  SGI
  Technical Publications
  46600 Landing Parkway
  Fremont, CA 94538

SGI values your comments and will respond to them promptly.

2 Overview of SGI InfiniteStorage NAS Storage Systems

The IS-NAS Server and the Titan Server are enterprise-class network storage servers that provide high-performance read/write access to data through multiple protocols, such as CIFS, NFS, iSCSI, and FTP. These systems support:

- Stand-alone servers, or clusters of up to eight nodes
- Tiered storage
- A single name space with global access
- Virtualized Storage Pools, servers, file systems, and volumes
- Read caching
- Global symlinks
- Data protection through NVRAM buffering, snapshots, NDMP-based backup, virus scanner integration, automatic data migration, policy-based data replication, data relocation (with transfer of primary access), quorum devices
- User and group quotas
- Complete tools for managing the system and monitoring system status, including Web Manager (a browser-based graphical user interface) and a command line interface.

License keys are used to control the availability of some system features and functionality. For more information on licenses, see Managing License Keys, on page 535 or contact SGI Global Services.

Physical and Logical Components

The IS-NAS Server and the Titan Server are highly scalable and modular Network Attached Storage (NAS) servers, with multi-gigabit throughput from network to disk. They consist of the following elements:

- A System Management Unit (SMU), either internal or external, depending on model and system configuration
- IS-NAS Server(s) and/or Titan Server(s)
- Virtual servers (EVSs)
- Private management network
- Public data network

- Storage subsystem(s)

System Management Unit (SMU)

The SMU manages the storage servers and/or clusters, and provides front-end server administration and monitoring tools. It supports clustering, data migration, and replication, and acts as the Quorum Device in a cluster. Although integral to the server, the SMU is not in the data movement path between the network client and the server.

There are two kinds of SMU; external and internal.

- An external SMU can manage up to eight (8) storage servers/clusters in any combination. Each external SMU can manage both IS-NAS Servers and Titan Servers or clusters.

  An external SMU is a separate device in the storage server system. To eliminate the SMU as a single point of failure, you can configure your system with a second external SMU as a standby SMU.

- An internal SMU can manage a single stand-alone IS-NAS Server (an external SMU is required to manage more than a single IS-NAS Server).

  An internal SMU is a service that runs on the IS-NAS Server and provides the same management and monitoring functionality as an external SMU. When using an internal SMU, there is no way to configure a standby SMU.

Storage Server

The storage server's patented architecture is structured around bi-directional data pipelines and a hardware-based file system. It scales to 4 petabytes of data, supporting higher sustained access loads without compromising performance. It can be configured as a single server or as a server cluster. All network clients communicate directly with the server.

The server processes file access requests from network clients via Gigabit Ethernet. It reads and writes from/to one or multiple storage devices, connected through Fibre Channel (FC) links. The storage subsystem can be configured with a single server. It can also be configured with multiple servers clustered together; as they share the same storage devices, network requests can be distributed across cluster nodes.

Note: While the first generation blades supported clusters with up to two nodes, the Titan Server now supports clusters with up to eight nodes. Should one cluster node fail, its file services and server administration functions are transferred to other nodes. All nodes in a cluster must be of the same series. A cluster may not be made up of a mixture of Titan Server series.

The server is rack mountable and consists of a passive backplane (not removable), three hot-swappable fan trays and two hot swappable redundant power supplies. The front panel displays system status with a green power LED and an amber fault LED. The only Field Replaceable Units (FRUs) accessible from the front of the server chassis are the cooling fans. The unit is serviced from its rear panel, which includes additional status LEDs, connectors (power, Ethernet, Fibre Channel, RS-232), and FRUs, such as the

power supplies and server modules. For more information about the server hardware, see the Titan Server *Hardware Reference*.

Virtual Servers (EVSs)

All file services are provided by virtual server entities referred to as "EVSs". Each EVS is assigned unique network settings and storage resources, enabling administrators the flexibility to logically partition access to shared storage resources. In server clusters, EVSs are automatically migrated between servers when faults occur to ensure maximum availability. When multiple servers or clusters are configured with common storage access, they are referred to as Server Farms. EVSs can be manually migrated between servers in a Server Farm based on performance and availability requirements.

Private Management Network

The private management network connects the SMU and devices such as Fibre Channel (FC) switches, and uninterruptible power supply (UPS) units.

The private management network is isolated from the public data network by the SMU. The private management network connects the private management interface of the SMU, the Ethernet management interface on the server, and all of the Ethernet managed devices that make up the storage system.

Devices on the private management network are only accessible from the public data network through the SMU, which uses Network Address Translation (NAT) to convert the public IP address of a device into the device's address on the private management network.

An external SMU has two 10/100/1000 Mbps Ethernet interfaces. The *eth0* interface connects to the public data network, and *eth1* connects to the private management network.

The following illustration shows the private management network connections with an external SMU.

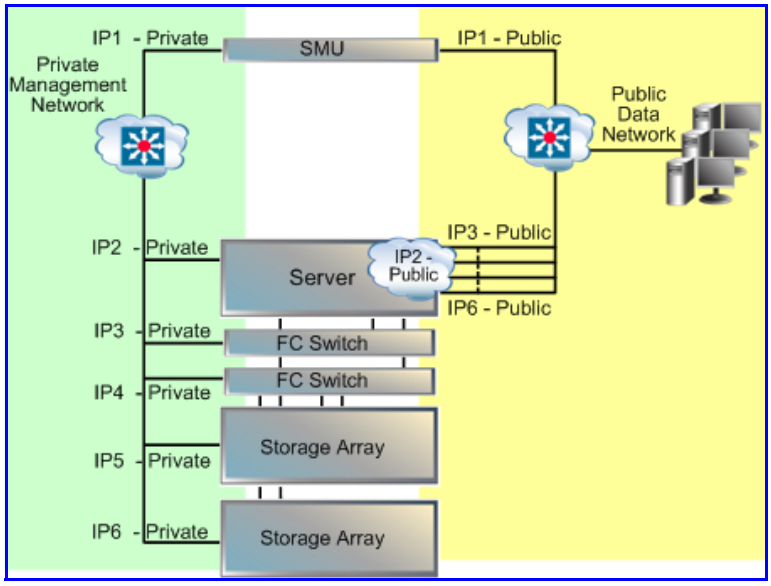

An internal SMU uses two 10/100/1000 Mbps Ethernet ports on the storage server motherboard as management interfaces. Like the ports on an external SMU, *eth0* connects to the public data network and *eth1* connects to the private management network.

The following illustration shows the private management network connections with an internal SMU.

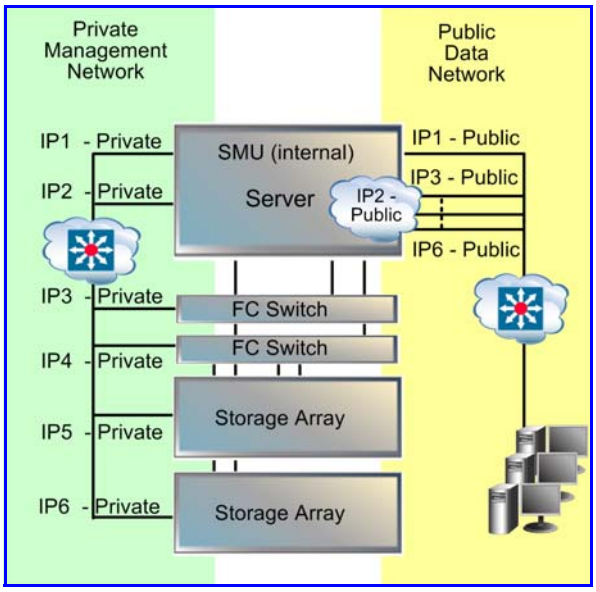

The diagram below shows how NAT isolates the private management network. The example shows a device with the IP address 192.0.2.13:80 accessible through HTTP, and a second device with IP address 192.0.2.14:443 accessible through HTTPS. These devices appear on the Public Data network as 10.1.1.13:28013 and 10.1.1.13:28014 (where 10.1.1.13 is the SMU's IP address on the public data network).

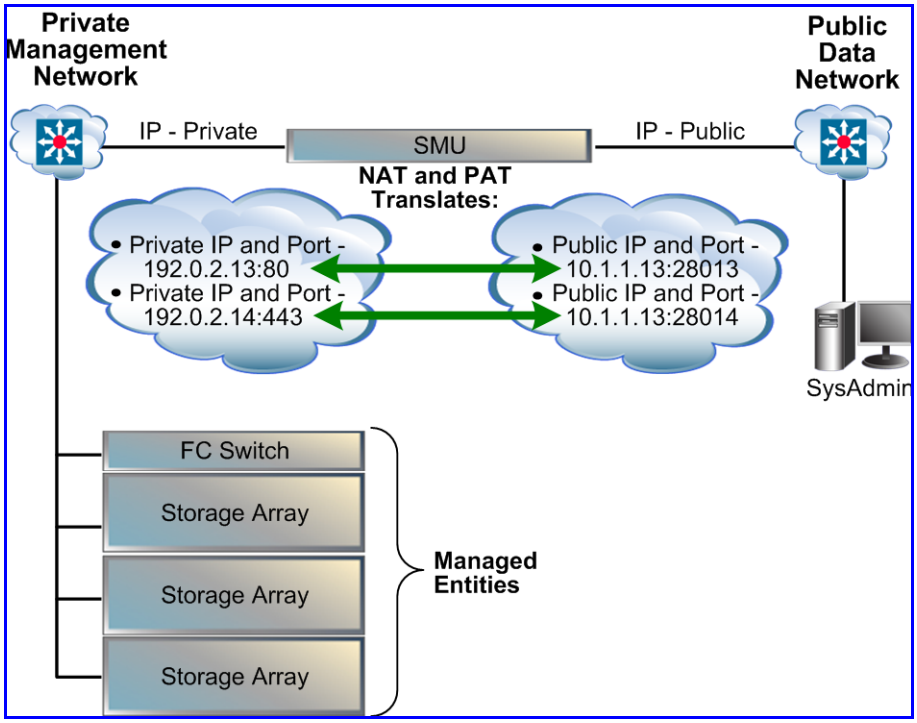

Public Data Network

The public data network, from the storage server perspective, consists of the first Ethernet port on the SMU (the public Ethernet interface). In addition, management access can be enabled on individual Gigabit Ethernet (GE) interfaces on the storage server. The public data network is the access point for getting data into and out of the storage server system.

Storage Subsystem

The storage subsystem contains the devices that store the data managed by the storage server. The server supports tiered storage, which simultaneously connects multiple diverse storage subsystems behind a single server unit (or cluster). Use tiered storage to customize the server to match the storage requirements of your applications, and your system requirements for performance and scaling.

The storage subsystem is made up of RAID controllers, storage devices and the Fibre Channel (FC) infrastructure (such as FC switches and cables) used to connect these devices to a single server or cluster.

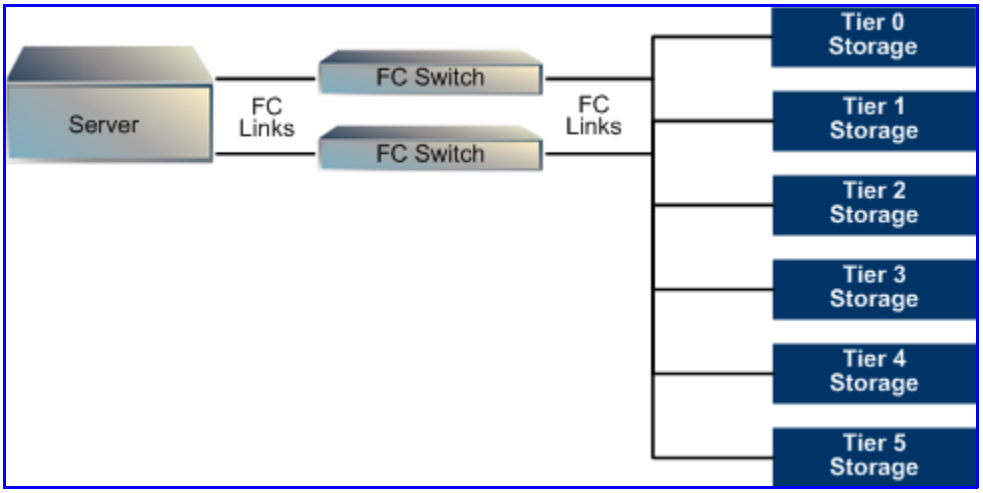

The IS-NAS Server and the Titan Server support tiers of storage, where each tier is made up of devices with different performance characteristics or technologies.

Tiers of storage and storage virtualization are fully supported by Data Migrator, an optional feature which allows you to optimize the usage of tiered storage by automatically migrating data among storage subsystems of primary and secondary storage. Based on user-defined policies, Data Migrator monitors file metadata such as size, type, duration of inactivity, access history, and so on. When the criteria of a policy are met, Data Migrator migrates files according rules specified in the policy as background tasks with minimal impact on server performance. From the perspective of the client workstation, primary versus secondary file location is transparent. Note that Data Migrator does not support migrating data to or from tape library systems.

Default User Name and Password

The following table provides all default system logins:

System Component	Username	Password
SMU Web Manager	admin	`nasadmin`
SMU CLI	manager	`nasadmin`
SMU Entering this specific username and password will provide unlimited access on the SMU.	root	`nasadmin`
Storage server (CLI)	supervisor	`supervisor`

Managing a Server/Cluster

Server/cluster management is performed through Web Manager (a browser-based graphical user interface) or through commands issued using a

command line interface (CLI). Most management functions can be accomplished through Web Manager, but some operations require using the CLI. The following sections describe these management interfaces in detail.

Using Web Manager

The Web Manager administration tool provides a browser-based interface for managing standalone or clustered server(s) and their attached storage subsystem(s). This tool allows you to perform most administrative tasks, from any client on the network using a suitable Web browser (Microsoft Internet Explorer 6.0 or later, or Firefox 1.5 or later).

Note: When accessing the SMU through Web Manager, some browsers may fail with an error message of "Secure Connection Failed" due to an invalid security certificate. This situation may occur the first time you access a new SMU, or when the SMU's IP address or host name have been changed. If you receive this error message, you must add the SMU's security certificate. Refer to the browser's help to find information on how to add a certificate, and make sure to permanently store the exception.

As an alternative to the Web Manager, you can use the Command Line Interface (CLI). For more information, see Using the Command Line Interface, on page 16. The CLI is documented in the *Command Line Reference*, which is available through a link on the Web Manager **Documentation** page.

Server Status Console

The **Home** page in the Web Manager's console displays summary status information pertaining to the *currently managed server*:

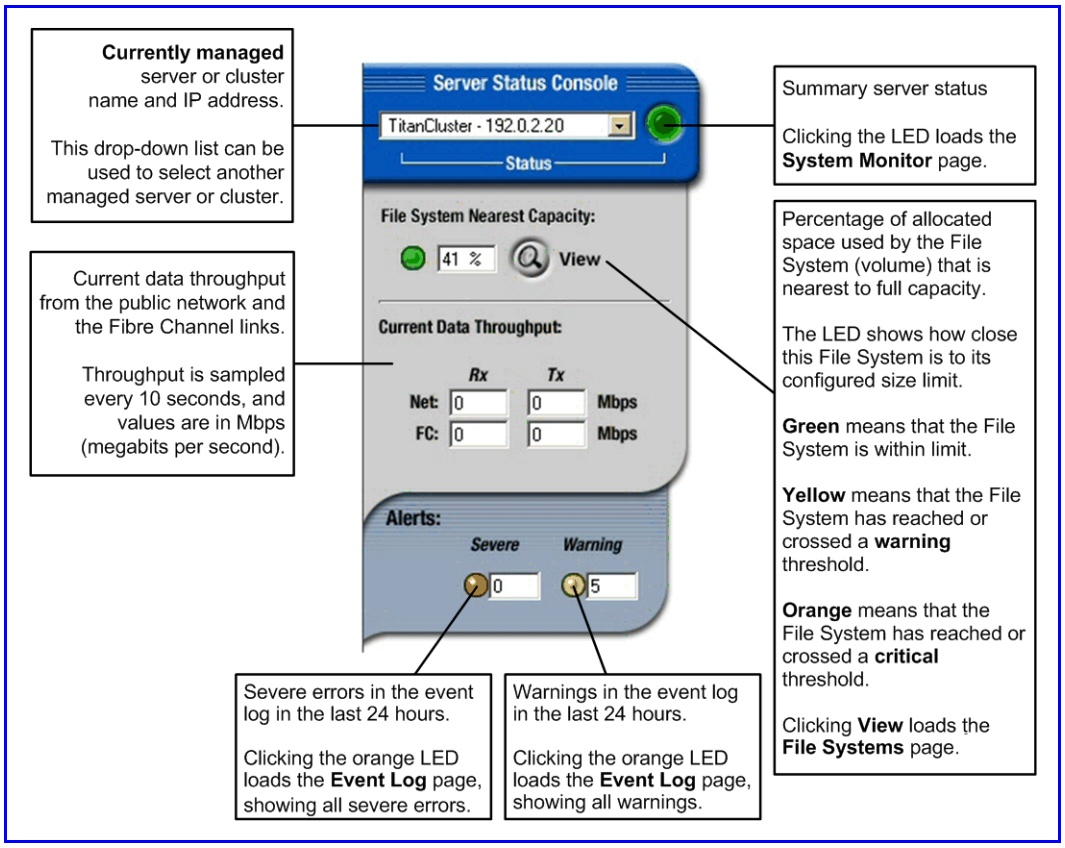

Interpret the color-coded **Status** indicator as follows:

- **Green**. Operating normally (server not showing an alert condition).
- **Amber**. Warning condition (server operating normally, but action should be taken to maintain normal operation).
- **Red**. Critical condition (server no longer functioning or failing in a way that presents a danger to the system).

Accessing Web Manager Pages (Navigation)

The Web Manager uses a two-level page structure, including *links* to specific functions and *categories* that wrap those functions. Clicking on a *link* starts its target function, while clicking on a *category* loads a page.

The top-level page is the Web Manager **Home** page:

This page groups functions associated with the *currently managed server* under page categories, including:

- **Status & Monitoring**: System Monitor, Event Log, Email Alerts Setup, SNMP, Management Access Statistics, etc.
- **Server Settings**: EVS Management, Server Setup Wizard, Cluster Configuration, etc.
- **Storage Management:** File Systems, Virtual Volumes, Quotas, System Drives, Data Migrator, etc.
- **Data Protection**: Virus Scanning, Replication, Snapshots, NDMP backup, etc.
- **File Services**: CNS, NFS Exports, CIFS Shares, iSCSI, FTP, User Mapping, Group Mapping, etc.
- **Network Configuration**: IP Addresses, Name Services, NIS/LDAP Configuration, IP Routes, Link Aggregation, etc.

Additional categories (not associated with the currently managed server):

- **SMU Administration:** Manages the SMU itself (for example, *currently managed server* selection, security, private management network).
- **Documentation:** Links to documentation, including the online help, this manual, and the following documents:
  - *Command Line Reference*
  - *Hardware Reference*
  - *Storage Subsystem Guide*

Web Manager Tables

Some of the pages in the Web Manager interface include tables:

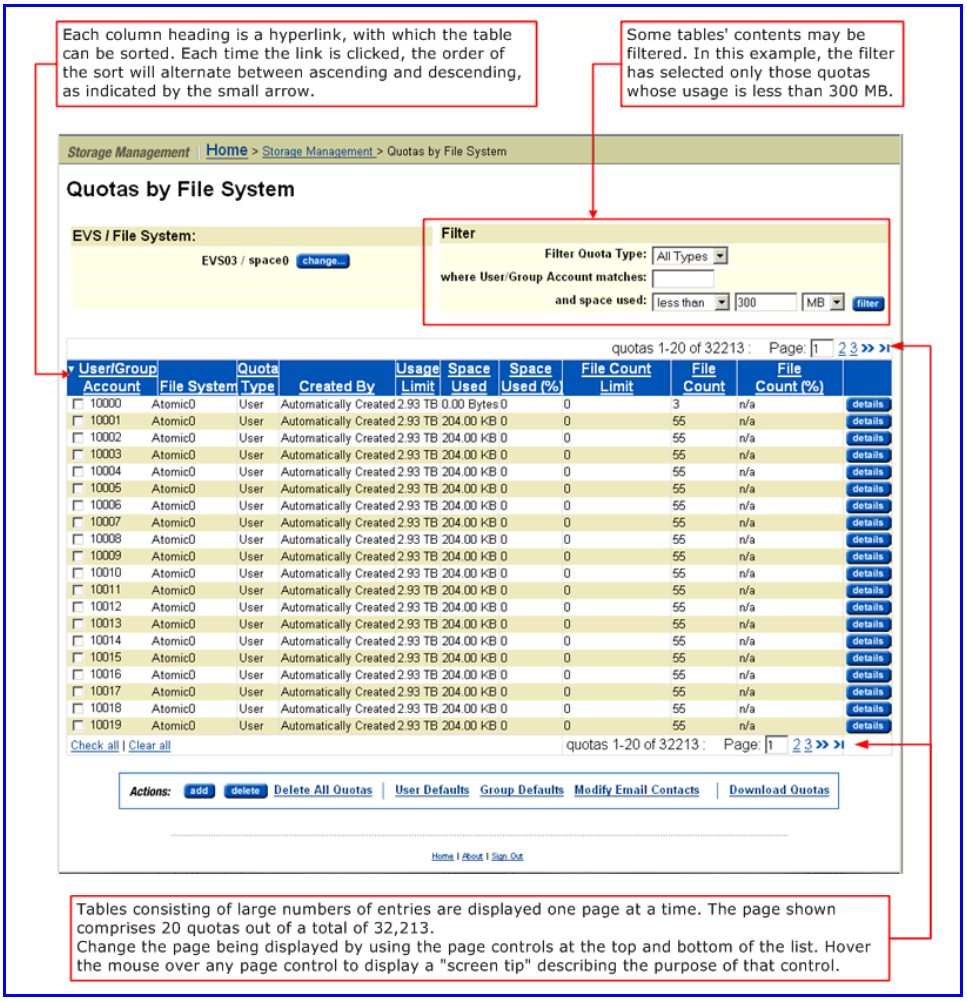

Using the Command Line Interface

The storage server and the System Management Unit (SMU) each come with a command line interface (CLI) for configuration and management. Both support secure connections, configurable passwords, and other security mechanisms.

Note: For more information on the Command Line Interface (CLI), refer to the *Command Line Reference*, which can be accessed through the **Documentation** page of **Web Manager** (see Using Web Manager, on page 13 for more information).

SMU Command Line Interface

Serial console connection

The SMU ships without a pre-configured network setup. To perform the initial setup, access the SMU through a direct serial connection. Once its

network configuration has been completed, access the SMU's CLI directly through SSH or through a Java-enabled SSH session running under Web Manager.

To connect using a serial console

1. **Connect serial cable.**

   Attach an RS232 null-modem cable (DB-9 Female to DB-9 Female) to the serial port on the SMU back panel. Attach the other end of the serial cable to a terminal (e.g. laptop).

2. **Configure terminal emulation.**

   Set the terminal emulation program (such as Windows HyperTerminal) to the following settings: *115,200 b/s, 8 bits/byte, 1 stop bit, No parity.*

3. **Log into the SMU.**

   - *If the SMU is being accessed to perform initial setup*, log in as user *"setup"*. When prompted, perform the configuration steps as directed.
   - *Otherwise*, login as the user *"manager"*. When prompted, enter the password for the user *"manager*".

4. **Once connected, launch the storage server CLI or select SMU shell.**

   From the SMU command line, access the server CLI using one of the methods in the displayed menu, or enter `"q"` to access the SMU shell.

SSH Connection

The SMU can be accessed using any SSH client. Note that the client should be configured to support the UTF-8 (Unicode) character set.

For Windows and others without an SSH client, access the SMU via *SSHTerm.* SSHTerm is a Java SSH client (applet) developed by 3SP and distributed under the General Public License (GPL). SSHTerm provides a convenient, cross-platform alternative to other SSH clients.

To connect using SSHTerm:

1. **Verify Java version and status.**

   SSHTerm requires Java 1.4.x or greater. Go to `http://www.java.com` to verify that your browser has the latest version of the Java Plug-in installed and enabled.

2. **Navigate to SSHTerm.**

   From the **Home** page, click **SMU Administration**; then select **SSHTerm**.

3. **Launch SSHTerm.**

   Click **Launch SSHTerm**. When the SSH client applet window pops up, accept the certificate registered to 3SP LTD, and click **Always** or **Yes** when asked to allow the host.

   SSHTerm automatically connects to the SMU as user *"manager"*.

4. **When prompted, enter the password for the user "*manager*".**
5. **Repeat as desired.**

   Multiple SSHTerm windows may be used at once. Just click **Launch SSHTerm** for each new SSH session.
6. **When finished, terminate the session.**

   When the SSH session has finished, just close the window.

Server Command Line Interface

The storage server ships with a comprehensive command line interface (CLI), documented separately in the *Command Line Reference.*

You can access the IS-NAS Server/cluster CLI through the Server Control (SSC) utility, available for Windows and Linux.

You can access the Titan Server/cluster CLI in the following ways:

- Secure Shell (SSH) connection into the server through the SMU.
- SSH or Telnet connection directly into the server.
- Server Control (SSC) utility, available for Windows and Linux.
- Perl Server Control (PSSC) utility, available for all operating systems with PERL support.

Note: To access the storage server CLI directly, using SSH, Telnet, SSC, or PSSC through the public network, a server administration IP address may be assigned to at least one of the Gigabit Ethernet interfaces. The server supports access to its CLI through any administrative IP address; by default, an administrative IP address is available on the private management network.

Note: When using the server's CLI, the console device should be configured to use UTF-8 (Unicode) encoding. When using the CLI through SSH, Telnet, SSC, or PSSC, the server sends and expects to receive data using UTF-8 encoding, so the device sending data to or receiving data from the server should be configured to use UTF-8 encoding.

SSH via the SMU

The SMU supports SSH. After logging into the SMU, the SMU can redirect connections directly to the server's CLI. This can be useful for two reasons:

1. It eliminates the need to assign a server administration IP address to the Gigabit interface of the server.
2. When prompted, enter the user name (*supervisor).*
3. It enhances the security of server by isolating administrative access to the private management network.

To SSH into the server, using the SMU as a proxy:

1. Connect to the SMU through SSH.

2. Log into the SMU as "*manager*".

   A list of servers will appear. Select the target server. The SMU automatically initiates a connection to the server's CLI.

SSH or Telnet

When connecting to a Titan Server/cluster through SSH or Telnet through the administrative services EVS IP address, log in using the user name *"supervisor"*. IS-NAS Servers/clusters do not support SSH or Telnet connections. Note that the administrator must have configured the server to accept SSH or Telnet connections, and the SSH/Telnet client should be configured to support the UTF-8 (Unicode) character set.

To SSH into the server:

1. Connect to the server's administrative services DNS name or IP address:
   `ssh supervisor@server_name_or_IP`

2. When prompted, enter the "*supervisor*" user's password.

To Telnet into the server:

1. Connect to the server's administrative name or IP address:
   `telnet server_name_or_IP`

2. When prompted, enter the *"supervisor"* user's password.

SSC

The server can connect to the SMU from Windows PCs and from Linux/Unix workstations via SSC, which provides a secure connection using a modified version of the Arcfour cipher for encryption and Sha-1 for authentication.

It comes in two varieties:

- SSC for Windows and Linux.
- PSSC, a Perl scripted version of SSC for Linux/Unix operating systems.

Use the SSC scripting utility to access the server's CLI. The server supports SSC access to its CLI through any administrative IP address. By default, the private management network has an administrative IP address available.

The syntax for SSC:

```
ssc [-u <username>] [-p <password>]
<host>[:<port>] [<command>]
```

The syntax for PSSC:

```
pssc [-u <username>] [-p <password>] <host>[:<port>]
[<command>]
```

The following table defines the variables:

Syntax	Description
Username	User account (typically *"supervisor"*).
Password	Password. If none specified, SSC/PSSC will prompt for one.
Host	The server's administration IP address or host name.
Port	If the SSC/PSSC port number has been changed from its default of 206, the port number configured for SSC must be specified in the command syntax.
Command	Command to execute. If no command is specified, SSC/PSSC allows interactive command entry.

3 Quick System Configuration

System installation and initialization tasks are usually performed by SGI Global Services. Following system initialization, basic configuration is required to ready the system for use. Once the system is operational, optional configuration changes can be made at any time, either by field service personnel or by system administrators.

This section explains how to use wizards to complete the basic configuration of the SMU and the server. See Optional Configuration Steps, on page 29 for information about optional configuration tasks and procedures.

Using the SMU Setup Wizard

For an external SMU, basic SMU configuration is usually performed as a part of system initialization. The SMU setup wizard is used to complete the basic configuration of an SMU. Using the SMU Setup Wizard, you can change access permissions to the SMU, set up name services for network operation, and configure the date and time.

To use the SMU setup wizard:

1. **Navigate to the SMU Setup Wizard page.**

   From the **SMU Administration** page, display the wizard by clicking **SMU Setup Wizard**:

   As the wizard progresses, enter requested information according to the following guidelines:

Item/Field	Description
Passwords	This page allows you to specify the default password for the SMU when the admin user accesses the SMU using HTTP (Web Manager). By default, the password is "`nasadmin`." You should change the default password as soon as possible.
DNS Server Setup	This page allows you to set IP addresses of the DNS servers and the domain search order that will be applied to the SMU.
SMTP Relay	This page allows you to specify the host name (not the IP address) of the email server to which the SMU can send and relay event notification emails.
Date & Time	This page allows you to set the clock on the SMU and select one or more NTP servers.
Private management network	This page allows you to configure the IP address of the SMU's *eth1* interface on the private management network.

2. **Save your changes.**

   Upon completion, the wizard displays a page summarizing parameters entered. To accept, click **finish**. Then click **OK** to reboot.

Selecting SMU-Managed Servers

The SMU manages multiple storage servers/clusters and their associated storage subsystems. Use the **Managed Servers** page to add information about each server; specifically, the IP Address and username/password of the server to be managed. Only one server, the *currently managed server,* may be managed at one time. From the Managed Servers list, any server can be selected as the *currently managed server*.

To display the servers managed by the SMU:

1. **Navigate to Managed Servers.**

   From the **SMU Administration** page, display the configuration page by clicking **Managed Servers**:

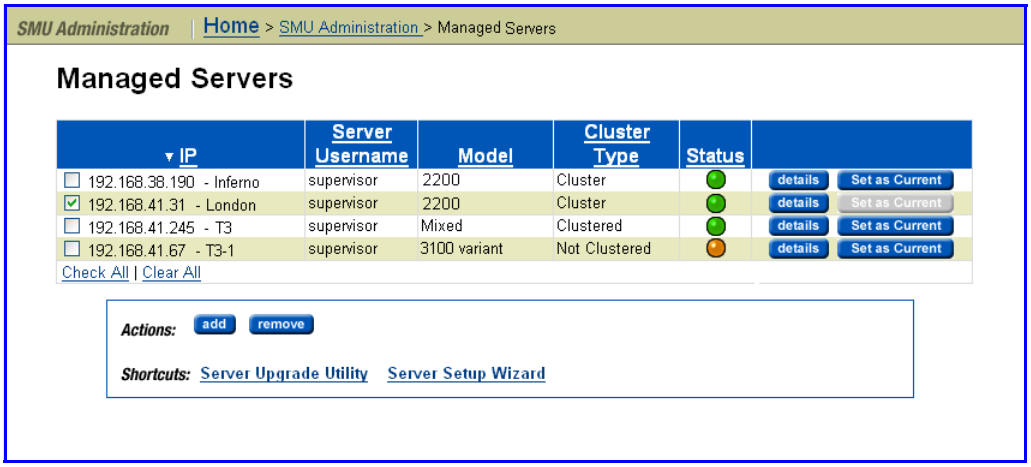

Interpret the color-coded **Status** indicator as follows:

- **Green**. Operating normally (server not showing an alert condition).
- **Amber**. Warning condition (server operating normally, but action should be taken to maintain normal operation).
- **Red**. Critical condition (server no longer functioning).

2. **Select an action.**

   The following **Actions** are available:

   - Click **Set as Current** for any server to make it the *currently managed server.* (Alternatively, you can use the drop-down box in the **Server Status** console on the **Home** page.)
   - Click **add** to add a new server to the **Managed Servers** list, then refer to the next step in this procedure.
   - Click **remove** to remove the selected servers from the **Managed Servers** list. Make a selection by filling a server's checkbox, or click *check all* to remove all servers from the **Managed Servers** list.

3. **Add a server (optional).**

   In the previous step, you clicked **add.** When the **Add Server** page displays,

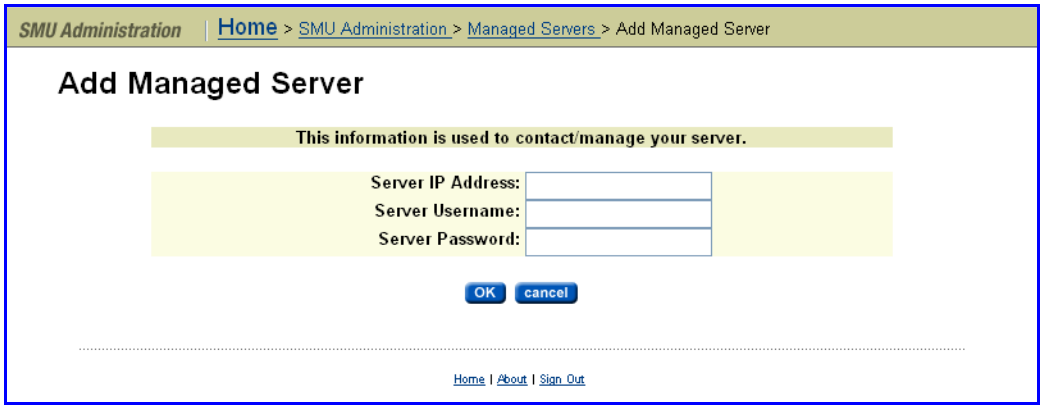

complete the fields, then click **OK**.

When the SMU adds a managed server, the following actions occur:

- If the server is managed through the private management network, the SMU's *eth1* IP address is added to the server's list of NTP servers.
- If the server is managed through the private management network, the SMU's *eth1* IP address is configured as the server's Primary SMTP server. If the server was already configured to use a mail server, this server will automatically become the Backup SMTP server.
- A user name and password are preserved on the SMU so that, when using Web Manager, you can select this server as the current managed server without causing the server to prompt for additional authentication.

Using the Server Setup Wizard

During the initial setup of the server performed using the SMU Setup Wizard (see Using the SMU Setup Wizard, on page 21) a number of configuration settings, such as system name and date/time were specified. You can change these settings using the procedures in the following sections. License keys also must be installed to enable the protocols and services purchased with the servers (see Managing License Keys, on page 535).

This wizard creates a basic server configuration, using user-defined values. At the end of the **Server Setup Wizard,** a confirmation dialog appears, allowing review of settings.

Note: An IP address must be assigned to the server before the setup wizard can be used. In addition, the server must be configured as a Managed Server. For more information, see Using the SMU Setup Wizard, on page 21 and Selecting SMU-Managed Servers, on page 22.

1. **Navigate to the Server Setup Wizard page.**

   From the **Server Settings** page, click **Server Setup** to display the **Clone Server Settings Wizard** page. This page allows you to copy certain

information from the SMU or another managed server, so you don't have to repeatedly enter the same configuration information:

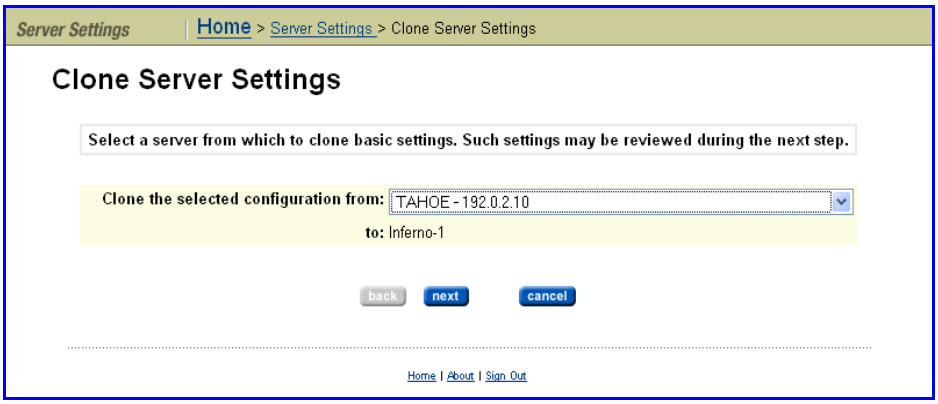

2. **Select a source for cloning configuration settings:**

   If this is the first server to be configured, the wizard can clone some settings from the SMU to the new server. Note that the settings that can be cloned from an SMU are a subset of the settings that can be cloned from another server. To clone settings from the SMU, select **SMU** from the drop-down menu, then click **next** to display the **Clone Server Settings** page.

   If the SMU is already managing another server, an expanded list of settings can be cloned from another server. To clone settings from another server select one of the managed servers from the drop-down menu, then click **next** to display the **Clone Server Settings** page, which allows you to clone more settings than are available from the SMU.

3. **Select configuration items to clone.**

   In the **Clone Server Settings** page, fill or clear the checkbox next to each of the configuration items you want to clone, then click **next.** Cloned settings are immediately applied to the server.

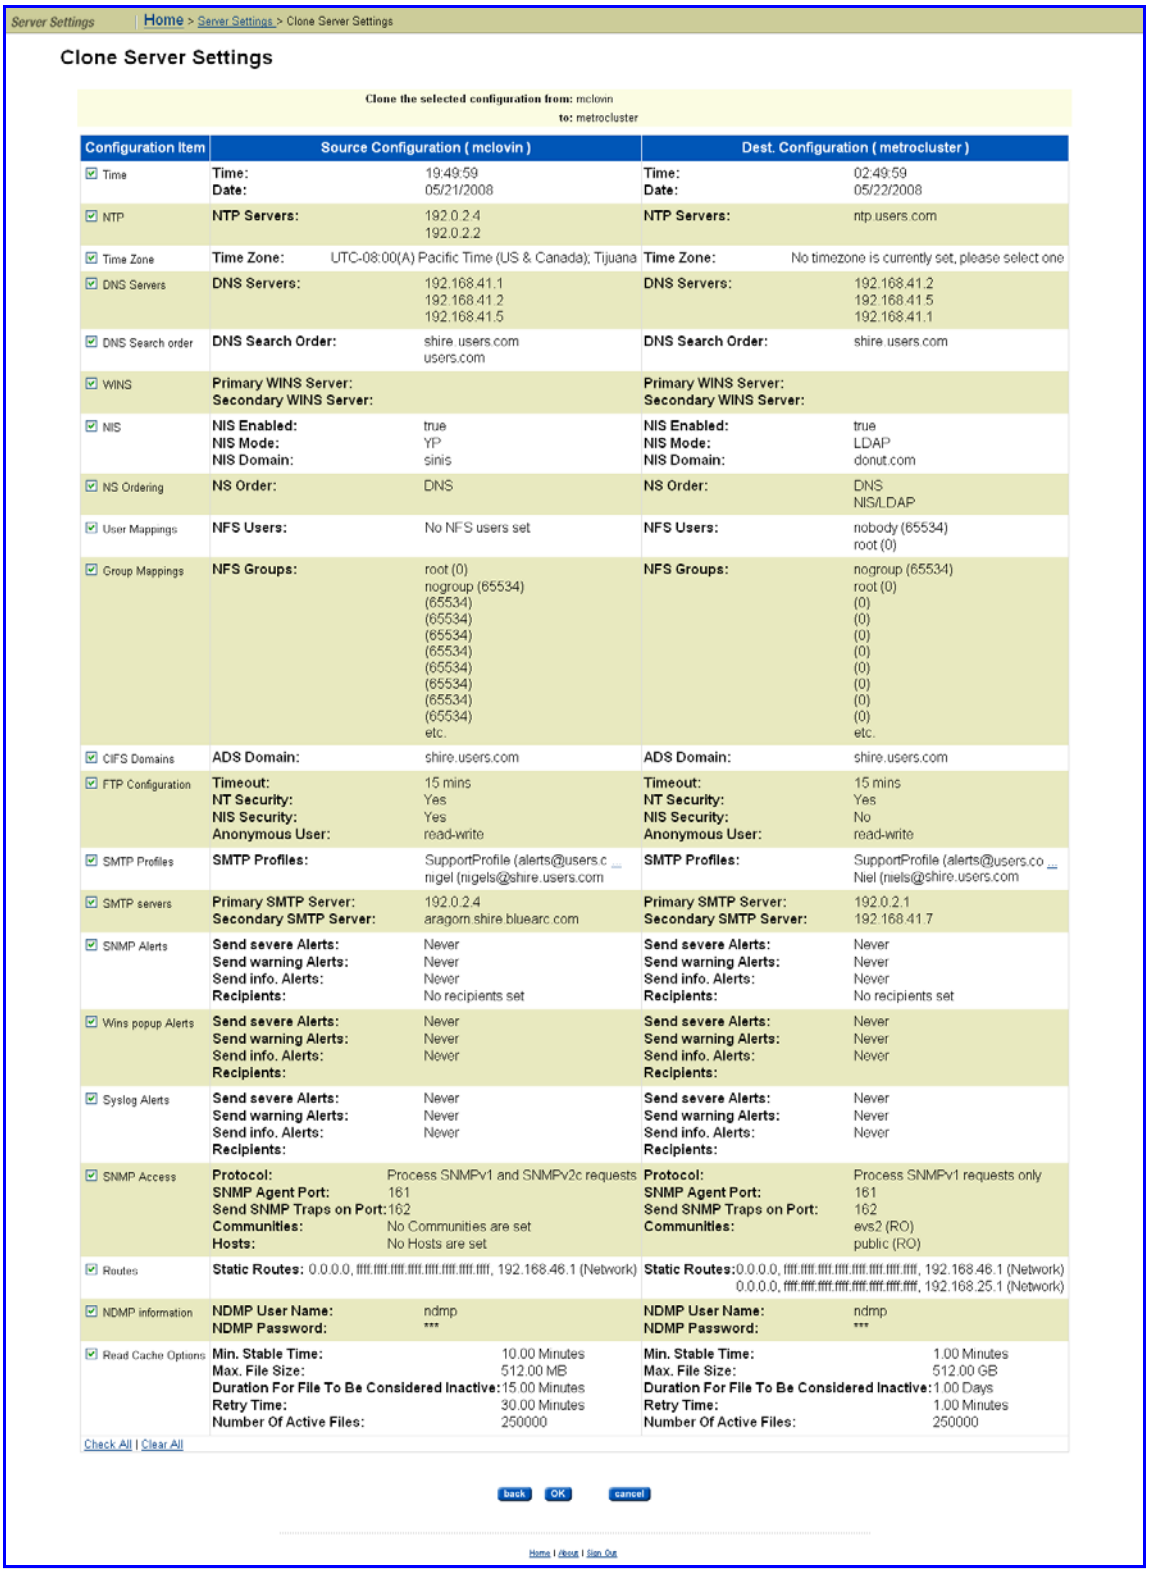

Server Settings | Home > Server Settings > Clone Server Settings

Clone Server Settings

Clone the selected configuration from: mclovin
to: metrocluster

Configuration Item	Source Configuration (mclovin)	Dest. Configuration (metrocluster)
Time	**Time:** 19:49:59 **Date:** 05/21/2008	**Time:** 02:49:59 **Date:** 05/22/2008
NTP	**NTP Servers:** 192.0.2.4 192.0.2.2	**NTP Servers:** ntp.users.com
Time Zone	**Time Zone:** UTC-08:00(A) Pacific Time (US & Canada); Tijuana	**Time Zone:** No timezone is currently set, please select one
DNS Servers	**DNS Servers:** 192.168.41.1 192.168.41.2 192.168.41.5	**DNS Servers:** 192.168.41.2 192.168.41.5 192.168.41.1
DNS Search order	**DNS Search Order:** shire.users.com users.com	**DNS Search Order:** shire.users.com
WINS	**Primary WINS Server:** **Secondary WINS Server:**	**Primary WINS Server:** **Secondary WINS Server:**
NIS	**NIS Enabled:** true **NIS Mode:** YP **NIS Domain:** sinis	**NIS Enabled:** true **NIS Mode:** LDAP **NIS Domain:** donut.com
NS Ordering	**NS Order:** DNS	**NS Order:** DNS NIS/LDAP
User Mappings	**NFS Users:** No NFS users set	**NFS Users:** nobody (65534) root (0)
Group Mappings	**NFS Groups:** root (0) nogroup (65534) (65534) (65534) (65534) (65534) (65534) (65534) (65534) (65534) etc.	**NFS Groups:** nogroup (65534) root (0) (0) (0) (0) (0) (0) (0) (0) (0) etc.
CIFS Domains	**ADS Domain:** shire.users.com	**ADS Domain:** shire.users.com
FTP Configuration	**Timeout:** 15 mins **NT Security:** Yes **NIS Security:** Yes **Anonymous User:** read-write	**Timeout:** 15 mins **NT Security:** Yes **NIS Security:** No **Anonymous User:** read-write
SMTP Profiles	**SMTP Profiles:** SupportProfile (alerts@users.c ... nigel (nigels@shire.users.com	**SMTP Profiles:** SupportProfile (alerts@users.co ... Niel (niels@shire.users.com
SMTP servers	**Primary SMTP Server:** 192.0.2.4 **Secondary SMTP Server:** aragorn.shire.bluearc.com	**Primary SMTP Server:** 192.0.2.1 **Secondary SMTP Server:** 192.168.41.7
SNMP Alerts	**Send severe Alerts:** Never **Send warning Alerts:** Never **Send info. Alerts:** Never **Recipients:** No recipients set	**Send severe Alerts:** Never **Send warning Alerts:** Never **Send info. Alerts:** Never **Recipients:** No recipients set
Wins popup Alerts	**Send severe Alerts:** Never **Send warning Alerts:** Never **Send info. Alerts:** Never **Recipients:**	**Send severe Alerts:** Never **Send warning Alerts:** Never **Send info. Alerts:** Never **Recipients:**
Syslog Alerts	**Send severe Alerts:** Never **Send warning Alerts:** Never **Send info. Alerts:** Never **Recipients:**	**Send severe Alerts:** Never **Send warning Alerts:** Never **Send info. Alerts:** Never **Recipients:**
SNMP Access	**Protocol:** Process SNMPv1 and SNMPv2c requests **SNMP Agent Port:** 161 **Send SNMP Traps on Port:** 162 **Communities:** No Communities are set **Hosts:** No Hosts are set	**Protocol:** Process SNMPv1 requests only **SNMP Agent Port:** 161 **Send SNMP Traps on Port:** 162 **Communities:** evs2 (RO) public (RO)
Routes	**Static Routes:** 0.0.0.0, ffff.ffff.ffff.ffff.ffff.ffff.ffff.ffff, 192.168.46.1 (Network)	**Static Routes:** 0.0.0.0, ffff.ffff.ffff.ffff.ffff.ffff.ffff.ffff, 192.168.46.1 (Network) 0.0.0.0, ffff.ffff.ffff.ffff.ffff.ffff.ffff.ffff, 192.168.25.1 (Network)
NDMP information	**NDMP User Name:** ndmp **NDMP Password:** ***	**NDMP User Name:** ndmp **NDMP Password:** ***
Read Cache Options	**Min. Stable Time:** 10.00 Minutes **Max. File Size:** 512.00 MB **Duration For File To Be Considered Inactive:** 15.00 Minutes **Retry Time:** 30.00 Minutes **Number Of Active Files:** 250000	**Min. Stable Time:** 1.00 Minutes **Max. File Size:** 512.00 GB **Duration For File To Be Considered Inactive:** 1.00 Days **Retry Time:** 1.00 Minutes **Number Of Active Files:** 250000

Check All | Clear All

back OK cancel

Home | About | Sign Out

Tip: Empty the checkboxes of those items you do not want to clone.

4. **Complete the wizard.**

   Refer to the following table for descriptions of fields encountered in the wizard:

Page	Description
Password	A password for the *"supervisor"* account on the server.
Server Identification	System name and other identifying information for use by SNMP and SMTP (email) and other protocols.
Name Services	Configure the server for one or more name services, such as DNS, WINS, and NIS.
SMTP	**Caution:** *Recommended profile alert!* Before configuring an SMTP profile, note that the wizard recommends a default profile (that alerts SGI Support). SGI strongly recommends creating this profile, so that SGI can respond quickly should a failure occur. Additional email profiles for administrative notification of failures can be set up once the Wizard is complete. This page configures primary and secondary SMTP servers to be used for sending Email Alerts.
Date & Time	Set the server's clock. Synchronize with one or more NTP servers. The SMU usually synchronizes its clock with an NTP server on the public data network, and it acts as an NTP server for devices on the private management network. Since a server typically resides on the private management network, add the SMU to the server's list of NTP servers.

5. **Reboot or shut down the server.**

   Once you have completed the wizard, you can either reboot the server or shut it down. When the server is restarted, it will use the new configuration.

4 Optional Configuration Steps

Optional configuration tasks are usually performed by a system administrator, either to change settings specified as a part of basic configuration (see Quick System Configuration, on page 21) or to set up additional system capabilities, such as managing administrative users, securing management interface access, setting up time synchronization, etc. The following sections explain how to perform these optional configuration steps.

Managing Users and Roles

The following sections describe the process of creating and editing global, system, and storage administrators:

- **Global Administrators** create Server Administrators, Storage Administrators, and other Global Administrators. They also control what servers an administrator can access. A Global Administrator can also change the name and/or administrator type and/or role of any user.
- **Server Administrators** have rights and privileges that allow management of servers, as specified in the administrator profile created by the Global Administrator. Additionally, Server Administrators may be able to manage storage, if the Global Administrator has given them that permission.

  A Server Administrator can manage file systems and file services such as CIFS Shares, NFS Exports, and they can manage file system related features and functions such as snapshots, quotas, and migration policies and schedules. If the Server Administrator can also manage storage, they can manage racks, physical disks, system drives (SDs), and Storage Pools. Server Administrators cannot manage users.
- **Storage Administrators** have rights and privileges that allow management of storage devices, as specified in the administrator profile created by the Global Administrator.

  Storage Administrators can manage only storage devices and their components (racks, physical disks, SDs, and Storage Pools). Storage Administrators cannot manage file systems, file services, or file system related features and functions, and they cannot manage users.

Using Advanced Mode Functions

Based on Advanced Mode settings, links to advanced configuration options and pages can be visible or hidden. When Advanced Mode is off, links to advanced configuration pages are not visible. To view these links, which are typically found on the **Web Manager Home** page, turn on Advanced Mode for the desired SMU user.

Caution: Advanced Mode functions can potentially degrade system performance and cause disruption to the existing services. Advanced Mode functions should only be used after consulting SGI Global Services.

Adding an SMU User - Global Administrator

To add an SMU User - Global Administrator:

1. **Navigate to the SMU Users page.**

   From the **Home** page, click **SMU Administration**, then click **SMU Users** to display the **SMU Users** page:

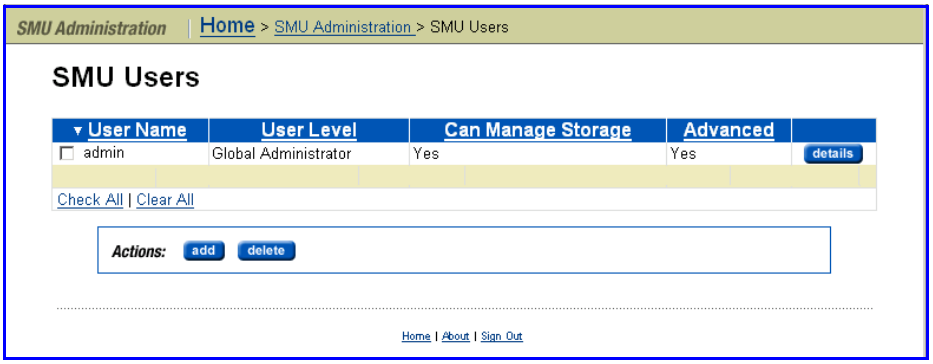

2. **Add a Global Administrator.**

   Click **add** to display the **Add SMU Users** page:

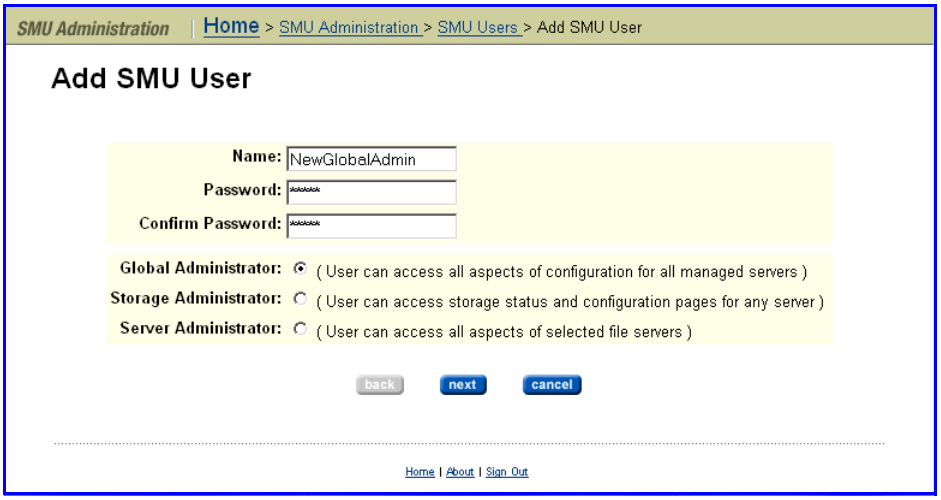

3. **Enter the requested information.**

   Enter identifying user information, select **Global Administrator**, then click **next** to display the Global Administrator version of the **SMU User Details** page:

4. **Specify whether the SMU User will have access to Advanced Mode.**

Caution: Advanced Mode functions can potentially degrade system performance and cause disruption to the existing services. Advanced Mode functions should only be used after consulting SGI Global Services.

Fill the **Advanced Mode** checkbox to enable advanced functions for this user (or not), then click **next** to display the new Global Administrator's profile:

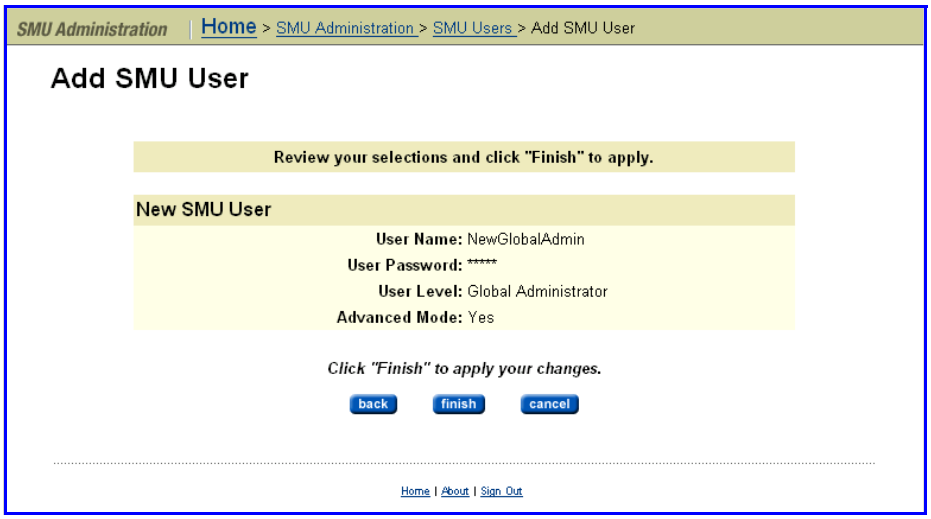

5. **Save the new profile.**

   Verify that the new profile is correct, then click **finish** to display the **SMU Users** page with the newly created Global Administrator listed.

Adding an SMU User - Storage Administrator

To add an SMU User - Storage Administrator:

1. **Navigate to the SMU Users page.**

   From the **Home** page, click **SMU Administration**, then click **SMU Users** to display the **SMU Users** page:

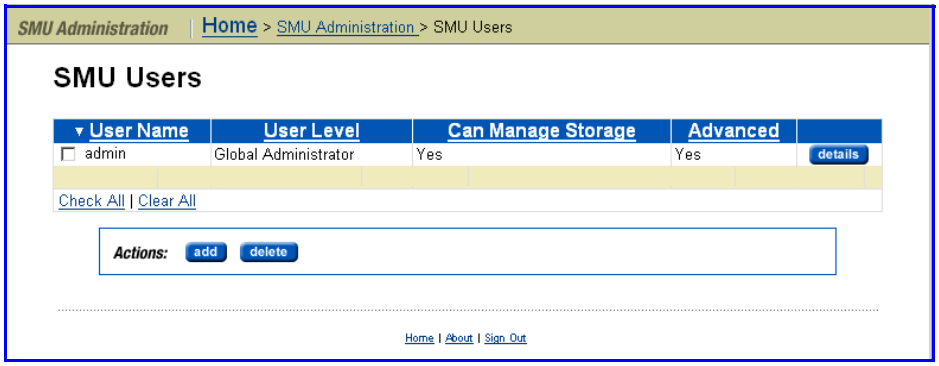

2. **Add a Storage Administrator.**

   Click **add** to display the **Add SMU User** page:

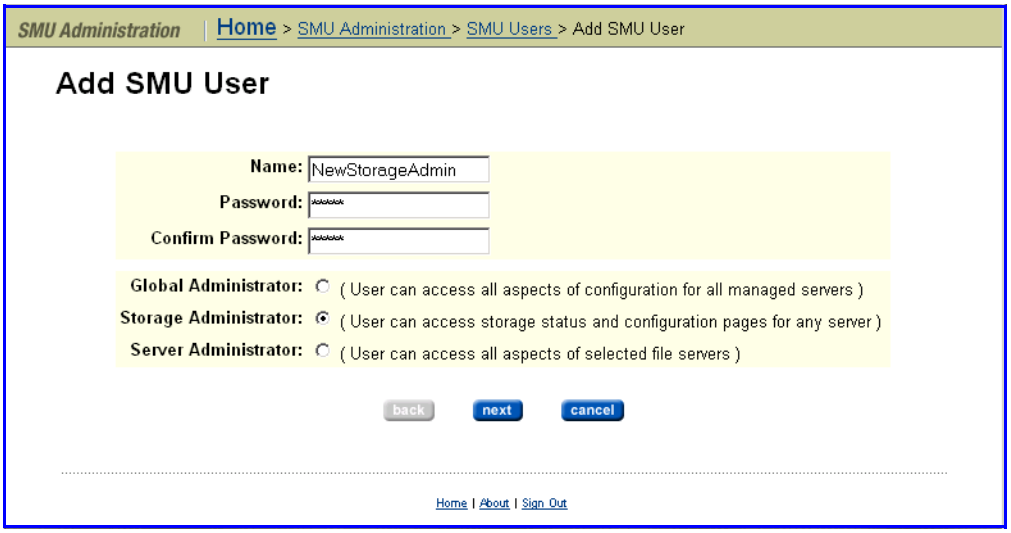

3. **Enter the requested information.**

   Enter identifying user information, select **Storage Administrator**, then click **next** to display the Storage Administrator version of the **Add SMU User** page:

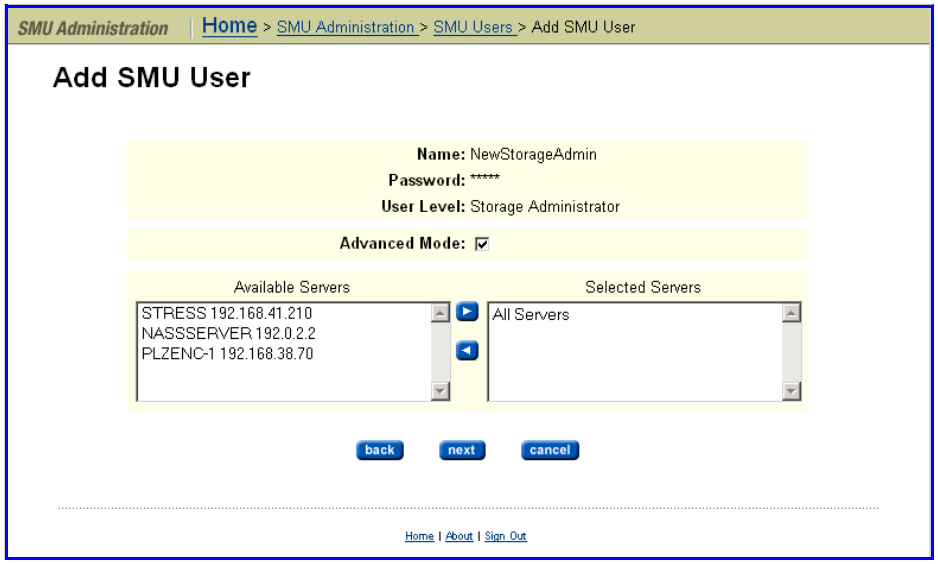

4. **Specify whether the SMU User will have access to Advanced Mode.**

Caution: Advanced Mode functions can potentially degrade system performance and cause disruption to the existing services. Advanced Mode functions should only be used after consulting SGI Global Services.

Fill the **Advanced Mode** checkbox to enable advanced functions for this user (or not), then click **next** to display the SMU user's modified profile:

5. **Specify server access.**

Highlight the servers that this SMU User has rights and privileges to manage from the **Available Servers** list and move them to the **Selected**

Servers list, then click **next** to display the new Storage Administrator's profile:

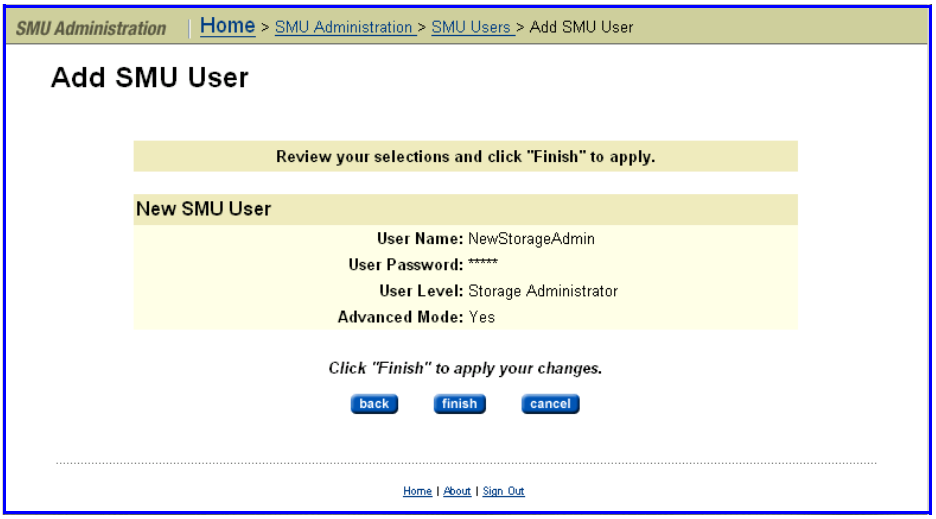

6. **Save the profile.**

   Verify that the new profile is correct, then click **finish** to display the **SMU Users** page with the newly created Storage Administrator listed.

Adding an SMU User - Server Administrator

To add an SMU User - Server Administrator:

1. **Navigate to the SMU Users page.**

   From the **Home** page, select **SMU Administration**, then click **SMU Users** to display the **SMU Users** page:

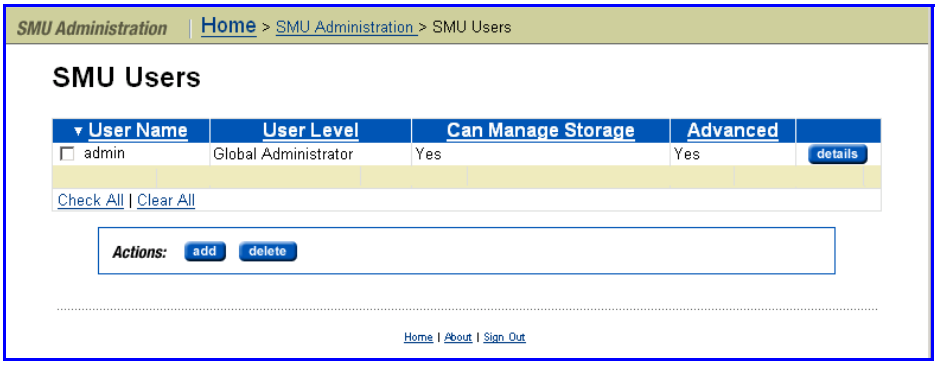

2. **Add a Server Administrator.**

   Click **add** to display the **Add SMU User** page:

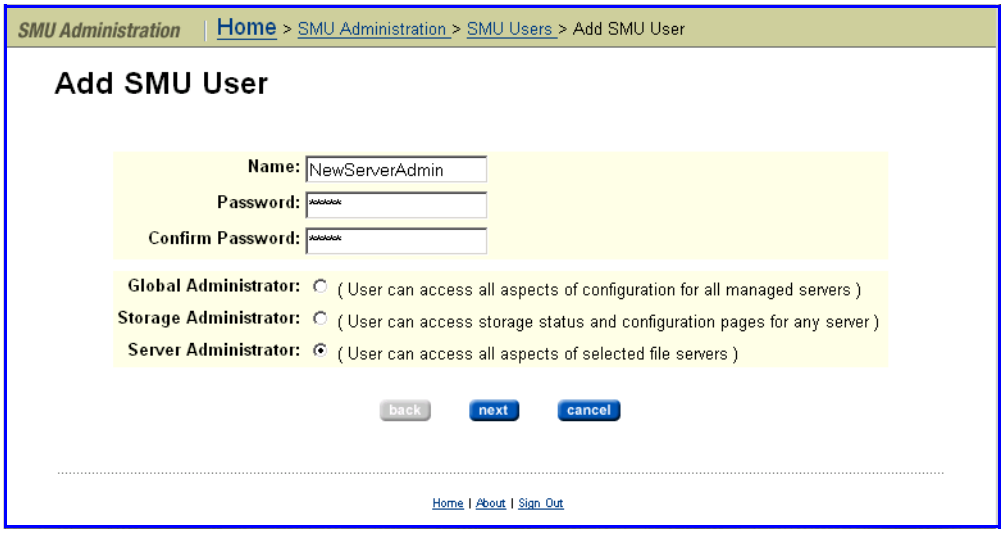

3. **Enter the requested information.**

   Enter identifying user information, select **Server Administrator**, then click **next** to display the Server Administrator version of the **Add SMU User** page:

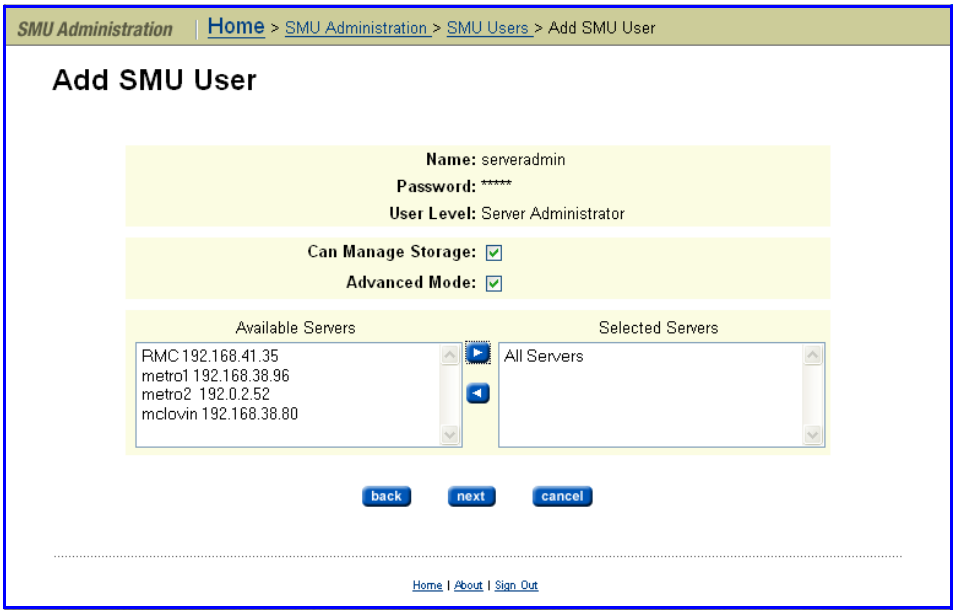

4. **Specify whether the SMU User can manage storage, and/or will have access to Advanced Mode.**

   Caution: Advanced Mode functions can potentially degrade system performance and cause disruption to the existing services. Advanced Mode functions should only be used after consulting SGI Global Services.

   Fill the **Advanced Mode** checkbox to enable advanced functions for this user (or not), then click **next** to display the new Server Administrator's modified profile:

5. **Specify server access.**

   Highlight the servers that this SMU User has rights and privileges to manage from the **Available Servers** list and move them to the **Selected Servers** list, then click **next** to display the user's modified profile:

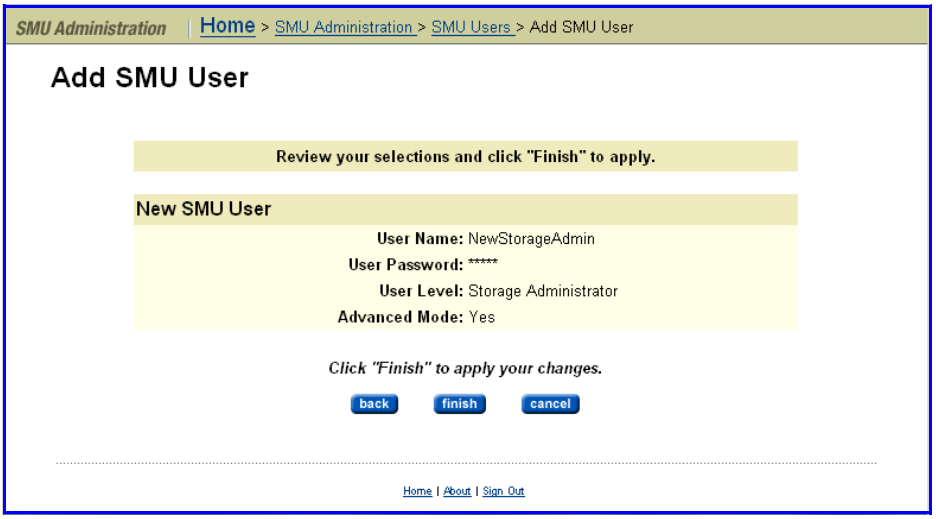

6. **Save the profile.**

   Verify that the new profile is correct, then click **finish** to display the **SMU Users** page with the newly created Server Administrator listed.

Editing an SMU User's Profile

To edit an SMU user's profile:

1. **Navigate to the SMU User Details page.**

   From the **Home** page, click **SMU Administration**, then click **SMU Users** to open the **SMU Users** page:

2. **Select the user profile you want to modify.**

   Click the **details** button to display the **SMU User Details** page for the user whose profile you want to modify.

   The **SMU User Details** page looks different for each type of administrative user:

- For a Global Administrator, the **SMU User Details** page looks like the following:

- For a Server Administrator, the **SMU User Details** page looks like the following:

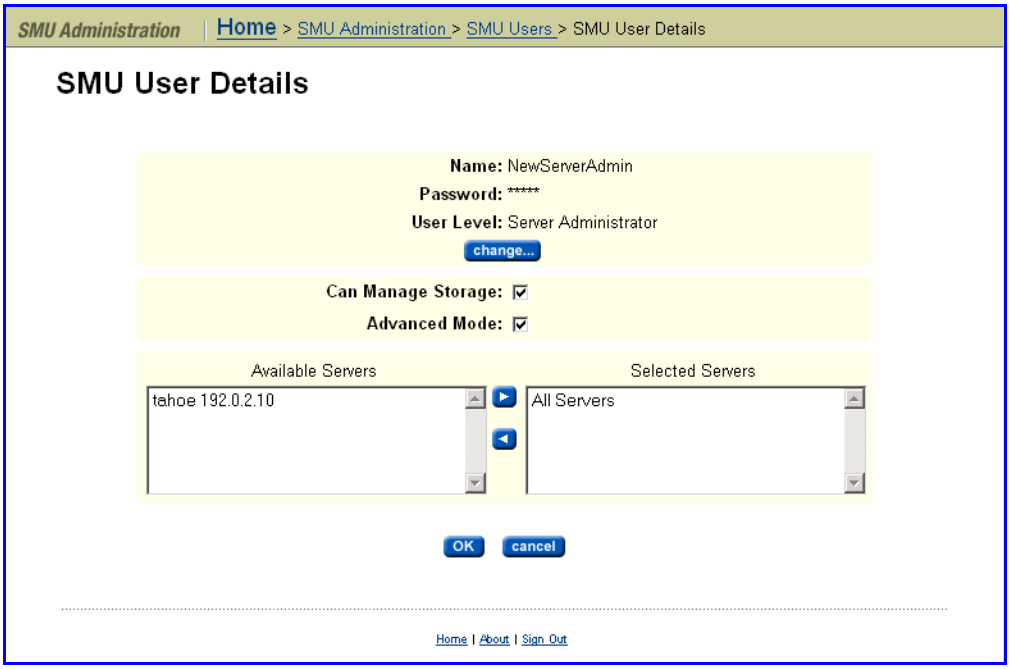

- For a Storage Administrator, the **SMU User Details** page looks like the following:

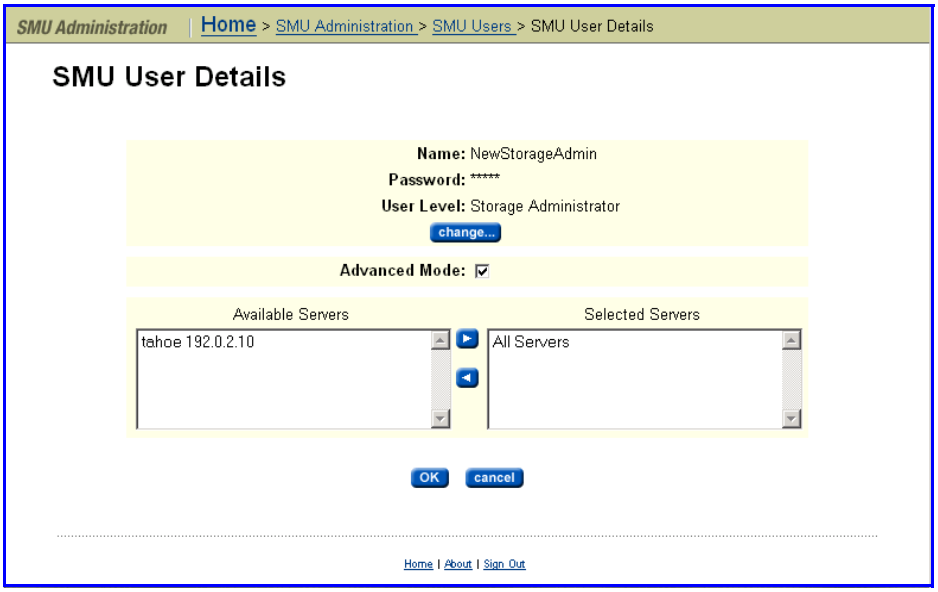

3. **Edit the SMU user information.**

   To edit the user's role or password, click **change.**

4. **Edit the SMU user password or role.**

   When the **SMU User Details** page appears, update the SMU user's password or role, then click **OK** to return to the **SMU User Details** page.

5. **Specify server and/or storage subsystem management rights.**

   Depending on the administrative role of the SMU user, the system displays a **Details** page for a Global, Server, or Storage Administrator. The possible fields displayed in this page are described in the following table:

Item/Field	Description
User Name	Administrator's user name.
User Level	Displays the user level or type of administrative role.

Item/Field	Description
Can Manage Storage (Displayed for Server Administrators only)	Displays if the Server Administrator can manage storage subsystems attached to the server(s) managed by this Server Administrator.
Advanced Mode	Displays if Advanced functions are available to this user.
Available Servers	For Server and Storage Administrators only, lists servers managed by the SMU: • For a Server Administrator, the servers in this list are those for which the Server Administrator does not have management rights. • For a Storage Administrator, the servers in this list have attached storage subsystems for which the Storage Administrator does not yet have management rights.
Selected Servers	For Server and Storage Administrators only, lists servers managed by the SMU and, once the profile is saved: • Storage Administrators will be able to manage storage subsystems attached to the listed servers, but they will not be able to manage the servers themselves. • Server Administrators will be able to manage the servers, and if the **Can Manage Storage** checkbox is filled, they will also be able to manage the storage subsystems attached to those servers.

Note: To move a server between the **Available Servers** list and the **Selected Servers** list, select the server, and use the arrow buttons between the lists.

When you are done making changes, click **OK** to save the profile and return to the **SMU Users** page.

Changing the Password for the Currently Logged in User

This alternate method for changing the password is provided because only Global Administrators can access the **SMU User** page.

To change the password for the currently logged in user:

1. **Navigate to the Current User Password page.**

   From the **SMU Administration** page, click **Current User Password** to display the **Current User Password** page:

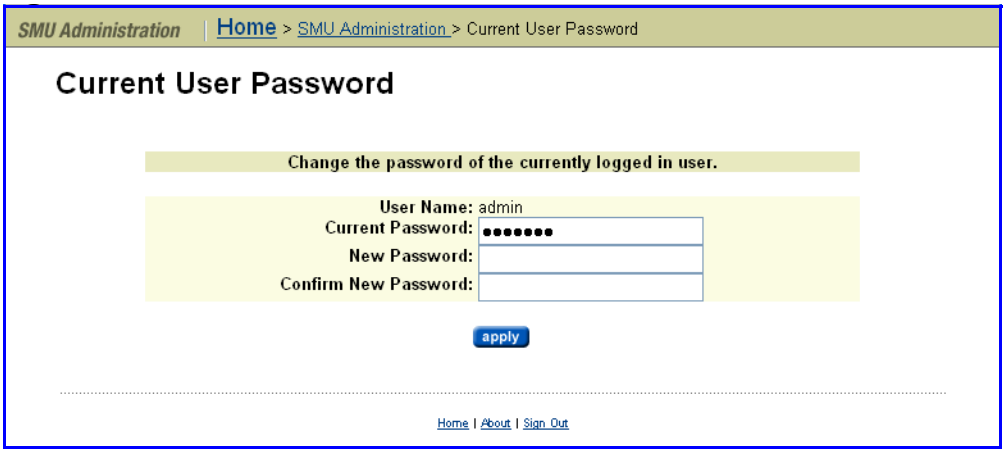

2. **Enter the requested information.**

   Enter the new password in both the **New Password** and **Confirm New Password** fields.

3. **Apply changes.**

   When finished, click **apply**.

Configuring the System Management Unit (SMU)

The System Management Unit (SMU) manages the storage servers/clusters and controls data migration and replication policies and schedules. After completing the steps in Quick System Configuration, on page 21 you can enable additional SMU capabilities. For example, you can:

- Secure the SMU, so that only certain predefined hosts can access the SMU for management purposes.
- Configure the SMU to act as an SMTP relay to the public network.

Note: For HTTP (Web Manager) access, the SMU ships with the default user name *admin* and the password `nasadmin`.

Configuring SMU Security

The SMU can be configured to control host access and auxiliary devices managed by the SMU.

To configure SMU security options:

1. **Navigate to the SMU Setup Wizard page.**

   From the **SMU Administration** page, display the configuration page by clicking **Security Options**:

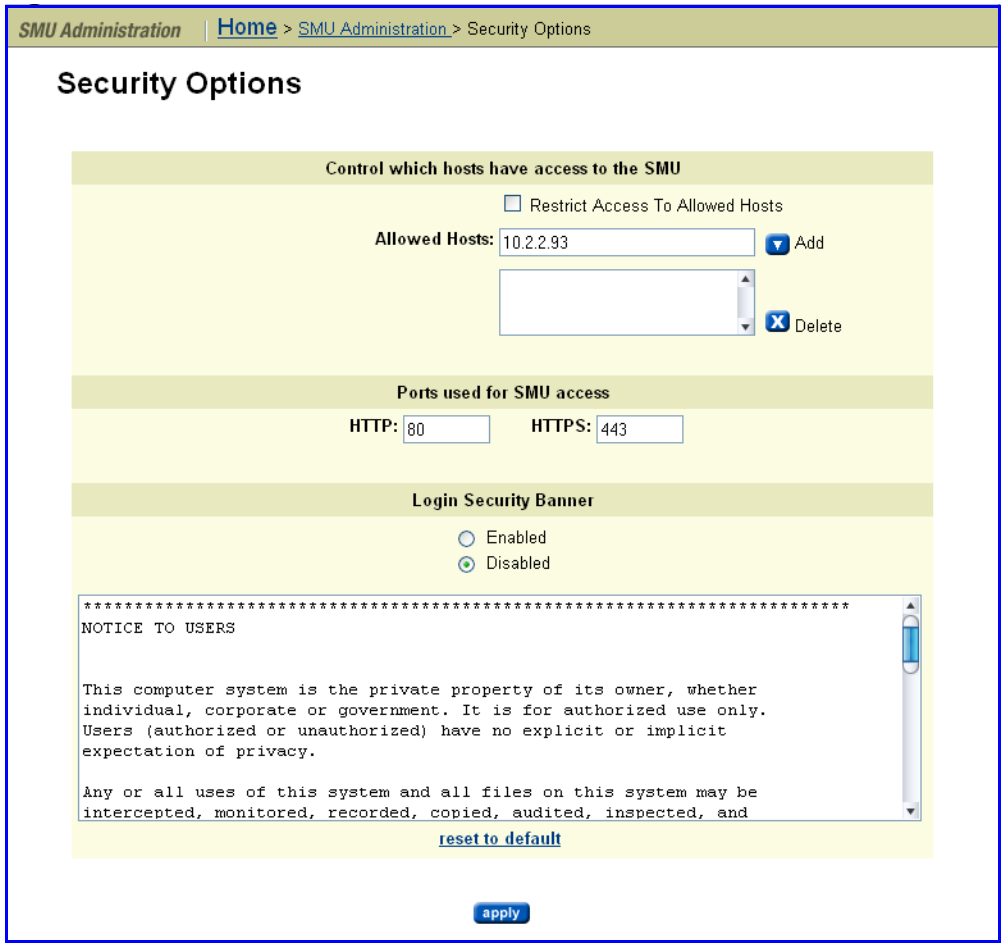

2. **Specify allowed hosts.**

   Enter the IP address of each allowed host, then click **add**. When the list is complete, make sure the **Restrict Access To Allowed Hosts** checkbox is filled, then click **apply**.

   Note: To prevent lockout of a host that is currently being used to manage the SMU, make sure to include the IP address of that host in the list of allowed hosts.

3. **Specify login security banner text.**

   Optionally, you can have the server display a message when logging in to the SMU using Web Manager, the serial console, or SSH.This message can be informational, a legal warning, or any other text you may want displayed. A default security banner is provided as a sample security message to users. You can customize this banner text by editing the text on this page. You can also click **reset to default**, which resets the banner text to the default text that is shipped with the SMU. By default, the security banner is disabled. Select the **Enabled** radio button to display the banner on the SMU login screen.

Changing the IP Address for a Managed Server

If the IP address of a managed server has been changed without using the Web Manager interface (for example, if the server's IP address was changed using the CLI or the console), you can update the IP address used by the SMU to communicate with the managed server.

Note: Updating the IP Address of a managed server does not actually change the IP address of the server, rather it tells the SMU the new IP address of the server. Updating the managed server's IP address does not interrupt management or delete completed replications or data migrations.

To update the IP address for a managed server:

1. **Navigate to the Managed Servers page.**

   From the **Home** page, display SMU-managed servers by clicking **SMU Administration**, then click **Managed Servers** to display the servers managed by the SMU.

2. **Select a managed server.**

   In the **Managed Servers** page, click **details** for a managed server to display the **Modify Server** page:

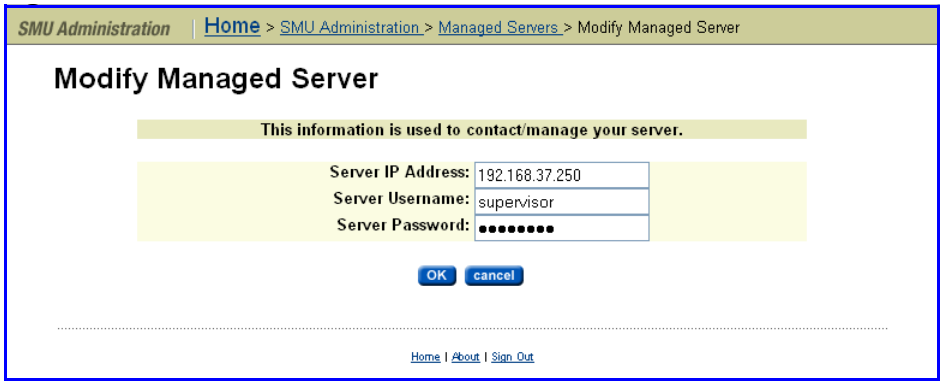

3. **Modify the IP address of the managed server.**

   Enter the requested information, then click **OK**.

   Note: Typically this IP address is assigned to the 10/100 management port.

Configuring an SMTP Relay for the SMU

The SMU can be configured to forward emails to the public network from the servers and auxiliary devices on the private management network, via SMTP relay, as illustrated here:

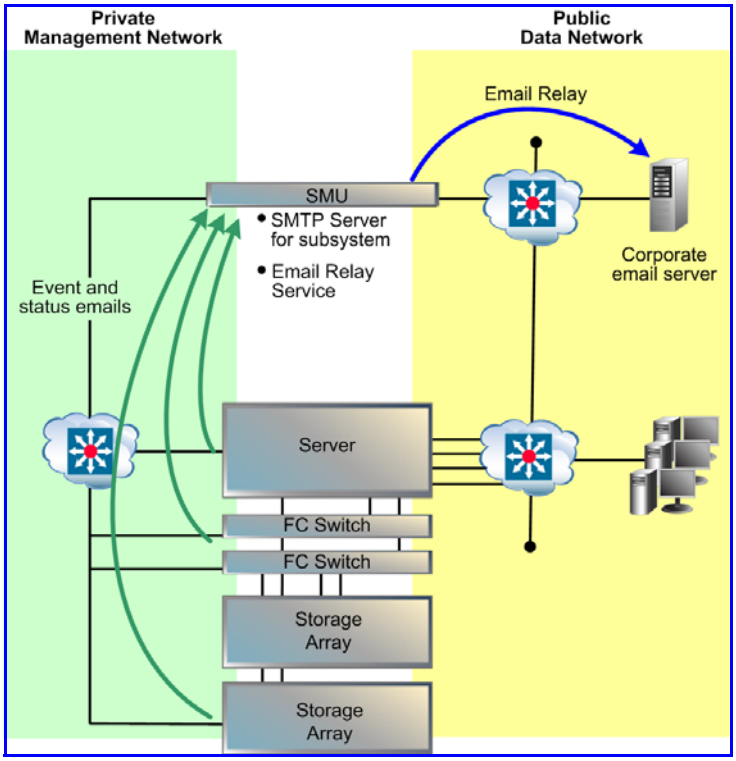

To configure an SMTP relay for the SMU:

1. **Navigate to the SMTP Configuration page.**

   From the **SMU Administration** page, click **SMTP Configuration** to display the configuration page:

2. **Specify an SMTP server on the public data network.**

   Enter the host name of an SMTP server on the public network, then click **apply**. The SMU will then relay emails to the public network from the devices on the private management network.

3. **Verify that the SMTP server IP address specified on the Email Alert Configuration page is set to the SMU's *eth1* IP address.**

   View the server's email configuration via the **Email Alerts Setup** link found on the **Status & Monitoring** page.

Configuring the Storage Server

Configuring Server Identification

This page defines system name and other identifying information for use by SNMP, SMTP (email) and other protocols. To configure the server identification:

1. **Navigate to the Server Identification page.**

   From the **Server Settings** page, display the configuration page by clicking **Server Identification**:

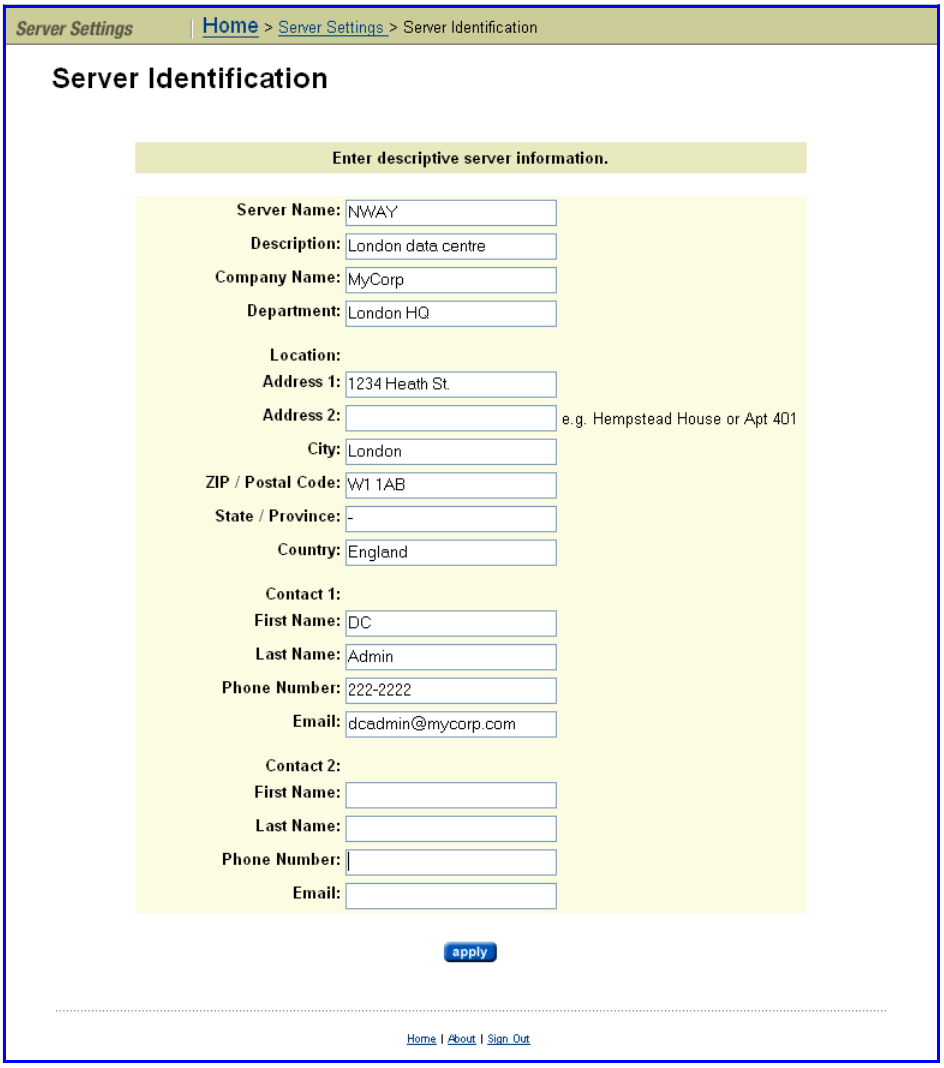

2. **Configure server identification.**

   Enter the requested information, then click **apply.**

Note: If configured for Microsoft Windows networking, the value entered for **Description** becomes the server's **Comment** for all configured CIFS names.

Configuring Date and Time

Administrators configure the server's current date, time, time zone, and NTP Server for synchronization.

Caution: Proper server operation requires time synchronization with a reliable time source. For example, Kerberos authentication (required when operating with Active Directory) depends on the current time. Clock 'drift' may also cause inaccurate reporting of file access and modification times, with unexpected results in data migrations. NTP provides the best and most reliable

method for maintaining the server's time accuracy.

1. **Navigate to the Date and Time page.**

   From the **Server Settings** page, click **Date and Time** to display the **Date and Time** page:

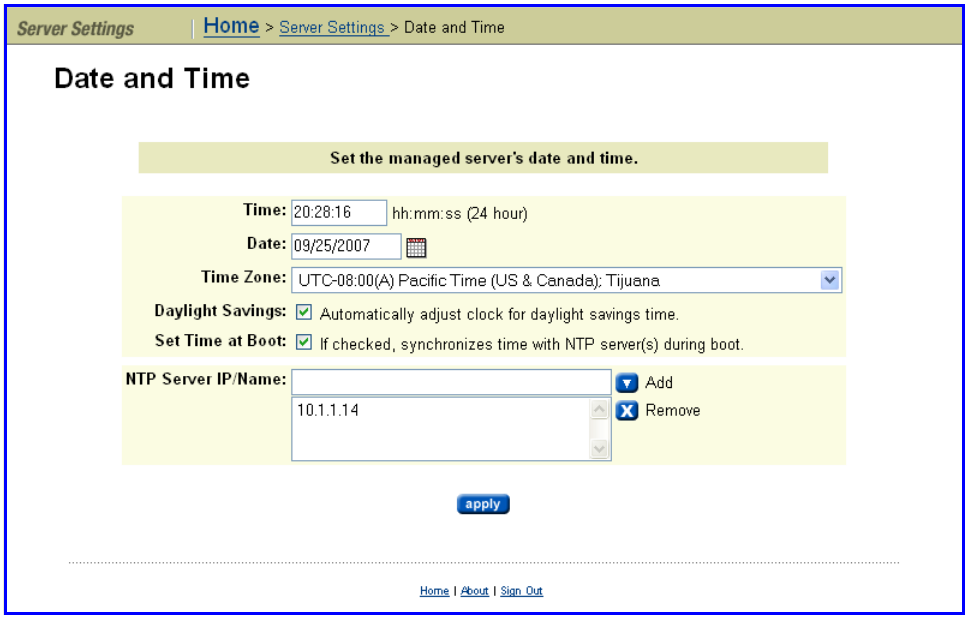

2. **Configure date and time.**

   Enter the requested information, then click **apply.**

   The following table describes the fields on this page:

Item/Field	Description
Time	In 24-hour format.
Date	Select from the calendar popup.
Time Zone	Select from the drop-down list. **Note:** For guidance on which zone to select, see: `http://www.worldtimeserver.com/`
Daylight Savings	Toggle enabled/disabled to auto-adjust. **Note:** Never try to compensate for daylight savings time by changing the time zone or the time.

Item/Field	Description
Set Time at Boot	Toggle enabled/disabled to synchronize time with NTP server on reboot: • If *disabled,* NTP aligns the server's time with the configured time server gradually and offsets of more than 15 minutes cause NTP updates not to register. • If *enabled* when the NTP service starts, the time synchronizes immediately, not gradually, and without regard for the current time offset. **Note:** A CLI command allows you to connect/ disconnect the server from a particular NTP service.
NTP Servers	Enter the IP address of the NTP server(s) you want to use to synchronize the server's time. You can specify several NTP server addresses, and the system will qualify and compare all listed NTP servers to determine and set the most accurate time. **Note:** For servers set up on the private management network, add the SMU's *eth1* IP address to the list of NTP servers.

Using the SMU for NTP

The SMU is configured as an NTP server. This ensures that every device on the private management network can synchronize with at least one NTP server. In turn, the SMU synchronizes with an NTP server on the public network. The following diagram illustrates this relationship:

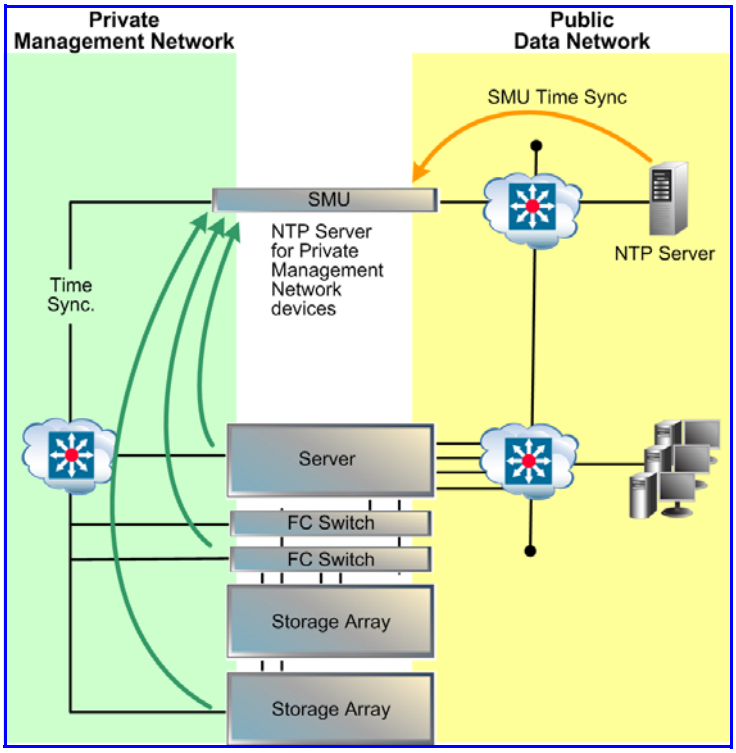

NTP Server Interaction

When using NTP, the server first verifies that the specified servers are legitimate; then, over a period of a few hours, gradually adjusts its clock to the time provided by the NTP server. This gradual adjustment is normal, and is designed to minimize the effects of changing the server's clock on utilities that use file timestamps.

If the time initially set on the storage server differs from the time returned by the NTP servers by more than 15 minutes, the server does not try to synchronize to the NTP time; instead, it records a *Warning* event in the event log, indicating that the date and time must be manually changed to within 15 minutes of the NTP time.

Configuring Server Management Access

The Web Manager provides the primary management interface for managing the server. In certain circumstances, however, an administrator may wish to use one of the following alternatives:

- For a IS-NAS Server:
  - The SSC utility, available for both Windows and Linux/Unix.
  - SNMP (Simple Network Management Protocol).
- For a Titan Server:
  - The command line interface (CLI), accessible through SSH and Telnet.
  - The SSC utility, available for both Windows and Linux/Unix.
  - SNMP (Simple Network Management Protocol).

To protect the server from unauthorized access, various safeguards have been built in. Statistics are available to monitor access through these various methods. The following sections detail the configuration options that secure the server's management interfaces and ports.

Note: To prevent unauthorized access to the storage system, SGI Global Services recommends configuring the server to respond only to predefined (authorized) management hosts on the network, based on the management access method (Telnet, SSC and SNMP) and defined port number.

Setting the Server Password

A password is required to authenticate direct management connections to the server. The password is required when adding a server to the SMU's list of managed servers, or when accessing a server directly through the command line interface.

1. **Navigate to the Change Password page.**

   From the **Server Settings** page, click **Change Password** to display the **Change Password** page.

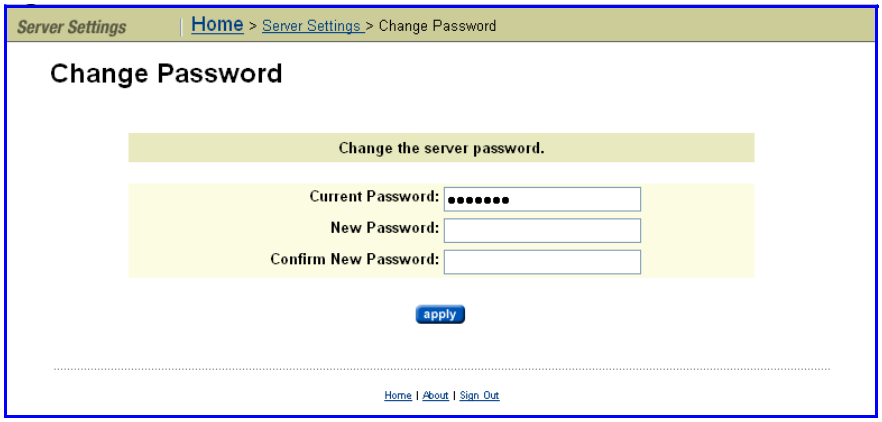

2. **Update password.**

   In the **Change Password** page, enter the requested information, then click **apply.**

Configuring Server Access Protocols

In addition to requiring a *user name* and a *password* for logon, both the IS-NAS Server and the Titan Server allow configuration of the following security settings:

- Ports for management protocol communications.
- Devices (hosts) that can be accessed using the management protocol.
- Disabling of unused management protocols.

To configure server access:

1. **Navigate to a protocol access page.**

   From the **Home** page, click **Server Settings**. Then, display a protocol access configuration page by selecting one of the access configuration links to display the page. The IS-NAS Server supports access through SSC and SNMP, and the Titan Server supports access through SSC, SNMP, TelNet, and SSH (SSC access configuration is illustrated here):

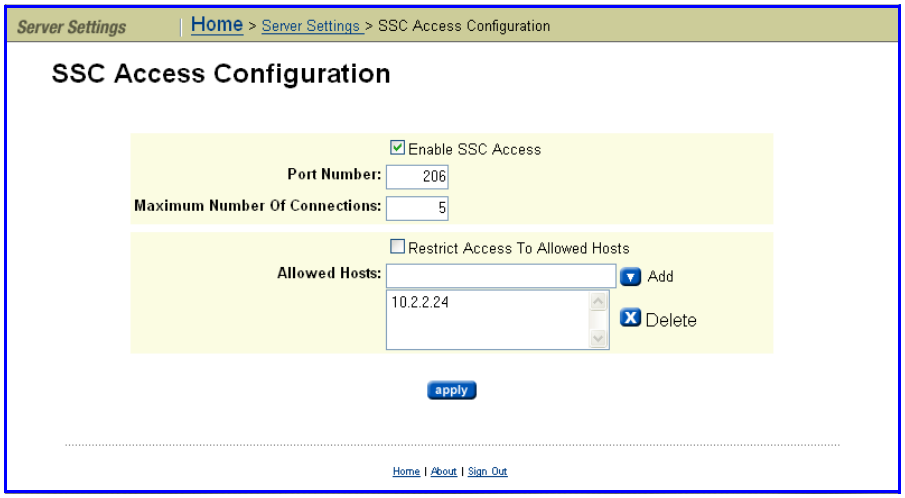

2. **Specify access configuration settings.**

   Using the access configuration page, enter the required information. Refer to the following table as needed:

Item/Field	Description
Enable Access checkbox	Fill the checkbox to allow access by the protocol or make sure the checkbox is empty to disable access using that protocol.
Port number	Enter the port number that the storage server should monitor for communication through the protocol. (Port 23 is the default Telnet access port, port 206 is the default SSC access port, and port 22 is the default SSH access port.)
Maximum number of connections	Specifies the maximum number of simultaneous connections to the server. For: • **Telnet**, you can allow up to 5 simultaneous connections. • **SSC**, you can allow up to 5 simultaneous connections. • **SSH**, you can allow up to 5 simultaneous connections.
Restrict Access to Allowed Hosts	Fill the checkbox to restrict protocol access to the hosts specified on this page. Make sure the checkbox is empty to enable the protocol to access any host.
Allowed Hosts	If protocol access is restricted to specified to hosts, use these fields to specify the hosts to which the protocol has access. **If protocol access is restricted to specified to hosts, make sure the SMU is an allowed host.** • **Allowed Hosts** (field). In the **Allowed Hosts** field, enter the IP address of a host that the protocol is allowed to access, then click **Add** to insert that host into the list of allowed hosts. When specifying IP addresses, you can specify an IP address range using the * as a wildcard character. For example: `10.168.*.*` or `172.*.*.*`. **Note:** If the system has been set up to work with a name server, you can identify allowed hosts by IP address or hostname. • **Allowed Hosts** (list). This list displays the IP address or hostname of each of the hosts that the protocol is allowed to access. To delete a host, select its IP address or hostname from the list and click **Delete.**
apply	Click **apply** to save the protocol access configuration settings.

3. **Save the access configuration.**

   Once you have entered and verified all the protocol access configuration settings, click **apply.**

Configuring the Private Management Network

The storage server operates in conjunction with a number of auxiliary devices, including Fibre Channel switches, power management units, primarily managed through Ethernet. In order to minimize the impact on an enterprise network, the SMU uses Network Address Translation (NAT) and Port Address Translation (PAT) to isolate the storage server from the main network through the SMU:

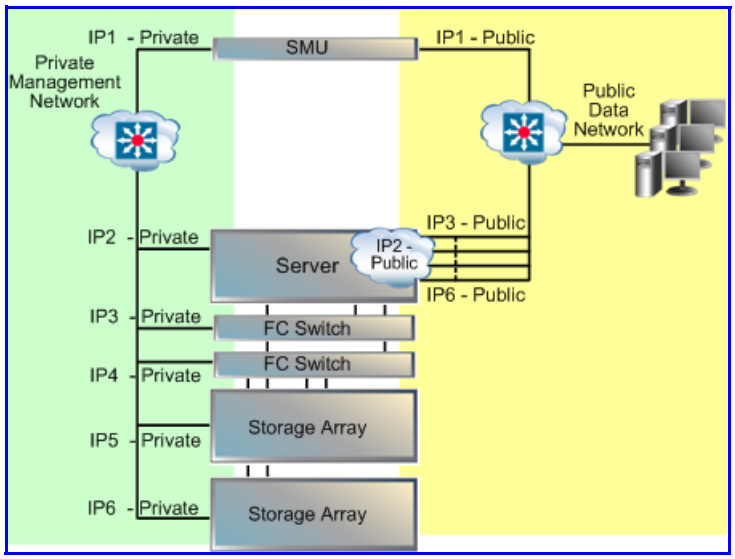

For example, an HTTP request for a device in the private management network would actually be made to the public IP address on the SMU's *eth0* interface, on a NAT-ed port (i.e., *192.168.1.124:28013*). The SMU translates this request to the private IP address and actual HTTP port of the device on the private management network (i.e. *192.0.2.13:80*), the NAT port.

The IP address range of the private management network includes only those IP addresses sharing the first three octets of the SMU's private (*eth1*) management network IP address. For example, for an SMU private management network IP address of *192.0.2.1*, devices on the private management network must have addresses in the range of *192.0.2.2 – 192.0.2.254*:

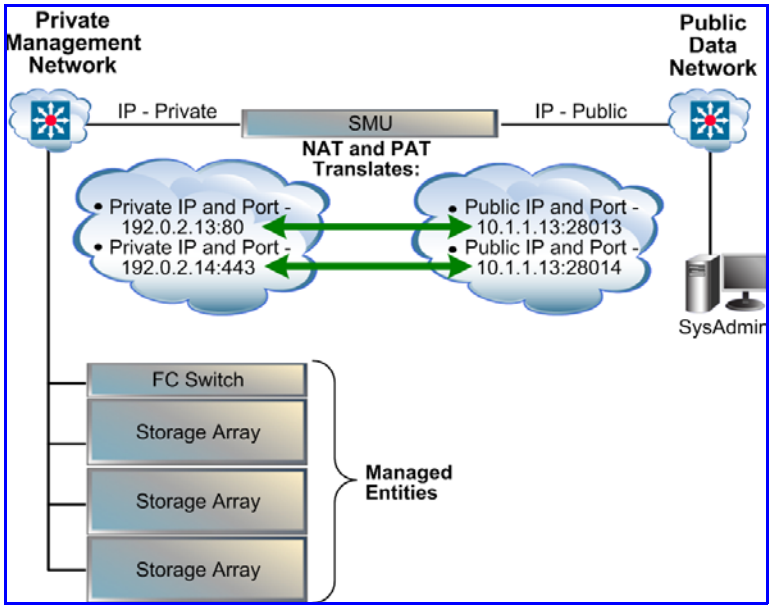

Significant advantages occur with a separate private management network:

- Network traffic required for normal SMU monitoring of the server and auxiliary devices will not be on the enterprise network.
- Devices on the private management network will not take up valuable IP addresses on the public data network.
- The SMU can discover all devices on the private management network, aiding setup.
- The private management network is more secure than the public data network.

As an alternative to the private management network, some or all of the auxiliary devices can be placed on the public data network. Such a configuration allows mixed systems, with some auxiliary devices isolated on the private management network, and others on the public data network.

Note: Devices on the public network require static IP addresses within the network.

Configuring the Management Network

To configure the management network:

1. **Navigate to the Management Network.**

   From the **SMU Administration** page, click to display the **Management Network** page:

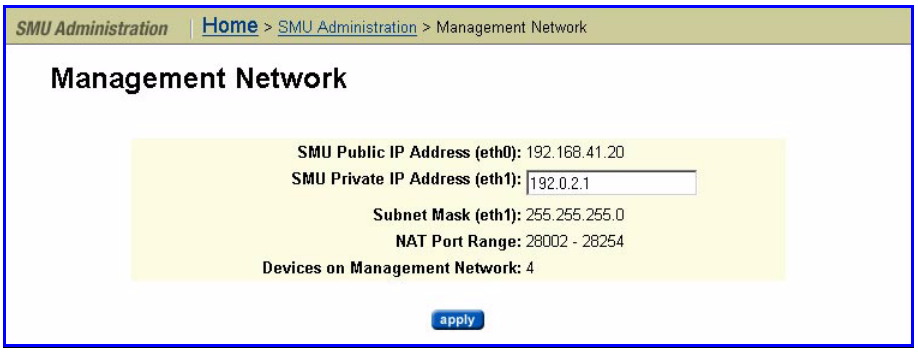

2. **Enter the requested information.**

   The **Management Network** page allows you to configure the *private management network address* of the SMU's *eth1* interface. The default address for the SMU's *eth1* port is *192.0.2.1.* Because the public network does not include this address, it falls into a distinctly different range than the SMU's public *eth0* address.

   Note: The *private management network address* must end with *.1,* to simplify the management relationship of the SMU with secondary devices.

   Note: The NAT Port range is provided for information only. It is rare that these values will ever need to be known.

3. **Apply the new settings.**

   Once defined, record the IP address settings separately for future reference when configuring the server's Administration Services IP address and subnet mask, then click **apply**.

Configuring Devices on the System Monitor

The system monitor allows you to easily display and monitor the devices that make up your storage system.

To configure devices on the System Monitor:

1. **Navigate to the System Monitor page.**

   From the **Home** page, click **System Monitor** to display the **System Monitor** page:

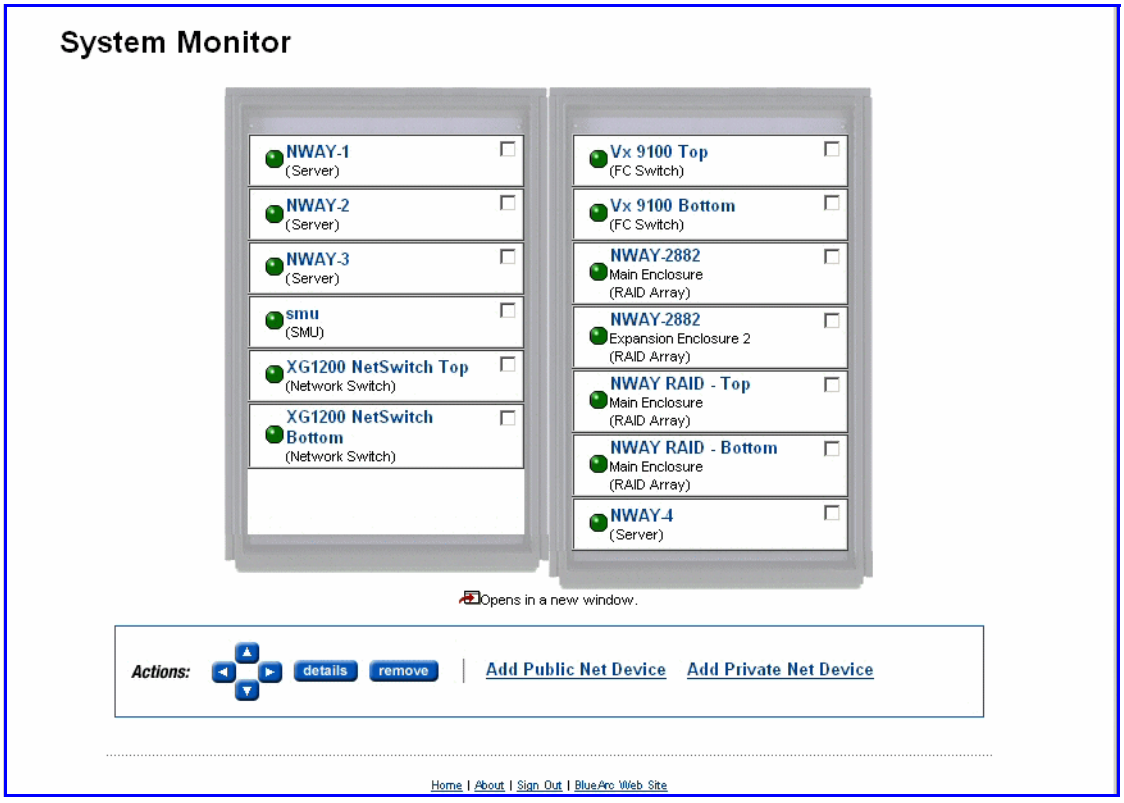

2. **Optionally, rearrange the sequence of components in the System Monitor.**

   To change the position of any component, fill its checkbox to select, then use the arrows in the **Action** section.

3. **Optionally, display status or details for any component in the System Monitor.**

   The rows in the following table list the basic components that make up a IS-NAS Server/Titan Server system. This table indicates what happens when you click on a component's name in the component list:

Component/Description	Clicking the component	Clicking the details button
Storage Server This component provides multiple Gigabit Ethernet interfaces to the network and multiple Fibre Channel interfaces to the main enclosure. In a cluster configuration, there are up to four nodes (servers).	Loads the **Server Status** page.	
Main Enclosure Contains dual power supplies, and dual RAID drive controllers. Depending on the model, the main enclosure may contain disk drives.	Loads the **Enclosure Status** page.	Loads the **System Drives** page.

Component/Description	Clicking the component	Clicking the details button
Expansion Enclosure Expansion enclosures contain disk drives and power supplies, but do not contain any RAID controllers.	Loads the **Enclosure Status** page.	Loads the **System Drives** page.
SMU The System Management Unit	Loads the **SMU System Status** page.	
System Power Unit This component is also known as an uninterruptible power supply (UPS).	Loads the **UPS Status** page.	Loads the **UPS Configuration** page.
NDMP Backup Devices The server automatically detects and adds backup devices to the system monitor. Since the storage server could be connected into a FC network shared with other servers, it does not automatically make use of backup devices found on its FC links. Backup devices are automatically discovered and added to the Status Monitor.	Loads the **NDMP Devices** page.	Loads the **NDMP Details** page for the device if the device can be contacted, or loads the **NDMP Device List** page if the device cannot be contacted.
FC Switches FC switches (and cables) connect FC devices, generally storage arrays, to the server(s).  **Note:** Upon adding an FC switch through the FC Switches page, it is automatically added to the System Monitor.	Loads either the embedded management utility for the switch, or the FC Switch details page for the switch, depending on the protocol specified when the switch was added. For more information, see Adding FC Switches, on page 531).	Loads the **FC Switch Details** page.
Other Components Any component can be added to the system monitor. If the device supports a web-based management interface, the management interface can be launched directly from the server management interface.	Loads the embedded management utility for the device.	Loads either the **Add Public Net Device** or the Add **Private Net Device** page. Settings for the component can be changed from this page.

4. **Optionally, add, remove, or display details about a device.**

   The following **Actions** are available and apply to selected components:

   - Click **remove** to delete a component.
   - Click **details** to display details regarding a particular component.
   - Click **add Public Net Device** to add a device residing on the public (data) network.
   - Click **add Private Net Device** to add a device residing on the public (data) network.

Note: Devices on the private management network are "hidden" from the data network through Network Address Translation (NAT).

Once a device has been added, to the System Monitor, clicking its name in the System Monitor:

- Opens its embedded management utility in the Web browser, using either HTTP, HTTPS, or Telnet.
- The SMU periodically checks for device activity and connectivity with the server; if a device fails to respond to network "pings", the System Monitor changes its color to red and the SMU issues an alert (devices can also be configured to send SNMP traps to the SMU).
- Events from the device will be added to the event log if the SMU has a MIB for the device.

Adding a Device from the Public (Data) Network

To add a device from the public data network:

1. **Navigate to the add Public Net Device page.**

   From the **System Monitor** page, click to display the **Add Public Net Device** page:

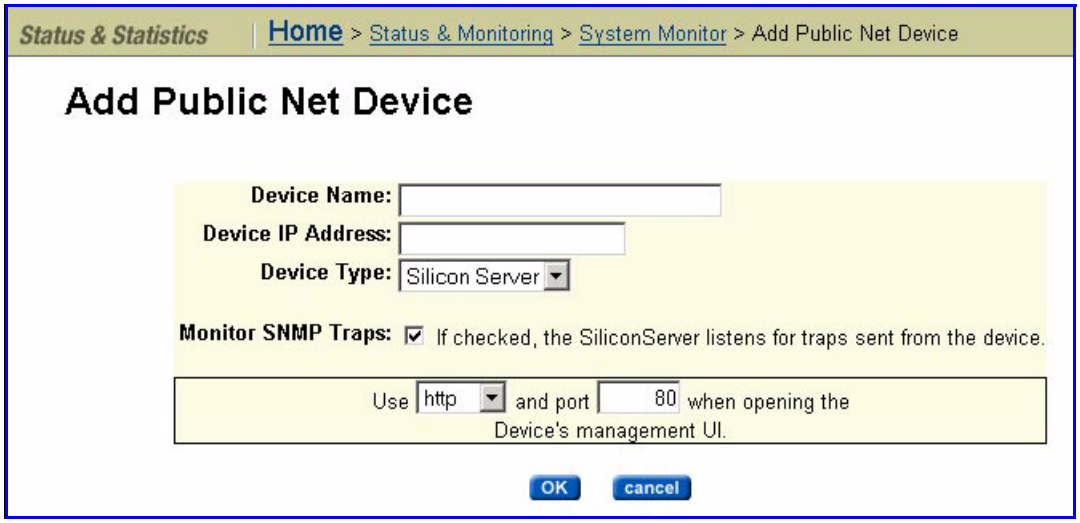

2. **Enter the requested information.**

   The table below describes the fields on this page:

Note: FC Switches are added to the System Monitor automatically, after being added through the **FC Switches** page. For more information, see Adding FC Switches, on page 531.

Item/Field	Description
Device Name	Enter any descriptive name to represent this device in the System Monitor.
Device IP Address	The IP address for the device.
Device Type	Select a device type that best describes the device. This is used purely as a label to help distinguish components in the System Monitor, and does not affect any functionality. Examples include *RAID Array* and *System UPS*.
Use Protocol and Port	Specify a protocol (e.g. HTTP) and port number (e.g. 80) to be used for accessing the device's management UI. For a device directly accessed for management by clicking on its name in the System Monitor, select HTTP, HTTPS, or Telnet and enter the corresponding port number. This information will be used to generate a link to the device. When you click on the device's name in the System Monitor, the browser uses the generated link to access the device's embedded management UI.  **Note:** If you specify "other," no access to an embedded management UI is configured. No link to the device's embedded management UI will be generated and no access to the embedded UI is available through the device name displayed on the System Monitor.

3. **Apply changes.**

   When changes are complete, click **apply.**

Adding a Device from the Private Management Network

To add a device from the private management network:

1. **Navigate to the add Private Net Device page.**

   From the **System Monitor** page, click to display the **Add Private Net Device** page:

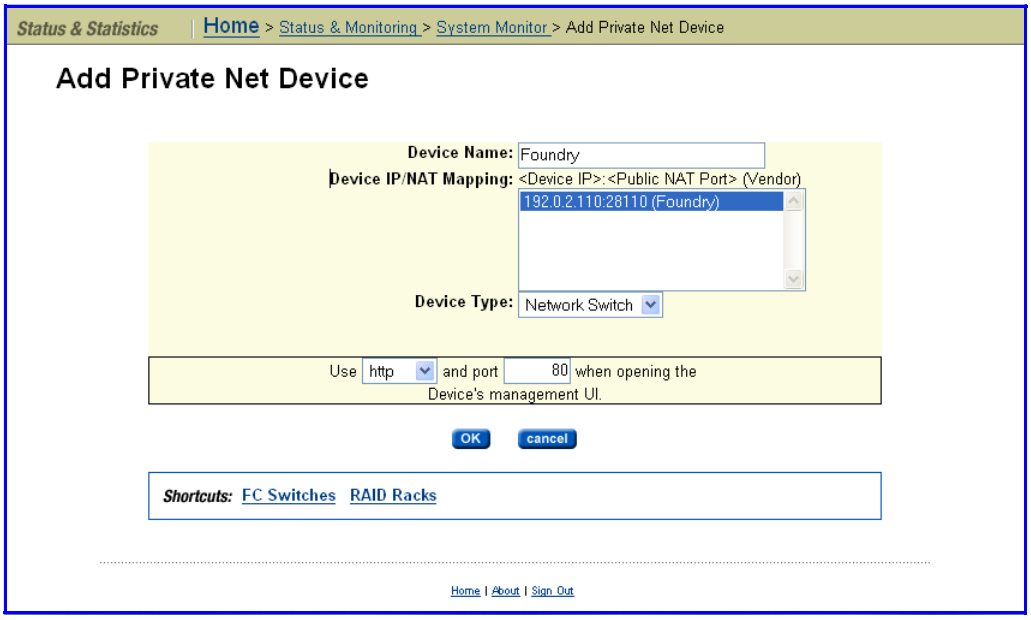

2. **Enter the requested information.**

Note: FC Switches are added to the System Monitor automatically, after being added through the **FC Switches** page. For more information, see Adding FC Switches, on page 531.

The table below describes the fields on this page:

Item/Field	Description
Device Name	Enter any descriptive name to represent this device in the System Monitor.
Device IP/NAT Mapping	Excluding devices already displayed in the System Monitor, devices discovered on the management network by the SMU are displayed here. The following Information is displayed for these devices: • **IP Address:** The private management network IP Address of the device (not directly accessible from the data network of the SMU). • **Public NAT Port:** This port on the SMU interface (*eth0*) accesses the management port on the device through the public network. • **Vendor:** Name of vendor corresponding to the device's MAC address. If the vendor is recognized, a Device Type is pre-selected; otherwise, “Generic” is displayed. Failure to recognize a Vendor or MAC address does not affect any functionality.
Device Type	Select a device type that best describes the device. This is a label used purely to help distinguish components in the System Monitor, and does not affect any functionality. Examples include RAID Array and System UPS.

Item/Field	Description
Use Protocol and Port	Specify a protocol (e.g. HTTP) and port number (e.g. 80) to be used for accessing the device's management UI. For a device directly accessed for management by clicking on its name in the System Monitor, select HTTP, HTTPS, or Telnet and enter the corresponding port number. This information will be used to generate a link to the device. When you click on the device's name in the System Monitor, the browser uses the generated link to access the device's embedded management UI.  **Note:** If you specify "other," no access to an embedded management UI is configured. No link to the device's embedded management UI will be generated and no access to the embedded UI is available through the device name displayed on the System Monitor.

Note: SGI recommends adding the SMU's *eth1* IP address to the device's list of NTP servers. Also, if the device supports email notification, and if email forwarding is configured on the SMU, the SMU's *eth1* IP can also be configured as the device's mail server.

Receiving SNMP Traps through the SMU

SNMP traps are alert messages sent by devices on the network. These traps provide information about failures or other conditions on those devices. Set the SMU's *eth1* IP address as the receiving target for SNMP traps sent by managed devices on the private management network. When a supported device sends a trap, the SMU decodes and registers it in each managed server's event log, detailing the trap's name and the contents of the trap's variable binding list.

The SMU supports and decodes traps from devices that support the following MIB modules:

- Fibre Alliance
- Brocade Silkworm
- DataDirect Networks

Note: Devices that do not support any of the Management Information Block (MIB) modules in the SMU's list can register traps in the storage server's event log by setting a server Administrative IP address as the receiving target for SNMP traps. Traps registered from APC devices will be properly decoded. Traps from any other device will be registered in unencoded form.

Managing Uninterruptible Power Supply Usage (Titan Server only)

An uninterruptible power supply (UPS), also known as a system power unit, isolates servers from loss of power, by providing power from a battery. Should the loss of power last long enough for the batteries to drain, the UPS will notify the server, which in turn conducts an orderly shutdown before power runs out. The server only supports Ethernet-connected APC SMART UPS and the APC Symmetra UPS devices.

Note: Currently, integrated UPS support and management functions are provided only for the BlueArc Titan Server, the SGI InfiniteStorage NAS Server does not include this functionality.

In order to receive alerts from the UPS, a server must be registered with the UPS as a PowerChute client. In a cluster, if the UPS is on the same subnet, only one IP address (the Administrative IP) needs to be registered; otherwise, each server must be registered individually.

Note: Each server has its own NVRAM, which it uses to buffer file system writes. In the event of a loss of power, the server will use its NVRAM to complete any disk transactions that were not saved to disk.

Adding a System Power Unit

To add a system power unit:

1. **Navigate to the UPS Configuration page.**

   From the **Server Settings** page, select **UPS Configuration** to display the **UPS Configuration** page:

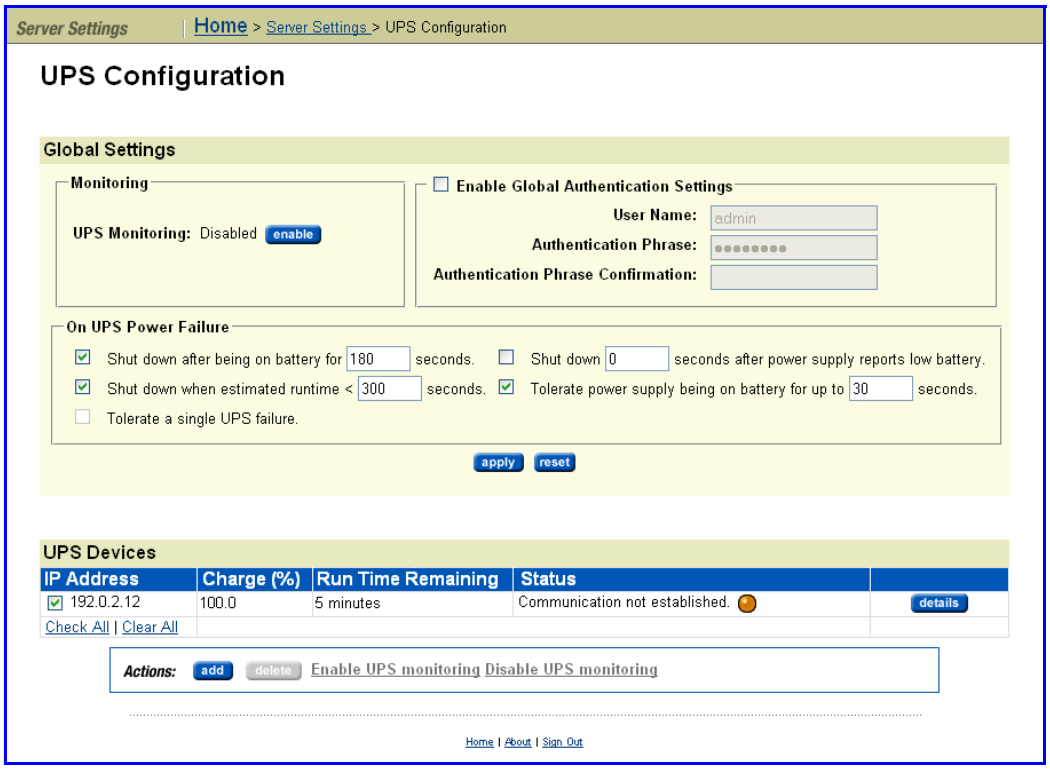

The table below describes the major sections on this page:

Item/Field	Description
Global Settings	This section displays global settings which, if enabled, the server applies to all configured UPS devices. The fields in this section allow you to: • Enable or disable UPS monitoring. • Enable or disable global authentication for UPS devices, and specify the global settings that are used when global authentication is enabled. • Specify response(s) to power failure events. For more information about these settings, see Configuring Power Failure Settings, on page 66.
UPS Devices	This section lists all of the configured UPS devices, and displays information about each device. For each configured UPS, the information in this section includes: • IP Address of the UPS. • The percentage of charge remaining in the UPS. • The run time remaining. This is the estimated amount of time that the server can operate on the battery power remaining in the UPS. • UPS device status, if available. For more information about these settings, see Adding a System Power Unit, on page 61.

2. **Navigate to the Add UPS Device page.**

   From the **UPS Configuration** page, click **add** to display the **Add UPS Device** page:

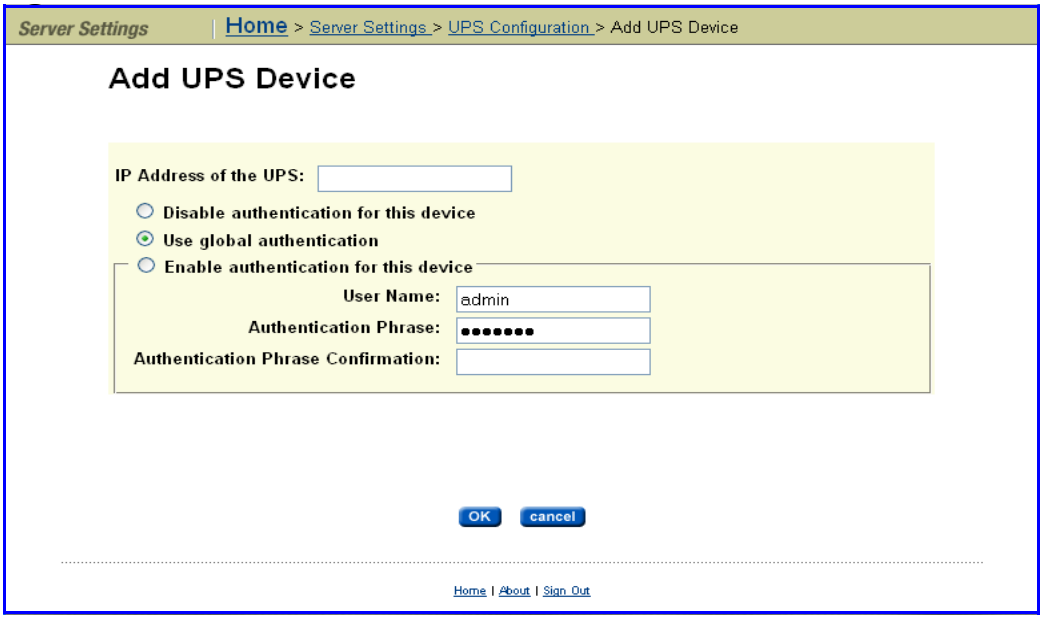

3. **Enter requested configuration.**

   The table below describes the fields on this page:

Item/Field	Description
IP Address of the UPS	Specify the IP address of the UPS to be connected to the server.
Disable Authentication for this device	Select this radio button to disable authentication for this device. The UPS must be configured not to require authentication.
Use Global Authentication	Select this radio button to enable the usage of the Global Authentication settings defined on the **UPS Configuration** page.
Enable Authentication for this device	Select this radio button to enable the usage of the authentication settings defined on this page. These settings may be different than the Global Authentication settings defined on the **UPS Configuration** page. • **User Name**. The user name may be a maximum of 8 characters. • **Authentication Phrase**. The authentication phrase must be between 15 and 32 characters in length, and it may contain only alphanumeric characters, spaces, and underscores. When you enter the authentication phrase, note that only asterisk (*) characters are displayed. • **Authentication Phrase Confirmation**. As with the originally entered authentication phrase, when you enter the confirmation authentication phrase, only asterisks are displayed.

4. **Save settings.**

   Click **OK** to save the settings, or click **cancel** to return to the **UPS Configuration** page without saving the settings.

Viewing or Changing UPS Configuration

To view or change the configuration of a UPS:

1. Navigate to the **UPS Configuration** page.

   From the **Server Settings** page, click **UPS Configuration** to display the **UPS Configuration** page.

2. Navigate to the **Edit UPS Configuration** page.

   Click on **details** for the UPS device to display the **Edit UPS Configuration** page.

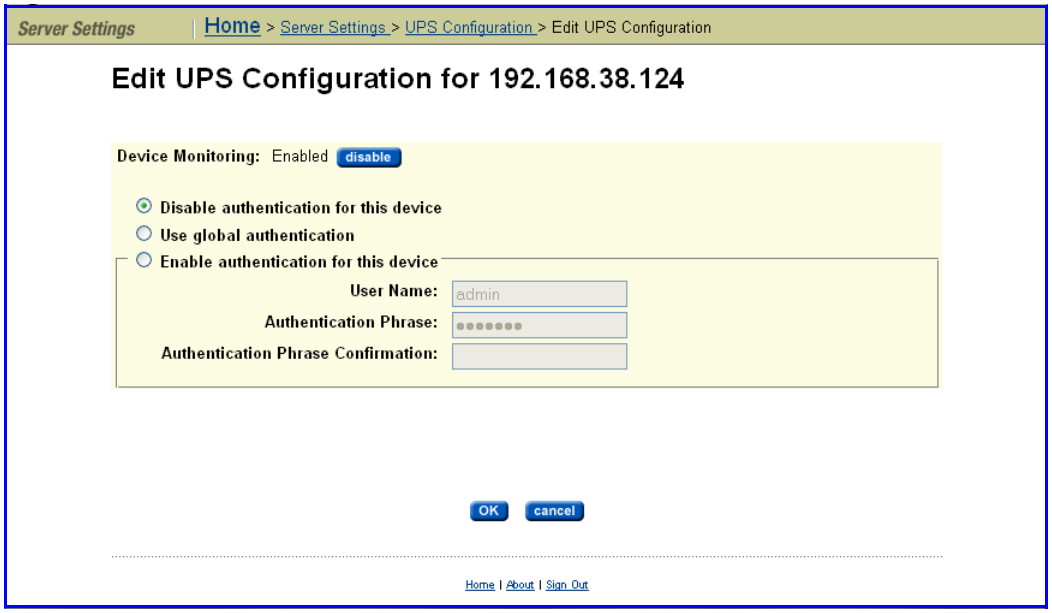

The table below describes the fields on this page:

Item/Field	Description
Device Monitoring	Indicates if monitoring is enabled or disabled for this UPS. Click **enable**/**disable** to start/stop monitoring this device.
Disable Authentication for this device	Select this radio button to disable authentication for this device. The UPS must be configured not to require authentication.
Use Global Authentication	Select this radio button to enable the usage of the Global Authentication settings defined on the **UPS Configuration** page.
Enable Authentication for this device	Select this radio button to enable the usage of the authentication settings defined on this page. These settings may be different than the Global Authentication settings defined on the **UPS Configuration** page. • **User Name**. The user name may be a maximum of 8 characters. • **Authentication Phrase**. The authentication phrase must be between 15 and 32 characters in length, and it may contain only alphanumeric characters, spaces, and underscores. When you enter the authentication phrase, note that only asterisk (*) characters are displayed. • **Authentication Phrase Confirmation**. As with the originally entered authentication phrase, when you enter the confirmation authentication phrase, only asterisks are displayed.

3. **Review and/or change the UPS Settings.**

Configuring Power Failure Settings

When a storage server is used with a UPS, you can specify actions to perform in the event of power loss to the UPS and what to do it the UPS is getting low or running out of power. To configure these settings:

1. **Navigate to the UPS Configuration page.**

   From the **Server Settings** page, select **UPS Configuration** to display the **UPS Configuration** page:

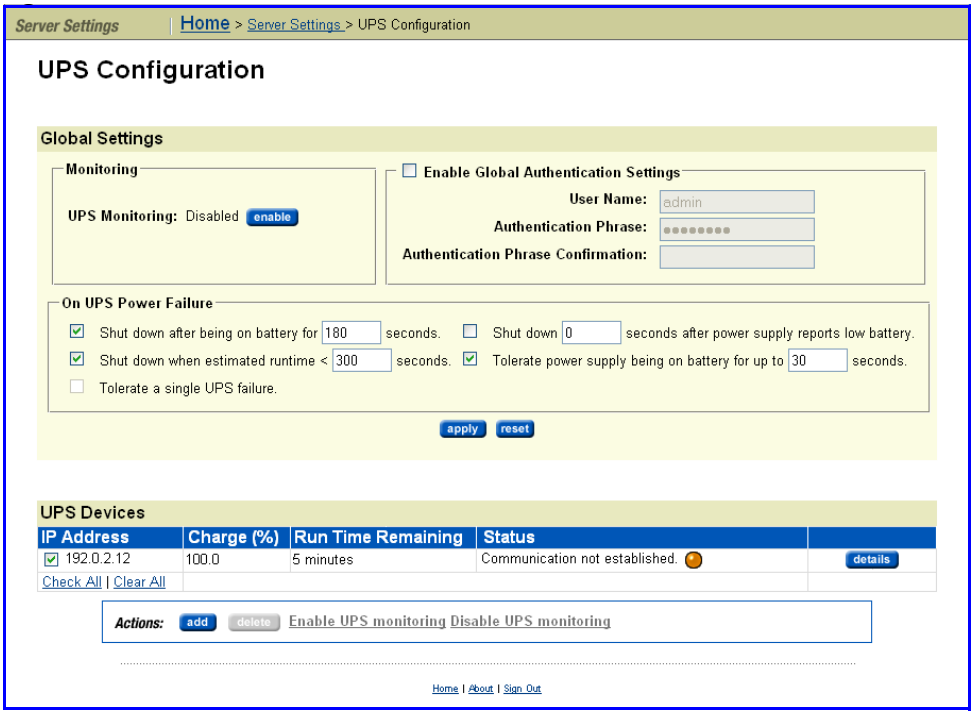

   The table below describes the fields on this page:

Item/Field	Description
Global Settings	
UPS Monitoring	Displays if UPS monitoring is enabled or disabled and allows you to start or stop monitoring. If monitoring is disabled, click **enable** to start monitoring the configured UPS units. If monitoring is enabled, click **disable** to stop monitoring the configured UPS units.
Enable Global Authentication Settings (checkbox)	Fill this checkbox to enable the usage of a single user name and authentication phrase for one or more UPS units. Leave the checkbox empty if you do not need to authenticate when connecting to the UPs or if you plan to specify the user name and authentication phrase manually for each UPS. **Note:** If you enable global authentication, and you want to use the global authentication settings to connect to a UPS, you must choose to use the global authentication connection settings when you add the UPS. See Adding a System Power Unit, on page 61 for more information about adding a UPS unit.

Item/Field	Description
User Name (field)	Enter user name to use for global authentication. The user name may be a maximum of 8 characters.
Authentication Phrase (field)	Enter the authentication phrase to use for global authentication. The authentication phrase must be between 15 and 32 characters in length, and it may contain only alphanumeric characters, spaces, and underscores. When you enter the authentication phrase, note that only asterisk (*) characters are displayed.
Authentication Phrase Confirmation (field)	Enter the authentication phrase again to confirm the phrase to use for global authentication. As with the originally entered authentication phrase, when you enter the confirmation authentication phrase, only asterisks are displayed.

Item/Field	Description
On UPS Power Failure	Us the following checkboxes to specify the way the server/cluster responds to power loss events. For each of the actions, fill the checkbox to enable the action or clear the checkbox to disable the action. Most actions allow you to specify a maximum time limit or a cutoff point. When the time limit or cutoff is reached, the action will be performed.  **Note:** The following settings are global, and they apply to all configured UPS units. Customizing actions to be taken un power failure on a per-UPS basis is not supported. • **Shut down after being on battery for x seconds.** Shut down the server if it has been running on UPS power for the specified number of seconds. • **Shut down x seconds after power supply reports low battery.** Shut down the server after a low battery event has been detected. Specify the duration (in seconds) for how long server will operate after the server receives low battery notification is received. After the specified number of seconds, the server will shut down. • **Shut down when estimated runtime < (is less than) x seconds.** Shut down the server when estimated runtime is less than the specified number of seconds. Use this option to shut down the server before the UPS runs out of power. The server estimates the amount of power remaining in the UPS and shuts down when the estimated run time is less than the specified number of seconds. • **Tolerate power supply being on battery for up to x seconds.** Tolerate power supply being on battery for up to a specified number of seconds. The server does not take any action on power failure for a specified number of seconds. This may be used to prevent unintended shutdowns due to UPS battery tests or maintenance. • **Tolerate a single UPS failure**. Where each UPS provides sufficient power to run the server, and more than one UPS is present, the administrator can configure the server *not* to shut down when one of the UPS units fails. This option appears only after the first UPS has been added.  **Note:** If you specify conflicting time limits or actions, the more conservative setting is followed. For example, if you specified to "Shutdown after being on battery for 60 seconds" and "Shutdown when estimated runtime is less than 300 seconds," the server/cluster will shut down after being on battery power for 60 seconds, even if there is more than 300 seconds of estimated runtime available.
UPS Devices	
IP Address	Displays the IP address of the UPS unit.
Charge	Displays the percentage of charge remaining in the UPS.
Runtime Remaining	Displays the estimated amount of time the UPS can continue to provide power to the server/cluster, based on charge remaining and power draw by the server/cluster.

Item/Field	Description
Status	Displays the status of the UPS unit. For an explanation of the status message, refer to the documentation for the UPS unit

2. **Specify response(s) to power failure events.**

   Identify what the server should do in the event of a power failure by customizing the settings in the **On UPS power failure** section (described above).

3. **Apply settings.**

   Click **Apply** to save the settings.

5 Network Configuration

Overview

This chapter presents storage server system networking concepts and procedures for configuring the *public data network* and the *private management network,* in the following sections:

- IP routing, including *static routes*, *default gateways,* and *dynamic routes,* with a brief discussion of *routing precedence*.
- Overview of the network interfaces, including the usage of jumbo frames and IP addressing for the *public data network*, the *private management network, clustering*, and VLAN support.
- Network statistics, historical and near-real-time.
- Name services, including *DNS, NIS, WINS,* and *LDAP.*

Network Interfaces

Each storage server is equipped with either Gigabit Ethernet (GE) ports or 10 Gbps Ethernet (10 GbE) ports and 10/100 Ethernet ports:

- Up to six GE ports, that support copper and fiber SFPs (Small Form-factor Pluggables). These ports support jumbo frames, and may be configured either individually or trunked together using IEEE 802.3ad link aggregation to provide high-performance access to the public data network.
- Two 10 GbE ports, that support copper and fiber XFPs (10 Gigabit Small Form-factor Pluggables). These ports support jumbo frames, and may be configured either individually or trunked together using IEEE 802.3ad link aggregation to provide very high-performance access to the public data network.
- 10/100 Ethernet ports use standard RJ-45 connectors, and are used to connect to the storage server's private management network.

Note: If your Titan Server is equipped with a NIM3 module, there are two 10 GbE ports and a single 10/100 port. The 10 GbE ports are used to connect to the public data network, and the 10/100 port is used to connect to the private management network. If your Titan Server is equipped with a NIM2 module, there are two 6 GE ports and four 10/100 ports. The GE ports are used to connect to the public data network, and the 10/100 ports are used to connect to the private management network.

For more information about the Titan Server modules and connections, refer to the Titan Server *Hardware Reference*.

Network clients use either the Gigabit Ethernet (GE) data interfaces or the 10 Gbps Ethernet (10 GbE) interfaces, configured for *diverse routing* or *link aggregation,* to access the storage server:

- With *diverse routing,* the administrator configures each port to support an IP subnet, to support physically connecting a server to a maximum of six separate IP subnets.
- With *link aggregation* (or trunking), the administrator configures multiple GE ports or 10 GbE ports together into an aggregation, which is then assigned to one EVS or several EVSs.

Note: Note that all ports in an aggregation must be of the same type/speed (either all GE ports or all 10 GbE ports).

An aggregation has a single MAC address, and at least one IP address. An aggregation assumes the IP address of each EVS which uses that aggregation, meaning that an aggregation will have multiple IP addresses if it is assigned to more than one EVS. Physical ports can be aggregated in any combination, as long as all ports in each aggregation are of the same type/speed. The server is initially configured with a single port aggregation containing GE port 1.

Link aggregation increases the network interface bandwidth for individual connections. Link aggregation also isolates the server from network infrastructure failure; for example, if some of the links in an aggregation fail, the other links in the aggregation share the traffic.

The server supports *Link Aggregation Control Protocol* (LACP), which automatically configures link aggregation settings when the server is connected to a switch that supports LACP.

The system supports mixed environments with simultaneous *diverse routing* and *link aggregation* on the same server or cluster.

IP Addressing

The server distinguishes between IP address requirements for the *public data network, the private management network,* and *clustering:*

- **File services (public data network).** Network clients access the server's file services through file service IP addresses, which are accessible only through the server's GE ports. Multiple IP addresses can be assigned for file services (these IP addresses may be on the same or different networks, but must be unique).
- **Administration services (private management network).** These IP addresses are used when managing a server or cluster, through the Web Administration Manager or using the server's embedded management interfaces. The server requires at least one IP address, which is assigned to

the 10/100 Ethernet port. Additional administrative IP addresses can be assigned to GE ports, so that management functions may be performed directly through these network ports using SSC (IS-NAS Servers/clusters and Titan Servers/clusters) or Telnet (Titan Servers/clusters only).

Note: When configuring an Administration Services IP address on the private management network, verify that the subnet mask for the IP address matches that of the SMU's private management network (*eth1* port); for example, *255.255.255.0*. Also, choose an IP address that resides within the private management network's range; for example, *192.0.2.2-254*. This should be the Administrative Services IP address used when configuring a server as the managed server on the SMU.

- **Clustering.** When configured as a cluster, each node requires a unique IP address for the 10/100 management port connected to the private management network. These unique addresses enable cluster node to communicate with each other and with the Quorum Device (QD).

VLAN Support

A VLAN (virtual LAN) makes it appear that a series of devices seem to be operating on a single dedicated network, regardless of their physical location or the network to which they are physically attached. These devices form a separate logical network regardless of the physical connections (a logical network within a physical network), and they disregard all traffic that is not for nodes of "their" VLAN.

To support a VLAN, each of the switches that link the devices involved in the VLAN network must be configured to properly support the VLAN traffic and make (or keep) the VLAN traffic separate from other traffic on the network. Very complex configurations are possible, including (but not limited to) having a dedicated physical subnet as a part of the VLAN, and having the switch(es) serving that subnet remove the VLAN tags for inbound traffic and add the VLAN tags for outbound traffic.

Because VLAN configuration can rapidly become complex, this section describes only how the storage server supports VLAN tagging.

The storage server supports VLAN tagging for the public data network interfaces (the GE and 10 GbE ports). The storage server's VLAN support is compatible with the 802.1Q standard, with the following exceptions:

- The use of 802.1Q to mark priority.
- Token Ring settings (ignored on receive and cleared on transmit).
- The storage server does not support stacked VLANs (consecutive VLAN tags).

The storage server VLAN implementation assumes that a VLAN corresponds to a single virtual network. As a result, each IP subnet may belong only to a single VLAN. However, many IP subnets may be configured within a VLAN. Note that VLANs are optional, and VLAN tags are not added by default.

To configure a VLAN, add a VLAN identifier (ID) for each source subnet, and use the `vlan` command on the storage server to associate particular subnets

with VLAN tags. The same command is used to display the current VLAN settings.

When the storage server is configured to use VLAN tags:

- Traffic sent through a port that belongs to a subnet in a VLAN is tagged, meaning that the VLAN tag (the subnet ID) is added in a header to the outgoing data frame. All transmitted packets are tagged with the VLAN tag (ID) of the subnet to which the source device belongs.
- All traffic received by the storage server is evaluated. Packets (data frames) are discarded if they do not contain a VLAN tag, or if the VLAN tag does not match one of the subnet IDs with which the storage server has been associated.

Jumbo Frames

All GE interfaces of a server support jumbo frames, which enable transmission of Ethernet frames larger than the Ethernet standard of 1,518 bytes. By reducing the number of frames required for large transfers, jumbo frames effectively increase transfer rate. Jumbo frames co-exist with standard frames on an Ethernet network.

All GE interfaces receive jumbo frames unconditionally, without any configuration changes. A GE interface can be configured to transmit jumbo frames by specifying an MTU size of between 1,519 and 9,000 bytes. To configure jumbo frame transmission, see Modifying Advanced IP Network Settings, on page 90, and configure the following settings:

- IP MTU for off-subnet transmits - bytes
- TCP MTU
- Other Protocol MTU

Caution: Networking equipment lacking the jumbo frames extension may drop jumbo frames and record an oversize packet error. Before configuring Jumbo Frame transmission, verify that all network equipment along the route (and at each end point) supports jumbo frames. If you enable jumbo frames and either network equipment or clients on the subnet do not support jumbo frames, you may experience a loss of communication with the server/cluster.

Successful IP data transmission using jumbo frames depends on the destination IP address or sub-network. The maximum MTU size for a destination IP address or sub-network is configured as an attribute in the IP routing table.

IP Routing

Depending on configuration, the storage server can route IP traffic in three ways: through *Default Gateways*, *Static Routes*, and *Dynamic Routes*. The illustration below shows how a server may be configured to communicate with various IP networks through routes:

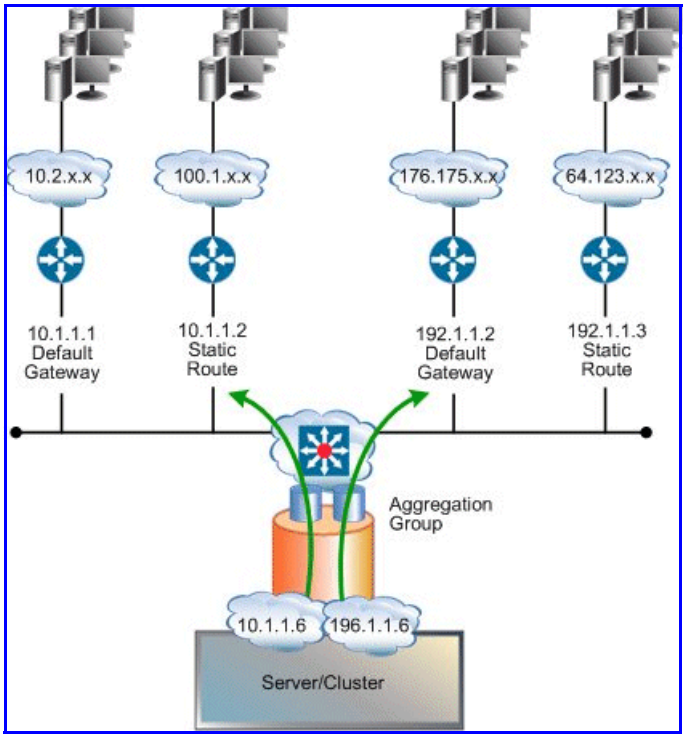

Default Gateways

The server supports multiple default gateways for routing IP traffic. When connected to multiple IP networks, add a default gateway for each network to which the server is connected. This configuration allows the server to direct traffic through the appropriate default gateway by matching source IP addresses specified in outgoing packets with the gateway on the same subnet.

With multiple default gateways, the server routes IP traffic logically, reducing the need to specify static routes for every network that connects with a particular server.

Static Routes

Static routing provides a fixed path for data in a network. When a server on a network is connected to additional networks through a router, communication between that server and the remote network(s) can be enabled by specifying a static route to each network.

Static routes are set up in a routing table. Each entry in the table consists of a destination network ID, a gateway address, and (sometimes) a subnet mask. Entries for static routes in the server's routing table are persistent, meaning that, if a server is restarted, the route table preserves the static routing entries.

The server supports both network- and host-based static routes. Select the **Network** option to set up a route to address all of the computers on a specific network. Select the **Host** option to address a specific computer on a different network than its usual router address. The maximum possible number of static routes is 127 (default gateways also count against this total).

Dynamic Routes

The server supports *ICMP redirects* and *RIP version 2 (RIPv2),* which allows it to dynamically add routes to its route table:

- **ICMP redirects** is an industry standard for routers to convey routing information back to the server. When one router detects that another router offers a better route to a destination, it sends the server a redirect that temporarily overrides the server's routing table. Being router-based, dynamic redirects do not require any configuration, but they can be viewed in the routing table.

  The server also supports ICMP router discovery, which allows it to discover the addresses of routers. ICMP routers periodically multicast their addresses; when the server receives these multicasts, it incorporates the routers into its routing table. Once a router appears in the server's routing table, it can be used as a gateway.

  ICMP router discovery is controlled using the CLI command `irdp`. For more information, refer to the *Command Line Reference*.

- **RIPv2** is also an industry standard, allowing servers to automatically discover routes and then update routes in the route table based on updates provided by other network devices. RIPv2 is controlled using the CLI command `rip`. For more information, refer to the *Command Line Reference*.

The server stores dynamic host routes in its route cache for ten minutes. When the time has elapsed, packets to a selected destination use the route specified in the routing table until the server receives another ICMP redirect.

Network Statistics

Fibre Channel, Ethernet and TCP/IP statistics for the server (per port in ten-second timeslices) are available. These statistics pages show activity since the previous reboot or since the point when statistics were last reset.

Name Services

The administrator can configure the server to work with a local name server and to support the following name resolution methods:

- Domain Name System (DNS)
- Dynamic Domain Name System (DDNS)
- Windows Internet Naming Service (WINS)

These methods associate computer identifiers (e.g., IP addresses) with computer names. This allows you to specify computer names rather than IP addresses in dialog boxes.

DNS and DDNS

On TCP/IP networks, the Domain Name System (DNS) is used to resolve host names into IP addresses.

With DNS, records must be created manually for every host name and IP address. Starting with Windows 2000, Microsoft enabled support for Dynamic DNS, a DNS database which allows authenticated hosts to automatically add a record of their host name and IP address, eliminating the need for manual creation of records.

Registering a CIFS Name

When an EVS goes online, the server registers one entry with the configured DNS servers (in both the forward and reverse lookup zones) for each configured *ADS CIFS name* and *IP address* associated with the EVS. Thus, the EVS records one entry in DDNS for every configured IP address. If a server has more than one configured ADS CIFS name, an entry for each IP address for each configured CIFS name is registered.

Each hostname registered with the DNS server has a Time To Live (TTL) property of 20 minutes, which is the amount of time other DNS servers and applications are allowed to cache it. The record's TTL dwindles with passing time and when the TTL finally reaches zero, the record is removed from the cache. After the 20-minute expiration point, the client must execute a fresh name lookup for more information.

The hostname is refreshed every 24 hours. This refresh commences after the first successful registration. For example, if the server registers its name at bootup, then every 24 hours after the bootup it refreshes its DNS entry. If the server cannot register or refresh its name, it goes into recovery mode with an attempt to register every 5 minutes. Once it successfully registers, it will resume the 24 hours-per-refresh cycle.

Secure DDNS Updates

The storage server supports both secure and insecure DDNS updates. By default, Microsoft Windows 2000, 2003, and 2008 DDNS servers only accept “secure”, Kerberos-authenticated registrations. To support both Microsoft and non-Microsoft DDNS servers, the server will first attempt to register with DDNS insecurely. If the insecure registration fails, the server will attempt a secure registration.

WINS

WINS resolves NetBIOS names to IP addresses, and is used by the server to communicate with CIFS clients on the network.

Note: WINS is deprecated in Windows 2008.

Directory Services

The administrator can configure the server to work with a local directory server and to support the location, administration, and management of network resource. The following directory service methods are available:

- Network Information Service (NIS)

- Lightweight Directory Access Protocol (LDAP)

These services associate identifiers with users, groups, devices, volumes, folders, and other network resources. These services associate an identifier of some kind with a resource, allowing you to specify policies for access on a broad basis, rather than explicitly on a per-resource basis, and to have this information accessible throughout your network.

NIS (for NFS)

NIS databases provide simple management and administration of Unix-based networks. These databases can provide details about users and groups, also individual client machines (including IP address and host name, to facilitate authentication for users logging in to clients on the network).

The server supports NIS and, when configured to use NIS, can provide the following:

- NFS user and group account information retrieval;
- Name services for resolving host names to IP addresses;
- (FTP) authentication.

LDAP

Many organizations are replacing their existing NIS infrastructure with the more reliable, scalable and secure system LDAP. In addition to providing the same services as NIS (user and group information retrieval, name service resolution, and FTP user authentication), LDAP also provides the following advantages:

- Improved accuracy, due to LDAP's more frequent data synchronization of current and replicated data.
- Communications encryption using Secure Sockets Layer (SSL) and Transport Layer Security (TLS).
- Authentication of connections to the LDAP database, instead of anonymous access to NIS databases.

The server supports LDAP version 2, including two of the most common LDAP service implementations:

- Sun Directory Server
- OpenLDAP

Note: LDAP cannot be used to resolve NIS Netgroups. If Netgroups are required, local Netgroups must be used.

Configuring the Gigabit Ethernet Data Interfaces

GE (Gigabit Ethernet) and 10 GbE (10 Gigabit Ethernet) port configuration requires setting up the following components:

1. Link aggregations.

2. IP Addressing, including Advanced IP Settings and Router Table Settings, for file and block services provided by the server.

Link Aggregations

Viewing Link Aggregation Status

Link aggregations combine multiple GE or 10 GbE ports into a single logical link for increased bandwidth. Aggregations allow you to increase the capacity and availability of the communications channel between the server/node and remote devices using the server's GE or 10 GbE ports.

In an aggregation, two or more like (GE or 10 GbE) ports are grouped, forming a single logical unit, to increase bandwidth capability and create resilient and redundant links. An aggregation also provides load balancing where the processing and communications activity is distributed across several links in a trunk so that no single link is overwhelmed. Aggregations provide higher link availability and increased link capacity.

To view status of an aggregation, navigate to the **Link Aggregation** page:

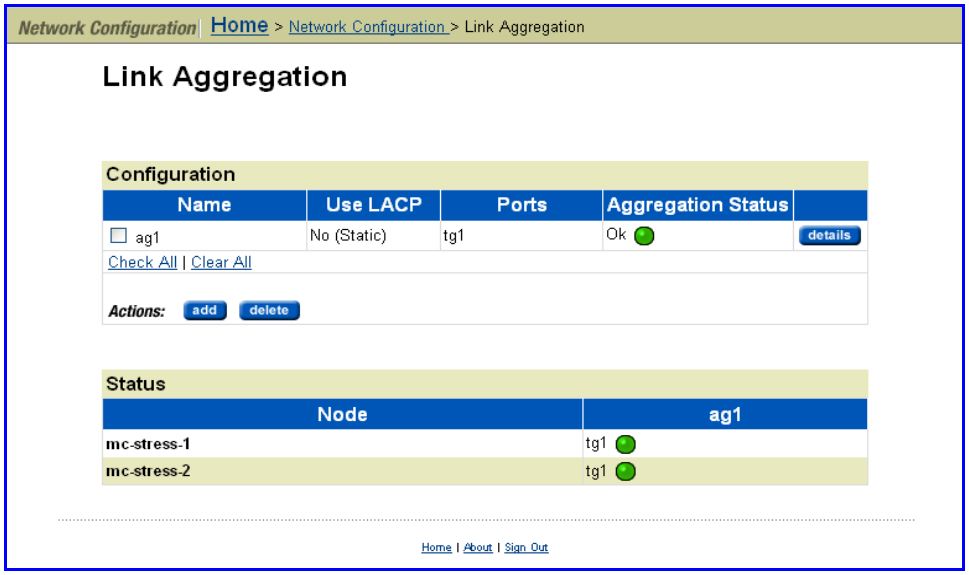

The following table describes the fields in this page:

Item/Field	Description
Name	Name of the aggregation (ag1, ag2, ag3, ag4, ag5, ag6).
Use LACP	Indicates if the aggregation uses LACP or not. If the aggregation does not use LACP, it is static. If the aggregation does use LACP, it is dynamic.
Ports	The list of ports used in the aggregation. Ports named "ge*x*" are Gigabit Ethernet ports, and ports named "tg*x*" are 10 GbE (10 Gigabit Ethernet) ports.
Status	Status of the aggregation.
Node	For each server/node, the current status of each configured aggregation is displayed.

Viewing or Changing the Aggregation Configuration

To view or change the configuration of an aggregation:

1. **Navigate to the Link Aggregation page.**

   From the **Network Configuration** page, select **Link Aggregation** to display the **Link Aggregation** page, which will list all currently configured aggregations.

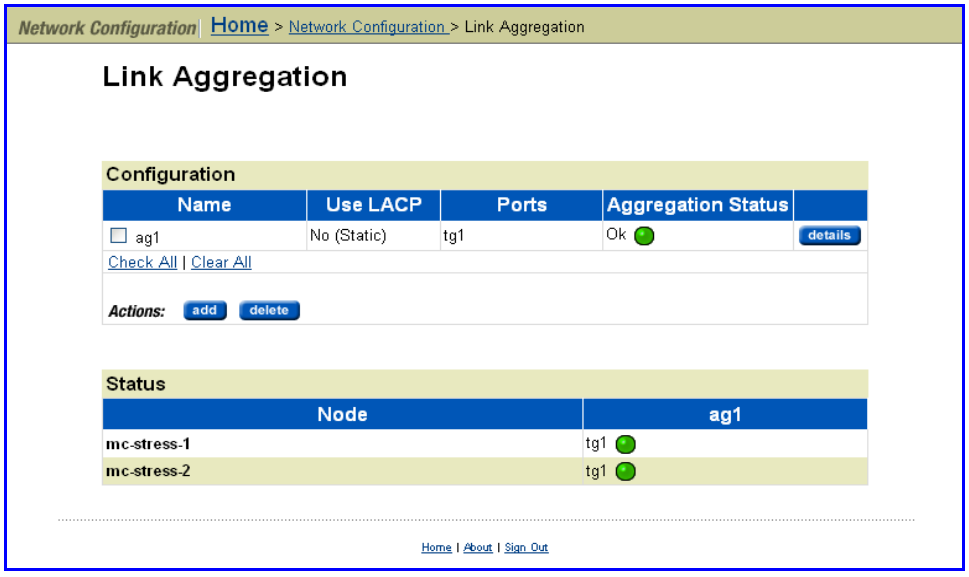

2. **Select the aggregation with the configuration you want to view or modify.**

To view or change an aggregation's configuration, click **details** to display the aggregation's **Edit Link Aggregation Details** page.

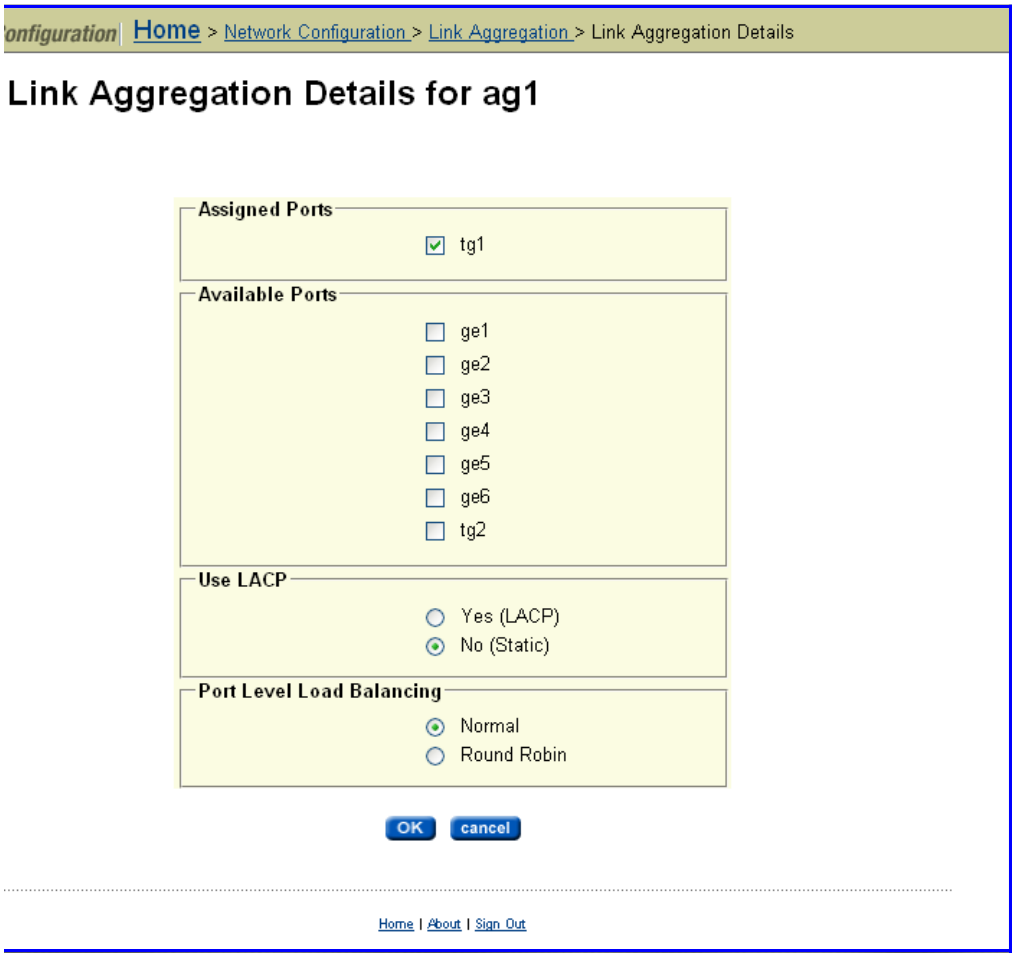

The following table describes the fields in this page:

Item/Field	Description
Assigned Ports	Lists the ports currently assigned to this aggregation. Ports named "ge*x*" are Gigabit Ethernet ports, and ports named "tg*x*" are 10 GbE (10 Gigabit Ethernet) ports. To remove a port from the aggregation, empty the checkbox next to the name of the port you want to remove.
Available Ports	The available GE (Gigabit Ethernet) and tg (10 GbE) ports that may be added to the aggregation. Ports named "ge*x*" are Gigabit Ethernet ports, and ports named "tg*x*" are 10 GbE (10 Gigabit Ethernet) ports. To add a port to the aggregation, fill the checkbox next to the name of the port you want to add.
Use LACP	Specify if this aggregation is to use LACP or not. An aggregation that does not use LACP is called a static aggregation, and an aggregation that does use LACP is called a dynamic aggregation.

Item/Field	Description
Port level Load Balancing	Displays the port load balancing scheme used for all ports in the aggregation. • **Normal** means that the server routes all traffic for a given "conversation" through one of the physical ports in the appropriate aggregation. The server's hash and routing functions determine which packets use which physical ports of the aggregation. For example, all traffic for a particular TCP connection will always be routed through the same physical port (unless the link drops). • **Round Robin** means that the packets making up the traffic are routed through the ports in sequential order. For example, the first packet goes down the first port, the second packet goes down the next port and so on until all ports have been used. Then the traffic starts again at the first port. This routing scheme ensures that all the ports are more or less equally used, to provide maximum link throughput. The disadvantage of round robin is that the clients must be able to cope with out of order TCP traffic at high speed. The LACP specification (802.3ad) requires that an implementation must follow the appropriate rules to minimise out of order traffic and duplicated packets. Round robin load balancing directly contravenes this requirement. However, there are situations where the server's hash functions cannot balance the conversations across physical ports very well, resulting in poor link utilisation and reduced throughput. In these cases, round robin load balancing can improve link utilisation and improve throughput. Select the radio button next to the port loading scheme you want the aggregation to use.

3. **Make any required changes.**

   If you want to change the aggregation's configuration, you can use the **Edit Link Aggregation Details** page to:

   - Remove ports from the aggregation.
   - Change the type of load balancing used in the aggregation.
   - Add ports to the aggregation.

4. **Save the changes.**

   Click **OK** to save the changes, or click **cancel** to return to the **Link Aggregation** page.

Adding Aggregations

To add aggregations:

1. **Navigate to the Add Link Aggregation page.**

From the **Network Configuration** page, select **Link Aggregation**, then click **add** to display the **Add Link Aggregation** page:

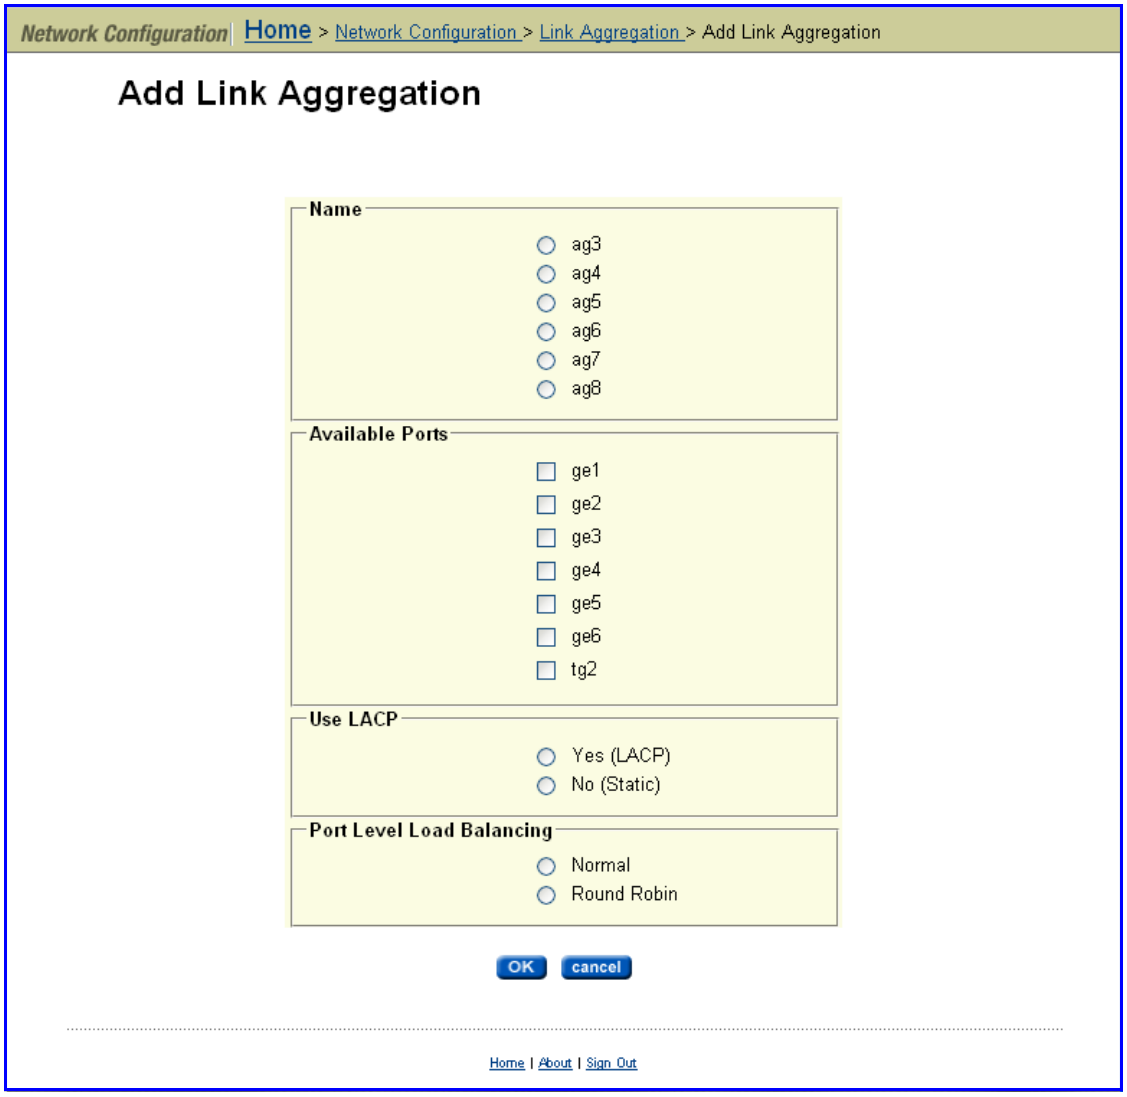

2. **Specify the aggregation configuration.**

   Using the checkboxes and radio buttons on the **Add Link Aggregation** page, specify the configuration of the aggregation.

Item/Field	Description
Name	Lists the available aggregation names (aggregation names not currently in use). To select the name for this aggregation, fill the radio button next to the name you want the aggregation to use.

Item/Field	Description
Available Ports	The available GE (Gigabit Ethernet) and tg (10 GbE) ports that may be added to the aggregation. Ports named "ge*x*" are Gigabit Ethernet ports, and ports named "tg*x*" are 10 GbE (10 Gigabit Ethernet) ports. To add a port to the aggregation, fill the checkbox next to the name of the port you want to add.
Use LACP	Specify if this aggregation is to use LACP or not. An aggregation that does not use LACP is called a static aggregation, and an aggregation that does use LACP is called a dynamic aggregation.
Port level Load Balancing	Displays the port load balancing scheme used for all ports in the aggregation. • **Normal** means that the server routes all traffic for a given "conversation" through one of the physical ports in the appropriate aggregation. The server's hash and routing functions determine which packets use which physical ports of the aggregation. For example, all traffic for a particular TCP connection will always be routed through the same physical port (unless the link drops). • **Round Robin** means that the packets making up the traffic are routed through the ports in sequential order. For example, the first packet goes down the first port, the second packet goes down the next port and so on until all ports have been used. Then the traffic starts again at the first port. This routing scheme ensures that all the ports are more or less equally used, to provide maximum link throughput. The disadvantage of round robin is that the clients must be able to cope with out of order TCP traffic at high speed. The LACP specification (802.3ad) requires that an implementation must follow the appropriate rules to minimise out of order traffic and duplicated packets. Round robin load balancing directly contravenes this requirement. However, there are situations where the server's hash functions cannot balance the conversations across physical ports very well, resulting in poor link utilisation and reduced throughput. In these cases, round robin load balancing can improve link utilisation and improve throughput. Select the radio button next to the port loading scheme you want the aggregation to use.

3. **Apply settings.**

   Verify your settings, then click **OK** to apply the settings or **Cancel** to decline.

Deleting Aggregations

To delete an aggregation:

1. **Navigate to the Link Aggregation page.**

Caution: *Aggregation deletion alert!* Before deleting an aggregation, all IP addresses, GE, and 10 GbE ports associated with the aggregation must be removed.

From the **Network Configuration** page, click to display the **Link Aggregation** page.

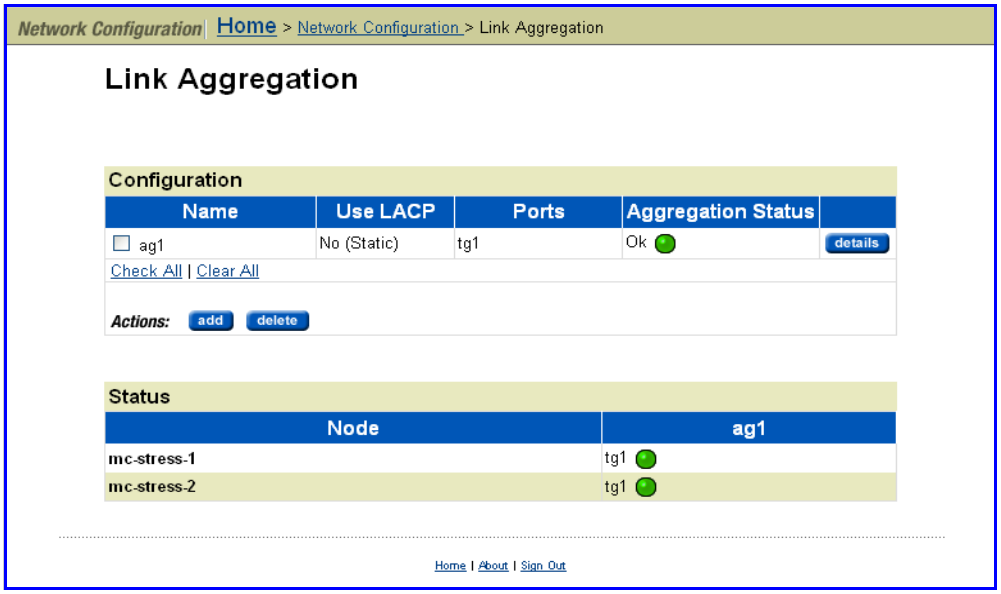

2. **Select the aggregation to delete.**

   Fill the checkbox by the name of the aggregation you want to delete.

3. **Click delete to immediately remove the aggregation.**

   Note: When deleting an aggregation, there is no confirmation required. When you click **delete**, the aggregation will be deleted immediately.

IP Addressing

At least two IP addresses are required to configure the server for access through the public data network. These are:

- A public IP address (an IP address on the public data network) on the System Management Unit (SMU) for server administration.
- A public IP address on at least one of the aggregation groups, to be used for file services. If all the GE ports use link aggregation, a single IP address supports all ports in an aggregation group.

Note: If configured for diverse routing (independent configurations per GE port), each port must have its own IP address.

Viewing Existing IP Addresses

To view existing IP addresses, navigate from the **Network Configuration** page to the **IP Addresses** page:

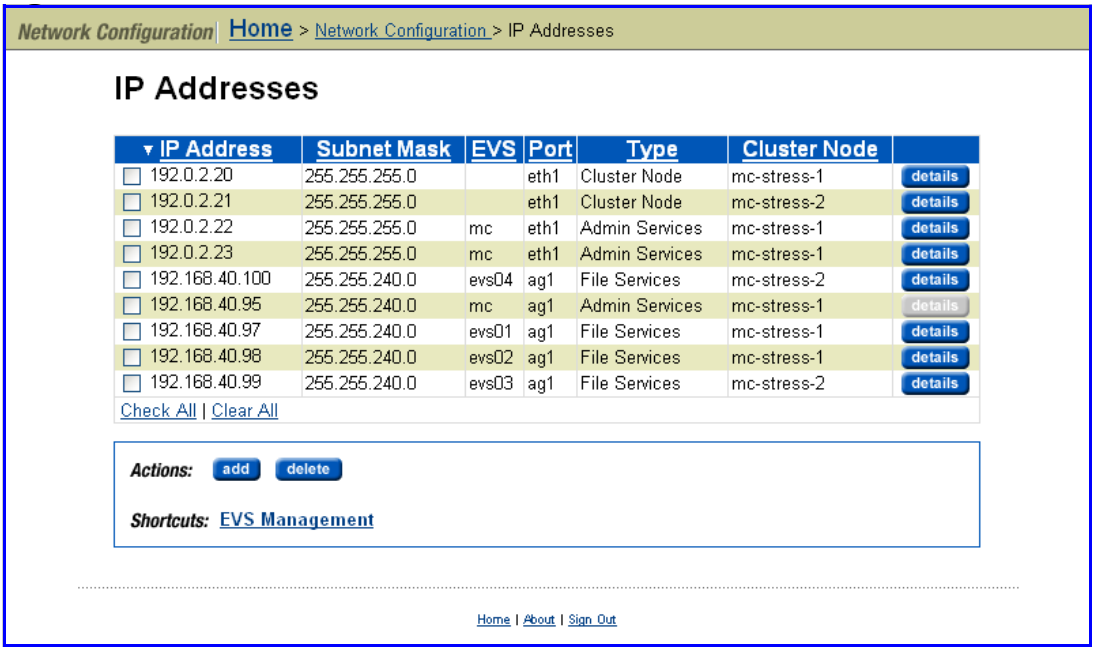

The following table describes the fields in this page:

Item/Field	Description
IP Address	IP address used for Admin or File services or for server/cluster node management.
Subnet Mask	Subnet mask of the services or cluster node.
EVS	One of the following: • In a stand-alone server configuration, if the server name is displayed, then the IP address is an administrative IP for the server. • In a cluster configuration, the IP address for an administrative services EVS. • The label of the EVS (Virtual Server) to which the file services IP is bound. • If there is no label displayed, the IP address is for server/node management.
Port	The interface used by the IP address: • `agX` identifies one of the GE aggregations • *eth0* or *eth1* identifies a 10/100 port for a IS-NAS Server • `mgmnt1` identifies the 10/100 management port for a Titan Server
Type	Type of services or configuration of the server: • **Admin Services:** an IP address associated with the Administrative Services for the cluster. Administration Services IP address may be on the public data network or on the private management network. • **File services:** an IP address associated with the File Services for the cluster. File Services IP addresses must be on the public data network. • **Cluster node:** the IP address associated with the physical cluster node. Because File and Administrative services may migrate between nodes, the Cluster Node IP address is used to communicate with the node instead of a service.

Item/Field	Description
Cluster Node	If configured as a cluster, the name of the cluster node to which the IP address is currently assigned.

Adding an IP Address

To add an IP address to a port or an aggregation:

1. **Navigate to the add IP Address page.**

   From the **Home** page, click **Network Configuration**, then click **IP Addresses** to display the **IP Addresses** page, and finally click **add** to display the **Add IP Address** page:

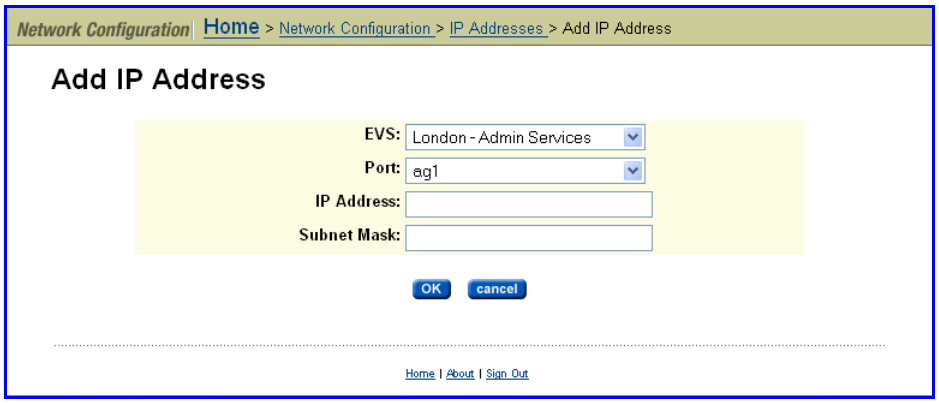

2. **Select a Virtual Server (EVS) to which to assign the IP Address.**

   From the drop-down list, select the EVS to which the IP will be assigned. Alternatively, specify that the IP address should be used for Admin Services.

3. **Select a port:**

   Select an aggregation or management port.

   From the drop-down list, select an aggregation (*agX*), or a management port (*mgmnt1* for a Titan Server, or *eth0* or *eth1* for a IS-NAS Server).

   Note: When assigning an IP address to an EVS, an Ag port must be specified.

4. **Define IP addressing:**

   Enter the IP address and Subnet Mask for the selected port.

5. **Apply settings.**

   Verify your settings, then click **OK** to apply the settings or **cancel** to decline.

Removing an IP Address

Caution: *IP Address deletion alert!* Before following the instructions in this step

to *delete* an IP Address, disable the EVS to which the IP Address is assigned. Once the IP address has been removed, the EVS should be re-enabled. This ensures that IP addresses are not in use at the time they are removed.

To remove an IP Address:

1. **Disable the EVS:**
   a. **From the Server Settings page, click to display the EVS Management page:**

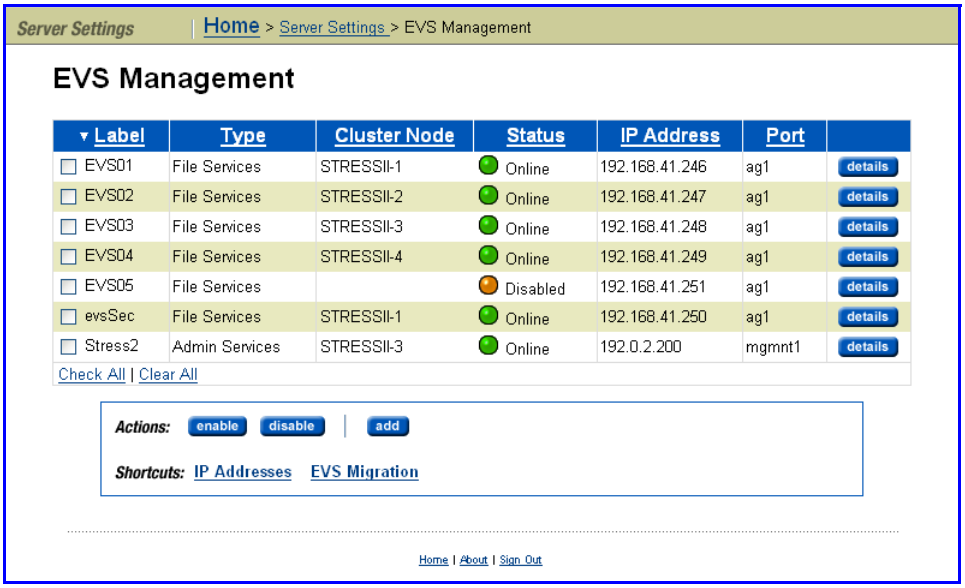

   b. **Select the EVS to which the IP is assigned, then click disable.**
2. **Delete the IP address.**

a. From the Network Configuration page, click IP Addresses to display the IP Addresses page:

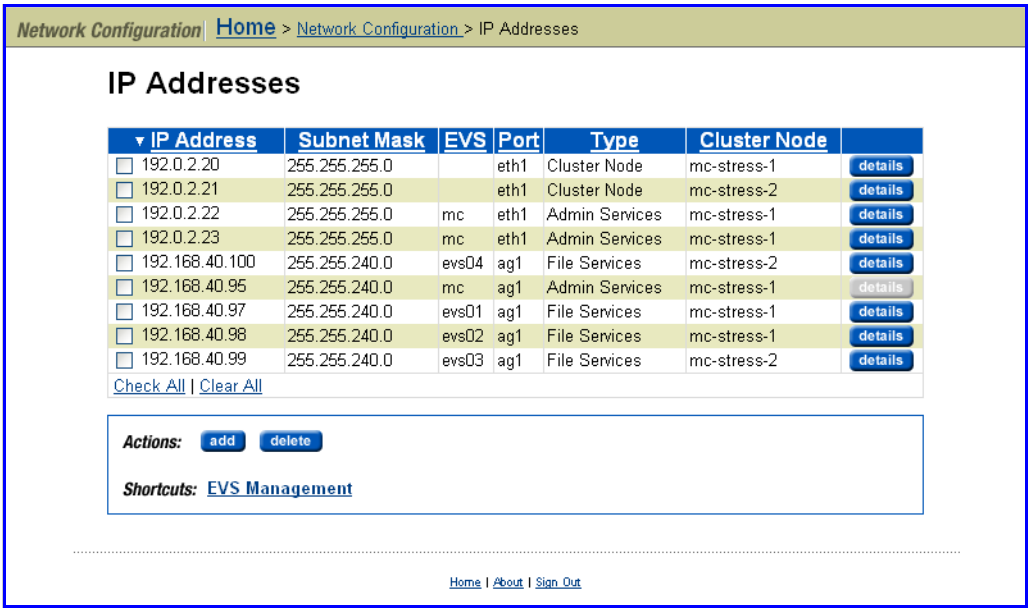

b. Select the IP Address to delete, then click delete.

3. Re-enable the EVS:

a. From the Server Settings page, click to display the EVS Management page:

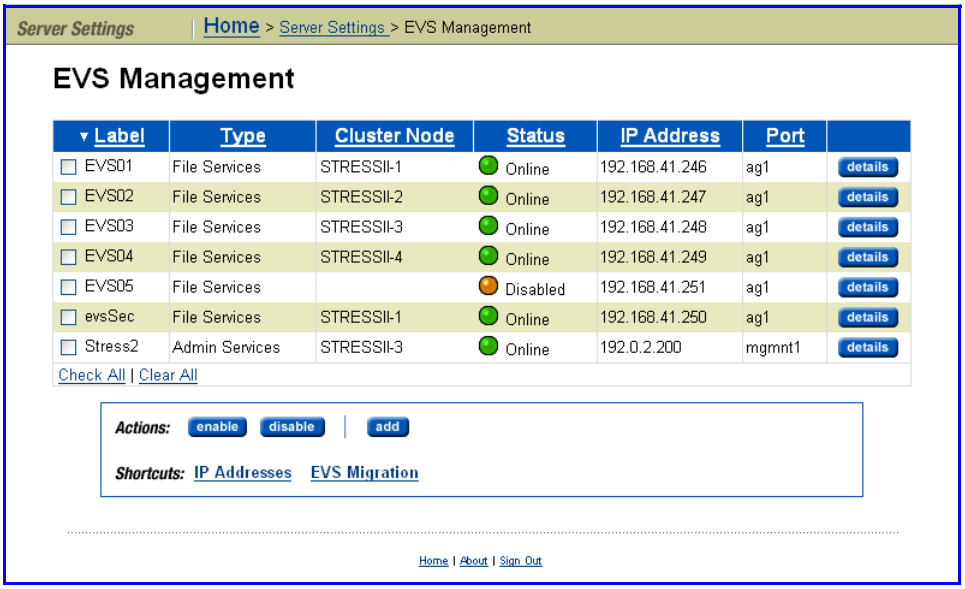

b. Select the EVS to be reactivated and click enable to re-enable the EVS.

Modifying Advanced IP Network Settings

To access additional configuration, Advanced Mode must be enabled (see Using Advanced Mode Functions, on page 29).

1. **Navigate to the Advanced IP Configuration page.**

   From the **Network Configuration** page, click **Advanced IP Configuration** to display the page:

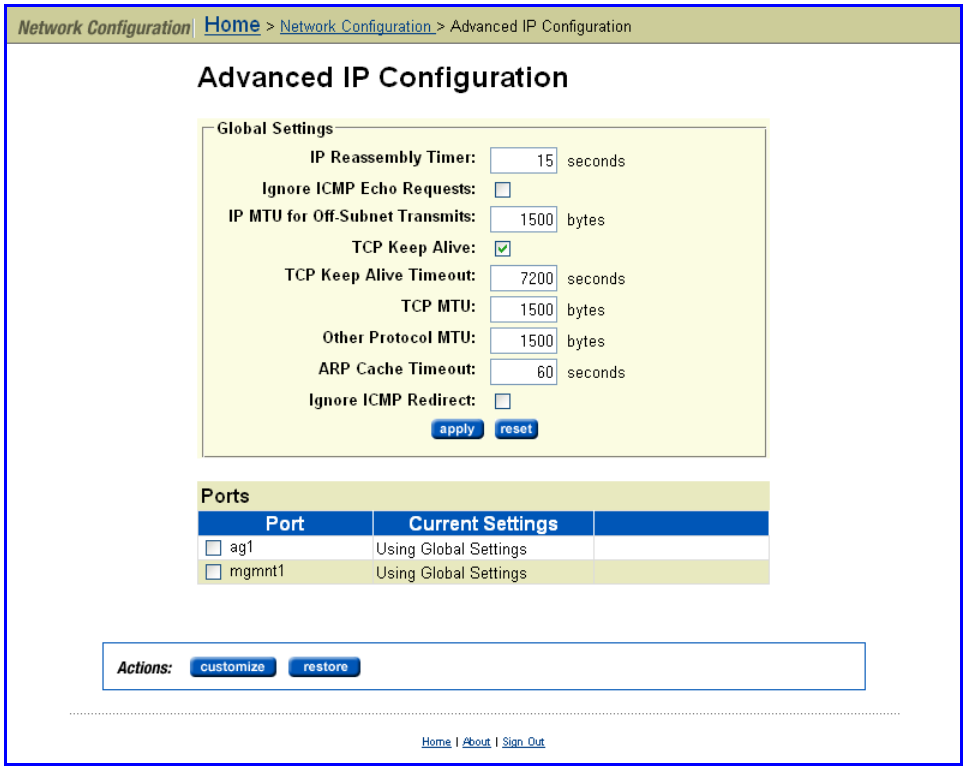

Note: In a cluster configuration, IP address settings on the **Advanced IP Configuration** page apply to all nodes in the cluster. You cannot configure nodes independently.

The **Global Settings** area contains the fields and entries that make up the global configuration, which become the default settings for all aggregations and ports.

Global Settings	Default
IP Reassembly Timer (seconds)	15
Ignore ICMP Echo Requests	No (empty)
IP MTU for Off-Subnet Transmits (bytes)	1500
TCP Keep Alive	Yes (filled)
TCP Keep Alive timeout (seconds)	7200

Global Settings	Default
TCP MTU (bytes)	1500
Other Protocol MTU (bytes)	1500
ARP Cache Timeout (seconds)	60
Ignore ICMP Redirect	No (empty)

Note: The Global Settings are applied at the server/cluster; that is, the values supplied as global settings are initially used for all aggregations (and the GE ports that make up the aggregations). Later, individual configuration settings may be defined for each defined aggregation (port) on the server/cluster.

The **Ports** lists all the aggregations and ports that have been configured, and indicates if the aggregation/port is using the global configuration or a customized configuration.

Ports	Meaning and Default
Port	Lists the name of each currently configured aggregation or port in the server/cluster. Default: `ag1 - agx`, `eth0`, `eth1`, and `mgmnt1`.
Current Settings	Indicates if the aggregation or port is using the default (global) settings, or customized settings. If the aggregation/port is using customized settings, the **details** button is also displayed. Click **details** to edit the configuration of an aggregation that is already using a customized configuration.

2. **Modify global settings:**

   For global settings, the following **Actions** are available:

   - **To customize the global settings**, specify the values you want to use for the global configuration settings by changing the values of the fields in the **Global Settings** area. All aggregations (ports) will use the global settings by default. Once you have made the changes you want in the global settings values, click **apply** to save your changes.
   - **To restore the global settings to the factory default values**, click **reset**.

   After completing the IP address configuration, you may have to reboot the server. If instructed to do so, follow the instructions to reboot the server.

3. **Customize or restore per-port settings:**

   For specific aggregations, the following **Actions** are available:

- **To customize settings** for the currently selected aggregation (the aggregation selected in the **Ports** field), click **customize** to display the **Advanced IP Per-Port Configuration** page:

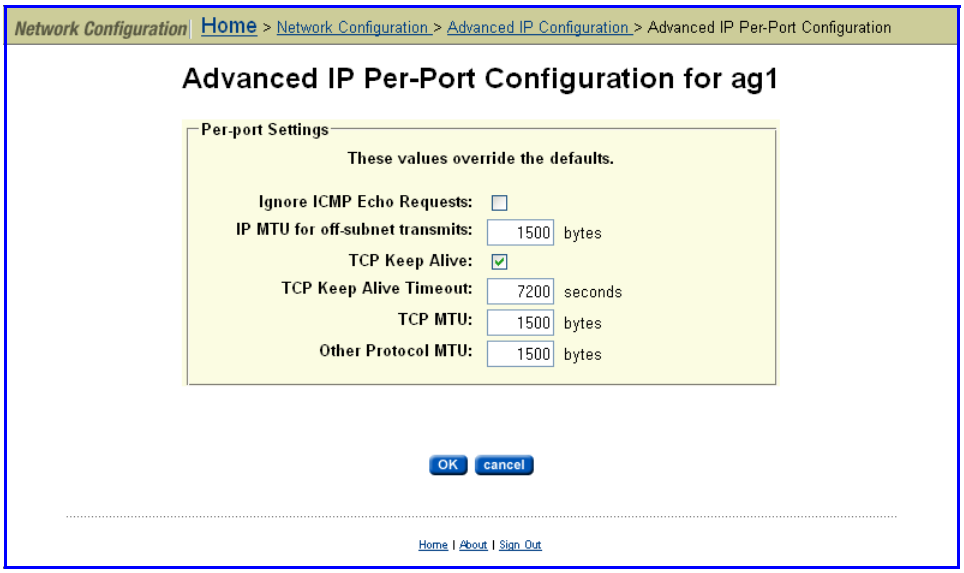

The default IP configuration settings for this page are detailed below:

Per-port Settings	Default Settings
Ports, a list of available Ethernet or aggregation ports	First port in list
Ignore ICMP Echo Requests	No (empty)
IP MTU for Off-Subnet Transmits (bytes)	1500
TCP Keep Alive	Yes (filled)
TCP Keep Alive Timeout (seconds)	7200
TCP MTU	1500
Other Protocol MTU	1500

Enter the new values in the fields, and click **OK**. The new settings will override the global settings.

- **To restore the settings of an aggregation (port) to the global configuration**, select an aggregation in the **Ports** field, then click **restore**. The settings for the aggregation selected in the **Ports** field, and all of its GE interfaces, will be erased, and will revert to the default (global settings).

- **To change the settings of an aggregation that uses a customized configuration**, click **details** to display the **Advanced IP Per-Port Configuration** page (described above).

  Enter the new values in the fields, and click **OK**. The new settings will override the global settings.

After completing the IP configuration, you may have to reboot the server. If instructed to do so, follow the instructions to reboot the server.

Managing the Server's Route Table

The server chooses the *most specific* route available for outgoing IP packets. The *host route* is the *most specific,* since it targets a specific computer on the network. The *network route* is the next most specific, since it targets a specific network. A *gateway* is the *least specific* route, hence the third routing option for the server. Therefore, if a server finds a host route for an outgoing IP packet, it will choose that route over a network route or gateway. Similarly, when a host route is not available, the server will choose a corresponding network route or, in the absence of host and network routes, the server will send the packet to a default gateway.

To manage a server's Routes Table, from the **Network Configuration** page, click **IP Routes** to display the **IP Routes** page:

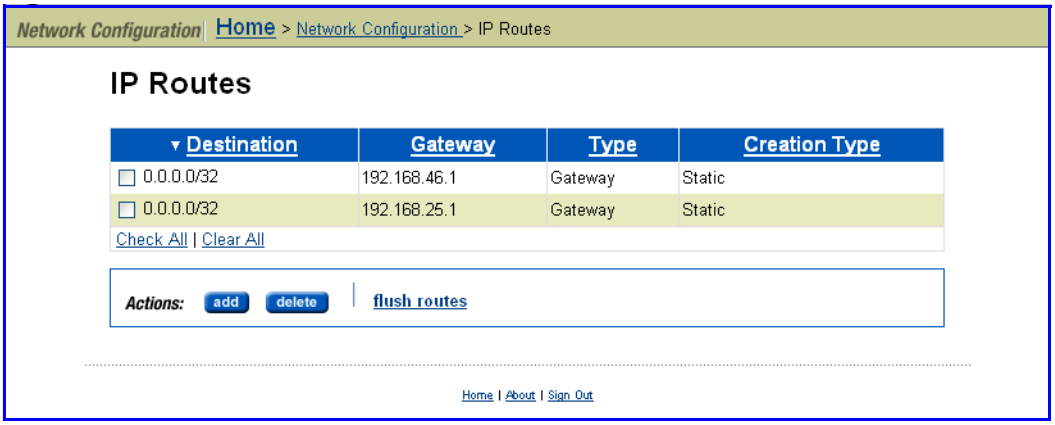

- **To delete a route:** Fill the checkbox next to the route you want to delete, and click **delete**.

  Dynamic routes cannot be deleted individually. To delete all dynamic routes, flush the cache by clicking **flush.**

- **To add a route:** Click **add** to display the **Add IP Route** page:

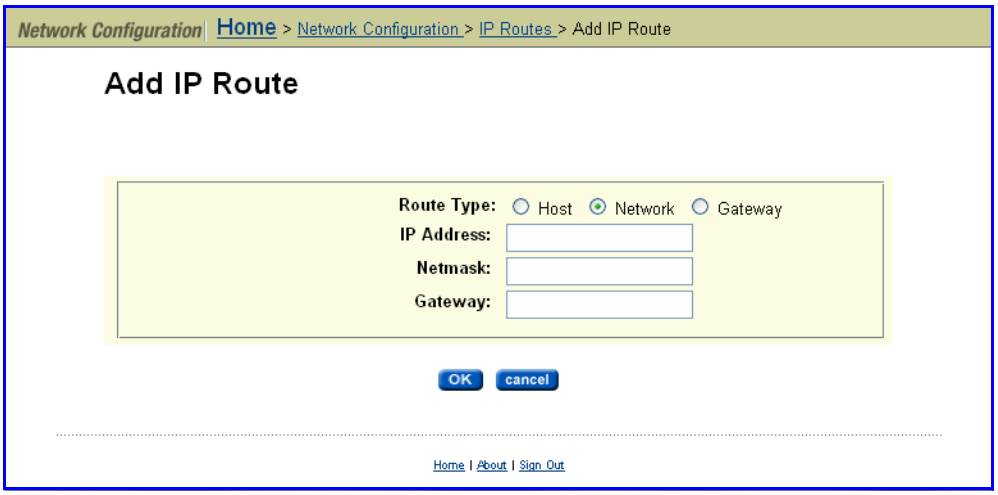

  - **For host-based static routing, select the Host radio button** and enter the IP address of the destination device and the gateway through which the host should be accessed. For host-based routes, the netmask will always be 255.255.255.255. This netmask is filled in automatically when the "host" route type is selected.
  - **For network-based static routing, select the Network radio button** and specify the target network based on the IP address and netmask; also, the gateway through which the host should be accessed.
  - **For gateways, select the Gateway radio button** and enter the IP address of the gateway in the **Gateway** field after selecting the route type. Note that the **IP** and **Netmask** fields are completed automatically.
- **To flush the route cache:** Click **flush routes**.

Configuring Name Services

Name Services configuration requires specifying and/or prioritizing name services. The following section provides information on how to complete these tasks.

Specifying and Prioritizing Name Services

To specify and prioritize name services:

1. **Navigate to the Name Services page.**

From the **Network Configuration** page, click to display the **Name Services** page:

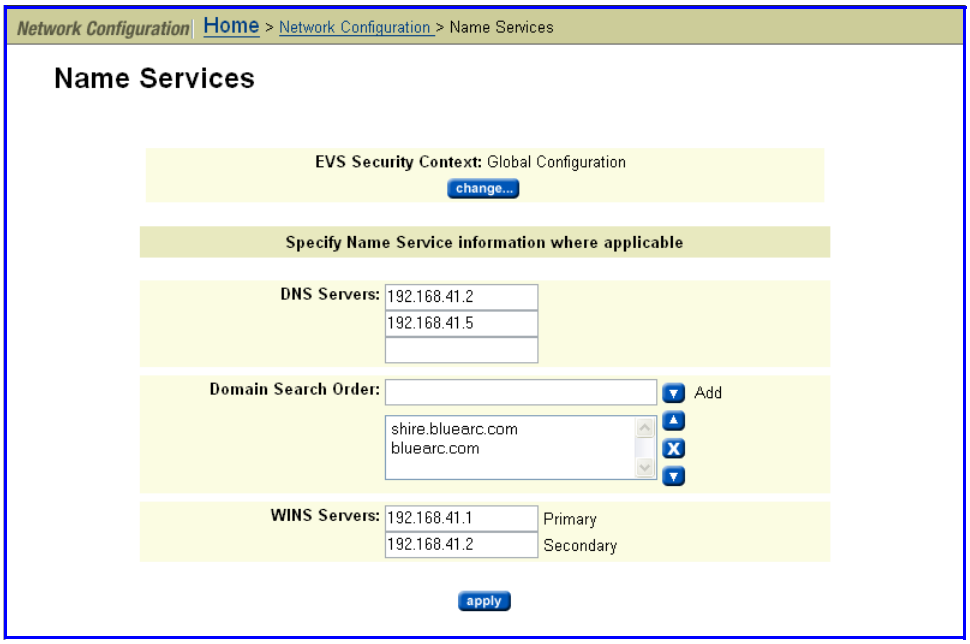

2. **Enter the requested information.**

   The following table describes the fields in this page:

Item/Field	Description
EVS Security Context	Displays the currently selected EVS security context. Changes to the name services using this page apply only to the currently selected EVS security context. • If an EVS uses the **Global Configuration**, any changes made to the global configuration settings will affect the EVS. • If an EVS uses an individual security context, changes made to the global configuration settings will not affect the EVS. To change the name services settings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context. Click **change** to select a different EVS security context or to select the global configuration.
DNS Servers	IP addresses of up to three DNS servers. If more than one DNS server is entered, the search will be performed using the DNS servers in the order listed.

Item/Field	Description
Domain Search Order	Enter a Domain suffix (e.g. *ourcompany.com*) to use as a search keyword. When searching for a computer name, the DNS server searches using suffix order. For example, if the server contains the entries *uk.ourcompany.com* and *us.ourcompany.com*, a request for the IP address of a host named *author* generates a query for *author.uk.ourcompany.com* and then for *author.us.ourcompany.com*. However, the system does not search the parent Domain *ourcompany.com*. **Note:** The suffix, combined with a computer's host name, makes up a fully qualified domain name. **Note:** **To append a suffix** to the displayed list, click **Add**. **To delete a suffix,** select it from the displayed list, then click **X**. **When using multiple domain suffixes,** select the search order for the suffixes by using the up and down arrows to change their order within the list box.
WINS Servers	To setup a primary WINS server, enter the IP address in the **Primary** WINS server field. If there is a secondary WINS server, enter the address in the **Secondary** WINS server field.

3. **Save your changes.**

   Click **apply** to save.

4. **For instances of just one name service, verify that the name service appears in the Name Services Order configuration page:**

   a. **Navigate to the Name Services Ordering page.**

   From the **Network Configuration** page, click **Name Services Order** to display the **Name Services Ordering** page, which lists **Available Name Services** and **Selected Name Services** in separate sections:

   b. **Use the change button to change the security context, if needed.**

c. **Select and deselect name services to create a list of Selected Name Services.**

   Use the left/right arrow keys to select name services from the **Available Name Services** box and move them to the **Selected Name Services** box, and vice-versa to deselect name services.

d. **Adjust the order of usage for selected name services.**

   Use the up/down arrow keys to change the order of usage for selected name services in the **Selected Name Services** box.

e. **Apply settings.**

   Verify settings, then click **OK** to apply the settings, or **cancel** to decline.

Configuring Directory Services

Directory Services configuration requires enabling services, as well as specifying directory servers, configuring, and/or prioritizing directory servers. The following sections provide information on how to complete these tasks.

Enabling and Configuring NIS and LDAP Services

This section discusses how to enable and configure NIS and LDAP services using the Web Manager.

Configuring NIS services includes the following tasks:

- Enabling and Disabling NIS
- Viewing the NIS Configuration
- Adding NIS Servers
- Modifying the NIS Configuration
- Changing the Priority of Configured NIS Servers
- Configuring LDAP to Provide NIS Services

Note: The Titan Server supports LDAP version 2.

Enabling and Disabling NIS

After deciding which network information service to use (see Name Services, on page 76), follow these instructions to enable NIS or LDAP:

1. **Navigate to the NIS/LDAP Configuration page.**

From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page:

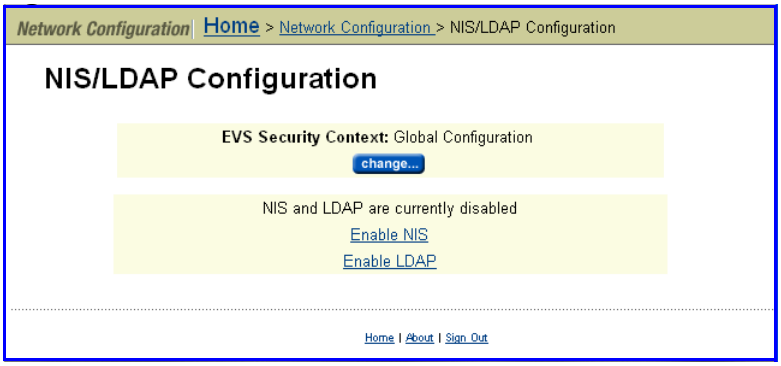

2. **Enable NIS or LDAP.**

   Click **Enable NIS** or **Enable LDAP**.

Viewing the NIS Configuration

View the NIS configuration for the current domain as follows:

1. **Navigate to the NIS/LDAP Configuration page.**

   From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page:

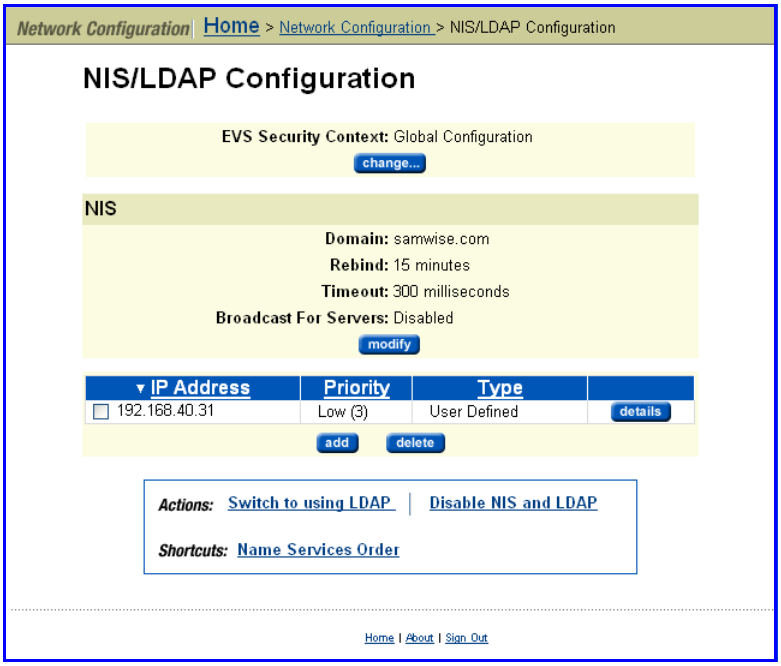

The following table describes the fields in this page:

Item/Field	Description
EVS Security Context	Displays the currently selected EVS Security Context; either an individual security context or the global security context. Click **change** to select a different EVS Security Context or to select the global configuration. Selecting a different EVS Security Context changes the context to which the NIS/LDAP configuration settings apply.
Domain	Name of the NIS Domain for which the system is a client.
Rebind	Frequency of server attempts to connect to its configured NIS servers. Enter a value from 1 to 15 minutes.
Timeout	Amount of time (in milliseconds) to wait for a response from an NIS server when checking the Domain for servers. Enter a value from 100 to 10,000 milliseconds. The default value is 300 milliseconds.
Broadcast For Servers	Enables server to discover the available NIS servers on the network. Servers must be in the same NIS domain and present on the server's network.
IP Address	Displays the IP addresses of the NIS servers which are currently configured.
Priority	Priority level for the selected NIS server (lowest value is highest priority). If the NIS Domain contains multiple servers, the system will try to bind to the server with the highest priority level whenever it performs a rebind check.  **Note:** Servers discovered by broadcast do not have a priority. If you assign a priority after clicking the **details** button, the NIS server type becomes "User Defined," and "User Defined" NIS servers are prioritized before servers discovered through broadcast.
Type	Type of NIS server.  **Note:** Servers can be automatically discovered through the **Broadcast for Servers** option. They may be defined by the user, and user defined servers, regardless of priority, are tried before servers found by broadcast.

2. **Add and delete servers, view server details and change server priorities, or modify the NIS configuration:**
   - **To add servers:** Click **add**, then refer to the instructions in Adding NIS Servers, on page 100.
   - **To delete servers:** Select a server, then click **delete.**
   - **To modify the configuration:** Click **modify**, then refer to the instructions in Modifying the NIS Configuration, on page 100.
   - **To view detailed properties and/or change server priority:** Select a server, then click **details**, and refer to the instructions in Changing the Priority of a Configured NIS Server, on page 102.

- **To switch to LDAP:** Click **Switch to using LDAP**. The change applies to all servers.
- **To disable NIS:** Click **Disable NIS and LDAP**. The change applies to all servers.
- **To modify the name services order:** Click **Name Services Order** to navigate to the **Name Services Ordering** page, where name service ordering is specified. See Changing Name Services Order, on page 107 for more information.

Adding NIS Servers

To add an NIS server:

1. **Navigate to the Add NIS Server page.**

   From the **Home** page, select **Network Configuration**, then **NIS Configuration,** then click **add** to display the **Add NIS Server** page:

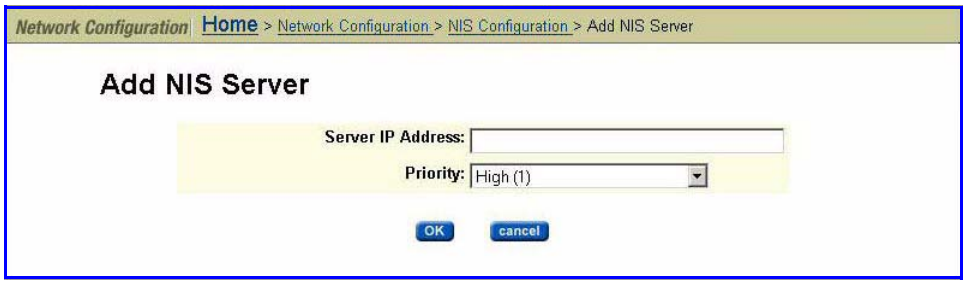

2. **Enter the requested information.**

   In the **Server IP Address** field, enter the IP address of the NIS server you want to add.

   In the **Priority** field, select a priority level for this NIS server from the drop-down list (lowest number is highest priority).

   Note: If the NIS Domain contains multiple servers, the system will try to bind to the server with the highest priority level whenever it performs a rebind check.

3. **Apply the addition of the new NIS server.**

   Click **OK.**

Modifying the NIS Configuration

To modify the NIS configuration:

1. **Navigate to the Modify NIS Configuration page.**

From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page, then click **modify** to display the **Modify NIS Configuration** page:

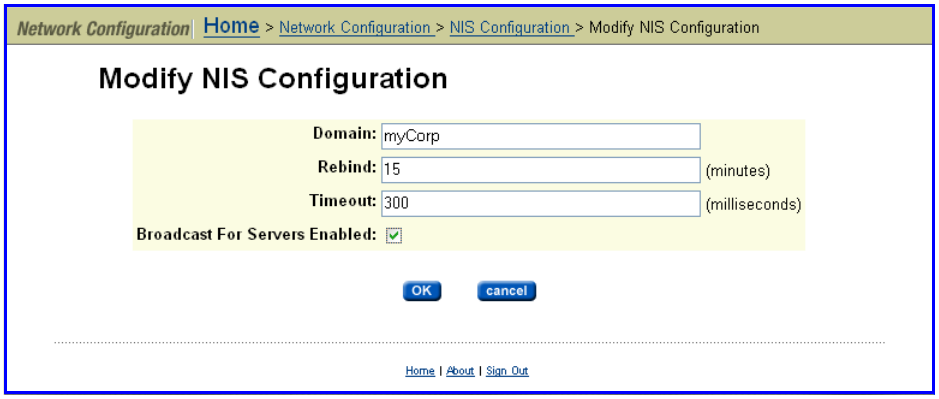

2. **Enter the requested information.**

   a. **Edit the values in the Domain, Rebind and Timeout fields.**

   b. **Enable/disable Broadcast For Servers.**

   Fill the "Broadcast for Servers Enabled" checkbox to enable the server to discover and automatically bind to NIS servers in the domain. Once enabled, the server will search for NIS servers in its configured NIS domain. These servers are found by broadcast and therefore must be on the same logical network as the server.

   Note: After a server has been found by broadcast, click **details** to configure that server. If you later clear the **Broadcast for Servers** checkbox, the server configuration is not deleted; it is retained for possible later use.

   NIS servers found by broadcast are regularly polled for responsiveness and, when a request for NIS lookup is made, the most responsive server is selected.

   To remove NIS servers found by broadcast, disable "Broadcast for Servers" (clear the "Broadcast for Servers" checkbox). If "Broadcast for Servers" is enabled, an attempt to remove NIS servers found by broadcast results in the following error message:

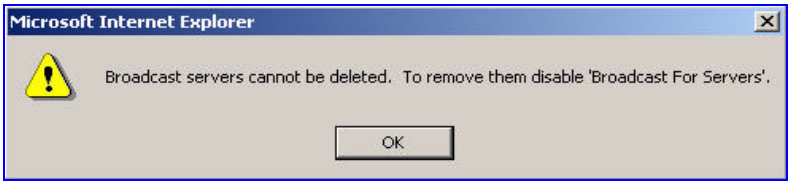

3. **Apply the configuration.**

   Verify that the configuration is correct, then click **OK** to apply its settings or **cancel** to decline.

Changing the Priority of a Configured NIS Server

To change the priority of a configured NIS server:

1. **Navigate to the Modify NIS Server page.**

   From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page, then click **details** to display the **Modify NIS Server** page:

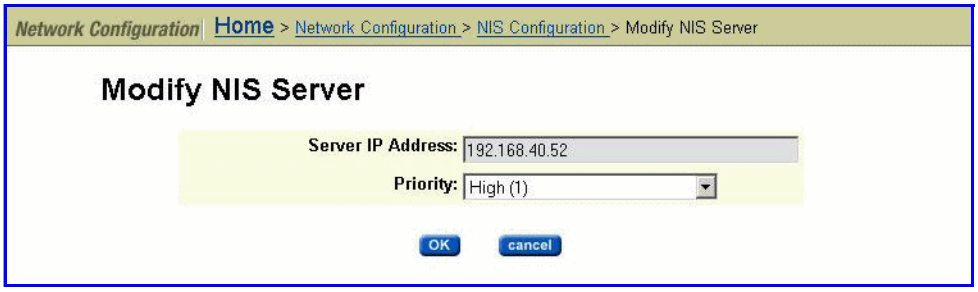

2. **Enter the requested information.**

   In the **Priority** field, select the priority level for this NIS server from the drop-down list (lowest number is highest priority).

   Note: If the NIS Domain contains multiple servers, the system will try to bind to the server with the highest priority level whenever it performs a rebind check.

3. **Apply settings.**

   Verify your settings, then click **OK** to apply or **cancel** to decline.

Configuring LDAP to Provide NIS Services

To configure LDAP to provide NIS Services:

1. **Navigate to the NIS/LDAP Configuration page.**

From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page:

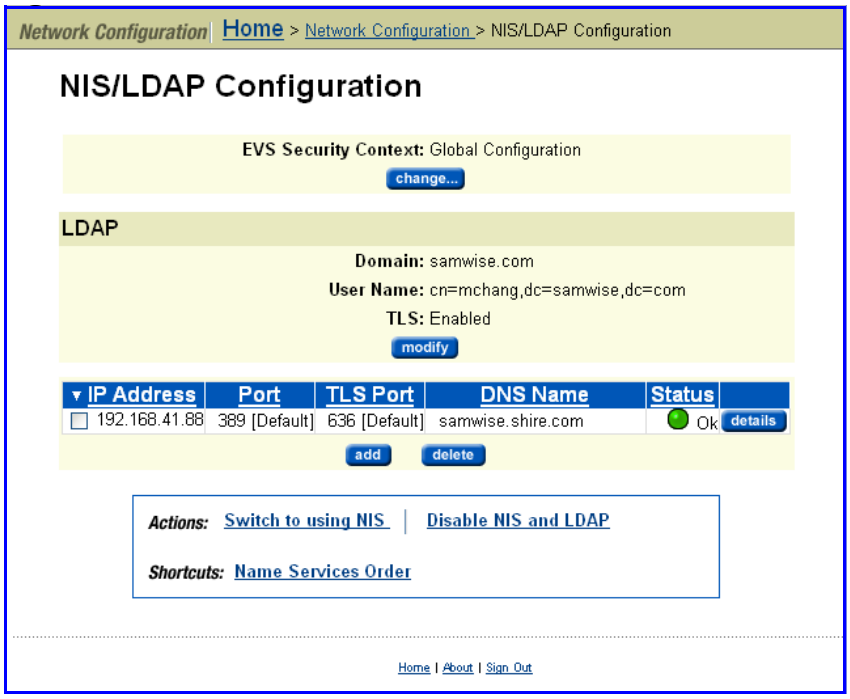

The following table describes the fields in this page:

Item/Field	Description
EVS Security Context	Displays the currently selected EVS Security Context; either an individual security context or the global security context. Click **change** to select a different EVS Security Context or to select the global configuration. Selecting a different EVS Security Context changes the context to which the NIS/LDAP configuration settings apply.
Domain	Name of the LDAP Domain for which the system is a client. For example: SGI*.com*
User Name	User name of the administrator who has rights and privileges for this LDAP server. The name can be up to 256 characters in length; however, if it includes spaces, the name must be enclosed in double quotes. For example: `cn="Directory Manager",dc=server1,dc=com`
TLS	Enable/disable the TLS and SSL connection.
IP Address	IP address of the NIS servers to which the server is currently bound.
Port	Standard port that is configurable by the administrator. The default port is 389.

Item/Field	Description
TLS Port	The secure port that is configurable by the administrator. The default port is 636.
DNS Name	Fully qualified hostname of the LDAP server.
Status	Displays the status of the LDAP server.

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes to the NIS/LDAP configuration using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the NIS/LDAP configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global NIS/LDAP configuration settings will not affect the EVS. To change the NIS/LDAP configuration settings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

   Click **Change** to select a different EVS security context or to select the global configuration.

3. **Create/modify/delete the configuration.**

   The following **Actions/Shortcuts** are available:

   - **To add servers:** Click **add**, then refer to the instructions in Adding an LDAP Server, on page 104.
   - **To delete servers:** Select a server, then click **delete.**
   - **To view or modify the configuration:** Click **details**, then refer to the instructions in Modifying the LDAP Configuration, on page 105.
   - **To view detailed properties and/or change server properties:** Select a server, then click **details** and refer to the instructions in Modifying the LDAP Server, on page 106.
   - **To switch to NIS:** Click **Switch to using NIS**. The change applies to all servers.
   - **To disable NIS and LDAP:**
     Click **Disable NIS and LDAP**. The change applies to all servers.
   - **To modify the name services order:** click **Name Services Order** to navigate to the **Name Services Order** page, where name service ordering is specified.

Adding an LDAP Server

The Titan Server supports LDAP version 2. To add an LDAP server:

1. **Navigate to the Add LDAP Server page.**

From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page, then click **add** to display the **Add LDAP Server** page:

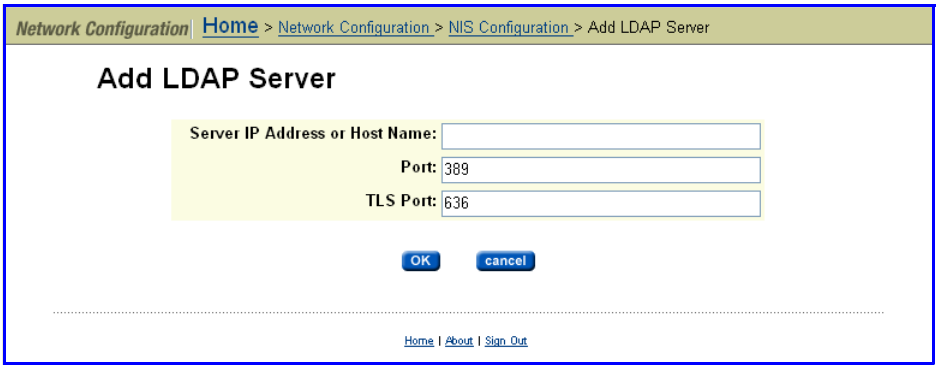

2. **In the Server IP Address or Host Name field, enter the IP address or the host name of the LDAP server.**

   Enter the IP address or a resolvable host name for the LDAP server.

3. **In the Port field, enter the new standard port number for the LDAP server.**

   The standard port used to communicate with the LDAP server. The default port is 389.

4. **In the TLS Port field, enter the new secure port number for the LDAP server.**

   The secure port used to communicate with the LDAP server. The default port is 636.

5. **Save the new LDAP server information.**

   Click **OK.**

Modifying the LDAP Configuration

To modify the LDAP configuration:

1. **Navigate to the Modify LDAP Configuration page.**

From the **Network Configuration** page, click **NIS/LDAP Configuration** to display the **NIS/LDAP Configuration** page, then click **modify** to display the **Modify LDAP Configuration** page:

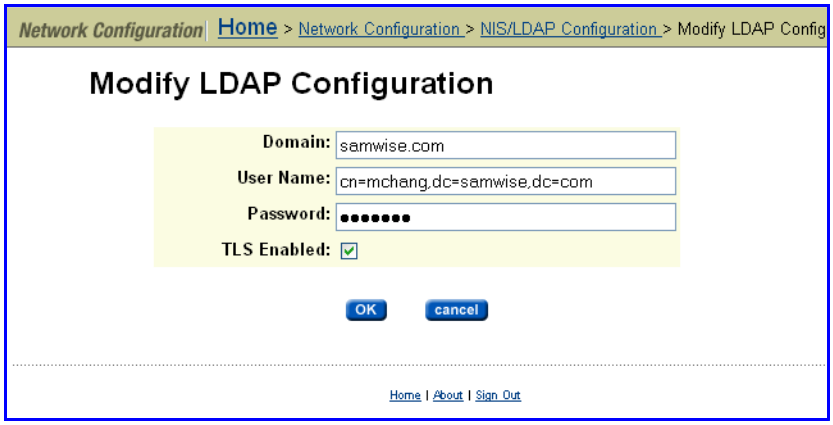

Note: This option supports both registered and anonymous login of users.

2. **Enter the requested information:**
   - Edit the **Domain, User Name** and **Password** fields.
   - Fill or clear the **TLS Enabled** checkbox to enable/disable TLS.
3. **Apply the configuration.**

   Verify that the configuration is correct, then click **OK** to apply the settings or **cancel** to decline.

Modifying the LDAP Server

To modify the LDAP server properties:

1. **Navigate to Modify LDAP Server.**

   From the **Network Configuration** page, select **NIS/LDAP Configuration,** then select the LDAP server configuration you want to change, and click **details** to display the **Modify LDAP Server** page:

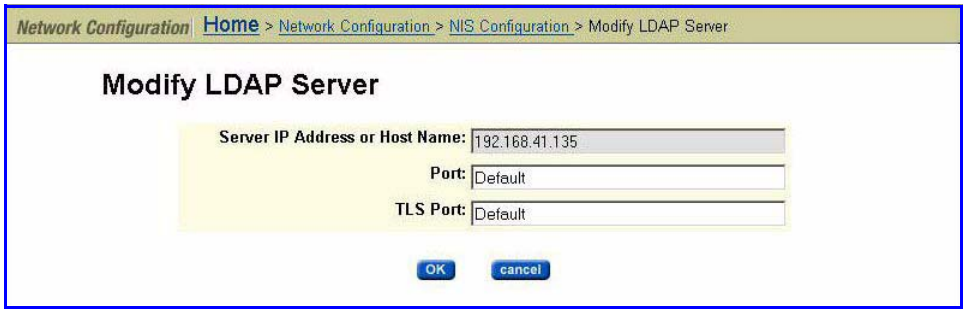

2. **Change/update current configuration information.**
   a. **In the Server IP Address or Host Name field, enter the new IP address or the new host name of the LDAP server.**

b. **In the Port field, enter the new standard port number for the LDAP server.**

   The standard port used to communicate with the LDAP server. The default port is 389.

c. **In the TLS Port field, enter the new secure port number for the LDAP server.**

   The secure port used to communicate with the LDAP server. The default port is 636.

3. **Apply the configuration.**

   Verify that the configuration is correct, then click **OK** to apply the settings or **cancel** to decline.

Changing Name Services Order

To change the order in which name services are used:

1. **Navigate to the Name Services Ordering page.**

   From the **Home** page, click **Network Configuration**, then click **Name Services Order**.

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes to the name services order using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the name services ordering settings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

Click **Change** to select a different EVS security context or to select the global configuration.

3. **Select the name services to be used.**

   From the **Available Name Services** list, select the name services you want to use, and click the right arrow.

4. **Set the name services order.**

   The **Selected Name Services** list displays the name services in the order in which they will be used. Name services higher in the list are used before services lower in the list.

   To change the position of a name service in the list, select the name service, and click the up arrow or the down arrow to change the order in which that name service will be used.

5. **Apply settings.**

   Verify your settings, then click **OK** to apply or **cancel** to decline.

6 Storage Management

Storage Management Concept	Conceptual Overview	Associated Tasks
System Drives	System Drives, on page 110	Using System Drives, on page 128
Storage Pools	Storage Pools, on page 110	Using a Storage Pool, on page 128
File Systems	File Systems, on page 111	Using File Systems, on page 138
Usage Quotas	Usage Quotas, on page 115	To Set User and Group File System Quota Defaults, on page 165
Data Migrator	SGI Data Migrator, on page 122	Configuring the SGI Data Migrator, on page 182

Overview

This chapter presents IS-NAS Server and Titan Server storage management concepts, as well as instructions for configuring and managing *file systems* (including *WORM file systems*), managing allocation of storage space using *quotas* and *virtual volumes,* and optimizing data storage using the *Data Migrator* to provide rules-based migration of data to primary and secondary storage.

Note: *File systems*, *Storage Pools,* and *system drives* are logical divisions of the physical storage subsystem.

Storage Management Components

The storage server architecture includes system drives, Storage Pools, file systems and Virtual Servers (EVSs), supplemented by a flexible quota management system for managing utilization, and the Data Migrator, which optimizes available storage. This section describes each of these storage components and functions in detail.

The following diagram illustrates a simplified view of the architecture:

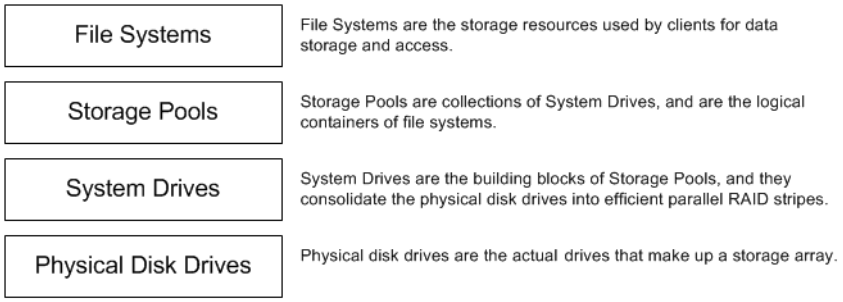

System Drives

System drives (SDs) are the basic (logical) storage element used by the server. Storage subsystems use RAID controllers to aggregate multiple physical disks into SDs. For more information about configuring and using SDs, refer to the *Storage Subsystem Guide*.

Storage Pools

A Storage Pool (known as a "span" in the command line interface) is the logical container for a collection of one or more system drives (SDs). Storage Pools can be expanded as additional SDs are created in the storage subsystem, and grow to a maximum capacity of 256 TB. Expanding a Storage Pool does not interrupt network client access to storage resources. By allocating a shared pool of storage for multiple users and allocating space dynamically (thin provisioning), a server cluster supports “over-subscription,” sharing space that accommodates the peak requirements of individual clients, saving the overhead associated with sustaining unnecessary storage. See Thin Provisioning File Systems, on page 118 for more information on thin provisioning.

A Storage Pool can hold up to 128 file systems, centralizing and simplifying management of its component file systems. For example, the settings applied to a Storage Pool can either allow or restrict the expansion of every file system in the Storage Pool.

Storage Pools are made up of multiple small allocations of storage called “chunks.” The size of the chunks in a Storage Pool is defined when the Storage Pool is created. A Storage Pool can contain up to a maximum of 16,384 chunks. In turn, an individual file system can contain up to a maximum of 1023 chunks. Planning the chunk size is an important consideration when creating Storage Pools, for two reasons:

- Chunks define the increment by which file systems will grow when they expand.
- As a file system contains a finite number of chunks, the chunk size places a limit on the future growth of file systems in a Storage Pool.

Note: To add more than one file system to a Storage Pool, a Storage Pool license is required; without this license, only a single file system is permitted. However, even without the license, Storage Pools and file systems can be expanded as long as the Storage Pool contains only a single file system.

File Systems

The file system is the main storage component of the IS-NAS Server/Titan Server. All other features on the server either directly or indirectly support the file system.

File systems have the following attributes:

- File system format; WFS-1 or WFS-2. WFS-1 is the original file system format, and is still supported by all IS-NAS Server and Titan Server servers. WFS-2 is the default file system format on IS-NAS Server servers and on the Series 3000 Titan Server servers. See File System Formats, on page 111 for more information.
- Maximum size of 256 TB, depending only on the number and size of available chunks of storage.
- Features for control and monitoring of capacity, allocation, and performance:
  - **Quotas** control the allocation of storage by client.
  - **Graphs** display traffic and usage activity.
  - **Virtual Volumes** divide a file system into discrete storage areas that appear to clients as independent file systems.
  - **Policy-based movement of data**, a feature of Data Migrator, policies control storage reallocation routines, keeping some data on high-performance storage devices while migrating other data onto low-performance, lower cost, storage devices. Policies determine which data is to be moved based on a variety of user-defined criteria.

File System Formats

The WFS-2 file system format provides an alternative file system format to the WFS-1 format (the original file system supported by the Titan Server). Note that the WFS-2 file system format is supported only on the Series 3000 Titan Server and later hardware platforms (including the IS-NAS Server).

File systems formatted as WFS-2 offer several improvements over file systems formatted using WFS-1. These benefits include:

- Improved file system resiliency following storage subsystem failures.
- Fast file system roll back to a specific checkpoint.
- Fewer disks are required to achieve a given level of performance in file systems that have a high churn rate. The churn rate is higher when file systems perform intensive file system object modifications, such as

creating new files/directories, deleting files/directories, or changing the content of files.

Note: By default, when WFS-2 is supported by the hardware platform, new file systems are formatted using the WFS-2 format. You can, however, choose to format a new file system using the WFS-1 format (see Creating a File System, on page 138and Formatting a File System, on page 151 for more information.

File System Block Size

File system block size affects performance, storage size, and the efficiency of storage utilization:

- A file system with a 32 KB block size provides higher throughput when transferring large files. However, a file system with a 4 KB block size performs better than a file system with a 32 KB block size when subjected to a large number of smaller I/O operations.
- If the file system contains many relatively small files, a 4 KB file system block size provides more efficient space utilization.

  For instance, when saving a 42 KB file:

  - In a file system with a 32 KB block size, the 42 KB file takes up two 32 KB blocks, for a total of 64 KB used (2 x 32 KB = 64 KB).
  - In a file system with a 4 KB block size, the 42 KB file takes up eleven 4 KB blocks, for a total of 44 KB used (11 x 4 KB = 44 KB).
  - In this case, the 32 KB block size wastes 22 KB of space while the 4 KB block size wastes only 2 KB of space.

One advantage of configuring multiple file systems within the same Storage Pool is that applications requiring a 4 KB block size can share storage with applications that require a 32 KB block size.

WORM File Systems

The storage server supports Write Once Read Many (WORM) file systems. WORM file systems are widely used to store crucial company data in an unalterable state for a specific duration.

Note: A license is required to use WORM file systems. Contact SGI Global Services to purchase a WORM license.

The server supports two types of WORM file systems: *strict* and *non-strict*:

- **Strict WORM** file systems cannot be deleted or reformatted and should be used once strict compliance measures are ready to be deployed.
- **Non-strict WORM file systems** can be reformatted and so should only be used for testing purposes. Should a non-strict WORM file system need to be deleted, it must first be reformatted as a non-WORM file system.

Caution: WORM file systems should be created only if you need to use the file system for regulatory compliance purposes to ensure that your company's data retention policies comply with government regulations.

WORM Characteristics

WORM file systems have several very important characteristics that differentiate them from regular file systems:

- **No changes once marked WORM.** Network clients can access files on a WORM file system in the same way they access other files. However, once a file is marked as WORM, it is "locked down." WORM files cannot be modified, renamed, deleted, or have their permissions or ownership changed. These restrictions apply to all users including the owner, Domain Administrators, and 'root.'
- **Once a WORM, always a WORM.** Once marked as WORM, the file remains a WORM file until its retention date has elapsed. In contrast, files not marked as WORM can be accessed and used just as any normal file.
- **OK to add/expand, not OK to shrink, reclaim used space or delete file system.** You can expand and add more storage to a WORM file system, but you cannot shrink a WORM file system, or reclaim unused space, and you cannot delete the WORM file system.

Read Caches

A *read cache* is a special read-only file system that stores copies of individual files outside of their local file systems, enabling a server or a node to have a cached copy of the file. When NFS v2 or NFS v3 clients submit a read request for a file in the read cache, the server/node can serve the read request from the copy in the read cache. Note that a read cache does not benefit CIFS clients, and that read caches have special characteristics and limitations. For information about read caches, refer to Read Caching, on page 412.

Controlling File System Space Usage

The server can monitor space allocation on a file system and trigger alerts when pre-set thresholds are reached; optionally, users can be prevented from creating more files once a threshold has been reached. Alternatively, the file system can be expanded either manually or automatically while online.

Two activities consume system space:

- **Live file system.** Refers to the space consumed when network users add files or increase the size of existing files.
- **Snapshots.** Refers to consistent file system images at specific points in time. Snapshots are not full copies of the live file system, and snapshot sizes change depending on the live file system. As the live file system uses more space, snapshots use more space, and as the live file system uses less space, snapshots require less space.

Note: Deleting files from the live file system may increase the space taken up by snapshots, so that no disk space is actually reclaimed as a result of the delete operation. The only sure way to reclaim space taken up by snapshots is to delete the oldest snapshot.

The server tracks space taken up by:

- The live file system

- Snapshots
- Entire file system

For each of these slices, both a *warning* and a *severe* thresholds can be configured. Although they differ from system to system, the following settings should work in most cases:

	Warning	Severe
Live file system	70%	90%
Snapshots	20%	25%
Entire file system	90%	95%

When the storage space occupied by a volume crosses the warning threshold, a **Warning** event is recorded in the event log. When the *Entire File System Warning* threshold has been reached, the space bar used to indicate disk usage turns yellow:

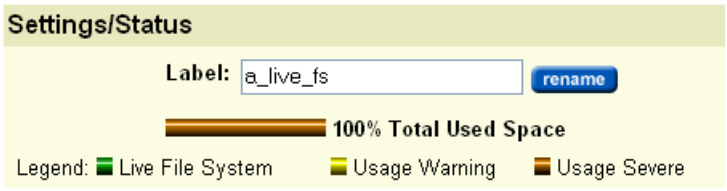

When the space reaches the *severe* threshold, a **Severe** event is recorded in the Event Log, generating corresponding alerts. If the *Entire File System Severe* threshold has been reached, the space bar used to indicate disk usage turns amber.

If file system auto-expansion is disabled, you can limit the growth of the live file system to prevent it from crossing the severe threshold, effectively reserving the remaining space for use by snapshots. To limit the live file system to the percentage of available space defined as the severe threshold, fill the "Do not allow the live file system to expand beyond its **Severe** limit" checkbox on the **File System Details** page. See To View the Details of a File System, on page 147 for more information about the **File System Details** page.

Note: To track and control space, or the number of files in the live file system, or to configure quotas for users and groups, or to create virtual volumes, see Usage Quotas, on page 115 and Virtual Volumes, on page 117.

Monitoring File System Load

The server's performance can be measured by how many operations per second (ops/sec) it is performing. Through the Web Manager, a graphic representation of the number of ops/sec can be viewed. For details on

displaying or downloading these statistics, see Server and File System Load (Ops per second), on page 488.

Increasing the Size of a File System

There are two methods to expand the amount of storage allocated to a file system:

- Manual expansion
- Automatic expansion

Note: Once storage is allocated to a file system, that storage becomes dedicated to that file system, meaning that once a file system is expanded, its size may not be reduced. Unused space in the file system cannot be reclaimed, allocated to another file system, or removed. To reclaim the storage space, the file system must be relocated to different storage (see Moving a File System, on page 120) or deleted.

Increasing the amount of storage allocated to a file system (manually or automatically) does not require that the file system be taken off line. For more information on managing the expansion of file systems, see Expanding a File System, on page 153.

Manual Expansion of a File System

Manually expanding a file system allows you to add storage capacity to a file system immediately. You specify the new size of a file system, and the storage is allocated immediately. The maximum size that a file system can attain is specified, and the file system size can be set to the maximum size supported by the Storage Pool in which the file system was created. For more information on manual file system expansion, see Expanding File Systems Manually, on page 153.

Automatic Expansion of a File System

File system auto-expansion allows a file system to grow to by adding chunks of storage on an as-needed basis. When auto-expansion is enabled, and the file system reaches approximately 80% of its allocated capacity, one or more additional chunks are allocated (see Storage Pools, on page 110 for a discussion of chunks). The maximum size that a file system can attain can be specified, or the file system size can be allowed to grow to the maximum size supported by the Storage Pool in which the file system was created. For more information on file system auto-expansion, see Expanding File Systems Automatically, on page 154.

Usage Quotas

By applying quotas, you can control disk usage to prevent network users from consuming more disk space (or creating more files) than allowed. The server supports the following types of quotas:

- **User and group quotas.** Monitor and control disk usage for individual users or groups of users.
- **Virtual volume quotas.** Monitor and control disk usage on a per-directory basis. Virtual volumes allow management of directory tree usage independently of users or groups; furthermore, user and group quotas can be created within the virtual volumes.

Note: In this section, the terms *user* and *group* indicate NFS or CIFS users and groups.

Understanding Quotas

Quotas track the number and total size of all files. When these reach specified thresholds, the system sends an email alert to each recipient on the list of contacts associated with the file system and, optionally, logs *Quota Threshold Exceeded* events. Operations that would take the user or group beyond the configured limit can be disallowed by setting hard limits.

Note: Where both *Usage* and *File Count* thresholds are defined, the server enforces whichever is reached first.

Quota Thresholds

The configuration settings defining quota restrictions are called thresholds, and are described in the following table:

	Space Usage	Number of Files
Soft Limit	Total size of all files should not exceed this value.	Total number of files should not exceed this value.
	If a *Soft Limit* is exceeded, the operation will be allowed, but an alert will be issued.	
Hard Limit	The server blocks any operation that may cause a *Hard Limit* to be exceeded.	
Warning	When total size of all files reaches this value, an *Information* alert is issued.	When total number of files reaches this value, an *Information* alert is issued.
Severe	When total size of all files reaches this value, a *Warning* alert is issued.	When total number of files reaches this value, a *Warning* alert is issued.
Reset	Alerts are disabled once a certain threshold is crossed and an alert is issued; no other alerts are issued until a reset level (threshold) is reached, after space (or a number of files) is recovered on disk. This means that the server does not continually issue alerts stating that a threshold has been crossed. *Quota* alerting is re-enabled once the used space (or number of files) drops a certain amount below the threshold. The default value for this reset is 5% of the limit.	

The following caveats apply in measuring the file system status against quota thresholds:

- **Metadata and snapshot files.** While quotas track used disk space and the number of files, neither file system metadata nor snapshot files count towards the quota limits.
- **File size calculation.** File sizes are computed based on the number of used file system blocks; for example, with a 32 KB file system block size, a 55 KB file will get reported as 64 KB.
- **Symbolic link calculation.** Files with multiple hard links pointing to them are included only once in the quota calculation. A symbolic link adds the size of the symbolic link file to a quota and not the size of the file to which it links.

Types of File System Quotas

There are two types of file system quotas:

- **Explicit User/Group Quotas.** A quota explicitly created to impose restrictions on an individual user or group, defining a unique set of thresholds.
- **Default User/Group Quotas.** A quota set automatically for each user or group that does not have explicit quotas. This is done by defining a set of *quota defaults* (thresholds), which will apply a quota automatically when a file is created or modified:
  - **User Quota Defaults.** A set of thresholds used to create a quota *for a user* the first time that user saves a file in the file system.
  - **Group Quota Defaults.** A set of thresholds used to create a quota *for a group* the first time a user in that group saves a file in the file system.

  Initially, all quota defaults are not set. When activity occurs in the file system, it is tracked, but quotas are not automatically created. When at least one quota default threshold is set to a non-zero value, a user or group quota (as appropriate) will be created for the owner of the directory at the root of the file system.

Virtual Volumes

A file system can be divided into discrete areas of storage called virtual volumes. From a client's perspective, a virtual volume appears to be a normal file system. Virtual volumes provide a simple method for allocating and controlling directories for projects, users, or groups. Capacity and number of files within the virtual volume can be controlled using quotas.

The server treats the virtual volume 'root' directory, together with all its sub-directories, as a self-contained file system. The virtual volume tracks its usage of space and number of files, to provide a way of monitoring file system usage. This tracking allows quotas to be imposed on disk space usage, as well as the total number of files.

Quotas can be imposed for the entire virtual volume, including individual users, and groups of users. When many users or groups will have access to a virtual volume, a set of quota defaults can be defined. In the absence of explicit user or group quotas, default quotas apply.

Note: In this section, the terms *user* and *group* are used to indicate NFS or CIFS users and groups.

Virtual volumes have the following characteristics:

- **Name:** A name or *label* by which the virtual volume is identified. This will often be the same as a CIFS share or NFS export rooted at the virtual volume's root directory.
- **File System:** The file system in which the virtual volume is created.
- **Path:** The directory at the *root* of the virtual volume.
- **Email Contacts:** A list of email addresses, to which information and alerts about virtual volume activity are sent. The list can also be used to send emails to individual users.

Note: When it comes to *moving* or *hard linking* files, virtual volumes behave the same as real file system volumes. Moving or linking files across different virtual volumes returns a *cross volume link* error. For a move operation, most CIFS or NFS clients will suppress this error and, instead, will copy the files to the target location and delete the original ones.

Important Information about Virtual Volumes

The following caveats apply in measuring the virtual volume status against quota thresholds:

- **Metadata and snapshot files.** Neither file system metadata nor snapshot files count towards the quota limits.
- **Symbolic link calculation.** Files with multiple hard links pointing to them are included only once in the quota calculation. A symbolic link adds the size of the symbolic link file to a virtual volume and not the size of the file to which it links.

Thin Provisioning File Systems

Thin provisioning is a method of controlling how a file system's free space is calculated and reported. Administrators use thin provisioning to optimize the utilization of storage and to plan resource acquisition in a way that helps minimize expenses, while ensuring that there is enough storage for all users' needs.

Thin provisioning allows you to oversubscribe the storage connected to the storage server. As long as the available storage is not completely allocated to file systems, the oversubscription cannot be noticed by storage system users.

Note: When thin provisioning is enabled and storage is oversubscribed, if a client attempts a write operation and there is insufficient storage space, the

client will receive an insufficient space error, even though a query for the amount of free space will show that space is still available. When storage is oversubscribed, the storage server's administrator must ensure that this situation does not occur; the storage server does not prevent this situation from occurring. To resolve this situation, the storage server's administrator must either disable thin provisioning or add storage.

When thin provisioning is enabled, the storage server reports the amount of free space for a file system based on the file system's expansion limit (its maximum configured capacity), rather than on the amount of free space based on the amount of storage actually allocated to the file system. Because file systems can be allowed to automatically expand up to a specified limit (the expansion limit), additional storage is allocated to the file system as needed, instead of all the storage being allocated to the file system when it is created. For more information about file system auto-expansion, see Increasing the Size of a File System, on page 115.

For example, a file system has an expansion limit of 20TB, with 6TB already used and 8TB currently allocated. If thin provisioning is enabled, the server will report that the file system has 14TB of free space, regardless of how much free space is actually available in the Storage Pool (for more information about Storage Pools, see Storage Pools, on page 110). If thin provisioning is disabled, the server will report that the file system has 2TB of free space.

By default, thin provisioning is disabled for existing file systems and for newly created file systems. Enable and disable thin provisioning using the `filesystem-thin` command (currently there is no way to enable or disable thin provisioning through the Web Manager GUI).

Thin provisioning works on a per file system basis, and does not affect reports from the `span-list --filesystems` and `filesystem-list` commands. Also, Web Manager (the GUI) displays the actual file system size. As a result, the administrator can perform proper capacity planning.

When enabled, thin provisioning information is returned by the following CLI commands:

- `cifs-share list`
- `df`
- `filesystem-limits`
- `filesystem-list –v`
- `fs-stat`
- `nfs-export list`
- `query`

For more information about CLI commands, refer to the Titan Server *Command Line Reference*.

If thin provisioning is enabled and you disable file system auto-expansion for a Storage Pool, the free space reported for each of the file systems in that

Storage Pool is the same as if thin provisioning were not enabled. This means that the free space reported becomes equal to the difference between the file system's current usage and the amount of space in all Storage Pool chunks currently allocated to that file system. If you re-enable file system auto-expansion for file systems in the Storage Pool, free space is again reported as the difference between the file system's current usage and its expansion limit, if an expansion limit has been specified. For more information on file system auto expansion, see Increasing the Size of a File System, on page 115.

Note: When thin provisioning is enabled, and the aggregated file system expansion limits of all file systems exceeds the amount of storage connected to the server/cluster, warnings are issued to indicate that storage is oversubscribed. These warnings are issued because there is an insufficient amount of actual storage space for all file systems to grow to their expansion limit.

Moving a File System

Moving a file system (or several file systems) may be necessary to improve performance or balance loads, to move data to different storage resources, to support changing network topography, or other reasons.

There are two basic methods of moving a file system:

- **File System Relocation**

  File system relocation changes the EVS (virtual server) that hosts the file system, but it **does not move file system data**. Moving the file system from one EVS to another changes the IP address used to access the file system, and also changes CIFS shares and NFS Exports for that file system. For information on how to relocate a file system using File System Relocation, see Relocating a File System, on page 154.

  If the file system to be relocated is linked to from within a CNS, and clients access the CNS using a CIFS share or an NFS export, the relocation can be performed with no change to the configuration of network clients. In this case, clients will be able to access the file system through the same IP address and CIFS share/NFS export name after the relocation as they did before the relocation was initiated. For more information on CNS, see Cluster Name Space (CNS), on page 410.

  Caution: Whether or not the file system resides in a CNS, relocating a file system will disrupt CIFS communication with the server. If Windows clients require access to the file system, the file system relocation should be scheduled for a time when CIFS access can be interrupted.

- **Transfer of primary Access**

  A transfer of primary access is a replication-based method of copying data from a portion of a file system (or an entire file system) and relocating the access points for that data (copying the data and metadata). A transfer of primary access causes very little down time, and the file system is live and servicing file read requests during most of the relocation process. For a short period during the relocation process, access is limited to read-only.

For more information on relocating a file system using Transfer of Primary Access, see Transfer of Primary Access, on page 337.

The method you use to relocate a file system depends, in part, on what you want to move, and what you want to accomplish by relocating the file system.

- If you want to **move the file system's access points, but not the actual data**, using file system relocation is the most appropriate method.
- If you want to **move the file system's data and access points**, using a Transfer of Primary Access is the most appropriate method.

File System Relocation

Before it can be shared or exported, a file system must be associated with a Virtual Server (EVS), thereby making it available to network clients. The association between a file system and an EVS is established when the file system is created. Over time, evolving patterns of use and/or requirements for storage resources may make it desirable to relocate a file system to a different EVS.

Note: Read caches cannot be relocated.

A **file system hosted by an EVS on a cluster node** *may* be relocated to:

- An EVS on the same cluster node, or
- An EVS on a different node in the same cluster.

but *may not* be relocated to:

- An EVS on a stand-alone server, or
- An EVS on a node of a different cluster.

A **file system hosted by an EVS on a stand-alone server** *may* be relocated to

- An EVS on the same server

but *may not* be relocated to:

- An EVS on a different server, or
- An EVS on a node in a cluster.

Typically, File System Relocation is used to move a file system from an EVS on a cluster node to an EVS on a different cluster node in order to improve throughput by balancing the load between cluster nodes.

File system relocation performs the following operations:

- Re-associates the file system with the selected EVS.
- Transfers explicit CIFS shares of the file system to the new EVS.
- Transfers explicit NFS exports of the file system to the new EVS.
- Migrates FTP users to the new EVS.
- Migrates snapshot rules associated with the file system to the new EVS.

- Migrates the iSCSI Logical Units and targets.

File system relocation may require relocating more than just the specified file system. If the file system is a member of a data migration path, both the data migration source file system and the target file system will be relocated. It is possible for the target of a data migration path to be the target for more than one source file system. If a data migration target is relocated, all associated source file systems will be relocated as well.

If more than one file system must be relocated, a confirmation dialog will appear indicating the additional file systems that must be moved. Explicit confirmation must be acknowledged before the relocation will be performed.

File System Relocation will affect the way in which network clients access the file system in any of the following situations:

- The file system is linked to from the CNS tree, but is shared or exported outside of the context of the CNS.
- The cluster does not use a CNS.

In each of the above cases, access to the shares and exports will be changed. In order to access the shares and exports after the relocation, use an IP address of the new EVS to access the file service.

Relocating file systems that contain iSCSI Logical Units will interrupt service to attached initiators, and manual reconfiguration of the IP addresses through which targets are accessed will be required once the relocation is complete. If relocating a file system with Logical Units is required, the following steps must be performed:

- Disconnect any iSCSI Initiators with connections to Logical Units on the file system to be relocated.
- Unmount the iSCSI Logical Unit.
- Relocate the file system as normal. This procedure is described in detail in Relocating a File System, on page 154.
- Re-connect the new Targets with the iSCSI Initiators. Be aware that the Targets will be referenced by a new name corresponding to the new EVSs.

Note: All iSCSI Logical Units on a target must be associated with file systems hosted by the same EVS.

SGI Data Migrator

The IS-NAS Server and the Titan Server support multiple storage technologies, with different performance capacity and cost characteristics. In order to take full advantage of tiered storage, data should be organized using a tiered hierarchy of importance and need. Data Migrator makes it easier to move data among different tiers of storage.

There are five key reasons to use Data Migrator with the server:

1. **Cost-Efficient Storage Utilization:** Using Data Migrator, newer or routinely accessed data can be retained on primary storage, while older, less-accessed, or less performance-critical data is migrated to cost-efficient secondary storage.

2. **Easy Policy-Based Configuration:** Data Migrator uses logical policies that invoke simple building blocks of rules to classify files as available for migration. Its rules and pre-conditions, which can include a file's size, type, access history, creation date, or owner, among other criteria, can make files available for migration.

3. **Discreet Migration:** Migrations from primary to secondary storage are handled as automated background tasks with minimal impact on server performance. While migrations are in progress, all data can continue to be accessed normally.

4. **Client Transparency:** Files migrated off primary storage are replaced by a link. The link looks and functions identically to the original file, while using only 1 KB of space. When the link is accessed, the contents of the associated file are retrieved transparently from their location on secondary storage. To the client workstation, they appear indistinguishable.

5. **Maximizing Storage Efficiency through Migration Reports:** Migration reports are created at the end of each migration cycle. These reports detail file usage and space consumption patterns, revealing opportunities to create more aggressive migration policies, freeing up more primary space. Further migration possibilities can be gauged by scheduling Data Migrator test runs where reports can be produced without an actual migration taking place.

Note: A license is required to use Data Migrator. Contact SGI Global Services to purchase a Data Migrator license, and see Managing License Keys, on page 535 for information on adding a license key.

Data Migration Paths

Before any Data Migration is run, the location of the migration target must be defined by creating a data migration path. A data migration path is a long term relationship between a migration source (which may be a file system or a virtual volume) and a migration target which may be a local file system, a set of file systems, a remote location, or a list of locations. (See Types of Migration Targets, on page 124 for details.) Once a migration path has been used, it may not be deleted until files migrated through that path have been deleted.

The advantages of having this long term relationship between migration source and target are as follows:

1. Other system software can stop file systems being destroyed when they are actively used by a migration path. This avoids migrated files becoming inaccessible.

2. Where snapshots may be taken on the target, synchronized source and

target snapshots may be taken to maintain snapshot images of migrated files.

3. When recovering (from tape) or replicating a file system that included migrated data, data which was originally migrated may be placed back on the migration target.

If using virtual volumes individually as migration sources within migration paths, the file system containing the virtual volumes may not be used as a migration source itself. Currently, it is only possible to define one migration path for a given migration source.

Types of Migration Targets

Data Migrator can move data to secondary storage attached directly to the storage server/cluster (a local migration), or to secondary storage attached to an external server that is not connected to the storage server/cluster (a remote migration).

Local migrations provide the benefits described previously, and remote migrations extend the functionality of Data Migrator to allow storage administrators to free up local storage resources by migrating data to storage attached to a remote NFSv3 server. Data may also be migrated to a remote server for a variety of other reasons, including archival, deduplication, or policy-based retention, compliance, and access control. As with files migrated locally, when a client tries to read a file migrated to a remote server, the file is retrieved and sent to the client, so that there is no indication to the client that the file is not in their local file system.

Note: A single migration path or operation can be to local targets **or** remote targets, but not both local **and** remote targets.

Remote migrations are controlled by user defined policies, just like the policies created for local migrations; only the paths to the secondary storage are different. Local migrations have paths to secondary storage that is attached to the same server/cluster that hosts the primary file system, while remote migrations have external paths (the secondary storage is attached to a remote server).

Cross Volume Links in Data Migrator

Data Migrator allows you to move data from primary to secondary storage based on policies that you define. When a file is migrated, a cross volume link is left, indicating the new location of the file's data. A cross volume link is a special file on the local file system that "points" to the file on a remote file system. The cross volume link stores the migrated file's metadata and is used to construct a handle for the file on the remote file system.

When a read request for a migrated file is received, the storage server reads the cross volume link, constructs the file handle, retrieves the file from the secondary storage, and sends the file to the requesting client. In this way, the read request is serviced and the client need not be aware that the file is not actually stored on the local file system.

The original cross volume link format (CVL-1) required that file systems must be bound to the same EVS. Enhanced cross volume links offer the following benefits:

- Remote file systems may be on a storage device attached to a remote server (not necessarily another IS-NAS Server or Titan Server) accessible through the NFSv3 protocol. This capability, called remote migration, allows the storage server to migrate files to a separate storage device, such as content archival or compliance products.

  

  Access to files located on the external storage device is a licensed feature, requiring an external volume link (XVL) license and a Data Migrator license. See Managing License Keys, on page 535 for information on adding a license key.

- Local access to migrated file attributes increases performance for getattr/ lookup/readdir+ requests.

  For original cross volume links, some file attributes for a migrated file are stored on the local file system and some are stored on the remote file system. When an operation requires the attributes of a migrated file, the storage server combines locally stored attribute information with attributes it retrieves from the remote file system.

  For enhanced cross volume links, all file attributes for a migrated file are stored on the local file system. When an operation requires the attributes of a migrated file, the storage server responds with locally stored attribute information, which provides better performance.

  In addition, local read caching can be used to improve performance when accessing migrated files. See Read Caching, on page 412 for information on local read caching.

- Enhanced cross volume links provide improved quota management.

  With original cross volume links, file length is tracked on the remote file system. As a result, quotas are tracked independently on the local file system and on the remote file system. By storing attributes on the local file system, enhanced cross volume links make it possible to track quotas entirely on the local file system, because file space and file count quotas are managed and calculated using local attributes. This simplifies quota management, but does not allow storage administrators to set up separate quotas for data based on the data's location.

  As a result of differences in how quotas are calculated when using original cross volume links or enhanced cross volume links, mixing of cross volume link formats is not supported within the same file system. By default, original cross volume links are created for local migrations, and enhanced cross volume links are created for all remote (external) migrations, but you can configure the storage server to create only original cross volume links. Contact SGI Global Services for more information about limiting the creation of cross volume links to the original format.

Cross Volume Link Format Considerations

When choosing which format of cross volume link to use, the following are important considerations:

- Files migrated to storage devices attached to remote servers (remote migrations) always use enhanced cross volume links.
- Files migrated locally (to storage attached to the same EVS) may use either original cross volume links or enhanced cross volume links, but enhanced cross volume links are used by default.

  If the file system has files that were migrated in previous releases, original cross volume links should be used, because original and enhanced cross volume links may not be mixed within the same file system. You can, however, convert original cross volume links to enhanced cross volume links using the `cvl-convert` command. If migrations are being performed on this file system for the first time, you should use enhanced cross volume links.
- When NDMP encounters a cross volume link to an external server (an enhanced cross volume link), it includes the link in the backup stream, but does not include the data.

Data Migrator Considerations

The server uses Data Migrator with the following considerations:

- **Snapshots and local migrations.** If files are migrated locally (to storage attached to the same EVS), when snapshots are created on the primary file system, corresponding snapshots are automatically created on the secondary file system. This preserves snapshot protection on migrated files. Likewise, when a snapshot is deleted on the primary file system, the corresponding snapshot on the secondary file system will also be automatically deleted.

  When attempting to access a locally migrated file through a snapshot on primary storage, the server will look for the corresponding snapshot on secondary storage and retrieve the migrated data from that snapshot. If the secondary file system does not contain any snapshots, the file contents will be retrieved from the live file system.
- **Snapshots and remote migrations.** If files are migrated to storage attached to a different server (a remote migration), when snapshots are created on the primary file system, corresponding snapshots are **not** created on the secondary file system.

  To preserve snapshot protection on migrated files for remote migrations, you must ensure that snapshots are taken of the storage attached to the remote server. Snapshots on the secondary file system are not managed, used, or accessed by the storage server.

  When a snapshot is accessed, and the snapshot contains a file system with a cross volume link, no special processing of the cross volume link is performed if the file in the snapshot is equivalent to the live file. If the file in the live file system has been modified since the snapshot was taken (if it

differs from the file in the snapshot), attributes from the file in the snapshot are returned for getattr/lookup/readdir+ requests, but an error is returned for read requests.

- **Virtual volumes.** If files are migrated locally, either enhanced cross volume links or original cross volume links may be used depending on your configuration. When files are migrated to a remote server, enhanced cross volume links are always used.
  - **If enhanced cross volume links are used**, virtual volumes are not recreated at all on the secondary storage.
  - **If original cross volume links are used**, virtual volumes are present on primary storage, they will be automatically recreated on the secondary storage when the data is moved during the first scheduled run of the data migration policy.
- **Quota space tracking.** Quotas are enforced only on the file system or virtual volume on which they were created. When a file is migrated through Data Migrator, however, the contents are moved from one file system to another file system or virtual volume, which may be on a remote server. Cross volume links are used to link the data from its original location to its new location. Quota tracking is different based upon the type of cross volume link being used:
  - **When enhanced cross volume links are used**, such as when files are migrated to a file system on a remote server, quotas are tracked just as if the file had remained in its original location. Quotas are tracked entirely on the local file system, because file space and file count quotas are managed and calculated using local attributes. This simplifies quota management, but does not allow storage administrators to set up separate quotas for data based on the data's location.
  - **When original cross volume links are used**, such as when files are migrated to another file system or virtual volume on the same server/cluster, quotas on primary storage are only effective on files that have not been migrated. To track space utilization of migrated data, quotas must be manually defined on secondary storage. Quota restrictions on virtual volumes cannot be set until after the policy has been completed.
- **Backup, restore, and replication of migrated Files.** While backing up a migrated file, NDMP will backup the entire contents of the file by retrieving it from secondary storage. Additionally, the backed up file will be identified as having been a migrated file. In this way, if the file is restored to a file system or virtual volume that has been configured as primary storage in a data migration path, the contents of the file will automatically be restored to secondary storage, leaving a cross volume

link on the primary storage. If the restore target is not part of a data migration path, the file will be restored in its entirety.

Alternatively, the NDMP environment variable `NDMP_BLUEARC_EXCLUDE_MIGRATED` can be used to prevent migrated data from being backed up. This can also be useful if the effective data migration policies are configured to migrate non-critical files such as music and video files from home directories or aged data. It can also improve backup and replication time, and isolate the backup data set to include only the critical information on primary storage.

You can back up a file system that is the target of a data migration. This is accomplished by performing backup of the primary file system, and selecting an option to back up only the files that have been migrated to the secondary file system. This functionality is controlled via the `NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED` NDMP environmental variable, which does the opposite of the `NDMP_BLUEARC_EXCLUDE_MIGRATED`. See NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED, on page 554 and NDMP_BLUEARC_EXCLUDE_MIGRATED, on page 551 for more information.

It is important to remember that Data Migrator extends the maximum available capacity of primary storage by migrating data to secondary storages tor age. This means that the capacity of the backup solution, whether tape libraries or a replication target, must also support the new maximum available capacity. To maintain a reliable backup and recovery system, ensure that the capacity of the deployed backup solution is at least equal to the combined capacity of Primary and secondary storage. Alternatively, use `NDMP_BLUEARC_EXCLUDE_MIGRATED` to isolate the backup dataset to only those files that are hosted natively on primary storage.

- **Files with hard links.** Files with hard links are not migrated.
- **Migrated file access.** Files that have been migrated should not be accessed directly by clients on the secondary file system. All access to migrated files should be done through the storage server.

Using System Drives

Creation and management of System drives (SDs) occurs at the storage subsystem level and subsystem-specific procedures are provided in the *Storage Subsystem Guide*.

Using a Storage Pool

Storage Pools contain one or more file systems, which consume space from the Storage Pool upon creation or expansion. A Storage Pool can also be used to control the auto-expansion policy for all of the file systems created in the

Storage Pool. The following procedures describe how to create, delete, expand, remove from service, and rename a Storage Pool.

Note: Once access is allowed to one System Drive (SD) in a Storage Pool, that Storage Pool becomes visible in the Web Manager. If access is denied to all SDs in a Storage Pool, the Storage Pool is not visible in the Web Manager.

Creating a Storage Pool

Given availability of SDs, administrators can create a Storage Pool at any time. At the time of creation, a Storage Pool can contain up to 32 SDs; later, Storage Pools can expand by adding more SDs until the total size of the pool has reached up to 256 TB. For instructions, see Expanding a Storage Pool, on page 132.

To attain optimal performance with a new Storage Pool, use as many disk drives as possible and keep SDs as large as possible. After the Storage Pool has been created, smaller file systems can be created in the pool for more granular storage provisioning. This method should also be applied when expanding a Storage Pool.

To create a Storage Pool:

1. **Navigate to the Storage Pools page.**

   From the **Storage Management** heading, click to display the **Storage Pools** page:

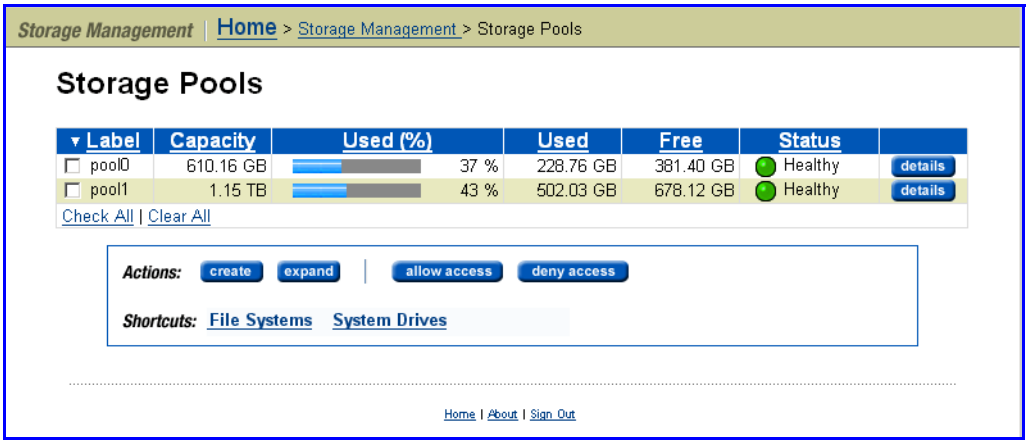

2. **Open the Storage Pool wizard.**

Click **create** to view the **Storage Pool Wizard** page:

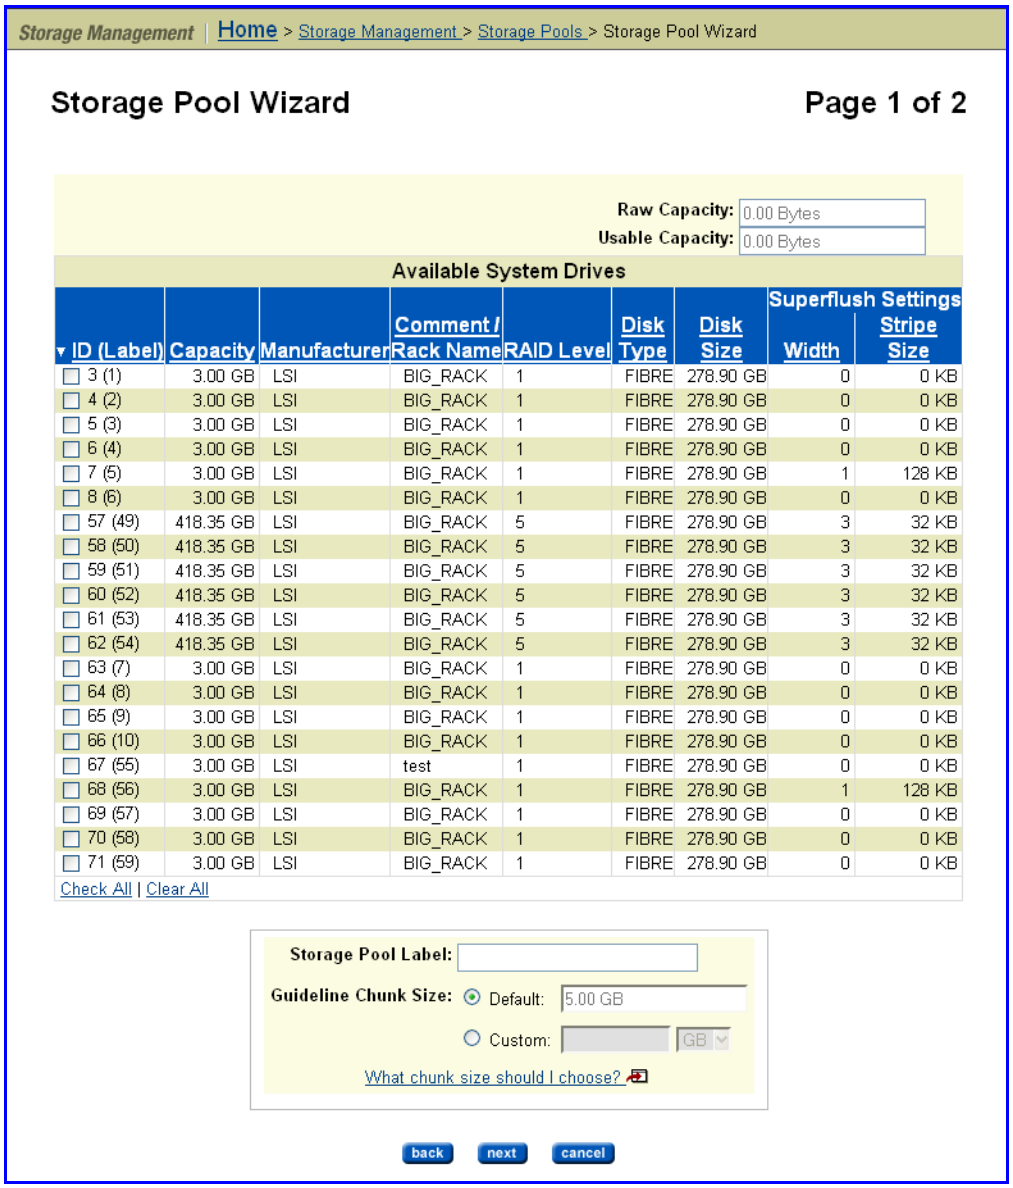

The following table describes the fields in this page:

Item/Field	Description
Raw Capacity	Running total of selected SD sizes.
Usable Capacity	Capacity of the Storage Pool that will be created, based on the selected SDs. Ideally, this quantity and the **Raw Capacity** quantity would be equal.
ID (Label)	SD number assigned by the server and the SD label. Appears in the event log.

Item/Field	Description
Capacity	SD's size.
Manufacturer	Manufacturer of the RAID array containing the disks on which the SD was created. For information about the supported RAID Arrays, refer to the *Storage Subsystem Guide*.
Comment / Rack Name	Rack name and any additional information.
Disk Type	Type of SD; for example, *Fibre*, *SATA*, and *SAS*.
Width	Number of physical disks in the SD.
Stripe Size	Data format size used for writing to a SD.
Storage Pool Label	The name of the new storage pool.
Guideline Chunk Size	The chunk size that will be used for storage pool expansion. You can use the default chunk size suggested by the system, or you can choose a custom chunk size.

3. **Select at least one SD.**

   Select one or more SDs for use in building the new Storage Pool.

   A Storage Pool cannot contain SDs on RAID arrays with different manufacturers, disk types, or RAID levels. Any attempt to create a Storage Pool from such dissimilar SDs will be refused.

   For the highest level of performance and resiliency, SGI Global Services strongly recommends that all SDs be of the same capacity, width, and stripe size, and consist of disks of equal size; however, creating a Storage Pool with SDs that are not identically configured is allowed after first acknowledging a warning prompt.

4. **Enter the requested information.**

   Two fields merit special mention:

   - **Storage Pool label.** Storage Pool labels are not case sensitive, and they may not contain spaces or special characters.
   - **Guideline Chunk size.** *To select the default chunk size,* click **Default.** The default chunk size will be shown in the adjacent text box, and is automatically be calculated to equal [Storage Pool size]/256. *To set a custom chunk size,* click **Custom** and enter a size between 512 MB and 1 TB. SGI Global Services does not recommend chunk sizes less than 5 GB.

Note: Storage Pool labels must be unique within a server or cluster. Also, a Storage Pool cannot have the same label as a file system.

Tip: Label Storage Pools for the data they will contain. Also, make sure your Storage Pool labels are unique across your entire site, in case you need

to move them.

5. **Apply settings.**

   Verify your settings, then click **next**; in the confirmation dialog, click **create.**

Note: After the Storage Pool has been created, it can be filled with file systems. For instructions, see Creating a File System, on page 138.

Deleting Storage Pools

A Storage Pool that does not contain file systems can be deleted at any time; otherwise, delete the file systems first. After the pool has been deleted, its SDs become free and available for use by new or existing Storage Pools. For instructions about deleting a file system, see Deleting a File System, on page 150.

To delete a Storage Pool:

1. **Navigate to the Storage Pools page.**

   From the **Storage Management** page, click to display the **Storage Pools** page.

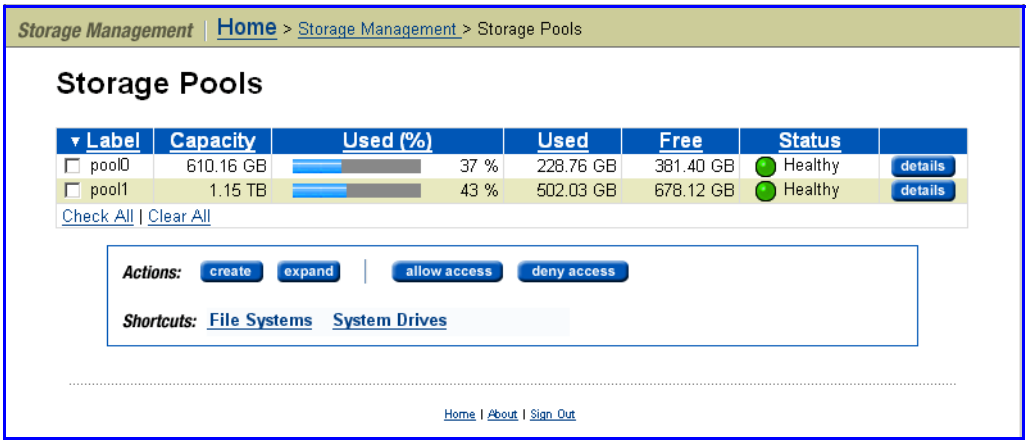

1. **Select a Storage Pool.**

   For the selected **Storage Pool,** click **details.**

2. **Delete the Storage Pool.**

   Click **delete,** then **OK** to confirm.

Expanding a Storage Pool

Expand a Storage Pool to provide more usable storage for file systems. A storage pool can be expanded at any time and without interrupting service to clients by adding more SDs to the pool. A Storage Pool can be expanded until its total size has reached 256 TB, but it is limited to a maximum of 16,384 chunks.

Tip: For optimal performance, when expanding a Storage Pool, add as many SDs as possible (the maximum number of SDs you can add at one time is 32). When adding SDs to a Storage Pool, all SDs added at the same time are placed

into a single stripe set, so they should have the same capacity.

To expand a Storage Pool:

1. **Navigate to the Storage Pools page.**

   From the **Storage Management** page, click to display the **Storage Pools** page:

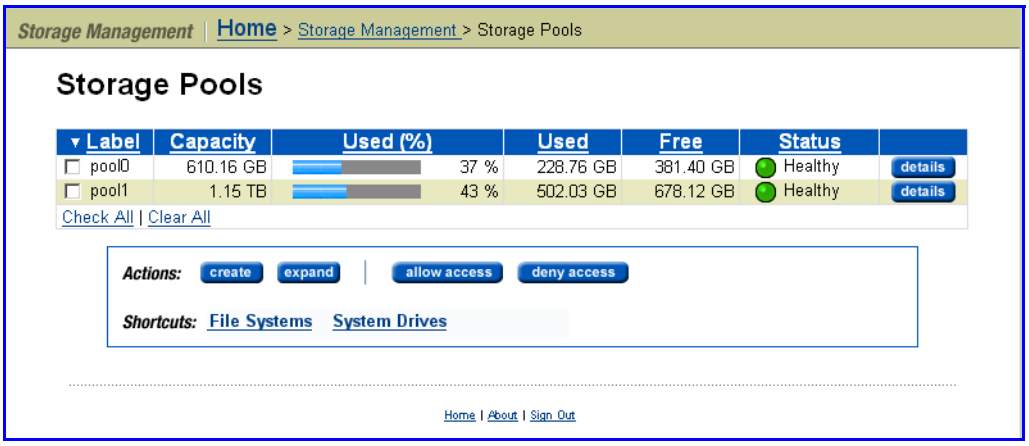

2. **Display and select from SDs available for the expansion.**

   Select a Storage Pool and click **expand** to view a list of available SDs. Select one or more SDs and click **next;** in the confirmation dialog, click **expand** to add the SDs to the pool.

Reducing the Size of a Storage Pool

The size of a Storage Pool cannot be reduced.

Denying Access to a Storage Pool

Denying access to a Storage Pool:

- Removes the association between the file systems hosted by the pool and their Virtual Server (EVS). Once access to the pool is allowed, file systems must be reassociated with an EVS. Access to a Storage Pool should only be denied during a planned maintenance window and only at the direction of SGI Global Services.
- Removes the Storage Pool and its file systems from memory, meaning that those file systems are no longer visible using Web Manager.
- Denies access to the SDs that host the Storage Pool, with the effect that the pool is not reloaded the next time the cluster node/server boots.
- Irrevocably deletes all data associated with file systems in the NVRAM buffer, with the effect of destroying data unless all file systems hosted by the Storage Pool were unmounted successfully before access to the Storage Pool was denied.

Caution: The technique of denying access to a Storage Pool can be used to

prepare a storage array for physical relocation. The Storage Pool and its contents are not lost, nor deleted, and access can be returned when needed.

The procedure for denying access to a Storage Pool involves unmounting each file system in the pool, then changing the pool's access mode. For instructions on unmounting file systems, see Unmounting a File System, on page 152.

To deny access to a Storage Pool:

1. **Navigate to the File Systems page.**

   From the **Storage Management** page, click to display the **File Systems** page:

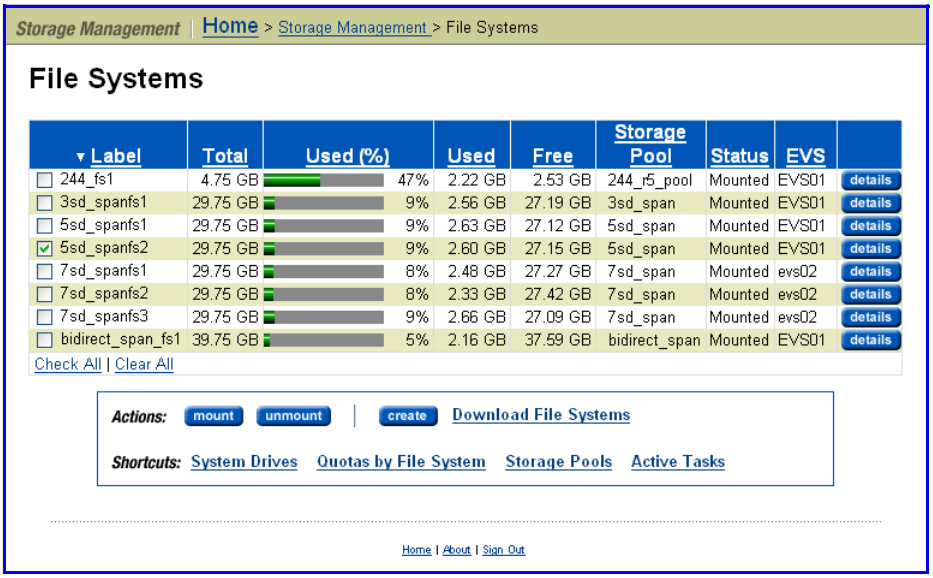

1. **Unmount every file system in the Storage Pool.**

   Select each file system, then click **unmount**; in the confirmation dialog, click **OK.**

2. **Select a Storage Pool and deny access.**

   Click the **Storage Pools** shortcut to display a list of all pools, then select a particular Storage Pool and click **Deny Access**; in the confirmation dialog, click **OK.**

   This will also remove the pool from the Storage Pools list, but it will not be deleted.

Allowing Access to a Storage Pool

This procedure restores access to a Storage Pool, but can also be used when a storage array previously owned by another server has been physically relocated to be served by another server. The process restores access to the SDs that belong to the Storage Pool, then restores access to the pool itself.

To allow access to a Storage Pool:

1. **Navigate to the System Drives page.**

From the **Storage Management** page, click to display the **System Drives** page:

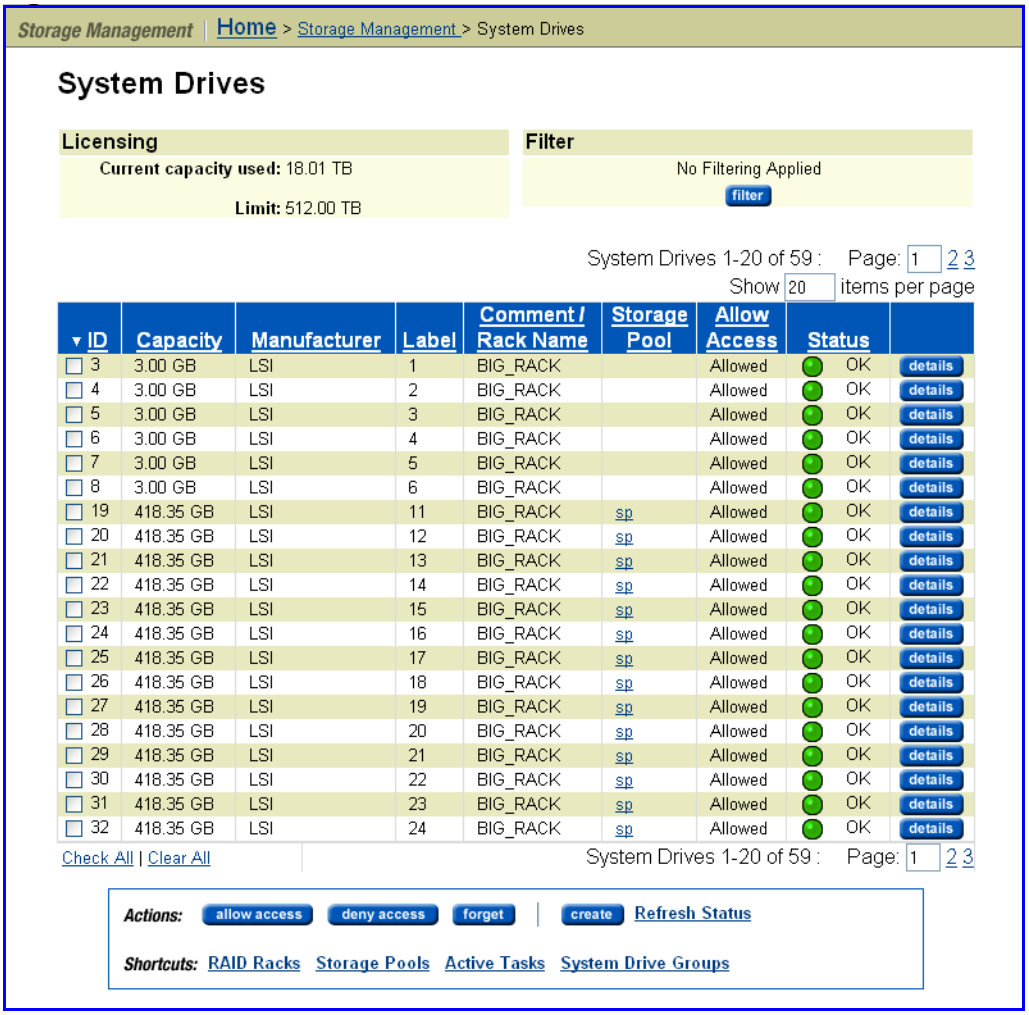

2. **Modify SD access.**

   Select one of the SDs belonging to the Storage Pool, then click **Allow Access.**

3. **Modify Storage Pool access.**

   Select a pool and click **details**. In the **Details** page for that Storage Pool, click **Allow Access;** then, in the confirmation page, click **OK.**

Note: To become accessible, each file system in the Storage Pool must be associated with an EVS. To do this, navigate to the **Details** page for each file system in the Storage Pool and assign it to an EVS.

Renaming a Storage Pool

The name for a Storage Pool can be changed at any time, and without affecting any clients.

To rename a Storage Pool:

1. **Navigate to the Storage Pools page.**

   From the **Storage Management** page, click to display the **Storage Pools** page:

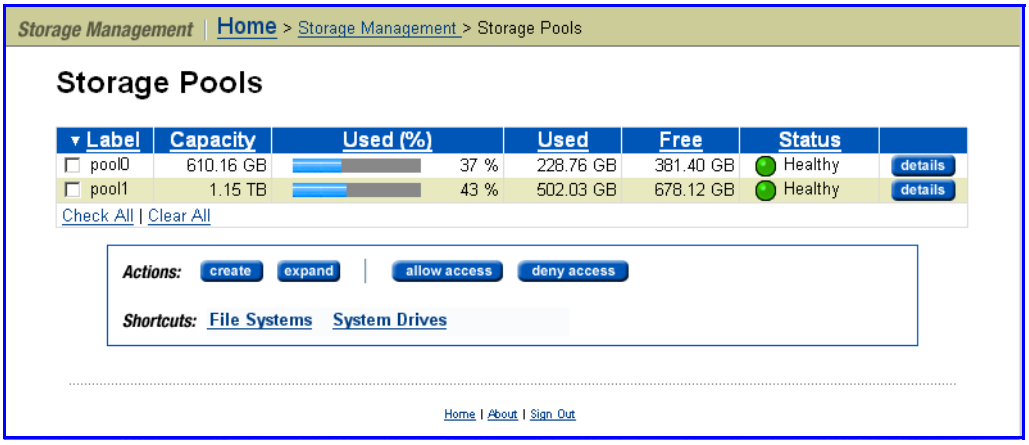

2. **Rename the Storage Pool.**

   Select a Storage Pool and click **details.** In the **Details** page, enter a new name in the **Label** text box, then click **rename.**

   Note: Storage Pool labels are not case sensitive and they may not contain spaces or special characters.

Configuring Automatic File System Expansion for an Entire Storage Pool

Use this procedure to allow or prohibit automatic expansion of all file systems in the specified Storage Pool. This setting will override any auto-expansion setting on individual file systems in the Storage Pool. This setting only affects auto-expansion; manual expansion of file systems in the Storage Pool is not affected by this setting.

To configure automatic file system expansion for all file systems in an entire Storage Pool:

1. **Navigate to the Storage Pools page.**

From the **Storage Management** page, click **Storage Pools** to display the **Storage Pools** page:

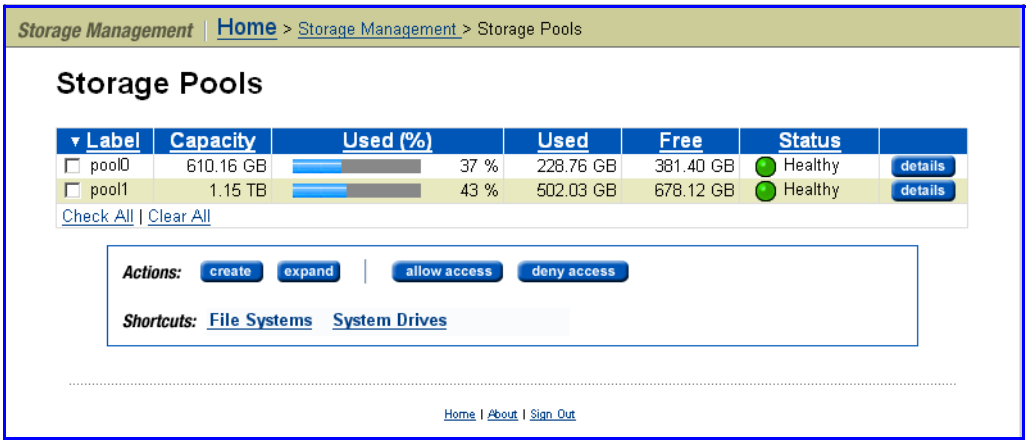

2. **Navigate to the Storage Pool Details page.**

   Select a Storage Pool and click **details** to display the Storage Pools Details page.

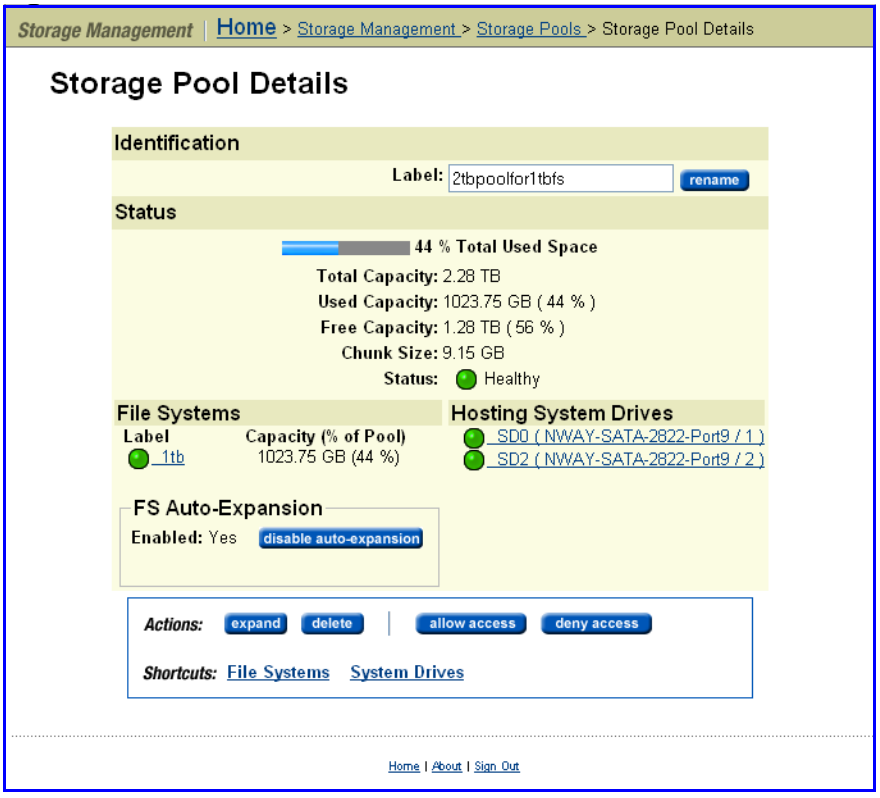

3. **Configure auto-expansion.**

   You can configure file system auto-expansion at the Storage Pool level in one of two ways:

- **Enable auto-expansion**

  Even if the Storage Pool is configured to allow its file systems to automatically expand, the file systems must also be configured to support automatic expansion. For more information, see Expanding File Systems Automatically, on page 154.

  Note: After a file system has expanded, its size cannot be reduced.

  If file system auto-expansion is currently disabled, you can enable it by clicking **enable auto-expansion** in the **FS Auto-Expansion** option box.

- **Disable auto-expansion**

  When automatic expansion of a file system is disabled, manual expansion of file systems in the Storage Pool is still possible. For more information, see Expanding File Systems Manually, on page 153.

  If file system auto-expansion is currently enabled, you can disable it by clicking **disable auto-expansion** in the **FS Auto-Expansion** option box.

Using File Systems

The storage server provides extensive facilities to monitor and manage file systems, and supports robust techniques for data protection. The quantity of space in use can be monitored and restricted; should a file system become full, it can be expanded.

A *read cache* is a special read-only file system that increases read performance for NFS clients. Read caches have special characteristics and limitations. For information about read caches, refer to Read Caching, on page 412.

Creating a File System

This procedure creates a new file system or read cache within an existing Storage Pool. If a Storage Pool will contain more than one file system, or a file system *and* a read cache, a Storage Pool license must be installed. A Storage Pool is required before a file system or a read cache can be created.

To create a file system:

1. **Navigate to the Create page.**

From the **Storage Management** section in the **Home** page, select **File Systems**; then, from the **File System** page, click **create** to view the **Create** page:

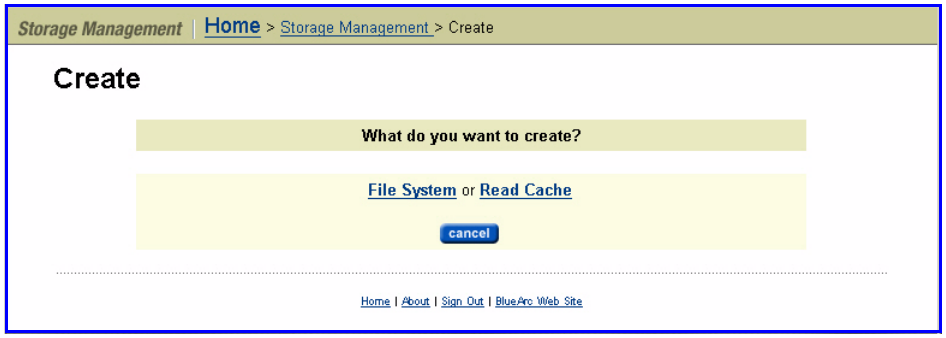

2. **Select *file system* or *read cache*:**
   - To add a normal file system (a file system that is not a read cache), click **File System** and continue to the next step.
   - To add a read cache, click **Read Cache**. For more information about creating read caches see Configuring Read Caching, on page 447.
   - To return to the **File Systems** page without creating a file system, click **Cancel**.
3. **Indicate new or existing Storage Pool location.**

   In the **Create File System** page, specify whether to use a new Storage Pool (that you will create) or an existing Storage Pool:

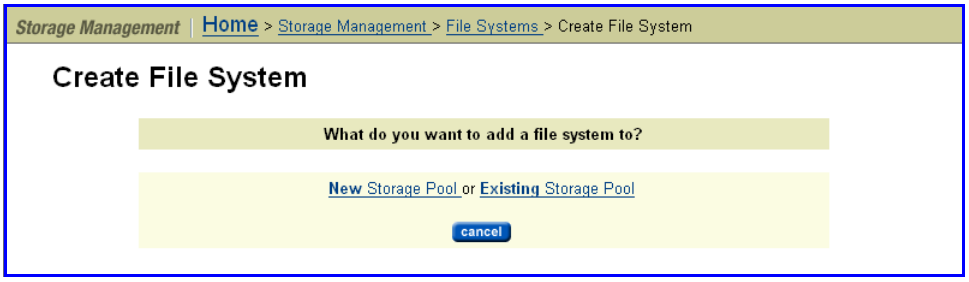

   - *To create a new Storage Pool* for the file system or read cache, click **New Storage Pool** to start the **Storage Pool** wizard. For more information about the Storage Pool Wizard, see Creating a Storage Pool, on page 129.
   - *To add the file system or read cache to a pre-existing Storage Pool,* click **Existing Storage Pool**. Select a Storage Pool to contain the file system, then click **next**.
4. **Create a normal file system.**

Caution: This step assumes that you have chosen to create a *normal file system*. To create a *read cache,* skip this step and go to the next step.

Enter the requested information in the normal file system-specific instance of the **Create File System** page:

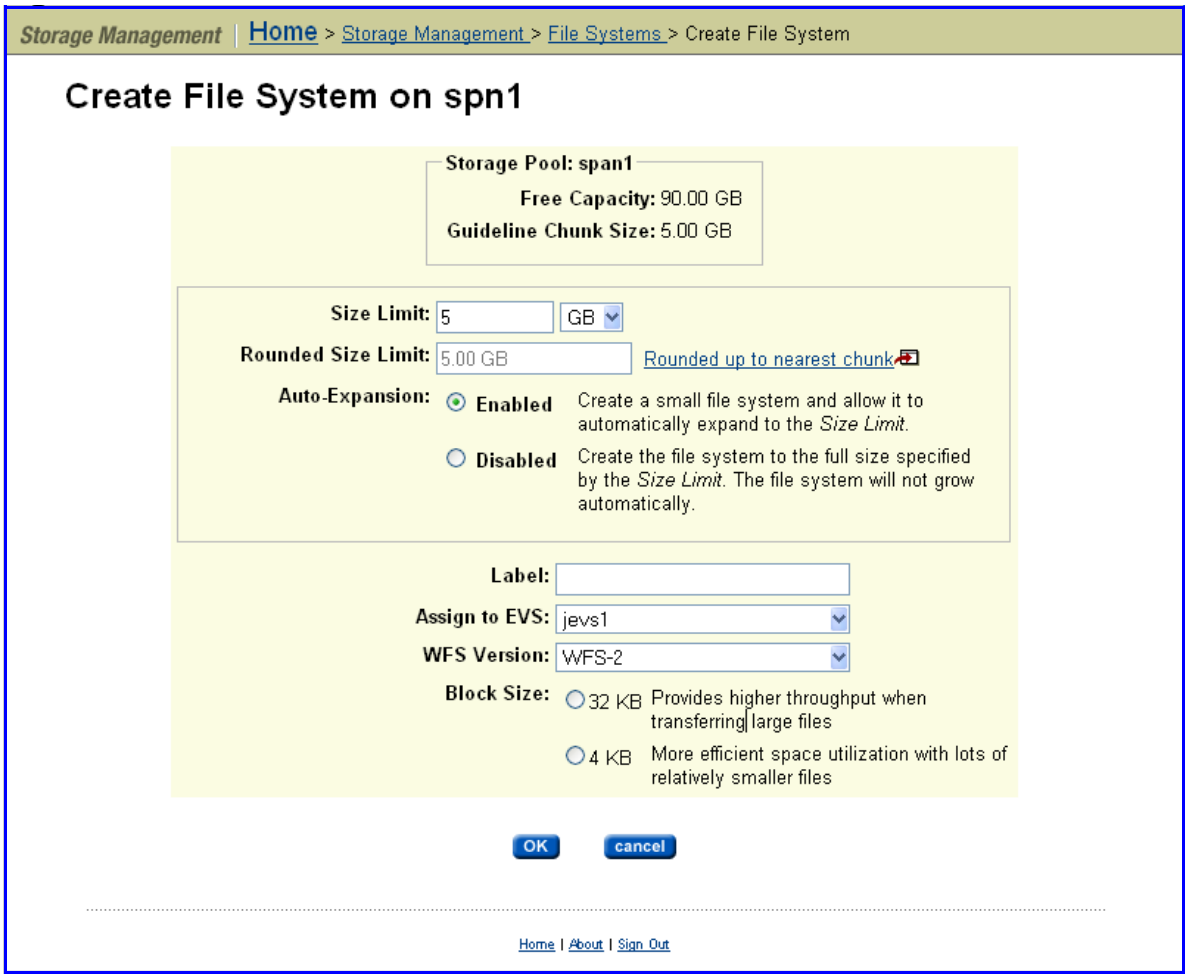

The following table describes the fields in this page and provides a column with special instructions pertaining to normal file systems:

Item/Field	Description	Settings for a Normal File System
Storage Pool	Name of the Storage Pool in which the file system is being created.	(Display only)
Free Capacity	Amount of available space available in the Storage Pool that can be used by file systems or a read cache.	(Display only)
Guideline Chunk Size	Approximate size of the chunks used in the selected Storage Pool.	(Display only)

Item/Field	Description	Settings for a Normal File System
Size Limit	*If Auto-Expansion is enabled,* this will automatically be the maximum size to which a file system will be allowed to expand. *If Auto-Expansion is disabled,* this specifies the capacity with which the new file system should be created.	Enter a Size Limit for the file system. This defines the maximum size to which the file system will grow through Auto-Expansion. Once the file system has been created, this value can be changed on the **File System Details** page. This limit is not enforced for manual file system expansions performed through the CLI.
Rounded Size Limit	Approximate size limit, based on the defined Size Limit and the chunk size defined for the Storage Pool. For more information, click **Rounded to nearest chunk**. If the specified size is not a multiple of the chunk size, then the server rounds down to the nearest chunk boundary.	This setting is calculated automatically by the system, but may be changed.
Auto-Expansion	Enable or disable Auto-Expansion. Use this to allow or constrain growth of this file system.	Be aware that Storage Pools can be configured to prevent the growth of file systems. A file system can never shrink; once space is allocated to a file system, the space cannot be recovered, and the file system cannot be reduced in size. When expanding, the file system will use the Storage Pool's chunk size as its growth increment. File systems configured to automatically expand will do so when they are about 80% full. File systems can be expanded manually through the CLI. File system expansion does not interrupt file services or require the file system to be unmounted.
Label	Enter the label (name) by which the file system should be referenced. Label file systems for the data they will contain. Also, make sure your file system labels are unique across your entire site, in case you need to move them.	 **Note:** File system labels are not case sensitive, and they may not contain spaces or special characters. File system labels must be unique within a server or cluster. Also, a file system cannot have the same label as a Storage Pool.
Assign to EVS	Select the EVS to which the file system should be assigned.	From the EVS drop-down list, select the EVS to which the file system should be assigned.

Item/Field	Description	Settings for a Normal File System
WFS Version	Use to specify the file system format. The WFS-2 file system format provides an alternative file system format to the WFS-1 format (the original file system supported by the Titan Server). Note that the WFS-2 file system format is supported only on the Series 3000 Titan Server and later hardware platforms (including the IS-NAS Server). For more information, see File System Formats, on page 111.	When supported by the server/node, WFS-2 is the default file system format. Select the desired file system format for the new file system. **Note:** Once a file system has been formatted, you cannot change file system vomits without formatting the file system, which will erase all the data in the file system.
WORM	Use to enable retention control. For more information, see Using WORM File Systems, on page 162.	Select whether the file system should be a normal or WORM file system. Unless the file system is to be used for regulatory compliance purposes, select Not WORM.
Block Size	Use to configure optimal block size for the file system. For more information, see File System Block Size, on page 112.	Select the desired file system block size.

When finished, verify the configuration, then click **OK.** A normal file system will be created.

Note: Before you can use the file system, it must be formatted and mounted. If you create a file system using Web Manager, the file system is formatted automatically. By default, when WFS-2 is supported by the hardware platform, the new file system is formatted using the WFS-2 format. You can, however, choose to format a new file system using the WFS-1 format (see Creating a File System, on page 138and Formatting a File System, on page 151 for more information.For information on mounting the file system, see Mounting a File System, on page 152.

5. **Create a read cache.**

Caution: This step assumes that you have chosen to create a read cache. To create a normal file system, refer to the previous step.

Enter the requested information in the read cache-specific instance of the **Create File System** page:

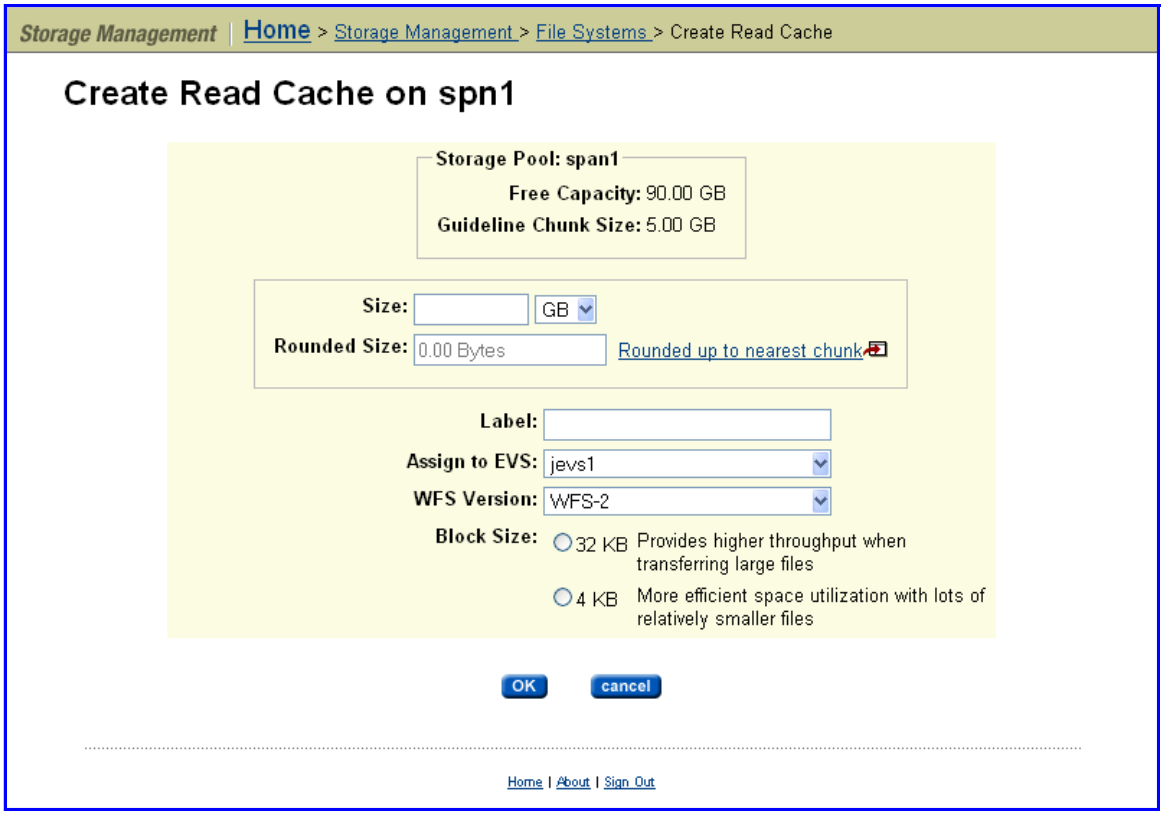

The following table describes the fields in this page and provides a column with special instructions pertaining to read caches.

Item/Field	Description	Settings for a Read Cache
Storage Pool	The name of the Storage Pool in which the read cache is being created.	(Display only)
Free Capacity	The amount of available space available in the Storage Pool that can be used by file systems or a read cache.	(Display only)
Guideline Chunk Size	The approximate size of the chunks used in the selected Storage Pool.	(Display only)
Size	The size of the read cache being created.	Enter a Size for the read cache. This defines the initial size of the read cache. This value can be changed on the **File System Details** page once the read cache has been created.

Item/Field	Description	Settings for a Read Cache
Rounded Size Limit	This shows the approximate size limit, based on the defined Size Limit and the chunk size defined for the Storage Pool. For more information, click **Rounded to nearest chunk.**	This setting is calculated automatically by the system, but may be changed.
Label	Enter the label by which the read cache should be referenced.  **Note:** File system/read cache labels are not case sensitive, and they may not contain spaces or special characters.	In the Label text box, type in a name for the new read cache.
Assign to EVS	Select the EVS to which the read cache should be assigned. The read cache should be assigned to an EVS that has a preferred mapping to a cluster node.	From the EVS drop-down list, select the EVS to which the read cache should be assigned.
WFS Version	Use to specify the file system format. The WFS-2 file system format provides an alternative file system format to the WFS-1 format (the original file system supported by the Titan Server). Note that the WFS-2 file system format is supported only on the Series 3000 Titan Server and later hardware platforms (including the IS-NAS Server). For more information, see File System Formats, on page 111.	When supported, WFS-2 is the default file system format. Select the desired file system format for the new file system.  **Note:** Once a file system has been formatted, you cannot change file system vomits without formatting the file system, which will erase all the data in the file system.
Block Size	Use to configure optimal block size for the read cache. For more information, refer to File System Block Size, on page 112.	Select the desired file system block size.

When finished, verify the configuration, then click **OK.** A read cache will be created, and the **File System Details** page appears.

Caution: Only one read cache should be created on any cluster node. If you create more than one read cache, only one will be used.

Note: Before you can use the read cache, it must be formatted and mounted. If you create a read cache using Web Manager, it is formatted automatically. For information on mounting the read cache, see Mounting a File System, on page 152.

Viewing Available File Systems

To view available file systems, navigate to the **Storage Management** page, then click to display the **File Systems** page:

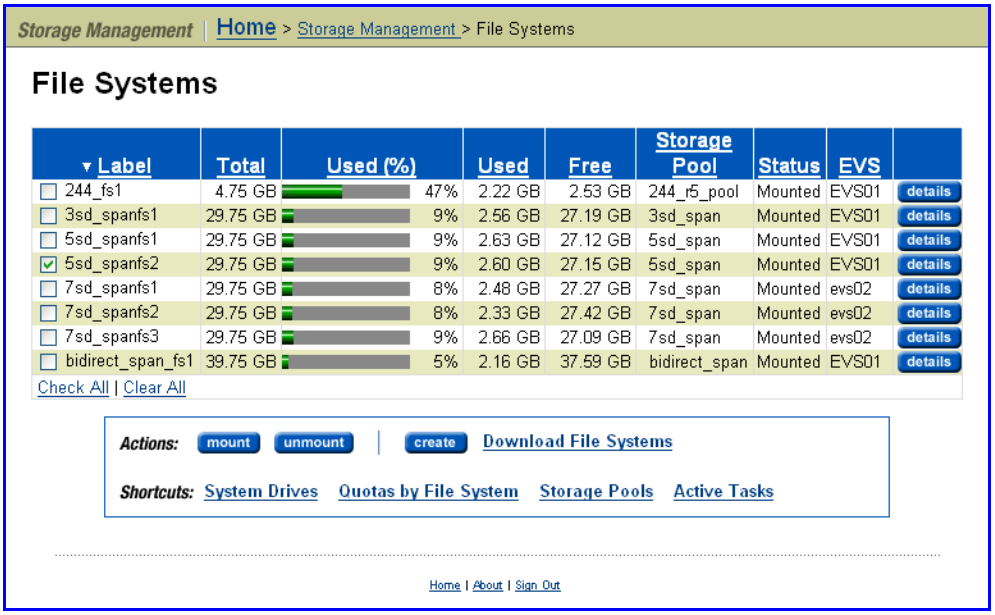

The following table describes the columns in this page:

Item/Field	Description
Label	Name of the file system, assigned upon creation and used to identify the file system when performing particular operations; for example, creating an export or taking a snapshot.
Total	Size of the file system (GB).
Used	Amount of space used (GB).
Free	Amount of free space available (GB).
Storage Pool	Name of the Storage Pool of which the file system is a member.

Item/Field	Description
Status	**Checking:** The file system is being checked; during this check, the approximate percentage of completion is displayed. **Failing:** The file system has failed, but is being checked, fixed, or recovered. **Formatting:** The file system is being formatted. **Initializing System Drive:** The System Drive is initializing. **Mounted:** The file system has been mounted and is available for service. **Mounted as Readonly:** The file system mounted, but is in read-only mode. **Mounting:** The file system is being mounted and available for service. **Not Assigned to EVS:** The file system is not currently assigned to an EVS. **Not Available for Mounting:** The file system is not available. Make sure to enable the EVS to which the file system is assigned and make sure the file system is not marked 'hidden'. **Not Mounted:** The file system is not mounted. **Not Mounted (System Drive initialization status unknown):** The file system is not mounted, and the SD initialization status is not known. **Not Mounted (System Drive is not initialized):** The file system is not mounted, the SD initialization status is known, but SD initialization is not complete. **Syslocked:** The file system is syslocked. For more information on Syslocked mode, see Using System Lock on File Systems, on page 156.
EVS	EVS to which the file system is assigned.
mount	Select an unmounted file system and click **mount** to mount the file system.
unmount	Select a mounted file system and click **unmount** to unmount the file system.

The following **Usage Alerts** criteria apply:

- *If Usage Alerts are enabled on the entire file system,* the sliding bar turns yellow when the *warning limit* is exceeded and amber when the *severe limit* is exceeded.
- *If Usage Alerts are not enabled*, the sliding bar turns yellow when *85% capacity* is reached and amber when the file system is *full*.

The following **Actions** are available:

- Click **Download File Systems** to download a spreadsheet containing information about all of the listed file systems.
- Click **System Drives** to display the **System Drives** page.
- Click **Quotas by File System** to display the **Quotas by File Systems** page described in Setting User/Group Defaults, on page 175.

- Click **Storage Pools** to display the list of Storage Pools on the server.
- Click **Active Tasks** to view the **Active Tasks** page for more information about active tasks.

Note: The server remembers which file systems were mounted when it shuts down, and mounts them automatically during system startup.

To View the Details of a File System

From the **File Systems** page, select a file system and click **details** to display the **File System Details** page:

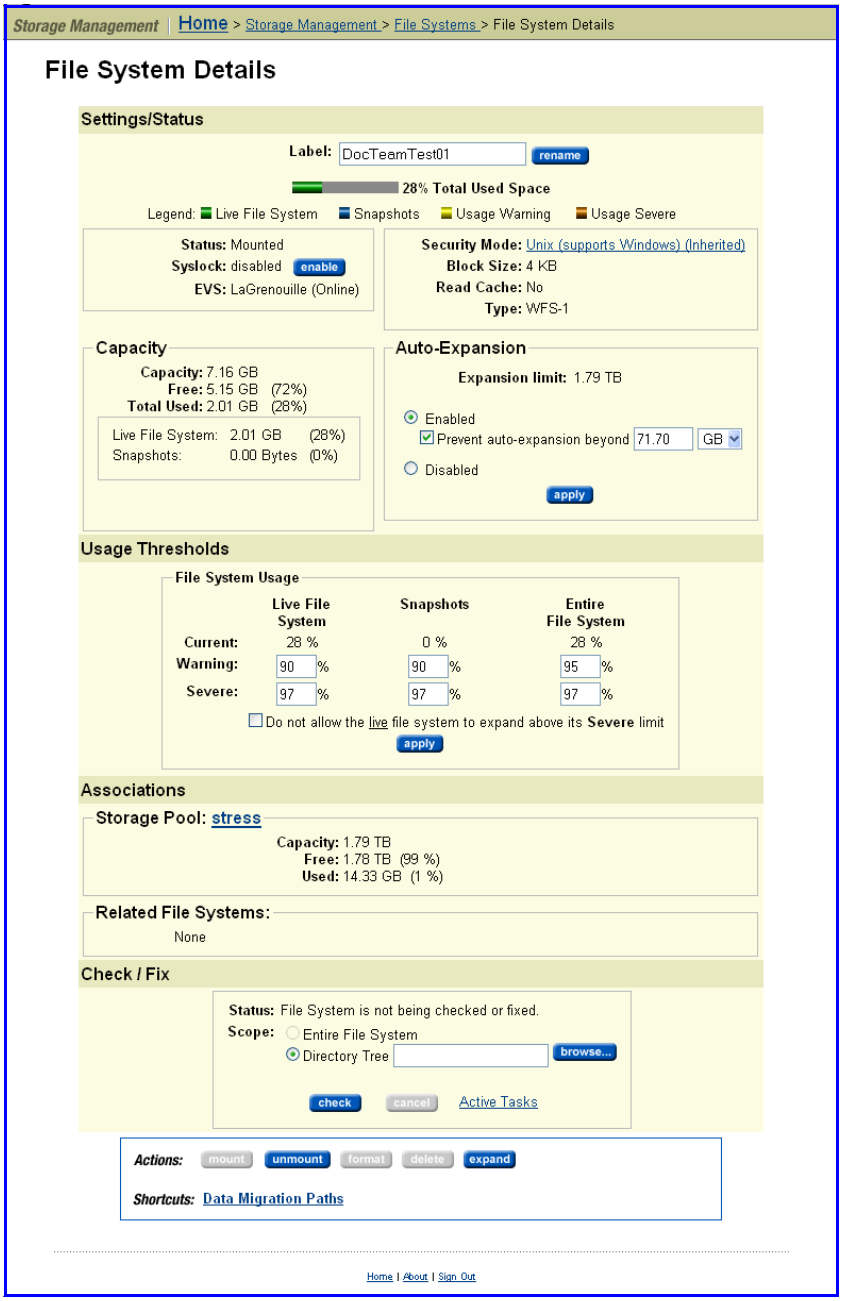

The following table describes the fields in this page:

Item/Field	Description
Settings/Status	
Label	Name of the file system, assigned upon creation and used to identify the file system when performing particular operations; for example, creating an export or taking a snapshot. You can rename the file system by typing a new name in the Label box, and clicking **rename**. Label file systems for the data they will contain. Also, make sure your file system labels are unique across your entire site, in case you need to move them.  **Note:** File system labels must conform to the following rules: • File system labels are not case sensitive, and they may not contain spaces or special characters. • File system labels must be unique within a server or cluster. • A file system cannot have the same label as a Storage Pool.
% Total Used Space	Percentage of the file system's total allocated space that has been used. This total reflects data and snapshots, if any.
Status	Current status of the file system, showing total used space and whether the file system is mounted or unmounted.
System Lock	Indicates whether the file system is in Syslocked mode (System Lock *enabled*), or if the file system is not in Syslocked mode (System Lock *disabled*). When System Lock is enabled for a file system, NDMP has full access to the file system and can write to it during a backup or replication, but the file system remains in read-only mode to clients using the file service protocols (NFS, CIFS, FTP, and iSCSI). *To enable/disable the System Lock for a file system,* click **enable/disable**. When viewing the details of a read cache, the System Lock's **enable/disable** button is not available.
EVS	EVS to which the file system is assigned. If the file system is not currently assigned to an EVS, a list of EVSs (to which the file system can be assigned) appears.
Security	Displays the file system security policy defined for the file system.
WORM	Indicates whether this file system is a "Write Once Read Many" file system or a "regular" file system. The values may be any of the following: • **Strict WORM,** meaning that the file system *is* a WORM file system that *cannot* be reformatted. • **Non-strict WORM,** meaning that the file system *is* a WORM file system that *can* be reformatted. • **Not WORM,** meaning that the file system *is not* a WORM file system, and that it *can* be reformatted. For more information on WORM file systems, see WORM File Systems, on page 112.
Block Size	File system block size: 32 KB or 4 KB, as defined when the file system was formatted.

Item/Field	Description
Read Cache	Indicates whether this file system is a read cache (*Yes*) or a regular file system (*No*).
Type	Indicates whether this file system has been formatted as a WFS-1 file system or a WFS-2 file system. For more information on WFS-s file systems, see File System Formats, on page 111.
Capacity	
Capacity	Total amount of formatted space (free + used space).
Free	Total amount of file system space unused (free), in GB and as a percentage of the total.
Total Used	Total amount of file system space in use, in GB and as a percentage of the total.
Live File System	Total space used by the file system data, in GB and as a percentage of the total.
Snapshots	Total space used by snapshots of the file system, in GB and as a percentage of the total.
Auto-Expansion	
Expansion limit	Indicates the capacity of the file system (including data and snapshots). When this limit is reached, auto-expansion of the file system occurs, if enabled.
Enabled/Disabled	Enables or disables auto-expansion of the file system. When enabled, you can limit the maximum size of the file system by clicking the checkbox and specifying a maximum size in the edit box and drop-down list box.  **Note:** If you change any of the auto-expansion settings, you must click **apply** to make the changes effective.
Usage Thresholds	Usage thresholds are expressed as a percentage of the space that has been allocated to the file system. When a threshold is reached, an event is logged and, depending on quota settings, an email may be sent. This area displays information on current usage, and you can use the edit boxes to specify the *Warning* and *Severe* thresholds: • The **Warning** threshold should be set to indicate a high, but not critical, level of usage. • The **Severe** threshold should be set to indicate a critical level of usage, a situation in which an out-of-space condition may be imminent. You can define both **Warning** and **Severe** thresholds for any or all of the following: • Live file system (data). • File system snapshots. • Total of the live file system and snapshots. To verify that the live file system does not expand beyond its Severe threshold setting, which would cause snapshots to be lost, fill the **Do not allow the live file system to expand above its Severe limit** checkbox.

Item/Field	Description
Check/Fix	

Status	Indicates when the file system was last checked, and displays its status since its last reboot. File system status messages may be any of the following:
	• File system fixed.
	• File system checked.
	• File system fix aborted by the user.
	• File system check aborted by the user.
	• Could not find directory tree to fix.
	• Could not find directory tree to check.
	• File system being fixed.
	• File system being checked.
	• File system not fixed since reboot.
	• File system not checked since reboot.
	• File system fix failed.
	• File system check failed.
	  • File system check does not cease after failing initially.
	  • You can start a check of the file system, or just part of the file system, using the Scope settings and the browse and check buttons.
	If one or more checks is in progress, click **Active Tasks** to view the **Active Tasks** page for more information about active tasks.
	Click **cancel** to abort a check in progress.

Scope	The scope controls allow you to set the scope of a check, by selecting either the **Entire File System** radio button or the **Directory Tree** radio button. *To check the whole file system,* select the **Entire File System** radio button. *To check a part of the file system,* select the **Directory Tree** radio button, then use the browse button to navigate to the part of the file system you want to check. Once you have set the scope, click **check** to start the check.
Storage Pool	The name of the Storage Pool in which the file system or read cache was created.
Capacity	Total space allocated to the Storage Pool, in GB.
Free	Total Storage Pool free space, in GB and as a percentage of the total.
Used	Total Storage Pool used space, in GB and as a percentage of the total.

Deleting a File System

A file system can be deleted at any time, unless it is a Strict WORM file system. After a file system has been deleted, the free space is restored to its Storage Pool.

Caution: After a file system has been deleted, it cannot be recovered.

Note: A user must be either a Global Admin or Storage Admin, and must have Advanced Mode enabled to delete a file system. For more information, see Managing Users and Roles, on page 29.

To delete a file system:

1. **Select a file system.**

   In the **Home** page, from the **Storage Management** section, click **File Systems** to display a list of all file systems. For the file system to be deleted, click **details.**

2. **Delete the file system.**

   a. **Unmount.**

      If the file system is mounted, click **unmount.** Then, in the confirmation dialog, click **OK.**

   b. **Delete.**

      In the **Actions** section, click **delete.** Then, in the confirmation dialog, click **OK.**

Formatting a File System

Formatting a file system prepares it for use by clients for data storage. File systems created through the Web UI will be formatted and mounted automatically. Therefore, this procedure should rarely, if ever, be used.

Caution: Before using this procedure, please observe the following:

- Any existing data in the file system will be lost.
- User must be either a Global Admin or Storage Admin with Advanced Mode enabled. For more information, see Managing Users and Roles, on page 29.
- This procedure assumes that the file system has already been mounted. For more information, see Mounting a File System, on page 152.

To format an existing file system:

1. **Select a file system.**

   In the **Home** page, from the **Storage Management** section, click **File Systems** to display a list of all file systems. For the file system to be formatted, click **details.**

2. **Unmount the file system.**

   If the file system is mounted, do the following to unmount it:

   - From the **Label** column, select the file system.
   - From the **Actions** section, click **unmount.**
   - In the confirmation dialog, click **OK**.

3. **Navigate to the Format File System page.**

   Click **details** for the file system to be formatted. Click **format** to display the **Format File System** page.

4. **Select the WFS version and block size for the file system.**

- Using the **WFS Version** drop-down list, select WFS-1 or WFS-2 as the format for the file system. See File System Formats, on page 111 for more information about file system formats.
- Using the radio buttons, select 32 KB or 4 KB as the block size for the file system. See File System Block Size, on page 112 for more information about block sizes and how they affect file system performance.

5. **Format the file system.**

   Click **OK** to proceed or **cancel** to return to the **File System Details** page.

Mounting a File System

Use this procedure to manually mount a file system. Mounting a formatted file system makes it available to be shared or exported, and thus accessible for use by network clients. This procedure may also be used when an auto-mount of a file system has failed, which may occur in the following situations:

- The file system was not mounted when the server was shut down.
- The command line interface was used to disable auto-mounting.
- A storage system failure caused the server to restart.
- A storage subsystem failure caused the server to restart three times within a 30 minute period.

To mount a file system:

1. **Select a file system.**

   In the **Home** page, from the **Storage Management** section, click **File Systems** to display a list of all file systems. Fill the checkbox next to the label of the file system to be mounted.

2. **Mount the file system.**

   If the file system is unmounted, click **mount.**

Unmounting a File System

Unmount a file system when it needs to be removed from service. From a client point of view, it simply disappears. This will not harm the file system, nor affect any of the data in it.

1. **Select a file system.**

   In the **Home** page, from the **Storage Management** section, click **File Systems** to display a list of all file systems. Fill the checkbox next to the label of the file system to be unmounted.

2. **Unmount the file system.**

   If the file system is mounted, click **unmount.**

   In the confirmation dialog, click **OK.**

Expanding a File System

A file system can be expanded at any time, without interruption of service, if the following conditions exist:

- **Available space.** Sufficient available free space and chunks remain in the Storage Pool.
- **Chunk limit.** The file system expansion will not cause the file system to exceed the maximum allowable number of chunks in a file system.

A file system can be expanded in two ways:

- *Manually,* by specifying either the new size or how much space to add (CLI only). Manual expansion is necessary only if Auto-Expansion has been disabled specifically for that file system, or if Auto-Expansion has been disabled for all file systems in a Storage Pool.
- *Automatically,* by specifying the maximum size for the file system. When the file system is approximately 80% full, it is expanded automatically.

Note: The size of a file system cannot be reduced.

Expanding File Systems Manually

Manual file system expansion is supported through the Web Manager and through the CLI.

To expand a file system *manually:*

1. **Select a file system.**

   In the **Home** page, from the **Storage Management** section, click **File Systems** to display a list of all file systems. For the file system to be expanded manually, click **details.**

2. **Navigate to the Expand File Systems page.**

   Click **expand** to display the **Expand File System** page.

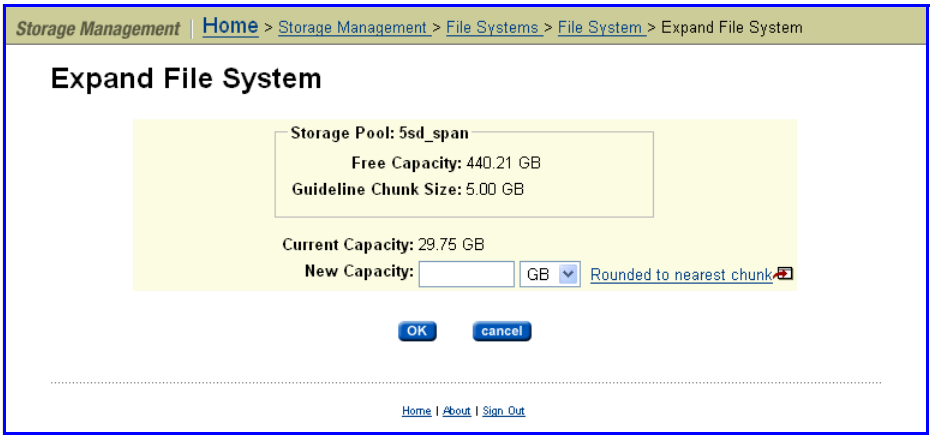

3. **Specify new capacity.**

   In the **New Capacity** edit box, specify the new file system capacity and use the drop-down list to select *MB, GB,* or *TB.*

4. **Start the process.**

   Click **OK** to expand the file system, or click **cancel** to return to the **File System Details** page without making the change.

   Note: Because space is always allocated in multiples of the chunk size set when the Storage Pool containing the file system was created, the final size of the file system may be slightly larger than you request.

   Manual expansion of file systems is also supported through the Command Line Interface. For detailed information on this process, run `man filesystem-expand` at the CLI.

Expanding File Systems Automatically

File system size can be automatically increased when needed. An overall size limit can also be set to control the amount of space it takes in a pool. The change can be applied without removing the file system from service.

1. **Select a file system.**

   In the **Home** page, from the **Storage Management** section, click **File Systems** to display a list of all file systems. For the file system to be expanded automatically, click **details.**

2. **Enable auto-expansion.**

   From the Auto-Expansion options box, select **Enabled**.

3. **Optionally, enable upper size limit.**

   If the file system must not expand beyond a specific size, do the following:

   - Fill the **Prevent Auto-expansion Beyond** checkbox. To disable the upper size limit, clear the checkbox.
   - Use the **Prevent Auto-expansion Beyond** text box and drop-down list to set the size limit.

4. **Start the expansion.**

   Verify that the configuration is correct, then click **apply** to start the process or **cancel** to decline.

   Note: The size of a file system cannot be reduced. This is also the case with a Storage Pool.

Relocating a File System

To relocate a file system:

1. **Navigate to the Storage Management page.**

   From the **Home** page, click on the **Storage Management** heading to view the **Storage Management** page.

2. **Navigate to the File System Relocation page.**

From the **Storage Management** page, click on **File System Relocation** to display the **File System Relocation** page.

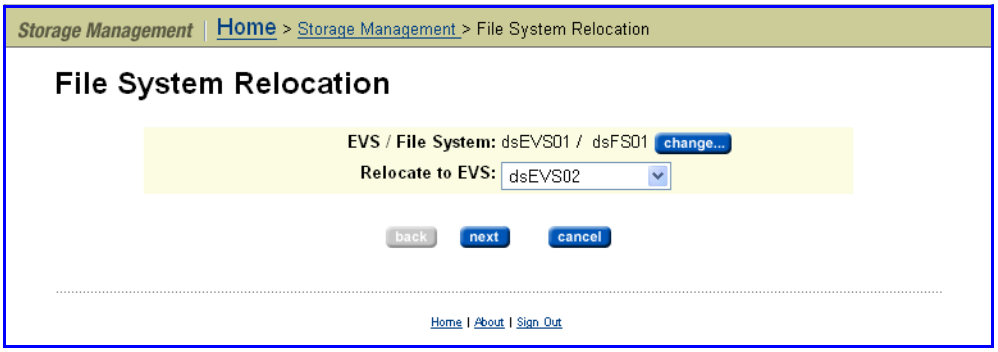

The following table describes the contents of this page:

Item/Field	Description
EVS/File System	Name of the currently selected EVS and file system. The **change** button allows you to select a different file system to relocate.
Relocate to EVS	A drop-down list contain a list of the EVSs available as a destination for the file system being relocated.
next	Click **next** to start the file system relocation process.
cancel	Click **cancel** to return to the **Storage Management** page.

3. **Select the file system to relocate.**

   To select a different file system, click **change** to display the **Select a File System** page, then click on the name of the file system you want to relocate.

4. **Select the destination EVS.**

   From the **Relocate to EVS** drop-down list, select the EVS to which you want to move the file system.

5. **Submit the relocation request.**

   Click **next**. A message box appears, listing any dependencies that exist in relocating this file system. Dependencies may include:

   - Additional file systems
   - CIFS shares
   - NFS exports
   - FTP users
   - Snapshot rules
   - iSCSI Logical Units
   - iSCSI Targets:

Acknowledge the message by clicking **OK**, or cancel the relocation by clicking **cancel**.

If you click **OK**, the **Relocating File Systems** page is displayed.

Note: If an iSCSI target contains Logical Units associated with another file system, an error message is displayed and the relocation will be canceled.

6. **Unmount any iSCSI Logical Units.**

   If the file system being relocated contains Logical Units, they must be unmounted before starting the file system relocation operation. Use the `iscsi-lu` command to unmount the iSCSI Logical Units (see the *Command Line Reference* for more information on the `iscsi-lu` command).

7. **Start, change, or cancel the relocation.**

   Click **OK** to start the relocation operation.

   To change the file system being relocated or the destination EVS, click **back**.

   To cancel the relocation, click **cancel**.

8. **Monitor the relocation.**

   After you start the relocation, the **Relocating File Systems** page displays the progress of the relocation.

   Once the relocation is complete, click **OK** to return to the **Storage Management** page.

Using System Lock on File Systems

System Lock mode protects file systems during replication and transfer of primary access operations. Four important distinctions apply:

- **NDMP (Network Data Management Protocol) versus File Service Protocols.** When **System Lock** is enabled for a file system:
  - NDMP has full access (including writes) during backups, replication, and transfer of primary access.
  - The file system remains in read-only mode to clients using the file service protocols (NFS, CIFS, FTP, and iSCSI).
- **System Lock versus Read Only:**
  - *When a file system is Syslocked,* NDMP still has full access to that file system and can write to it.
  - *When a file system is mounted as read-only,* NDMP (like all other protocols) has read-only access to that file system, and cannot write to it. To ensure that a file system remains completely unchanged, you should mount it as read-only.
- **Replication versus Transfer of Primary Access:**
  - *During replication operations,* the destination file system is put into System Lock mode.

- *During transfer of primary access operations*, both the source file system and the destination file system are put into System Lock mode.

- **Read Cache Exception.** A read cache may not be put into System Lock mode.

Enabling and Disabling System Lock for a File System

To enable/disable system lock for a file system:

1. **Navigate to the File System Details page.**

   From the **Home** page, click Storage Management, then click the **File Systems** link to display the **File System** page. Next, select a file system and click **details** to display its **File System Details** page:

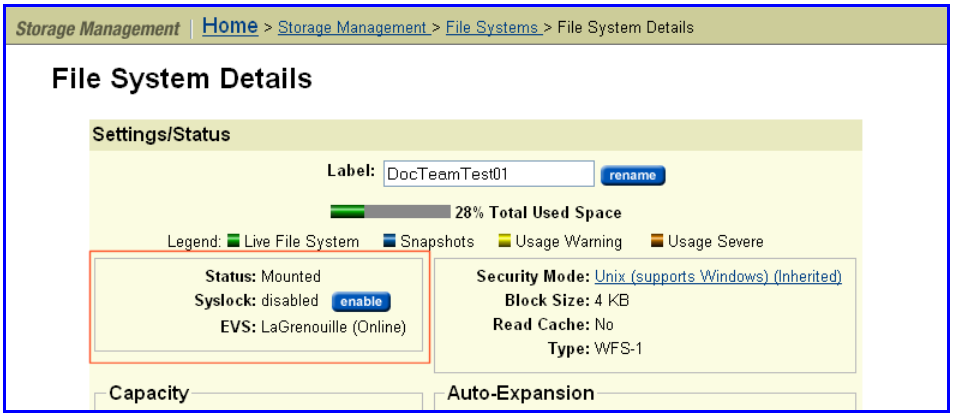

2. **Enable/disable system lock.**

   In **System Lock** field in the red section, toggle the **Enable/Disable** button.

   When the file system is in System Lock mode, the **Status** changes to "Syslocked," the **System Lock** status becomes "enabled," and the **Enable** button becomes **Disable**.

Recovering File Systems

Following some system failures, a file system may require recovery before mounting. If required, such a recovery is performed automatically when you mount the file system. Performing recovery rolls the file system back to its last checkpoint and replays any data in NVRAM.

In extreme cases, when you mount a file system after a system failure, the automatic recovery procedures may not be sufficient to restore the file system to a mountable condition. In such a case, you must *forcefully mount* the file system, which discards the contents of NVRAM before mounting the file system.

To recover a file system:

1. **Navigate to the File Systems page.**

   From the **Storage Management** page, click **File Systems** to display the **File Systems** page:

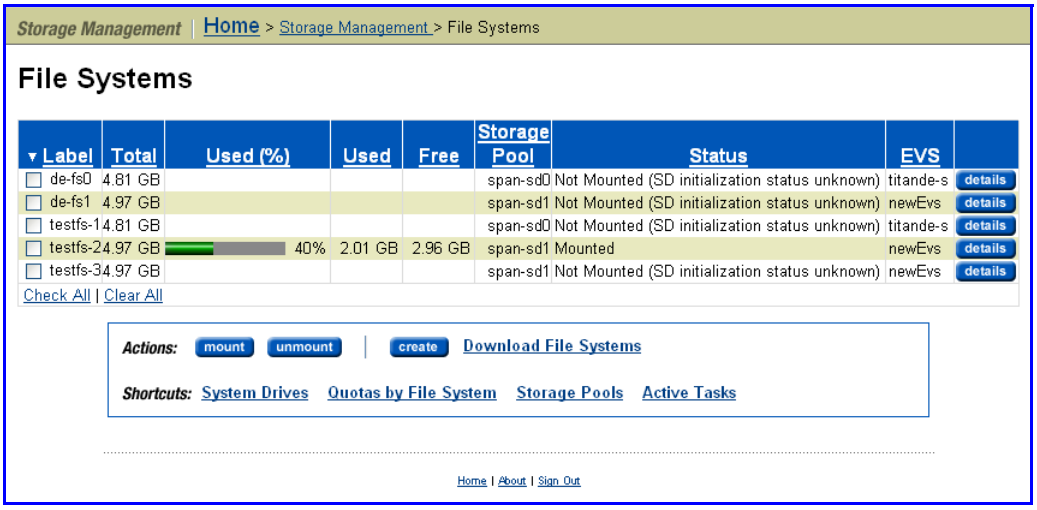

If a file system displays "Not Mounted" in the **Status** column, click **mount** to try to mount the file system.

2. **Mount the file system.**

   If a file system displays "Not Mounted" in the **Status** column, click **mount** to try to mount the file system.

   If necessary, the automatic recovery processes will be invoked automatically. If the automatic recovery fails, the file system will not mount, and the File Systems page will reappear, indicating that the file system was not mounted.

3. **For the File System that failed to mount, navigate to the File System Details page.**

   For the file system that failed to mount, click **details** to display the **File System Details** page. In the **Settings/Status** area of the page, the file system label will be displayed, along with the reason the file system failed to mount (if known), and suggested methods to recover the file system, including the link for the **Forcefully mount** option.

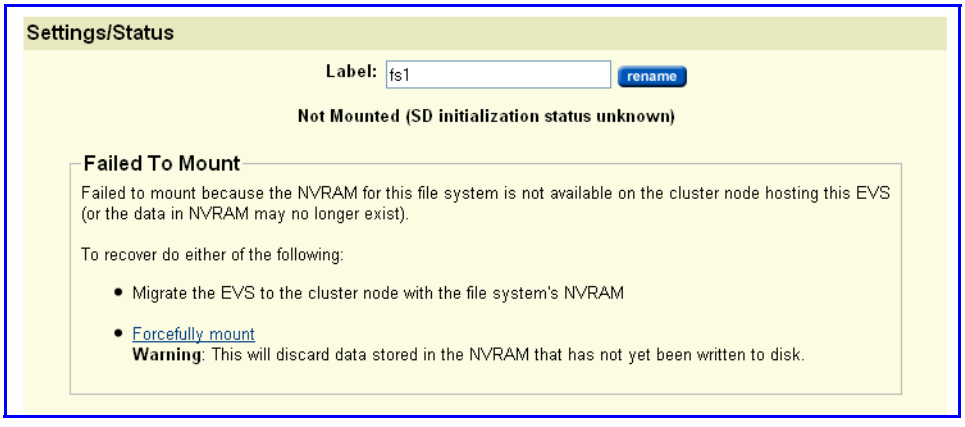

4. **Initiate recovery.**

Depending on the configuration of your system, and the reason the file system failed to mount, you may have several recovery options:

- *If the server is part of a cluster,* you may be able to migrate the assigned EVS to another cluster node, then try to mount the file system. This can become necessary when another node in the cluster has the current available data in NVRAM that is necessary to replay write transactions to the file system following the last checkpoint. An EVS should be migrated to the cluster node that mirrors the failed node's NVRAM (for more information on NVRAM mirroring, see Buffering in a Cluster Configuration, on page 328. For more details on migrating EVSs, see Migrating Virtual Servers (EVSs) within a Cluster, on page 422.
- *If the first recovery attempt fails*, click the **Forcefully mount** link. This will execute a file system recovery without replaying the contents of NVRAM.

Caution: Using the Forcefully mount option discards the contents of NVRAM, data which may have already been acknowledged to the client. Discarding the NVRAM contents means that all write operations in NVRAM (those write operations not yet committed to disk) are lost. The client will then have to resubmit the write request. **Use the Forcefully mount option only upon the recommendation of** SGI Global Services**.**

Recovering a File System

Following a storage subsystem failure, it may be necessary to recover file systems. WFS-2 file systems have better file system resiliency following storage subsystem failures than WFS-1 file systems. File systems formatted using WFS-2 have underlying redundant metadata structures, which allow them to be recovered from checkpoints or snapshots. File systems formatted using WFS-1 do not have the same underlying redundant metadata structures, and so cannot be recovered from checkpoints or snapshots.

Restoring a File System from a Checkpoint (WFS-2 Only)

File system corruption due to an event (such as RAID controller crash, storage system component failure, or power loss) often affects objects that were being modified around the time of the event.

WFS-2 preserves multiple checkpoints for a file system. By default, a WFS-2 file system is configured to keep up to 128 checkpoints. The maximum number of checkpoints supported is 1024. The number of checkpoints preserved is configurable when the file system is formatted, but, once set, the number of checkpoints cannot be changed.

When a checkpoint completes, rather than immediately freeing the storage used for the previous checkpoint, WFS-2 maintains a number of old checkpoints. As each new checkpoint completes, the oldest checkpoint is overwritten. This means that there can be multiple checkpoints on-disk, each of which is complete and internally consistent point-in-time view of the file

system. If necessary, the file system can be restored to any of these checkpoints.

In the case of file system corruption, if there are enough checkpoints on disk, it may be possible to roll back to a previous checkpoint, pre-dating the event that caused the corruption and restoring the file system using the uncorrupted checkpoint.This may be possible even if this event occurred up to a few minutes before the file system was taken off-line.

To restore a file system to a previous checkpoint, use the `fs-checkpoint-health` and the `fs-checkpoint-select` commands. Refer to the *Command Line Reference* for more information about these commands.

Note the following:

- Restoring a file system using a checkpoint does not affect snapshots taken prior to the checkpoint being restored, but, like any other file system update, snapshots taken after that checkpoint are lost.
- After restoring to a checkpoint, it is possible to restore again, to an older checkpoint and, if the file system has not been modified, restore again, to a more recent checkpoint. So, for example, it is possible to mount the file system in read only mode, check its status, and then decide whether to re-mount the file system in normal (read/write) mode or to restore to a different checkpoint.

Caution: Once you mount a restored file system in normal (read/write) mode, you cannot restore to a later checkpoint.

File System Recovery From a Snapshot (WFS-2 Only)

It is possible that, although corruption has occurred in the live file system, a good snapshot still exists. If so, it may be preferable to go back to this snapshot, with some loss of data, rather than incur the downtime that might be required to fix the live file system.

Recovering a file system from a snapshot makes it possible to roll back the file system to the state that it was in when a previous snapshot was taken.

File system recovery from a snapshot is a licensed feature, which requires a valid "FSRS" license on the server/cluster.

Note: You can recover a file system from a snapshot only when **at least** the configured number of preserved file system checkpoints have been taken since that snapshot was taken. For example, if a file system is configured to preserve 128 checkpoints (the default), then you can recover the file system from a snapshot only after a minimum of 128 checkpoints have been taken after the snapshot. If less than the configured number of checkpoints have been taken since the snapshot, you can either recover from an earlier snapshot or recover the file system from a checkpoint (as described in Restoring a File System from a Checkpoint (WFS-2 Only), on page 159).

The following file system rollback considerations apply:

- File system rollback can be performed even if the live file system is corrupted.
- All snapshots are lost after the rollback.
- Even though the file system recovery happens very quickly, no new snapshots can be taken until all previous snapshots have been discarded. The time required before a new snapshot can be taken depends on the size of the file system, not on the number of files in the file system.

Note: Once you have recovered a file system from a snapshot, you cannot undo the recovery or recover again to a different snapshot or checkpoint.

To roll back a file system from a snapshot, use the `snapshot-recover-fs` command. Refer to the *Command Line Reference* for more information about this command.

Automatic File System Recovery (WFS-1 and WFS-2)

The command line interface for file system recovery accommodates WFS-2 file systems. The `fixfs` utility is the main file system recovery tool for WFS-1 and WFS-2 file systems, but it should only be used under the supervision of SGI Global Services personnel. An additional tool is available to kill all current snapshots, that is the `kill-snapshots` command (refer to the *Command Line Reference* for more information about this command).

`Fixfs` is capable of repairing a certain amount of non-critical metadata, for example performing orphan recovery. At all stages that have the potential to last longer than a few minutes, `fixfs` provides accurate progress reporting, and the option to abort the fix. For some operations, `fixfs` will also provide an estimate of time until the completion of the operation.

The strategy used by `fixfs` to repair file systems can be summarized as:

- `Fixfs` is the only recovery tool to be used if a file system is experiencing corruption. The default `fixfs` behavior may be modified by various command line switches, but often the required switch is suggested by `fixfs` during or at the end of a previous run.
- Where possible, `fixfs` will run with the file system in any state (there will be no need to perform file system recovery first, so that there's no need to worry about what happens if recovery cannot complete due to corruption). Where not possible (for example, if the file system is already mounted, or is marked as "failed" or "requires expansion"), `fixfs` will not run. When `fixfs` does not run, it will give a clear indication of what needs to be done to get to the point where it can run.
- By default, `fixfs` will only modify those objects which are lost or corrupted.
- By default, `fixfs` will only bring the file system to the point where it can be mounted.
- Snapshots are considered expendable and are deleted.

In addition to `fixfs`, the `kill-snapshots` utility is available to help fix a file system. Note that, when `kill-snapshots` is used on WFS-2 file systems, the benefits of multiple file system checkpoints are lost for 5-10 minutes after `kill-snapshots` is invoked.

Using WORM File Systems

A file system is designated as WORM at the time of creation. When creating a WORM file system, start with a non-strict WORM file system to test your retention policy; then, when you are ready to deploy a compliance solution, create a strict WORM file system. For more information, see Creating a File System, on page 138.

Note: Any existing non-WORM file system can be reformatted as a WORM file system. Reformatting a file system to use a different file system type must be done at the CLI. For detailed information on this process, run `man format`.

Designating Retention Date

Before marking a file as WORM, designate its retention date. To configure the retention date, set the file's "last access time" to a specific time in the future. The "last access time" can be set using the Unix `touch` command; for example:

```
touch -a MMDDhhmm[YY] ./filename
```

Should the retention date be less than or equal to the current time, it will never expire.

After marking a file as WORM, permissions cannot be altered until the file reaches its retention date. Once a WORM file reaches its retention date, you can read the contents of the file, or delete the entire file, but its permissions *cannot* be changed to allow read-write access. When write access is granted, the file can be deleted; however, the contents of the file will still remain unavailable for modification.

Marking a file as WORM

Once the retention date has been set, a file can be marked as WORM. To mark a file as WORM, set the permissions of the file to *read-only*:

- *From a Unix client,* remove the write attribute permissions.
- *From a Windows client,* mark the file as read-only through the file's properties:

- 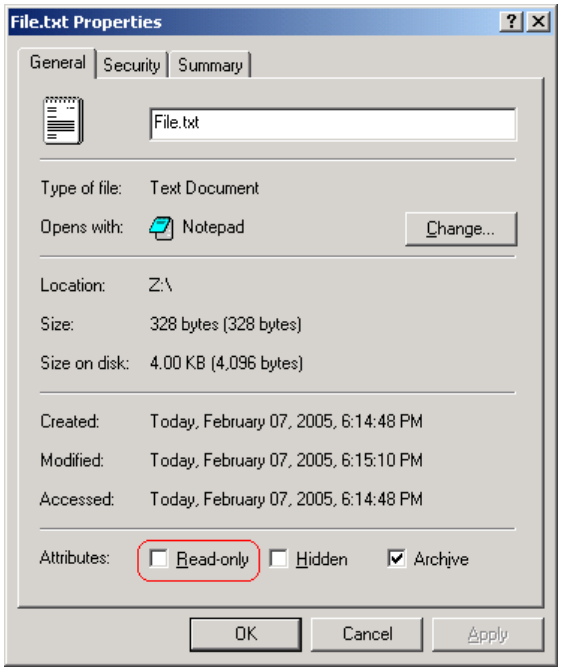

Managing Usage Quotas

To view usage quotas, navigate to the **Storage Management** page, then click **Quotas by File System** to display the page:

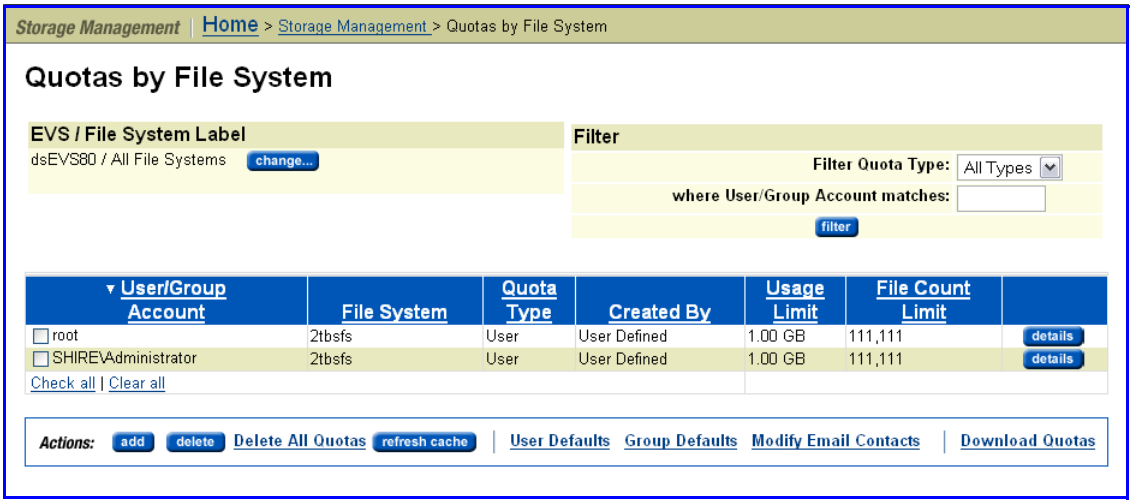

which lists usage quotas for the selected file system:

- *To sort in ascending or descending order by any column,* click the column header.
- *To display a different set of quotas,* click **change** and select a different file system.

The following table describes the fields in this page:

Item	Description
EVS/File System	The name of the selected EVS and file system. Click **change** to select an different File System.
Filter	Filters can be defined to reduce the number of quotas or virtual volumes displayed on the page. Click **filter** to list quotas based on the filter criteria specified.
User/Group Account	A quota name may consist of: • A CIFS domain and user or group name, such as `bb\Smith` or `bb\my_group` (where `bb` is a domain, `Smith` is a user and `my_group` is a group). • An NFS user or group such as `richardb` or `finance` (where `richardb` is an NFS user and `finance` is an NFS group). A name may be '0' (if the quota was created for the owner of the directory at the root of the virtual volume).
File systems	The file system on which the quota applies
Quota Type	Type of file system activity. Possible values are *User, Group*, or *Virtual Volume*. Virtual volume applies to anyone initiating activity in the entire virtual volume, and only one quota with this target type may exist on each virtual volume.
Created By	Method of quota creation. Possible values are *Automatically Created* (created using a quota default) or *User Defined* (where the quota was set uniquely for one particular quota).
Usage Limit	Overall limit set for the total size of all files in the file system owned by the target of the quota.
File Count Limit	Overall limit set for the total number of files in the file system owned by the target of the quota.
Actions	**details** allows you to edit details about the selected quota. Add a quota by clicking **add**. Delete a quota (or selection of quotas) by filling its checkbox and clicking **delete**. **Delete All Quotas** deletes all of the current quotas for the Virtual Volume. **refresh cache** clears the SMU's cache and repopulates it with relevant objects. (This is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.) **User Defaults** allows you to set, edit, or reset the user defaults. **Group Defaults** allows you to set, edit, or reset the group defaults **Modify Email Contacts** allows you to edit the list of contacts who are notified when quotas are reached. **Download Quotas** (Not File System Quotas): The quotas for this Virtual Volume may be downloaded as a .csv file.

To Set User and Group File System Quota Defaults

From the **Quotas by File System** page, click **User Defaults** (or **Group Defaults**) to display the **User** (or **Group**) **File System Quota Defaults** page (illustrated for User Defaults):

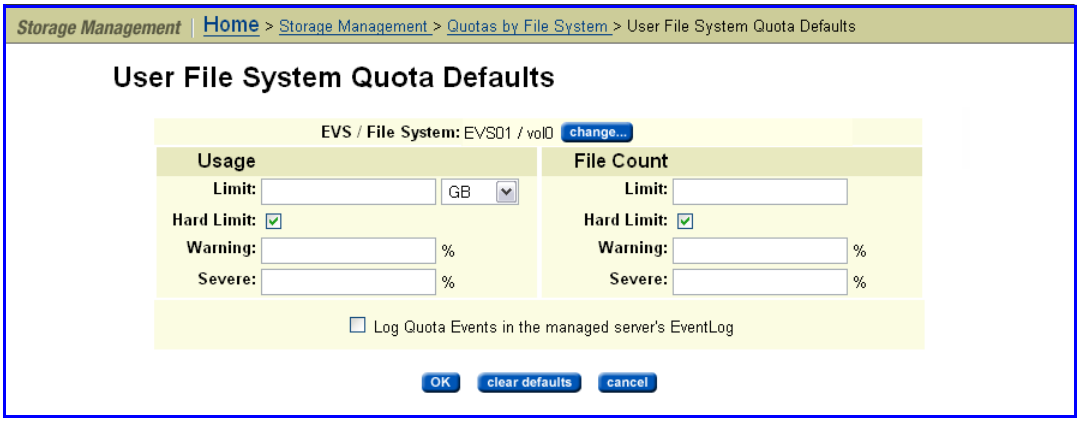

Note: The **Group File System Quota Defaults** page is the same as those detailed in the previous table for users, with the addition of the checkbox **Automatically create quotas for Domain Users,** which allows the creation of default quotas for the group "Domain Users". By default, every NT user belongs to the group "Domain Users", so enabling this option effectively includes every NT user in the quota, unless each user's primary group has been explicitly set.

The following table describes the fields in this page:

Item/Field	Description
EVS/File System	The EVS and file system on which the user file system quota applies.
Usage	This section displays the current quota settings that are related to usage.
Limit	Amount of space to enable in Bytes: *KB*, *MB*, *GB* or *TB*.
Hard Limit	If selected, the amount of space specified in the **Limit** field may not be exceeded.
Warning	Percentage of the amount of space specified in the **Limit** field at which a *Warning* alert will be sent.
Severe	Percentage of the amount of space specified in the **Limit** field at which a *Severe* alert will be sent.
File Count	This section displays the current quota settings that are related to file count.
Limit	Enter a maximum number of files to enable for this quota.
Hard Limit	If selected, the number of files specified in the **Limit** field may not be exceeded.

Item/Field	Description
Warning	Percentage of the amount of space specified in the **Limit** field at which a *Warning* alert will be sent.
Log Quota Events in the managed server's Event Log	Filling this check box sets the default for all users or groups to have quota events logged in the server's event log.

Note: If zero (or nothing) is left in a field, that entry will be regarded as 'not set'.

When all necessary fields have been completed, click **OK**.

Note: If the User Defaults are to be cleared, so that no further default quotas will be created in the file system, click **clear defaults**. This will convert any existing "Automatically Created" user quotas into "User Defined" user quotas.

To Add a File System Quota

From the **Quotas by File Systems** page, display the **Add File System Quota** page by clicking **add.**

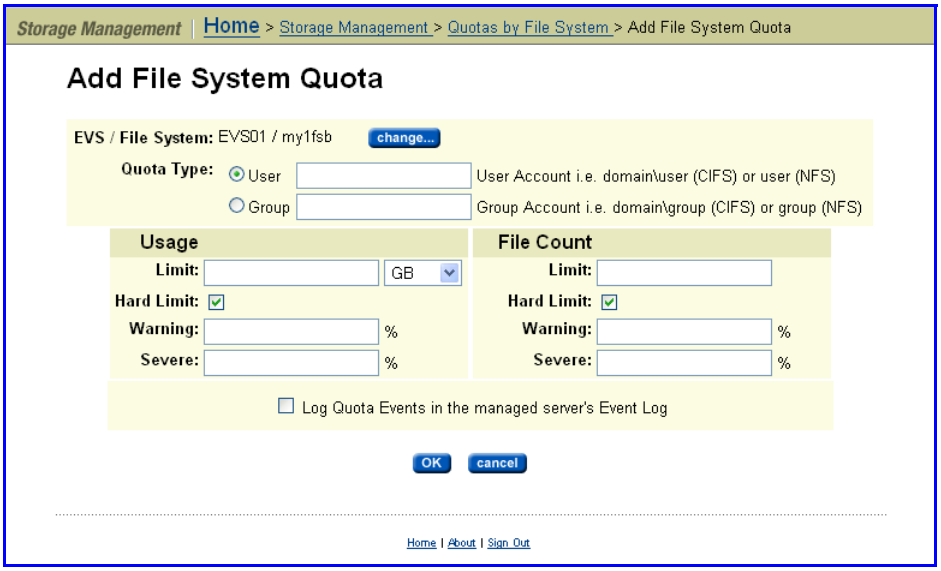

The table below describes the fields on this page:

Item/Field	Description
File systems	The File System on which the quota applies
Quota Type	Type of source of virtual volume activity. Possible values are *User, Group,* or *Virtual Volume.* Virtual volume applies to anyone initiating activity in the entire virtual volume, and only one quota with this target type may exist on each virtual volume.

Item/Field	Description
User/Group Account	A quota name may consist of: • A CIFS domain and user or group name, such as `bb\Smith` or `bb\my_group` (where `bb` is a domain, `Smith` is a user and `my_group` is a group). • An NFS user or group such as `richardb` or `finance` (where `richardb` is an NFS user and `finance` is an NFS group). A name may be empty (if the quota is a virtual volume quota) or '0' (if the quota was created for the owner of the directory at the root of the virtual volume).
Usage	
Limit	Amount of space to enable in Bytes: *KB*, *MB*, *GB* or *TB*.
Hard Limit	If selected, the amount of space specified in the **Limit** field may not be exceeded.
Warning	Percentage of the amount of space specified in the **Limit** field at which a *Warning* alert will be sent.
Severe	Percentage of the amount of space specified in the **Limit** field at which a *Severe* alert will be sent.
File Count	
Limit	Enter a maximum number of files to enable for this quota.
Hard Limit	If selected, the number of files specified in the **Limit** field may not be exceeded.
Warning	Percentage of the amount of space specified in the **Limit** field at which a *Warning* alert will be sent.
Log Quota Events in the managed server's Event Log	Filling this check box sets the default for all users or groups to have quota events logged in the server's event log.

Note: If zero (or nothing) is left in a field, that entry will be regarded as 'not set'.

When all necessary fields have been completed, click **OK**.

To Modify a File System Quota

In the **Quotas by File Systems** page, click **details** for the quota to be modified. The details of this page are the same as those detailed in the previous table for users, with the addition of the checkbox **Automatically create quotas for Domain Users** that allows the creation of default quotas for the group "Domain Users". By default, every NT user belongs to the group "Domain Users", so enabling this option effectively includes every NT user in the quota (unless each user's primary group has been explicitly set). When all necessary fields have been completed, click **OK**.

To Delete a File System Quota

On the **Quotas by File Systems** page, select the quota or quotas, then click **delete**.

Managing Virtual Volumes

To view virtual volumes, navigate to the **Storage Management** page, then click **Virtual Volumes & Quotas** to display the **Virtual Volumes** page:

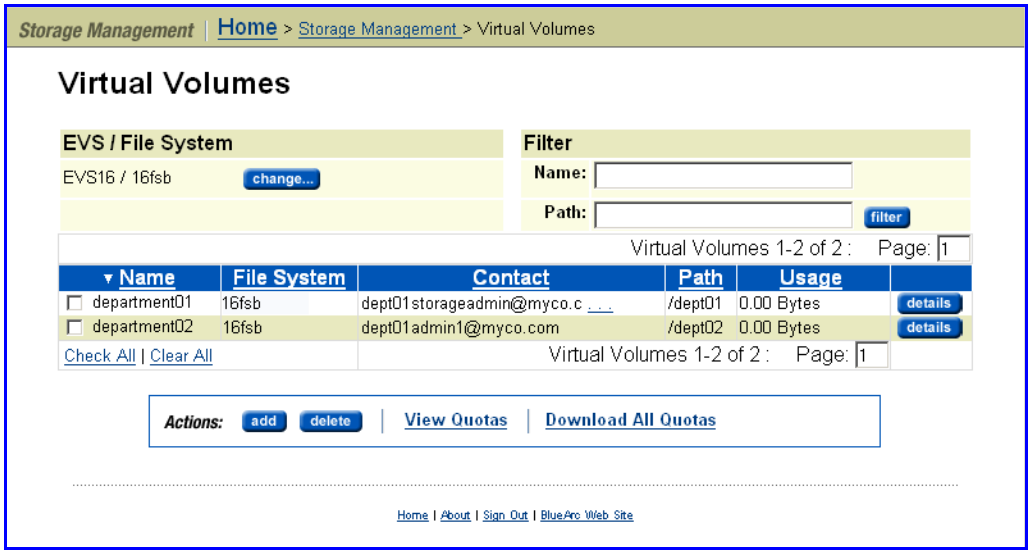

which lists virtual volumes for the selected file system:

- *To sort in ascending or descending order by any column,* click column header.
- *To display a different set of virtual volumes,* click **change** and select a different EVS/file system.

The following table describes the fields in this page:

Item/Field	Description
EVS/File System	The name of the selected EVS and file system.
Filter	Filters can be defined to reduce the number of virtual volumes displayed on the page and can be configured based on the name or the path.
Name	Name of the virtual volume.
File System	Name of the file system.
Contact	Contact email address for information and alerts about virtual volume activity.  **Note:** Only the first contact email address is shown; to view the full set of contacts, or otherwise modify the virtual volume, click **details.**
Path	Directory on which the virtual volume has been created.
Usage	Amount of data in the virtual volume.

The following **Actions** are available:

- Click **details** to display the **Virtual Volume** page for the selected virtual volume.
- Click **add** to display the **add Virtual Volume** page and see Adding a Virtual Volume, on page 169.
- Click **delete** to delete the selected virtual volume(s) and see Deleting a Virtual Volume, on page 171.

The following **Shortcuts** are available:

- Click **View Quotas** to display the **Quotas** page for the selected virtual volume.
- Click **Download All Quotas** to download a CSV (comma separated values) file listing all virtual volumes' configured quotas.

  The saved quota information includes: *Quota Type, Created By, Usage, Usage Limit, Usage Hard Limit, Usage Reset (%), Usage Warning (%), Usage Severe (%), File Count, File Count Limit, File Count Hard Limit, File Count Reset (%), File Count Warning (%),* and *File Count Severe (%).*

Adding a Virtual Volume

To add a virtual volume:

1. **Navigate to the add Virtual Volume page.**

   From the **Virtual Volumes** page, click **add** to display the **Add Virtual Volume** page:

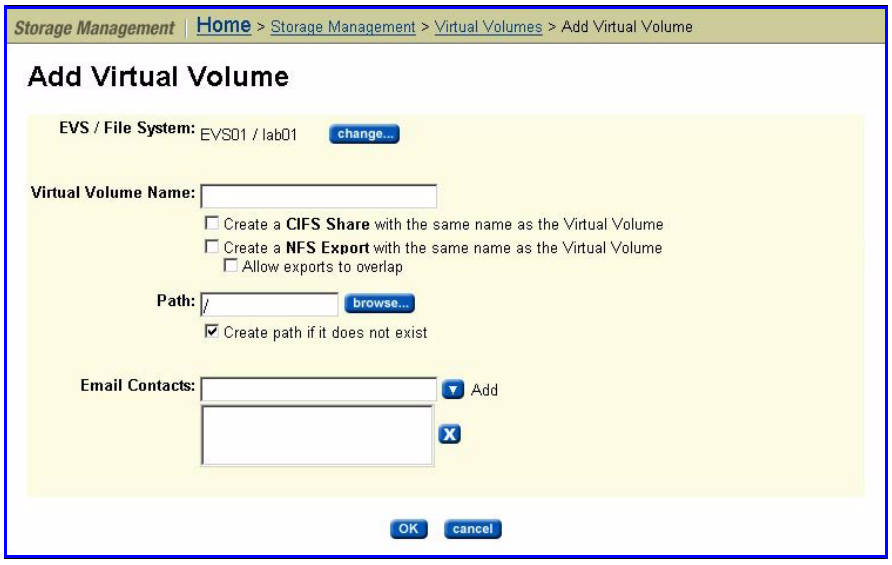

2. **Enter the requested information.**

   The following table describes the fields on this page:

Item/Fields	Description
EVS/File System	The EVS and the file system to which to add this virtual volume. If the virtual volume will be added to a different EVS/file system, click **change** and select an EVS/file system.
Virtual Volume Name	The name may consist of up to 128 characters. The following characters are not supported: ?*=+[];:/,<>\ | The name "A$" is reserved for the Windows Event Viewer, and may not be used.
Create a CIFS Share or NFS Export with the same name as the Virtual Volume	If a share or export with the same name as the virtual volume does not exist, selecting this checkbox ensures its creation. This is only intended for convenience in accessing the virtual volume through CIFS or NFS.  **Note:** The CIFS Share or NFS Export name may not exceed 80 characters, and the CIFS Share name "A$" may not be used, as it is reserved for the Windows Event Viewer.
Allow exports to overlap	As overlapping exports can potentially expose security loopholes, the condition can be tested for and, if found, the export creation can be denied. Select this checkbox to prevent this check and allow overlapping NFS exports to be created.
Path	Directory in the file system that will be the 'root' of the virtual volume; for example, `/company/sales`. All subdirectories of this path will be a part of this virtual volume. Please note: • **Unchanging path.** Once created, the path may not be changed. • **File system root is off limits.** Virtual volumes cannot be created at the root of the file system (/). They must be applied to the directories in the file system. • **Empty host file system directory required.** Virtual volumes can only be created and assigned to empty directories. To create a virtual volume on a directory that contains data, first move the data out of the directory; once empty, the virtual volume can be created and assigned to that directory, then the data can be moved back.
Email Contacts	Email contacts to receive information about virtual volume usage. Enter each email address in the box, then click **Add** to append it to the list. If an address is entered in error, select it from the list, and click **X** to delete. *To configure email notification of threshold alerts,* designate explicit email recipients (e.g. admin@company.com) to receive email notification any time a defined threshold has been reached. *To send email to all affected user accounts when their user quota has been reached,* add an email address beginning with * to the Email Contacts list (e.g. *@company.com). Email lists are limited to a maximum of 512 characters.  **Note:** If no email contacts are specified for the virtual volume, the server generates events for quota warnings. To generate events *in addition to* email alerts, go to the server's command line interface and issue the command `quota-event--on`.

1. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

   Note: The virtual volume may be subsequently modified by clicking **details** in the **Virtual Volume** list page.

Modifying a Virtual Volume

To modify a virtual volume:

1. **Navigate to the Virtual Volumes page.**

   From the **Storage Management** page, click **Virtual Volumes & Quotas** to display the **Virtual Volumes** page:

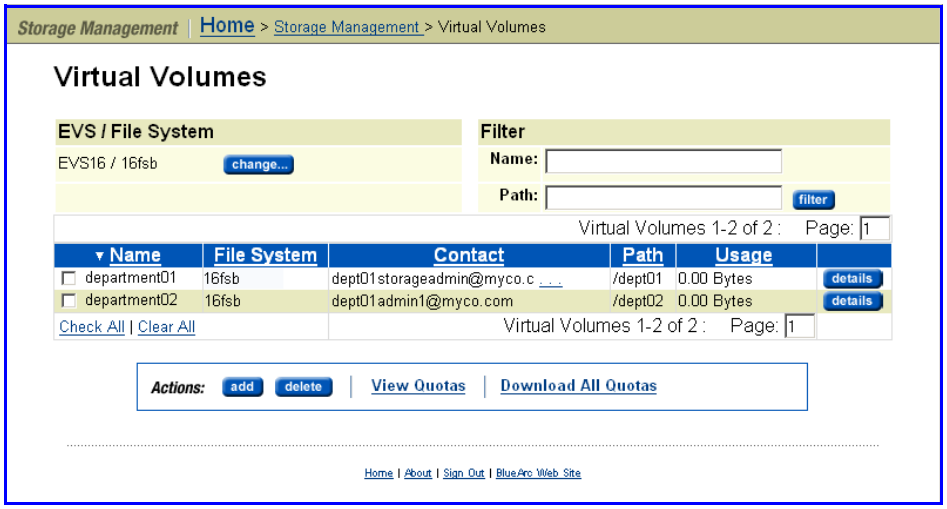

2. **Modify settings.**

   Click **details** for a virtual volume, then modify settings.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Deleting a Virtual Volume

To delete a virtual volume:

1. **Navigate to the Virtual Volumes page.**

   From the **Storage Management** page, click **Virtual Volumes & Quotas** to display the **Virtual Volumes** page.

2. **Select a virtual volume.**

   Select a virtual volume or **Check All** for all virtual volumes.

   Note: A virtual volume can only be removed from a directory when the directory is empty. To delete a virtual volume which is assigned to a directory that contains data, first remove the data, then delete the virtual volume.

3. **Delete.**

   On clicking **delete,** a warning will displayed asking for confirmation that this action is definitely required. Click **OK** to continue deleting the virtual volumes.

Managing Quotas on Virtual Volumes

Three types of quotas are maintained for each virtual volume:

- **Explicit User/Group Quotas.** A quota explicitly created to impose restrictions on an individual user or group, defining a unique set of thresholds.
- **Default User/Group Quotas.** A quota set automatically for all users and groups that do not have explicit quotas, set by defining a set of **Quota Defaults** (thresholds) for creating a quota automatically when a file is created or modified in the virtual volume.

  Default quotas for virtual volumes operate in the same manner as those defined for file systems. User (Group) quota defaults define a set of thresholds for creating a quota for a user (or group) the first time that user (or group) saves a file in the virtual volume.

  Initially, quota defaults are not set. When activity occurs in the virtual volume, it is tracked, but quotas are not automatically created. When at least one threshold is set to a non-zero value, a User or Group quota (as appropriate) will be created for the owner of the directory at the root of the virtual volume.
- **Virtual Volume Quotas.** A Virtual volume quota tracks the space used within a specific directory on the virtual volume. A quota can be explicitly created to define a set of thresholds restricting all operations in the virtual volume, unrelated to which user or group initiated them.

Note: Quotas track the number and total size of all files. At specified thresholds, emails alert the list of contacts associated with the virtual volume and, optionally, *Quota Threshold Exceeded* events are logged. Operations that would take the user or group beyond the configured limit can be disallowed by setting hard limits.

When *Usage* and *File Count* limits are combined, the server will enforce the first quota to be reached.

Viewing/ Modifying Virtual Volume Quotas

To view/modify virtual volume quotas:

1. **Navigate to the Quotas page.**

From the **Storage Management** page, select **Virtual Volumes & Quotas**, then select the virtual volume for which you want to see quotas, and click **View Quotas** to display the **Quotas** page:

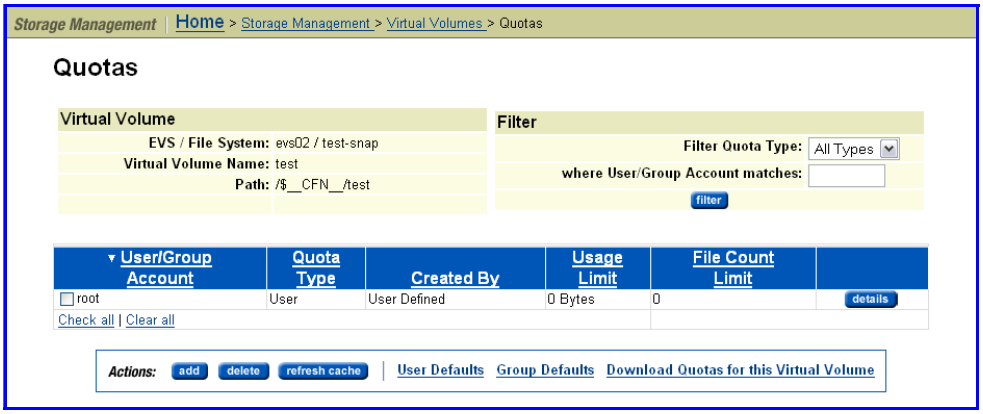

2. **Filter to display a quota or group of quotas.**

   The **Quotas** page displays only 20 quotas. Navigate to additional pages of quota listings using the links at the top and bottom of the list. Hovering over these links displays screen tips that describe their use.

   The table below describes the contents of the **Quotas** page:

Section	Description
Virtual Volume	Identifies the virtual volume to which these quotas apply: • **EVS/File System:** EVS and file system on which the virtual volume resides. • **Virtual Volume Name:** Name of the virtual volume. • **Path:** Directory on which the virtual volume has been created.
Filter	Since many user/group quotas can exist on a virtual volume, the server provides a way to filter the list. • **Filter Types**: You can select from *All Types (default)*, *Users* or *Groups*. • **where User/Group Account matches**: In this field, enter a name to be matched. The wildcard character * may be used. Click **filter** to limit the quotas displayed in the list below to those matching the filter criteria specified above.

Section	Description
User/Group Account (also known as the target)	A quota name may consist of: • A CIFS domain and user or group name, such as `bb\Smith` or `bb\my_group` (where `bb` is a domain, `Smith` is a user and `my_group` is a group). • An NFS user or group such as `richardb` or `finance` (where `richardb` is an NFS user and `finance` is an NFS group). A name may be empty (if the quota is a virtual volume quota) or '0' (if the quota was created for the owner of the directory at the root of the virtual volume).
Quota Type	Type of source of virtual volume activity. Possible values are *User, Group,* or *Virtual Volume.* Virtual volume applies to anyone initiating activity in the entire virtual volume, and only one quota with this target type may exist on each virtual volume.
Created By	Method of quota creation. Possible values are *Automatically Created* (created using a quota default) or *User Defined* (where the quota was set uniquely for one particular quota).
Usage Limit	Overall limit set for the total size of all files in the virtual volume owned by the target of the quota.
File Count Limit	Overall limit set for the total number of files in the virtual volume owned by the target of the quota.

3. **Optionally, modify settings.**

   The following **Actions** are available:

   - Click **add** to add a new quota, and see the instructions in Adding a Quota, on page 176.
   - Click **delete** to delete a selected quota.
   - Click **refresh cache** to clear the SMU's cache and repopulate the cache with the relevant objects. (This is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.)

   The following **Shortcuts** are available:

   - Click **User Defaults** to go to the **User Quota Defaults** page, where you can set or change the defaults for users.
   - Click **Group Defaults** to go to the **User Quota Defaults** page, where you can set or change the defaults for groups.
   - Click **Download Quotas** to download a CSV (Comma Separated Value) file containing all available quota information for the virtual volume.

Setting User/ Group Defaults

To set user /group defaults:

This procedure illustrates the **User Default** page. The **Group Default** page is identical, except for the **Automatically create quotas for Domain Users** checkbox. This option allows default quotas for the group "Domain Users" to be created. By default, every NT user belongs to the group "Domain Users", which includes every NT user in the quota unless each user's primary group has been explicitly set.

1. **Navigate to the User Quota Defaults page.**

   On the **Quotas** page, click **User Defaults** to display the **User Quota Defaults** page:

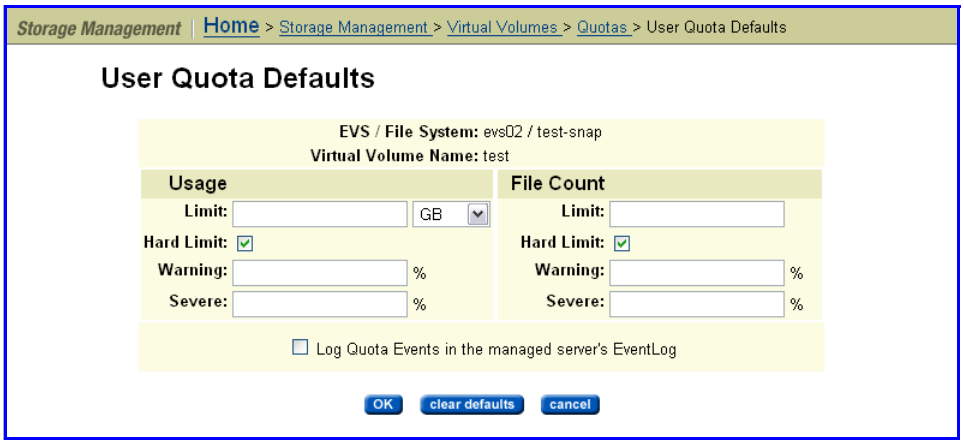

   On the **User Quota Defaults** page, an EVS/file system and a virtual volume Name will be displayed for the defaults.

2. **Enter the requested information.**

   If a zero (or nothing) is left in a field, that entry will be considered not set. For example, a File Count Limit of zero means that a quota created will not have a limit on the number of files it may contain. The *Warning* and *Severe* thresholds will also be considered not set.

   Note: To clear the user quota defaults, click **Clear Defaults,** which prevents additional user quota defaults from being created in the virtual volume. It also converts any existing "Automatically Created" user quotas into "User Defined" user quotas.

   The following table describes the fields in this page:

Item/Field	Description
EVS/File System	The EVS and file system to which the user quota applies.
Virtual Volume Name	Name of the virtual volume to which a user quota created using these defaults is assigned.
Usage	

Item/Field	Description
Limit	Amount of space to enable in Bytes: *KB*, *MB*, *GB* or *TB*.
Hard Limit	If selected, the amount of space specified in the **Limit** field may not be exceeded.
Warning	Percentage of the amount of space specified in the **Limit** field at which a *Warning* alert will be sent.
Severe	Percentage of the amount of space specified in the **Limit** field at which a *Severe* alert will be sent.
File Count	
Limit	Enter a maximum number of files to allow in this virtual volume.
Hard Limit	If selected, the number of files specified in the **Limit** field may not be exceeded.
Warning	Percentage of the number of files specified in the **Limit** field at which a *Warning* alert will be sent.
Severe	Percentage of the number of files specified in the **Limit** field at which a *Severe* alert will be sent.
Log Quota Events in the managed server's Event Log	Filling this check box sets the default for all users or groups to have quota events logged in the server's event log.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Adding a Quota

To add a quota:

1. **Navigate to the Quotas page.**

   From the **Storage Management** page, select **Virtual Volumes & Quotas**, then select the virtual volume for which you want to see quotas and click **View Quotas** to display the **Quotas** page:

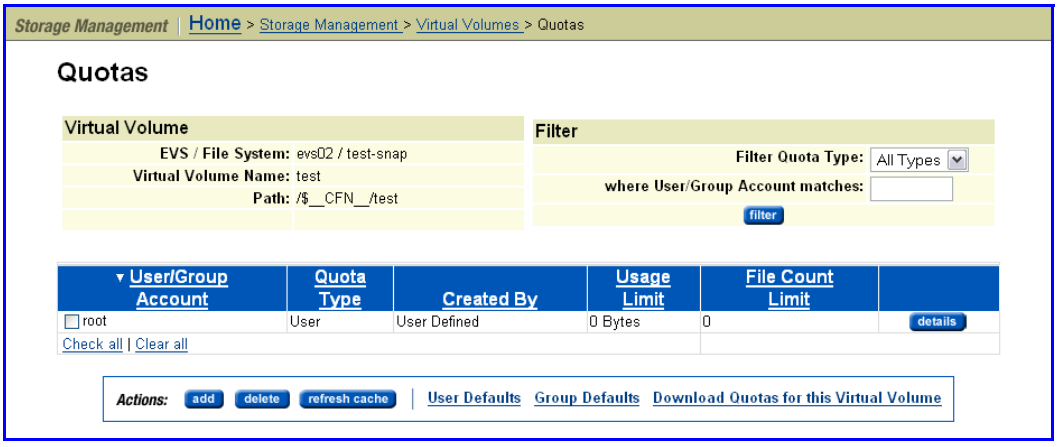

2. **Navigate to the add Quota page.**
   From the **Quotas** page, click **add**:

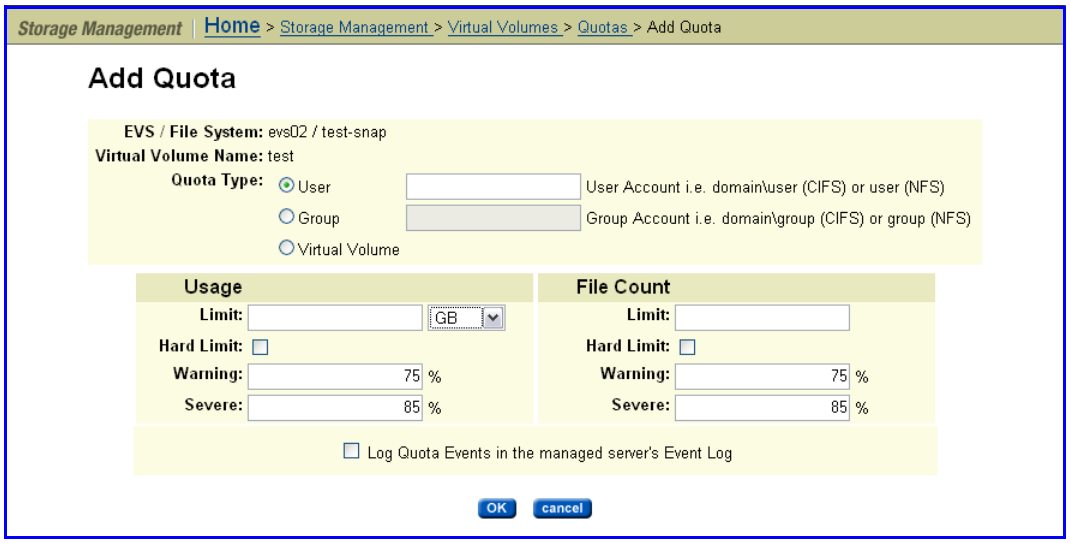

The following table describes the fields in this page:

Item/Field	Description
EVS/File System	Name of the EVS and the file system hosting the virtual volume to which the quota is being added.
Virtual Volume Name	Name of the virtual volume to which the quota is being added.
Quota Type	Type of source of virtual volume activity. Possible values are *User*, *Group*, or *Virtual Volume*.
User/Group Account	User/Group Account name consisting of: • A CIFS domain and user or group name, such as `bb\Smith` or `bb\my_group` (where `bb` is a domain, `Smith` is a user and `my_group` is a group). • An NFS user or group such as `richardb` or `finance` (where `richardb` is an NFS user and `finance` is an NFS group). If the virtual volume has been selected as the quota Target Type, it will not be possible to specify a Quota Name.
Usage	
Limit	The amount of space to enable in Bytes: *KB*, *MB*, *GB* or *TB*.
Hard Limit	If selected, the amount of space specified in the **Limit** field may not be exceeded.
Warning	Enter the percentage of the amount of space specified in the **Limit** field at which a *Warning* alert will be sent.

Item/Field	Description
Severe	Enter the percentage of the amount of space specified in the **Limit** field at which a *Severe* alert will be sent.
File Count	
Limit	Enter a maximum number of files to enable for this quota.
Hard Limit	If selected, the number of files specified in the **Limit** field may not be exceeded.
Warning	Enter the percentage of the number of files specified in the **Limit** field at which a *Warning* alert will be sent.
Severe	Enter the percentage of the number of files specified in the **Limit** field at which a *Severe* alert will be sent.
Log Quota Events in the managed server's Event Log	Fill this check box to have events due to this quota logged in the server's event log.

3. **To add a quota, enter the requested information.**

   If a zero (or nothing) is left in a field, that entry will be considered not set. For example, a File Count Limit of zero means that a quota created will not have a limit on the number of files it may contain. The *Warning* and *Severe* thresholds will also be considered not set.

4. **Save your changes.**

   Verify your additions or deletions, then click **OK** to save or **cancel** to decline.

Deleting a Quota

To add a quota:

1. **Navigate to the Quotas page.**

   From the **Storage Management** page, select **Virtual Volumes & Quotas**, then select the virtual volume for which you want to see quotas and click **View Quotas** to display the **Quotas** page:

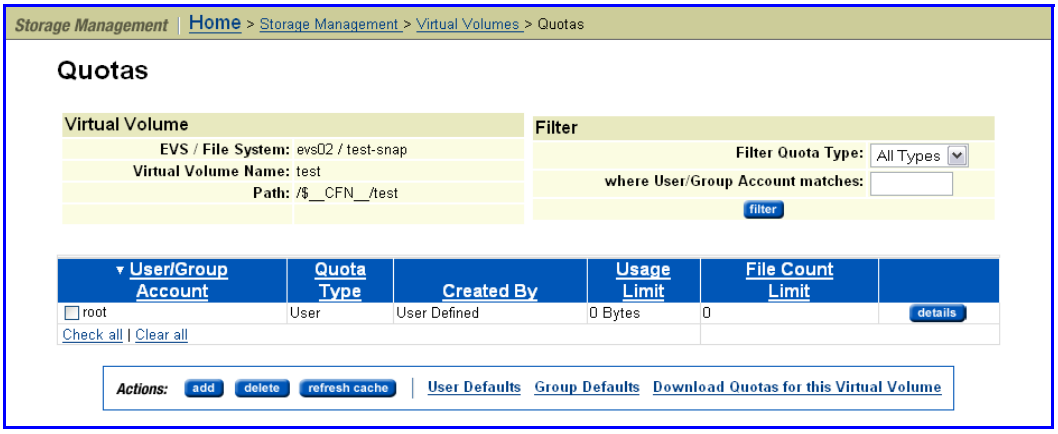

2. **Select the quota you want to delete.**

   Fill the check box(es) for the quota(s) you want to delete.

3. **To delete a selected quota, click delete.**

Note: Default quotas for the owner of the virtual volume's root directory will automatically reappear in the quota list after an explicit quota is deleted.

Exporting Quotas for All Virtual Volumes

To export quotas for all virtual volumes:

1. **Navigate to the Quota Management page.**

   From the **Storage Management** page, select **Virtual Volumes & Quotas** to display the **Virtual Volumes** page.

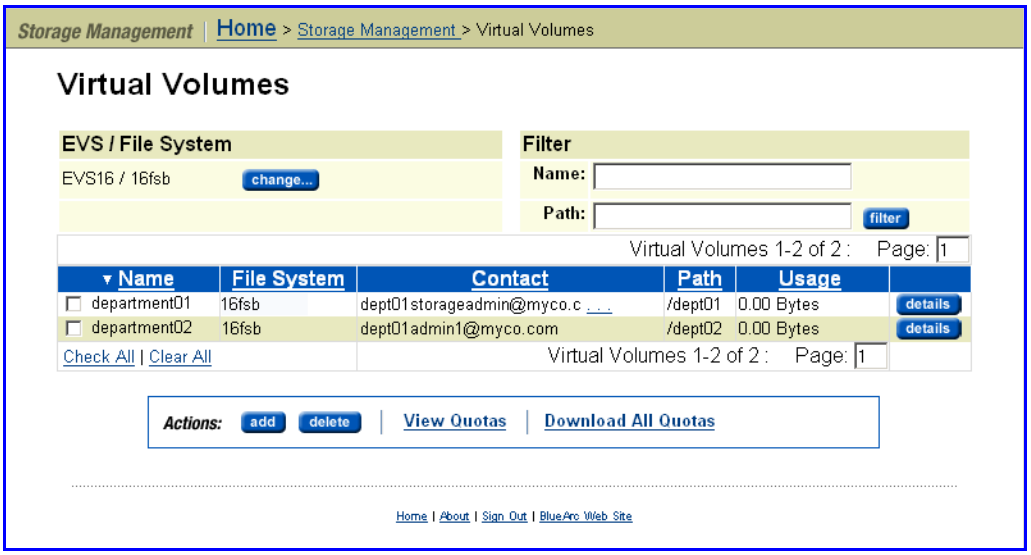

2. **Invoke the export.**

   Click **Download All Quotas.** A dialog box appears allowing you to save the quota information as a CSV (comma separated value) file. You can also choose to display the quota information in an application.

Exporting Quotas for a Specific Virtual Volume

To export quotas for a specific virtual volumes:

1. **Navigate to the Quota page.**

From the **Storage Management** page, select **Virtual Volumes & Quotas** to display the **Virtual Volumes** page, then select a virtual volume and click **View Quotas** to display the **Quotas** page:

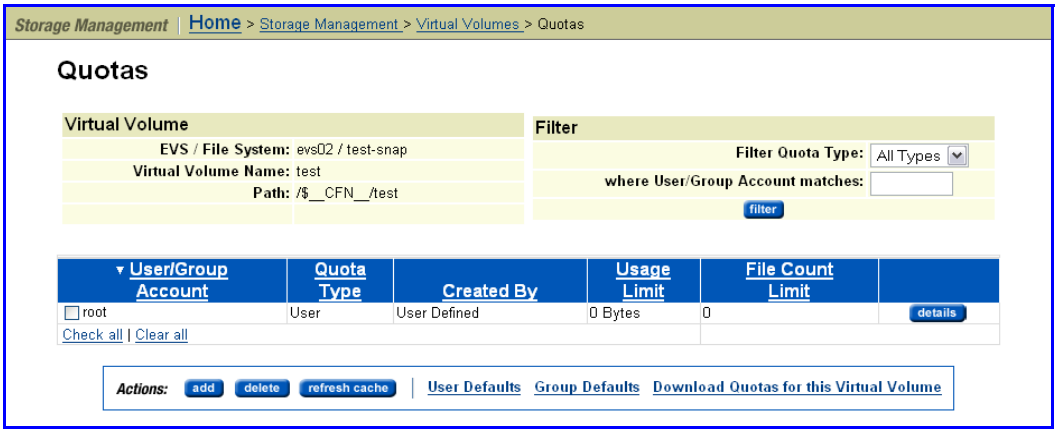

2. **Invoke the export.**

   Click **Download Quotas for this Virtual Volume**. A dialog box appears allowing you to save the quota information as a CSV (comma separated value) file. You can also choose to display the quota information in an application.

About the rquotad Service

The rquotad protocol has been implemented as a service on the storage server. It functions as a read-only protocol and is only responsible for reporting information about user and group quotas. Quotas can be created, deleted, or modified through the **Storage Management** section of the Web Manager.

A Unix/Linux NFS client can issue the `quota` command to retrieve information regarding quota usage of a user or group, based on their ID. The retrieved report contains *block count*, *file count*, *quota limits on both*, and other information (based on options invoked with the command). For accurate syntax, refer to the client's `man` pages, as implementation varies between client operating systems.

The server reports only **Hard Limit** quota information through rquotad. Three different quota limitations can be defined:

- User and group quotas to limit **space** and/or **file quantity** for **individuals** and/or **groups** *within a virtual volume.*
- User and group quotas to limit **space** and/or **file quantity** for **individuals** and/or **groups** *within an entire file system.*
- Virtual volume quotas to limit **space** and/or **file quantity** by *a virtual volume as a whole.*

Note: rquotad reports quota usage information on explicitly defined quotas and automatically created (default) quotas. Default quota information will be

reported if an explicit quota has not been defined.

rquotad Service Settings

The rquotad service can be configured to report quota information using one of two modes:

- **Restrictive mode.** For the user or group specified in the client-side `quota` command, the rquotad service reports the quota information for the quota with the most constraints.
- **Matching mode.** For the user or group specified in the client-side `quota` command, the rquotad service reports the quota information for the first quota that meets the parameters defined by the client-side `quota` command.

Note: If the rquotad service is disabled, all requests are rejected with an error code of "EPERM".

Restrictive Mode Operation

When in Restrictive mode, the rquotad service picks the first applicable quota threshold crossed. It enables the user to determine the amount of data that can be safely recorded against this quota before reaching its Hard Limit. This is the default configuration option for rquotad on the server.

Note: The restrictive mode option returns quota information combined from the quota that most restricts usage and the quota that most restricts file count. For example:

If the user quota allowed 10K of data and 100 files to be added, and the virtual volume quota allowed 100K of data and 10 files to be added, rquotad would return information stating that 10K of data and 10 files could be added. Similarly, if the user quota is 10K of data of which 5K is used, and the virtual volume quota is 100K of data of which 99K is used, rquotad would return information stating that 1K of data could be added.

The console command `rquotad` is provided to change between the two options, and also to disable access to quota information. For information on how to configure rquotad, please refer to the *Command Line Reference.*

Matching Mode Operation

When in Matching mode, the rquotad service follows a specific order to find a match for relevant quota information:

- If rquotad returns quota information for a user, it returns the *user's quota within the virtual volume* if it exists;
- Otherwise, it moves to the *user's file system quota* if that exists;
- If no file system quota exists for the user, then it will move to the *virtual volume quota*.

In this manner, rquotad keeps checking until a quota is found for the specified user or group. Once the **first** matching quota is found, rquotad stops searching and returns the quota information.

If a user does not have a specifically defined quota in a virtual volume, or in a file system, and the virtual volume quota allows all users 100K of data and 10 files, rquotad would return information stating that user's quota is 100K of data and 10 files. Similarly, if the user has a specified virtual volume quota that is 200K of data and 20 files, and a file system quota that is 400K of data and 40 files, rquotad would return information about only the first quota, stating that 200K of data and 20 files could be added.

Configuring the SGI Data Migrator

In order to use Data Migrator, you must define the following:

- **Data migration paths** from primary to secondary storage.

  Data migration paths define the relationship between primary and secondary storage. The primary and secondary storage defined in the data migration paths must be assigned to the same EVS.

  

  Note: Data Migrator is a licensed feature. For information about adding licenses, see Adding a License Key, on page 538.

- **Data migration rules,** which define the properties of files that will be migrated.
- **Data migration policies,** which define rules to apply to specific data migration paths based on the available free space on the source file system or virtual volume.

  Free space is calculated as follows:

  - For a file system, free space is the amount of unused space allocated to the file system (before it automatically expands, if automatic expansion is enabled for the file system).
  - For a Virtual Volume, if a quota has been defined, free space is the amount of unused space before reaching the usage limit of the quota for that Virtual Volume. If a quota has not been defined for the Virtual Volume, free space is the same as the free space for the file system.

- **Schedules** in which the frequency with which the data migration policies will be run.

Configuring Data Migrator Paths

Before Data Migrator can be used, primary and secondary storage must be identified:

- **Primary storage,** typically Fibre Channel disk arrays, will be the source for data migrations.

  

  Note: WORM file systems may not be specified as a Data Migrator source.

- **Secondary storage,** typically SATA disk arrays, will be the target for data migrations. Note that there are two types of paths to secondary storage:

  

  Note: WORM file systems may not be specified as a Data Migrator path.

- **Local paths**, which are paths to secondary storage attached to the same EVS, storage server, or cluster. Local paths may be added using the Web Manager interface, as described in Adding a Local Data Migration Path, on page 183.
- **External Paths**, which are paths to secondary storage that is attached to a remote server (a IS-NAS Server, a Titan Server, or another server using the NFS protocol). External paths may **not** be added using the Web Manager interface. Instead, you must use CLI commands, as described in Adding External Data Migration Paths, on page 186.

Once Data Migrator has been configured, data will be migrated from primary to secondary storage based on the *data migration rules,* freeing up space and extending the capacity of the primary storage.

Caution: *Dysfunctional backups alert!* Accessing files directly on secondary storage may alter access and modification times of the files, resulting in unexpected results when performing backups. The organizational structure of migrated data on secondary storage does not mirror that of primary storage.

Caution: *Lost access to migrated files alert!* If only the primary or only the secondary file system is moved to a different EVS, access to migrated files will be lost. If both the primary and the secondary file systems are moved to the same EVS, access to migrated files will be retained. When moving file systems, File System Relocation is the recommended method, because, when using File System Relocation, if the file system being moved is a member of a data migration path, both the data migration source file system and the target file system are relocated. See File System Relocation, on page 121 for more information.

Caution: *Exclusive migration pathing!* Once a migration path has been assigned to a virtual volume, a subsequent migration path cannot be created to its hosting file system. Also, once a migration path has been assigned to a file system, subsequent migration paths cannot be created from virtual volumes hosted by that file system.

Note: When defining data migration paths, specify a file system or virtual volume as the primary storage. Once a file system is selected as primary storage, that entire file system, including all virtual volumes, is included as a part of the data migration policy. Therefore, in order to create individual policies for different parts of a file system, create virtual volumes and assign each virtual volume a unique migration path.

Adding a Local Data Migration Path

To add a local data migration path:

1. **Navigate to the Data Migration Paths page.**

From the **Storage Management** page, click to display the **Data Migration Paths** page:

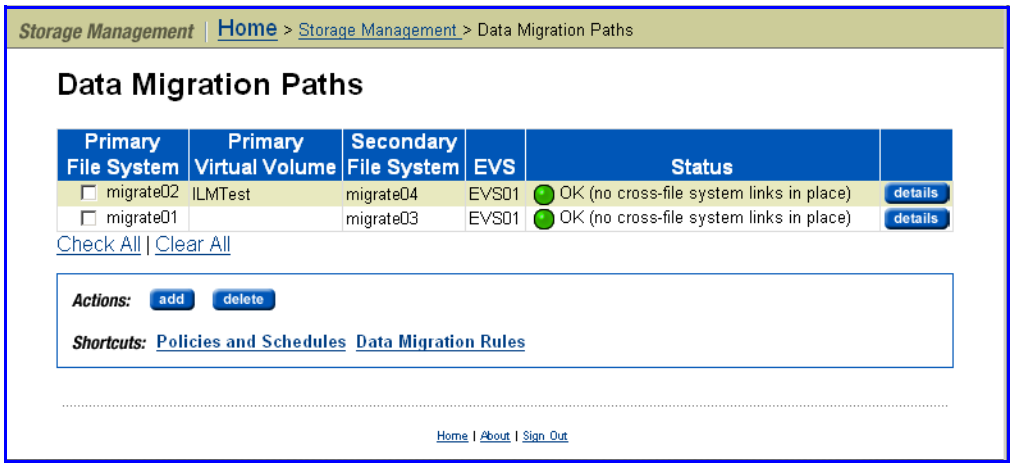

The fields on this page are described in the table below:

Item/Field	Description
Primary File System	Displays the file system from which data will be migrated.
Primary Virtual Volume	If a virtual volume has been selected as primary storage, this field displays the name of the virtual volume from which data will be migrated.
Secondary File System	The destination file system (on secondary storage) to which the data will be migrated.
EVS	Displays the EVS hosting the file system from which data will be migrated.
Status	Status of the data migration path. The status should always be OK; if otherwise, migrated files may be inaccessible.

2. **Add a Migration Path.**

   Click **add** to display the **add Data Migration Path** page:

Note: WORM file systems may not be specified in a Data Migrator path.

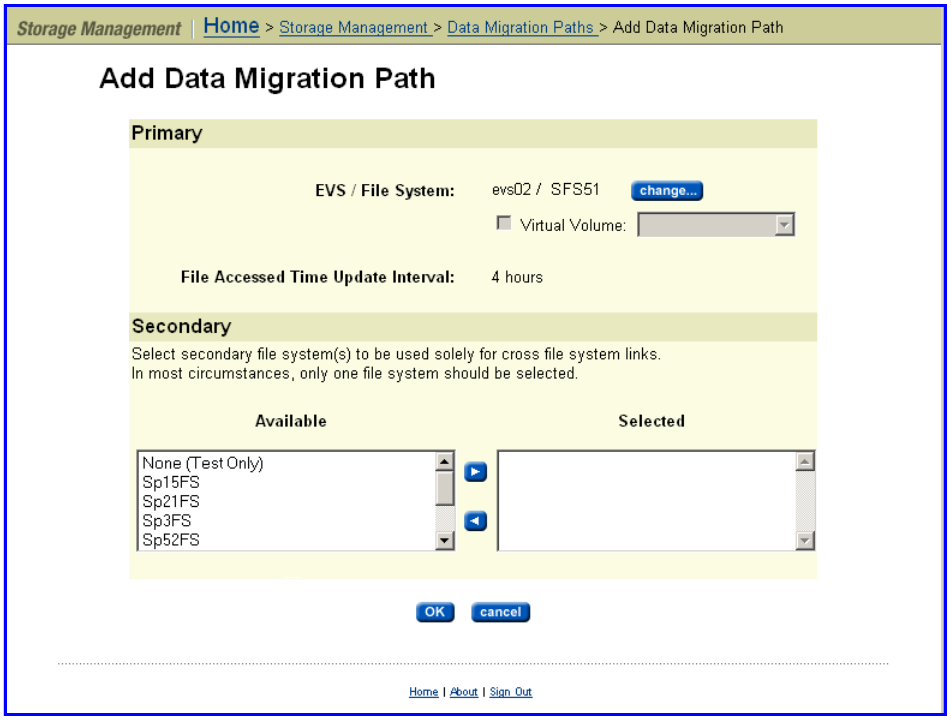

The fields on this page are described in the table below:

Section	Item/Field	Description
Primary	EVS/File System	EVS and file system on primary storage. This defines the *source* for the data migration path. To change the currently selected EVS and file system, click **change.**
	Virtual Volume	By default, data migration policies include the entire file system. To configure migrations on a per virtual volume basis, fill this checkbox and select the virtual volume to be used as the primary storage for this data migration path.
	File Accessed Time Update Interval	Currently configured Accessed Time Update Interval, defining the maximum elapsed time between file access and updating of the "last accessed time" field on the file. This value is important where migrations are based on an aggressive file access policy. For configuration details, see the *Command Line Reference*.

Section	Item/Field	Description
Secondary	Available	File systems to which the data could be migrated (the destination file system). Select the destination file system from the list. The file system(s) you select should be on secondary storage.  **Note:** When creating a policy for testing purposes, select "None (Test Only)." Running this policy will then determine the outcome of the migration operation without actually moving data. In most cases you should specify a single destination file system to create a "single-target" migration path. However, if the amount of data will be too large for a single target file system, you may want to nominate multiple file systems as targets to create a "multi-target" migration path. For "multi-target" migration paths, you should be aware of the following: • Data is distributed between the destination file systems based on the amount of free space available on those file systems. If the destination file system is expandable, the data distribution algorithm calculates free space not based on the file system's current size, but on the maximum size to which a file system can be expanded. • Once specified, multi-target paths may not be modified through Web Manager. If you need to change the migration path targets, for instance to add an additional destination file system, you must use the CLI command `migration-expand-target`.
	Selected	The file system(s) selected as the destination of the migration.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Adding External Data Migration Paths

External data migration paths are not defined through Web Manager. Instead, CLI commands are used to specify the path to external secondary storage. These commands are:

- `migration-add-external-path`
- `migration-change-external-path`
- `migration-delete-external-path`
- `migration-expand-external-path`

For information about these commands, refer to the *Command Line Reference*, or the man page for each command.

You should specify a unique external path for each file system being migrated to a remote server.

Note: Once an external migration path has been defined, it will be visible and available for selection in the Web Manager **Data Migration Paths** page.

Once the external migration path has been configured via the CLI, all remaining external migration management tasks may be performed through Web Manager, including specifying migration policies, rules, and schedules.

Note: When adding external migration paths, make sure that the remote server's target IP address or host name is correct and, if using a host name, make sure that the host name is resolvable (fully qualified domain names are also acceptable).

Data Migration Rules

The **Data Migration Rules** page lists all existing rules, and provides for removal of selected rules and creation of new rules. Data migration rules are used in conjunction with data migration paths to set up data migration policies.

Viewing the Data Migration Rules

To view the Data Migration Rules, navigate from the **Storage Management** to the **data migration Rules** page:

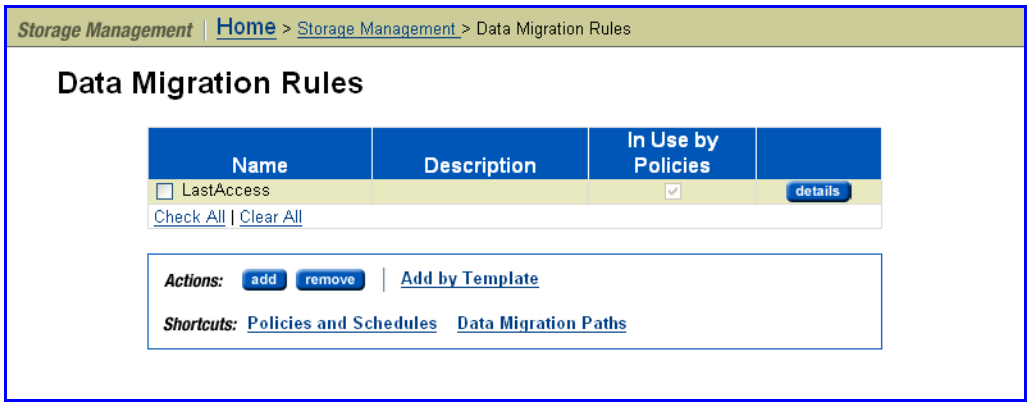

The fields on this page are described in the table below:

Item/Field	Description
Name	Displays the name given to the Rule. This is assigned when the Rule is created, and is used to identify the Rule when creating or configuring policies.
Description	A description given to the rule to help in identifying the criteria to be applied.
In Use by Policies	Fill the checkbox to indicate that the rule is being used by one or more policies.

The following **Actions** are available:

- Click **details** for a selected migration rule to display complete details.
- Click **delete** to remove a selected migration rule.
- To create custom rules that will exactly define the criteria by which files will be migrated, click **add** and refer to the instructions in Adding a Custom Data Migration Rule, on page 193.
- To create simple rules using predefined templates, click **Add by Template** and refer to the instructions in Adding a Data Migration Rule by Template, on page 188

The following **Shortcuts** are available:

- Click **Policies and Schedules** to display the corresponding page.
- Click **Data Migration Paths** to display its corresponding page.

Caution: Once created, do not change a data migration Rule without verifying whether it is in use by existing policies, as such changes may result in unintentional changes to existing policies.

Adding a Data Migration Rule by Template

To add a data migration rule by template:

1. **Navigate to the Data Migration Rule Templates page.**

   From the **Storage Management** page, select **Data Migration Rules,** then click **Add by Template** to display the **Data Migration Rule Templates** page:

2. **Specify a template.**

Select a **Rule Template**, then click **next.**

The following table describes each of the available rule templates:

Rule Template	Description
By Last Access	Migrates all files that have remained inactive (or have been active) within a certain period of time.
By Creation Time	Migrates all files created before or after a specific point in time.
By File Name	Migrates all files with the same extension; for example, *.mp3*, *.html*, or *.doc*.
By Path	Maps a path to a particular directory and migrates all files under it.
By User	Migrates all files of the specified user(s).
By File Name and Last Access	Migrates files of a certain type (file extension) that have remained inactive for a certain period of time.
By Path and Last Access	Migrates all files under a certain directory that have remained inactive for a certain period of time.

3. **Enter requested template-specific information:**
   - **If you chose "By Last Access":** the **Rules Template: By Last Access** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Include Criteria	To specify the maximum period a file can be inactive before being migrated to a secondary file system, do the following: • From the drop-down menu, select *inactive.* The drop-down menu includes an option for selecting the opposite scenario; that is, to choose *active within* to specify files that have been active within the specified period. • From the drop-down menu, select the period (days, hours, or minutes). • Enter the threshold quantity period. See Rule Syntax, on page 195 for important information about rule criteria.

- **If you chose "By Creation Time":** the **Rules Template: By Creation Time** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Include Criteria	To specify the point in time for the migration rule, do the following: • From the first drop-down menu, select **more than** or **less than**. • Enter the threshold number. • From the second drop-down menu, select one of the following: **month(s)**, **week(s)**, **day(s)**, **hour(s)**, or **minute(s)**. See Rule Syntax, on page 195 for important information about rule criteria.

- **If you chose "By File Name":** the **Rules Template: By File Name** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Case sensitive pattern checks	To specify case-sensitive rule checking, select this checkbox.
Include Criteria	To specify the type of files (based on their file extension) to be migrated to a secondary file system, do the following: • From the drop-down menu, select *include*. The drop-down menu also has an option for selecting the opposite scenario; that is, selecting to exclude files not of the specified type. • In the **all files named** field, enter the file extension. More than one file type can be named in this field separated by commas. For instance, `*.jpg, *.bmp, *.zip.` See Rule Syntax, on page 195 for important information about rule criteria.

- **If you chose "By Path":** the **Rules Template: By Path** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Case sensitive pattern checks	To specify case-sensitive rule checking, select this checkbox.
Include Criteria	To specify the path to the files under a certain directory, do the following: • From the drop-down menu, select *include.* The drop-down menu also has an option for selecting the opposite scenario; that is, to select *exclude* to select all files that are not in the path. • In the **all files in the path** field, enter the directory file path. See Rule Syntax, on page 195 for important information about rule criteria.

- **If you chose "By User":** the **Rules Template: By User Name** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Case sensitive pattern checks	To specify case-sensitive rule checking, select this checkbox.

Item/Field	Description
Include Criteria	To specify the user name(s) for the owner(s) of the files to be migrated to a secondary file system, do the following: • From the drop-down menu, select *include*. The drop-down menu also has an option for selecting the opposite scenario; that is, selecting to exclude files from owners other that the specified owner(s). • In the **all files owned** field, enter the UNIX or Windows user name(s) for the owner(s) of the files you want to migrate. More than one user name can be listed in this field, but names must be separated by commas. For instance, `jjames`, `myco\smithr`, `myco\wsmith` Windows user names are specified in the form `domain\username`, and backslashes in user names should **not** be escaped (double backslashes are not required). See Rule Syntax, on page 195 for important information about rule criteria.

- **If you chose "By File Name and Last Access":** the **Rules Template: By File Name and Last Access** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Case sensitive pattern checks	To specify case-sensitive rule checking, fill this checkbox.
Include Criteria	To migrate inactive files from a specified directory to a secondary file system, do the following: • In the **All files named** field, enter the file name extension of the files to be migrated. For example .txt, .or mp3. • In the **All files not accessed within ___** field, enter the threshold quantity. • Select the period from the drop-down list. You can choose days, hours, or minutes. See Rule Syntax, on page 195 for important information about rule criteria.

- **If you chose "By Path and Last Access":** the **Rules Template: By Path and Last Access** page contains the fields described in the following table:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Case sensitive pattern checks	To specify case-sensitive rule checking, fill this checkbox.
Include Criteria	To migrate inactive files from a specified directory to a secondary file system, do the following: • In the **All files in the Path** field, enter the directory file path. • In the **All files not accessed within ___** field, enter the threshold quantity. • Select the period from the drop-down list. You can choose days, hours, or minutes. See Rule Syntax, on page 195 for important information about rule criteria.

4. **Save the rule template.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Adding a Custom Data Migration Rule

To add a custom data migration rule:

1. **Navigate to the add Data Migration Rule page.**

From the **Data Migration Rules** page, click **add** to display the **add Data Migration Rule** page:

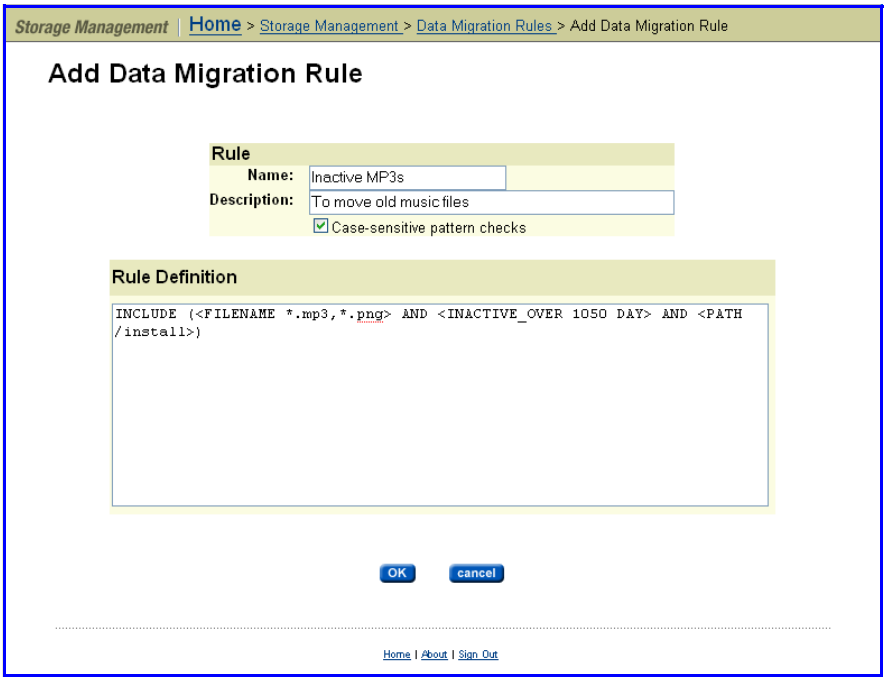

The following table describes the fields in this page:

Item/Field	Description
Name	Enter a name for the new rule. The rule name may include only alphanumeric characters, hyphens, and underscores.
Description	Enter a description of what the rule does.
Case sensitive pattern checks	To specify case-sensitive rule checking, select this checkbox.
Rule Definition	Insert the syntax for the data migration rule. See Rule Syntax, on page 195 for important information about rule criteria.

2. **Save the custom rule.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Creating Specific and Detailed Rules

Before building migration rules, refer to the following several reference sections regarding *Syntax*, *Keywords*, *Connectors*, *Conditionals*, and *Statement Order*. The following example provides a three-step process for assembling simple, yet specific and detailed rules:

- Start with a simple INCLUDE statement that is specific about what should be migrated, such as:

  ```
  INCLUDE (<PATH /Temp> AND <FILENAME *.mp3>)
  ```

- Refine the INCLUDE statement by adding exceptions to the rule with restrictive EXCLUDE statements. But add these EXCLUDE statements above the INCLUDE, such as:

  ```
  EXCLUDE (<FILE_SIZE_UNDER 2MB>)

  EXCLUDE (<ACTIVE_WITHIN 14>)
  ```

- The rule should finally appear this way:

  ```
  EXCLUDE (<FILE_SIZE_UNDER 2MB>)

  EXCLUDE (<ACTIVE_WITHIN 14>)

  INCLUDE (<PATH /Temp> AND <FILENAME *.mp3>)
  ```

Rule Syntax

Data migration rules compares a series of INCLUDE and EXCLUDE statements, each qualified by expressions stating the criteria for data migration. The following guidelines govern rule building:

- **At least one INCLUDE or EXCLUDE.** Each rule must contain at least one INCLUDE or EXCLUDE statement. Rules consisting only of EXCLUDE statements imply that everything on primary storage should be migrated except what has been specifically excluded.

  

  Note: If a rule contains only INCLUDE statements, all items not specified by the INCLUDE statements are excluded.

- **Wildcards.**The asterisk "*" can be used as a wildcard character to qualify PATH and FILENAME values.
  - When used in a PATH value, "*" is only treated as a wildcard if it appears at the end of a value; for example, `<PATH /tmp*>`.
  - In a FILENAME value, a single "*" can appear either at the beginning or the end of the value.
  - Multiple instances of the wildcard character are not supported and additional instances in a value definition will be treated as literal characters.

- **Bracketed keyword/value pairs.** Expressions identifying migration criteria should be enclosed in brackets. All criteria contain a keyword, defining the condition for data migration, followed by a single value of a list of values; for example, `<FILENAME *.doc>`.

- **Evaluation of statement sequence.**When using multiple INCLUDE or EXCLUDE statements, they are evaluated using top-down ordering. For more information on ordering, refer to Statement Order, on page 200.

- **Grouping criteria within statements.**Parentheses are used to group the criteria in INCLUDE and EXCLUDE statements; for example, `INCLUDE (<PATH /Temp/*>).`
- **Number of INCLUDE or EXCLUDE statements per line.** When using multiple INCLUDE or EXCLUDE statements in a rule, each INCLUDE or EXCLUDE statement must be placed on its own line (multiple INCLUDE and/or EXCLUDE statements may not be put on the same line).
- **Separating multiple listed values.**When defining multiple values in a FILENAME list, use a comma to separate values; for example, `INCLUDE (<FILENAME *.mp3,*.wav,*.wmv>).`
- **Characters requiring escaping.**The following characters need to be escaped with a backslash (\) when used as a part of PATH or FILENAME values: \ *(backslash),* > *(greater than),* and , *(comma)*; for example, `INCLUDE (<FILENAME *a\,b> OR <PATH /tmp/\>ab>).`

Note: Backslashes used when specifying a domain and user name should **not** be escaped (double backslashes are not required when specifying `domain_name\user_name`).

- **Forward slash (/) reserved as a path separator.**The forward slash (/) is used as a path separator; as such, it must not be used in a FILENAME list.
- **Evaluation of absent PATH.**If a PATH element is not specified in a statement, the statement will apply to the entire file system or virtual volume defined in the data migration path.
- **Quotation mark usage.**Quotation marks (“) are not allowed around a FILENAME or PATH list.

Keywords

The following table describes the keywords and their related values that can be used to build rule statements. Each keyword can be defined in the rule with an INCLUDE or EXCLUDE statement to indicate how the keyword values are to be applied.

Keyword	Value(s)
FILENAME	Names and types of files contained in the rule. Separate multiple names by commas. FILENAME values may *start or end* with a “*” wildcard character to indicate all files starting/finishing with specific characters. Usage: FILENAME will often be used with an INCLUDE statement to ensure that non-essential files are migrated to secondary storage. It can also be used with an EXCLUDE statement to prevent specific important data sets from being migrated. For example: `<FILENAME *.mp3,*.txt,filename*>`

Keyword	Value(s)
PATH	Specifies literal paths to which a rule applies. Values must be full paths, starting with a forward slash (/). Separate multiple paths by commas. PATH values may *end* with a "*" wildcard character to indicate all subdirectories under the specified path. Usage: When used in an INCLUDE statement, PATH specifies directories to migrate. This is useful when migrating less-critical directories such as `temp` or `home` directories. When used in an EXCLUDE statement, directories can be excluded from migration, leaving all the files within on primary storage. For example: `<PATH /temp/*,/home*,/other/dir*>`
USERNAME	Specifies user names to which a rule applies. Values must be valid Windows or UNIX user names. Separate multiple names by commas. Usage: When used in an INCLUDE statement, USERNAME specifies the user name of file owners whose files are to be migrated. This is useful when migrating files owned by a particular user. When used in an EXCLUDE statement, users can be excluded from migration, leaving all the files owned by the specified user(s) on primary storage. Windows user names are specified in the form `domain\username`.  **Note:** Backslashes in user names should **not** be escaped (double backslashes are not required). For example: `jjames`, `myco\smithr`, `myco\wsmith`
FILE_SIZE_OVER	Identifies a subset of files in a rule with sizes crossing an upper threshold. The threshold value is appended to the keyword and defined by the threshold size in *B, KB, MB,* or *GB.* Usage: This will likely be used with INCLUDE statements to ensure files of very large sizes are migrated to secondary storage. For example: `<FILE_SIZE_OVER 4GB>`
FILE_SIZE_UNDER	Identifies a subset of files in a rule with sizes crossing a lower threshold. The threshold value is appended to the keyword and is defined by the threshold size in *B, KB, MB,* or *GB.* Usage: This will usually be used in an EXCLUDE statement to ensure that very small files are not migrated en masse. Migrating small files that take up little space provides minimal value in extending the efficiency of primary storage. For example: `<FILE_SIZE_UNDER 10KB>`

Keyword	Value(s)
OLDER_THAN	Identifies files that were created more than a specified number of days in the past (files older than *x* days). The value appended to the keyword defines the minimum age (number of days) of a file before the rule is applied to that file. Usage: Used primarily in INCLUDE statements to ensure that older files are migrated. For example: `<OLDER_THAN 28>`
NEWER_THAN	Identifies files that were created less than a specified number of days in the past (files newer than *x* days). The value appended to the keyword defines the maximum age (number of days) of a file before the rule is applied to that file. Usage: Used primarily in EXCLUDE statements to ensure that newer files are not migrated. For example: `<NEWER_THAN 14>`
INACTIVE_OVER	Identifies files that have not been accessed within a specific number of days. A file's last access time is updated whenever the file is read or modified. The value is appended to the keyword and defines the number of days of inactivity. Usage: Used primarily in INCLUDE statements to ensure that older, less frequently used files are migrated. For example: `<INACTIVE_OVER 21>`
ACTIVE_WITHIN	Identifies files accessed within a specific number of previous days. A file's last access time is updated whenever the file is read or modified. The value is appended to the keyword and defines the number of days within which the activity has occurred. Usage: Used primarily in EXCLUDE statements to prevent actively used files from being migrated. For example: `<ACTIVE_WITHIN 30>`
UNCHANGED_OVER	Identifies files not modified within a specific number of previous days. A file's modification time is updated whenever the file's contents have been changed. The value is appended to the keyword and defines the number of days of inactivity. Usage: Used primarily in INCLUDE statements to ensure that older, less frequently used files are migrated. For example: `<UNCHANGED_OVER 14>`

Keyword	Value(s)
CHANGED_SINCE	Identifies files modified within a specific number of previous days. A file's last access time is updated whenever the file's contents have been changed. The value is appended to the keyword and defines the number of days of inactivity. Usage: Used primarily in EXCLUDE statements to prevent actively used files from being migrated. For example: `<CHANGED_SINCE 7>`

Connectors

Statements can combine multiple criteria, as follows:

- **AND** indicates that both statements must be satisfied. For example, in the statement:

  ```
  INCLUDE (<FILENAME *.mp3> AND <FILE_SIZE_OVER 5GB>)
  ```

  both conditions must be true in order for the statement to be true.

- **OR** indicates that only one statement needs to be satisfied. For example, for the same statement, replacing AND with OR:

  ```
  INCLUDE (<FILENAME *.mp3> OR <FILE_SIZE_OVER 5GB>)
  ```

  only one condition needs to be true for the statement to be true.

While **AND** requires *both* conditions to be true, **OR** only requires that *either* condition to be true.

Conditionals

The following table shows a set of rules with explanations. The syntax can easily be broken down into cause and effect statements, identified by IF and THEN connectors.

Rule	Description
`INCLUDE (<FILENAME *.doc>)`	IF the *file is* a `.doc` file, THEN include it for migration.
`EXCLUDE (<PATH /mydir/*>)`	IF the *path is* the `/mydir` directory THEN exclude it from migration.
`INCLUDE (<FILENAME *.prj> AND <FILE_SIZE_OVER 4GB>)`	IF the *file is* a `.prj` file **AND** the `.prj` file is *over 4 GB* in size, THEN include it for migration.
`INCLUDE (<PATH /unimportant>)`	IF the path is the `/unimportant` directory THEN include it for migration.

Rule	Description
`EXCLUDE (<FILE_SIZE_OVER 100GB>) INCLUDE (<FILE_SIZE_OVER 12GB>)`	IF files are *larger than 12 GB* but *smaller than 100GB* in size, THEN include them for migration.

Statement Order

Statement order is critical. Statements are evaluated top-down, starting with the first statement defined; therefore, as the following example illustrates, best practice usually specifies EXCLUDE statements at the top of the rule:

Rule Scenario A :

```
INCLUDE (<PATH /Temp> AND <FILENAME *.mp3>)

EXCLUDE (<ACTIVE_WITHIN 14>)

EXCLUDE (<FILE_SIZE_UNDER 2MB>)
```

The above rule is interpreted as:

- IF path name includes `/Temp` **AND** filename is `*.mp3` THEN MIGRATE.
- IF file is *active less than 14 days* **AND** *less than 2MB* in size THEN EXCLUDE.

In scenario A, all the `.mp3` files under `/Temp` will be migrated based on the first INCLUDE statement. Statements 2 and 3 are disregarded since they are evaluated after the more inclusive INCLUDE statement that has already added what rules 2 and 3 are trying to exclude.

Rule Scenario B :

If the same rules were ordered differently:

```
EXCLUDE (<FILE_SIZE_UNDER 2MB>)

EXCLUDE (<ACTIVE_WITHIN 14>)

INCLUDE (<PATH /Temp> AND <FILENAME *.mp3>)
```

The above rule is interpreted as:

- IF file is *less than 2 MB* in size AND *active less than 14 days* THEN EXCLUDE.
- IF path name includes `/Temp` AND filename is `*.mp3` THEN MIGRATE.

In this Scenario, only `.mp3` files *greater than 2 MB* in size that have been *inactive for greater than 14 days* will be migrated.

Data Migration Policies

Having created both data migration paths and data migration rules, data migration policies can now be created. Policies assign a rule or set of rules to a specific data migration path. They also define the conditions that initiate data migrations.

Viewing Data Migration Policies

To view data migration policies, navigate from the **Storage Management** page to the **Data Migration** page:

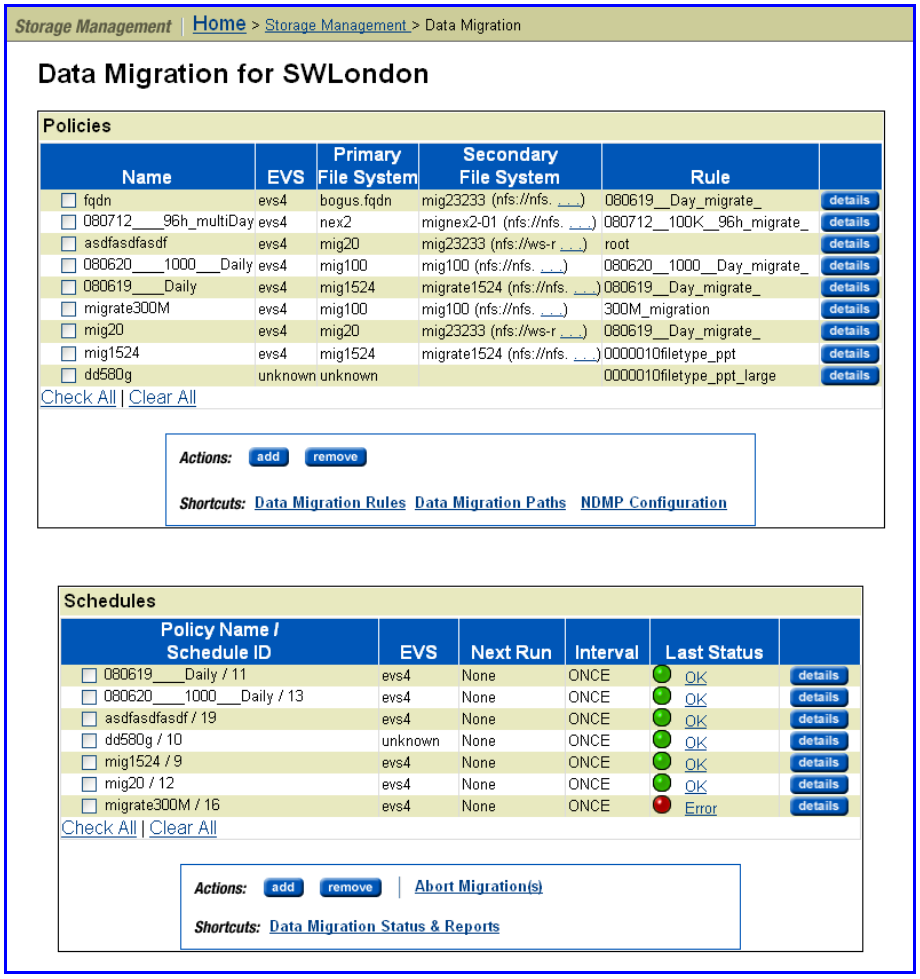

The fields in the **Policies** list on this page are described in the table below (for information about the **Schedules** list, see Migration Schedules, on page 205):

Item/Field	Description
Name	Name of a data migration policy.
Server/EVS	Primary EVS from which the migration originates.
Primary File System	Primary file system or virtual volume that will be migrated.

Item/Field	Description
Secondary File System	Secondary file system, to which all data will be migrated.  **Note:** If the path to the secondary file system is an external path, the name or IP address of the server hosting the secondary file system is also displayed in parentheses. The displayed server name/IP address is a link, and you can click on the link to display the full path.
Rule	Rules which may be triggered in this migration policy.

The following **Actions** are available:

- Click **detail** to display detailed information about the specified migration policy.
- To create a new migration policy, click **add** and refer to Adding a data migration policy, on page 202.
- To delete a specified migration policy, click **remove**.

You can also select shortcuts to the following pages:

- **Data Migration Rules**
- **Data Migration Paths**
- **NDMP Configuration**

Adding a data migration policy

To add data migration policies:

1. **Migrate to the add Data Migration Policy page.**

   From the **Storage Management** page, select **Data Migration,** then click **add** to display the **Add Data Migration Policy** page:

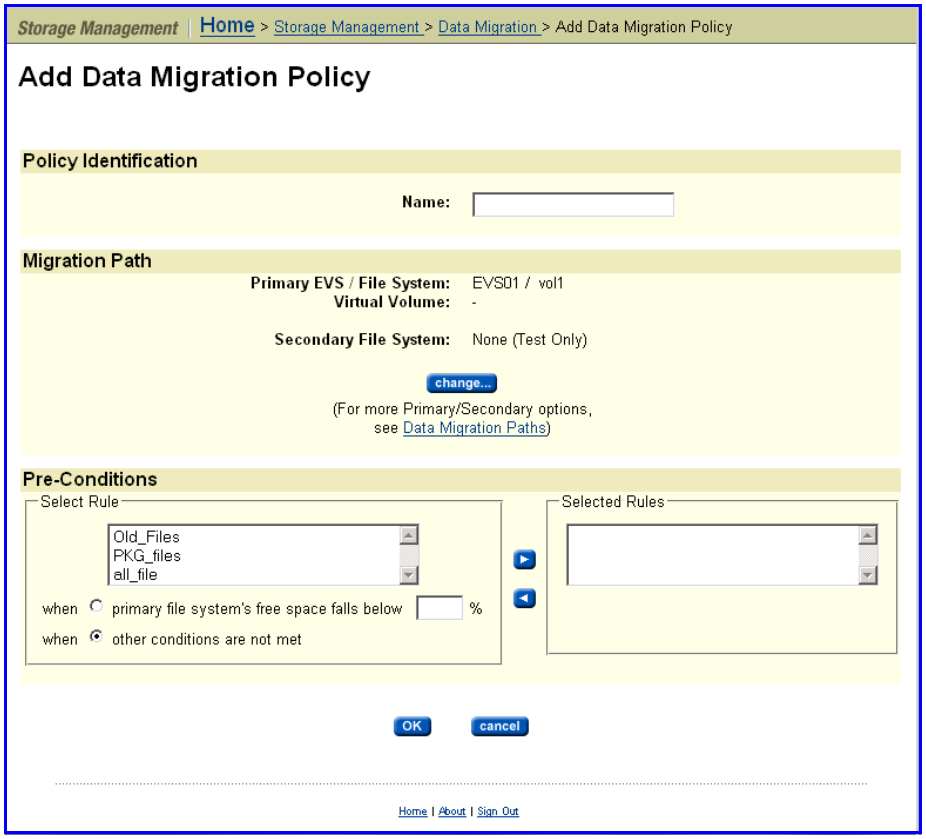

The following table describes the fields in this page:

Item	Description
Name	Name for the new data migration policy.
Primary EVS/File System	EVS and file system names for primary storage (migration source).
Virtual Volume	If a virtual volume has been selected as primary storage, the virtual volume name will be displayed.
Secondary File System	The file system on secondary storage that will host the migrated data (migration target). To change the selected migration path, click **change.** A list of paths will appear in the **Select a Path** page. Select a path for the migration, then click **OK**.

Item	Description
Pre-Conditions	Rules with specific threshold limits are displayed here. This list of rules define the set of conditions by which file migrations are triggered: • Add a Pre-Condition to the Selected Rules list by selecting it and clicking the right arrow (>). • Remove a rule from the Selected Rules list by selecting it and clicking the left arrow (<). Select *when* the Selected Rules are applied. You can choose either list to be applied if either of the following conditions are met: • When primary's free space falls below ***X***% (set the percentage level for the condition) • When other conditions are not met once defined, add the rule and threshold to the list of Pre-Conditions by clicking **add.**

2. **Complete the requested information.**

Note: If you are adding a policy to perform a test migration, a valid migration path is not required.

3. **Save your settings.**

   Verify your settings, then click OK to save or **cancel** to decline.

Using Pre-Conditions

When a migration policy is scheduled to run, it evaluates the percentage of available free space in the Policy's primary storage. Based on this analysis, one rule may be triggered to define the data set subject to migration. Migrations of data from primary storage then occurs based on the statements in the rule that was triggered. Only a single rule will be engaged during any particular migration operation.

When defining pre-conditions, SGI Global Services recommends aggressive tiering; specifically, it may be desirable to migrate `.mp3` files and the contents of the directory `/tmp` regardless of the available free space. Then, if free space on primary storage is reduced to less than 50%, also to migrate all files not accessed within the last sixty days. Finally, if available free space is reduced to less than 15%, also to migrate the contents of users' `home` directories.

The following will illustrate this scenario:

Rule	Statement
Rule 1:	`INCLUDE (<FILENAME *.mp3>) OR <PATH /tmp/*)`
Rule 2:	`INCLUDE (<FILENAME *.mp3>) OR <PATH /tmp/*)`

Rule	Statement
	`INCLUDE (<INACTIVE_OVER 60>)`
Rule 3:	`INCLUDE (<FILENAME *.mp3>) OR <PATH /tmp/*)`
	`INCLUDE (<INACTIVE_OVER 60>)`
	`INCLUDE (<PATH /home/*>)`

Related pre-conditions:

- Rule 3 if free space is less than 15%.
- Rule 2 if free space is less than 50%.
- Rule 1 if no other condition applies.

When the migration policy is scheduled to run, different rules may be triggered based on the available free space on primary storage. When a migration policy is engaged, only a single rule will be triggered to run.

For example:

- If free space is at 80%, then Rule 1 will be used.
- If free space is at 40%, then Rule 2 will be used.
- If free space is at 10%, then Rule 3 will be used.

When percentage thresholds are specified, they are evaluated based on whole number percentages. This means that if two rules are specified, one that will take effect at 8% of free space and one at 9% of free space, if the file system has 8.5% free space available, then the rule with the 8% pre-condition will apply.

Note: If the primary storage defined in the migration path is a virtual volume, free space will be based on the limit defined by the virtual volume quota. If a virtual volume quota has not been defined, then free space available will be based on the free space of the file system hosting the virtual volume.

Migration Schedules

Once a data migration policy has been defined, it must be scheduled. The decision how often to run a Policy may be affected by the Rules selected in this policy. For example:

- A policy with a single Rule to migrate all `.mp3` files may be scheduled to run once every month.
- Another policy, used to archive a working `/project` directory once the project is complete, may be scheduled as a Once Only Schedule.
- Other policies which migrate based on various Pre-conditions, which are triggered on available free space, may be scheduled to run every week.

When planning migration schedules, SGI recommends scheduling during off-peak times, such as evenings and weekends.

Once a data migration has begun, additional data migrations for the same policy cannot be started until the current one has completed; however, it is possible to start multiple concurrent data migrations, each for its own policy.

Viewing Scheduled Migrations

To view scheduled migrations, navigate from the **Storage Management** page to the **Data Migration** page:

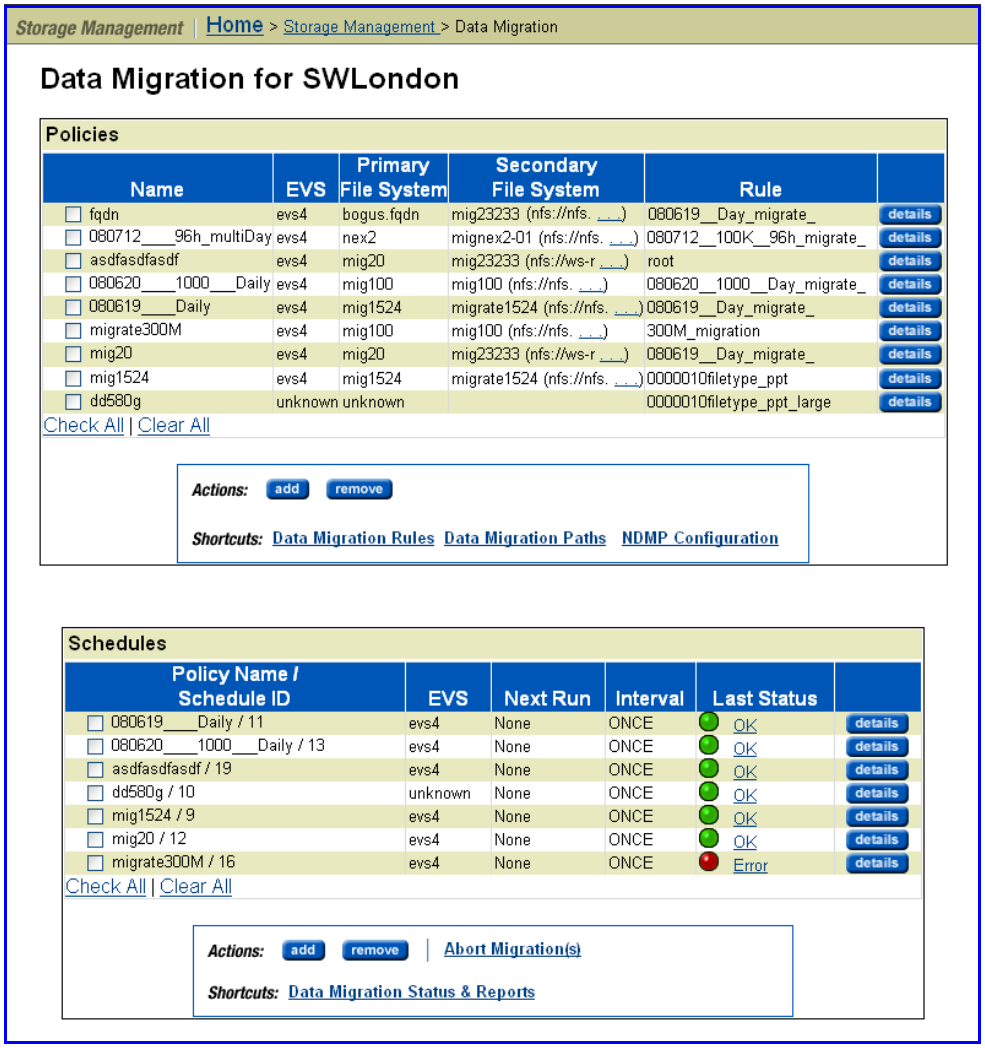

The fields in the **Schedules** list on this page are described in the table below (for information about the **Policies** list, see Viewing Data Migration Policies, on page 201):

Item/Field	Description
Policy Name/Schedule Id	Name of the data migration policy.
Server/EVS	Primary server and EVS from which the migration will originate.
Next Run	Month, date, year and time for the next scheduled data migration run for this policy.
Interval	Frequency at which the data migration has been scheduled to run.
Last Status	If a migration operation has been scheduled and has run, this column displays the status of the last migration operation that was run according to the schedule. Click on the status link to see the migration report for the operation.

The following **Actions** are available:

- Click **details** to display detailed information about the specified migration schedule.
- Click **add** to create a new migration schedule and refer to Adding a Data Migration Schedule, on page 207.
- Click **remove** to delete a specified migration schedule.
- Click **Abort Migrations** to abort a selected, in-process migration. Only in-progress migrations can be aborted.

You can also select shortcuts to the following pages:

- **Data Migration Status & Reports**

Adding a Data Migration Schedule

Note: You must create a migration policy before you can schedule it.

To add a data migration schedule:

1. **Navigate to the add Data Migration Schedule page.**

   From the **Storage Management** page, select **Data Migration,** then click **add** to display the **add Data Migration Schedule** page:

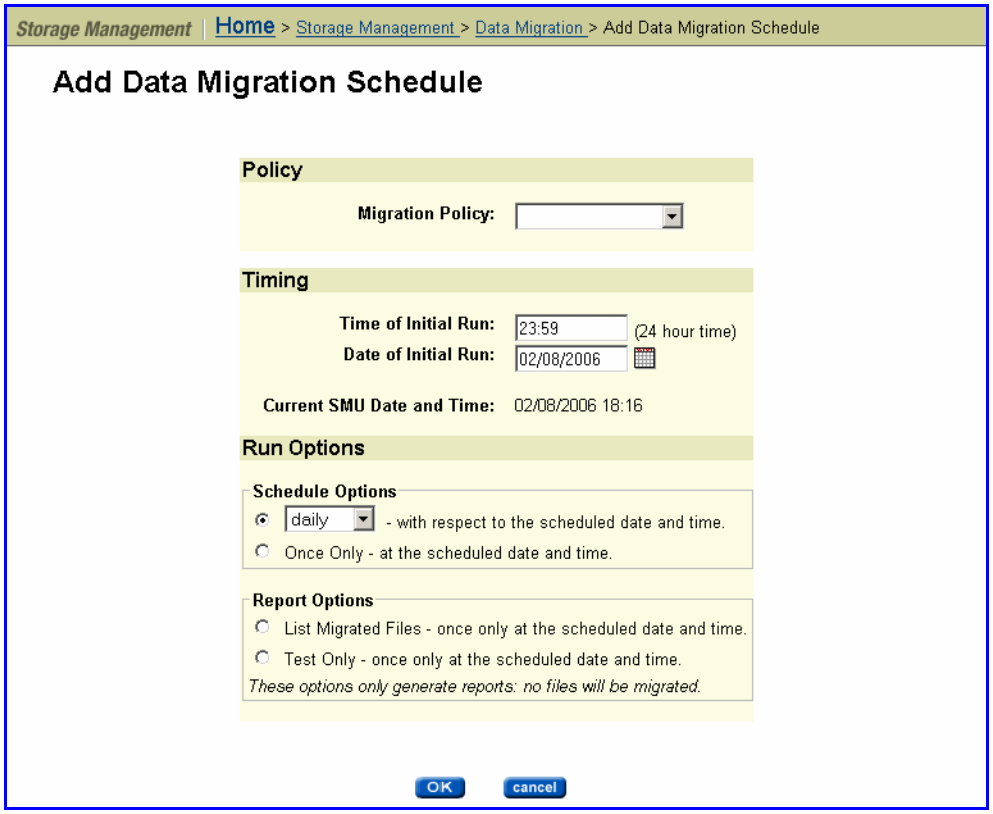

2. **Enter the requested information.**

   The data migration policy needs to be set up before it can be scheduled.

Item	Description
Migration Policy	Select a migration policy from the drop-down menu.
Time of Initial Run	Scheduled run time on a 24 clock (i.e. 11:59 PM will be entered as 23:59). The current SMU date and time are provided for reference.
Date of Initial Run	From the calendar, select a start date for the policy's initial run. The selected date appears on the field.
Schedule	When selecting the first option, pick a pre-set rule of *daily, weekly,* or *monthly* from the drop-down menu. This will be applied at the same time of day as the Initial Run. Selecting *Once Only* indicates that the policy is scheduled to run only once, the initial run.
Report Options	Select **List Migrated Files** to generate a report of all migrated files in the selected data migration path. Select **Test Run** to initiate a one-time test. The files are not migrated under this option, but reports can be generated to provide valuable insights about the validity of using a certain policy and its rules.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Modifying a Data Migration Schedule

Once defined, schedules can be easily modified to meet the changing requirements of the data migration policies. When modifying a schedule, the scheduled date and time, as well as the interval in which the schedule will run can be changed.

To modify a data migration schedule:

1. **Navigate to the Data Migration page.**

   From the **Data Storage** page, select **Data Migration**, then click **details** for a particular schedule to display its **Modify Data Migration Schedule** page:

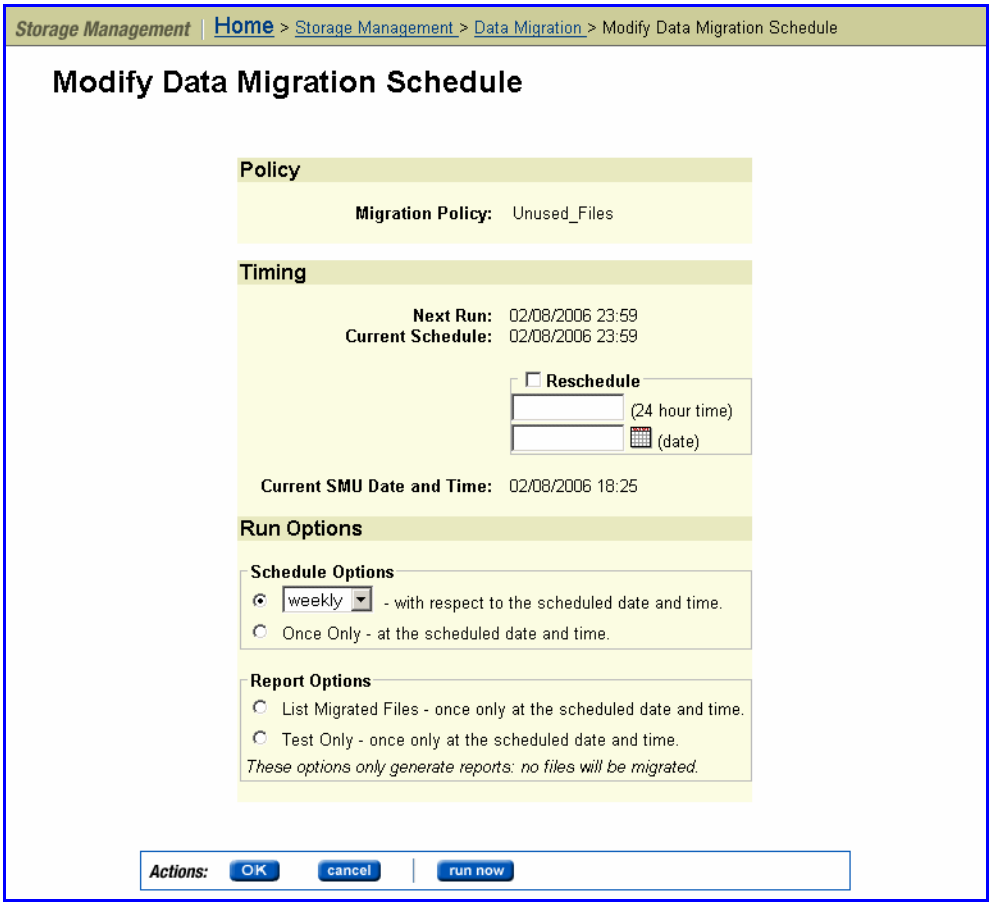

2. **Modify the schedule.**

   The following modifications are available:

   - *To define a new starting date and time for the selected schedule,* fill the **Reschedule** box and enter the new values in the appropriate fields.
   - *To change the schedule's interval,* configure the schedule to repeat either daily, weekly, or monthly, or configure the schedule to run Once Only.

- *To change the schedule to run a report,* click **List Migrated Files** to list all migrated files in the selected data migration path, or **Test Only** to generate a report of what files would be migrated if the specified migration policy were run.

3. **Save, run immediately, or decline.**

   Verify your settings, then click **OK** to save, **run now** to run the schedule immediately, or **cancel** to decline.

Migration Reports

Once a data migration policy has completed a cycle, it generates a data migration report that includes details about files migrated, including available free space before and after the migration. Reports of the last five scheduled migrations are routinely saved; the rest are purged. If a schedule is deleted, so are its reports.

Migration reports can be saved and printed. They are useful in studying the system access patterns, file storage tendencies, the efficiency of rules, paths, policies and schedules. By gauging file and space usage statistics of Primary and secondary storage, Data Migrator reports can be used to refine a rule or pre-condition. The more precise and aggressive the rule, the better Data Migrator serves the storage system.

Viewing Completed Migrations

To view completed migrations, navigate from Storage Management to the **Data Migrations Status and Reports** page:

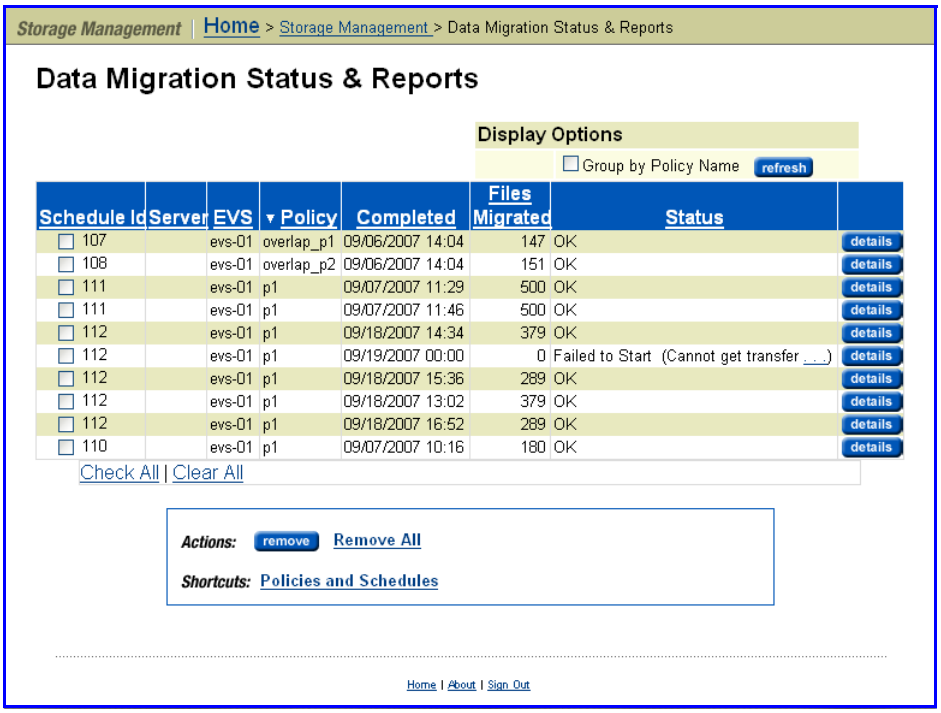

The following table describes the fields in this page:

Item	Description
Schedule ID	ID number for the completed migration.
Server	Primary file system's server.
EVS	Primary file system's EVS.
Policy	Policy's name.
Completed	Month, date, year and time when the migration was completed.
Files Migrated	Number of files that were migrated.
Status	Migration completion status.

The following **Actions** are available:

- Click **detail** to display detailed information about the specified **Completed Migration**.
- Click **remove** to delete a specified migration report.

Viewing Data Migration Reports

To view data migration reports, from the **Storage Management** page, select **Completed Data Migrations,** then **Completed Migration**, and click **details** to display the **Completed Data Migration Details** page:

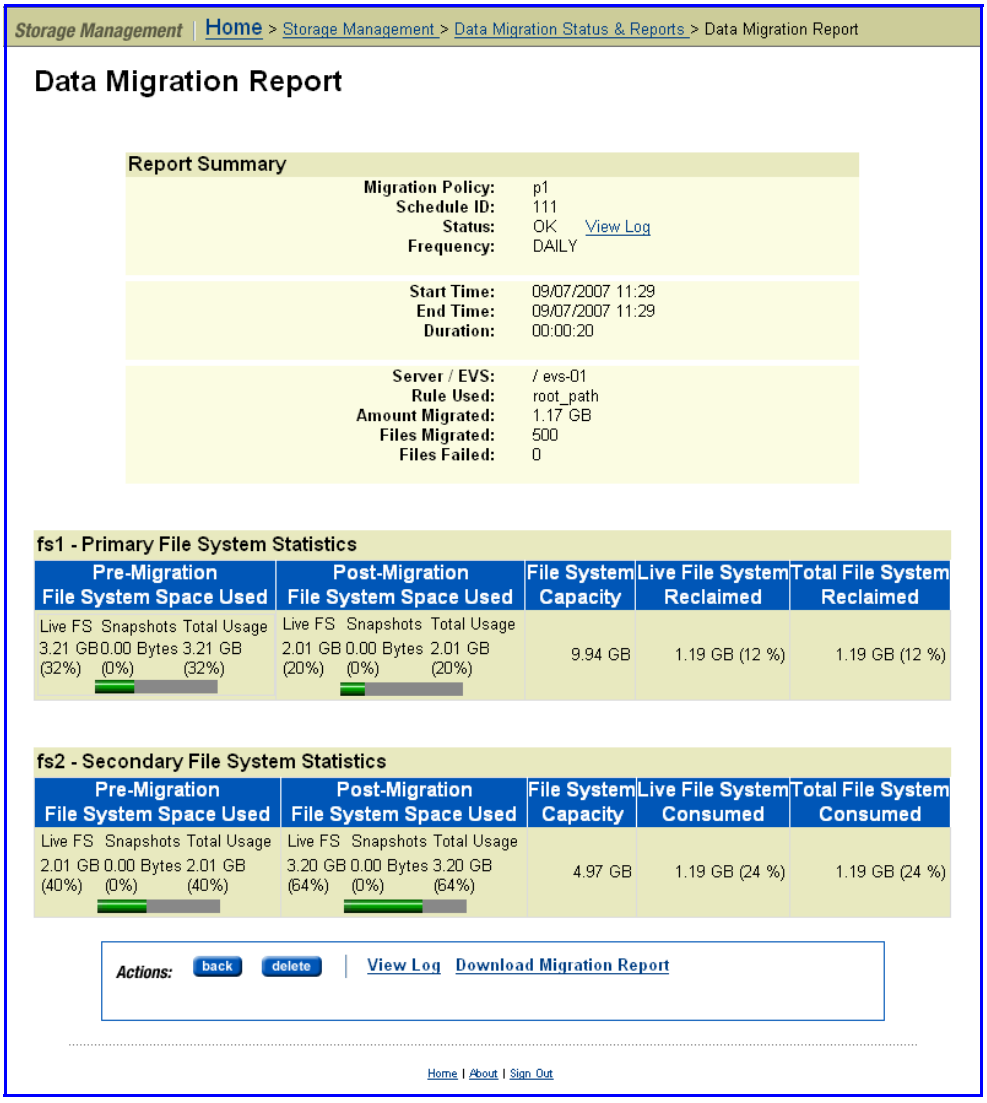

The following table describes the contents of this page:

Item	Description
Report Summary	
Migration Policy	Completed migration policy's name.
Schedule ID	Migration schedule ID.
Status	Migration completion status.
Frequency	How often the Policy is scheduled to run.
Start Time	Date and time when the migration began.
End Time	Date and time when the migration ended.

Item	Description
Duration	Duration of migration.
Server/EVS	EVS on which the Primary and secondary storage reside.
Rule	Rule used by the policy.
Amount Migrated	Migrated data quantity in GB.
Files Migrated	Quantity of files that were migrated. If files have been migrated, click this to view a list of the files that were migrated. The list provides details on their path, size, and their start and end times.
Files Excluded	Number of files that should have been migrated but could not. For example, files in use at the time of the migration may not be migrated.
Primary File System Statistics	
Pre-Migration File System Space Used	File system size, snapshot size, and the total used space before the migration.
Post-Migration File System Space Used	File system size, snapshot size, and the total used space after the migration.
File System Capacity	File system's total capacity.
Live File System Reclaimed	Reclaimed space in the live file system, defined as the usable space on the file system; that is, the part of the file system not reserved or in use by snapshots.
Total File System Reclaimed	Reclaimed space in the total file system, defined as the entire capacity of the file system and includes usable space and space that is reserved or in use by snapshots.
Primary Virtual Volume Statistics	
Pre-Migration Virtual Volume Space Used	Details the virtual volume's size and the total space used before the migration.
Post-Migration Virtual Volume Space Used	Details the virtual volume's size and the total space used after the migration.
Virtual Volume Reclaimed	Displays the virtual volume space gained due to the migration.
Secondary File System Statistics	
Pre-Migration File System Space Used	File system size, snapshot size, and the total used space before the migration.
Post-Migration File System Space Used	File system size, snapshot size, and the total used space after the migration.
File System Capacity	File system's total capacity.

Item	Description
Live File System Consumed	Space taken up due to the migration.
Total File System Consumed	Total space used in the file system by migration.
Secondary Virtual Volume Statistics	
Pre-Migration Virtual Volume Space Used	Details the virtual volume size and the total space used before the migration.
Post-Migration Virtual Volume Space Used	Details the virtual volume size and the total space used after the migration.
Virtual Volume Consumed	Displays the virtual volume space taken up by the migration.

The following **Actions** are available:

- Click **View Log** to view a log file containing *time*, *duration* and *status* details of the migration. A **View Log** link is available at both the top and bottom of the page.
- Click **Download Migration Report** to view a report about the completed data migrations with details on the primary and secondary file systems and virtual volumes, including *status, space utilization before and after the migration,* the *duration, start,* and *end time* for the migrations.

  Included in the download are two other important reports: one that lists all the files that were migrated (list.gz) and the other that lists all the files that were not migrated (failed.gz).

Reclaimed Space

Reclaimed space is the difference in available space between the start and completion of the migration. It is not a report of the amount of data migrated from the source file system to the target. For this information, refer to Amount Migrated.

It is likely that the file system will be in use by network clients while the migration is in progress. As a result, the reclaimed space can be substantially different than the amount migrated. The value can even be negative if files were added to the source.

Once a data migration has completed, copies of the files may be preserved on the source file system in snapshots. For the space to be fully reclaimed, all snapshots on the source file system that reference the migrated files must be deleted.

Reversing Migration

Although the server does not support automatic reverse migration of files, it is possible to restore a migrated file in two different ways:

- **Reverse Migration Through the server CLI.** Individual files or whole directory trees can be reverse-migrated through the CLI. The files which

are included in the reverse migration can be identified by pattern or by last access time. For detailed information on this process, run `man reverse-migrate` at the CLI.

- **Reverse Migration From a Network Client.** A file can be restored from a network client by performing the following sequence of operations:
  - From a Windows or Unix client, make a copy of the file (using a temporary file name) on the primary storage. This copy of the file will reside fully on primary storage.
  - Delete the original file. This will delete the link on primary storage, and the migrated data from secondary storage.
  - Rename the copied file to its original name.

iSCSI Logical Units

Mounted iSCSI Logical Units cannot be migrated, regardless what has been defined in the data migration policy. Due to the types of applications typically hosted on iSCSI storage, SGI Global Services does not recommend migrating iSCSI Logical Units to secondary storage. However, if this is desired, it can be accomplished by performing the following:

- Disconnect any iSCSI Initiators with connections to Logical Unit.
- Unmount the iSCSI Logical Unit. This can be done through the **iSCSI Logical Unit Properties** page.
- Run the data migration policy to migrate the Logical Unit.
- Re-mount the iSCSI Logical Unit.
- Reconnect the Initiator to the iSCSI Target.

7 File Services

The IS-NAS Server and the Titan Server are file-serving products, and their principal use is to satisfy incoming file access requests issued from network clients.

File Service or Component	Conceptual Overview	Associated Tasks
Supported File System Protocols	File System Protocols, on page 218	Enabling and Disabling File Services, on page 221
Unicode Support	Unicode Support, on page 219	Changing the Character Set, on page 220
File System Security	Managing File System Security, on page 222	Viewing Security Configurations, on page 225 Changing Security Mode, on page 226
NFS	Enabling NFS Protocol Support, on page 244 Configuring NFS Exports, on page 245	Adding an NFS Export, on page 246 Viewing the Properties of an NFS Export, on page 250 Backing Up and Restoring NFS Exports, on page 252
CIFS	Configuring CIFS Security, on page 255 Using Windows Server Management, on page 280	Configuring Local Groups, on page 264 Configuring CIFS Shares, on page 268 Using the Computer Management Tool, on page 281
FTP	Configuring FTP Preferences, on page 283 Configuring FTP Users, on page 284 Setting Up FTP Audit Logging, on page 288	To Configure FTP Preferences, on page 284 Setting up an FTP User, on page 284 Configuring FTP Audit Logging, on page 289
iSCSI	Configuring iSCSI, on page 292 Configuring iSNS, on page 293 Configuring iSCSI Logical Units, on page 294 Configuring iSCSI Security (Mutual Authentication), on page 308 Accessing iSCSI Storage, on page 312	Creating and Deleting iSNS Servers, on page 293 Managing iSCSI Logical Units, on page 296 Configuring the Storage Server for Mutual Authentication, on page 308 Using iSNS to Find iSCSI Targets, on page 313

File System Protocols

The server supports the CIFS, NFS, and FTP protocols for client file access, as well as iSCSI for block-level access to storage. All supported protocols can be enabled or disabled.

The server allows NFS, CIFS, and FTP users to access the same file space; however, although iSCSI Logical Units reside on file systems, it is not possible to access folders and files located on an iSCSI target through the server's file services (for example, CIFS or NFS).

These protocols, with the exception of FTP, require a license key for activation.

Note: For more information about how the server resolves differences between protocols, see Mixed Mode Operation, on page 227.

Supported CIFS Versions

The storage server supports CIFS (SMB) versions 1 and 2. SMB version 2 (SMB2) support is provided to maintain compatibility with computers running the Windows Vista and Windows Server 2008 operating systems. SMB2 support is not enabled by default. When enabled, SMB2 support is on a per-EVS basis, it is not server or cluster-wide, meaning that SMB2 support must be specifically enabled on each EVS to which clients may need to connect using the CIFS version 2 protocol.

Note: SMB2 support is limited to the file server portion of the protocol. When acting as a client (for example, when making domain controller and virus scanner connections), the NAS server uses only the SMB version 1 client implementation.

To enable SMB2 support, use the CLI command `cifs-smb2-enable`. Once enabled, the EVS supports CIFS clients connecting through version 1 or version 2 of the protocol.

Note: A valid CIFS license is required in order to enable SMB2 support. For more information about license keys, see Managing License Keys, on page 535.

When SMB2 support is enabled, the type of CIFS connection established is determined based on the connection type that the client advertises it supports. Only clients advertising support for CIFS version 2 establish CIFS version 2 connections. Clients requesting a CIFS connection without specifying a version 2 connection establish CIFS version 1 connections.

Note: After SMB2 support is enabled on the storage server, some clients that support CIFS version 2 may continue to connect using CIFS version 1 connections until they have been restarted. This occurs because some clients cache connection type, and do not negotiate the connection type every time they connect. Clients that operate in this manner will continue to connect using CIFS version 1 until they have been restarted.

To disable SMB2 support, use the CLI command `cifs-smb2-disable`. When SMB2 support is disabled, only CIFS version 1 connections can be established. When a client advertises that it supports CIFS version 2 connection, the NAS server will establish a CIFS version 1 connection.

Statistics for both CIFS version 1 and 2 client connections are kept, and can be viewed on the **CIFS Statistics** page (see Viewing CIFS Statistics, on page 478 for more information).

Supported NFS Versions

The storage server supports NFS versions 2, 3, and 4. Both TCP and UDP are supported in versions 2 and 3, but version 4 support is for TCP only. By default, the maximum version supported is version 3, meaning that the server supports versions 2 and 3 by default. To change the maximum supported version to NFSv4, use the CLI command `nfs-max-supported-version`. By setting the maximum supported version to 4, you allow the storage server to support NFS versions 2, 3, and 4.

Unicode Support

The storage server (or cluster) stores metadata about the files, directories, migration paths, CIFS shares, NFS exports, user names, group names, log entries, mount points and so on for the virtual servers, file systems, and name spaces served by the server/cluster.

When interacting with another network device, the metadata transmitted to or received by the storage server/cluster must be encoded in a character set supported by the other network device. Typically, clients/devices using the CIFS (Windows) protocol encode data in the UCS-2 character set, and clients/devices that use the NFS protocol encode data in the UTF-8 character set.

Note: The data on storage subsystems attached to a storage server/cluster is not affected in any way by the character sets currently used by the server/cluster.

CIFS Unicode Support

When using the CIFS protocol to communicate with clients/devices, the storage server/cluster supports the UCS-2 character set.

FTP Unicode Support

When using the FTP protocol to communicate with clients/devices, the storage server/cluster supports the UTF-8 character set for user names, passwords, and file/directory names.

NFSv2/3 and NIS Unicode Support

When using NFSv2/3 to communicate with clients/devices and/or NIS servers, the default character set is ISO 8859-1 (Latin-1). This character set may not be sufficient to communicate with both NFS clients and NIS servers (for example, when characters outside the Latin-1 range are required, such as Chinese, Japanese or Korean), so an administrator can specify the character set(s) to be used when communicating with NFS clients and/or NIS servers using the `protocol-character-set` command. The `protocol-character-set` command specifies which character set is used when sending/receiving:

- File and directory names to/from NFS clients.
- User and group names to/from NIS servers.

Communication with all NFS clients and/or all NIS servers uses the same character set; you cannot specify that the storage server/cluster communicates with some NFS clients or some NIS servers using one character set, and other NFS clients or NIS servers using a different character set.

Changing the Character Set

By default, the storage server/cluster uses the ISO 8859-1 (Latin-1) character set when communicating with NFS clients and/or NIS servers. When NIS servers and NFS clients use different character sets, the administrator must specify which character set the NFS clients are using, and which character set the NIS servers are using.

The `protocol-character-set` command allows an administrator to specify the character set to be used when communicating with NFS clients and/or NIS servers. Refer to the *Command Line Reference* for more information on the `protocol-character-set` command.

Note: Once the `protocol-character-set` command is issued, the specified character set is put into use immediately, without the need to restart the server/cluster.

File System Security

The server's file system security enables concurrent access from multiple protocols. A file may have UNIX-like security (UID, GID, mode) and CIFS/NFSv4-like security (NT/NFSv4 owner/group, access control list).

The following security modes are supported:

Mode	Clients	Notes
Mixed	CIFS	The server authenticates CIFS sessions by communicating with a domain controller, which returns user security information. Accesses to files with NT permissions are checked against this security information. If a file has UNIX permissions, the security information is mapped to an equivalent UNIX identity and checked against the file permissions.
	NFS	NFS clients identify their users with either an unauthenticated UNIX credential (which provides the user's UID and one or more GIDs) or an authenticated Kerberos credential (which provides the user's Kerberos principal name). Accesses to files that have UNIX-only permissions can be checked directly against a UNIX credential, and Kerberos credentials are mapped to a UID/GID. If a file has NT permissions, the UID/GIDs in a UNIX credential are mapped to equivalent NT identities, and Kerberos credentials are mapped to an NT user/group.

Mode	Clients	Notes
UNIX	CIFS	The server authenticates CIFS sessions by communicating with a domain controller that returns user security information. All files have UNIX permissions, so the security information is mapped to an equivalent UNIX identity and checked against the file permissions.
	NFS	Accesses to files can be checked directly against a UNIX credential, Kerberos credentials are mapped to a UID/GID.

Note: FTP clients follow either the Windows or the UNIX security model depending on how they were authenticated. FTP clients authenticated by an NT domain appear as CIFS clients for the purpose of security. Similarly, FTP clients authenticated through NIS appear as NFS clients.

With both Mixed and UNIX security mode, it is necessary to configure user and group mappings between UNIX and Windows. However, NFS users do not require security mappings when in UNIX mode.

Enabling and Disabling File Services

Use the **Enable File Services** page to enable or disable the desired file services for the system.

To enable file services:

1. **Navigate to the Enable File Services page.**

   From the **Home** page, click **File Services**. Then, click to display the **Enable File Services** page:

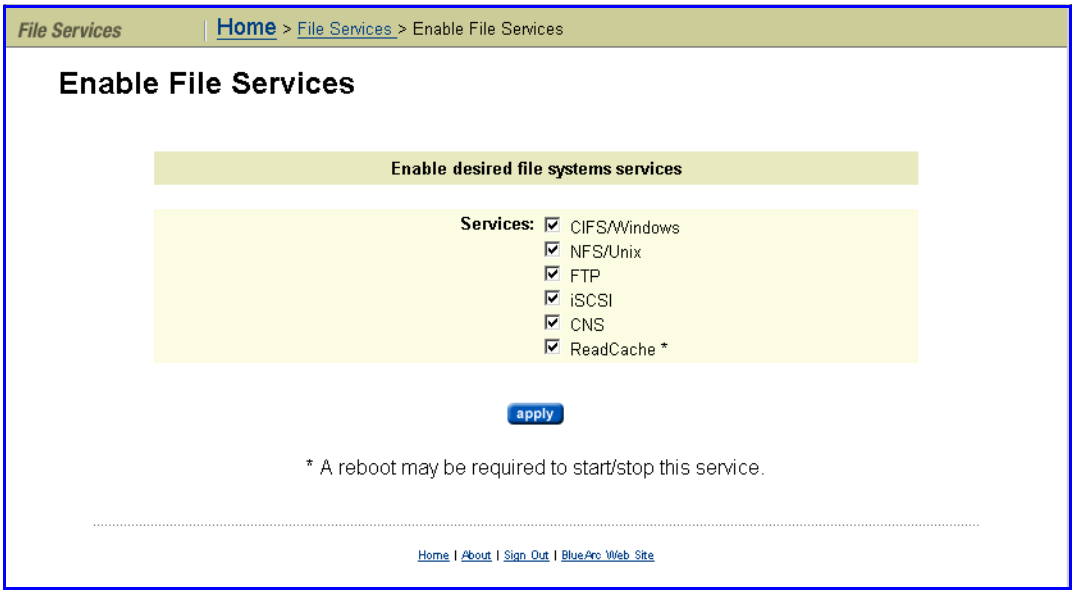

2. **Select/deselect one or more Services:**

- CIFS/Windows
- NFS/Unix
- FTP
- iSCSI
- CNS
- ReadCache

Note: With the exception of FTP, these services require a valid license to enable the service. For information about obtaining licenses, see Adding a License Key, on page 538.

3. **Save your configuration**

   Click **apply** to save.

   If ReadCache has been selected or deselected, a reboot may be required. If so, then follow the on-screen instructions to restart the server.

Managing File System Security

Security modes can be configured per-cluster/server, per-file system, or per-virtual volume. Selecting security modes on a tiered basis, rather than system-wide, enhances the granularity and convenience of managing system security.

NFS Security and Kerberos

The NAS server supports Kerberos to provide authentication, integrity, and privacy when using NFS v2, v3, and v4. Kerberos provides a mechanism for entities (principals) to authenticate to each other and securely exchange session keys. The NAS server supports RPCSEC_GSS using Kerberos v5.

Secure NFS requires configuration of the NFS server's Kerberos principal name, and secret key(s). Kerberos related configuration settings are setup both globally and on a per-EVS basis. The NFS hostname is configured on a per-EVS basis. For information on configuring Kerberos on the NAS server, see Kerberos Configuration, on page 245.

Setting Secure NFS

NFS supports three secure options: Authentication only (the default), Integrity (checksum on data), and Privacy (encryption of data). NFS exports can be set to accept only secure connections. This is done by specifying the appropriate security options in the **Access Configuration** field of the **Add Export** page or the **NFS Export Details** page.

The syntax for setting the secure option is described in IP Address Export Qualifiers, on page 248. Setting the type of secure connections can also be done using the CLI command `nfs-export` with the `mod -c` option. See the CLI Reference for more information.

Mixed Security Mode

The server's mixed security mode supports both Windows and UNIX security definitions. Security is set up uniquely on each file (or directory), based on which user created, or last took ownership of, the file (or directory). If a Windows user, the security definition will be native CIFS and subject to Windows security rules; likewise, if a UNIX user, the security definition will be native NFS and subject to UNIX security rules.

CIFS Access to Native CIFS Files

When a CIFS client tries to access a native CIFS file (that is, with Windows security information), the server checks the user information against the file's security information to determine whether an operation is permissible:

- **User Security.** This information is contained in an access token, which is made up of the user security identifier (SID), primary group SID, and other SIDs. The server receives the token from the domain controller and caches it for use throughout the user's session.
- **File Security.** This information is contained in a file's security descriptor, which is made up of the owner SID, group SID, and access control list (ACL). The ACL can contain several access control entries (ACEs), which specify the conditions for access.

  ACE entries can be modified or deleted using a set of CLI commands called the "cacls" commands. This set of commands includes `cacls-add`, `cacls-del`, `cacls-fields`, `cacls-mask-in`, `cacls-mask-out`, and `cacls-set`. For more information on these commands, refer to the Titan Server *Command Line Reference*.

NFS Access to Native NFS Files

When an NFS client tries to access a native NFS file (with UNIX security information), the server checks the user's UNIX credentials against the file's security information to determine whether or not an operation is permissible. The file security information is made up of a user ID, group ID, and read, write, and execute permissions.

Client Access to Non-Native Files

CIFS users may access files which have UNIX security information, and NFS users may access files which have Windows security information. The server supports this functionality with mapping tables, set up in the Web Manager, that associate the names of NFS users and groups with their Windows equivalents. For example, when a CIFS user tries to access a file that has UNIX-only security information, the server automatically maps the user name to the corresponding NFS name in the mapping table.

- The server automatically translates user security information from UNIX to Windows format, or vice-versa, and caches it for the duration of the session:

UNIX credential		NT access token
UID	User mapping table	User SID
GID	Group mapping table	Primary group SID
Other groups	Group mapping table	Other groups

- The system automatically converts file security attributes from Windows to UNIX format and stores the result in file metadata, henceforth making the files native to both CIFS and NFS clients. Although UNIX files are also converted to Windows format, the results are not stored in file metadata:

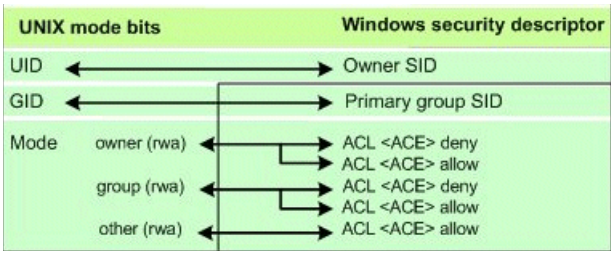

- Any changes that a user makes to a file's security attributes are applied equally to Windows and UNIX.

In summary, when a CIFS user tries to access a file that has UNIX-only security information, the server maps the user to an NFS name and converts the user's access token to UNIX credentials. It then checks these credentials against the file's security attributes to determine whether or not the operation is permissible.

Similarly, when an NFS user tries to access a file that has Windows-only security information, the server maps the user to a Windows name and converts the user's UNIX credentials to a Windows access token. It then checks the token against the file's security attributes.

UNIX Security Mode

When the server is configured in UNIX security mode, it supports UNIX security for CIFS and NFS clients. However, all security settings are saved with Unix file attributes. As a result, NFS clients are always accessing files in native mode, while CIFS clients are always accessing file non-native mode. For more information on both modes of operation, see Mixed Mode Operation, on page 227.

Note: With UNIX security mode, NFS users do not need to rely on the presence of a Windows domain controller (DC) in order to access files. As a result, they are fully isolated from potential DC failures.

Viewing Security Configurations

To view Security Configurations, navigate from the **Home** page to **File Services**, then click to display the **File System Security** page:

The following table describes the fields and columns in this page:

Item/Field	Description
EVS Security Context	Displays the currently selected EVS security context. Click **change** to select a different EVS security context. You can select either Global Configuration which applies to all EVSs, or select a specific EVS.
EVS Security Mode	Displays current EVS security mode settings, and allows you to change those settings.
Default File System Security Mode	Indicates the default security mode that is in effect for the entire EVS. Click the **Switch Mode** link to switch the security mode for the entire EVS. You can switch between Mixed mode and UNIX mode.
Filter	Allows you to control the information displayed in this page. In the **File Systems** field, select whether to show file systems. In the **Virtual Volumes** field, select whether to show virtual volumes. Click **filter** to refresh the page based on the criteria selected in these two fields.
EVS	List of all Virtual Servers (EVSs) defined by the filter.
File System	*If this column is blank*, the displayed security mode is associated with the EVS. *If this column displays a file system label*, the displayed security mode is associated with this specific file system.
Virtual Volume	Lists the virtual volumes found on the file systems defined by the filter.

Item/Field	Description
Mode	Security mode defined on the EVS or file system. File systems without an explicit security mode configuration inherit security mode from the EVS.
details	Click to advance to the **Security Configuration** page where the security mode for the selected EVS can be modified.

Changing Security Mode

Changing the Security Mode for a File System

By default, the file system inherits its parent EVS's security mode; for example, when the parent EVS has a Unix security mode, its child file system inherits that security mode.

To change a file system to use a different security mode:

1. **Navigate to the File System Security page.**

   From the **File Services** page, click to display the **File System Security** page.

2. **Click the details link to view the Security Configuration page.**

   From the **File System Security** page, make sure you are viewing the correct EVS and that your filter is properly set for you to view the desired file system. Click the **details** button to view the **Security Configuration** page.

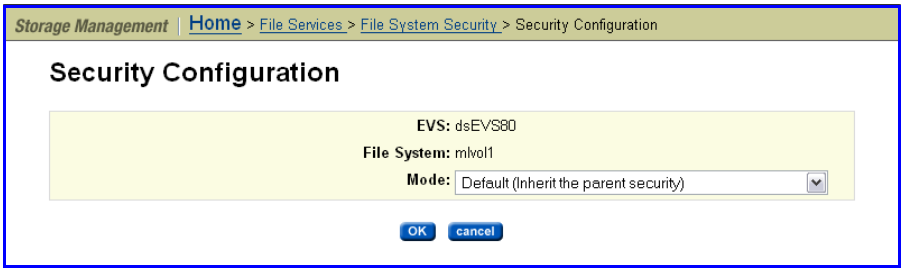

3. **Select a security mode.**

   From the drop-down menu, select a security mode, then click **OK** to select or **cancel** to decline.

Changing the Security Mode for a Virtual Volume

By default, the virtual volume inherits the parent file system's security mode; for example, if the parent file system has a Unix security mode, its child virtual volume inherits that security mode:

1. **Navigate to the File System Security page.**

   From the **File Services** page, click to display the **File System Security** page.

2. **Click the details link to view the Security Configuration page for the virtual volume.**

   From the **File System Security** page, make sure you are viewing the correct EVS and that your filter is properly set for you to view the desired virtual volume. Click the **details** button to view the **Security Configuration** page.

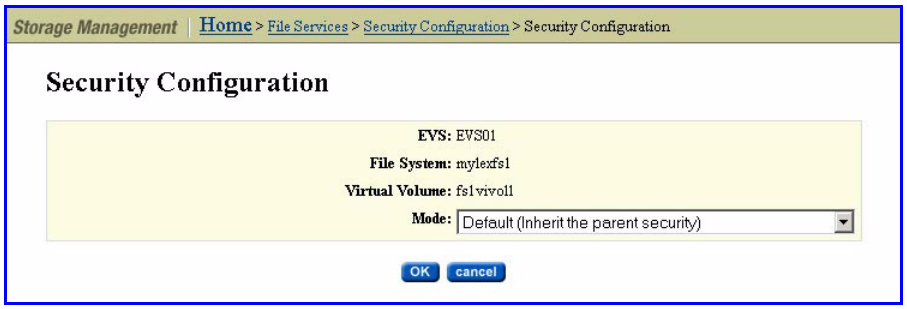

   The following table describes the fields in this table:

Item/Field	Description
EVS	The virtual server hosting the file system.
File System	The parent file system of the virtual volume.
Virtual Volume	The virtual volume which has the security mode you want to change.
Mode	Using the drop-down list, select a security mode for the virtual volume.

3. **Select a security mode.**

   From the drop-down menu, select a security mode, then click **OK** to select or **cancel** to decline.

Mixed Mode Operation

The storage server allows network clients to share a common pool of storage on both Windows and UNIX clients. This is called *mixed mode operation*. Although the server does this as seamlessly as possible, the two protocols are considerably different, so mixed mode operation presents some challenges, discussed in the *File Name Representation* and *Symbolic Links* sections that follow.

File Name Representation

The maximum length of a file name is 255 characters.

File names may contain any Unicode character. Windows NT 4.0, Windows 2000, Windows 2003, and Windows XP clients can make full use of Unicode,

but Windows 9x and NFS clients support only the Latin-1 version of extended ASCII.

Case-sensitivity in file names is significant to NFS and FTP clients, but not CIFS clients.

Symbolic Links

Symbolic links (symlinks) are commonly used:

- To aggregate disparate parts of the file system.
- As a convenience, similar to a shortcut in the Windows environment.
- To access data outside of a cluster. For example, a symlink can point to data in another server in a server farm or in a non-SGI storage system.

There are two types of symlinks:

- **Relative symlinks** contain a path relative to the symlink itself. For example, `../dst` is a relative symlink.
- **Absolute symlinks** contain a path relative to the root of the file system on the NFS client that created the symlink (not relative to the root of the server's file system). For example, `/mnt/datadir/dst` is an absolute symlink.

When accessing the file system through NFS, the server fully supports symlinks. NFS/UNIX clients assume that files marked as symbolic links contain a text pathname that the client can read and interpret as an indirect reference to another file or directory. Any client can follow a symlink, but accessing the target file (or directory) still requires permission.

CIFS clients, however, are not able to follow files marked as symlinks, so the server provides a server-side *symlink following capability*. When a CIFS or FTP client accesses a server-side symlink, the server reads the path from the link and attempts to follow it automatically:

- **For relative symlinks,** the link can be followed, because the server can follow the path from the link itself.
- **For absolute symlinks,** however, the server does not have access to the root of the file system on the NFS client that created the link, and is therefore not able to follow the link automatically.

  To overcome this problem, the server provides *global symlinks,* which allow CIFS clients to follow absolute symlinks via Microsoft's DFS mechanism. With global symlinks enabled, a CIFS client that accesses an absolute symlink is referred via DFS to the link's destination. The link's destination may be on the same file system as the link, on a different file system within a server farm, or on a remote CIFS server. In order to associate a global symlink with an absolute symlink, the server maintains a translation table between *absolute symlink paths* and *global symlink paths*.

Note: CIFS2 (SMB2) clients are not able to follow symlinks (relative or absolute) to files on storage accessed through the Titan Server.

When accessing server-side symlinks, CIFS clients cannot follow some symlinks which are perfectly valid for NFS, because the storage system follows the symlink on the CIFS client's behalf and presents the linked-to file instead of the symlink. In this case, in line with the behavior of Samba, the server hides the existence of the symlink entirely from the CIFS/FTP client. By default, a symlink that points outside of the scope of its own share (for example, to a different file system) are not followed.

Global symlinks (also called absolute symlinks) start with a slash character (/), and they allow you to set up links to data outside a cluster. NFS clients follow the global symlink directly and, for CIFS clients, the server maintains a server-side translation table, that allows those clients to access the symlink destination. Both NFS and CIFS clients can follow the same global symlink to the destination directory, when the global symlink, the exports, shares, and mount points are set up correctly. When a client encounters a global symlink:

- **For NFS clients,** the server simply returns the content of the global symlink, allowing the client to follow the link to the destination. This means that the NFS client's mount points and the NFS exports must be set up correctly.
- **For CIFS clients,** the server causes the client to request a symlink lookup from the local EVS translation table. Once the client requests the lookup, the server returns the destination server name, share name, and path to the CIFS client, allowing it to access the destination.

Caution: *Symlink Destination Directory Alert!* Once the CIFS client follows the path for the global symlink, it may not ask the server for another lookup for that symlink for an extended period of time. Because the symlink is not looked up every time the client follows the symlink, if the destination directory is changed or deleted, the CIFS client may attempt to connect to the wrong destination, possibly causing the client to report an error.

Using global symlinks with CIFS has a performance penalty. Therefore, global symlinks are disabled by default, but can be enabled by filling the **Follow Global Symbolic Links** checkbox on the **Add Share** page (when creating the share) or **CIFS Share Details** page (after the share has been created).

Symlink translation tables are maintained on a per-EVS basis, meaning that:

- **Table entries *do migrat*e with the EVS.** If an EVS is migrated, all of its table entries migrate along with the EVS.
- **Table entries *do not replicate* from the EVS.** When replicating data from one EVS to another, the mapping information for global symlinks is not automatically relocated, and it must be recreated in the translation table of the EVS into which the data was copied.
- **Table entries *do not move* with a file system. I**f a file system is moved from one EVS to another, the mapping information for global symlinks is not automatically relocated and must be manually adjusted, except for those symlinks that are relative to a CNS tree (those symlinks do not require adjustment).

- **Table entries *irrelevant* for symlinks that are relative to a CNS.** When an EVS is migrated, no adjustment is necessary for symlinks that are relative to a CNS because, when the client follows the symbolic link, it is first referred to the CNS tree, then from the CNS tree to a real file system when the path crosses a CNS link.

The symlink translation table is managed via the following CLI commands:

```
global-symlink-add
global-symlink-del
global-symlink-delall
global-symlink-list
```

Refer to the *Command Line Reference* for information on the CLI commands used to manage global symlinks.

Mixed Mode Operation and LDAP Servers

The storage server supports mixed mode access for file systems, meaning that a mapping is required between the file system permissions and owners in order to ensure consistent security and access. NIS/LDAP services allow the server to locate and map users and permissions based on an existing NIS/LDAP service on the network, instead of creating a local account on the storage server.

On an existing LDAP service, one of the following methods will typically be used for allowing the server to locate and map users and permissions:

- RFC 2307 Schema

  RFC 2307 defines a standard convention for the storage and retrieval of user and group mapping information from an LDAP server. If your site uses RFC 2307 schema, and you configure your storage server/cluster to support both mixed mode operations and LDAP services, it is assumed that you have already loaded the RFC 2307 schema into your directory, and that you have already provisioned the user objects appropriately.

- Services for Unix (SFU) schema

  If you have configured SFU (Services for Unix), you must explicitly enable NIS participation **for each account in the active directory (AD) domain**. You can enable NIS participation for an individual account on the UNIX Attributes tab of the user account properties in the Active Directory Users and Computers utility.

To ensure optimum performance when your server/cluster is configured to support both mixed mode operations and LDAP services, the most optimized configuration includes the creation of indexes in the LDAP service for attributes queried by the storage server. To ensure fastest responses to queries, exact-match indexes should be configured on the LDAP server for the

attributes to be searched. The LDAP server on your network should index at least the following attributes:

Objects that:	RFC 2307 Class	Services for Unix Class	Map to NIS Class
Describe user accounts	posixAccount	user	posixAccount
Describe the group identifier	posixGroup	group	posixGroup

Attributes for:	RFC 2307 Attribute	Services for Unix Attribute	Map to NIS Attribute
User ID/login name	uid	sAMAccountName	memberUid
User ID number	uidNumber	msSFU30UidNumber	uidNumber
Group name	cn	cn	memberNisNetgroup
Group ID number	gidNumber	msSFU30GidNumber	gidNumber

The RFC 2307 or Services for Unix attributes may or may not be previously indexed on the LDAP server, depending on the distributor of the director services.

To track indexing performance, you can use the `ldap-stats` command, which permits you to monitor response times for LDAP queries. It is necessary to first let the storage server complete some successful user lookups so that some statistical data can be gathered. In a short period of time, however, you should be able to determine whether any of the attributes are not indexed.

File Locks in Mixed Mode

When a CIFS client reads or writes to a file, it respects the locks that both CIFS and NFS clients have taken out. In contrast, an NFS client respects the locks taken out by CIFS clients only. NFS clients must therefore check for existing NFS locks with the Network Lock Manager (NLM) protocol. The server supports both monitored and non-monitored NFS file locks.

Opportunistic Locks (Oplocks)

An oplock is a performance-enhancing technique used in Microsoft networking (CIFS) environments. It enables applications to speed up file access and minimize network traffic by caching all or part of a file locally. As the data is kept on the client, read and write operations can be performed locally, without involving the server.

The server supports three categories of oplocks:

- **Exclusive.** An Exclusive oplock enables a single client to cache a file for both *reading and writing*. As the client that owns the oplock is the only

client accessing the file, it can read and modify all or part of the file locally. The client does not need to post any changes to the server until it closes the file and releases the oplock.

- **Batch.** A Batch oplock enables a single client to cache a file for both *reading and writing*, as in the case of an exclusive oplock. In addition, the client can preserve the cached information even after closing the file; file open and close operations are also performed locally. The client does not need to post any changes back to the server until it releases the oplock.
- **Level II.** A Level II oplock enables multiple clients to cache a file for *reading only.* The clients owning the oplock can read file data and attributes from local information, cached or read-ahead. If one client makes any changes to the file, all the oplocks are broken.

When dealing with oplocks, the server acts in accordance with the CIFS specification. Whether operating in a pure Windows environment or with a mix of CIFS and NFS clients, the server allows applications to take advantage of local caches while preserving data integrity.

Exclusive and Batch Oplocks

An Exclusive or Batch oplock is an exclusive (read-write/deny-all) file lock that a CIFS client may obtain at the time it opens a file. The server grants the oplock only if no other application is currently accessing the file.

When a client owns an Exclusive or Batch oplock on a file, it can cache part or all of the file locally. Any changes that the client makes to the file are also cached locally. Changes do not need to be written to the server until the client releases the oplock. In the case of an Exclusive oplock, the client releases the oplock when the server requests that it does so, or when it closes the file. In the case of a Batch oplock, the client may keep information (including changes) locally even after closing the file. While the client has an Exclusive or Batch oplock on a file, the server guarantees that no other client may access the file.

If a client requests access to a file that has an Exclusive or Batch oplock, the server asks the client with the oplock to release it. The client then writes the changes to the server and releases the oplock. Once this operation has finished, the server allows the second client to access the file. This happens regardless of the second client's network protocol.

In cases where a CIFS client requests an oplock on a file that has an Exclusive or Batch oplock, the server breaks the existing oplock and grants both clients Level II oplocks instead.

Level II Oplocks

A Level II oplock is a non-exclusive (read-only/deny-write) file lock that a CIFS client may obtain at the time it opens a file. The server grants the oplock only if all other applications currently accessing the file also possess Level II oplocks:

- *If another client owns an Exclusive or Batch oplock,* the server breaks it and converts it to a Level II oplock before the new client is granted the oplock.
- *If a client owns a Level II oplock on a file,* it can cache part or all of the file locally. The clients owning the oplock can read file data and attributes from local information without involving the server, which guarantees that no other client may write to the file.
- *If a client wants to write to a file that has a Level II oplock,* the server asks the client that has the oplock to release it, then allows the second client to perform the write. This happens regardless of the network protocol that the second client uses.

User and Group Names in NFSv4

In NFSv4 users and groups are identified by UTF-8 strings of the form: user@dns_domain and group@dns_domain. The NAS Server supports the following universal user/group identifiers:

OWNER@	The owner of a file.
GROUP@	A file's group.
EVERYONE@	The world.
NETWORK@	Accessed via the network.
ANONYMOUS@	Accessed without any authentication.
AUTHENTICATED@	Any authenticated user.

Configuring User and Group Mappings

When the server is operating in either mixed or UNIX security mode, it creates mappings between UNIX and Windows users and groups. For example, user John Doe could have a UNIX user account named `jdoe` and a Windows user account named `johnd`. These two user accounts are made equivalent by setting up a user mapping. Furthermore, the server assumes that equivalent user and group names are the same for both environments. For example, if no explicit mapping is found for user `janed`, the server assumes that the UNIX user account named `janed` is the same as the Windows user account with the same name.

There are two steps to follow when setting up user and group mappings on the server:

- **Specify each NFS user and group's name and ID.** Note that this step is not required for Windows users or groups, as the server obtains all of the information it needs from the domain controller (DC).
- **Map the NFS user (group) names to Windows NT user (group) names.**

Managing NFS User and Group Mapping

Windows access to a file created by a UNIX user (or vice-versa) is permitted when the UNIX name and Windows name are recognized as being the same

user. However, NFS clients present an NFS operation to an NFS server with numerical UNIX User ID (UID) and UNIX Group ID (GID) as credentials. The server must map the UID and GID to a UNIX user or group name prior to verifying the UNIX to Windows name mapping.

The server uses the following methods to map from a numerical UNIX UID or GID to a UNIX user name or group name:

- If the server is configured to use the Network Information Service (NIS) no special configuration steps are needed; the server automatically retrieves the user (group) names and IDs from the NIS server.
- NFS user and group names can be added manually. For information on adding user and group mappings manually, see Specifying NFS User Mappings Manually, on page 234 and Specifying NFS Group Mappings Manually, on page 237.
- NFS user and group names can be added by importing files. For example, the UNIX `/etc/passwd` file can be imported, providing the server with a mapping of user name to UID. The `/etc/groups` file should also be imported to provide the server with a mapping of Group name to GID. For information on importing user and group mappings from files, see Importing User or Group Mappings From a File, on page 239.
- You can import the numerical ID to Name mappings directly from a NIS server or an LDAP server if one has been configured. Every time a UID is presented to the server, it will issue an NIS request to an NIS server to verify the mapping. This mapping can remain cached in the server for a configurable time. A cached ID to name binding for a User or Group will appear as *Transient* in the NFS Users or Groups list. For information on importing user and group mappings from a NIS or an LDAP server, see Importing Users or Groups From an NIS or LDAP Server, on page 242.

Note: When a Windows user creates a file and the UNIX user or group mapping fails, the server sets the UID or the GID to 0 (root). In previous releases, the server sets the UID or GID to 0 (root) or to 65534 (nobody).

Specifying NFS User Mappings Manually

Each UNIX user name and numerical UID can be manually entered, along with its corresponding Windows user and domain name. Users configured manually will appear as *permanent* in the NFS users list.

To specify an NFS name manually:

1. **Navigate to the User Mapping page.**

   From the **File Services** page, click **User Mappings** to display the page:

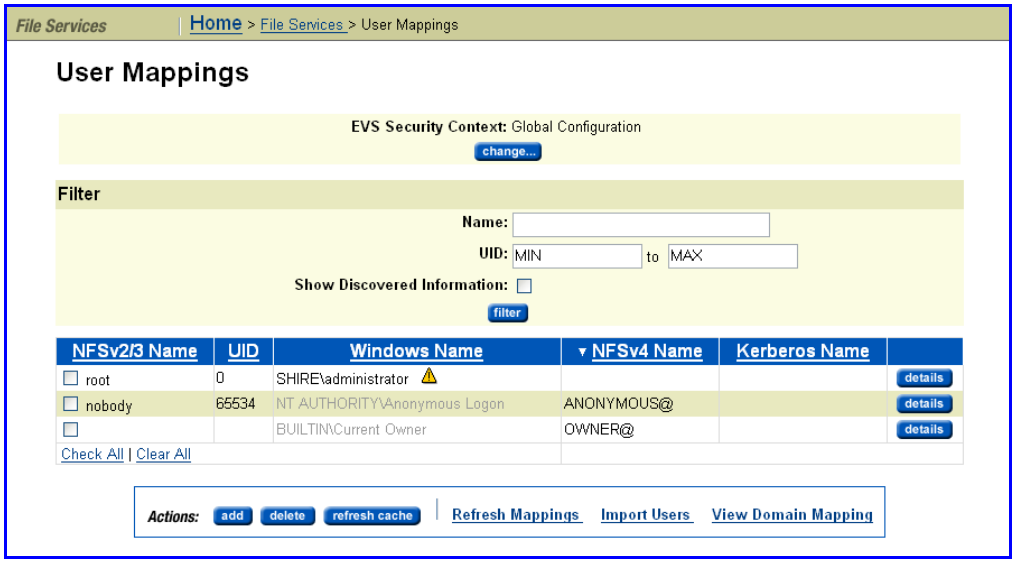

The fields on this page are described in the table below:

Item/Field	Description
EVS Security Context	Displays the currently selected EVS security context. Click **change** to select a different EVS Security Context or to select the global configuration. Selecting a different EVS Security Context changes the EVS to which the mapping applies.
Filter	Filter the list of user mappings using any of the following criteria: • **Name**, which applies to the NFSv2/v3, NFSv4 user names or the Windows user name. • **UID**, which can be used to specify a range of UID values to display, or a minimum/maximum UID value to display. • Fill the **Show Discovered Information** checkbox to display only information that has been discovered from NIS servers, LDAP servers, or domain controllers.
NFSv2/3 User Name	User name configured in the UNIX environment.
UID	User ID configured in the UNIX environment.
Windows User name	User name configured in the Windows environment.
NFSv4 Name	Displays the NFSv4 user name. For more information see User and Group Names in NFSv4, on page 233
Kerberos Name	Displays the Kerberos principal (of the form user@realm) for the user.

2. If necessary, change the EVS Security Context.

The **EVS Security Context** displays the currently selected EVS security context. Changes made to mappings using this page apply only to the currently selected EVS security context.

- If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
- If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the mappings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

Click **Change** to select a different EVS security context or to select the global configuration.

3. **To create, modify, or delete a user mapping.**
   - **To modify a user mapping:**

     To modify a user mapping, click **details**, and edit the **NFSv2/3 Name**, **UID**, **Windows Name**, **NFSv4 Name**, or the **Kerberos Name**, in the specified fields and click **OK**.

     Fill the **Stored** checkbox to store the mapping information locally as a part of the security context used by the virtual server. If you leave a field blank and fill the Stored checkbox, the server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers. If the **Stored** checkbox is empty, the server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.
   - **To delete a user mapping:**

     To delete a user mapping, fill the checkbox next to the NFSv2/3 Name of the user mapping you want to delete, then click **delete**.
   - **To create a user mapping:**

     Setting up NFS users requires the following steps:

     i Click **add** to display the **Add User Mapping** page.

     ii Specify the appropriate user information to be mapped (**NFSv2/3 Name**, **UID**, **Windows Name**, **NFSv4 Name**, or **Kerberos Name**).

     iii Fill the **Stored** checkbox to store the mapping information locally as a part of the security context used by the virtual server. If you leave a field blank and fill the **Stored** checkbox, the server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers. If the Stored checkbox is empty, the server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.

     iv Click **OK**.

4. **Refresh the SMU's cache.**

   To clear the SMU's cache and repopulate it with the relevant objects, click **refresh cache**. This is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.

Specifying NFS Group Mappings Manually

Each UNIX group name and numerical GID can be manually entered, along with its corresponding Windows group and domain name. Groups configured manually will appear as *permanent* in the NFS group list.

To specify NFS group names manually:

1. **Navigate to the Group Mapping page.**

   From the **File Services** page, click **Group Mappings** to display the page:

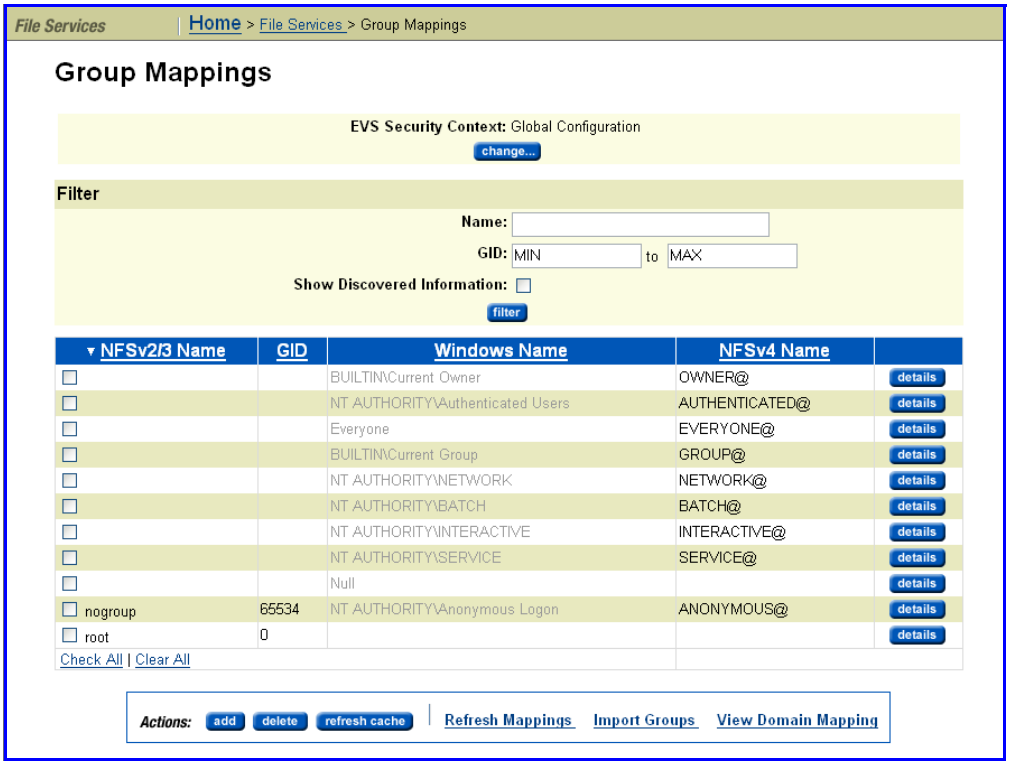

   The following table describes the fields in this page:

Item/Field	Description
EVS Security Context	Displays the currently selected EVS security context. Mappings are added to a particular security context; either the individual security context used by a secure EVS, or the global configuration. Click **change** to select a different EVS Security Context or to select the global configuration. Selecting a different EVS Security Context changes the EVS to which the mapping applies.

Item/Field	Description
Filter	Filters the list of group mappings based on: • **Name**, which applies to all of the name fields. • **GID**, which can be used to specify a range of GID values to display, or a minimum/maximum GID value to display. • Fill the **Show Discovered Information** checkbox to display only information that has been discovered from NIS servers, LDAP servers, or domain controllers.
NFSv2/3 Name	Displays the NFS group name configured in the UNIX environment.
GID	Displays information about the Windows group account, including the where the account is from (typically the domain) and the account name.
Windows name	Displays information about the Windows group account, including the where the account is from (typically the domain) and the account name.
NFSv4 Name	Displays the NFSv4 group name.

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes made to mappings using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the mappings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

   Click **Change** to select a different EVS security context or to select the global configuration.

3. **To create, modify, or delete a group mapping.**

- **To modify a group mapping:**

  To modify a group mapping, click **details**, and edit the **NFSv2/3 Name**, **GID**, **Windows Name**, or **NFSv4 Name**, in the specified fields and click **OK**.

  Fill the **Stored** checkbox to store the mapping information locally as a part of the security context used by the virtual server. If you leave a field blank and fill the Stored checkbox, the server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers. If the **Stored** checkbox is empty, the server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.

- **To delete a group mapping:**

  To delete a group, fill the checkbox next to the NFSv2/3 Name of the group mapping you want to delete, then click **delete**.

- **To create a group mapping:**

  Setting up NFS group mapping requires the following steps:

  i Click **add** to display the **Add Group Mapping** page.

  ii Specify the appropriate group information to be mapped (**NFSv2/3 Name**, **GID**, **Windows Name**, or **NFSv4 Name**).

  iii Fill the **Stored** checkbox to store the mapping information locally as a part of the security context used by the virtual server. If you leave a field blank and fill the Stored checkbox, the server will not try to retrieve that information from NIS servers, LDAP servers, or domain controllers. If the Stored checkbox is empty, the server relies on information discovered from NIS servers, LDAP servers, or domain controllers for this mapping.

  iv Click **OK** to save the mapping.

4. **Refresh the SMU's cache.**

   To clear the SMU's cache and repopulate it with the relevant objects, click **refresh cache**. This is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.

Importing User or Group Mappings From a File

You can specify user or group details by importing them from a file.

This section specifies importing NFSv2/3 data and mappings. NFSv4 user/group names are distinct from the Unix name associated with Unix UIDs/GIDs. However, in many environments a user/group's NFSv4 name can be derived from their Unix name by appending the NFSv4 domain. The storage server can perform this conversion automatically, based on the settings specified on the **Domain Mappings** page of Web Manager or through the CLI command `domain-mappings-add`. (To display the **Domain Mappings** page, go to the **File Services** page, click **User Mapping** or **Group Mapping**, then select the **View Domain Mapping** link, and for more information on the `domain-mappings-add` command, refer to the *Command Line Reference*.)

A UNIX /etc/passwd file can be imported, providing the server with a mapping of user name to UID. The /etc/groups file should also be imported to provide the server with a mapping of Group name to GID.

The server will ignore other fields from the passwd file, such as the encrypted password and the user's home directory. Users or Groups configured by importing from the /etc/passwd file will then appear in the appropriate list on the **User Mappings** page or the **Group Mappings** page.

Choose one of the three following formats and use it consistently throughout the file:

- **NFSv2/3 user/group data only.** The source of the user data can be a UNIX password file, such as /etc/passwd.

  When using Network Information Service (NIS), use the following command to create the file:

  ```
  ypcat passwd > /tmp/x.pwd
  ```

  The resulting file has the following format:

  ```
  john:x:544:511:John Brown:/home/john:/bin/bash
  keith:x:545:517:Keith Black:/home/keith:/bin/bash
  miles:x:546:504:Miles Pink:/home/miles:/bin/bash
  carla:x:548:504:Carla Blue:/home/carla:/bin/bash
  ```

- **NFSv2/3-to-Windows user/group mappings only.** Create a file with entries in the following format:

  ```
  UNIXuser="NT User", "NT Domain"
  ```

  with the following syntax rules:

  - NT domain is optional.
  - NFS user names cannot contain spaces.
  - NT names must be enclosed in quotation marks.
  - If the domain name is omitted, the server domain is assumed. If the empty domain name is required, it must be specified like this:

    ```
    users="Everyone", ""
    ```

    where the Everyone user is the only common account with an empty domain name.

- **Both NFSv2/3 user/group data and NFSv2/3-to-Windows user mappings.** Create a file with entries in the following format:

  ```
  UNIXuser:UNIXid="NT User", "NT Domain"
  ```

  with the same rules for NFS and NT names as for the NFSv2/3-to-Windows user mapping.

  The resulting file has entries in the following format:

  ```
  john:544="john", "Domain1"
  keith:545="keith", "Domain1"
  ```

```
miles:546=”miles”, “Domain1”
carla:548=”carla”, “Domain1”
```

To specify NFS user/group mappings by importing a file:

1. **Navigate to the User Mappings page or the Group Mappings page.**

   From the **File Services** page, click **User Mapping** to display the **User Mapping** page or click **Group Mapping** to display the **Group Mapping** page.

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes made to mappings using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the mappings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

   Click **Change** to select a different EVS security context or to select the global configuration.

3. **Navigate to the Import User Mappings page or the Import Group Mappings page.**

   - To import **users**, from the **User Mappings** page, click **Import Users** to display the **Import User Mappings** page:

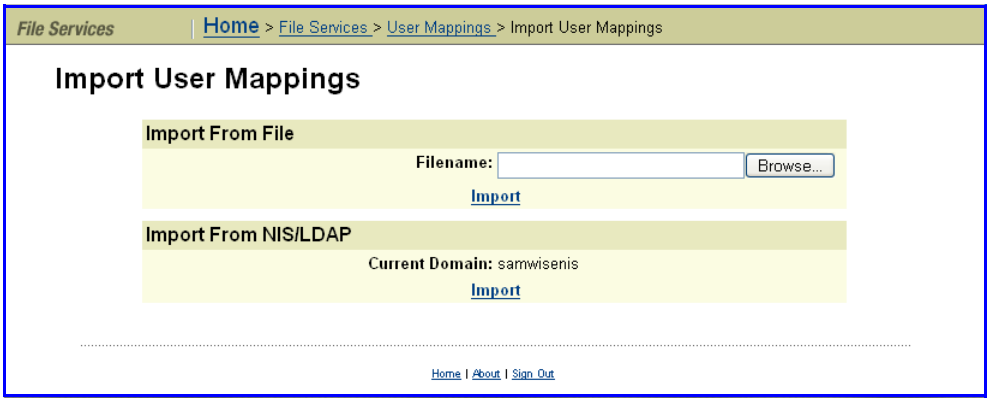

- To import **groups**, from the **Group Mappings** page, click **Import Groups** to display the **Import Groups Mappings** page:

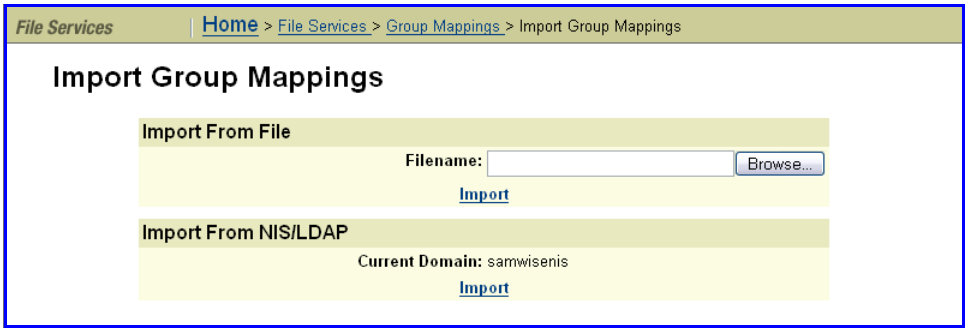

4. **Specify the file containing the user/group mappings to import.**

   Enter the filename in the **Filename** field, or **Browse** to locate the file. If any user/group names in the file already exist in the NFS list, Web Manager ignores them and displays a warning that it encountered errors or duplicate users.

5. **Import.**

   Click **Import**.

Importing Users or Groups From an NIS or LDAP Server

Administrators can import user or group detail from an NIS server or an LDAP server.

To import NFS user/group mappings from an NIS server or an LDAP server:

1. **Navigate to the User Mappings page or the Group Mappings page.**

   From the **File Services** page, click **User Mapping** to display the **User Mapping** page or click **Group Mapping** to display the **Group Mapping** page.

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes made to mappings using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To change the mappings of an EVS using an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

   Click **Change** to select a different EVS security context or to select the global configuration.

3. **Navigate to the Import User Mappings page or the Import Group Mappings page.**
   - To import **users**, from the **User Mappings** page, click **Import Users** to display the **Import User Mappings** page:

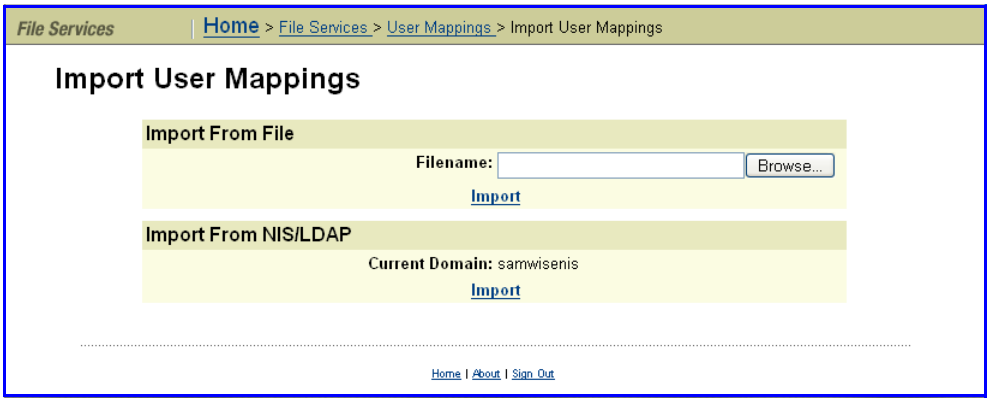

   - To import **groups**, from the **Group Mappings** page, click **Import Groups** to display the **Import Groups Mappings** page:

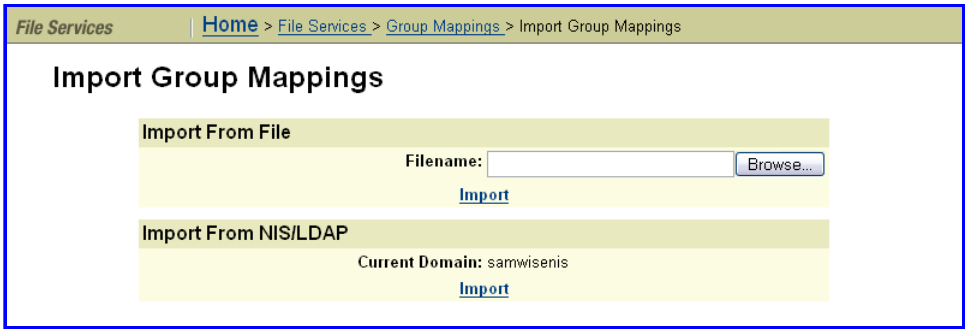

4. **Import user/group mappings.**

   Click **Import**. The NIS or LDAP server displayed in the **Current Domain** will be contacted, and mappings will be imported.

   To change the current domain, go to the **NIS/LDAP Configuration** page and click **modify** to change the domain.

Sharing Resources with NFS Clients

A fundamental component of most UNIX networks, the Network File System (NFS) protocol provides PC and UNIX workstations with transparent access to one another's files. This section describes how to set up NFS exports, and explains about NFS statistics, supported clients, and prerequisites.

The server implements the file-serving functions of an NFS server, providing normal file-serving functions, such as:

- Export manipulation
- File manipulation (read, write, link, create, etc.)

- Directory manipulation (mkdir, readdir, lookup, etc.)
- Byte-range file locking
- File access control (permissions)
- File and directory attributes (sizes, access times, etc.)
- Hard links and symbolic (soft) links

Enabling NFS Protocol Support

Prerequisites

To enable NFS access to the system:

- Enter an NFS license key.
- Enable the NFS service.

Supported Clients and Protocols

The following platforms and versions are supported:

Platform	Supported versions
Red Hat Linux	7, 8, 9
Fedora Linux	Core 1, Core 2, Core 3, Core 4
Solaris (SPARC)	5 through 9
Solaris (Intel)	8, 9, 10
Macintosh OS X	10.3 or later
FreeBSD	4.3, 4.7, 5.0
HP-UX	10, 11
Irix	6.5

It also supports the following UNIX protocols:

Protocol	Supported versions
NFS	2 and 3
Port Mapper	2
Mount	1 and 3
Network Lock Manager (NLM)	1, 3, and 4
Network Status Monitor (NSM)	1

NFS Statistics

NFS statistics for the storage server (in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Configuring NFS Exports

NFS exports are configured on mounted file systems. NFS exports can be configured manually or export details can be imported from a file (For more information, see Backing Up and Restoring NFS Exports, on page 252). NFS exports can be configured manually or by importing the details from a file.

The NFSv4 Pseudo File System

NFSv4 introduces the concept of the pseudo file system, where exports appear as directories. NFSv4 clients do not connect directly to NFS exports as in NFSv2/3. Instead all clients connect to the root of the pseudo file system, which is a virtual directory. The pseudo file system is generated automatically from the NFS exports, and is maintained automatically as exports are modified and removed. You can choose to present all the file systems in a single pseudo file system.

The NAS Server allows you to create views of many file systems from one point of contact, name spaces. These views are available on a per EVS basis or for the entire cluster.

Pseudo File System Example

A server named `numbers` has two exports: `/one` and `/two`. If a client wishes to get access to export `/one`, there are two ways to mount exports:

```
mount -t nfs4 numbers:/  /mnt
```
which mounts the pseudo file system at `/mnt`

```
mount -t nfs4 numbers:/one /mnt
```
which mounts the export `/one` at `/mnt`

The first method is only supported in NFSv4. The second method is supported in versions 2, 3, and 4. In the first method, the client can get to export `/one` with the command `cd /mnt/one`, and to the export `/two` with `cd /mnt/two`.

Kerberos Configuration

Configuring the NAS server requires three steps:

1. Export a keytab file from the KDC (Key Distribution Center). We recommend using MIT Kerberos version 5.
2. Import the keytab file into the NAS server.
3. Set the Kerberos realm for the NAS server.

After performing these steps, the NAS server is able to complete the configuration. However, you may choose to create mappings between the Kerberos users/groups and the Active Directory users/groups.

Export a keytab file from the KDC

The keytab file needs to contain the service principal for the NFS service for the EVS. Once the NFS service principal for the EVS has been added, create a keytab file specifically for the EVS. The type of key is critical. Create only one key for the NFS service of type `des-cbc-crc:normal`.

For example, with an EVS named "man" in the Kerberos realm "`AESIR.EXAMPLE.COM`", the keytab file for the NFS service on "man" should contain a principal "`nfs/man.aesir.example.com@AESIR.EXAMPLE.COM`". The format of the principal starts with the service (nfs), followed by a slash, then the fully-qualified-domain name of the EVS, then the symbol @, and finally the Kerberos realm. (Note that case is significant. Kerberos realms are always in uppercase. Also, there must be no trailing period after the Kerberos realm.)

Typically you will use the kadmin utility run from the master KDC to export a keytab file. For details on creating an appropriate keytab file, refer to the documentation for the tools supplied with your version of Kerberos. (We recommend using MIT Kerberos version 5.)

Import the Keytab File into the NAS Server

Transfer the keytab file to the flash of the NAS server.

For example: securely move the keytab file to the SMU and transfer it to the NAS server. Login with ssc, and do the following:

```
SERVER:$ ssput man.nfs.keytab man.nfs.keytab
```

The first name is the local file name on the SMU, the second name is the name to use on the NAS server. Once the file has been placed on the NAS server, import the keytab in the context of the EVS with:

```
SERVER:$ krb5-keytab import man.nfs.keytab
```

After the keytab has been imported, the uploaded keytab file can be safely removed with:

```
SERVER:$ ssrm man.nfs.keytab
```

Set the Kerberos Realm for the NAS Server

Set the realm by using the command `krb5-realm`. For example:

```
SERVER:$ krb5-realm AESIR.EXAMPLE.COM
```

Adding an NFS Export

To add an NFS export:

1. **Navigate to the add Export page.**

From the **File Services** page, select **NFS Exports,** then click **add** to display the **Add Export** page:

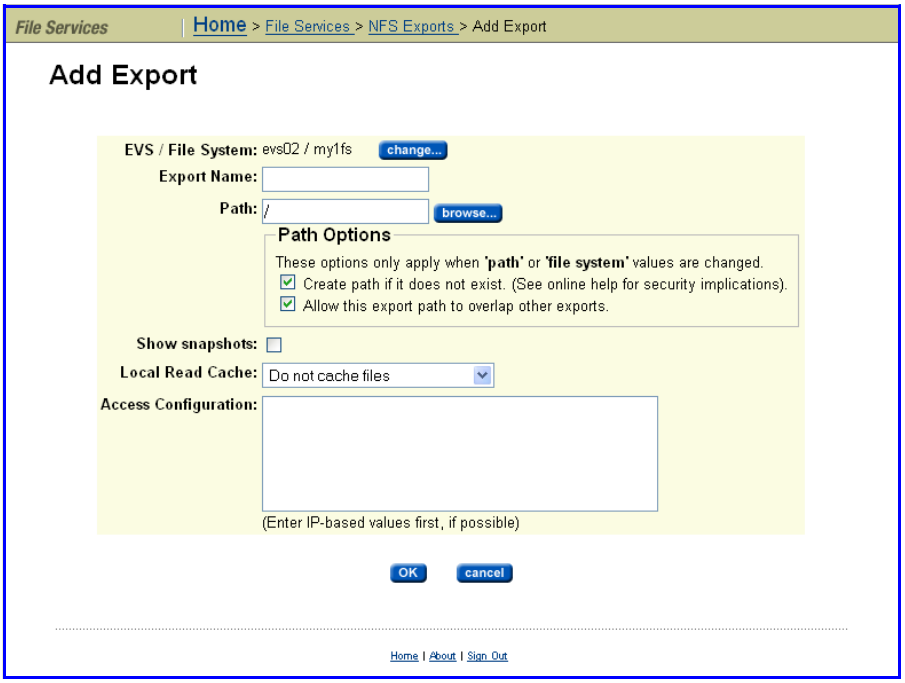

2. **Add the Export:**

   The fields in the **Add Export** page are described in the table below:

Item/Field	Description
EVS/File System	Currently selected EVS and file system, to which the NFS Export will link. To add an NFS Export to a different EVS or file system (or CNS link to a file system), click **change.**
Name	Name of the export.
Path	Path to the source directory for the export (case-sensitive). Click **browse** to locate the path.
Path Options	• Create the path if it does not exist: Fill the **Create path if it does not exist** checkbox to create the path entered in the **Path** field. • Allow this export to overlap other exports: When filled (default), nested NFS exports are allowed. Use this in situations where you want to export the root directory of a file system and make it available to only one group of users (for example managers), and then also to export the sub-directories of the root directory and make each sub-directory available to a different group of users.
Show snapshots	Fill to allow snapshot access from the NFS export.

Item/Field	Description
Local Read Cache 	To allow caching of files or cross file system links from the file system to which this export points, select one of the following: • **Cache all files**. Allows caching of files and cross file system links in the file system of the export. Cross file system links are local links that point to a data file in a remote file system. The remote file system may be on a remote server or storage device. • **Cache cross-file system links**. Allows only cross file system links to be cached. • To disallow read caching of files and cross file system links, do not change the default selection of **Do not cache files.** Local read caching is not supported for NFSv4 clients. For more information on read caching, see Read Caching, on page 412.
Access Configuration	IP addresses, host names, or the NIS netgroups of the clients who are allowed to access the NFS export (up to 2,000 characters). If the system has been set up to work with a name server, you can enter the NIS netgroup to which the clients belong, or the client's computer name rather than its IP address (not case sensitive). You can also specify the flavor of NFS security using the option `(sec=<mode>)`. The syntax is described in IP Address Export Qualifiers, on page 248. See IP Address Export Qualifiers, on page 248, Specifying Clients by Name (instead of IP Address), on page 250, and NFS Security and Kerberos, on page 222 for more information on specifying access configurations.

What to type	Means
Blank or *	All clients can access the export.
Specific address or name. Examples: **10.168.20.2**, **client.dept.company.com**	Only clients with the specified names or addresses can access the export.
A range of addresses using Classless Inter-Domain Routing (CIDR) notation. Example: **10.168.1.0/16**	Clients with addresses within the range can access the export.
Partial address or name using wildcards. Examples: **10.168.*.***, ***.company.com**	Clients with matching names or addresses can access the export.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

IP Address Export Qualifiers

The following table describes qualifiers that can be appended to IP addresses when specifying client access to an NFS export:

Qualifier	Description
read_write, readwrite, rw	Grants read/write access. This is the default setting.
read_only, readonly, ro	Grants read-only access.

Qualifier	Description
root_squash, rootsquash	Maps user and group IDs of 0 (zero) to the anonymous user or group. This is the default setting.
no_root_squash, norootsquash	Turns off root squashing.
all_squash, allsquash	Maps all user IDs and group IDs to the anonymous user or group.
no_all_squash, noallsquash	Turns off all squashing. This is the default setting.
secure	Requires requests to originate from an IP port less than 1024. Access to such ports is normally restricted to administrators of the client machine. To turn it off, use the insecure option.
insecure	Turns off the secure option. This is the default setting.
anon_uid, anonuid	Explicitly sets an anonymous user ID.
anon_gid, anongid	Explicitly sets an anonymous group ID.
noaccess, no_access	Denies the specified clients access to the export.
`(sec=<mode>)`	Allows you to specify the flavor of NFS security, where `<mode>` is a colon delimited list of allowed security flavors (sys:krb5:krb5i:krb5p).

Here are some examples:

- `10.1.2.38(ro)`

  Grants read-only access to the client whose IP address is 10.1.2.38.

- `10.1.2.0/24(ro)`

  Grants read-only access to all clients whose IP address is within the range 10.1.2.0 to 10.1.2.255.

- `yourcompanydept(ro)`

  Grants read-only access to all members of the NIS group yourcompanydept.

- `*.mycompany.com(ro, anonuid=20)`

  Grants read-only access to all clients whose computer name `ends.mycompany.com`. All squashed requests are to be treated as if they originated from user ID 20.

- `10.1.*.* (readonly, allsquash, anonuid=10, anongid=10)`

  Grants read-only access to all the matching clients. All requests are squashed to the anonymous user, which is explicitly set as user ID 10 and group ID 10.

- The order in which the entries are specified is important. Take the following two lines:

  ```
  *(ro)
  10.1.2.38(rw)
  ```

  The first grants read-only access to all clients, whereas the second grants read/write access to the specified client. The second line is redundant, however, as the first line matches all clients. These lines must be transposed to grant write access to 10.1.2.38.

- `10.1.1.*(sec=sys),10.1.2.*(sec=krb5:krb5i:krb5p),*(sec=krb5p)`
  - Clients in the `10.1.1.*` subnet use `sys` authentication.
  - Clients in the `10.1.2.*` subnet to use `krb5`, `krb5i`, or `krb5p`.
  - All other clients use `krb5p`.

Note: To improve performance, when specifying clients that can access an export, SGI recommends specifying IP addresses or IP address ranges, including those that include wildcards, before specifying host names or NIS netgroups.

Specifying Clients by Name (instead of IP Address)

- **Full Qualified Domain Name Required.** Be sure to specify the fully qualified domain name of the client. For example, use `aclient.dept.mycompany.com` rather than simply `aclient`.
- **Leading Wildcard Allowed.** To specify a partial name, a single wildcard, located at the start of the name, may be used.
- **Export Options Change Requires Remount.** When the client mounts the NFS export, it determines which export option to apply to a specific client. Subsequent changes to DNS, WINS, or NIS that would resolve the client's IP address to a different computer name are only applied to mounted exports when the client unmounts the exports and then remounts them.
- **Name Service Order is Significant.** Application of export options to a client's mount request may be affected by the order in which the system applies DNS, WINS, and NIS information to resolve IP addresses. The first service in name order sequence that can resolve the client name supplies the name *and* searches configuration options for the export.

Viewing the Properties of an NFS Export

To view the properties of an NFS export:

1. **Navigate to the NFS Exports page.**
   From the **File Services** page, click **NFS Exports** to display the page:

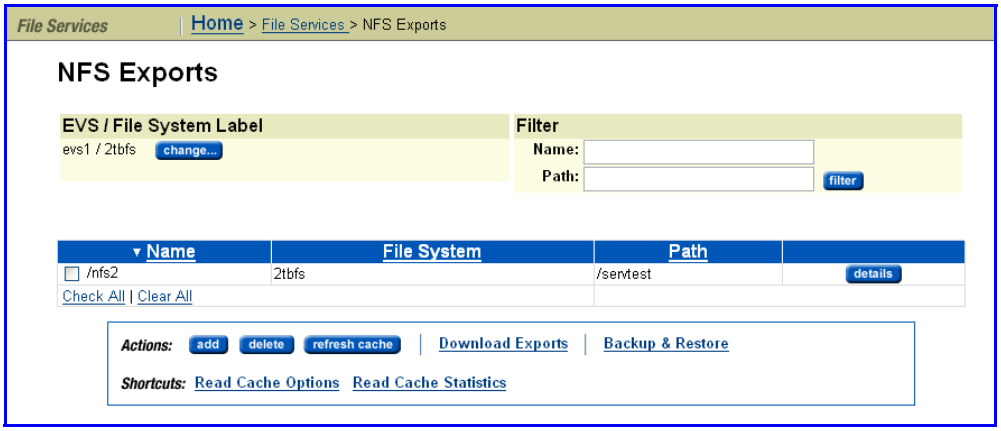

The fields on the **NFS Exports** page are described in the following table:

Item/Field	Description
EVS/File System:	The name of the EVS and the file system to which the NFS Export is assigned. To view the NFS Exports of a different EVS or file system, click **change.**
Filter	This allows the table to be filtered by name and path. Click **filter** to display the NFS Export table.
Name	The name of the NFS Export.
File System	The name of the file system (or CNS link to a file system) to which the NFS Exports is assigned.
Path	The path and directory to which the NFS Export is directed.
Details	Click **details** to display the **NFS Export Details** page, where you can view detailed information about the NFS Export.
Actions	
add	Select a file system and click **add** to add an NFS Export.
delete	Select an NFS Export and click **delete** to delete the NFS Export.
refresh cache	Clears the SMU's cache, and then repopulates it with the relevant objects. Note that this is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.
Download Exports	Click **Download Exports** to download a comma separated value (.csv) file containing a list of all configured NFS exports on the selected EVS and file system. Note that the downloaded file cannot be used to restore NFS exports (you must restore NFS exports from an NFS exports backup file). To download a list of exports from another file system, click **change.**

Item/Field	Description
Backup & Restore	Displays the **NFS Export Backup & Restore** page (for more information, see Backing Up and Restoring NFS Exports, on page 252).

2. **Navigate to the NFS Export Details page.**

   On the **NFS Export** page, select an Export and click **details.**

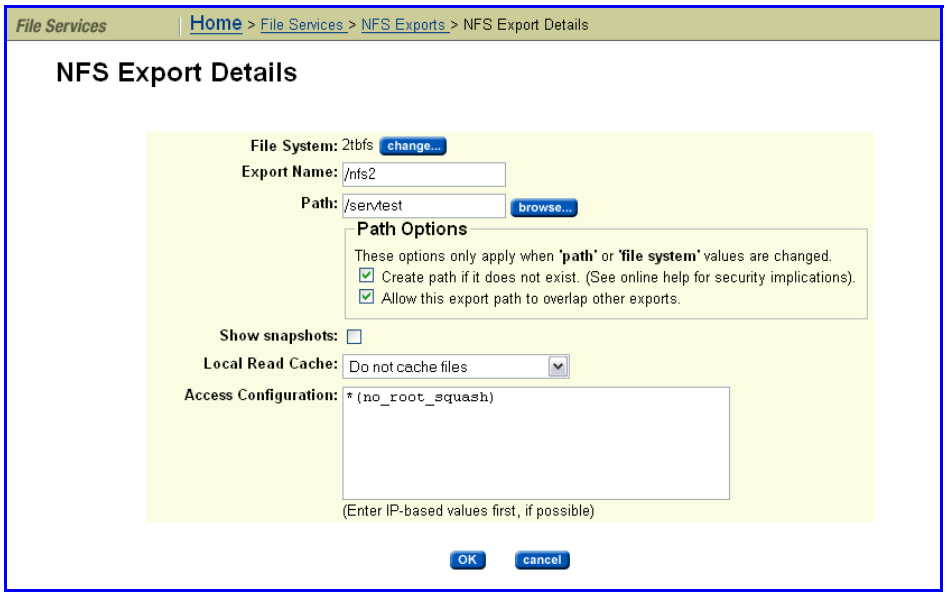

3. **Optionally, modify the settings for this Export.**

   For a description of the items and fields on the **NFS Export Details** page, see Adding an NFS Export, on page 246.

4. **Save your changes.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Backing Up and Restoring NFS Exports

When backing up NFS Exports:

- **Backup scope.** All NFS Exports in all EVSs are backed up (except for those in the CNS tree).
- **Backup format.**The NFS Export backup file is saved as a .txt file.

When you restore NFS Exports from a backup file:

- The restore operation does not modify or overwrite currently existing exports that have the same name.
- With the exception noted above, all exports in the selected backup file are restored.

Backing up or Restoring NFS Exports

To back up or restore NFS exports:

1. **Navigate to the NFS Export Backup & Restore page.**

   From the **File Services** page, select **NFS Exports,** then click to display the **NFS Exports Backup & Restore** page.

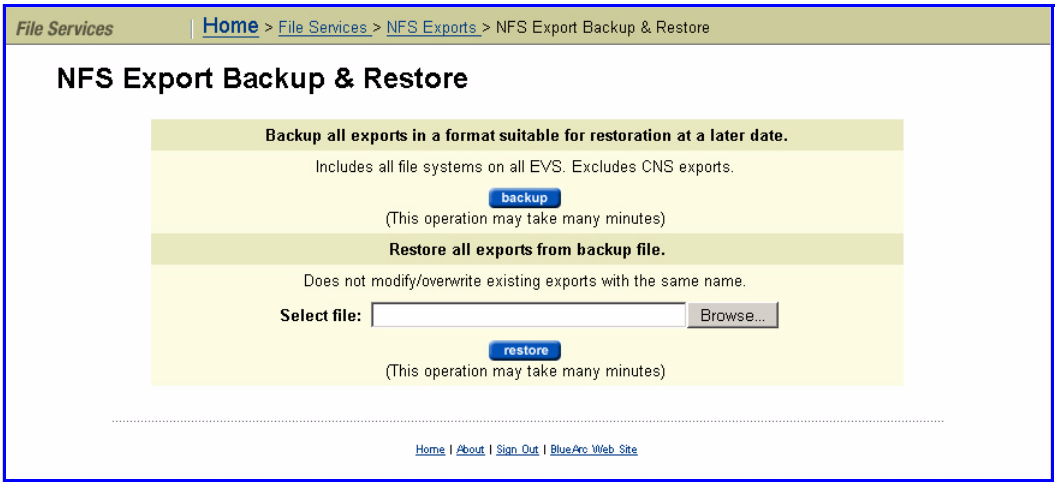

2. **Backup or restore:**

   - **To backup:** Click **backup.** In the browser, specify the name and location of the backup file, then click **OK/Save** (the buttons displayed and the method you use to save the backup file depend on the browser you use).

     A backup file name is suggested, but you can customize it. The suggested file name uses the syntax:

     `NFS_EXPORTS_date_time.txt`, where the following example illustrates the appropriate syntax:
     `NFS_EXPORTS_Aug_4_2006_11_09_22_AM.txt`

   - **To restore:** Click **restore.** In the browser, the backup text file (`NFS_EXPORTS_date_time.txt`) for the specific export(s) you want to restore, then click **Open**.

   When the **NFS Export Backup & Restore** page displays the name and location of the selected file, click **restore.**

Deleting an NFS Export

Caution: *Export Deletion Alert!* Before carrying out the instructions that follow for deleting an export, verify that it is not currently being accessed. If an export is deleted while users are accessing it, their NFS sessions will be terminated and any unsaved data may be lost.

When replacing a storage enclosure, delete all the exports associated with it. Then, when the replacement enclosure is available, add new exports on the new system drives.

Retrieving quota information

To retrieve Quota information, see About the rquotad Service, on page 180.

Using CIFS for Windows Access

Windows networks use the **Common Internet File System** (CIFS) protocol for file sharing between workstations and servers.

CIFS Protocol Support

The server emulates the file-serving functions of Windows NT 4.0, Windows 2000, Windows 2003, and Windows 2008 servers. From the client perspective, the server is indistinguishable from a Windows file server. It provides all of the normal file-serving functions, including:

- Share manipulation (for example: add, list, and delete).
- File manipulation (for example: read, write, create, delete, move, and copy).
- File locking and byte-range locking.
- File access control using standard Windows ACLs.
- File and directory attributes (for example: read-only, and archive).

Prerequisites

To enable CIFS access to the server:

- Enter a CIFS license key.
- Enable the CIFS service.
- Configure the server.

  Depending on the security model used on the CIFS network, configure the server using one of the following methods:

Security Model	Client Authentication	Configuration Method
NT Domain security	NT 4 only	Add server to NT domain
Windows 2000,2003, and 2008 Active Directory	NT 4 only	Add server to NT domain
	Kerberos and NT 4	Join Active Directory

When configured to join an Active Directory, the server functions the same way as a server added to an NT domain, except that after joining an Active Directory, the server can authenticate clients using the Kerberos protocol

as well as NT 4-style authentication. Most modern Windows clients support both authentication methods, though a number of older Windows clients only support NT 4-style authentication.

Supported Clients and Versions

The following table describes supported platforms and versions:

Platform	Supported versions
Windows 2008	SP1
Windows 2003	SP1, SP2, SP3
Windows XP	SP1, SP2, SP3
Windows 2000	SP1, SP2, SP3, SP4
Windows NT 4.0	SP4, SP5, SP6a
Windows 98	SE
Macintosh OS X	10.3 or later (native client only)

Domain Controller Interaction

The storage server relies on Windows domain controllers to authenticate users and to obtain user information (for example, group membership). The server automatically discovers and connects to the fastest and most reliable domain controllers. Since operating conditions may change over time, the server selects the "best" domain controller every 10 minutes.

By default, when authenticating clients in an Active Directory, the server uses the time maintained by the domain controller, automatically adjusting for any clock inconsistencies.

Dynamic DNS

The storage server supports DNS and DDNS. For more information, see DNS and DDNS, on page 76.

CIFS Statistics

CIFS statistics for the storage server (in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Configuring CIFS Security

The server integrates seamlessly into the existing domain and simplifies access control by performing all authentications against the existing domain user accounts.

Note: Only accounts that have been created in the domain or in a trusted domain can access the server.

When a user attempts to access a share, the server verifies appropriate permissions; once access is granted at this level, standard file and directory access permissions apply.

The server operates on a specific domain and can, optionally, join an Active Directory. It interacts with a domain controller (DC) in its domain to validate user credentials. The server supports Kerberos-based authentication to an Active Directory, as well as NTLM authentication (using pre-Windows 2000 protocols). In addition to users *belonging* to its domain, the server allows *members of trusted domains* to connect through CIFS.

The server automatically grants administrator privileges to domain administrators who have been authenticated by the DC. In addition, local administration privileges can be assigned, including backup operator privileges to selected groups (or users).

Assigning CIFS Names

Windows clients access the server through configured CIFS names. Traditional Windows servers have a single host name; in environments where multiple Windows servers are being consolidated, the server can be configured with multiple CIFS names.

In order to appear as a unique server on a Windows network, the server will do the following for each configured CIFS name:

- Allow administration through the Microsoft Server Manager (NT 4) or Computer Management (Windows 2000, 2003, or 2008) administrative tools.
- If NetBIOS is enabled, register each CIFS name with the domain Master Browser so each name appears as a unique server in Network Neighborhood.
- Register each CIFS name with DDNS or WINS for proper host name resolution.
- Support up to 256 CIFS names per EVS.

Joining an Active Directory (AD)

To join an Active Directory, a CIFS server name must be added. For each configured CIFS name, a corresponding computer account must exist in the Active Directory. Computer accounts can be pre-created in the desired folder using the “Active Directory Users and Computers” tool. If no computer account exists, the server will add a corresponding computer account to the “Computers” folder when the CIFS name is added to the server’s configuration.

Note: For security, the Microsoft AD requires that the time difference between the joining computer and the AD is not more than 5 minutes. Verify that the time on the server is configured properly and is in sync with the domain before attempting to join the Directory.

Note: To add a server/cluster to a Microsoft Windows 2008 AD, NetBIOS must be disabled on the server/cluster. After joining the AD, NetBIOS can be enabled. Note, however, that NetBIOS is deprecated in Windows 2008.

To join an Active Directory:

1. **Navigate to the CIFS Setup page.**

   From the **File Services** page, click **CIFS Setup** to display the **CIFS Setup** page:

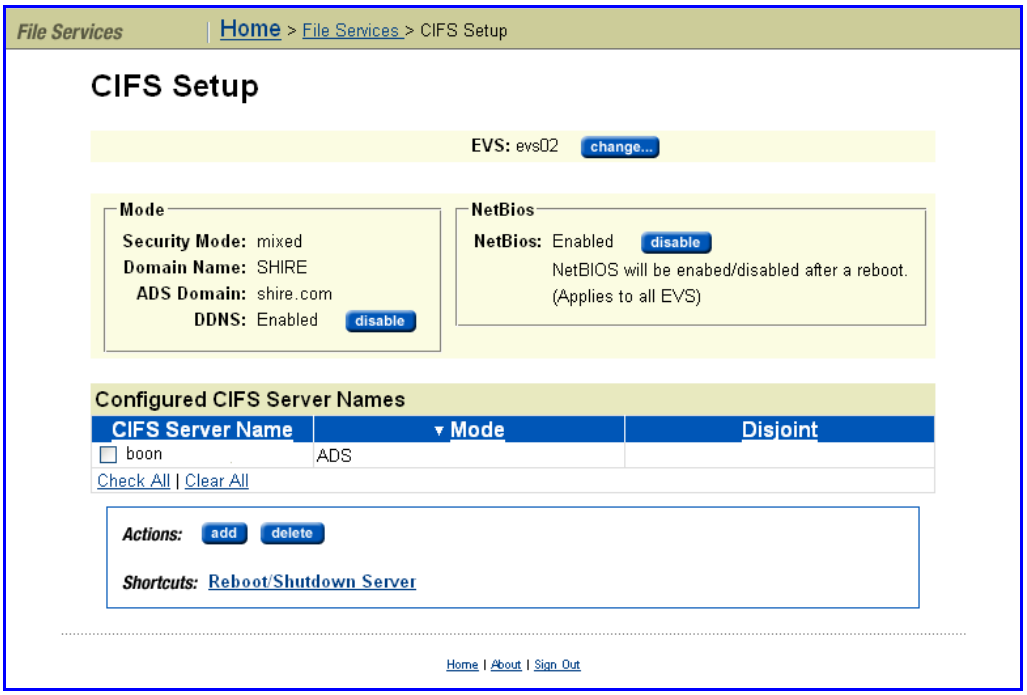

2. **Select a Virtual Server (EVS).**

   From the **EVS** drop-down menu, select an **EVS**.

3. **Navigate to the Add CIFS Server page.**

Click **add** to display the **Add CIFS Server Names** page.

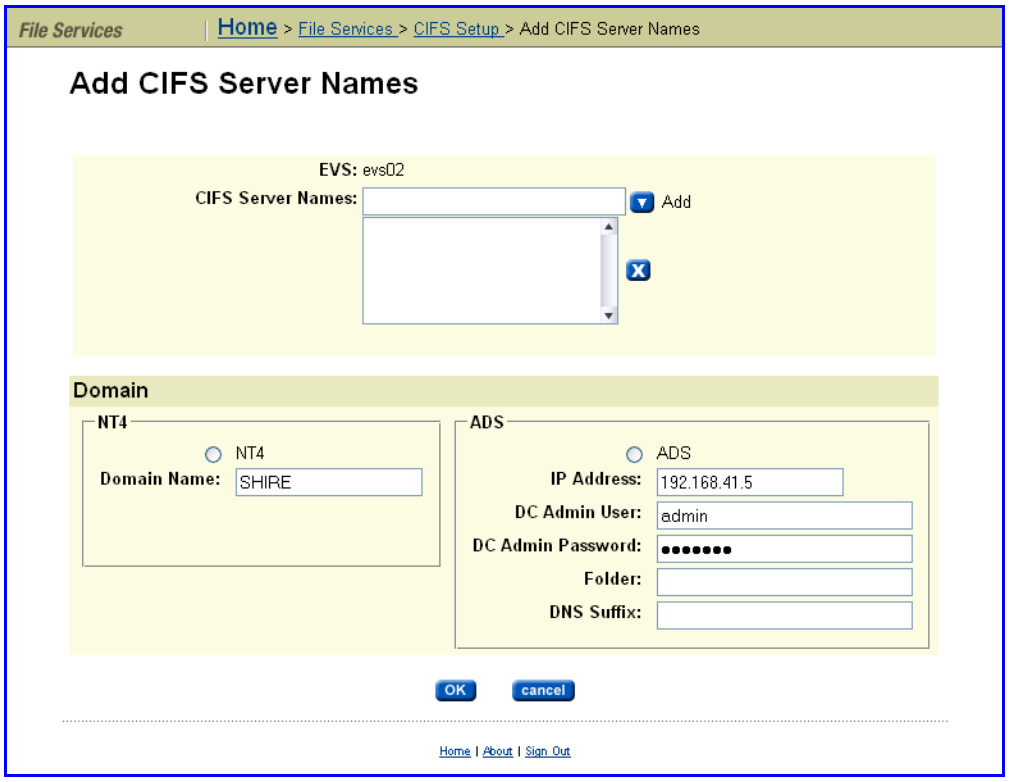

The following table describes the fields in this page:

Item/Field	Description
CIFS Server Name	The computer name through which CIFS clients will access file services on the server. In an ADS domain, the maximum number of characters for the CIFS server name is 63. In an NT4 domain, the maximum number of characters for the CIFS server name is 15.
Domain	Indicates if the CIFS server is to be a part of an NT4 domain or an ADS domain, and allows you to specify the settings required to be a part of the domain.
NT4	Select the NT4 radio button to indicate that the CIFS server is to be a part of an NT4 domain.
Domain Name	The name of the NT4 domain to which you want to add the CIFS server.
ADS	Select the ADS radio button to indicate that the CIFS server is to be a part of an ADS domain.
IP Address	The IP address of a domain controller in the Active Directory in which the server will be configured.

Item/Field	Description
DC Admin User	A user account that is a member of the Domain Administrators group. This privilege is necessary to create a computer account in the Active Directory.  **Note:** When specifying a user account from a trusted domain, the user account must be entered using the *Kerberos* format; that is, `administrator@ADdomain.mycompany.com`, not `ADdomain\administrator`.
DC Admin Password	Correct password for the Domain Administrator user.
Folder	The folder in the Active Directory in which the computer account should be created. By default, the computer account will be created in the Computers folder.
DNS Suffix	Use this option only if you need to set a DNS suffix other than the Active Directory domain's primary DNS suffix. (For example, set this if you have a disjoint domain.)

4. **Specify the CIFS server name.**

   To add a CIFS server, specify a CIFS server name in the **CIFS Server Name** field.

5. **Select the ADS radio button.**

   In the **Domain** area, select the ADS radio button.

   When joining an ADS domain, a CIFS server name must be added. For each configured CIFS name, a corresponding computer account must exist in the Active Directory. Computer accounts can be pre-created in the desired folder using the "Active Directory Users and Computers" tool. If no computer account exists, the server will add a corresponding computer account to the "Computers" folder when the CIFS name is added to the server's configuration.

   Note: For security, the Microsoft AD requires that the time difference between the joining computer and the AD is not more than 5 minutes. Verify that the time on the server is configured properly and is in sync with the domain before attempting to join the Directory.

6. **Specify the ADS domain settings.**

   Using the fields in the ADS domain area, specify the ADS domain configuration settings for the CIFS server you are adding':

   a. **Specify the IP Address of a DC in the ADS domain.**

      In the **IP Address** field, specify the IP Address of a Domain Controller in the Active Directory.

   b. **Specify the DC Admin User name.**

      In the **DC Admin User** field, specify a user account that is a member of the Domain Administrators group. This privilege is necessary to create a computer account in the Active Directory.

Note: When specifying a user account from a trusted domain, the user account must be entered using the *Kerberos* format; that is, `administrator@ADdomain.mycompany.com`, not `ADdomain\administrator`.

c. **Specify the password associated with the DC Admin User name.**

In the **DC Admin Password** field, specify the password associated with the DC Admin User name specified above.

d. **Specify the computer account folder.**

In the **Folder** field, specify the name of the folder in the Active Directory in which the computer account should be created. By default, the computer account will be created in the Computers folder.

7. **Save the configuration.**

Click **OK** to save the configuration, or click **cancel** to return to the **CIFS Setup** page.

Adding a Server to an NT 4 Domain

To enable access to a server in an NT 4 domain, a computer account must be created in the NT domain, and a corresponding NT 4 CIFS name must be created on the server:

1. **Navigate to the CIFS Setup page.**

From the **File Services** page, click **CIFS Setup** to display the **CIFS Setup** page:

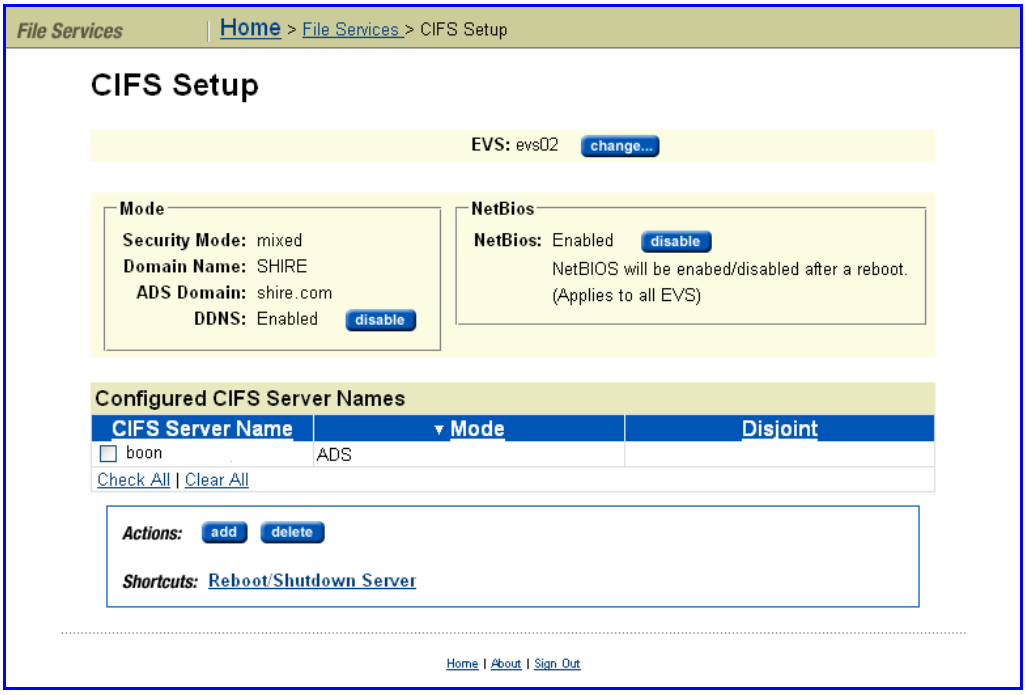

2. **Select a Virtual Server (EVS).**

From the **EVS** drop-down menu, select an **EVS**.

3. **Navigate to the Add CIFS Server Names page.**
   Click **add** to display the **Add CIFS Server Names** page.

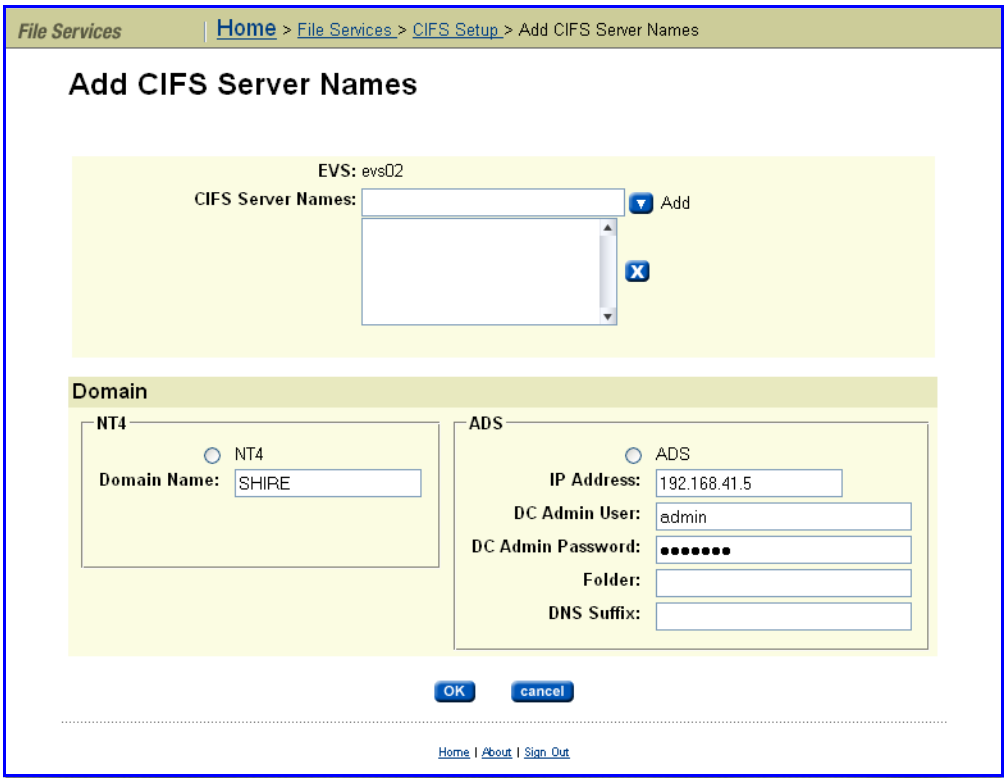

The following table describes the fields in this page:

Item/Field	Description
CIFS Server Name	The computer name through which CIFS clients will access file services on the server (maximum 63 characters).
Domain	Indicates if the CIFS server is to be a part of an NT4 domain or an ADS domain, and allows you to specify the settings required to be a part of the domain.
NT4	Select the NT4 radio button to indicate that the CIFS server is to be a part of an NT4 domain.
Domain Name	The name of the NT4 domain to which you want to add the CIFS server.
ADS	Select the ADS radio button to indicate that the CIFS server is to be a part of an ADS domain.
IP Address	The IP address of a domain controller in the Active Directory in which the server will be configured.

Item/Field	Description
DC Admin User	A user account that is a member of the Domain Administrators group. This privilege is necessary to create a computer account in the Active Directory. **Note:** When specifying a user account from a trusted domain, the user account must be entered using the *Kerberos* format; that is, `administrator@ADdomain.mycompany.com`, not `ADdomain\administrator`.
DC Admin Password	Correct password for the Domain Administrator user.
Folder	The folder in the Active Directory in which the computer account should be created. By default, the computer account will be created in the Computers folder.
DNS Suffix	Use this option only if you need to set a DNS suffix other than the primary Domain Name Service suffix. (For example, set this if you have a disjoint domain.)

4. **Specify the CIFS server name.**

   To add a CIFS server, specify a CIFS server name in the **CIFS Server Name** field.

5. **Select the NT4 radio button.**

   In the **Domain** area, select the NT4 radio button.

6. **Specify the NT4 domain name.**

   In the **Domain Name** field, specify the NT4 domain name. To add the CIFS server to an NT4 domain:

7. **Create an account in the NT4 domain.**

   To create an NT 4 domain account, run **Server Manager** from a domain controller in the NT 4 Domain and create a new "Windows NT Workstation or Server" account using the desired host name.

   Click **Add** to proceed or **Cancel** to decline.

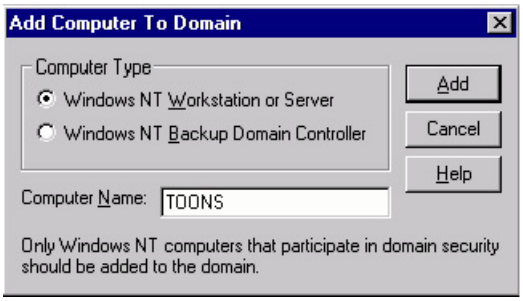

8. **Save the configuration.**

   Click **OK** to save the configuration, or click **cancel** to return to the **CIFS Setup** page.

Removing CIFS Server Names

CIFS server names can be removed from the server's configuration by deleting them from the list of configured CIFS server names. When ADS CIFS names are removed, the corresponding computer account in the Active Directory is also removed. Computer accounts in NT 4 Domains must be deleted manually through Server Manager.

Caution: *CIFS Name Deletion Alert!* At least one CIFS name must be configured on the server to support connections from Windows clients. As a result, if the last configured CIFS name is removed, Windows clients will no longer be able to access the server over CIFS.

Note: DNS entries do not de-register automatically after removing a CIFS server name, so the admin should delete the CIFS server name entry from DNS manually.

Using NetBIOS

When enabled, NetBIOS allows NetBIOS and WINS on this server. If this server communicates by name with computers that use older Windows versions, this setting is required. By default, the server is configured to use NetBIOS.

Disabling NetBIOS has some advantages:

- Simplifies the transport of SMB traffic.
- Removes WINS and NetBIOS broadcast as a means of name resolution.
- Standardizes name resolution on DNS for file sharing.

Removing CIFS Names (Disabling NetBIOS)

Caution: NetBIOS should only be disabled if a reliable DNS infrastructure is in place. Before disabling NetBIOS, verify that there is no need to use NetBIOS, WINS, or legacy NetBT-type applications for this network connection. Once disabled, clients will only be able to communicate with the server by its name through DNS. Disabling NetBIOS can also cause connectivity problems for users of older versions of Windows.

If this server communicates only with computers that run Windows 2000 or later versions of Windows, disabling NetBIOS will be transparent and may result in a performance benefit, as dynamic DNS registration of CIFS names and IP addresses is an easy way to ensure reliable connectivity.

To disable NetBIOS, go to the **File Services** page and click **CIFS Setup** to display the **CIFS Setup** page:

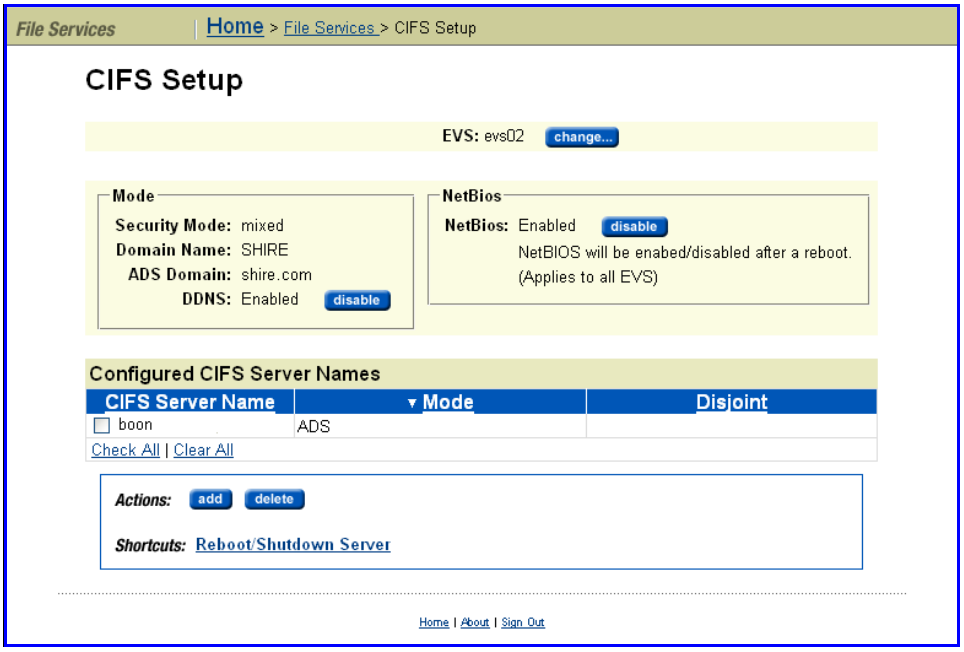

Disable NetBIOS by clicking **disable**. When prompted, reboot the server.

Configuring Local Groups

In a Windows security domain, users and groups identify users (for example, `vsmith`) and groups of users (for example, `software`) on the network. Apart from the user-defined network group names (for example, `software`, `finance`, and `test`), Windows also supports a number of built-in or "local" groups with each providing various privileges and levels of access to the server on which they have been configured.

These groups exist on every Windows computer. They are not network groups, but are local to each computer. So, the user `jamesg` may be granted Administrator privileges on one computer and not on another. Similarly in the server, the administrator can add users to any of these local groups, but only four of them are currently effective:

- **Root.** If a user is a member of the local Root group, the user bypasses all security checks, and can take ownership of any file in the file system.
- **Administrators.** If a user is a member of the local Administrators group, the user can take ownership of any file in the file system.
- **Backup Operators.** If a user is a member of the local Backup Operators group, the user bypasses all security checks, but cannot take ownership of a file in the file system. The privilege to bypass all security checks in the file system is required for accounts that run Backup Exec or perform virus scans. Virus scanner servers that are a part of the Backup Operators group can, however, take ownership of any file in the file system.

- **Forced Groups.** If a user is a member of the local Forced Groups group, when the user creates a file, the user's defined primary group is overridden and the user account will be used to indicate the file creator's name.

Adding a Local Group or Local Group Members

To add a local group:

1. **Navigate to the Local Groups page.**

   From the **File Services** page, click **Local Groups** to display the **Local Groups** page:

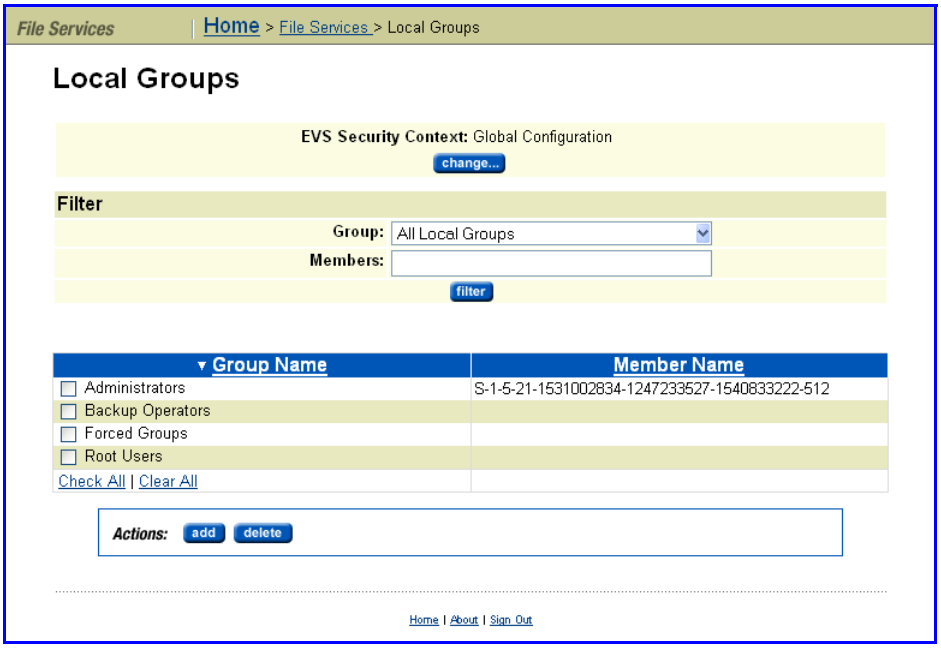

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes made to local groups using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To manage local groups for an EVS that uses an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

   Click **Change** to select a different EVS security context or to select the global configuration.

3. **Add the new local group or member.**

a. **Navigate to the Add a Local Group page.**

Click **add** to display the **Add a Local Group** page:

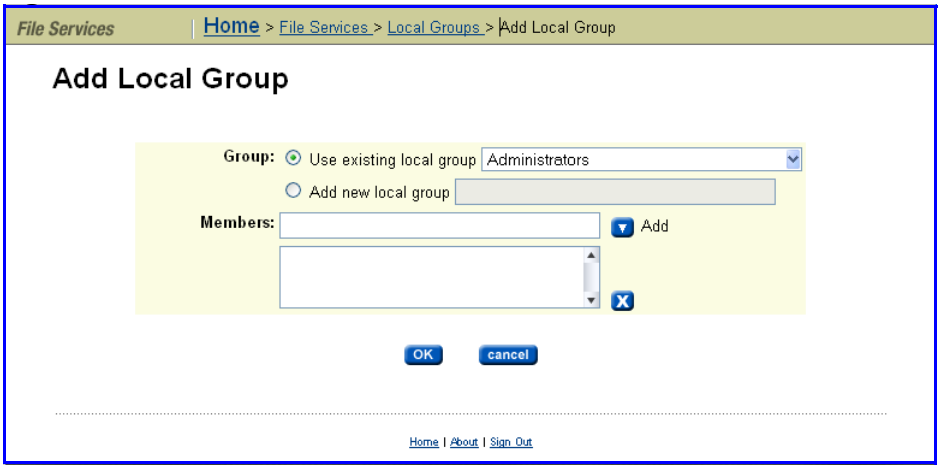

b. **Enter the name for the new group or select the group to which you want to add the user(s).**

- **To add a new local group**, select the **Add new local group** radio button, and enter the name of the new local group in the field.
- **To add a user to an existing local group**, select the **Use existing local group** from the drop down list.

c. **Specify group members.**

Enter the name of a user for the local group in the **Members** field, and click the down arrow to add the user to the membership list. You can repeat this step as often as necessary to add several members at the same time.

To delete a name from the membership list, select the name to delete, and click the **x**.

d. **Save the new local group and/or group members.**

Click **OK** to save the new local group/group members.

Note: Once created, group names may not be changed. To change a group name, you must delete the group, then create a new group, and add members to the new group.

Deleting a Local Group or Local Group Members

To delete a local group:

1. **Navigate to the Local Groups page.**

From the **File Services** page, click **Local Groups** to display the **Local Groups** page:

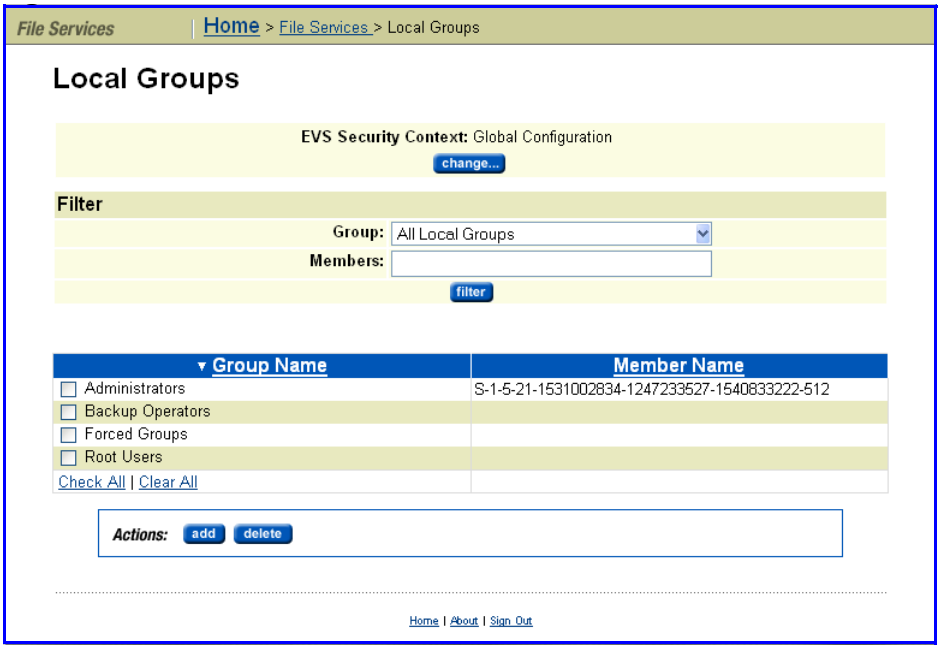

2. **If necessary, change the EVS Security Context.**

   The **EVS Security Context** displays the currently selected EVS security context. Changes made to local groups using this page apply only to the currently selected EVS security context.

   - If an EVS uses the Global configuration, any changes made to the global configuration settings will affect the EVS.
   - If an EVS uses an Individual security context, changes made to the global configuration settings will not affect the EVS. To manage local groups for an EVS that uses an individual security context, you must select the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.

   Click **Change** to select a different EVS security context or to select the global configuration.

3. **Delete all the members of the group.**

   Note: Deleting a local group is a two-stage process; you must delete all members of the group before you can delete the group itself.

   a. **Select the group members to delete.**

      Fill the checkbox next to all members of the group you want to delete.

b. **Delete the selected group members.**

Click **delete** to delete the selected group members.

A confirmation dialog will appear. Click **OK** to delete the users and return to the **Local Groups** page, or **Cancel** to return to the **Local Groups** page without deleting the users.

4. **Delete the local group.**

a. **Select the group to delete.**

Fill the checkbox next to the group you want to delete.

b. **Delete the group.**

Click **delete** to delete the selected group.

A confirmation dialog will appear. Click **OK** to delete the group and return to the **Local Groups** page, or **Cancel** to return to the **Local Groups** page without deleting the group.

Configuring CIFS Shares

CIFS shares can be set up on mounted volumes. The server can support more than 1,000 shares. However, the exact limit of any share allocation depends on the overall size of the server's configuration.

To add or delete a CIFS share:

1. **Navigate to the CIFS Shares page.**

From the **File Services** page, click **CIFS Shares** to display the page.

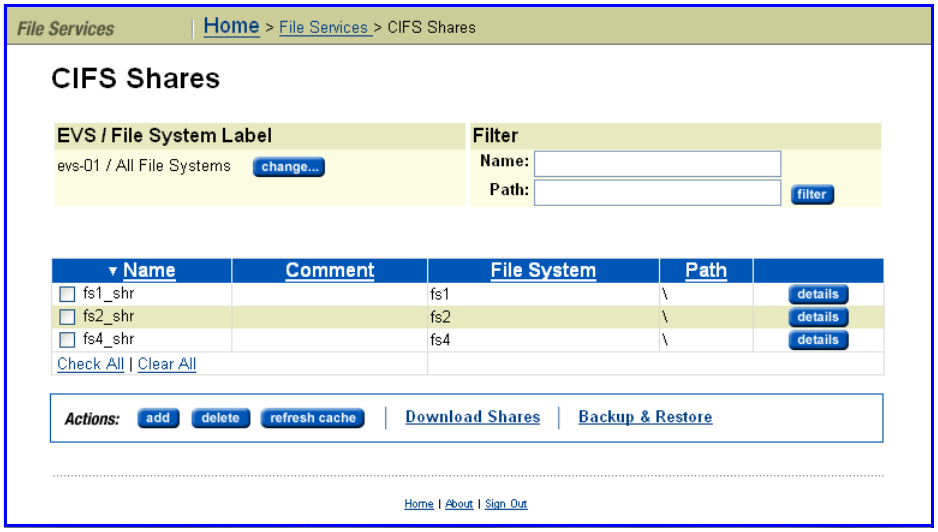

The following table describes the fields in this page:

Item/Field	Description
Cluster Name Space/ CNS Label or EVS/File System Label	Name of the currently selected context (either the cluster name space context or an EVS local context) for which shares are displayed. • When the Cluster Name Space context is displayed, the list shows shares that link to the CNS tree. Users accessing those shares can go to the directories at or below that point in the CNS tree. • When a local EVS context and a File System label are displayed, the list shows shares for file systems in the selected EVS. To change the context being viewed, click **change**.
Filter	Filters can be defined, based on share name or path, to constrain the list of shares displayed on the page.
Name	Name of the CIFS share.
Comment	Additional information associated with the CIFS share. This information is often displayed to clients along with the share name.
File System/Name Space	Name of the file system or CNS link on which the share is located.
Path	The directory to which the CIFS share points. Users accessing the share will be able to access this directory, and any directories under it in the directory tree.

The following **Actions** are available:

- Click **details** for a CIFS Share to display its details.
- Click **add** to add a CIFS share and proceed to the next step.
- Click **delete** to delete a selected CIFS share.
- Click **refresh cache** to clear the SMU's cache and repopulate it with the relevant objects. Note that this is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.

The following **Shortcuts** are also available:

- Click **Download Shares** to download a comma separated value (.csv) file containing a list of all configured CIFS shares on the selected EVS and file system. Note that the downloaded file cannot be used to restore CIFS shares (you must restore CIFS shares from an CIFS shares backup file). To download a list of shares from another file system, click **Change.**
- Click **Backup & Restore** to display the **CIFS Shares Backup & Restore** page. For more information, see Backing Up and Restoring CIFS Shares, on page 279.

2. **Add a CIFS Share.**

   Previously, you clicked **add** to display the **Add Share** page:

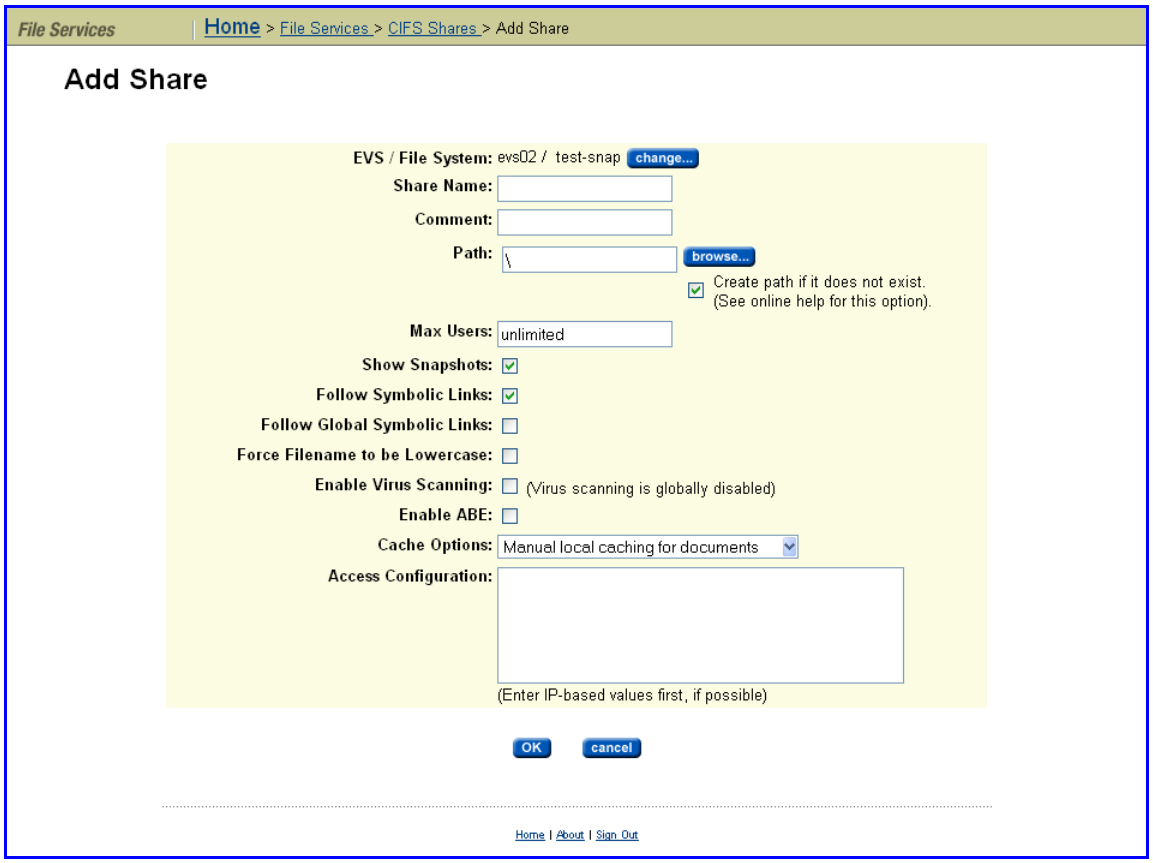

Now enter the requested information. The following table describes the fields in this page:

Item/Field	Description
Cluster Name Space/ CNS Label or EVS/File System Label	Name of the currently selected context (either the cluster name space context or an EVS local context) for which shares are displayed. • When the Cluster Name Space context is displayed, the list shows shares that link to the CNS tree. Users accessing those shares can go to the directories at or below that point in the CNS tree. • When a local EVS context and a File System label are displayed, the list shows shares for file systems in the selected EVS. To change the context being viewed, click **change**.
Name	Name of the CIFS share.
Comment	Enter additional information regarding the CIFS share. This information is often displayed to clients along with the share name.

Item/Field	Description
Path	The directory to which the CIFS share points. Users accessing the share will be able to access this directory, and any directories under it in the directory tree. To find a directory, click **browse.** Fill the **Create path if it does not exist** checkbox to create the path if it does not already exist. **Note:** The browse button only exists if the path being created is the path in a file system, not a namespace.
Max Users	The maximum number of users who can be associated with the CIFS share. The default is unlimited.
Show snapshots	Fill this checkbox to allow snapshot access from this share.
Follow Symbolic Links	Fill this checkbox to enable the following of symlinks for this share. **Note:** CIFS2 (SMB2) clients are not able to follow symlinks (relative or absolute) to files on storage accessed through the Titan Server.
Follow Global Symbolic Links	Fill this checkbox to enable CIFS clients to follow global (absolute) symlinks via Microsoft's DFS mechanism for this share. For more information on symbolic links, see Symbolic Links, on page 228.
Force Filename to be Lowercase	Fill this checkbox to force all filenames generated on this share to be lowercase. This is useful for interoperability of UNIX applications.
Enable Virus Scanning	If virus scanning is enabled and configured for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is enabled by default. If virus scanning is not enabled for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is not enabled by default, but you can enable it a per-EVS basis. Note that virus scanning is set up on a per-EVS basis, or for all EVSs using the global configuration context, but cannot be set up on a per-server or per-cluster basis. To disable virus scanning for a specific share, clear this checkbox.
Enable ABE	To enable ABE (access based enumeration) for this share, fill the check box. By default, ABE is disabled for shares and on the server/cluster as a whole. Before enabling ABE for a share, you must make sure ABE is enabled for the server/cluster as a whole (the CLI command to enable ABE support is `fsm set disable-ABE-support false`). When enabled, ABE filters the contents of a CIFS share, so that only the files and directories for which a user has "FileReadData" or "FileListDirectory" rights are visible to the user (for example returned in a directory listing or considered by a wildcarded delete). Note that enabling ABE may have a negative impact on CIFS performance.
Cache Options	Select the Cache options for the share: • Manual Local Caching for Documents • Automatic Local Caching for Documents • Automatic Local Caching for Programs • Local Caching Disabled

<table>
<tr><th>Item/Field</th><th>Description</th></tr>
<tr><td>Access Configuration</td><td>IP addresses of the clients who can access the share (up to 2,000 characters allowed in this field). For more information on specifying access configuration, see Controlling Access to Shares using Qualifiers, on page 275.
<table>
<tr><th>What to type</th><th>Means</th></tr>
<tr><td>Blank or *</td><td>All clients can access the share.</td></tr>
<tr><td>Specific addresses.
Example: 10.168.20.2</td><td>Only clients with the specified IP address can access the share</td></tr>
<tr><td>A range of addresses using CIDR notation.
Example: 10.168.20.0/16</td><td>Clients with addresses within the range can access the export.</td></tr>
<tr><td>Partial addresses using wildcards.
Example: 10.168.*.*</td><td>Clients with matching addresses can access the share.</td></tr>
</table></td></tr>
</table>

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

   Note: By default, when a CIFS share is created, the group “Everyone” is added to the share permissions list. For information on modifying share permissions, see Adding or Changing CIFS Share Access Permissions, on page 277.

Viewing and Modifying the Properties of a CIFS Share

To view and modify the properties of a CIFS share:

1. **Navigate to the CIFS Shares page.**

   From the **File Services** page, click **CIFS Shares** to display the page:

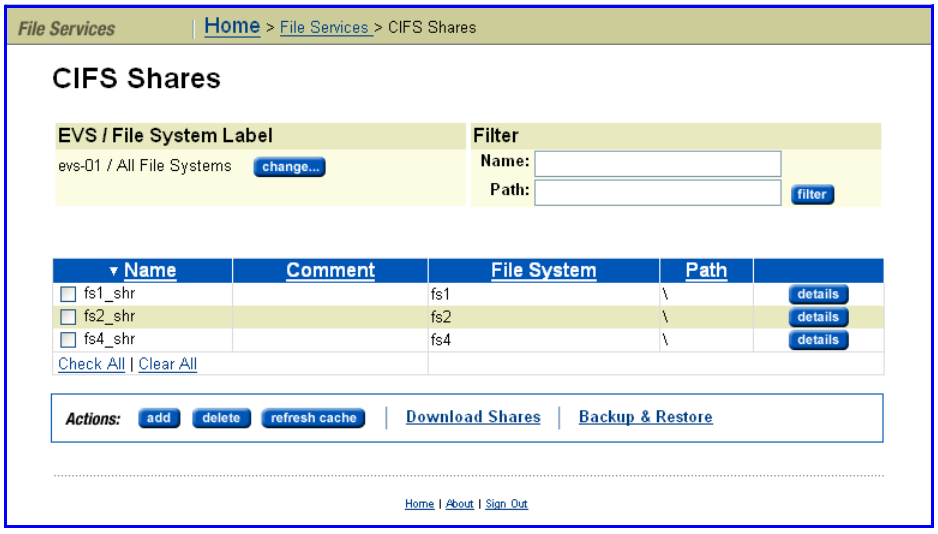

2. **Display details, and (optionally) modify, a CIFS Share.**

Click **details** for a CIFS Share to display its **CIFS Share Details** page:

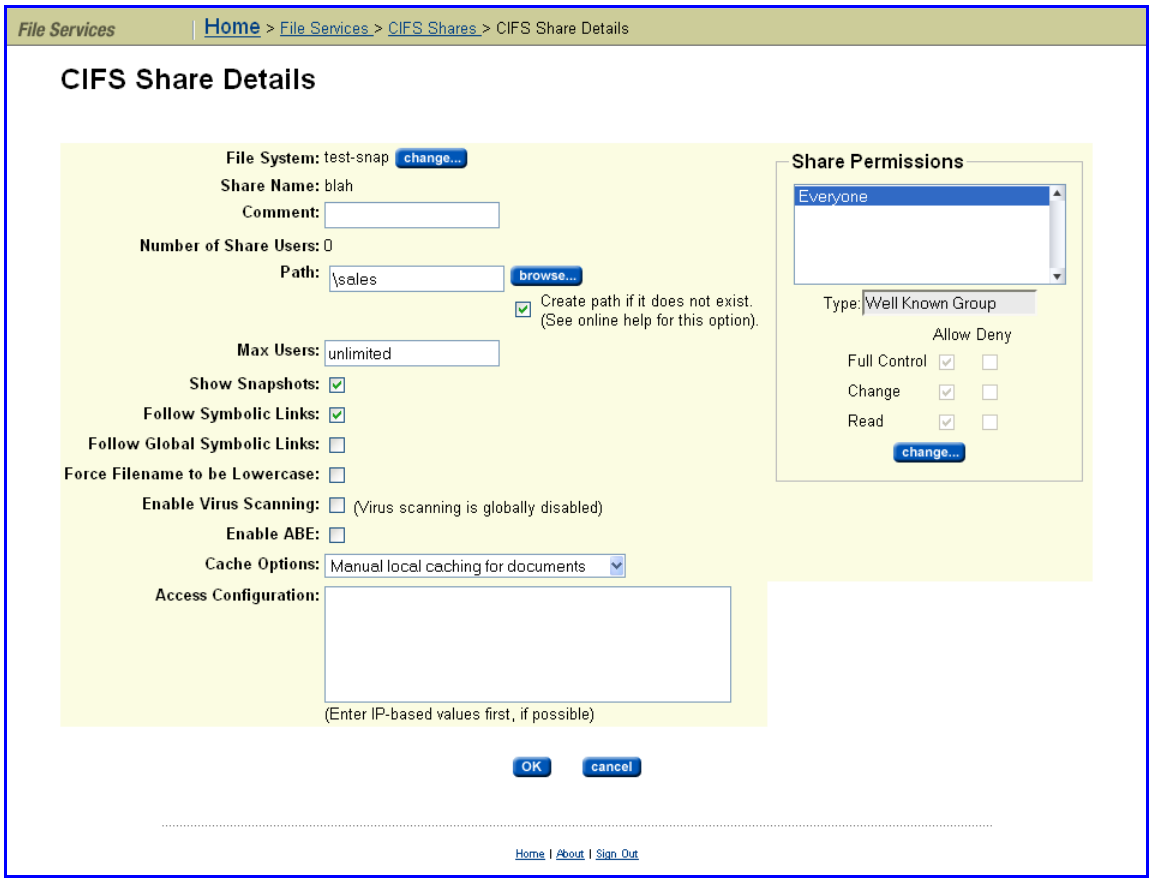

Modify the fields in this page as needed.

The following table describes the fields in this page:

Item/Field	Description
EVS/File System	Displays the currently selected EVS and a file system (or CNS link to a file system) to which the CIFS Share will link. To add the CIFS Share in a different EVS or file system, click **change.**
Share Name	Name of the CIFS share.
Comment	Enter additional information regarding the CIFS share. This information is often displayed to clients along with the share name.
Number of Share Users	Displays the number of users currently connected to the share.

Item/Field	Description
Path	The directory to which the CIFS share points. Users accessing the share will be able to access this directory, and any directories under it in the directory tree. If the share resides in a file system, click **browse** to help find an existing directory in the file system. If the share resides in a Cluster Name Space, click **change** to help find an existing directory. Fill the **Create path if it does not exist** checkbox to create the path if it does not already exist.
Max Users	The maximum number of users who can be associated with the CIFS share. The default is unlimited.
Show snapshots	Fill this checkbox to allow snapshot access from this share.
Follow Symbolic Links	Fill this checkbox to enable the following of symlinks for this share.  **Note:** CIFS2 (SMB2) clients are not able to follow symlinks (relative or absolute) to files on storage accessed through the Titan Server.
Follow Global Symbolic Links	Fill this checkbox to enable CIFS clients to follow global (absolute) symlinks via Microsoft's DFS mechanism for this share. For more information on symbolic links, see Symbolic Links, on page 228.  **Note:** CIFS2 (SMB2) clients are not able to follow symlinks (relative or absolute) to files on storage accessed through the Titan Server.
Force Filename to be Lowercase	Fill this checkbox to force all filenames generated on this share to be lowercase. This is useful for interoperability of UNIX applications.
Enable Virus Scanning	If virus scanning is enabled and configured for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is enabled by default. If virus scanning is not enabled for the global context or for the EVS hosting the file system pointed to by the share then, when the share is created, virus scanning is not enabled by default, but you can enable it a per-EVS basis. Note that virus scanning is set up on a per-EVS basis, or for all EVSs using the global configuration context, but cannot be set up on a per-server or per-cluster basis. To disable virus scanning for a specific share, clear this checkbox.
Enable ABE	To enable ABE (access based enumeration) for this share, fill the check box. By default, ABE is disabled for shares and on the server/cluster as a whole. Before enabling ABE for a share, you must make sure ABE is enabled for the server/cluster as a whole (the CLI command to enable ABE support is `fsm set disable-ABE-support false`). When enabled, ABE filters the contents of a CIFS share, so that only the files and directories for which a user has "FileReadData" or "FileListDirectory" rights are visible to the user (for example returned in a directory listing or considered by a wildcarded delete). Note that enabling ABE may have a negative impact on CIFS performance.

Item/Field	Description
Cache Options	Select the Cache options for the share: • Manual Local Caching for Documents • Automatic Local Caching for Documents • Automatic Local Caching for Programs • Local Caching Disabled
Access Configuration	IP addresses of the clients who can access the share (up to 2,000 characters allowed in this field). For more information on specifying access configuration, see Controlling Access to Shares using Qualifiers, on page 275.
Share Permissions	A display area, showing the share access permissions for users and groups. Click **change** to display the **CIFS Share Permissions** page, where you can change the user/group permissions. For information on the user/group settings see Controlling Access to Shares using Permissions, on page 276.

What to type	Means
Blank or *	All clients can access the share.
Specific addresses. Example: **10.168.20.2**	Only clients with the specified IP address can access the share
A range of addresses using CIDR notation. Example: **10.168.20.0/16**	Clients with addresses within the range can access the export.
Partial addresses using wildcards. Example: **10.168.*.***	Clients with matching addresses can access the share.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Note: By default, the group "Everyone" is included in the share permissions list. For information on modifying share permissions, see Adding or Changing CIFS Share Access Permissions, on page 277

Controlling Access to Shares using Qualifiers

To specify which clients have access to a CIFS share, qualifiers can be appended to the IP address(es):

Qualifier	Description
read_write, readwrite, rw	Grants read/write access. This is the default setting.
read_only, readonly, ro	Grants the specified client read-only access to the CIFS share.
no_access, noaccess	Denies the specified client access to the CIFS share.

Some CIFS share qualifier examples are:

- `10.1.2.38(ro)`

  Grants read-only access to the client with an IP address of 10.1.2.38.

- `10.1.2.0/24(ro)`

  Grants read-only access to all clients whose IP address is within the range 10.1.2.0 to 10.1.2.255.

- `10.1.*.*(readonly)`

  Grants read-only access to all clients with an IP address beginning with 10.1.

The order in which the entries are specified is important. For example,

```
*(ro)
10.1.2.38(noaccess)
```

where the first line grants read-only access to all clients, and the second denies access to the specified client. However, the second line is redundant, as the first line matches all clients. These lines must be transposed to ensure access is denied to 10.1.2.38.

Controlling Access to Shares using Permissions

Access to shares is restricted through a combination of share-level and file-level permissions. These permissions determine the extent to which users can view and modify the contents of the shared directory. When users request access to a share, their share-level permissions are checked first; if authorized to access the share, their file-level permissions are checked.

When the share-level permissions differ from the file-level permissions, the more restrictive permissions apply, as described in the following table, where [a] = “allowed” and [d] = “denied”:

Activity	Read	Change	Full
View the names of files and subdirectories	a	a	a
Change to subdirectories of the shared directory	a	a	a
View data in files	a	a	a
Run applications	a	a	a
Add files and subdirectories	d	a	a
Change data in files	d	a	a
Delete files and subdirectories	d	a	a
Change permissions on files or subdirectories	d	d	a

Activity	Read	Change	Full
Take ownership of files or subdirectories	d	d	a

When configuring access to a share, it is only possible to add users or groups that are:

- Known to domain controllers, and
- Seen by the server on the network.

Note: When a user is given access to a share, if the user has also a member of a group with a different access level, the more permissive level applies. For example, if a user is given *Read* access to a share, and that user also belongs to a group that has *Change* access to that same share, the user will have *Change* access to the share, because *Change* access is more permissive than *Read* access.

Adding or Changing CIFS Share Access Permissions

To add or change the access permissions for a CIFS Share:

1. **Navigate to the CIFS Shares page.**

   From the **File Services** page, click **CIFS Shares** to display the **CIFS Shares** page.

2. **Navigate to the CIFS Shares Details page.**

   Select the share for which you want to add/change access permissions, and click **details**.

3. **Navigate to the CIFS Share Permissions page.**

   In the **Share Permissions** area of the CIFS Share Details page, click **change**.

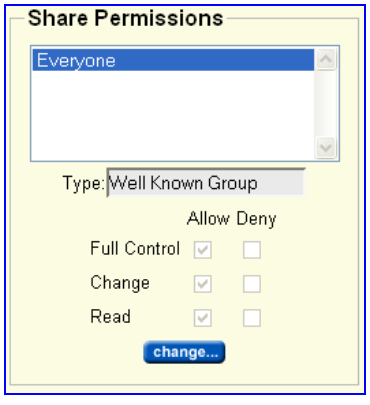

The **CIFS Share Permissions** page appears.

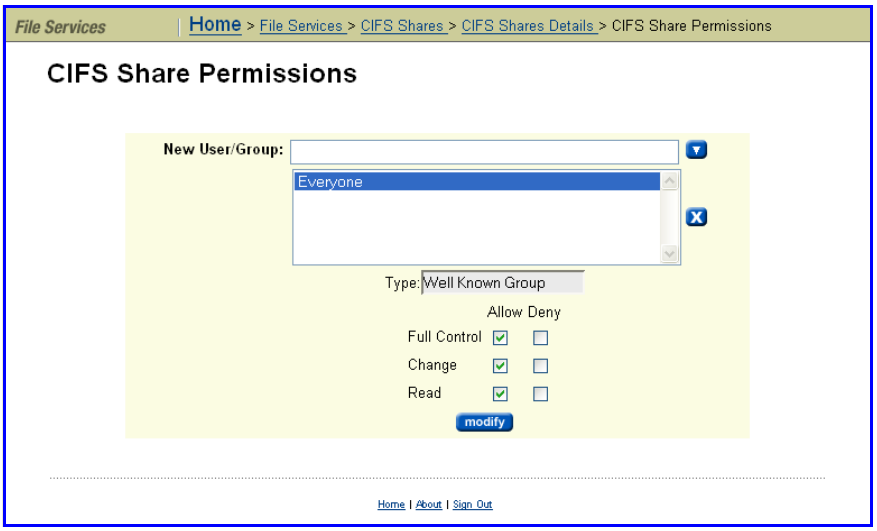

1. **Add a new user/group or select an existing user/group.**
   - To add a new user or group:
     i Enter the name for the new user or group in the **New User/Group** field.
     ii Add the new user/group to the user/group list by clicking the down arrow button following the **New User/Group** field.
     The Type field displays a standardized identifier for the security group to which user/group being added belongs. The value is set automatically, based on the well known security identifier for the user/group being added.
     iii Select the new user/group, by clicking on the user/group name in the user/group list.
   - To select an existing user/group, click on the user/group name in the user/group list.
2. **Specify permissions for the selected user/group.**
   Set the permissions for the user/ group by filling the **Allow** or the **Deny** checkboxes according to the guidelines above (see the table in Controlling Access to Shares using Permissions, on page 276).

Offline File Access Modes

The server supports Offline Files Access. This allows network clients to cache files that are commonly used from a network/file share. To use Offline Files, the client computer must be running Windows 2000 (or later). There are three different share caching modes (supporting all three modes of caching):

- **No Caching:** No caching of files or folders occurs.
- **Manual:** Allows user specification of individual files required for offline access. This operation guarantees a user can obtain access to the specified files whether online or offline.

- **Automatic:** Applies to the entire share. When a user accesses any file in this share, that file becomes available to that user for offline access. This operation does not guarantee a user can obtain access to the specified files, because only files that have been used at least once are cached. The Automatic mode can be defined for documents or programs.

Modifying or Deleting a Share

Before modifying or deleting a share, go to the **CIFS Shares Details** page for the share, and verify that the number of share users is zero.

Caution: If a share is modified or deleted while other users are accessing it, their CIFS sessions will be terminated and any unsaved data may be lost.

Note: If you are replacing storage, and need to modify shares to point to the new file system location, use the file system replication and transfer of primary access features, which will move the file systems and modify the shares automatically. See Data Replication, on page 333 and Transfer of Primary Access, on page 337 for more information on these features.

Backing Up and Restoring CIFS Shares

When backing up and restoring CIFS Shares:

- All CIFS Shares in all EVSs are backed up (except those in the CNS tree).
- A CIFS Share backup file is saved as a .txt file. The backup file contains the file system name and the share name, as well as most of the information about the share, including the settings for: *Ensure Path Exists, Show Snapshots, Follow Symbolic Links, Force Filename to Lowercase, Virus Scanning, Cache Options,* and *Max Users.*

When you restore CIFS Shares from a backup file:

- The restore operation does not modify or overwrite currently existing shares that have the same name.
- With the exception noted above, all shares in the selected backup file are restored.

To backup or restore CIFS shares:

1. **Navigate to the CIFS Shares Backup & Restore page.**

From the **File Services** page, select **CIFS Shares**, then click **Backup & Recovery** to display the **CIFS Shares Backup & Restore** page:

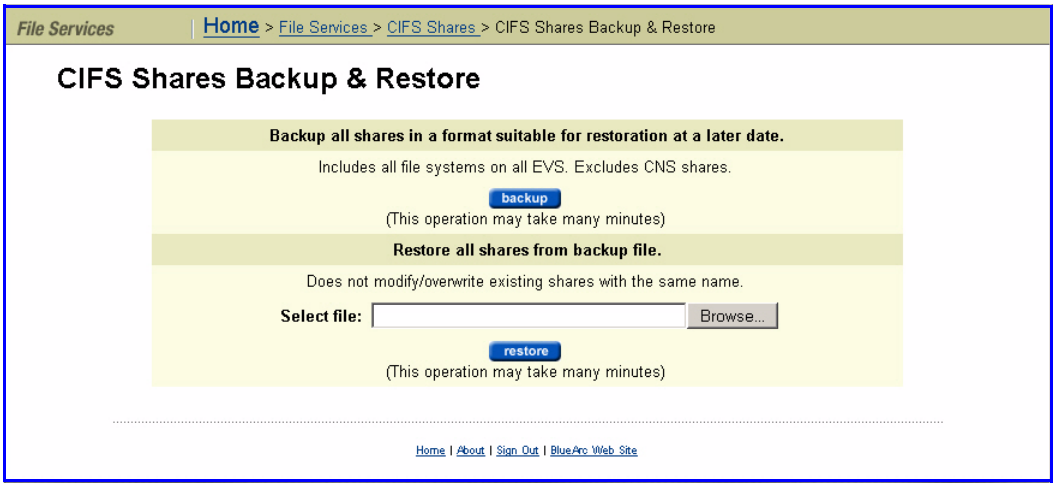

2. **Backup or restore.**
   - **To back up:** Click **backup.** In the browser, specify the name and location of the backup file, then click **OK/Save** (the buttons displayed and the method you use to save the backup file depend on the browser you use).

     A backup file name is suggested, but you can customize it. The suggested file name uses the syntax:

     `CIFS_SHARES_date_time.txt`, where the following example illustrates the appropriate syntax:
     `CIFS_SHARES_Aug_4_2006_11_09_22_AM.txt`
   - **To restore:** Click **restore.**In the browser, the backup text file (`CIFS_SHARES_date_time.txt`) for the specific share(s) you want to restore, then click **Open**. When the **CIFS Export Backup & Restore** page displays the name and location of the selected file, click **Restore**.

Using Windows Server Management

The Computer Management MMC tool, available for Windows 2000 or later, can perform share management tasks from any remote computer; for example:

- Viewing a list of all users currently connected to the system.
- Creating shares.
- Listing all shares on the system and the users connected to them.
- Disconnecting one or all of the users connected to the system or to a specific share.
- Closing one or all of the shared resources that are currently open.
- Viewing an event log.

Note: For older versions of Windows, the equivalent of this tool is provided by Server Manager.

Using the Computer Management Tool

To use the Computer Management tool:

1. **Open the Computer Management tool.**

   In the Windows interface, from **Administrative Services**, select **Computer Management**; then right-click on **Computer Management (Local)** to display a context menu, and select **Connect to another computer**:

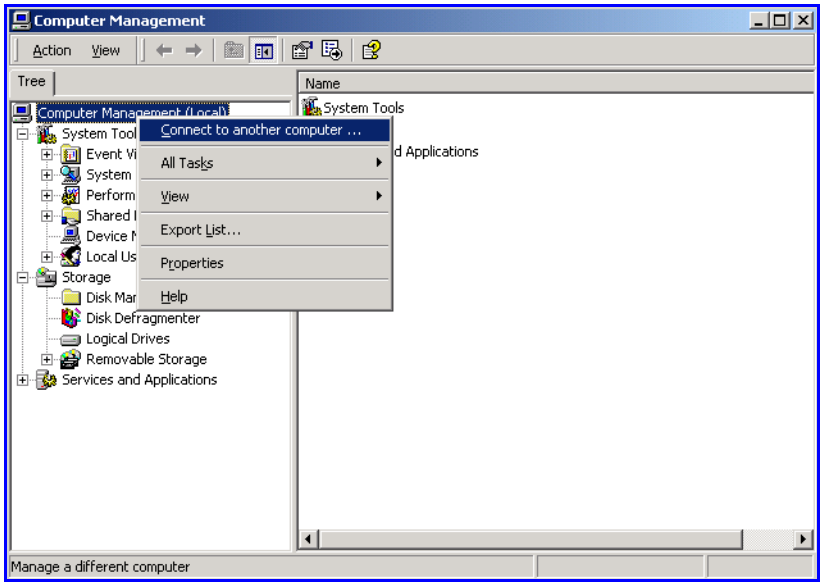

2. **Optionally, specify the domain; then select a name.**

   Select the domain from the drop-down **Look in** field,

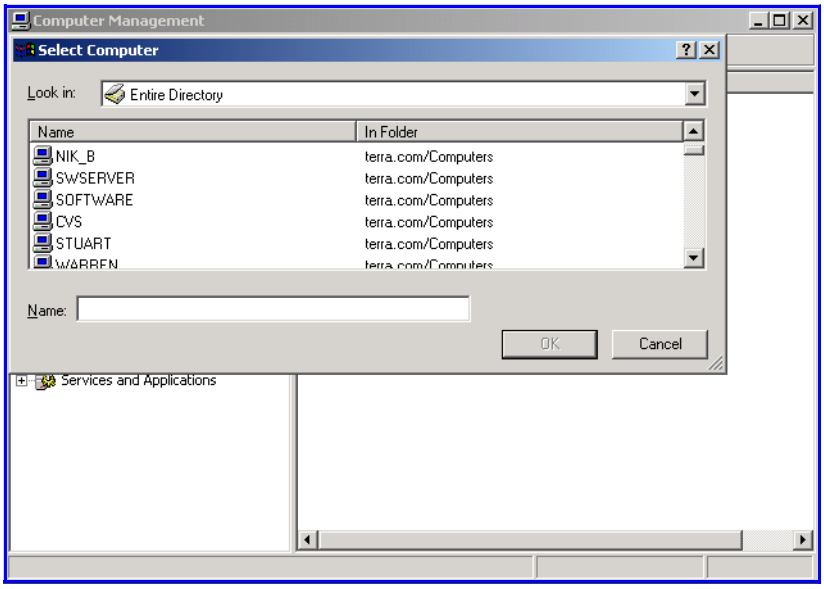

then click to highlight a name or an IP address to use for file services on the server, then click **OK.**

Note: Do not specify a server administration name or IP address for this purpose.

3. **View available information.**

   The following information is available:

   - Click **Event Viewer** to display the server's event log:

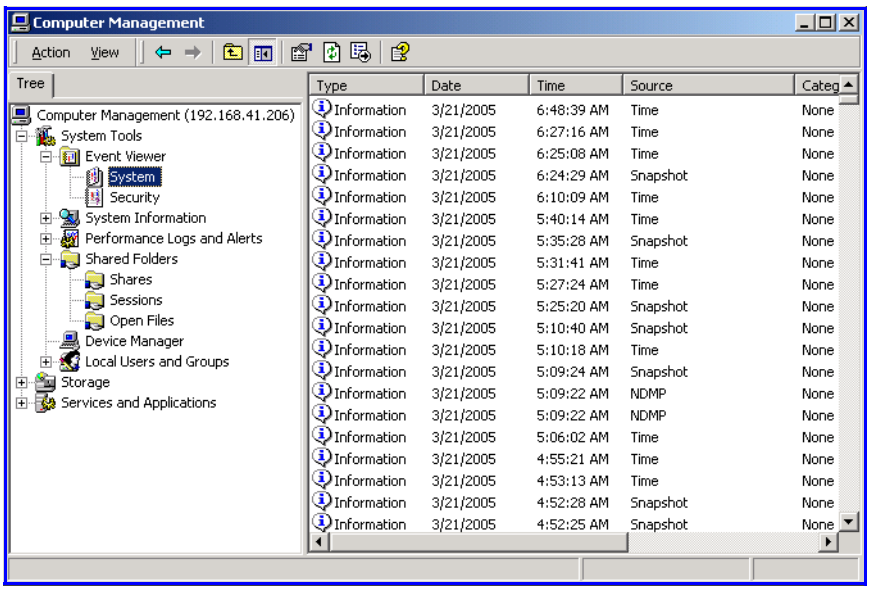

     On the event log window:

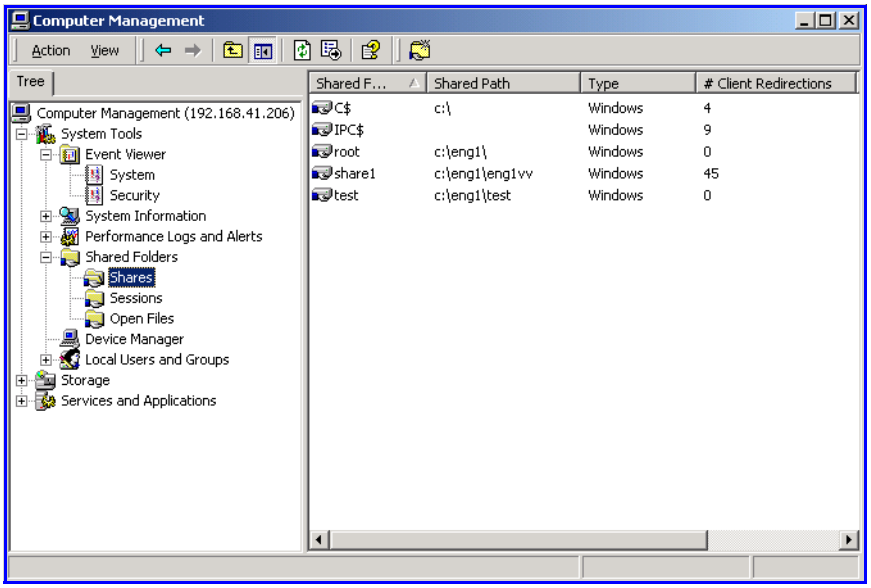

   - Click **Shares** to list all of the shares. Some or all of the users can be disconnected from specific shares.

- Click **Sessions** to list all users currently connected to the system. Some or all of the users can be disconnected.

  Click **Open Files** to list all the open shared resources. Some or all of the shared resources can be closed.

Creating or Managing Shares

When adding a share, enter the name of the folder to share in the text box provided. The format for this text box is, `c:\volname\pathname`, where `volname` is the volume on the server and `pathname` is the path on the volume to be shared. For example,

`c:\volname`

`c:\volname\pathname`

Paths must exist before they can be shared.

Note: Trailing backslashes are not permitted at the end of the manually entered string.

Transferring Files with FTP

This section explains how to set up File Transfer Protocol (FTP) so that users with FTP clients can access files and directories on the storage server.

FTP Protocol Support

The server implements the file-serving functions of an FTP server. The server provides the file-serving functions required for:

- File manipulation.
- Directory manipulation.
- File access control (for example, permissions).

Prerequisites

Prior to allowing FTP access to the system, the FTP service must be enabled. No license key is required for this protocol.

FTP Statistics

FTP statistics for the storage server (in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Configuring FTP Preferences

As part of the process of setting up FTP, choose a service for authenticating the passwords of the FTP users. Also, a timeout must be set with which to end FTP sessions that have been inactive.

To Configure FTP Preferences

1. **Navigate to the FTP Configuration page.**

   From the **File Services** page, click **FTP Configuration** to display the page:

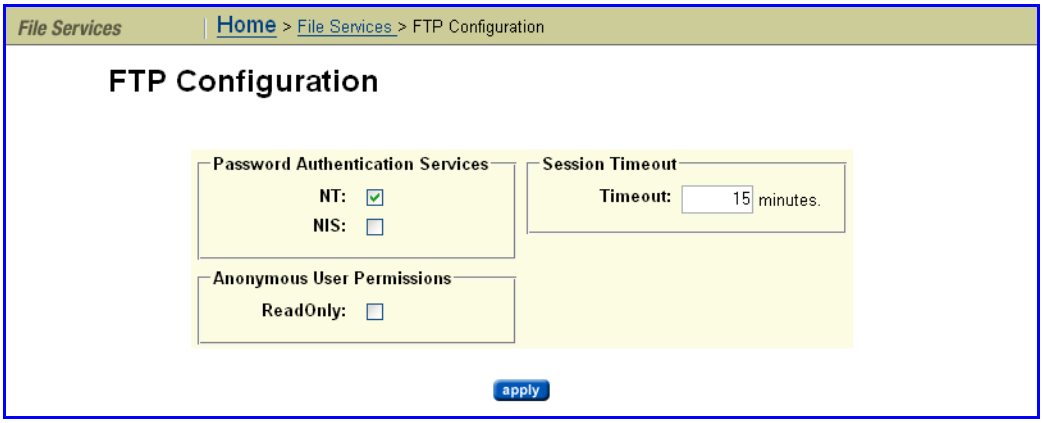

2. **Select a Password Authentication Service, a Session Timeout, and Anonymous User Permissions:**

   - **Password Authentication Service:** Used to authenticate FTP users. Select *NT* or *NIS*.

     The configured security mode determines what password authentication service methods can be used. For more information about security modes, see File System Security, on page 220.

     If operating in UNIX or Mixed security mode, both NT and NIS password authentication are supported. If both services are enabled, the FTP user will be authenticated against the configured NT domain first. If authentication fails, the server will attempt to authenticate the user against the configured NIS domain.

   - **Session Timeout:** The number of minutes of inactivity after which to end an FTP session automatically. The value must be at least 15 minutes.

   - **Anonymous User Permissions:** Specifies whether read-write is allowed for anonymous requests. Read-write permissions are the default, but you can fill the **ReadOnly** checkbox to limit anonymous requests to read only.

3. **Save your settings.**

   Verify your settings, then click **apply** to save.

Configuring FTP Users

FTP users can be manually set up or their details can be imported from a file.

Setting up an FTP User

To set up an FTP user:

1. **Navigate to the FTP Users page.**

From the **File Services** page, click **FTP Users** to display the page:

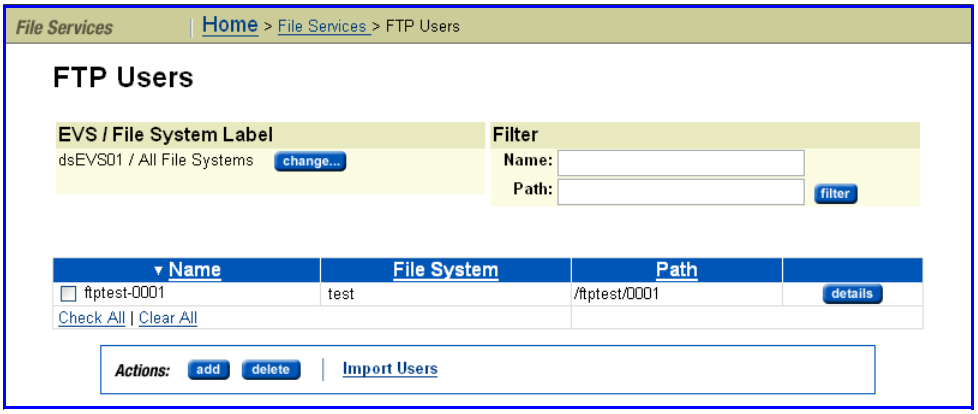

The following table describes the fields in this page:

Item/Field	Description
EVS / File System Label	This field displays the EVS and File System where the FTP users listed on the page have been configured. Click the **change** button to select a different file system.
Filter	The **filter** button allows you to filter the users based on user **Name** or **Path**.
Name	This column displays the existing FTP users. Up to 500 users can be listed, but the FTP user list displays a maximum of 20 users per page. Use the **filter** criteria to control the display of users.
File System	This column displays the label for the file system containing the user's initial path.
Path	Path of the directory in which the selected FTP user starts when logged in over FTP.
details	Opens the FTP User Details page, allowing you to modify certain details about selected user.
add	Opens the Add User page, allowing you to set up a new user.
delete	Deletes the selected user. To select a user, fill the checkbox next to the user **Name**.
Import Users	Opens the Import FTP Users page, allowing you to set up new users by importing them from a file. See Importing an FTP User, on page 286 for more information on importing FTP users.

2. **Add an FTP user.**

   Click **add** to display the **Add User** page.

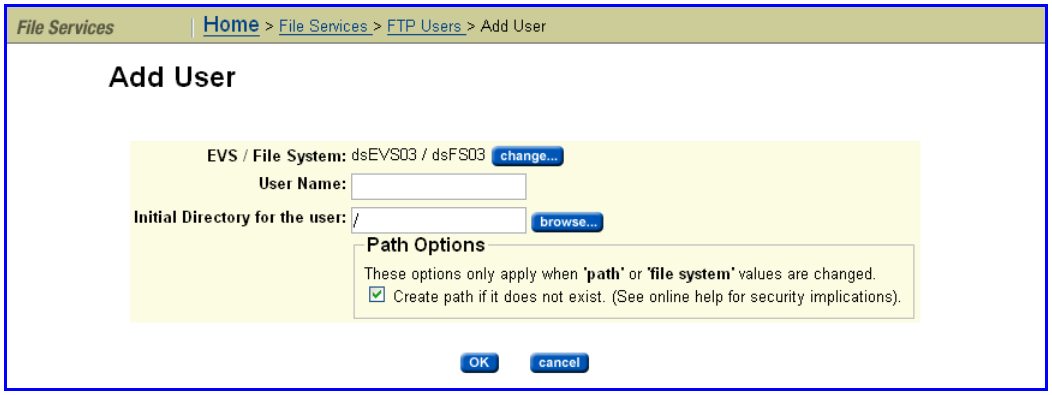

Provide the appropriate information:

- The **EVS / File System** field displays the selected file system. To change the file system, click the change button to open the Select a File System page.
- In the **User Name** field, type the name with which the user is to log in. To allow anonymous logins to the initial directory, specify the user name as "anonymous" or "ftp".

  The password authentication service that you use determines whether users must log in with their NT domain name or UNIX user name.
- In the **Initial Directory for the user** field, type the path to the directory in which the user starts when he or she logs in over FTP. To create the path automatically when it does not already exist, select the **Create path if it does not exist** checkbox.

  Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to rwxrwxrwx). It is recommended that such directories are created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created via this option.

3. **Save your changes.**

   Verify your settings, then click **OK** to save your settings and return to the FTP Users page, where the new user is added to the table on the page.

Importing an FTP User

To import an FTP User:

1. **Navigate to the Import FTP Users page.**

   From the **File Services** page, click **FTP Users** to display the **FTP Users** page, then click the **Import Users** link to display the **Import FTP Users** page.

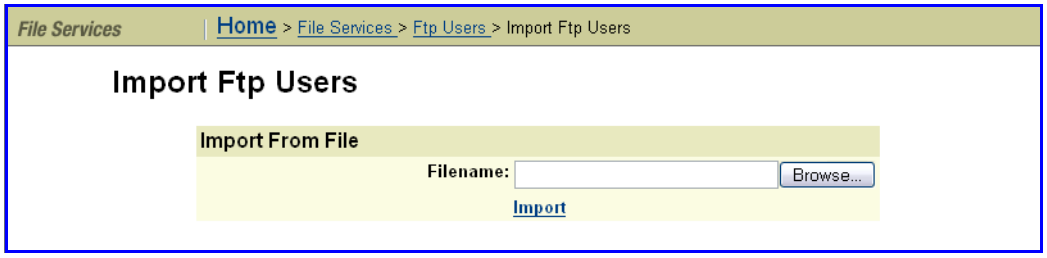

2. **Locate the import file.**

   In the **Filename** field, enter the file name that contains the user details, or click **Browse** to search for the filename.

   The user details in the import file have the following syntax:

   ```
   user_name file_system initial_directory
   ```

   Each entry must be separated by at least one space. If either the user_name or initial_directory contains spaces, the entry must be within double-quotes. For example:

   ```
   carla Sales /Sales/Documents
   miles Sales "/Sales/Sales Presentations"
   john Marketing /Marketing
   ```

   If you cannot be certain that the initial directory exists, you can create it automatically by specifying the option `ENSURE_PATH_EXISTS` on a separate line in the file. For example:

   ```
   ENSURE_PATH_EXISTS true
   carla Sales /Sales/Documents
   miles Sales "/Sales/Sales Presentations"
   ENSURE_PATH_EXISTS false
   john Marketing /Marketing
   ```

   In the first instance of the `ENSURE_PATH_EXISTS` option, the `true` attribute turns on the option, and it applies to the two following entries until the option is turned off by the second instance of the option, with the attribute `false`. The default for the `ENSURE_PATH_EXISTS` option is `true` so that the initial directory is automatically created.

   To insert a comment in the file, precede it with a hash character (#).

   Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to rwxrwxrwx). We recommend that such directories are created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created via this option.

3. **Import the file.**

   Click **Import.**

Viewing and Modifying FTP Users

To view and/or modify an FTP user:

1. **Navigate to the FTP Users page.**

   From the **File Services** page, click **FTP Users** to display the **FTP Users** page, then click **details** to display the **FTP User Details** page:

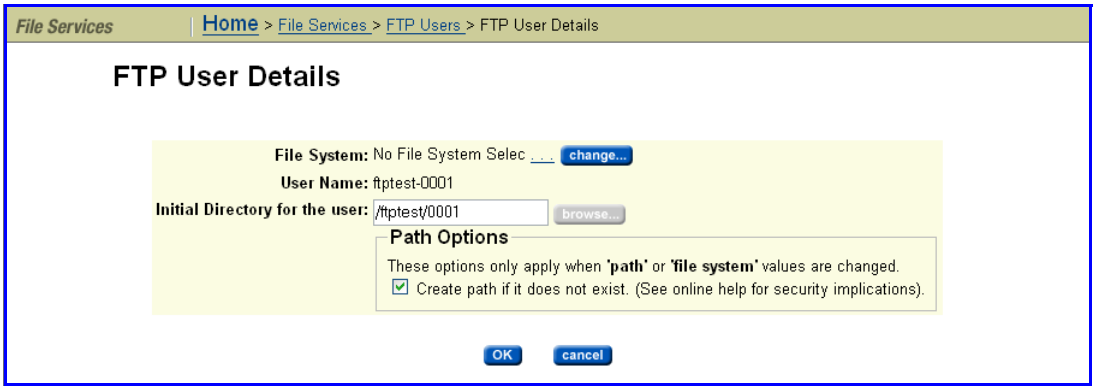

2. **Modify settings as needed:**
   - In the **File System** field, you can click **change** to select a different file system.
   - In the **Initial Directory for the user** field, you can change the directory by typing the path to the new directory. You can click the **browse** button to find the desired directory. This directory is the location where the user starts after logging in over FTP.
   - In the **Path Options** box, you can fill the checkbox **Create path if it does not exist** to create the path automatically when it does not already exist.

Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to rwxrwxrwx). It is recommended that such directories are created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created via this option.

3. **Start the process.**

   Verify your settings, then click **OK** to change the settings for this FTP User.

Setting Up FTP Audit Logging

FTP generates an audit log to keep track of user activity. The system will record the event when each time a user takes any of the following actions:

- Logging in or out
- Renaming or deleting a file
- Retrieving, appending or storing a file

- Creating or removing a directory

The system also records when a session timeout occurs.

Each log file is a tab-delimited text file containing one line per FTP event. Besides logging the date and time at which an event occurs, the system logs the user name and IP address of the client and a description of the executed command. The newest log file is called **ftp.log**, and the older files are called **ftp*n*.log** (the larger the value of ***n***, the older the file).

Configuring FTP Audit Logging

To view and modify the FTP Audit Logging configuration:

1. **Navigate to the FTP Audit Logging page.**

   From the **File Service** page, click **FTP Audit Log** to display the page:

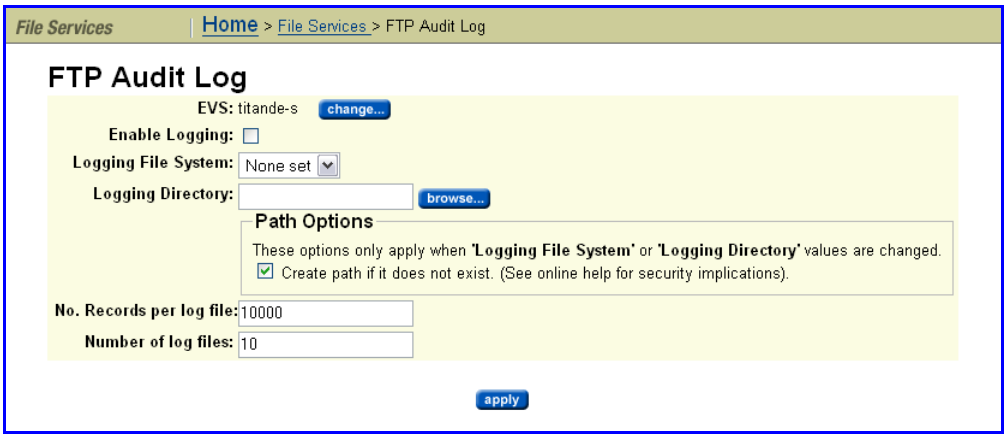

2. **View and/or modify settings as needed.**

   The following table describes the fields in this page:

Item/Field	Description
EVS	The **EVS** field specifies the EVS for the log files. To select a different EVS, click the **change** button.
Enable Logging	If this checkbox is filled, FTP audit logging is enabled.
File System	Select a file system on which to keep the log files.
Logging Directory	The directory on the specified file system in which to keep the log files.

Item/Field	Description
Path options	If this checkbox is filled, the Logging Directory is automatically created if it does not exist. **Note:** Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to rwxrwxrwx). It is recommended that such directories are created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created via this option.
No. Records per log file	The maximum number of records to store in each log file. For optimum performance, aim to produce a small number of large files rather than a large number of small files.
Number of log files	The maximum number of log files to keep. Once it has reached this limit, the server deletes the oldest log file each time it creates a new one.

3. **Save your settings.**
   Verify your settings, then click **apply** to save your changes.

Block-Level Access Through iSCSI

The storage server supports iSCSI, the Internet Small Computer System Interface (iSCSI) protocol enables block level data transfer between requesting applications and iSCSI Target devices. Using Microsoft's iSCSI Software Initiator (version 1.06 or later), Windows servers can view iSCSI Targets as locally attached hard disks. Windows can create file systems on iSCSI targets, reading and writing data as if it were on a *local* disk. Window server applications, such as *Microsoft Exchange* and *Microsoft SQL Server* can operate using iSCSI Targets as data repositories.

The server iSCSI implementation has attained the **Designed for Windows Server™ 2003** certification from Microsoft. The **Designed for Windows Server™ 2003** logo helps customers identify products that deliver a high quality computing experience with the Microsoft Windows Server 2003 operating system:

iSCSI Support

To use iSCSI storage on the server, one or more iSCSI Logical Units (LUs) must be defined. iSCSI Logical Units are blocks of SCSI storage that are accessed

through iSCSI Targets. iSCSI Targets can be found through an iSNS database or through a Target Portal. Once an iSCSI Target has been found, an Initiator running on a Windows server can access the Logical Unit as a "local disk" through its Target. Security mechanisms can be used to prevent unauthorized access to iSCSI Targets.

On the server, iSCSI Logical Units are just regular files residing on a file system. As a result, iSCSI benefits from file system management functions provided by the server, such as NVRAM logging, snapshots, and quotas.

The contents of the iSCSI Logical Units are managed on the Windows server. Where the server views the Logical Units as files containing raw data, Windows views each iSCSI Target as a logical disk, and manages it as a file system volume (typically using NTFS). As a result, individual files inside of the iSCSI Logical Units can only be accessed from the Windows server. Server services, such as snapshots, only operate on entire NTFS volumes and not on individual files.

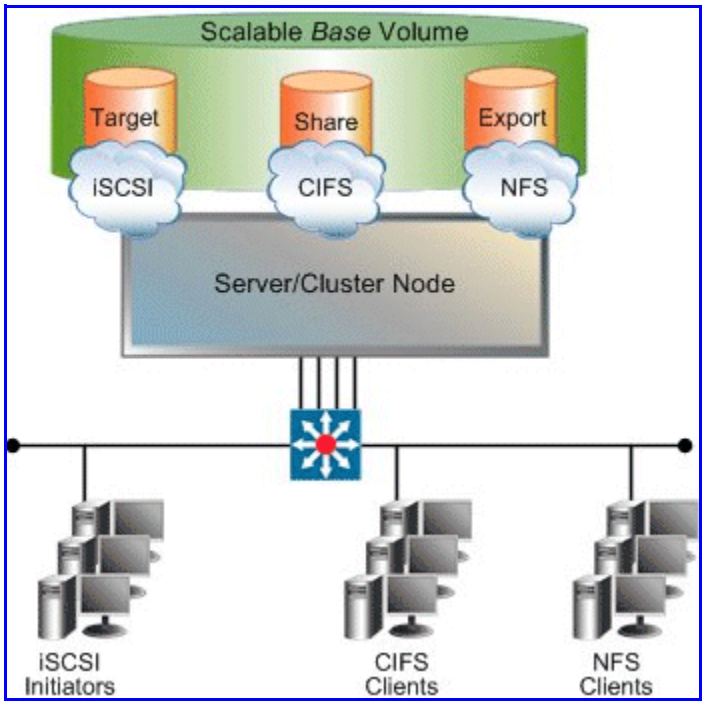

iSCSI MPIO

iSCSI MPIO (Multi-path Input/Output) uses redundant paths to create logical "paths" between the client and iSCSI storage. In the event that one or more of these components fails, causing the path to fail, multi-pathing logic uses an alternate path so that applications can still access their data.

For example, clients with more than one Ethernet connection can use logical paths to establish a multi-path connection to an iSCSI target on the server. Redundant paths mean that iSCSI sessions can continue uninterrupted in the event of the failure of a particular path. An iSCSI MPIO connection can also be used is to load-balance communication to boost performance.

If you intend to use an offload engine, make sure it is compatible with Microsoft multi-path and load-balancing.

iSCSI MPIO is supported by Microsoft iSCSI Initiator 2.0.

iSCSI Access Statistics

Statistics are available to monitor iSCSI activity since the server was last started or its statistics were reset. The statistics are updated every ten seconds.

Prerequisites

To enable iSCSI capability:

- Enter an iSCSI license key.
- Enable the iSCSI service.

Supported iSCSI Initiators

The server currently supports the following iSCSI initiators:

- Microsoft iSCSI Initiator version 1.06 (or later).
- Microsoft iSCSI Initiator version 2.05 (provides MPIO support).
- Linux iSCSI initiator versions 3.4.2, 3.6.2, 3.6.3, and 4.0.188.13 (available from the Linux iSCSI project on SourceForge).
- Solaris 10 U2 (64 bit) native initiator, iscsiadm v1.0.
- Macintosh (OS X 10.4 Tiger) ATTO Xtend SAN v3.10.
- Open iSCSI version 2.0.865.

Note: Other iSCSI initiators and/or versions of the initiators listed above may also work with the server, but have not been tested. Check with SGI Global Services for the latest list of supported iSCSI initiators.

Offload Engines

The server currently supports the use of the Alacritech SES1001T and SES1001F offload engines when used with the Microsoft iSCSI initiator version 1.06 or later. Check with SGI Global Services for the latest list of supported offload engines.

Configuring iSCSI

In order to configure iSCSI on the server, the following information must be specified:

- iSNS Servers.
- iSCSI Logical Units.
- iSCSI Targets (including iSCSI domain).
- iSCSI Initiators (if using mutual authentication).

Configuring iSNS

The Internet Storage Name Service (iSNS) is a network database of iSCSI Initiators and Targets. If configured, the server can add its list of Targets to iSNS, which allows Initiators to easily find them on the network.

The iSNS server list can be managed through the i**SNS** page. The server registers its iSCSI Targets with iSNS database when any of the following events occurs:

- A first iSNS server is added.
- An iSCSI Target is added or deleted.
- The iSCSI service is started.
- The iSCSI domain is changed.
- A server IP address is added or removed.

Creating and Deleting iSNS Servers

To create or delete an iSNS server:

1. **Navigate to the iSNS page.**

   From the **File Services** page, click **iSNS** to display the **iSNS Servers** page:

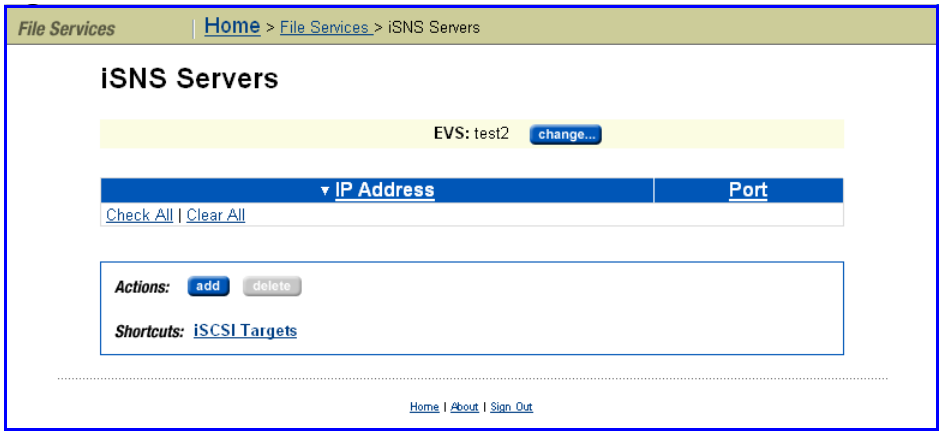

2. **If necessary, change the EVS.**

   The **EVS** name displayed indicates the EVS to which the iSNS server will be added. Click **change** to select a different EVS.

3. **Add or delete an iSNS server:**

- *To add an iSNS server* click **add** to display the **Add ISNS Server** page,

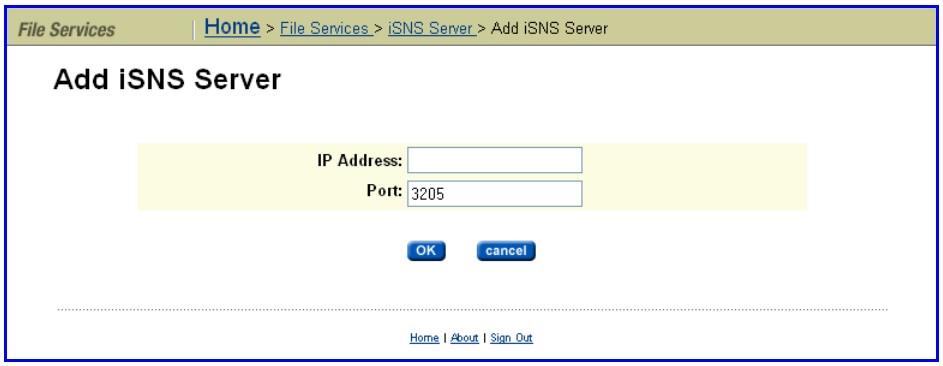

  then enter the iSNS server IP address and port number (default port number 3205) and click **OK**.

- *To delete an iSNS server,* select the iSNS server you want to delete (fill the checkbox next to the IP address of the iSNS server), then click **delete**.

Note: To download the latest version of Microsoft iSNS Server, visit: `http://www.microsoft.com`.

Configuring iSCSI Logical Units

Setting up iSCSI Logical Units

An iSCSI Logical Unit (LU) is a block of storage that can be accessed by iSCSI initiators as a locally attached hard disk. A Logical Unit is stored as a file on the server file system. Like any other data set on the file system, iSCSI Logical Units can be bound in size using the server's size management tools, including virtual volumes and Quotas. Logical Units are created with a specific initial size but can be expanded over time, as demand requires.

After a Logical Unit has been created and the iSCSI domain name has been set, an iSCSI Target must be created to allow access to the Logical Unit. A maximum of 32 Logical Units can be configured for each iSCSI Target.

Logical Unit Management

An iSCSI Logical Unit is a file within one of the server's file systems. Such a file must have an `.iscsi` extension to identify it as an iSCSI Logical Unit. However, apart from this extension there is no other way to determine that a file does indeed represent a Logical Unit.

Note: SGI Global Services recommends that all iSCSI Logical Units are placed within a well-known directory, for example `/.iscsi/`. This provides a single repository for the Logical Units in a known location.

Logical Unit Security

As Logical Units are files, they can be accessed over other protocols, such as CIFS and NFS. This renders Logical Units vulnerable to malicious users who can modify, rename, delete or otherwise affect them.

Caution: SGI Global Services recommends setting sufficient security on either

the Logical Unit file, the directory in which it resides, or both, to prevent unwanted accesses.

Concurrent Access to Logical Units

The server's iSCSI implementation allows multiple initiators to connect to a single Logical Unit, which is necessary for applications and operating systems that support, or rely upon, concurrent file system access. However, concurrent access can be detrimental to a client machine when the client is unaware of other clients accessing the file system. For example:

- **Simultaneous independent updates to the same files.** *Scenario:* Two independent Microsoft Windows clients can connect to the same Logical Unit, containing an NTFS file system. *Result:* If allowed to simultaneously and independently modify data, metadata, and system files, conflicting disk updates will quickly corrupt the file system.
- **Simultaneous access to separate partitions.** *Scenario:* A Logical Unit contains two distinct NTFS partitions, with one Microsoft Windows client connected only to the first partition, and another connected only to the second partition. *Result:* Because a Microsoft iSCSI client will attempt to mount each partition it encounters on the Logical Unit, a Microsoft Windows client mounting an NTFS partition updates system files on all partitions; therefore, even though the two clients are accessing *separate* partitions within the Logical Unit, both will update system files on both partitions, causing conflicting system file updates, causing one or both of the clients to fail.

Taking Snapshots of Logical Units

The contents of an iSCSI Logical Unit are controlled entirely by the client accessing it. The server cannot interpret the file systems or other data contained within a Logical Unit in any way. Therefore, the server has no knowledge of whether the data held within an iSCSI Logical Unit is in a consistent state. This introduces a potential problem when taking a snapshot of a Logical Unit.

For example, when a client creates a file, it must also insert the file name to the host directory. This means that more than one write is required to complete the operation. If the server takes a snapshot after the file object has been created, but before its name has been inserted into the directory, the file system contained within the snapshot will be inconsistent. If another client were to view the snapshot copy of the file system, it would see a file object without a name in a directory. This example provides only one possible scenario for snapshot inconsistency.

Caution: SGI Global Services recommends that prior to taking a snapshot of an iSCSI Logical Unit, all applications should be brought into a known state. A database, for example, should be quiesced. Disconnecting the iSCSI initiators from the Logical Units undergoing snapshot is also recommended. This guarantees that all pending writes are sent to the Logical Unit before the snapshot is taken.

Volume Full Conditions

Unexpected Volume Full conditions can occur with iSCSI Logical Units, as illustrated by the following two examples:

- **Directly Attached Disks.** When a client uses a directly attached disk, it can monitor the amount of available free space. If a partition contains no free space, the client can return a Volume Full condition. In this way, the client can ensure against file system corruption due to running out of disk space part way through an operation.
- **iSCSI Logical Unit.** By way of background, on iSCSI Logical Units with snapshots enabled, old data is preserved, not overwritten. Therefore, overwriting an area of a Logical Unit causes the server to allocate extra disk space, while using no extra disk space within the client's partition, causing a Volume Full condition to occur, even when partitions within the Logical Unit contain free space. Under this scenario, a client may receive a Volume Full condition part-way through an operation, causing file system corruption. Although this corruption should be fixable, this situation should be avoided.

Caution: *File system corruption on iSCSI Logical Units Alert!* SGI Global Services recommends allocating sufficient disk space on the server to contain all iSCSI Logical Units and snapshots, as well as careful monitoring of free disk space.

Managing iSCSI Logical Units

Viewing the Properties of iSCSI Logical Units

To view the properties of iSCSI Logical Units, navigate to the **File Services** page, then click **iSCSI Logical Units** to display the **iSCSI Logical Units** page:

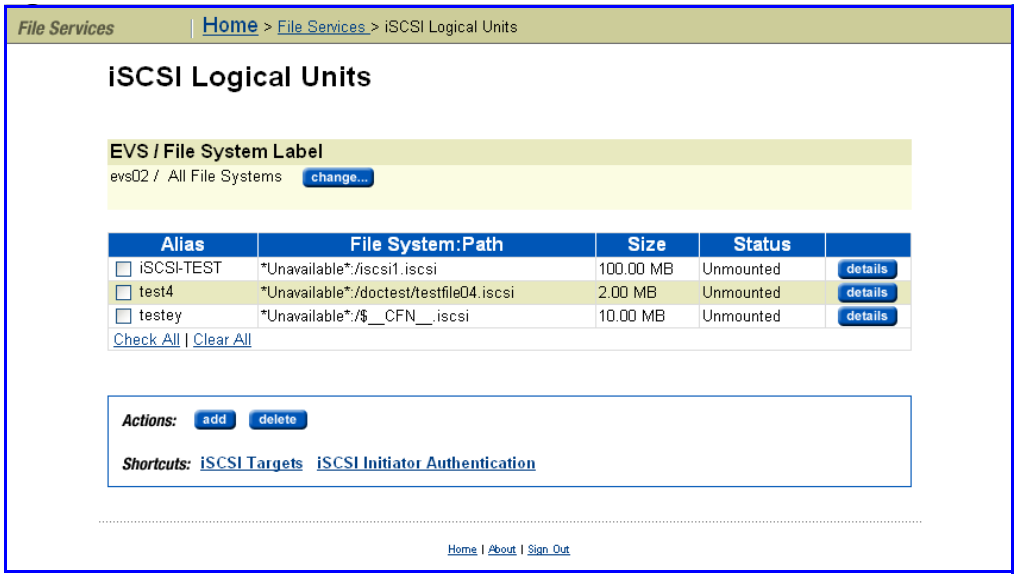

The table below describes the fields in this page:

Item / Field	Description
EVS/File System	Selector for EVS and File System where Logical Units reside, or where Logical Units can be created.
Alias	Name of the Logical Unit.
File System:Path	The file system and path for the Logical Unit. **Note:** Logical Units appear as regular files in server file systems.
Size	Size of the Logical Unit. **Note:** The maximum size of a Logical Unit is 2 TB. This limit is imposed by the SCSI protocol.
Status	Indicates of the Logical Unit status, usually whether the Logical Unit is mounted. **Note:** The status will display **Unmounted** while a LU is being created asynchronously, and will then display **Mounted** once the creation has completed.

Adding iSCSI Logical Units

1. **Navigate to the add iSCSI Logical Units page.**

   From the **File Services** page, select **iSCSI Logical Units**, then click **add** to display the **Add iSCSI Logical Units** page:

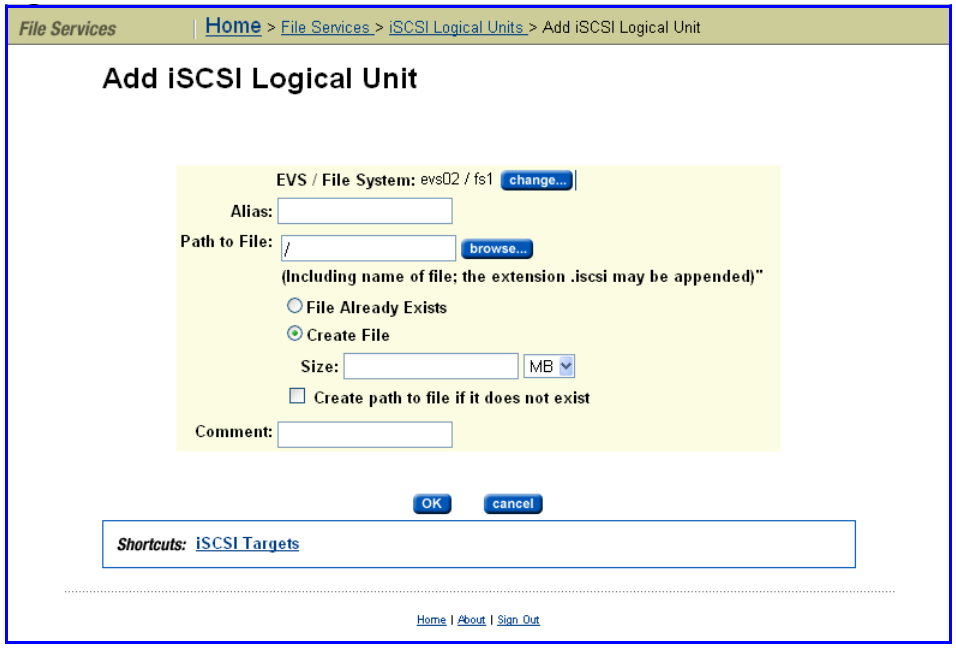

2. **If necessary, change the EVS and/or file system.**

The **EVS** name displayed indicates the EVS and file system to which the Logical Unit will be added. Click **change** to select a different EVS or file system.

3. **Specify the Logical Unit alias.**

   In the **Alias** field, enter a name for the Logical Unit.

4. **If the path to the file already exists, specify the path to the Logical Unit.**

   There are several steps to complete when entering a path for a Logical Unit file that already exists:

   a. **Choose the file.**

      Click **browse** to display a dialog that will allow you to select the file for the Logical Unit. Alternatively, you can enter the path name of the file (including the extension) and not use the **browse** button.

   b. **Select the File Already Exists radio button.**

   c. **Optionally, add a comment.**

      Using the comment field, you can provide descriptive information about the Logical Unit.

   d. **Save the Logical Unit definition.**

      Click **OK** to add the Logical Unit.

5. **If the path to the file does not already exist, specify to create the path to the Logical Unit.**

   There are several steps to complete when creating a new file for use as a Logical Unit:

   a. **Choose the path for the file.**

      Click **browse** to display a dialog that will allow you to select the directory for the Logical Unit file. The name of the file as well as the directory need to be specified. The file does not need an extension; `.iscsi` is appended automatically. Alternatively, you can enter the file name and path (including the extension) and not use the **browse** button.

   b. **Select the Create File radio button.**

   c. **Specify the Logical Unit size.**

      Using the Size field and the drop-down list, specify the size of the Logical Unit file.

   d. **Fill the Create path to file if it does not exist checkbox.**

   e. **Optionally, add a comment.**

      Using the comment field, you can provide descriptive information about the Logical Unit.

   f. **Save the Logical Unit definition.**

      Click **OK** to add the Logical Unit.

Modifying or Deleting an iSCSI Logical Unit

To modify or delete the properties of an iSCSI Logical Unit:

1. **Navigate to the iSCSI Logical Units page.**

   From the **File Services** page, click **iSCSI Logical Units** to display the **iSCSI Logical Units** page:

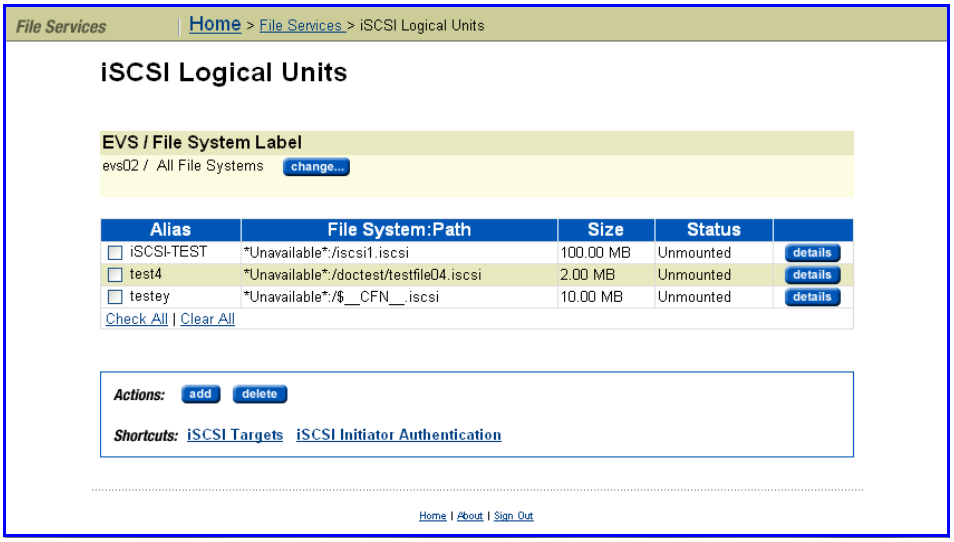

2. **If necessary, change the EVS and/or file system.**

   The **EVS** name displayed indicates the EVS and file system where the Logical Unit resides. Click **change** to select a different EVS or file system.

3. **Select the Logical Unit to modify or delete.**

   Fill the checkbox next to the Logical Unit you want to modify or delete.

4. **Modify or Delete.**

   - *To delete an iSCSI Logical Unit*, select the LU and click **delete**. The **Confirm Delete** pop-up window displays the Delete options. Make your selection and click **OK** to delete the LU, or **cancel** to close the pop-up without deleting the LU.

- *To modify an iSCSI Logical Unit:*

  Click **details** to display the **iSCSI Logical Unit Details** page:

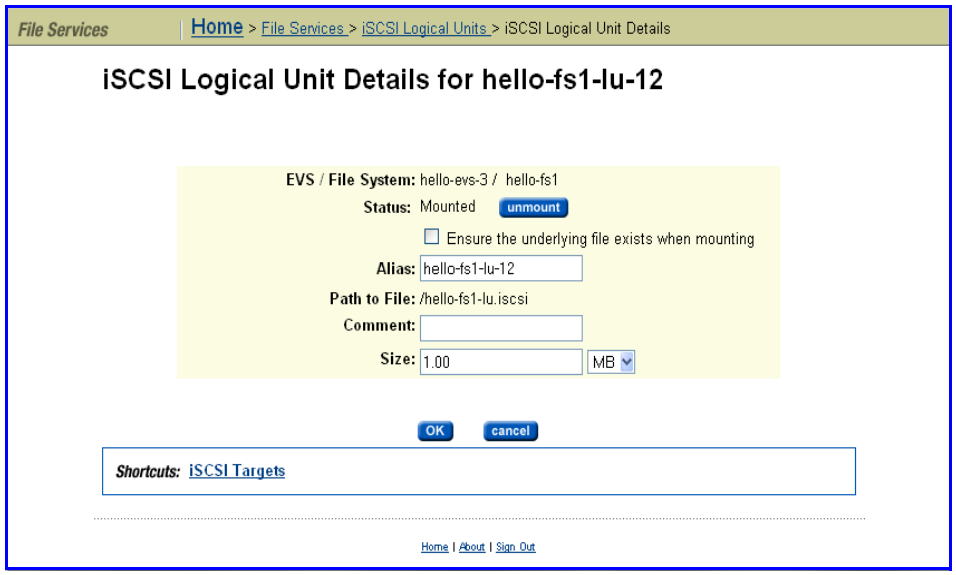

The table below describes the fields in this page:

Item / Field	Description
EVS/File System	Displays the EVS and file system hosting the Logical Unit.
Status	Indicates whether the Logical Unit is mounted or unmounted. It is possible to mount/unmount a Logical Unit while its underlying file system remains mounted. If the Logical Unit is not mounted, click **mount** to mount the Logical Unit. If the Logical Unit is mounted, click **unmount** to unmount the Logical Unit. Filling the **Ensure the underlying file system exists when mounting** checkbox will ensure that the underlying file system exists when the Logical Unit is mounted.
Alias	Name of the Logical Unit. You can change this name.
Path to File	The complete file system path to the Logical Unit file. **Note:** Logical Units appear as regular files in server file systems.
Size	Size of the Logical Unit. **Note:** The maximum size of a Logical Unit is 2 TB. This limit is imposed by the SCSI protocol.

5. **Modify the properties of the Logical Unit.**

   Modify the Logical Unit settings as needed.

6. **Save your settings.**

   Verify your settings, then click **OK** to save the changes, or click **cancel** to decline.

Backing up iSCSI Logical Units

Only a client connected to the Logical Unit through its Target can access and backup individual files and directories contained in the Logical Unit. Logical Units back up as normal files on a server file system.

Caution: If backing up the iSCSI Logical Unit from the server, ensure that the iSCSI initiators are disconnected, or make the backup from a snapshot.

To back up an iSCSI Logical Unit:

1. **Disconnect the iSCSI Initiator from the Target.**
2. **Unmount the iSCSI Logical Unit.**

   To unmount the iSCSI Logical Unit, you can use the following CLI command:

   `iscsi-lu unmount <name>`

   Where **<*name*>** is the name of the iSCSI Logical Unit.
3. **Back up the Logical Unit to a snapshot or backup device.**

   For safety, you should either back up the iSCSI Logical Unit to a snapshot or to another backup device. See Taking Snapshots of Logical Units, on page 295 for more information on this process.
4. **Mount the Logical Unit.**

   To mount the iSCSI Logical Unit, you can use the following CLI command:

   `iscsi-lu mount <name>`

   Where **<*name*>** is the name of the iSCSI Logical Unit.
5. **Reconnect to the iSCSI Target using the iSCSI Initiator.**
6. **If necessary, rescan disks.**

   You may have to use Window's Computer Manager rescan disks to make the Logical Unit reappear to clients. See Using Computer Manager to Configure iSCSI Storage, on page 318 for more information on this process.

Restoring iSCSI Logical Units

To ensure consistency of data on a Logical Unit, it may be necessary to restore it from a snapshot or a backup. To restore an iSCSI Logical Unit, perform the following steps:

To restore an iSCSI Logical Unit:

1. **Disconnect the iSCSI Initiator from the Target.**

2. **Unmount the iSCSI Logical Unit.**

   Use the following CLI command: `iscsi-lu unmount <name>`, where name is the name of the Logical Unit.

3. **Restore the Logical Unit from a snapshot or backup.**

4. **Mount the iSCSI Logical Unit.**

   Use the following CLI command: `iscsi-lu mount <name>`, where name is the name of the Logical Unit.

5. **Reconnect to the Target using the iSCSI Initiator.**

6. **If necessary, rescan disks in Computer Management.**

   Refer to Using Computer Manager to Configure iSCSI Storage, on page 318.

Setting Up iSCSI Targets

An iSCSI Target is a storage element accessible to iSCSI initiators. These targets appear to iSCSI initiators as different storage devices accessible over the network. The server supports a maximum of 32 iSCSI Targets per EVS and a maximum of 32 iSCSI sessions per Target.

Viewing the Properties of iSCSI Targets

On the **File Services** page, click **iSCSI Targets** to display the **iSCSI Targets** page:

The table below describes the fields on this page:

Item / Field	Description
EVS	Select the EVS on which the Target will be hosted. Click **change** to select a different EVS.

Item / Field	Description
EVS iSCSI Domain	Displays the iSCSI Domain, which is the DNS domain used when creating unique qualified names for iSCSI Targets.
Alias	Identifies the name of the Target.
Comment	Additional information related to the Target.
Globally Unique Name	The Target's name. The name is generated automatically by the server, and is unique across the globe.

Adding iSCSI Targets

To add an iSCSI Target:

1. **Navigate to the add iSCSI Target page.**

   On the **File Services** page, select **iSCSI Targets,** then click **add** to display the **Add iSCSI Target page**:

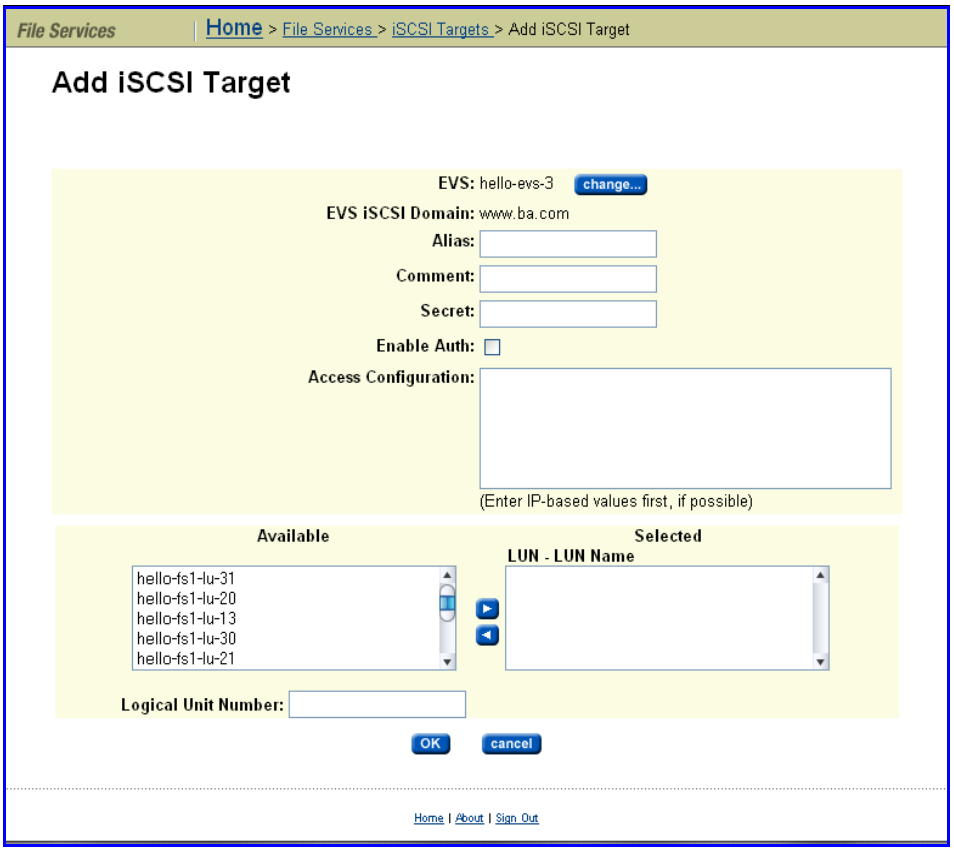

The table below describes the fields on this page:

Item / Field	Description
EVS	Selector for EVS where Logical Units reside. Click **change** to switch to a different EVS.
EVS iSCSI Domain	Displays the iSCSI Domain, which is the DNS domain used when creating unique qualified names for iSCSI Targets.
Alias	Name of the iSCSI Target.
Comment	Additional information about the iSCSI Target.
Secret	Password used to secure the Target from any unauthorized access. The initiator authenticates against this password when connecting to the Target. The secret should be greater than or equal to 12 characters, but less than 17 characters, in length. Although the secret may be between 1-255 characters in length, some iSCSI initiators will refuse to connect if the secret contains less than 12 characters or more than 16 characters.
Enable Auth	Enable authentication of the iSCSI Target. By default, the checkbox is not filled. Filling or clearing the checkbox enables or disables authentication. When authentication is disabled, initiators are permitted to connect to the target and its logical units without needing to know the target's secret.
Access Configuration	Enter access configuration parameters. Refer to the Access Configuration table below for syntax.
Available logical units	The list of Logical Units available for assignment to the iSCSI Target. This list includes all LUs on the EVS. Some of these LUs may already be assigned to other targets.
Selected LUN - LUN Name	The list of Logical Units selected to be part(s) of the iSCSI Target.
Logical Unit Number	The number assigned to the Logical Unit (the LUN). Enter a Logical Unit Number in the range of 0-255, then click **OK**.

2. **Enter the information about the iSCSI Target.**
   a. **Specify the required information.**

      The iSCSI Domain, Alias, Available Logical Units, and Logical Unit Numbers are required, all other fields are optional.

b. Optionally, specify the Comment, Secret, and/or Access Configuration for the Target.

The following table provides syntax for the **Access Configuration** field:

What to type	Means
Blank or *	All clients can access the target.
Specific address or name. Examples: 10.168.20.2, client.dept.company.com To deny access to a specific host, use the no_access or noaccess qualifier. For example, 10.1.2.38 (no_access) will deny access to the host with the IP address 10.1.2.38.	Only clients with the specified names or addresses can access the target.
Partial address or name using wildcards. Examples: 10.168.*.*, *.company.com To deny access to a specific host, use the no_access or noaccess qualifier. For example, 10.1.2.38(no_access) will deny access to the host with the IP address 10.1.2.38.	Clients with matching names or addresses can access the target.

3. Save your settings

Verify your settings, then click **OK** to save or **cancel** to decline.

Adding a LU to an iSCSI Target

To add a Logical Unit to an iSCSI target:

1. Navigate to the iSCSI Target Details page.

On the **File Services** page, click **iSCSI Targets** to display the **iSCSI Targets** page, then click **details** to display the **iSCSI Target Details** page for the iSCSI target you want to modify:

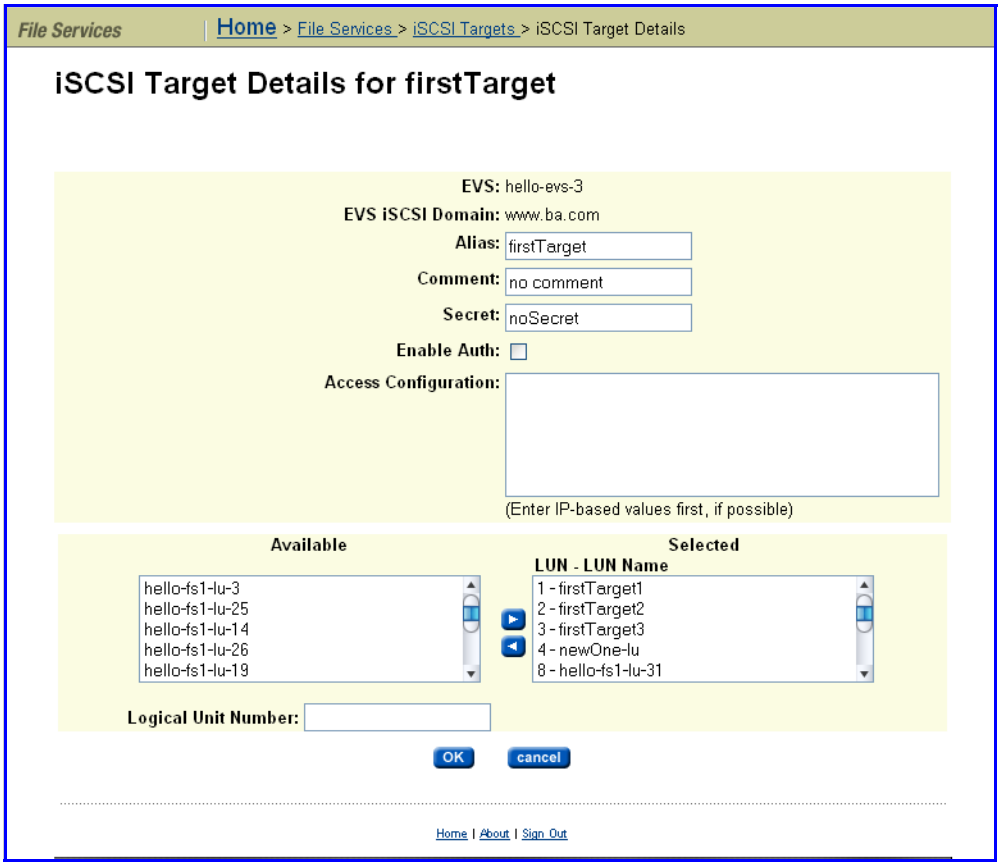

2. **Add the iSCSI LU.**

   Select a Logical Unit from the **Available Logical Units** list, specify a number (0-255) in the **Logical Unit Number** field, and click the right arrow to move the LU to the **Selected Logical Units** list.

   Note: You should make sure that the Logical Unit is not already assigned to a target.

3. **Save your changes.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Modifying the Properties of an iSCSI Target

To modify the properties of an iSCSI target:

1. **Navigate to the iSCSI Target Details page.**

On the **File Services** page, click **iSCSI Targets** to display the **iSCSI Targets** page, then click **details** to display the **iSCSI Target Details** page for the iSCSI Target you want to modify:

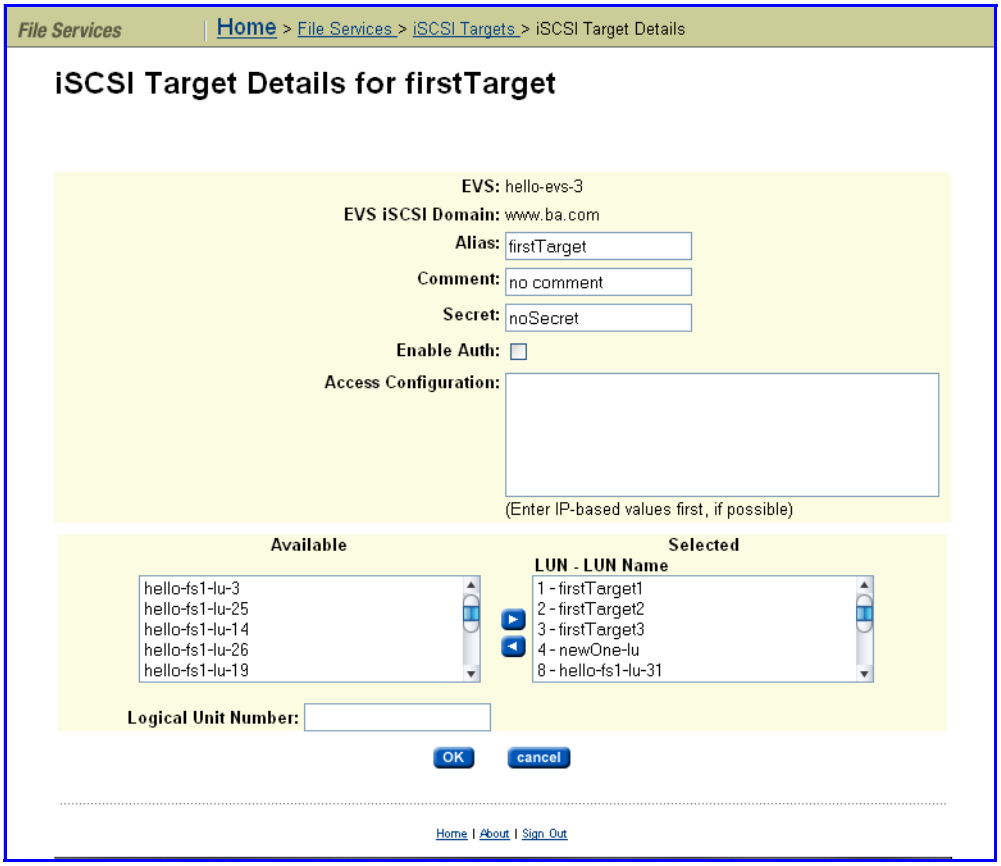

2. **Modify the iSCSI Target.**

   The iSCSI Domain, Alias, Available Logical Units, and Logical Unit Numbers are required. Optionally, you can specify the Comment, Secret, and/or Access Configuration for the Target. For a description of the items and fields on this page, see Adding iSCSI Targets, on page 303.

   Note: Once set, the iSCSI Domain cannot be changed, but it will be overridden/replaced if you later specify a new iSCSI Target with a different iSCSI Domain in the same EVS. **The most recently specified iSCSI Domain overrides all previously-specified iSCSI Domains** set for all previously added iSCSI Targets in the EVS.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Deleting iSCSI Targets

To delete an iSCSI Target:

1. **Navigate to the iSCSI Targets page.**

On the **File Services** page, click **iSCSI Targets** to display the **iSCSI Targets** page:

2. **Delete an iSCSI Target.**

   Select a Target from the list by filling the checkbox next to the Alias, then click **delete**. A confirmation dialog appears, and you can click **OK** to delete the iSCSI Target, or **cancel** to return to the iSCSI Targets page without deleting the Target.

Configuring iSCSI Security (Mutual Authentication)

The storage server uses the Challenge Handshake Authentication Protocol (CHAP) to authenticate iSCSI Initiators. CHAP requires a "shared secret" known by the Initiator and the Target. The server also supports mutual authentication where in addition to the Initiator authenticating against the Target on the server, the server must also authenticate against the Initiator.

To facilitate the mutual authentication process, the server must maintain a list of the Initiators with which it can authenticate and the shared secret for each Initiator.

Configuring the Storage Server for Mutual Authentication

To configure the server for mutual authentication:

1. **Navigate to the iSCSI Initiator Authentication page.**

   On the **File Services** page, click **iSCSI Initiator Authentication** to display the **iSCSI Initiator Authentication** page:

This page lists all the iSCSI Initiators currently specified for the selected EVS. The table below describes the fields on this page:

Item / Field	Description
EVS	Select the EVS on which to configure Initiator Authentication. Click **change** to select a different EVS.
Initiator Name	Identifies the initiator with a globally unique name.
Secret	Password used to secure the Initiator from any unauthorized access. The secret should be from 12 to 17 characters in length, but may be between 1-255 characters in length.

2. **Add an iSCSI Initiator.**

   a. **Navigate to the Add iSCSI Initiator Authentication page.**

   Click **add** to display the to the Add iSCSI Initiator Authentication page:

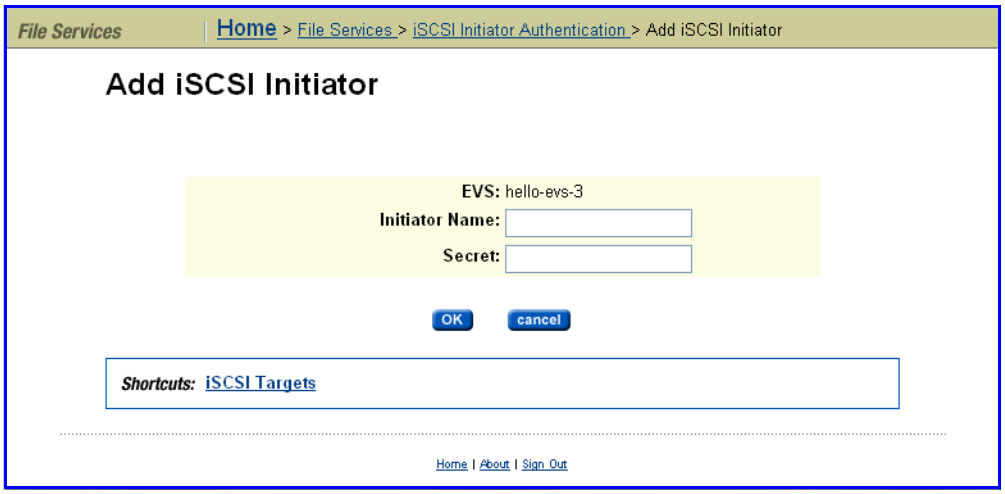

b. **Specify the Initiator name.**

Enter an **Initiator name**. This is the same name as will display in the **Change Initiator node name change** dialog of the **Microsoft iSCSI Initiator**.

c. **Specify the Secret.**

Enter the **Secret** for the Initiator. This is the secret which will be entered in the **Chap Secret Setup** dialog of the iSCSI Initiator.

d. **Save the configuration.**

Verify your settings, then click **OK** to save or **cancel** to decline.

Configuring the Microsoft iSCSI Initiator for Mutual Authentication

Note: For the latest version of Microsoft iSCSI Software Initiator, visit: http://www.microsoft.com/.

To configure the Microsoft iSCSI Initiator for mutual authentication:

1. **Open iSCSI Initiator Properties.**

   Start the Microsoft iSCSI Initiator, and open **Initiator Properties** showing the **General** tab:

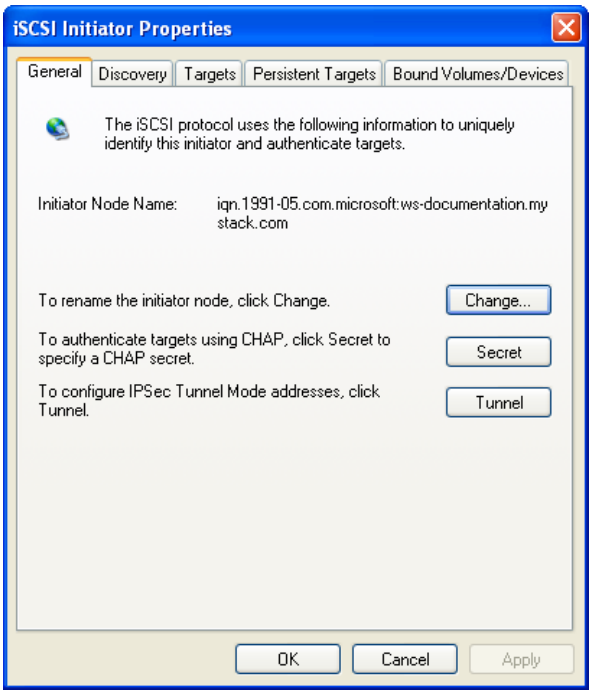

2. **Enter a secret.**

   Click the **Secret** button to display the **CHAP Secret Setup** dialog.

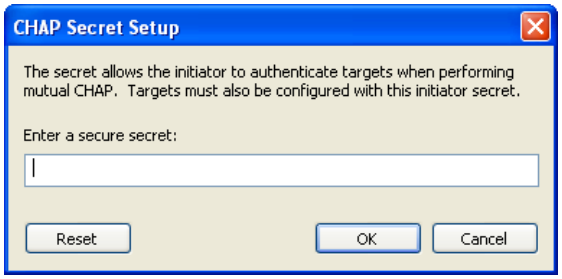

In the field, enter the secret which allows the Target to authenticate with Initiators when performing mutual CHAP, then click **OK**.

Note: The shared secret used to authenticate an Initiator with a server should be different from the secret specified when setting up the Target.

3. **Optionally, change the initiator node name.**

   If you want to change the initiator node name, click the **Change** button to display the **Initiator Node Name Change** dialog.

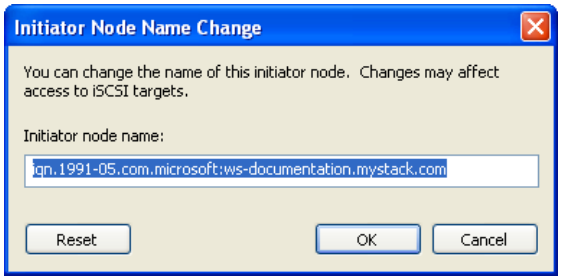

In the field, change the name and click **OK**.

Note: The **Initiator node name** is the name which should be used as the **Initiator Name** on the **iSCSI Initiator Authentication** page, found under the **File Services** page.

4. **Save your changes.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Changing the Storage Server's Mutual Authentication Configuration

To change the server's mutual authentication configuration, you can change the initiator's secret, but the initiator's name cannot be changed.

1. **Navigate to the iSCSI Initiator Authentication page.**

   On the **File Services** page, click **iSCSI Initiator Authentication** to display the **iSCSI Initiator Authentication** page.

2. **If necessary, modify an Initiator's secret.**

   If it is necessary to change an Initiator's secret:

a. **Navigate to the iSCSI Initiator Details page.**

Click **details** to display the **iSCSI Initiator Details** page:

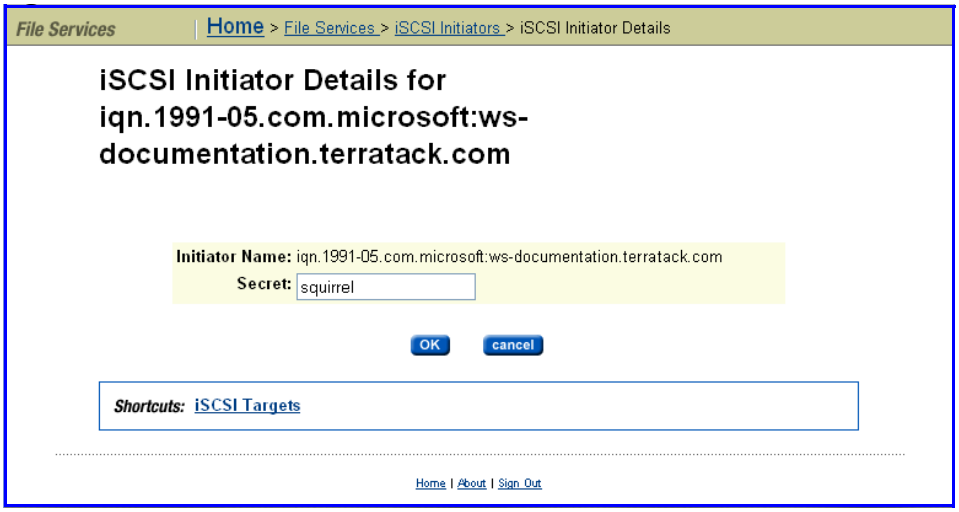

b. **Enter the new secret.**

In the **Secret** field, type the new secret for the initiator. The secret should be from 12 to 17 characters in length, but may be between 1-255 characters in length.

c. **Save the configuration.**

Verify your settings, then click **OK** to save or **cancel** to decline.

3. **If necessary, delete an iSCSI Initiator.**

To delete an Initiator and its Secret, fill the checkbox next to the **Initiator Name**, and click **delete**. A confirmation dialog appears, and you can click **OK** to delete the iSCSI Initiator, or **cancel** to return to the iSCSI Initiators page without deleting the Initiator.

Accessing iSCSI Storage

iSCSI Logical Units can be accessed through their Targets using the Microsoft iSCSI Initiator. Discovered through iSNS or through the Target Portal, all iSCSI Targets that are available will be displayed as available Targets by the Initiator.

Caution: *Dynamic disk alert!* Microsoft currently only supports creating a Basic Disk on an iSCSI Logical Unit. To ensure data integrity, do not create a dynamic disk. For more information, refer to the Microsoft iSCSI Initiator User Guide.

If its underlying volume is mounted read-only by the storage server, or if it is a snapshot copy of another Logical Unit, an iSCSI Logical Unit will also be read-only. In turn, if a Logical Unit is read-only, then any file systems contained within it will also be read-only. Clients accessing such read-only file systems will not be able to change any part of them, including file data, metadata and system files.

Although they can mount read-only FAT and FAT32 file systems, Microsoft Windows 2000 clients cannot mount read-only NTFS file systems. Microsoft Windows 2003 clients can mount read-only FAT, FAT32 and NTFS file systems. Therefore, if Microsoft Windows clients are required to access read-only NTFS file systems over iSCSI, Microsoft Windows 2003 must be used.

Using iSNS to Find iSCSI Targets

Using iSNS is the easiest way to find iSCSI Targets on the network. If the network is configured with an iSNS server, configure the Microsoft iSCSI Initiator to use iSNS.

To add an iSNS server:

1. **Open the iSNS Servers tab.**

   Within the Microsoft iSCSI Initiator, click the **Discovery** tab:

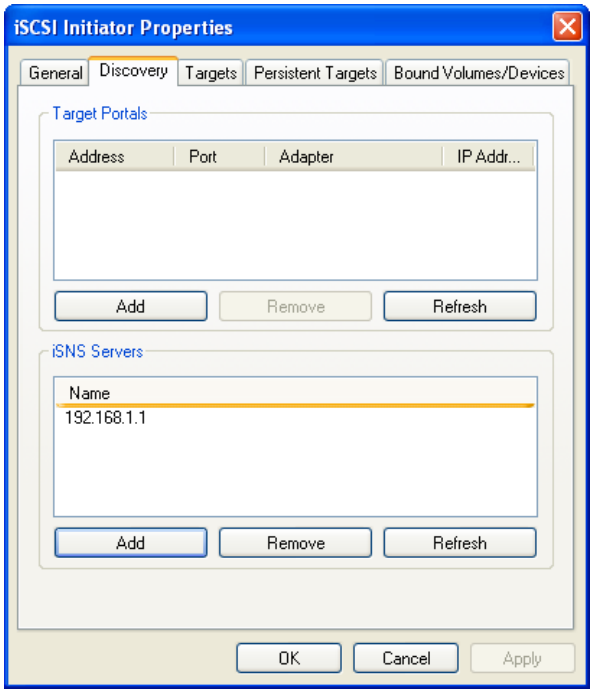

2. **Add the iSNS server.**

   In the **iSNS Servers** section of the tab, click the **Add** button to display the **Add iSNS Server** dialog:

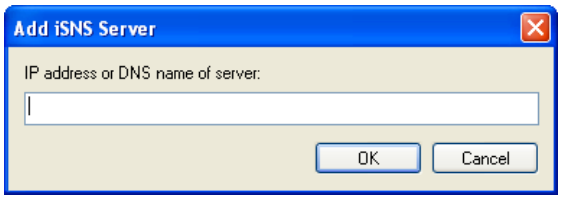

   Enter the iSNS Server's IP Address or DNS name, then click **OK.**

   Note: After the iSNS server(s) have been added, all available iSCSI Targets

that have been registered in iSNS will appear as available Targets.

3. **Save your changes.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Using Target Portals to find iSCSI Targets

If there are no iSNS servers on the network, iSCSI Targets can be found through the use of Target Portals. Add the server's EVS IP to the Target Portals list to find Targets associated with that server or EVS.

1. **Open the Target Portals tab.**

   Within Microsoft's iSCSI Initiator, click the **Discovery** tab:

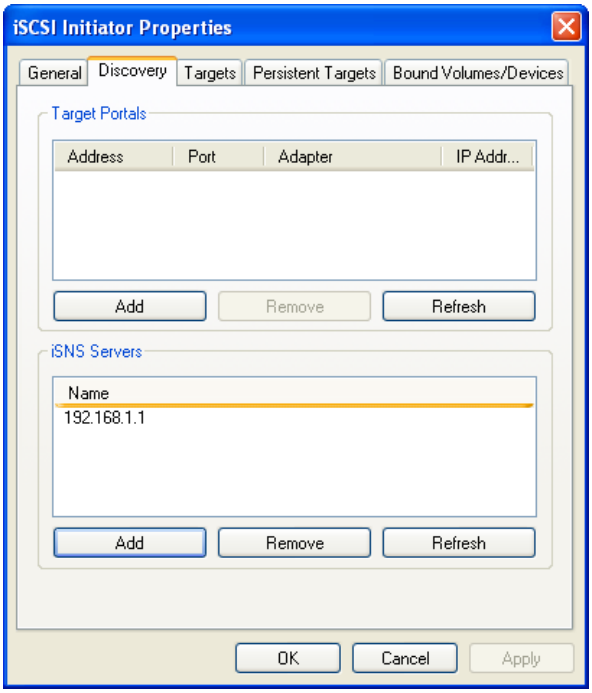

2. **Add a Target Portal.**

   In the **Target Portals** section of the tab, click the **Add** button to display the **Add Target Portal** dialog:

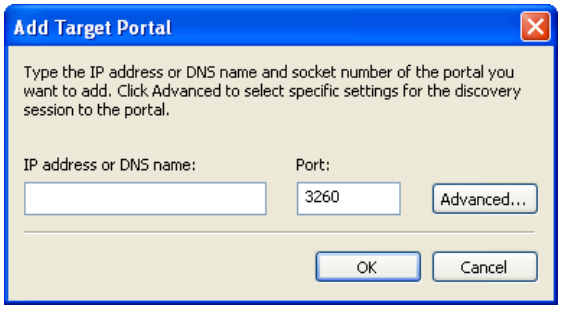

   Enter the server's file service IP address.

3. **Save your changes.**
   Verify your settings, then click **OK** to save or **Cancel** to decline.

Accessing Available iSCSI Targets

To access an available iSCSI Target:

1. **Open the Targets tab.**
   Within Microsoft's iSCSI Initiator, click the **Targets** tab:

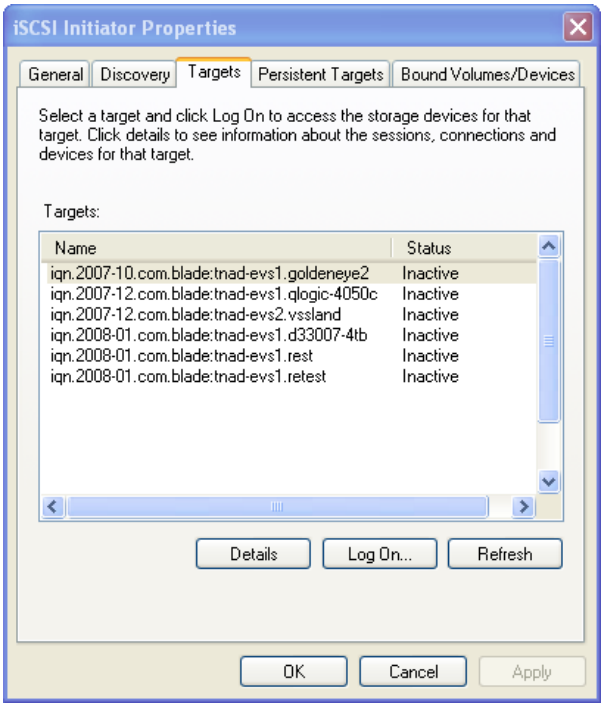

2. **Log on to an Available Target.**
   Each logon starts an iSCSI session:

   - Select a Target, then click the **Log On** button to display the **Log On to Target** dialog:

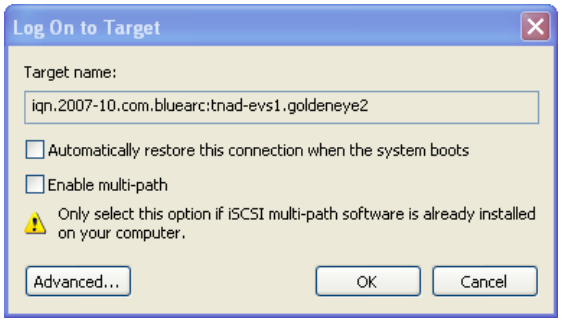

Note: A maximum of 32 iSCSI sessions are allowed per Target.

- If authentication is enabled on the Target, click **Advanced** to open the **General** tab in the **Advanced Settings** dialog:

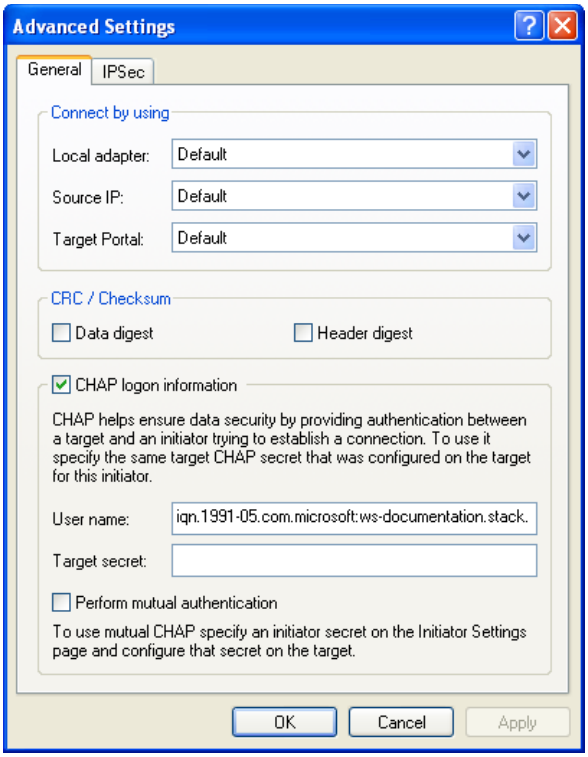

  Fill the **CHAP logon information** checkbox and enter the **Target secret** (the password configured when the iSCSI Target was created); and, if Mutual Authentication has been configured, fill the **Perform mutual authentication** checkbox; then click **OK.**

3. **Optionally, configure multi-pathing.**

   If multi-pathing is supported by the Microsoft iSCSI Initiator, and you want to use multiple paths to the Target, fill the **Enable Multi-Path** checkbox.

   To create multiple paths to the Target, you must later start another session to the Target. See iSCSI MPIO, on page 291 for more information.

4. **Establish the Connection.**

   In the **Log On** dialog, click **OK** to save, or **Cancel** to decline.

Verifying an Active Connection

After the connection has been established, click the **Targets** tab to view any details about the newly established connection.

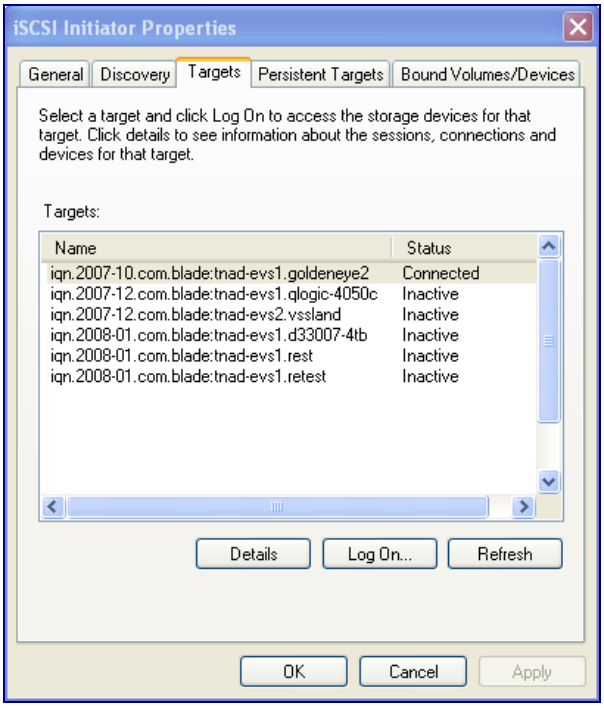

The **Status** column for the Target should display *Connected*.

Terminating an Active Connection

Once the connection has been established, the **Targets** tab displays the connection status:

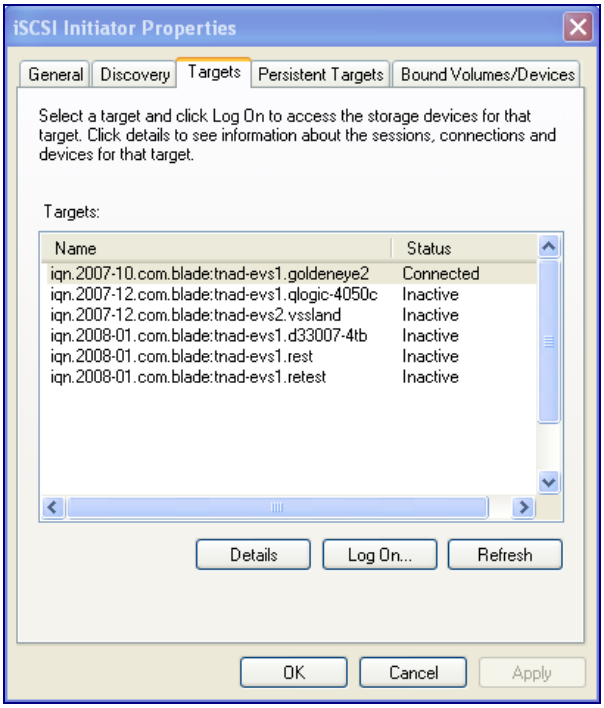

To end a connection:

1. **Select the connection you want to end.**
2. **Display the Target Properties dialog.**
   Click the **Details** button to display the Target Properties dialog:

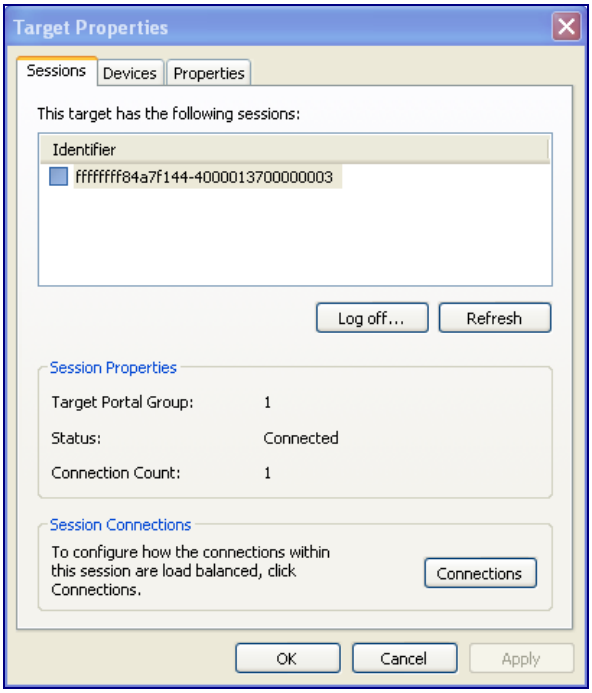

3. **Select the connection to terminate.**
   In the list of session identifiers, select the identifier for the session you want to end.
4. **Terminate the connection.**
   Click the **Log off...**button to terminate the session. The initiator will attempt to close the iSCSI session if there are no applications currently using the devices.

Using Computer Manager to Configure iSCSI Storage

The iSCSI "local disk" must be configured through Windows Disk Management tools, and Microsoft recommends that:

- If the LU is smaller than 2 Tb, it should be configured as a Basic Disk.
- If the LU is larger than 2 Tb, it should be configured as a GPT Disk.

Once the disk is configured, use the Windows Disk Management tools to create and format a partition on the disk. For information and instructions on using the Windows Disk Management tools, refer to the online help of the operating system on your Windows computer.

Accessing Snapshots Initiated by VSS

The VSS Hardware Provider DLL provides support for taking snapshots initiated by Microsoft's Volume Shadow Copy Service (VSS). The VSS Hardware Provider allows you to take snapshots of iSCSI LUs located on storage devices managed by NAS servers. VSS initiated snapshots are exported as iSCSI LUs. Once a snapshot has been taken and exported (or "surfaced"), a pointer is provided to the application that requested the snapshot. Using this pointer, the application can then access the snapshot to back up database-type applications such as Microsoft Exchange and SQL Server.

The VSS Hardware Provider runs on a Windows server (see Installing the VSS Hardware Provider, on page 321). Once installed, the VSS Hardware Provider registers with the Microsoft VSS Service.

When a VSS initiated snapshot is taken of a file system, the snapshot is added to the iSCSI target as one or more iSCSI LUs (one iSCSI LU is added for each source LU supplied to VSS). The snapshot LUs are then visible to the VSS host (the system on which the VSS Hardware Provider is installed) and are used as the backup source.

Note: If a VSS snapshot request contains LUs on different file systems, then only one snapshot will be created for all the LUs in each file system. However, copies of each requested LU are always created and made visible to the VSS host by the NAS server.

Each NAS server or cluster must be configured to allow VSS access. This is a multi-step process, and is described in Setting up the NAS Server for VSS Snapshots, on page 320

Removing VSS Initiated Snapshots

Snapshots initiated by the VSS should be managed through the application that requested the snapshot. Snapshots are either non-persistent or persistent.

For *non-persistent* snapshots, once the backup is complete, the snapshot LUs are removed from the target and the VSS initiated snapshot(s) is deleted.

Persistent snapshots should be deleted through the backup application whenever possible. Although it is possible to delete a VSS initiated snapshot via the CLI or Web Manager, care must be taken to ensure that a backup application is not active and an iSCSI host is not bound to the snapshot's LU(s). Properly deleting a snapshot will also result in the snapshot LUs being removed from the target.

Note: Using the CLI or Web Manager to delete VSS initiated snapshots, or to remove the snapshot LUs from their associated iSCSI target, will result in the unexpected removal of a disk from the VSS host system, and can cause the VSS host to crash.

VSS Restrictions

- VSS is not supported on iSCSI LUs formatted as dynamic disks, only basic disks are supported.

- Quick Snapshot Restore of iSCSI LUs used by VSS is not supported.
- VSS initiated snapshots may not be backward compatible between major firmware releases. For example, snapshots taken on a NAS server running firmware version SU 6.x cannot be accessed by the VSS host if the NAS server is returned to firmware version SU 5.x. For information on backward compatibility of VSS initiated snapshots between releases, contact your technical support representative.

Setting up the NAS Server for VSS Snapshots

Setting up the NAS server to support VSS snapshots requires several steps:

1. Configure each NAS server for VSS access.
2. Install the VSS Hardware Provider software.
3. Specify the NAS server connection information.

Configuring VSS Access to a Server

You can configure a storage server to allow VSS access using Web Manager or the CLI command `mscfg`.

To configure the storage server using Web Manager:

1. **Navigate to the VSS Access Configuration page.**

   From the **Server Settings** page, click **VSS Access Configuration** to display the page:

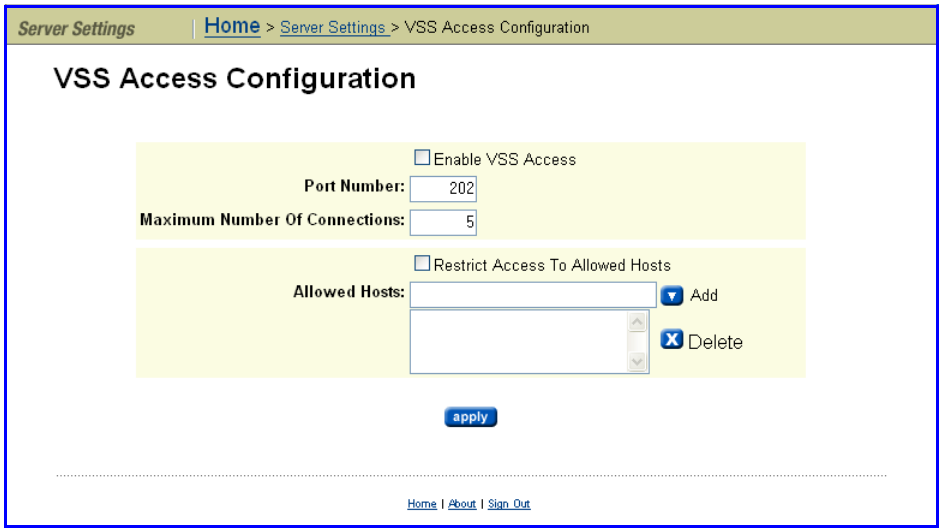

2. **Specify VSS access configuration settings.**

   Using the **VSS Access Configuration** page, enter the required information, as described in the following table:

Item/Field	Description
Enable VSS Access checkbox	Fill the checkbox to allow access by VSS or leave the checkbox empty to disable VSS access.
Port number	Enter the port number that the storage server will monitor for VSS communications. Port 202 is the default VSS access port.
Maximum number of connections	Specifies the maximum number of simultaneous VSS connections to the NAS server. You can allow up to 5 simultaneous connections.
Restrict Access to Allowed Hosts	Fill the checkbox to restrict VSS access to the hosts specified on this page. Leave the checkbox empty to enable VSS access from any host.
Allowed Hosts	If VSS access is restricted to specified hosts, use these fields to specify the hosts that are allowed VSS access. • **Allowed Hosts** (field). In the **Allowed Hosts** field, enter the IP address or DNS name of a host that is allowed VSS access, then click **Add** to insert that host into the list of allowed hosts. You can specify an IP address range using the * character as a wildcard. For example: 10.168.*.* or 172.*.*.*.  **Note:** If the NAS server has been set up to work with a name server, you can identify allowed hosts by IP address or DNS name. • **Allowed Hosts** (list). This list displays the IP address or DNS name of the hosts allowed VSS access. To delete a host, select it's IP address or DNS name from the list and click **Delete.**
apply	Click **apply** to save the VSS access configuration settings.

1. **Save the VSS access configuration.**

   Once you have entered and verified all the VSS access configuration settings, click **apply.**

Installing the VSS Hardware Provider

The VSS Hardware Provider is a DLL that must be installed on a host computer.

Installation Requirements

The VSS Hardware Provider software has the following requirements:

- The NAS server: Firmware version 5.1 or later.
- The VSS host:

- Windows Server 2003 with SP2 or later (32-bit or 64-bit version) or Windows Server 2008 (32-bit or 64-bit version).
- 16 MB of free disk space.

Installation Process

During installation, the installation program automatically installs the correct 32-bit or 64-bit executable for the operating system, and the installation program also installs:

- The Manage NAS Server Connections utility, which allows you to specify each NAS server or cluster that you want to be able to access using VSS. A shortcut is placed in the Start menu for easy access to the utility.
- Microsoft Visual Studio 2005 SP1 Redistributable runtime library, which is used by the VSS Hardware Provider. After installation this library is listed in the Add or Remove Programs Control Panel as Microsoft Visual C++ 2005 Redistributable. You can display the version number in the Control Panel by highlighting the application and clicking the "Click here for support information" link.

On the last dialog of the installation program, select **Manage** IS-NAS Server/ Titan Server **connections for VSS provider** to start the Manage NAS Server Connections utility. Note that you can choose to configure your server connections at a later time, and access the utility through the Start menu (under Titan Server/IS-NAS Server VSS Hardware Provider). If you receive a message prompting you to restart your computer to complete the installation, you must reboot before using the Manage NAS Server Connections utility.

Note: If a previous version of the Titan Server/IS-NAS Server VSS Hardware Provider is already present on the VSS host, we recommend that you uninstall it before installing the new version. To install a new version or uninstall a previous version, there must not be a connection open between the VSS Hardware Provider and the NAS platform. You can uninstall the application through the Start menu.

Specifying NAS Server Connections

Server connections are specified using the **Manage NAS Server Connections** utility, shown below. You can start this utility during the VSS Hardware Provider installation (by selecting **Manage** IS-NAS Server/Titan Server

connections for VSS provider on the last dialog of the installation program) or after installation through the Start menu.

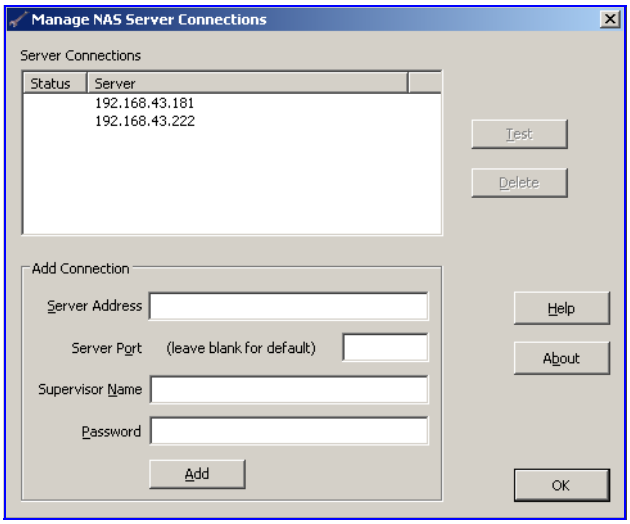

This utility allows you to specify the VSS connection information for each NAS server or cluster. The utility also displays any NAS server connections that have been configured.

To configure a VSS connection for a NAS server, specify the information in the Add Connection area (Server Address, Server Port, Supervisor Name, and Password) and click **Add**. This creates a unique VSS credential (discussed in About VSS Credentials, on page 325), which is saved on the target NAS server and on the VSS host. Once added, the NAS server's IP address or DNS name appears in the **Server** column of the Server Connections list in this dialog. The **Status** column lists the status of the VSS credential.

The **Manage NAS Server Connections** dialog contains the following fields and buttons:

Item/Field	Description
Server Connections	This table lists the DNS name or IP address of all configured NAS servers, along with their status. This table has two columns: Status and Server. The **Server** column lists the DNS name or IP address of all NAS servers that have had their connection information specified through this utility (that is, NAS servers for which this host has a VSS credential). If a non-default port number was used when the VSS credential was created, then that number is shown following the server's DNS name or IP address. The **Status** column lists the status of the VSS credential. A status is provided only after the server has been tested (by selecting the server and clicking the **Test** button). The possible status values are blank (no test run), OK, or Fail. A status of "Fail" indicates that the VSS host cannot connect to the NAS server. If this occurs, make sure your NAS server is running and that you can PING the server. You can also use the mscfg vss and vss-account CLI commands to ensure that VSS is enabled, and that the NAS server's copy of the credential has not been removed.
Server Address	Specify either the IP address or the DNS name of the NAS server or cluster.
Server Port	Leave blank to use the VSS default port (202). If the server has been configured to use a non-default VSS management port, specify that port number.
Supervisor Name	Specify the name of a management account with supervisor privileges on the NAS server. (The supervisor name is not saved by the Titan Server/IS-NAS Server VSS Hardware Provider.)
Password	Specify the password of the Supervisor Name provided. (The password is not saved by the Titan Server/IS-NAS Server VSS Hardware Provider.)
Test	Verifies that the selected server and its VSS credential are still valid. The test establishes a connection to the NAS server's "VSS management server" and sends a loopback message to verify functionality. The test returns either OK or Fail, which is displayed in the Status column.
Delete	Removes the selected NAS server's credential. The credential is removed from the VSS host, and, if server connectivity is possible, the credential is also removed from the NAS server.
Help	Displays help information.
About	Displays version information.

Item/Field	Description
Add	After filling in the fields in the Add Connection area, clicking **Add** creates a unique VSS credential for the NAS server. The credential is saved on the NAS server and on the VSS host (the system running the Titan Server/IS-NAS Server VSS Hardware Provider).
OK	After adding one or more connections, clicking **OK** closes the dialog. While entering information in the Add Connection area, clicking **OK** steps you through the fields, and pressing Escape on the keyboard closes the dialog.

About VSS Credentials

A VSS credential is saved on both the VSS host and the NAS server. The server address and port number are saved along with the credential. This means that if the NAS server's VSS management port setting is changed, any existing VSS credentials for that server must be removed and new credentials must be created. If a DNS name is used for a NAS server, then changes to the server's IP address alone will not require removing and recreating the credential.

A VSS credential has limited rights on the server: it can only be used to perform VSS-related operations using the VSS management interface. In particular it cannot be used to gain access to the normal NAS server management console, either locally or remotely.

8 Data Protection

The IS-NAS Server and the Titan Server architecture provide data protection across multiple levels, including:

Data Protection Component	Conceptual Overview	Associated Tasks
Hardware-based file system consistency	Hardware-Based File System Consistency, on page 327	Automatic, hardware-based functionality.
Snapshots	Snapshots, on page 329 Snapshots and the Volume Shadow Copy Service (VSS), on page 330	Using Snapshots, on page 342
NDMP support	NDMP Support, on page 331	Using NDMP Backups, on page 348
Data replication	Data Replication, on page 333	Using Data Replication, on page 362
Primary access transfer	Transfer of Primary Access, on page 337	Transferring Primary Access, on page 389
Virus scanning	Virus Scanning, on page 341	Using Virus Scanning, on page 392

The system administrator can configure any of these components to protect all data on the system.

Hardware-Based File System Consistency

Checkpoints and NVRAM Buffering

To guarantee file system consistency, the system periodically writes complete and consistent *checkpoints* (file system images) to the storage subsystem. This process ensures that file system metadata is always internally consistent, even after a system failure. In the event of a system failure, the file system can be rolled back to the last successful checkpoint, thus ensuring that file system consistency is never lost. The system uses NVRAM to buffer any file system modifications in process during checkpoint writes (with acknowledgement of the operation returned to the client only after all resulting file system modifications have been successfully buffered in NVRAM and, if in a cluster, mirrored).

Once modification requests have been acknowledged by the server (to the client) they cannot be lost. Furthermore, every request to modify the file system is buffered in NVRAM until the checkpoint containing the modification has been successfully written to the storage subsystem. Buffering requests in local NVRAM ensures that software failure or power failure will not result in data loss. Additionally, nodes in a server cluster will mirror the contents of their NVRAM to a partner node so that even a single

server hardware failure will not result in data loss. In the event of a power failure, the contents of NVRAM are preserved (using dedicated battery backup) for a maximum of 72 hours. Once power is restored, the original server can roll back all its file systems and replay modifications stored in the NVRAM log.

Buffering in a Cluster Configuration

When configured as a cluster, the cluster buffers all file system modifications, and the NVRAM contents of each cluster node are mirrored on another cluster node. This process ensures data integrity in the event of a cluster node failure. When a cluster node takes over for a failed node, it plays back the contents of the NVRAM mirror to complete all file system modifications that were not yet committed to the storage by the failed server.

- In a two-node cluster configuration, each cluster node mirrors the NVRAM contents and buffers the file system modifications of the other cluster node. In a two-node configuration, NVRAM data and file system modifications are mirrored between the cluster nodes as shown in this diagram:

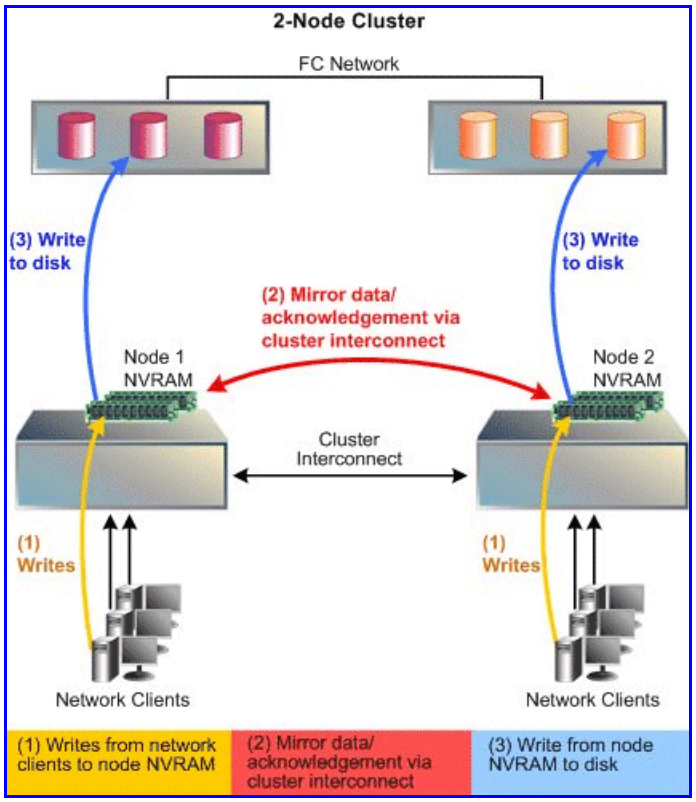

- In a cluster configuration with more than two nodes, every cluster node mirrors the NVRAM contents and buffers the file system modifications to another cluster node. Mirroring and buffering are done in a round-robin fashion, with node 1 being mirrored on node 2, node 2 being mirrored on node 3, etc., until the last node, which is mirrored on node 1. In a cluster configuration with four nodes, NVRAM data and file system

modifications are mirrored between cluster nodes as shown in the following diagram:

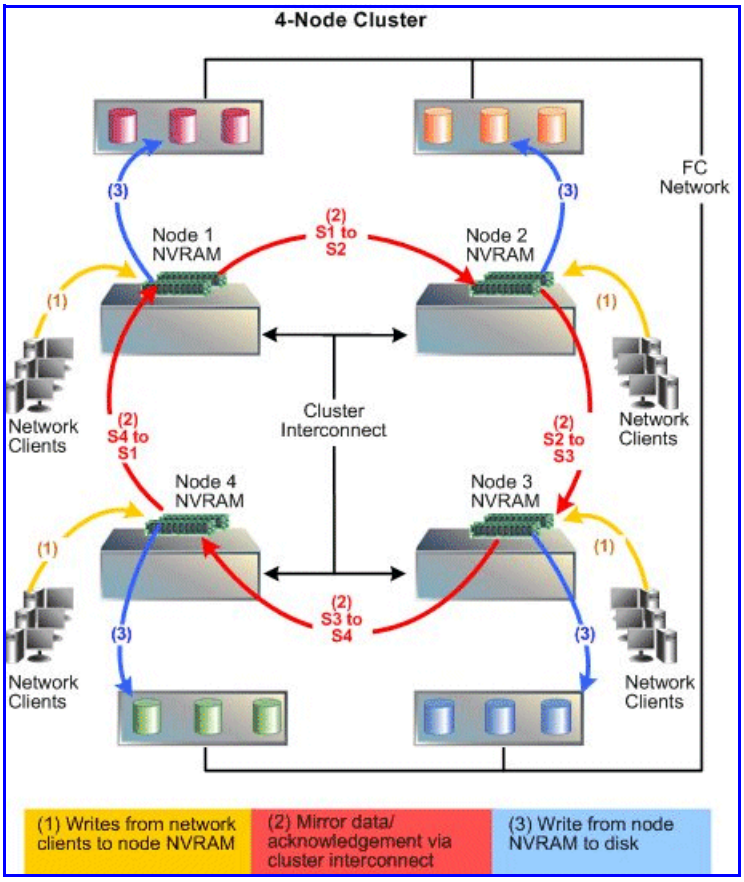

NVRAM Statistics

Statistics for NVRAM activity on the server (in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Snapshots

For users whose data availability cannot be disrupted by management functions such as system backup and data recovery, snapshots create near-instantaneous, read-only images of an entire file system at a specific point in time. Snapshots safely create backups from a running system, allowing users to easily restore lost files without having to retrieve the data from backup media, such as tape.

Snapshots capture a moment in time for a live file system. They contain only those blocks that have changed since the snapshot was created, such that the disk space occupied by a snapshot is a fraction of that used by the original file system. However, over time, the space occupied by a snapshot grows, as the live file system continues to change.

Snapshots solve the problem of maintaining consistency within a backup; specifically, during a system backup, users continue to modify its component files, resulting in backup copies that may not provide a consistent set. Since a snapshot provides a frozen image of the file system, a backup copy of a snapshot (rather than of the live file system) provides a usable, consistent backup that appears to a network user like a directory tree. Users with appropriate access rights can retrieve the files and directories that it contains through CIFS, NFS, FTP, or NDMP.

Snapshots and the Volume Shadow Copy Service (VSS)

Snapshots of storage attached to the storage server may be initiated by Microsoft's Volume Shadow Copy Service (VSS). VSS is available on servers running Windows Server 2003 or 2008, and it provides a coordination point for enabling consistent backups of online storage. Snapshots initiated by VSS are exported as iSCSI LUNs.

Storage writers (for example MS Exchange or a backup application) first register with VSS. Then, when a backup application wishes to back up a piece of storage (a "volume"):

1. The backup application requests that VSS take the snapshot.
2. VSS requests that all registered writers flush their data to make sure that all of their on-disk data files are in a consistent state.
3. Once the writers report completion of this step, VSS takes the snapshot.
4. VSS then returns a pointer to the snapshot to the backup application so that the backup application can back up a stable view of the storage (the snapshot).
5. Once the backup is completed, the backup application notifies VSS so that the snapshot may be deleted.

VSS may also be used to take "point in time" copies for later reference. The process is similar, except in this case no automatic deletion of the snapshot is performed by VSS. The storage server supports this mechanism by means of a VSS "hardware provider," a DLL which registers with VSS in order to support snapshots of volumes attached to a storage server.

Note: Snapshots initiated by the VSS service only contain images of iSCSI LUNs attached to the storage server. Non-iSCSI volumes attached to the storage server are not included in snapshots initiated by VSS.

For information about configuring VSS Access to the storage server, see Accessing Snapshots Initiated by VSS, on page 319.

Latest Snapshot

The storage server provides a file system view that can be used to access the *latest snapshot* for a file system. This view automatically changes as new snapshots are taken, but is not affected by changes in the live file system. The latest snapshot is the most recent snapshot for the file system, and is accessible

through `.snapshot/.latest` (or `~snapshot/.latest`). The latest snapshot can be exported to NFS clients with the path `/.snapshot/latest`. Latest snapshots can also be shared to CIFS clients. When accessing files via the latest snapshot, NFS operations do not use autoinquiry or autoresponse.

Note: The `.latest` (`~latest`) file designation is a hidden snapshot directory and does not show up in directory listings.

Quick Snapshot Restore

Quick Snapshot Restore is a licensed feature for rolling back one or more files to a previous version of a Snapshot. For more information about this command line procedure, open the CLI and run `man snapshot`, or refer to the *Command Line Reference.*

If a file has been moved, renamed or hard linked since the snapshot was taken, Quick Snapshot Restore may report that the file cannot be restored. If the file cannot be Quick Restored, it must be copied from the Snapshot to the live file system normally.

Accessing Snapshots Through NFS Exports and CIFS Shares

NFS exports and CIFS shares can easily access Snapshots, so that users can restore older versions of files without intervention.

- The root directory in any NFS export contains a `.snapshot` directory which, in turn, contains directory trees for each of the snapshots. Each of these directory trees consists of a *frozen* image of the files that were accessible from the export at the time the snapshot was taken (access privileges for these files are preserved intact).
- Similarly, the top-level folder in any CIFS share contains a `~snapshot` folder with similar characteristics. Both with NFS and with CIFS, each directory accessible from the export (share) also contains a hidden `.snapshot` (`~snapshot`) directory which, in turn, contains *frozen* images of that directory. A global setting can be used to hide `.snapshot` and `~snapshot` from NFS and CIFS clients.

Note: Backing up or copying all files at the root of an NFS export or a CIFS share can have the undesired effect of backing up multiple copies of the directory tree (that is, current file contents plus images preserved by the snapshots; for example, a 10GB directory tree with four snapshots would take up approximately 50GB).

Administrators can control access to snapshot images by disabling Snapshot access for specific NFS exports and CIFS shares. For example, by creating one set of shares for users with snapshots disabled, and a second set of shares with restricted privileges (for administrator access to Snapshot images).

NDMP Support

The storage server supports Network Data Management Protocol (NDMP), an open standard protocol for network-based backups, with two significant advantages:

- It enables a storage management application to control backup and recovery on another device without transfer of the backup data across the network.
- NDMP backups can preserve security settings in a mixed protocol environment, including virtual volume and quota information.

A standard NDMP configuration is shown in the following diagram:

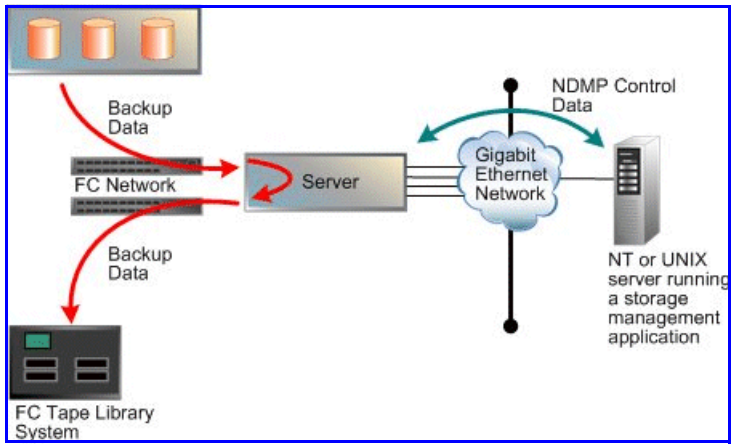

In the diagram, the storage management application sends backup instructions to the server, which makes a backup copy of data onto tapes in the tape library. The data travels through the Fibre Channel (FC) network, not the Ethernet network. Details of the backup data are sent to the storage management application, which initiates recovery of the data if necessary.

NDMP transfers data between disks and tapes attached to the same server. Data can also be transferred between two separate NDMP servers over an Ethernet connection (in NDMP this is known as a 3-way backup or recovery):

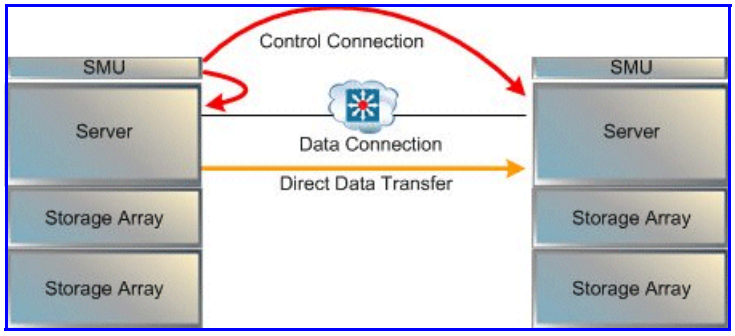

Some common applications of NDMP include:

- Backing up (or recovering) data on a server to (or from) a Fibre Channel attached NDMP tape library.
- Backing up (or recovering) data on a server without a tape library to (from) a second storage server that has a tape library attached.
- Using a utility, such as Accelerated Data Copy (ADC) or Data Replication to copy file systems between storage servers.

While the server supports backups done over network protocols such as NFS or CIFS, only NDMP will preserve security settings in a mixed protocol environment, including virtual volume and quota information.

When using NDMP, the server uses snapshots to backup data consistently and without being affected by ongoing file activity. Snapshots also facilitate incremental backups. However, if so desired, data can be backed up without using snapshots.

Storage Management Applications

The storage server acts as an NDMP server, operating with leading storage management applications. It supports NDMP Version 2, 3 and 4. The storage server implementation of NDMP can back up and restore:

- Both Windows and UNIX files from a single storage management application.
- The full attributes of each Windows and UNIX file (including Windows ACLs), saving and restoring whole volumes and preserving all file attributes.

The server supports recovery of single files or subdirectories, associated lists, or complete backup images. The Direct Access Recovery (DAR) mechanism can be used, provided the Storage Management Application supports it. DAR allows NDMP to go directly to the correct place in the tape image to find the data, rather than reading the whole image. This can dramatically reduce recovery times.

Data Replication

Data replication allows you to copy or relocate both file data and file system metadata. Storage servers provide manual and automatic mechanisms for data replication, supporting the replication of data and, when using the transfer of primary access feature, also supporting the replication of file system settings. When using replication with the transfer of primary access feature, you can relocate file system data and CNS links, CIFS shares, permissions and all other file-level metadata (for more information about the transfer of primary access, see Transfer of Primary Access, on page 337). Administrators can use Web Manager to configure policy-based replication jobs independently from other backup strategies.

Policy Based Replication

Policy-based replication comprises:

- **Replication Policy:** A replication policy identifies the data source, the replication target, and optionally a replication rule. Pre-replication and post-replication scripts can also be set up in the **Policy** page.
- **Replication Rules:** Optional configuration parameters that allow tuning of replications to enable/disable specific functions or to optimize performance.
- **Replication Schedule:** Defines all aspects of automated timing.

Incremental Replication

Storage servers can also perform incremental data replication, which works as follows:

- Upon establishing a replication policy, the SMU performs an *initial copy* of the source volume (or directory) to a target destination.
- Once a successful *initial copy* has occurred, the system performs *incremental copies* (replications) at scheduled intervals. During an *incremental data replication,* the system copies to the target, in full, those files that have been changed since the last scheduled replication.
- To replicate large files more efficiently, the server also supports *incremental block level replication*. With incremental block-level replication, only the changes in files are replicated and not the whole file, thus reducing the amount of data replicated and the overall replication time. Note that, in order to use block-level replication, a Replication license is required.

A replication policy defines the properties of a replication, including a *replication rule* (source and target), and a *replication schedule*. Replication rules can be expanded to include optional settings. Pre-replication and post-replication scripts can also be configured.

Incremental Data (File Level) Replication

The server supports incremental data replication, performed under control of the System Management Unit (SMU). Incremental replication means that, after the initial copy, only changes in the source volume or directory are actually replicated on the target. Snapshots ensure the consistency of the replication.

Note: If the snapshot that was copied by the last successful replication copy is deleted, an incremental copy cannot be performed, so the full data set is replicated.

Incremental data replication uses the same data management engine as NDMP to copy:

- The contents of an entire file system,
- A virtual volume, or
- An individual directory tree to a replication target.

Upon configuration of a replication policy and schedule, the incremental data replication process takes place automatically at the specified interval. The replicated data can be left in place (and used as a standby data repository). In addition, the replicated file system or directory can be backed up through NDMP to a tape library system for long-term storage (which can also be automated).

Incremental data replication supports the following targets for replication:

- A file system or directory within the same server.

  Tiered storage technology ensures that replications taking place within a server are performed efficiently, without tying up network resources.

- A file system, virtual volume, or directory on another server.
- A file system, virtual volume, or directory on another Server model.

Although the SMU schedules and starts all replications, replicated data flows directly from source to target without passing through the SMU.

Incremental Block-Level Replication

By default, *incremental data replication* copies files that have changed since the last replication. With the *Block-Level Replication* feature enabled, only data blocks in large files that have been written since the last replication are copied. Depending on the use of files within the source volume, this could substantially reduce the amount of data copied.

Note: Block-Level Replication copies the entire file if the file has multiple hard links.

Note: The Block-Level Replication feature is automatically enabled if the Replication license is present.

Multiple Stream Replication

Multiple replication streams are created by adding TCP connections between the source and target systems of a replication or ADC copy operation. Each additional connection is used for an additional data stream by the replication application.

Multi-stream replication helps to alleviate some latency problems found with single-stream replication by running multiple independent streams in parallel. When latency from sequentially executed functions limits performance, multiple independent streams can produce a significant performance improvement.

Multi-stream replication should also alleviate performance problems caused by high speed WAN connections with high latencies. Connections with high latencies limit the throughput of a single TCP connection, because no data is sent during the time spent waiting for acknowledgements. These pauses in the sending of data result in an under-utilization of high speed WAN links. By using multiple TCP connections (one per stream), multi-stream replication addresses the problem of under utilization of the high speed WAN connections.

Multi-stream replication is only supported if both the source and destination systems are using software Release 6.1 or later.

For policy-based replication operations, multi-stream replication is controlled via the replication **Add Rule** or **Modify Rule** pages of Web Manager (see Adding a Replication Rule, on page 370 for more information).

For ADC copies, multi-stream support is enabled by setting the number of additional connections requested as the value of the environment variable

NDMP_BLUEARC_MULTI_CONNECTION (see NDMP_BLUEARC_MULTI_CONNECTION, on page 554 for more information).

Note the following:

- When using software release 6.1 or later, and using multi-stream replication or embedded inline hard linked files, if a replication fails part way through, it will not be possible to restart replication if the server is downgraded to an earlier release. See NDMP_BLUEARC_EMBEDDED_HARDLINKS, on page 550 for more information about NDMP support for embedded hard links.
- Multi-stream replication features are not enabled using the ndmp-option CLI command; instead, the invoking NDMP command must request multiple streams.
  - For policy based replications the multi-stream feature is configured via replication rules (see Adding a Replication Rule, on page 370 or Modifying a Replication Rule, on page 374 for more information.
  - For individual ADC copies, multiple streams are specified by adding the NDMP_BLUEARC_MULTI_CONNECTION environment variable to the ADC script (see NDMP_BLUEARC_MULTI_CONNECTION, on page 554for more information).
- NDMP has two ways of copying data from files with hard links. The NDMP_BLUEARC_EMBEDDED_HARDLINKS environment variable controls this behavior (see NDMP_BLUEARC_EMBEDDED_HARDLINKS, on page 550) for more information on this variable.

Note: When multiple connections/streams are used, the data from files with hard links is embedded within the hierarchical path data, regardless of the setting of the NDMP_BLUEARC_EMBEDDED_HARDLINKS variable.

Relocating File Systems

Storage servers support relocation of file systems, or parts of file systems, *including both file system data and file system metadata* from one server to another. Metadata refers, for example, to CNS links, CIFS shares, NFS mount points, FTP users, Snapshot rules, backup files, and other file system-level settings.

Note: Unlike other file system metadata, iSCSI configuration settings remain with the original EVS, as an iSCSI target may contain Logical Units from multiple file systems. In this instance, **Relocation** page displays a message, reminding the Administrator to properly configure iSCSI Domains, Targets, iSNS, and Initiator Authentications on the new EVS.

Allowable destinations for a relocation may be:

- Another EVS on the same cluster node,
- Another node in the cluster, or
- An EVS on another server or cluster.

The following list includes some examples of file system relocations:

- Moving data to a new storage system.
- Dividing a single large file system into several smaller file systems within a Storage Pool.
- Load balancing, by moving data from one file system to another, or by moving a file system from one EVS to another.
- Moving an EVS (and all its file systems) to another server to gain access to other storage devices or to change the structure of the data.

From a high level, relocating file systems requires two steps:

1. Replicate online data while the system is live and in normal use. This may require several incremental replications, to synchronize the data on the source and the target as much as possible. Synchronizing the data shortens the amount of time required for the next step.

2. Perform a final replication with source data (file system) in *Syslocked mode*. When in Syslocked mode, the data is write-protected, so the data can be accessed and read, but data cannot be changed or added. At the end of this stage, the target is brought online in place of the source.

Transfer of Primary Access

A transfer of primary access copies data from a portion of a file system and relocates the access points for that data, or relocates an entire file system and its access points (copying the data and metadata), with very little down time, while the file system is live and servicing file read requests. For a short period, access is limited to read-only.

A transfer of primary access can be performed on any replication policy as long as the following conditions are met:

- A full replication has completed. Preferably an incremental replication should also have completed.
- The snapshot required to support another incremental replication must still be available.

The Process of Transferring Primary Access

A single transfer of primary access operation may be in progress at any time for any given replication policy, and the process for the transfer of primary access is as follows:

1. **Place the source file system into "Syslock" mode, allowing read-only access to the file system, but no write access.**

   The storage server ensures that the target file system data is consistent with the source file system data before primary access is transferred. This involves making the source read-only for a short time. Although any

arbitrary directory can be relocated, the entire source file system is set to syslocked mode while the final replication operation is in progress.

Note: Clients should be notified that a short period of read-only access will be necessary while data and file system settings are relocated.

2. **Replicate the data and file system settings to the new location.**

   Once a transfer of primary access has been started, the SMU monitors the replication to determine when it is complete. When the replication is complete, the SMU starts moving configuration information automatically. The following table describes how network access points on the source file system are moved or deleted:

Source File System Setting/Network Access Point Being Moved	Destination		
	Within the EVS	**Another EVS in the Same Cluster**	**An EVS on Another Server or Cluster**
CIFS Shares (if within replicated path)	Moved (path is modified). Clients that had the share mounted before the transfer of primary access do not have to remount the share after the transfer.	Moved (deleted from source EVS then added on target EVS). Clients that had the share mounted before the transfer of primary access must remount the share after the transfer only if the share was not to a directory in the CNS.	Moved (added to target EVS then deleted from source EVS). Clients that had the share mounted before the transfer of primary access must remount the share after the transfer.
NFS Exports (if within replicated path)	Moved (path is modified). **Note:** Clients that had the export mounted before the transfer of primary access must mount the export again after the transfer.	Moved (deleted from source EVS then added on target EVS).	Moved (added to target EVS then deleted from source EVS).
FTP Initial Directories/ Users (if within replicated path)	If all users within an initial directory can be moved, the initial directory is also moved. If all users within an initial directory cannot be moved, no users are moved and the initial directory is not moved.	If all users within an initial directory can be moved, the initial directory is also moved. If all users within an initial directory cannot be moved, users are moved where possible, and the initial directory is duplicated.	
Snapshot Rules	Not moved if replicating only part of a file system. Moved only if file system will be a standalone file system when replicating data and access points to root (/), meaning that the target file system will be a standalone file system when the transfer of primary access is complete.		

Source File System Setting/Network Access Point Being Moved	Destination		
	Within the EVS	**Another EVS in the Same Cluster**	**An EVS on Another Server or Cluster**
CNS Links	If CNS entries already point to the replication source, then the CNS link is removed and a link to the new file system is added at the corresponding path. Note, however, that if the file system is linked to the cluster name space at a point higher in the directory structure than the root directory for the file system path being replicated, moving the CNS link is not possible. In such cases, the CNS link is reported as an error in the list of successful/failed transfers, and the administrator must manually create a CNS link to the file system in the new location. After a transfer of primary access, network clients will not be able to access the file system through the a CNS name space if any of the following are true: • The file system did not have CNS links. • The file system's CNS links were not moved. • The file system was replicated to another server or cluster. To access to the file system in its new location, network clients must reconnect through CIFS shares or NFS exports pointing to the relocated file system or to a CNS name space into which the file system is linked. NFS clients pointing to a CNS name space will not experience any interruption.  **Note:** If clients will not access the relocated file system using CNS links, they must access it using new IP addresses.		If CNS links exist, the relocation is not allowed to proceed, and a message advises the administrator to remove the links before proceeding.
iSCSI Targets	Not moved.		
Global Symlinks	Not moved.		

Note: For CIFS shares and NFS exports: if possible, a text file backup of the moved shares and export will be left on the SMU.

3. **After replicating the data and file system settings, bring the target file system online.**

   The system administrator receives instructions to bring the target file system on-line, by allowing read/write access. Read/write access is re-enabled on the entire source file system unless it was syslocked originally).

   The SMU tracks/records the progress of the final replication. Status of the network access point relocation is available through the **Status and Reports** page; replication failures are logged and can be viewed by following a link from the Replication Report.

4. **Begin servicing file service requests from the relocated file system.**

5. **Unless it was offline when the transfer of primary access was started, take the source file system out of syslocked mode, putting it back online.**

Note: If the SMU is rebooted during a Transfer of Primary Access, the source file system may not be returned to its original online state. If the SMU is rebooted during a Transfer of Primary Access, you may have to take the file system out of syslock manually, from the **File System Details** page. For more information on Syslocked mode, see Using System Lock on File Systems, on page 156.

After the final replication has completed, the original source data is still present on the source. This data can be accessed (and modified) through access points configured higher up in the directory tree, and should be deleted manually.

Note: After the successful completion of a transfer of primary access, the original server data should be removed or made inaccessible.

All replication schedules configured for the replication policy are set to inactive once the transfer of primary access is completed and these inactive policies should then be deleted manually.

How a Transfer of Primary Access Moves CNS Links

Using the diagram below as a sample file system:

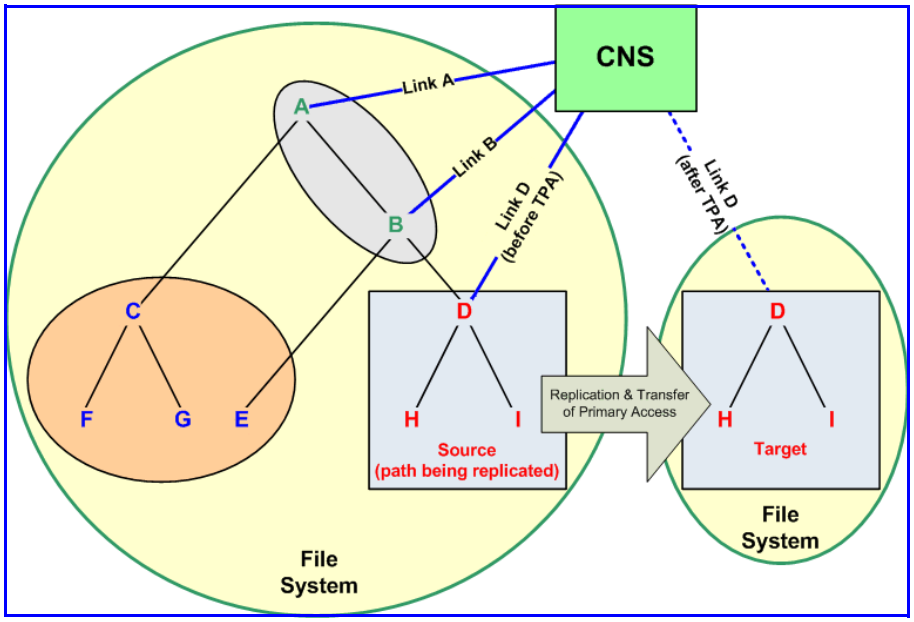

When replicating the file system path beginning at "D," CNS links are transferred as follows:

- If the file system is linked to the CNS tree at "A" or "B," the CNS link is not moved and is listed in the replication report as an error. The CNS link is not moved because doing so would deny access to the file system at point "E."

Caution: In this situation, users will be able to access the "old" data after the replication and transfer of primary access are complete. After a successful transfer of primary access, the original source data should either be removed or made inaccessible by network clients (permissions should be changed).

- If the replication is within the cluster, and the file system is linked to the CNS tree at "D" or below, then the CNS link is moved and is listed in the replication report as having been successfully moved.
- If the file system is linked to the CNS tree at "C" or "E," the CNS link is not moved and it is not listed in the replication report, because it is not relevant to the path being replicated.

Virus Scanning

The storage server architecture reduces the effect of a virus because the file system is hardware-based. This prevents viruses from attaching themselves to (or deleting) system files required for server operation. However, viruses can still propagate and infect user data files that are stored on the server. Therefore, Silicon Graphics International Corporation works with industry leading anti-virus (AV) software vendors to ensure that the server integrates into an organization's existing AV solutions and without requiring special installations of AV software and servers. To reduce the effect that a virus may have on user data, SGI Global Services recommends that AV be configured for the server and that AV software run on all user workstations.

The server itself does not perform any scanning of the files., but rather provides a connection with configured Virus Scan Engines on the network:

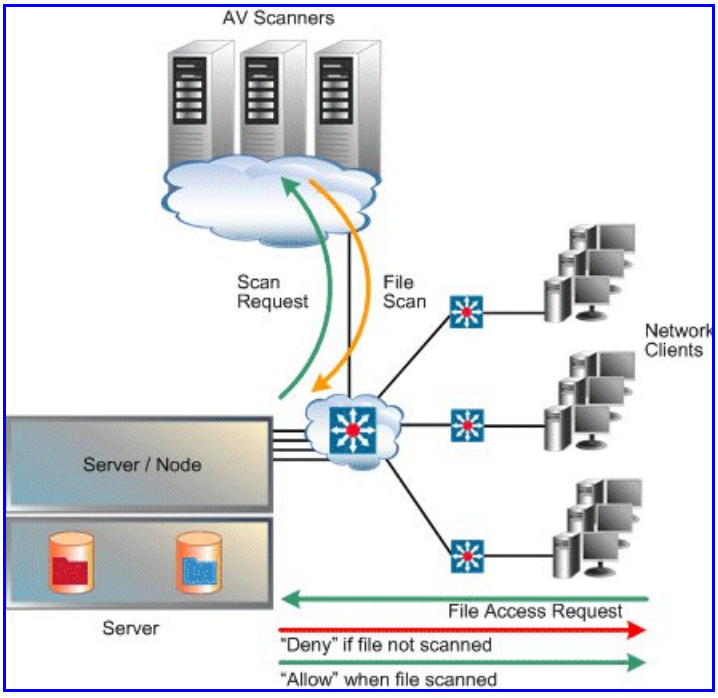

Virus Scanning is enabled and configured at the EVS level. Only files accessed using the CIFS protocol are scanned. If a file has not been verified clean by a Virus Scan Engine, it requires scanning before access. As virus scanning can cause delay when a client requires access, files are automatically queued for scanning as soon as they are closed (after creation or modification). Queued files are scanned promptly, expediting detection of viruses in new or modified files and making it unlikely that a virus infected file will remain dormant on the system for a long period of time.

When a virus is detected, a severe event is placed in the Event Log, identifying the path of the infected file and the IP address of the client machine that wrote the file. For information on accessing the event log, see Event Logging and Notification, on page 497.

Multiple Virus Scan Engines can be configured to enhance performance and high-availability of the server. If a Virus Scan Engine fails during a virus scan, the server automatically redirects the scan to another Virus Scan Engine.

Using Snapshots

Snapshots create near-instantaneous, read-only images of an entire file system at a specific point in time. A conceptual overview of Snapshots can be found at Snapshots, on page 329.

Managing Snapshot Rules

Snapshot rules define *scope* (that is, what file system), while Snapshot schedules define *frequency.* This section describes how to use Web Manager to create rules and schedules and to assign schedules to rules, in the following sections:

- Creating Snapshot Rules, on page 342
- Modifying Snapshot Rules, on page 345
- Deleting Snapshot Rules, on page 346
- Managing Individual Snapshots, on page 346

Note: This section does not cover setting up specific storage management applications or tape libraries. Consult the documentation that accompanies the application and tape library for setup instructions.

Creating Snapshot Rules

To create a snapshot rule:

1. **Navigate to the Snapshot Rules page.**

From the **Data Protection** page, select **Snapshot Rules** to display the page:

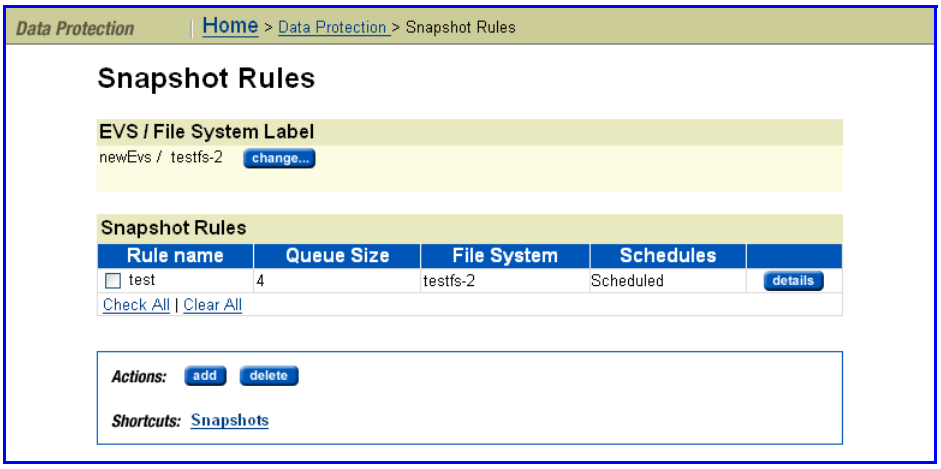

2. **Navigate to the Add Snapshot Rule page.**
   Click **add** to display the **Add Snapshot Rule** page:

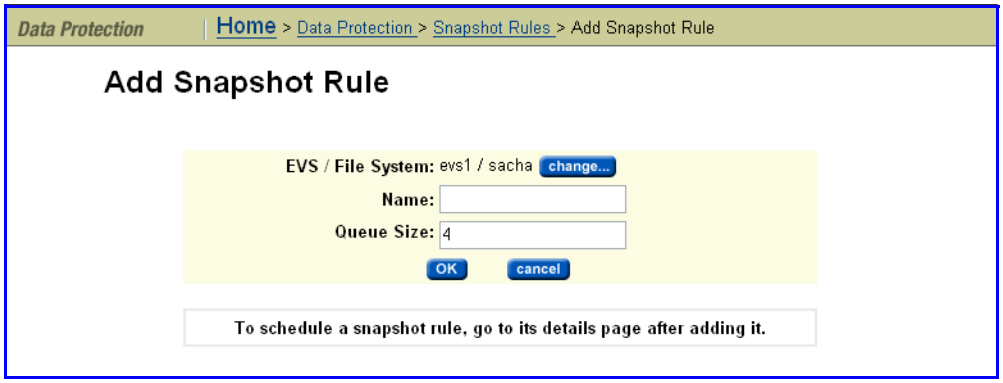

3. **Enter the required information.**
   Define the Snapshot Rule and select a file system:
   - Click the **change** button to select the file system.

- In the **Name** field, type a name for the rule (containing up to 30 characters). Do not include spaces or special characters in the name.

  The name of the rule determines the names of the snapshots that are generated with it. For example,

  `YYYY-MM-DD_HHMM[timezone information].rulename.`

  where *date and time* are expressed in the indicated format, *timezone information* is a placeholder for the offset from Greenwich Mean Time, and *rulename* is the name of the file.

  If more than one snapshot is generated per minute by a particular rule, the names will be suffixed with .a, .b, .c etc.

  For example, a rule with the name *frequent* generates snapshots called:

  `2002-06-17_1430+0100.frequent`

  `2002-06-17_1430+0100.frequent.a`

  `2002-06-17_1430+0100.frequent.b`

  and so on.

- In the **Queue Size** field, specify the number of snapshots to keep before the system automatically deletes the oldest snapshot. The maximum is 32 snapshots per rule.

Note: The system automatically deletes the oldest snapshot when the number of snapshots, associated to a snapshot rule, reaches the specified queue limit. However, any or all of the snapshots may be deleted at any time, and new snapshots can be taken.

4. **Assign a schedule.**

   a. **Select the rule to which you want to add a schedule, and click details.**

   Fill the checkbox next to the name of the rule to which you want to add a schedule, then click **details** to display the **Snapshot Rule Details** page.

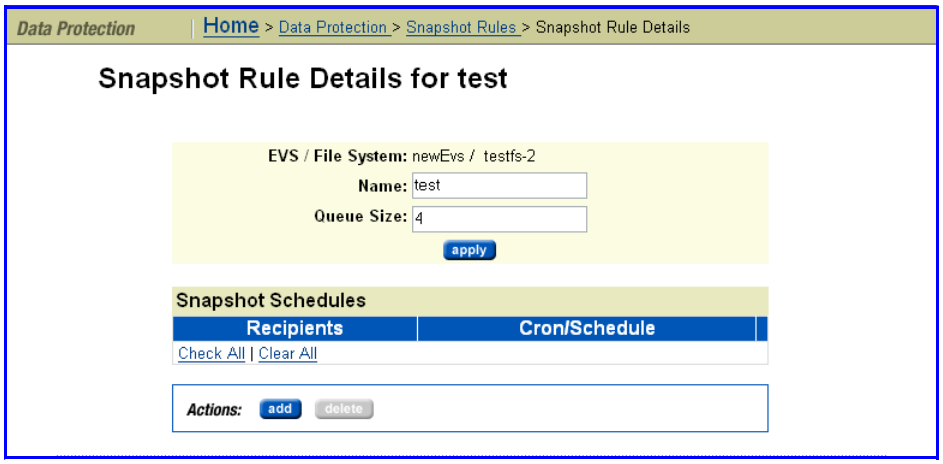

b. **Click add to display the Add Snapshot Schedule for rule page.**

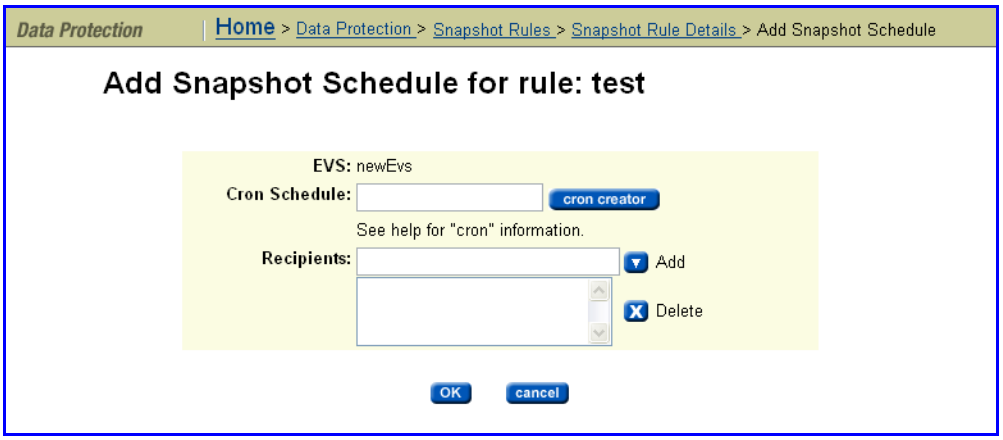

c. **Specify the schedule for the rule.**

You can click **cron creator** and build your schedule, or you can specify the schedule directly in the **Cron Schedule** field.

For more information on the cron syntax, refer to the UI help page and the `crontab` command in the *Command Line Reference*.

d. **Enter an email address to be notified upon completion of each snapshot.**

In the **Recipients** field, you can enter a single email address or multiple email addresses. Multiple addresses should be separated with a semicolon (**;**). SGI Global Services recommends sending Snapshot notifications to at least one user.

5. **Save your changes.**

Verify your settings, then click **OK** to save or **Cancel** to decline. You are returned to the **Snapshot Rules** page, which summarizes properties for the rule you just created.

Modifying Snapshot Rules

To modify a rule:

1. **Navigate to the Snapshot Rules page.**

From the **Data Protection** page, click **Snapshot Rules**. From this page you can perform any of the following tasks.

2. **Modify Rule properties.**

From the **Snapshot Rules** table, click the **details** button for a snapshot rule, which opens the **Snapshot Rule Details** page. As needed, modify the **Name** and **Queue Size** fields, then click **apply** to save and return to the **Snapshot Rules** page.

3. **Modify a Rule schedule.**

From the **Snapshot Rules** table, click the **details** button for a snapshot rule, which opens the **Snapshot Rule Details** page. Click the **details**

button for a snapshot schedule, which opens the **Snapshot Schedule Details** page. As needed, modify the **Cron Schedule** and **Recipients**, then click **OK** to save or **cancel** to decline.

4. **Delete a Rule Schedule.**

   From the **Snapshot Rules** table, click the **details** button for a snapshot rule, which opens the **Snapshot Rule Details** page. From the **Snapshot Schedules** table, select a snapshot schedule and click **delete**.

Deleting Snapshot Rules

To delete a rule:

1. **Navigate to the Snapshot Rules page.**

   From the **Data Protection** page, click **Snapshot Rules.**

2. **Delete a Rule.**

   From the **Snapshot Rules** table, select a Snapshot Rule, then click **delete**.

Managing Individual Snapshots

To manage individual Snapshots:

1. **Navigate to the Snapshots page.**

   From the **Data Protection** page, click **Snapshots** to display the **Snapshots** page:

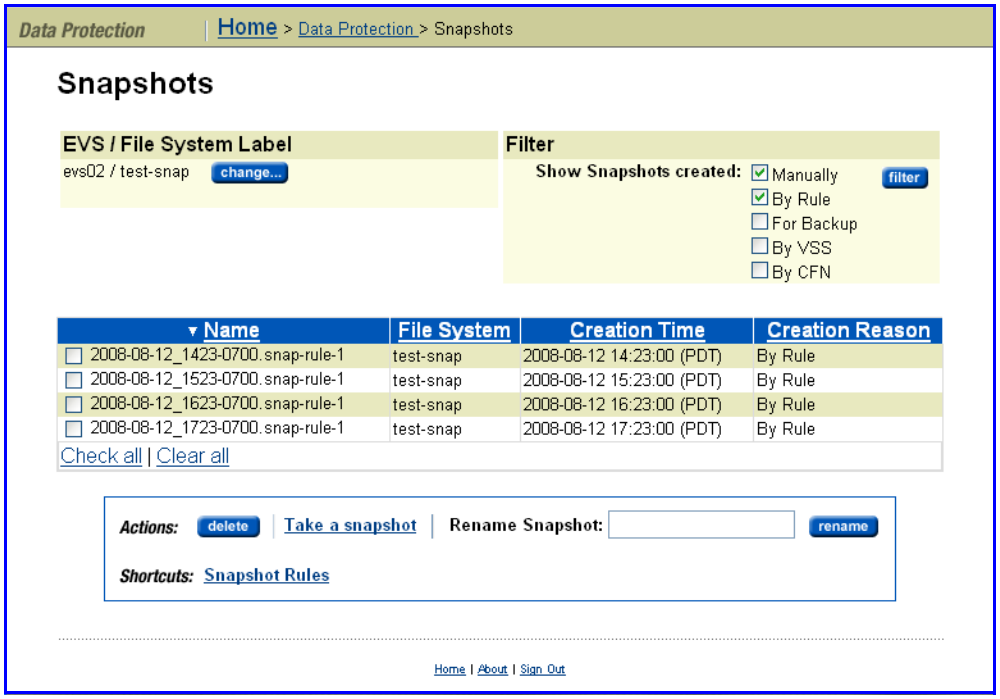

2. **Select a file system.**

In the **EVS/file system** section, click **change** to select a specific file system and display a list of snapshots.

3. **Filter the snapshots.**

   In the **filter** section, fill the appropriate checkbox(es) to filter the snapshots you want to display, then click **filter**. Snapshot filters allow you to limit which snapshots are displayed based on your selection of the reason(s) or mechanism(s) that can cause snapshots to be created. Select one or more of the following:

   - **Manually**, to display snapshots created manually.
   - **By Rule**, to display snapshots created by snapshot rules.
   - **For Backup**, to display snapshots as a part of the backup process.
   - **By VSS**, to display snapshots initiated by VSS (the Microsoft Volume Shadow Copy Service).
   - **By CFN**, to display snapshots created by the Changed File Notification feature.

4. **Manage the snapshots:**

   The following actions are available:

   - **Delete an individual snapshot,** by selecting it, then clicking **delete**.

     Note: Snapshots initiated by the Microsoft Volume Shadow Copy Service (VSS) should be managed through the application that requested the snapshot. You can, however, delete these snapshots through the **Snapshots** page.

   - **Delete all the snapshots,** by selecting **Check all**, then clicking **delete**.
   - **Rename an individual snapshot,** by selecting it, entering the new name in the **Rename Snapshot** text field, then clicking **rename.**
   - **Take a new snapshot** by clicking **Take a Snapshot** to display the **Take a Snapshot** page,

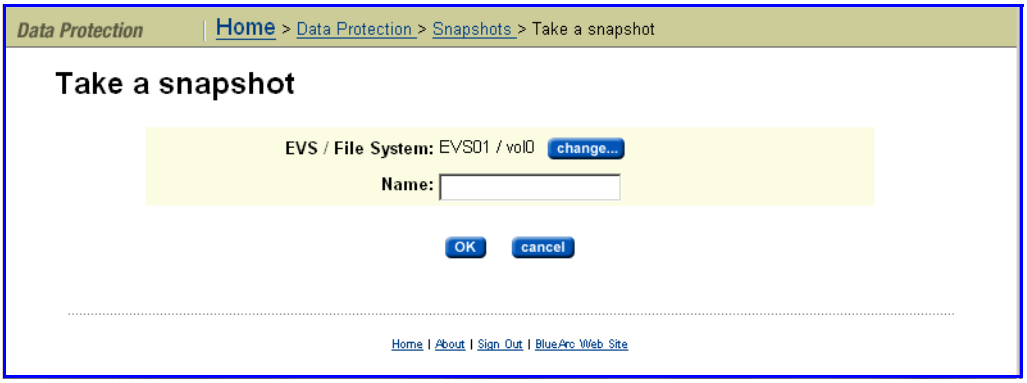

   then enter a **Name** for the snapshot (up to 30 characters, no spaces or special characters). Click **OK** to take the snapshot or **Cancel** to decline.

   Note: Users with permission can also take spontaneous rule-associated snapshot, without waiting for the next scheduled time. This can be

done from the command line interface.

Managing Snapshots Initiated by VSS

Snapshots initiated by the Microsoft Volume Shadow Copy Service (VSS) should be managed through the application that requested the snapshot. These snapshots may not be deleted by rule, but if necessary, you can delete these snapshots through the **Snapshots** page.

Using NDMP Backups

NDMP enables a storage management application to control backup and recovery on another device without transfer of the backup data across the network. A conceptual overview of the NDMP Protocol can be found at NDMP Support, on page 331.

Configuring NDMP

This section describes how to configure NDMP, under the following topics:

Note: This section does not cover setting up specific storage management applications or tape libraries. Consult the documentation that accompanies the application and tape library for setup instructions.

NDMP Version

By default, the storage server uses NDMP version 4 for NDMP backup and recovery; if required, it can be configured for version 2 or 3 of the NDMP protocol.

Note: Both incremental data replication and ADC require NDMP version 3 or 4. Set NDMP to version 2 only if required by your backup software.

Enabling and Disabling NDMP

NDMP processing status can be started or stopped at any time. NDMP can also be enabled or disabled to start on Boot.

To enable or disable NDMP:

1. **Navigate to the Backup Status page.**

From the **Data Protection** page, click **NDMP Configuration** to display the page:

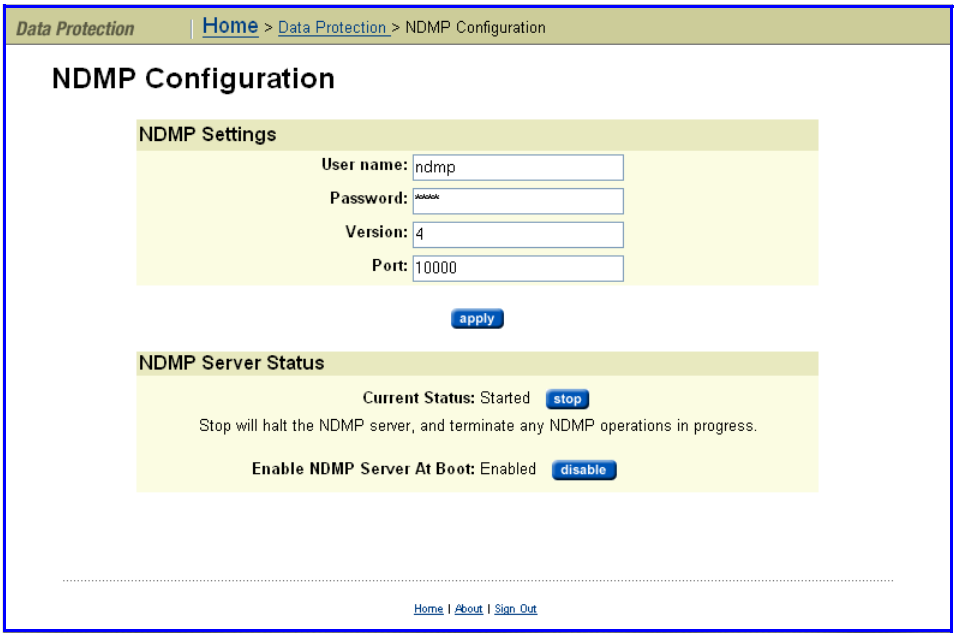

2. **Start or stop the NDMP process:**

Caution: *Read this caution before following instructions to start and stop!* Clicking **stop** terminates all NDMP processes immediately, leaving any tapes in use in an untidy state. It may also confuse the storage management application. Therefore, SGI Global Services recommends terminating NDMP transfers using the storage management application *before* clicking **stop**.

- To stop NDMP processing, click **stop**. If any NDMP operations are in progress when you click **stop**, those operations will be aborted.
- To start NDMP processing, click **start**.

3. **Enable or disable the NDMP process at Boot:**
   - To automatically enable NDMP processing at Boot, click **enable.**
   - To automatically disable NDMP processing at Boot, click **disable.**

Specifying the NDMP User Name, Password, and Version

A storage management application must successfully authenticate a configured NDMP user before starting a backup or recovery.

Note: Any user with NDMP *username* and *password* knowledge can access an NDMP-enabled storage management application to access data on the system. Therefore, SGI Global Services recommends taking measures to keep the information secure.

An administrator can specify two types of users:

- **NDMP Primary User.** For an *NDMP Primary User,* an account *username* and *password* provide full access to the files on the system, supporting most backup, recovery and replication activities.
- **Restricted NDMP Users.** The SSC command `ndmp-ruser` can create less trusted *NDMP Restricted Users* with access to a restricted set of files (and possibly devices). An administrator could assign these *usernames* to various users to allow them to use the ADC utility to copy data within limited areas of the file systems. The SSC command `ndmp-ruser-pwd` can also change the password for a selected restricted user.

  For more information about `ndmp-ruser` and `ndmp-ruser-pwd`, see the *Command Line Reference*.

To specify an NDMP *username* and *password:*

1. **Navigate to the NDMP Configuration page.**

   From the **Data Protection** page, select **NDMP Configuration** to display the **NDMP Configuration** page:

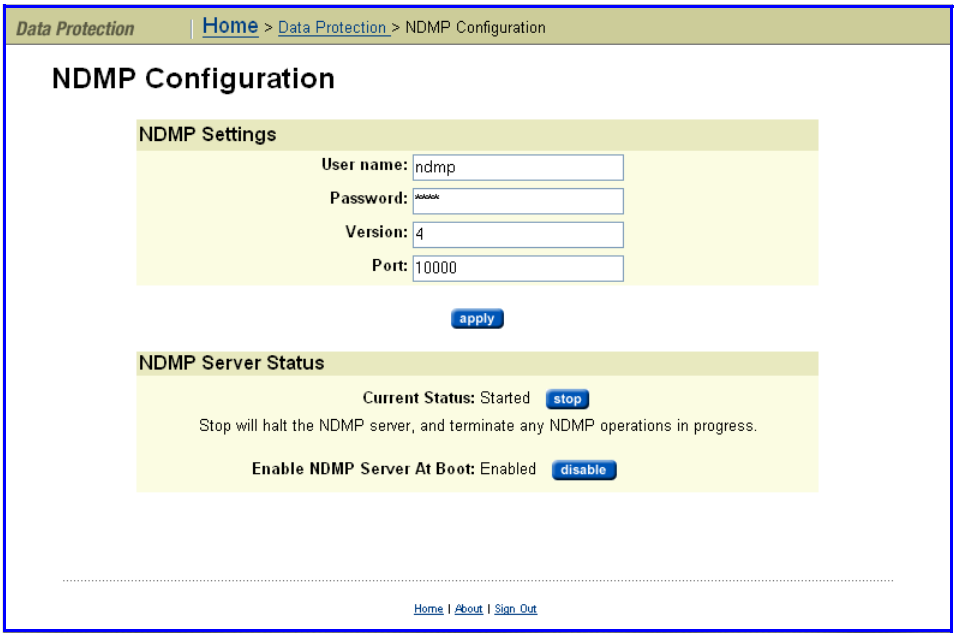

2. **Enter the requested information.**

   Recall that the server uses NDMP version 4 for NDMP backup and recovery; if required, it can be configured for version 2 or 3 of the NDMP protocol however, both incremental data replication and ADC require NDMP version 3 or 4, so SGI recommends using version 2 only if required for your backup application.

   Note: Additional configuration of NDMP can be performed using the `ndmp-option` CLI command. For more information, refer to the *Command Line Reference*.

3. **Save your changes.**
   Click **apply**.

Configuring NDMP Devices

NDMP backup devices, such as tape libraries and auto-changers, require special configuration. The server monitors its Fibre Channel (FC) links periodically and automatically detects the presence of backup devices. Since the server may be connected into a Storage Area Network (SAN) shared with other servers, it does not automatically make use of backup devices it detects on its FC links.

To use Web Manager to refresh the list of backup devices (for example, tape libraries and autochangers), modify server assignments, and to enable/disable access and visibility:

1. **Navigate to the NDMP Device List page.**
   From the **Data Protection** page, click **NDMP Device List** to display the page:

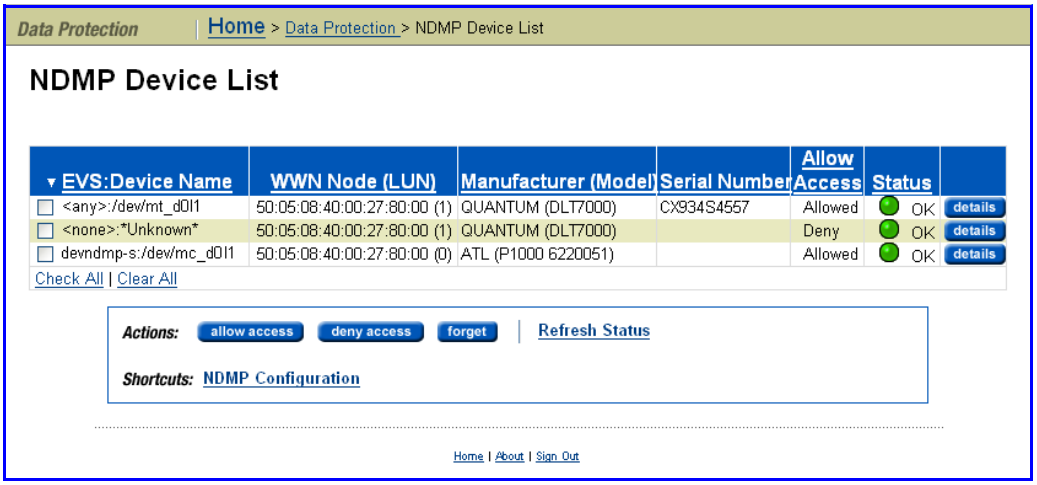

The following table describes the fields in this page:

Item/Field	Description
EVS:Device Name	Displays the EVS or EVSs allowed to use the device, and the ID of the device. This ID is generated by the system and cannot be changed.
WWN Node (LUN)	Displays the WWN (World Wide Name) and LUN ID of the Fibre Channel node.
Manufacturer (Model)	Displays the manufacturer and model of the device, if detected.
Serial Number	Serial number of the device.

Item/Field	Description
Allow Access	Displays if access is allowed to the device. If access is not allowed, then NDMP will not attempt to use the corresponding device. **Note:** An NDMP device must be assigned to an EVS before access can be allowed to the device. If access is not allowed to a device, fill the checkbox next to the device, and click **allow access**. To deny access to a device, fill the checkbox next to the device, and click **deny access**. A request to deny access will be rejected if an NDMP client has opened the device. The backup application configuration should be changed to avoid use of the device before denying access.
Status	Current status of the selected device.

2. **Enable/disable access to devices:**

   The following **Actions** are available:

   - Click **deny access** to disable access to a device, which prevents NDMP from attempting to use the device.

     Note: While an NDMP server has the device open, a **deny access** request will be rejected. Therefore, the storage management application configuration should be changed to avoid use of the device before the current configuration process.

   - Click **allow access** to enable access to a device, which allows NDMP to use the device.

     Note: Before using an NDMP device, you must first allow access to it, then it must be assigned to an EVS. NDMP Devices are assigned to an EVS using the **NDMP Device Access Details** page described in step 3, on page 352.

   - Click **forget** to remove the selected device from the list (only available for devices that have been disconnected from the FC).
   - Click **Refresh Status** to discover any changes in the Fibre Channel connection; that is, to find any newly attached devices and discover whether any previously discovered devices that are no longer accessible. If new devices are plugged into the Fibre Channel, use Refresh to identify them.

3. **Modify device access configuration as needed:**

From the **Data Protection** page, click **NDMP Device List**, then click **details** to display the page:

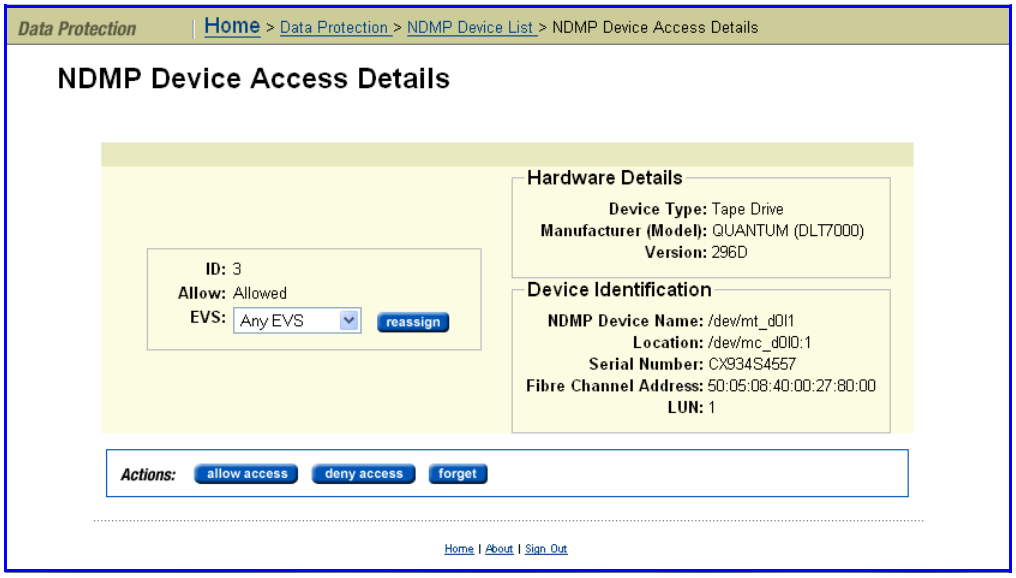

The following table describes the fields in this page:

Item/Field	Description
ID	Displays the server-assigned device identifier.
Allow	Indicates if device access is allowed (Allow) or denied (Deny). **Note:** An NDMP device must be assigned to an EVS before access can be allowed to the device. If access is not allowed to a device, click **allow access** to enable access. To deny access to a device, click **deny access**. A request to deny access will be rejected if an NDMP client has opened the device. The backup application configuration should be changed to avoid use of the device before denying access.

Item/Field	Description
EVS	Indicates the specific EVS to which the device is assigned, or indicates that the device is assigned to all EVSs. To change the device assignment, select the EVS to which you want to assign the device, or select **All EVS** to assign the device to all EVSs hosted by the server/cluster, and click **reassign**. Tape devices can be shared among EVSs under the following conditions: • The EVSs must be within the same cluster. • The tape device is not shared with another SGI server. • The tape device is not shared with another non-SGI device.  **Note:** If the device is to be shared between this server and another non-clustered server, additional sharing logic is required. Some backup applications automatically allow such sharing without any extra configuration. For other backup applications it is necessary to use the SCSI Reserve/Release protocol, which can be enabled using the CLI command `ndmp-option reserve_devices`.  **Note:** When one EVS is currently using a tape device, any attempt to use it through a different EVS will prompt a notification that the device is currently in use (that is, the operation will not be queued).  **Note:** When a tape device is currently assigned to a specific EVS (but not to *All EVS*), any attempt to access it through a different EVS will prompt notification that the device has not been found.
Hardware Details	This section displays hardware-related details about the device, including: • Device Type, which can be either tape drive or autochanger. • Manufacturer (Model), which are the device manufacturer and model detected when the device is discovered. • Version, which indicates the version of the firmware currently on the device, if it was detected when the device was discovered.

Item/Field	Description
Device Identification	This section displays identification information about the device, including: • NDMP Device Name, which displays the name by which the device can be addressed by the server. • Location, which displays the name of the autochanger that holds the drive and the position of the drive in the autochanger. For example, the location of the first drive in autochanger /dev/mc_d0l0 is /dev/mc_d0l0 : 1. • Serial Number, which indicates the device's serial number, if it was detected when the device was discovered. • Fibre Channel Address, which indicates the device's Fibre Channel node name. • LUN, which indicates the LUN identifier for the device. When the Web Manager cannot determine the location of a tape drive, it displays **unknown** . When this occurs, check for the following conditions and follow the troubleshooting instructions: • The tape library is offline. • The autochanger does not support the server's mechanism for querying the tape drive location, or the autochanger has not been set up to accept this query. Where this is the case, compare the serial numbers of the tape drives with displays available in the tape library to verify the drive locations. • The autochanger and a tape drive within it are attached to different servers. In this case, use the tape drive serial numbers to match the device name shown by one server with the location shown on the other.  **Note:** Devices will not be available or visible if access to them has not been enabled.

The following **Actions** are available:

- Click **deny access** to disable access to a device, which prevents NDMP from attempting to use the device.

Note: While an NDMP server has the device open, a **deny access** request will be rejected. Therefore, the storage management application configuration should be changed to avoid use of the device before the current configuration process.

- Click **allow access** to enable access to a device, which allows NDMP to use the device.

Note: Before allowing access, NDMP devices must be assigned to an EVS.

- Click **forget** to remove the selected device from the list on the **NDMP Device List** page (only available for devices that have been disconnected from the FC).

Displaying NDMP Device Information

Device information can be viewed through the **NDMP Device Access Details** page. From the **Data Protection** page, click **NDMP Device List**, then click **details** to display the **NDMP Device Access Details** page.

To configure your storage management application to work with NDMP devices, enter the names of the autochangers and tape drives.

Using NDMP with Snapshots

The server uses snapshots to backup data consistently and without being affected by on-going file activity. Snapshots also facilitate incremental backups.

This section describes how to use NDMP with snapshots, under the following topics:

- Backing Up Snapshots, on page 356
- Incremental Backups and Snapshots, on page 358
- Configuring NDMP Snapshot Options, on page 358
- Backing Up Virtual Volumes and Quotas, on page 360
- Clearing the Backup History, on page 361

Note: It is also possible to backup data without using snapshots.

Backing Up Snapshots

The following options should be considered when planning a backup strategy:

- **Back up automatically created snapshots.**

  When backing up a file system that is being actively updated, a snapshot of the file system is much more likely to produce a fully consistent image than backing up the live file system. As a result, NDMP is configured by default to automatically create a snapshot for backup.

- **Back up pre-created snapshots**

  A backup can be taken from a specific snapshot that has been *created by a rule* or *created spontaneously by user request:*

  - **To back up the latest snapshot created under a snapshot rule,** use the environmental variable `NDMP_BLUEARC_USE_SNAPSHOT_RULE`. For more information about environmental variables, refer to Supported Environment Variables, on page 547.

- **To back up the latest snapshot created spontaneously by user request,** request a specific snapshot by explicitly including the snapshot name in the path to back up. Where the path is based on a CIFS share name, indicate the snapshot using `/~snapshot/snapshot_name`; for paths based on an NFS export name, indicate the snapshot using `/.snapshot/snapshot_name`. CIFS shares and NFS exports may also include a snapshot name.

- **Backing up databases and iSCSI Logical Units**

  The internal structures of Databases and iSCSI Logical Units are tightly coupled with the state of the client software (database manager/iSCSI Initiator) that is controlling the files. For example, backing up such files during a client operation may produce inconsistencies in the backup that would prevent recovery.

  Therefore, any backup of databases and iSCSI Logical Units must ensure that files are in a consistent state at the time of back up. Snapshots can be used to achieve this. Snapshot rules provide the most convenient mechanism, as this avoids having to explicitly specify the name of the snapshot used.

  Note: When configuring snapshot rules, ensure that snapshots have a sufficiently long shelf life, and before initiating a backup, verify that the snapshot is not scheduled to be replaced during the anticipated time of the backup, as such replacement would cause the backup to fail.

  For more information on backing up and restoring iSCSI Logical Units, refer to Backing up iSCSI Logical Units, on page 301 and Restoring iSCSI Logical Units, on page 301.

All four of the following steps for backing up a database can be scripted for automatic backup creation at pre-defined times:

1. **Bring files into a consistent state.**

   Shut down the database or use a database-specific command to bring the database files into a consistent state.

2. **Take a snapshot of the file system.**

   For information about configuring snapshot rules and managing individual snapshots, refer to Managing Snapshot Rules, on page 342 and Managing Individual Snapshots, on page 346.

   Note: Users with permission can also take spontaneous rule-associated snapshot, without waiting for the next scheduled time. This can be done from the command line interface.

3. **Restart the database.**

4. **Make a backup copy of the snapshot.**

   Backing Up Snapshots, on page 356.

Incremental Backups and Snapshots

Because of the time consumed by full backups, regular incremental backups may be required to complement less frequent full backups.

However, incremental backups may fail to capture all of the changes in a file system; for example, where the modification time of a file is the determining factor in whether to back it up, a backup program will not archive the contents of a directory that has been moved, as the times/dates of the files remain unchanged.

Snapshots provide a solution to this problem. After taking the initial, base backup using a snapshot image, it can be used at incremental backup time to obtain a better picture of changes in the file system. In order to use the snapshots in this way, the snapshot must be kept around for as long as the associated backup may be used as the basis for an incremental backup.

Configuring NDMP Snapshot Options

By default, the server automatically creates a snapshot before it starts a backup operation. The backup then proceeds from the snapshot image rather than the file system. However, if the file system cannot take the snapshot for any reason, the backup proceeds directly from the live file system.

To configure NDMP snapshot options:

1. **Navigate to the NDMP History and Snapshots page.**

From the **Data Protection** page, click **NDMP History and Snapshots** to display the **NDMP History and Snapshots** page:

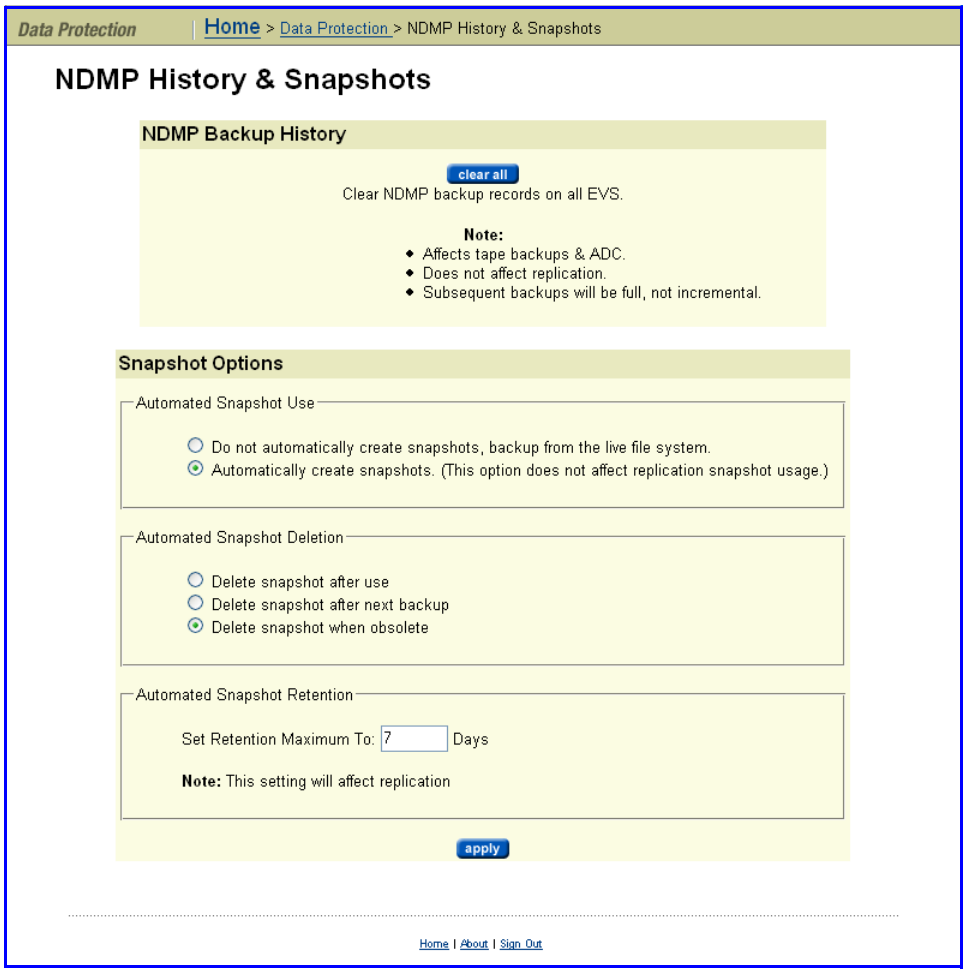

2. **Configure automated snapshot use.**

   This selection only affects backups or ADC copies where the path refers to the live file system. If the backup path already specifies a snapshot, or if the backup is using a snapshot rule, this selection has no effect. In the **Automated Snapshot Use** section, select whether NDMP should automatically create a snapshot to be backed up:

   - **Do not automatically Create Snapshots.** Backups of the live file system will use the live file system directly. If this option is selected, *do not* configure **Automated Snapshot Deletion** (next step), but *do* click **apply** to save or **cancel** to decline.

     Note: If a backup path explicitly contains a snapshot reference then the system does not take a new snapshot, regardless of this setting.

   - **Automatically Create Snapshots** (recommended). A backup of a path referring to the live file system will cause a snapshot to be taken for use in the backup.

3. **Configure automated snapshot deletion.**

By default, NDMP keeps the snapshot to make incremental backups more accurate. In the **Automated Snapshot Deletion** section, select whether to delete the snapshot:

- **Delete snapshot after use** deletes an automatic snapshot after completion of the backup for which it was taken. To prevent accumulation of unneeded snapshots, select this option for full backups or if the file system is changing very rapidly.
- **Delete snapshot after next backup** deletes an automatic snapshot after it has been used as the basis of a new incremental backup. With an exception for full backups, this option supports "incremental" backup schedules based on the immediately preceding backup.
- **Delete snapshot when obsolete** deletes an automatic snapshot upon next backup at the same level. For example, a snapshot taken for a full backup will only be deleted when the next full backup is completed. This option supports "differential" backup schedules based on a common base backup.

Note: Snapshots initiated by the Microsoft Volume Shadow Copy Service (VSS) may not be deleted by rule. These snapshots should be managed through the application that requested the snapshot. You can, however, delete these snapshots through the **Snapshots** page.

4. **Set the maximum retention time for automated snapshots.**

   Usually, the system deletes automatically created snapshots according to the rule selected in the previous step; however, after a sequence of backups using automatically created snapshots is stopped, snapshots may be left over. The maximum retention time provides a way of tidying up in these circumstances.

   Note: This setting applies to snapshots automatically taken by replications. Set the retention time to be long enough to make sure that a snapshot from a replication copy is not deleted until after the next successful copy is complete. This means that the maximum time set here must be longer than the time taken to run two replication copies, including the interval between the replication copies and the time required to make the copies.

   In the **Set Retention Maximum To** box, enter the number of days (1-40) to keep snapshots before auto-deletion.

5. **Save your changes.**

   Verify your settings, then click **apply** to save the settings.

Backing Up Virtual Volumes and Quotas

While CIFS and NFS backups will not backup or restore any virtual volume or quota information, NDMP backups or copies retain virtual volume and Quota information (if not disabled). Use NDMP in the following circumstances:

- NDMP backups to tape libraries.
- File system copies using the ADC utility.

- Data Replication controlled by the SMU.

Additional information regarding the virtual volume backups or copies:

- Setting the NDMP environment variable `NDMP_BLUEARC_QUOTAS` to *No* disables quota processing on NDMP backups/recoveries and ADC copies. See Supported Environment Variables, on page 547.
- Data Replication will copy virtual volume and quota information.
- Configuration information for a virtual volume will only be copied if its root is in the `backup/copy` path.
- Incremental backups and copies, including replication copies, transfer updates of a virtual volume if the `backup/copy` path includes the virtual volume root. However, copies and recoveries will not delete virtual volumes.
- If a recovered/copied source is merging its contents into an existing virtual volume, the virtual volume settings will also be merged. If a virtual volume is recovered/copied to an existing non-empty directory that is not part of the same virtual volume, the existing on-disk settings will be kept.
- Quotas associated with the file system (as opposed to quotas on virtual volumes) are not be recovered or replicated. Optionally, these quotas may also be written. This feature is controlled by the `vlqb` and `vlqr` options of the `ndmp-option` CLI command. See the *Command Line Reference* for more information on the `ndmp-option` command.

Clearing the Backup History

When necessary, you can clear the records of completed tape-based backups and either scripted or command line-based incremental ADC copies.

Note: Clearing the history does not affect replication operations (replication history is managed separately) or data migration operations (migration is not an incremental operation).

When performing incremental backups, the server uses the records of old backups to determine the date and time after which it must back up modified files. If you have lost a backup for any reason, you can clear the records, which forces the next backup to be a full backup instead of an incremental backup.

Note: To force a full backup for replication, delete the snapshot that was automatically created at the start of the last replication.

To clear the backup history, from **Data Protection** page, click **NDMP History and Snapshots** to display the **NDMP History & Snapshots** page:

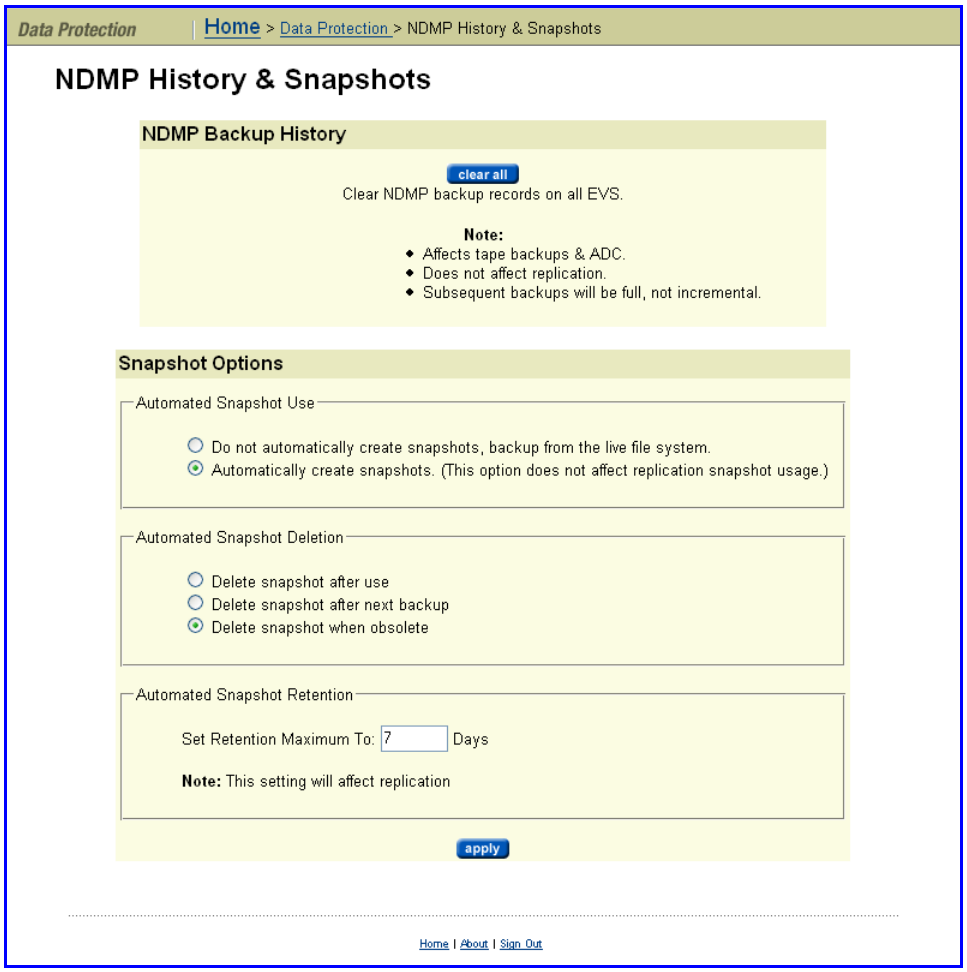

Click **clear all** to clear records of old backups. The next backup will be full rather than incremental.

Using Data Replication

Data replication provides a mechanism, manual or automatic, for copying or relocating both file data and file system metadata. IS-NAS Servers and Titan Servers support replication of data and, when using the transfer of primary access feature, of file system settings. When using replication with the transfer of primary access feature, you can relocate file system data and CNS links, CIFS shares, permissions and all other file-level metadata (for more information about the transfer of primary access, see Transfer of Primary Access, on page 337). Administrators can use Web Manager to configure policy-based replication jobs independently from other backup strategies.

A conceptual overview of Replication can be found at Data Replication, on page 333.

This section provides a deeper conceptual understanding of the components of data replication and instructions for configuring and implementing replication, in the following sections:

Configuring Policy-Based Data Replication for Managed and Unmanaged Servers

Before administrators can add a replication policy, the type of server that will be used for storing the replicated data must be determined. You can choose from one of the following policy destination types:

- **Managed Server:** For a server to be considered as managed server, it needs to be entered in the SMU configuration.
- **Not a Managed Server:** A non-managed server is one where the IP Address and username/password of the server is not known by the SMU. Administrators can still select a non-managed server as the target by specifying the IP address along with the username and password.

To configure policy-based data replication:

1. **Navigate to the Policy Destination Type page.**

   From the **Data Protection** page, select **Replication**, page, then click **add** to display the **Policy Destination Type** page:

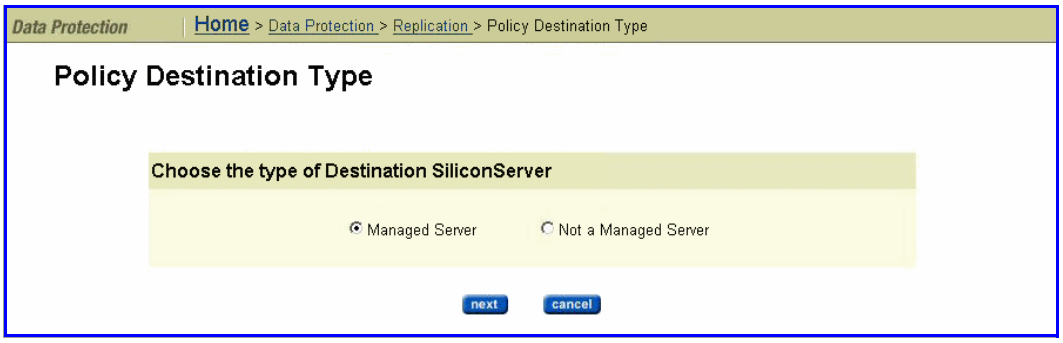

2. **Select server type.**

   Fill the **Managed Server** or **Not a Managed Server** box, then click **next** to display a server type-specific **Add Policy** page.

The **Add Policy** page for *managed* servers appears here:

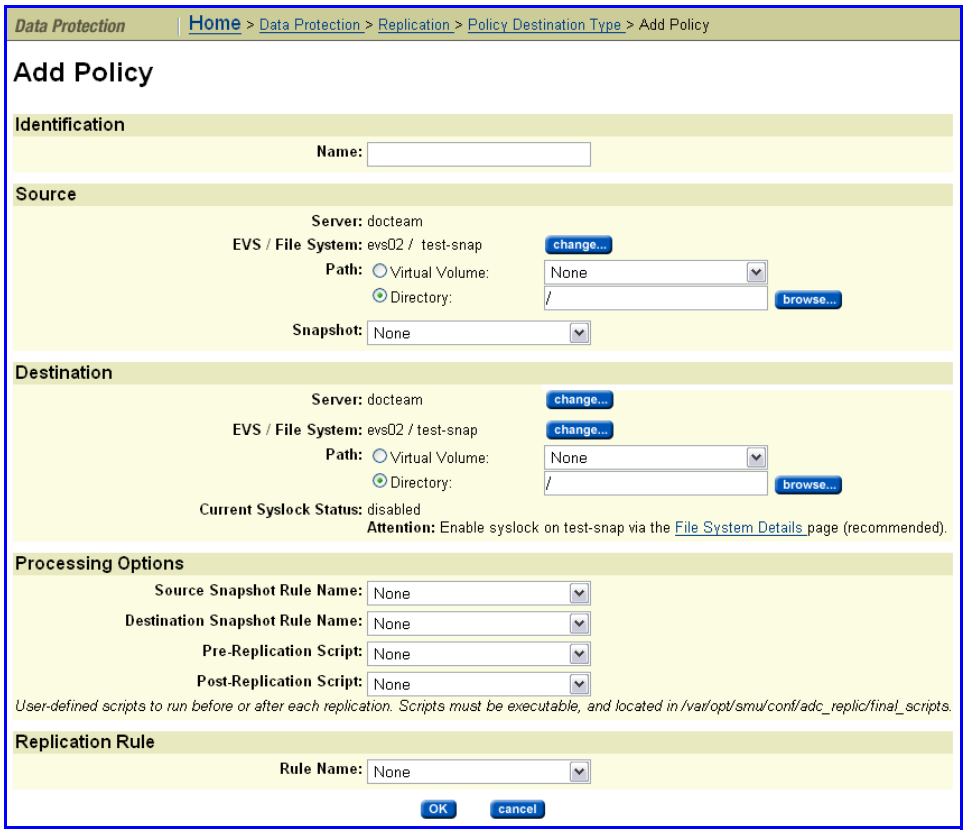

The **Add Policy** page for *unmanaged* servers is the same, with additional fields in the **Destination** section to specify the target server and NDMP username/password:

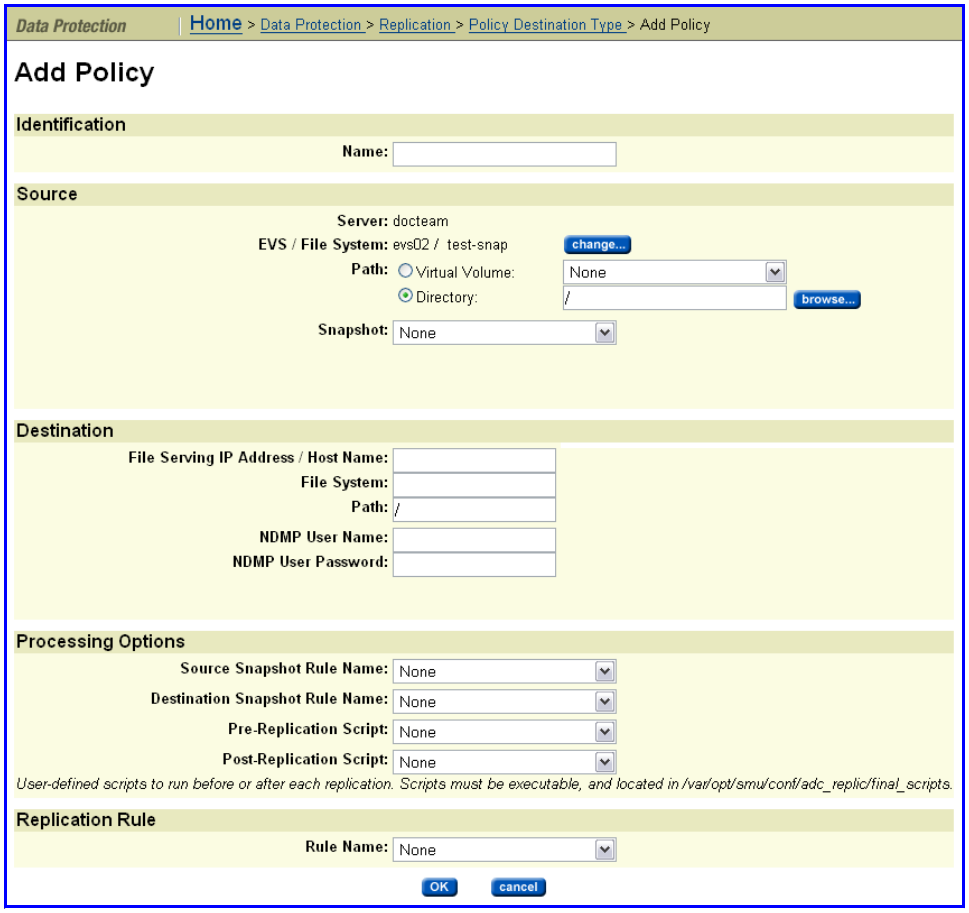

Note: Administrators should be authorized to use a non-managed server to access and store replication data.

3. **Enter the requested information.**

   The following table describes the fields in this page:

Item/Field	Description
Identification	Name of the replication, which must not contain spaces or any of the following characters: `\/<>"'!@#£$%^%&*(){}[] +=?:;,~`|.'`

Item/Field	Description
Source	Source of the replication. Set this field only if you want to make a simple copy of a specific snapshot. Do not set this field if you are intending to run incremental replications. The source is identified using the following fields: • **Server:** Name of the server where the replication will be created. • **EVS/file system:** Name of the EVS and file system to which the replication source is mapped. Click **change** to change the EVS or file system. • **Path:** Select a virtual volume from the drop-down list. Or select **Directory** and enter the path. • **Snapshot:** Select a snapshot to migrate a file system from a snapshot taken at a specific point in time. Using a snapshot as a source allows you to replicate the snapshot rather than the live file system, eliminating the possibility of file changes during the replication.
Destination (for managed servers)	Destination of the replication (managed server): • **Server:** Name of the server containing the target EVS file system. Click **change** to change the destination to a different server. • **EVS/file system:** Name of the Virtual Server and file system to which the replication is mapped. Click **change** to change the EVS/file system. • **Path:** Specify the directory path. Note that you may not specify a virtual volume as a path. • **Syslock file system**: Fill the checkbox to put the target file system into Syslocked mode. When System Lock is enabled for a file system, NDMP has full access to the file system and can write to it during a backup or replication, but the file system remains in read-only mode to clients using the file service protocols (NFS, CIFS, FTP, and iSCSI). For more information on Syslocked mode, see Using System Lock on File Systems, on page 156.
Destination (for non-managed servers)	Destination of the replication (non-managed server): • **File Serving IP Address / Host Name:** Name of the server containing the target EVS/ file system. Click **change** to change the destination to a different server. • **File System:** Name of the file system to which the replication is mapped. Click **change** to change the file system. • **Path:** Specify the directory path. Note that you may not specify a virtual volume as a path. • **NDMP User Name:** Name of the NDMP user for which the replication target was created. • **NDMP Password:** Password for the selected NDMP user.
Processing Options	• **Source Snapshot Rule Name:** The Snapshot Rule for replication of the source file system. • **Destination Snapshot Rule Name:** The Snapshot Rule to use for the snapshot of the destination file system following a successful replication. • **Pre-/Post-Replication Script:** A user-defined script to run before or after each replication. Scripts must be located in `/opt/smu/adc_replic/final_scripts`. The permissions of the scripts must be set to “executable”.
Replication Rule	Optional configuration parameters that allow tuning of replications to enable/disable specific functions or to optimize performance. For information about creating Snapshot Rules.

4. **Save your changes.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Understanding Snapshot Rules

By default, replications automatically create and delete the snapshots they require to complete consistent copies. That being the case, snapshot rules are not usually required. However, there are cases where the snapshots must be taken or used by external software. In these cases, snapshot rules are used so that the external software and the replication can be sure they are using the same snapshot.

Note: Snapshot creation is normally synchronized with a specific event. The snapshot is explicitly created at this time, so the snapshot rule should not have an associated snapshot schedule.

Specific instances where snapshot rules may be used include:

1. Replications which copy databases or iSCSI LUNs. A snapshot taken automatically at the start of a replication will not capture a consistent image of a database or an iSCSI LUN that is actively in use. In order to capture a consistent image, the database/iSCSI LUN needs to be brought into a quiescent state before the snapshot is taken. These actions are normally be executed by a script, which then takes a snapshot within the snapshot rule so that the replication can identify which snapshot to copy.

   The script could be invoked as part of a pre-replication script (see Understanding Custom Replication Scripts, on page 368). Alternatively the script could be independently scheduled. If scheduled independently, however, the schedule must allow the script to complete before the replication starts.

2. Linked, two-stage replications, which copy a file system from server A to server B and then copy on from server B to server C. These types of replications can use snapshot rules to synchronize the copies.

   The replication from server B to server C may start while a copy from server A to server B is running. If a snapshot was taken at this point, an inconsistent file system state would be captured. One way to avoid this is to use a specific snapshot rule as both the destination snapshot rule of the server A to server B copy and the source snapshot rule of the copy from B to C. Then the B to C copy will always copy a snapshot taken at the end of the last complete copy from A to B.

Two kinds of rules define snapshot use during replication:

- **Source Snapshot Rules** determine which snapshot to use as the replication source.

  For Replication Policies configured to use a *source snapshot rule,* the most recent snapshot associated with the rule becomes the replication source.

  Source snapshot rules are particularly useful when the replication includes a database or other system that must be stopped in order to capture a

consistent copy. Based on an external command (perhaps issued by a pre-replication script), the data management engine expects that a snapshot will be taken.

To perform *incremental replications*, the data management engine requires that the snapshot used during the previous successful replication still exist when a new replication is made. If you are using the snapshot rule queue length to control the deletion of snapshots, you must take this requirement into account and set the queue length long enough to allow for keeping the snapshot used during the previous successful replication. Also, you must take into account the possibility of intermediate failed replications, which may also create snapshots.

The following actions are taken if the required snapshots do not exist:

- **If no snapshot exists in the rule,** then the data management engine issues a warning message and performs a full replication, using an automatically created snapshot that it deletes immediately after the copy.
- **If the snapshot taken during the previous replication has been deleted,** the data management engine cannot take an incremental snapshot and therefore performs a *full* copy.

- **Destination Snapshot Rules** govern the snapshot taken after a successful replication operation.

Understanding Custom Replication Scripts

Under normal conditions, pre- and post-replication scripts are not required. Where required to perform specific functions (for example, to stop an application to facilitate a snapshot of its files in a quiescent, consistent state), these custom scripts can be run *before* or *after* each instance of a replication.

In the case of databases or other applications that require a consistent state at the time of a snapshot, best practices indicate using scripts and snapshot rules together:

- **Pre-replication scripts** are executed to completion before the replication is started.
- **Post-replication scripts** are executed after a successful replication.

Potential uses of scripts are illustrated in the following examples:

- **Database replication.** A *pre-replication* script can be used to enable the replication of a consistent copy of the database. Typically, this pre-replication script will need to:

  1. Shut down the database to bring it into a consistent or quiescent state.
  2. Take a snapshot of the file system using a snapshot rule.
  3. Restart the database.

- **Backing Up Data from the Replication Target.** A *post-replication* script can initiate incremental (or full) backups from a replication target after each incremental replication has completed. Backing up from the replication

target (rather than the original volume or directory) minimizes the performance impact on network users.

Using Replication Rules

The **Replication Rules** page lists all existing rules and allows creation of new rules. Replication rules comprise optional configuration parameters that allow replications to be tuned to enable/disable specific functions or to optimize performance.

Replication Rules control values like the *number of read-ahead processes*, *minimum file size used in block replication*, *when snapshots are deleted* and *whether replications will include migrated files.* The server's default values should be optimal in most cases; however, these values can be changed to customize replication performance characteristics based on the data set.

Viewing Replication Rules

To view replication rules, navigate to the **Data Protection** page, then click **Replication Rules** to display the **Replication Rules** page:

The fields on this page are described in the table below:

Item/Field	Description
Rule Name	Displays the name given to the Rule when created, and referenced when creating or configuring replication policies.
In Use by Policies	Select to indicate that the rule is being used by one or more policies.
Details	Click **details** for a rule to view its complete details. Select a rule and click **remove** to delete it.

The following **Actions** are available:

- Click **add** to add a new rule, then refer to Adding a Replication Rule, on page 370
- Click **remove** to delete a selected rule.

- Click **details** for a rule to display its properties. This page contains the same fields as the Add Rule page, as displayed and defined under Adding a Replication Rule, on page 370.

The following **Shortcut** is available:

- Click **Policies and Schedules** to display the **Policies and Schedules** page.

Adding a Replication Rule

To add a replication rule:

1. **Navigate to the add Rule page.**

   From the **Data Protection** page, select **Replication Rules,** then click **add** to display the **Add Rule** page:

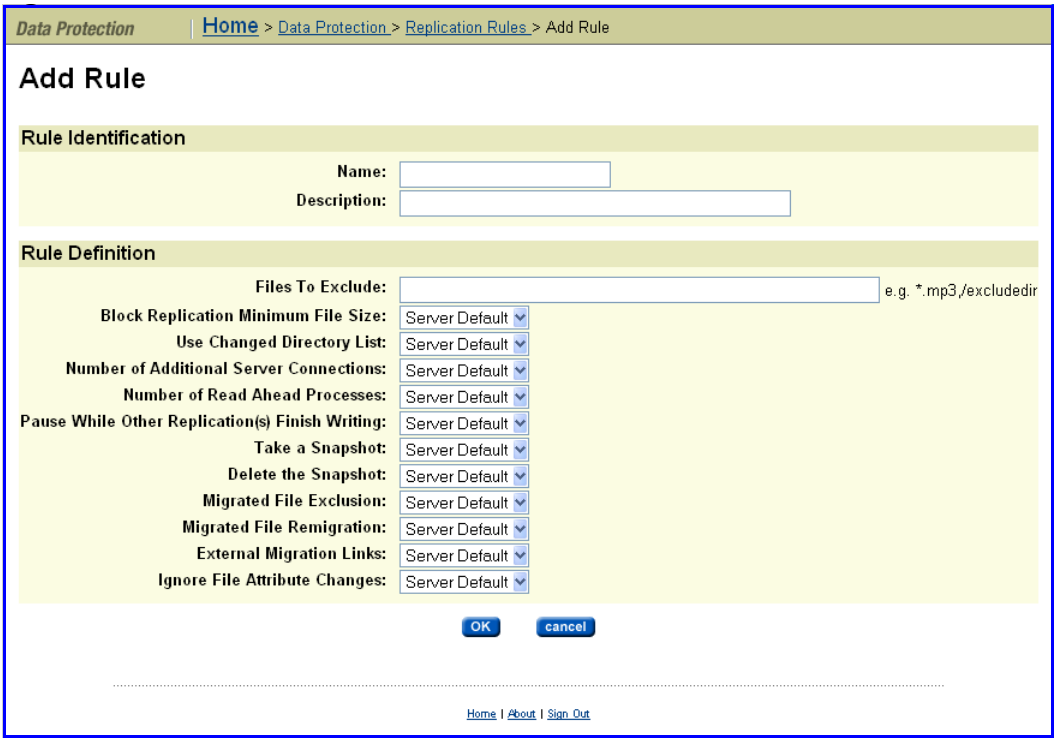

2. **Enter the requested information.**

   The fields on this page are described in the table below:

Item/Field	Description
Name	Name of replication rule. The rule name is may include only alphanumeric characters, hyphens, and underscores.
Description	Free form description of what the replication rule does.

Item/Field	Description
Files to Exclude	Specifies files or directories to exclude from a replication. By default, none are excluded. When specifying a file or directory, enter either: • *A full path name,* relative to the top-level directory specified in the replication path. The path name must begin with a forward slash (/); at the end, an asterisk (*) can be entered as a wildcard character. • *A terminal file or directory name,* which is simply the last element in the path. The name must not contain any forward slash (/) characters; at the beginning or end, an asterisk (*) can be entered as a wildcard character.
Block Replication Minimum File Size	Block replication minimum file size controls the minimum file size that is used for block replication. The drop-down list options available are: 256 or 512 K, and 1, 2, 4, 8, 16, 32, 64 or 128 MB If this option is set to 64 MB: • For a source data file of 63 MB, for which the system determines that only 1 MB has changed, the entire source file (63 MB) will be replicated. • For a source data file of 65 MB, for which the system determines that only 1 MB has changed, only the delta will be replicated.  **Note:** This option is only functional if the *Replication license* is present.
Use Changed Directory List	Indicates whether incremental replications will use a *changed object list* to direct the search for changed files. Processes not using the *changed object list* must search the entire directory tree looking for changed files. When using the changed object list, the search only passes through those directories that contain changed files.  **Note:** Using the *change object list* is likely to improve performance in some cases; for example, where there are sparse changes. However, it can degrade performance where there are many changes throughout the directory structure.
Number of Additional Server Connections	Controls the number of additional server connections that will be established during a replication operation. See Setting NDMP Performance Options, on page 385 for more information.  **Note:** Each additional server connection consumes system resources, and best practices indicate limiting the number of additional server connections to situations where they improve performance. Also, as the number of additional server connections is increased, more read-ahead processes are required. See Setting NDMP Performance Options, on page 385 for more information.

Item/Field	Description
Number of Read Ahead Processes	Controls the number of read-ahead processes used when reading directory entries during a replication. While the default number of read-ahead processes is suitable for most replications, file systems made up of many small files increase the amount of time spent reading directory entries proportionately. In such cases, adding additional processes may speed up the replication operation. See Setting NDMP Performance Options, on page 385 for more information.  **Note:** As each additional read-ahead process consumes system resources, best practices indicate limiting the number of additional processes to situations where they improve performance.
Pause While Replication(s) Finish Writing	By default, the data management engine imposes an interlock to stop *NDMP backups* and *ADC copies* from the destination of a replication during active replication writes. This function supports installations that replicate to a particular volume, then back up from that volume. However, as the lock is held at the volume level, it may be useful to override this action in the case of directory-level replication. To make use of this replication interlock, specify this rule option on both the replication *that waits* and the replication that is *waited upon*. As a best practice: • Create one rule with this option enabled and have each participating Replication Policy enable the same rule. • Then, schedule the Replication Policy *that waits* to run after the Replication Policy that is *waited upon.*
Take a Snapshot	Overrides the Backup configuration option **Automatic Snapshot Creation**. The setting for this option should be left as the server default in almost all cases. The only case where it might be useful is when taking a single, non-incremental copy of a file system or a directory. If there is insufficient space on the file system to take a snapshot, the copy may be taken from the live file system by selecting "Disable." However, it should be noted that copying the live file system while it is changing may give an inconsistent copy. • Enable this option to support incremental replication copies. • This option should only be disabled for full replication copies or when making a complete copy of a directory. Different files will be copied at different times, so if the source file system is changing and there are dependencies between different files on the system, then inconsistencies may be introduced.  **Note:** Snapshots are an integral part of the algorithm for incremental replication, and disabling snapshot usage will affect the ability to run incremental replications. This option must be enabled in order to make incremental replication copies.

Item/Field	Description
Delete the Snapshot	Determines when snapshots are deleted. The setting for this option should be left as the server default in almost all cases. The only case where it might be useful is when taking a single, non-incremental copy of a file system or directory. If the file system is short on space, it may be useful to request the immediate deletion of the snapshot taken for the replication. The deletion options are: • `IMMEDIATELY` gives the same effect as **Delete snapshot after replication is done,** • `LAST` preserves snapshot for use with incremental replications, and • `OBSOLETE` deletes an automatically created snapshot when the next backup of the same level is taken.  **Note:** As changing these settings can adversely effect the replication process, SGI Global Services recommends that this option be changed only at the direction of SGI Global Services.
Migrated File Exclusion	Indicates whether replications will include files whose data has been migrated to secondary storage. • The default setting of *disable* means that migrated files and their data will be replicated as normal files. • If set to *enable,* the replication will not include files whose data has been migrated to another volume using the Data Migrator facility.
Migrated File Remigration	Controls the action at the destination when the source file had been migrated. • If set to *enabled,* the file will be re-migrated when written to the destination volume, provided the volume or virtual volume has a Data Migrator path to indicate the target volume. • If set to *disabled,* all the files and their data will be written directly to the replication destination volume.
External Migration Links	Controls what happens when a replication operation encounters a cross volume link (a link to a file that has been migrated to an external server). • If set to *Server Default*, the replication operation uses the default setting, which is *re-migrate*. • If set to *re-migrate*, the replication operation copies the file contents but marks the file as having been externally migrated. The destination re-migrates to secondary storage if there is an existing data migration path. This is the default behavior. Use this setting when the replication is between a main site and a disaster recovery site, where the disaster recovery site includes a similar data migration configuration. • If set to *ignore*, the replication operation copies only the files on the primary (migrated files are not copied). Use this setting when files have been migrated because they are less useful, so they are not replicated in order to save time. • If set to *re-create link*, the replication operation copies only the details of the cross volume link. The cross volume link is recreated on the destination if the relevant external migration data path is in place and the migrated file is accessible. Use this setting when the replication is between storage servers or clusters on the same site, and there is a single external migration target server.
Ignore File Attribute Changes	Specifies that files where the only change is an attribute change, are not included in a replication. You should only enable this option if you are certain that you do not want to replicate files with only attribute changes. This option is disabled by default.

Modifying a Replication Rule

To modify a replication rule:

1. **Navigate to the Modify Rule page.**

   From the **Data Protection** page, select **Replication Rules**, select the rule you want to modify, and click **details** to display the **Modify Rule** page:

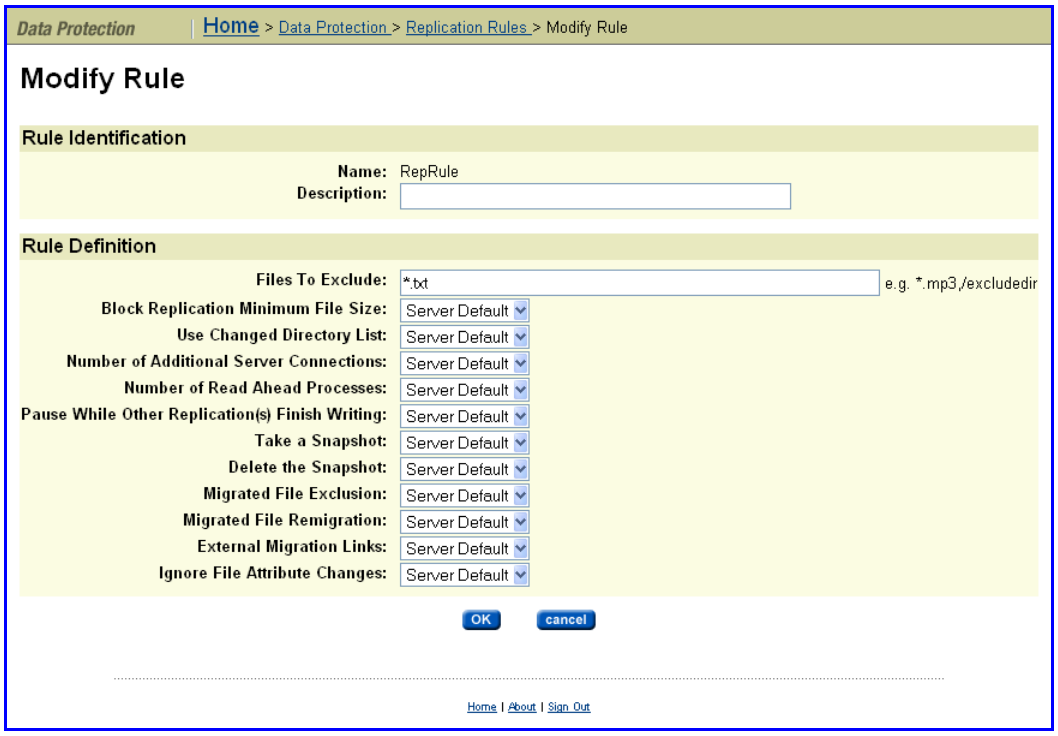

2. **Enter the requested information.**

   The fields on this page are the same as those on the **Add Rule** page. For a description of the fields on this page, see Adding a Replication Rule, on page 370.

3. **Save the modified rule.**

   Once you completed making changes to the rule, click **OK**.

Understanding Files to Exclude Statements

Files to Exclude statements contain expressions identifying directories or files to exclude from the replication. They can be written using the following guidelines:

- The asterisk "*" can be used as a wildcard character to qualify path and filename values.

  In a path, "*" is only treated as a wildcard if it appears at the end of a value, e.g. `/path*`.

  In a filename, a single * can appear at the beginning and or at the end of the value; for example, `*song.mp*, *blue.doc, file*.`

- Parentheses `()`, spaces, greater than `(>)`, and quotation marks (“) are allowed around a filename or path list, but they will be treated as literal characters.
- Path and filename can be defined together but must be separated by a comma (`,`); for example, `subdir/path*,*song.doc,newfile*,/subdir2`
- The forward slash (`/`) is used as a path separator. As such, it must not be used in a filename list.

Note: SGI Global Services recommends creating the *Files to Exclude* list before the initial replication copy, and not changing it unless necessary. When running incremental updates, changes in the list do not act retrospectively. For example, for a list initially excluding `*.mp3 that drops this exclusion,` any new or changed mp3 files will now be replicated; however, any `.mp3` files than haven't changed since the previous replication copy will not be replicated.

Replication Schedules

This section describes how to create and modify replication schedules, in the following sections:

- Viewing Scheduled Replications, on page 376
- Adding a Replication Schedule, on page 377
- Modifying a Replication Schedule, on page 378

Overview

After a Replication Policy has been defined, it must be scheduled to run. Replications can be scheduled and rescheduled at any time and with any of the available scheduling options.

Replication Schedules Overview:

- **Periodic replication:** Replications occur at preset times. Periodic replications can be set up to run daily, weekly, monthly or at intervals specified in numbers of hours or days.
- **Continuous replication:** When a replication policy specifies continuous replication, as soon as the replication job completes, the same replication job starts again.
- **One time replication:** A new replication job starts after the previous job has ended. The new replication job can start immediately or after a specified number of hours.

When planning Replication Schedules, SGI Global Services recommends scheduling during off-peak times such as nights or weekends. Once a replication has started, additional replications for the same policy cannot start until the current replication has completed; however, multiple concurrent replications are allowed for replications by different policies.

Viewing Scheduled Replications

To view scheduled replications, navigate to the **Data Protection** page, select **Data Protection,** then click **Replication** to display the **Replication for *NAS Server Name*** page:

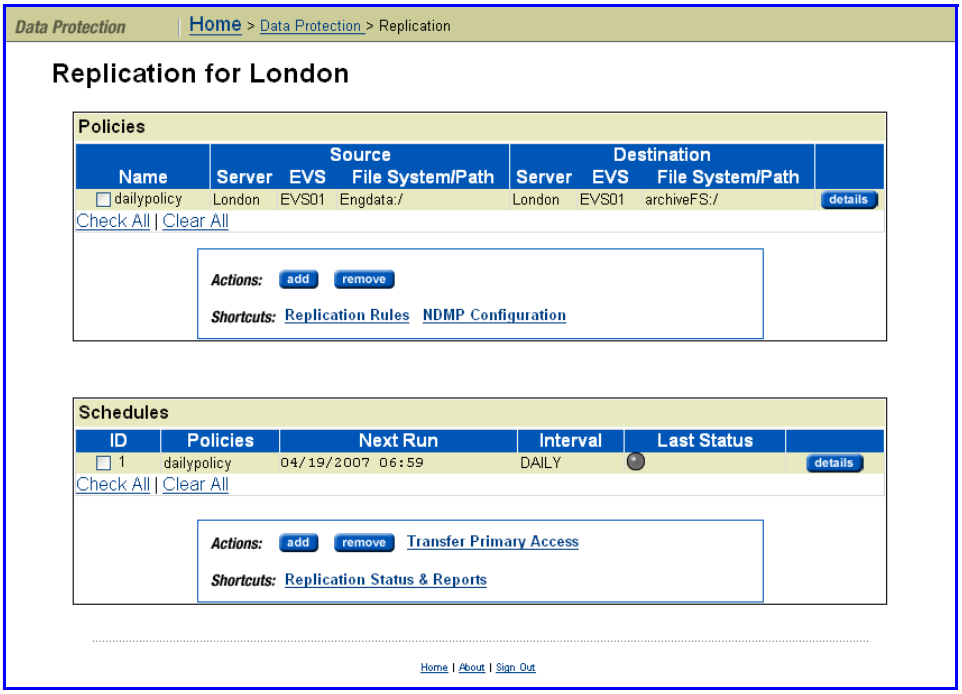

The fields on this page are described in the table below:

Item/Field	Description
Id	ID assigned to the Replication Policy.
Policies	Name of the Replication Policy.
Next Run	Month, date, year and time for the next scheduled replication run for this policy.
Interval	Frequency at which the replication has been scheduled to run.
Last Status	Light indicator for successful and failed replication jobs: • *Green* indicates that a successful replication job has completed. • *Red* indicates a failed replication job and lists the reason for failure. **Note:** In case of a replication failure, the next time a replication starts, the data management engine attempts to restart the failed replication instead of starting a new replication.

The following **Actions** are available:

- In the **Policies** section,
  - To create a new replication policy, click **add**.
  - To delete a selected replication policy, click **remove**.
- In the **Schedules** section,
  - To create a new schedule, click **add** and refer to Adding a Replication Schedule, on page 377.
  - To delete a selected schedule, click **remove**.

Adding a Replication Schedule

A replication policy must be set up before a replication job can be scheduled.

To add a replication schedule:

1. **Navigate to the add Schedule page.**

   From the **Data Protection** page, select **Replication,** then click **Add** to display the **Add Schedule** page:

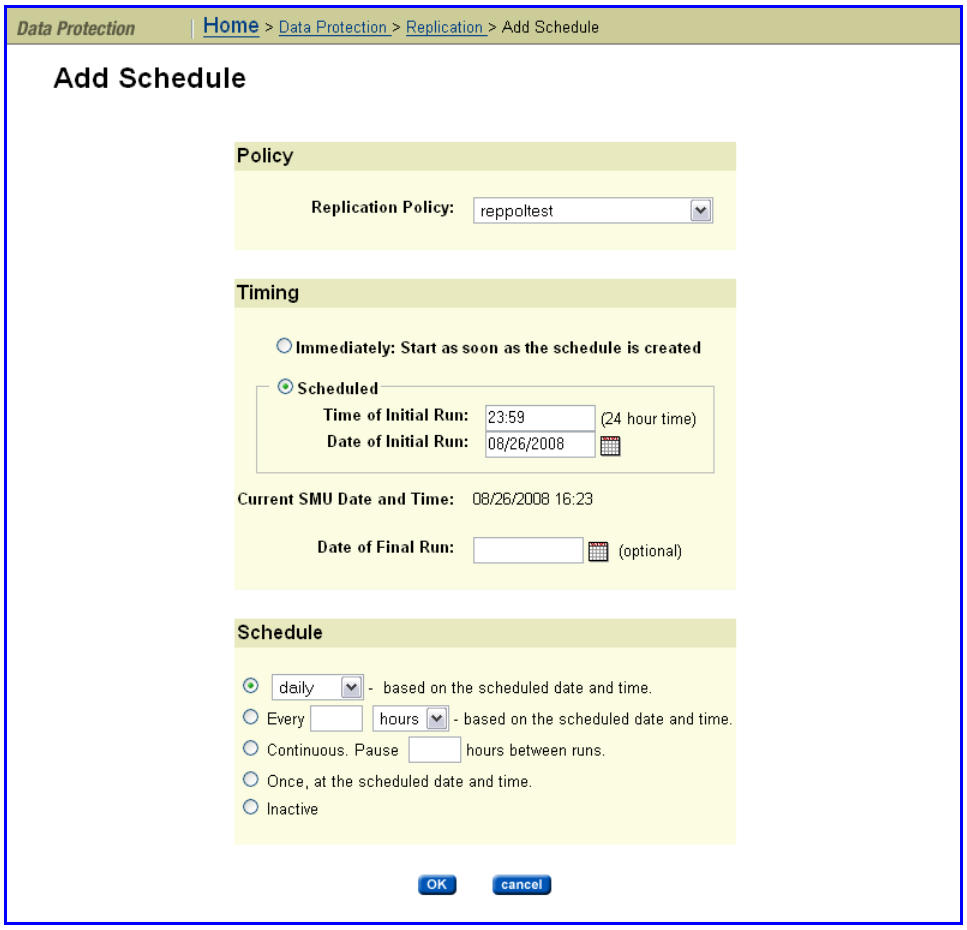

2. **Enter the requested information.**

   The following table describes the fields in this page:

Item	Description
Replication Policy	Select a replication policy from the drop-down menu.
Time of Initial Run	Scheduled run time on a 24 hour clock (such that 11:59 PM will be entered as 23:59). The current SMU date and time are provided at the bottom of the section for reference.
Date of Initial Run	Click the calendar next to the field, then select the start date for the policy's initial run. The selected date appears on the field.
Date of Final Run	Click the calendar next to the field, then select the start date for the policy's final run. The selected date appears on the field. This is an optional setting.
Schedule	Select one of the five radio buttons: • *From the dropdown, select Daily, monthly, or weekly* **based on the scheduled date and time.** • Enter a *quantity,* then from the dropdown select *hours* or *days* **based on the scheduled date and time.** • Enter a quantity to complete the label: **Continuous. Pause *quantity* hours between runs.** The new replication job can start immediately or after a specified number of hours. • Selecting **Once, at the scheduled date and time** guarantees that the policy is scheduled to run only once. • Selecting **Inactive** causes the replication schedule to be placed on pause.  **Note:** If an excess amount of time elapses between replication runs, snapshots may take up a larger amount of space. By default, replication-defined snapshots are purged after 7 days (configurable to 40 days). Waiting 8 or more days between replication runs could result in a full replication.

3. **Save your changes.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Modifying a Replication Schedule

Once defined, schedules can be easily modified to meet the changing requirements of the Replication Policies. When modifying a schedule, the scheduled date and time, as well as the interval in which the schedule will run, can be changed.

To modify a replication schedule:

1. **Navigate to the Modify Schedule page.**

   From the **Data Protection** page, click **Replication** to display the **Replication** page, then select a schedule and click **details** to display its properties in the **Modify Schedule** page:

2. **Enter the requested information.**

   The following table describes the fields in this page:

Item	Description
Policy	This section displays information about the replication policy being scheduled.
Replication Policy	Displays the name of the replication policy being scheduled.
Next Run	Displays the date and time of the next replication run specified by this schedule.
Last Status	Displays the status of the last run of this schedule. Click the **View Latest Report** to display the Replication Report for the last replication run according to this schedule.

Item	Description
Immediate Actions	Click **Run now** to run the replication policy immediately, regardless of schedule.  **Note:** A replication job cannot be started if a previous instance of the same policy is still in progress. In this case, the replication is skipped, and an error is logged. Click **Abort** to stop an in-progress replication.
Recovery Actions	Click **restart** to restart the replication if the previous replication attempt failed. Click **rollback** to roll back a failed or aborted replication. The target file system is rolled back to the last good snapshot. Note that a snapshot is taken after every successful replication.  **Note:** Rollback should only be used when the target will be used as the live file system. If the replication's source file system cannot be used as the live file system (either permanently or temporarily), users can access the latest available data on the replication target (the file system created by the last successful replication). **If the target file system will be used as the live file system permanently**, delete the replication policy and all related schedules (since the source will not be used for this replication again). You can then create new replication policies and schedules. **If the target file system will be used as the live file system temporarily**, contact SGI Technical Support for assistance in synchronizing the "old" (source) and the "new" (target) file systems before transferring access and resuming replication operations as implemented prior to the "rollback."
Timing	This section displays information about the execution timing for the replication policy, and it allows you to reschedule the next (or final) execution of the replication policy.
Schedule Time	Time of the next replication specified by this schedule.
Schedule Date	Date of the next replication specified by this schedule.
Date of Final Run	Date of the final replication specified by this schedule.
Reschedule	This section allows changing the schedule of the next or final replication specified by this schedule: • To change the schedule of the next replication, fill the **Reschedule** box, then enter the new values for the **Time** and/or **Date**. • To change the schedule of the final replication, fill the **Reschedule** box, then enter the new value for the **Final Run** in the appropriate fields.
Current SMU Date and Time	Current date and time as set on the SMU.
Schedule	This section allows you to specify how often the replication policy is to be executed.

Item	Description
	Select one of the five radio buttons: • *From the dropdown, select Daily, monthly, or weekly* **based on the scheduled date and time.** • Enter a *quantity,* then from the dropdown select *hours* or *days* **based on the scheduled date and time.** • Enter a quantity to complete the label: **Continuous. Pause *quantity* hours between runs.** The new replication job can start immediately or after a specified number of hours. • Selecting **Once, at the scheduled date and time** guarantees that the policy is scheduled to run only once. • Selecting **Inactive** causes the replication schedule to be placed on pause.
Actions	
OK	Click **OK** to save changes to the replication policy schedule, and return to the **Replication** page.
cancel	Click **cancel** to return to the **Replication** page without saving changes to the replication policy schedule.

3. **Save your changes**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Understanding Incremental Replications

Incremental replications rely on the existence of the snapshot taken during the previous replication. If this snapshot no longer exists, the data management engine performs a full replication. The data management engine automatically preserves the snapshots it needs for replication. However, there is an age limit applied to snapshots that are automatically taken by the NDMP system (including during a replication).

Snapshots older than the age limit are automatically purged from the system. The default limit is 7 days, but the limit can be configured through the **NDMP History and Snapshots** page. If the replication copy time is very long, or the interval between replications is long, then the default age limit must be extended.

Viewing Replication Status & Reports

To view replication status and reports, navigate to the **Data Protection** page, then click **Replication Status & Reports** to display the **Replication Status & Reports** page:

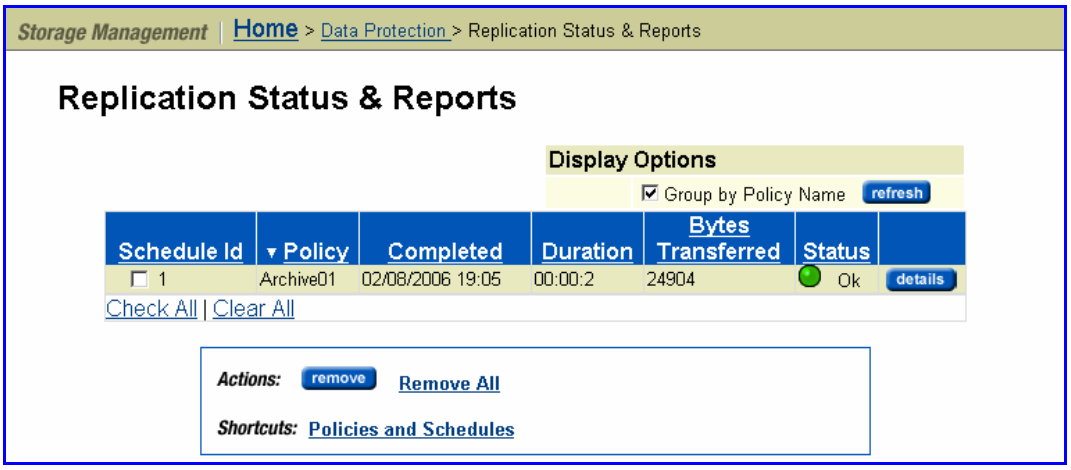

This page displays a list of replication jobs *in progress* or *completed*. It also includes reporting details on files replicated, amount of data replicated, and success or failure status. If a schedule is deleted, the reports associated with it are also deleted.

The replication report **Status** column displays results of a replication job (*green* for OK, *red* for failed). Reports can also be beneficial for analyzing the effects of a particular incremental replication policy. The information in the **Report Summary** page provides a detailed view of the replication job results. This information can be used to make performance adjustments to the replication policy and schedule.

The following table describes the fields in this page:

Item	Description
Schedule ID	ID number for the completed replication.
Policy	Policy name.
Completed	Month, date, year and time when the replication was completed.
Duration	Duration of a replication schedule run.
Bytes Transferred	Volume of replicated data in bytes.
Status	Status of replication completion.

To view replication reports, navigate to the **Data Protection** page, select **Replication Status & Reports,** then **Replication Reports,** then click **details** for a selected replication to display its properties:

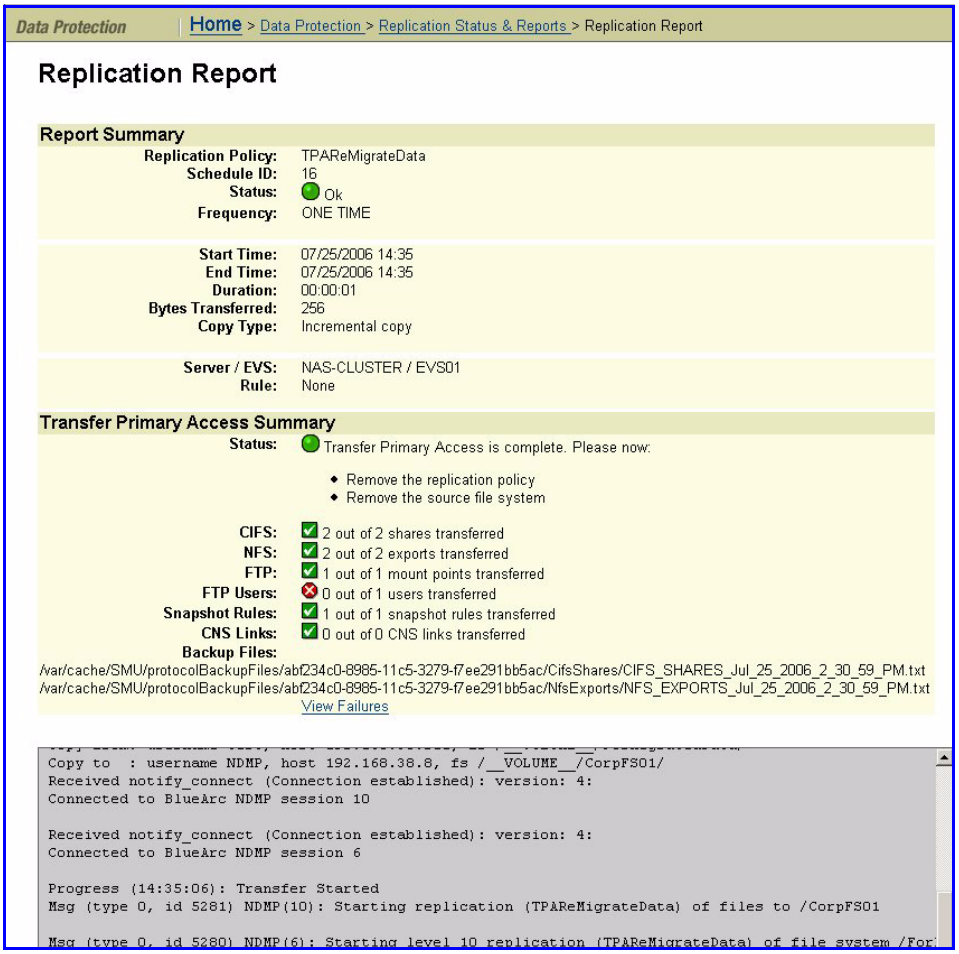

The following table describes the fields in this page:

Item	Description
Replication Policy	Completed replication policy name.
Schedule ID	Completed replication schedule ID.
Status	Indicates whether the replication was successfully completed.
Frequency	How often the Policy is scheduled to run.
Start Time	Date and time when the replication began.
End Time	Date and time when the replication ended.
Duration	Duration of replication.
Bytes Transferred	Volume of data replicated, in bytes.

Item	Description
Copy Type	Type of replication performed. May be any of the following: • **Full Copy**: A complete initial replication of the entire source to the target. • **Incomplete Copy**: The replication did not complete. • **Incremental Copy**: A replication of the changes on the source file system to the target. • **Restart Copy**: The replication started from the point of failure of the previous replication. • **Rollback Copy**: After a failed replication run, the target file system was rolled back to its state following the last successful replication
Server/EVS	EVS on which the Source and Destination file systems reside.
Rule	The name of the rule used by the policy.
Transfer Primary Access Summary This section appears in the replication report only after a transfer of primary access.	
Status	Indicates whether the transfer of primary access was successfully completed, and indicates any actions that should now be taken.
CIFS	Number of CIFS shares that were successfully transferred to the new location.
NFS	Number of NFS exports that were successfully transferred to the new location.
FTP	Number of FTP initial directories that were successfully transferred to the new location.
FTP Users	Number of FTP users that were successfully transferred to the new location.
Snapshot Rules	Number of snapshot rules successfully transferred to the new location.
CNS Links	Number of CNS links successfully transferred to the new location.
Backup Files	List of CIFS shares backup files and NFS exports backup files that were successfully transferred to the new location.
View Failures	Click **View Failures** to view a list of items not transferred during the transfer of primary access.

Enabling Multiple Replication Streams

You can add additional server connections to a replication rule using the "Number of Additional Server Connections" field of the replication Add Rule page (see Adding a Replication Rule, on page 370) or the Modify Rule page (see Modifying a Replication Rule, on page 374).

Select the number of additional connections to add for use by the replication/ADC copy operation. You can specify between 0 and 30 additional server connections. Note that these are **additional** server connections; if the number

of additional connections is set to 0, the replication operation will have a single connection. The default is 4 additional connections, along with 12 read-ahead processes.

If the number of additional server connections has been set to non-default and more than zero, then the number of read-ahead processes must also be set to a non-zero value that is appropriate for the specified number of additional server connections. See Setting NDMP Performance Options, on page 385 for more information about read-ahead processes.

Setting NDMP Performance Options

NDMP Performance options are set using the **Add Rule** or **Modify Rule** page of Web Manager. On these pages, you can set the number of additional server connections and the number of read-ahead processes (which are the options with the biggest effects on replication performance), as well as other replication options. See Adding a Replication Rule, on page 370 for a description of the fields and settings available on the **Add Rule** page and the **Modify Rule** page.

Note: The `readahead_procs` setting of the `ndmp-option` CLI command is no longer used for replications.

The number of additional server connections and the number of read-ahead processes should be coordinated to get the best performance. Each additional connection causes the creation of a separate process at the source and one at the destination, and these processes are connected by their own separate TCP connection. These two processes work together as an independent replication stream which can process subdirectories in parallel with other replication processes. Read-ahead processes are used only at the replication source; these processes pre-read directory entries and file details from the storage media (typically disks) so that the main replication processes can use them immediately without being delayed by disk read latencies.

Although allocating more processes to a replication can improve its performance, the extra processes take up system resources. Using these resources for replication operations may negatively impact the performance of other processes for protocols (such as NFS or CIFS), features, or even other replications. Also, the performance improvement per additional process reduces as the number of processes increases, and at some point there will be no further improvement (there may be a reduction in performance if too many processes are used). With these points in mind, you should not request a very high number of processes, except in very special cases.

The optimal settings for these values depend on many factors, including:

1. File system size and layout of files. Typically, to get best performance when replicating file systems with smaller files and fewer files per directory, you should dedicate more read-ahead processes and connections to the replication.

2. The number of replications that will run simultaneously. When running many replications simultaneously, each replication should be allocated fewer read-ahead processes so that the overall load from replication

processes is not too high.

3. The relative priority of the replication and other uses of the system (such as NFS and CIFS access). If replications appear to be adversely affecting the performance of system functions or user access, then reducing the number of read-ahead processes and connections used by replications should alleviate the problem.

4. The number of physical disks containing data for the file system (which can be found by looking at the number of physical disks in the system drives used by the file system). If the data of the file systems being replicated is stored on relatively few physical disks, then increasing the number of connections (data streams) used in the replication operation will not improve performance. Refer to the BlueArc Titan Server *Storage Subsystem Guide* for information on system drives.

5. The properties of the network route between the source and destination machines. When the connection between source and destination machines has high bandwidth available, long latency connections (high speed cross-continental or intercontinental links), then the long latency may impose an artificially low data rate over a single TCP connection. Using parallel connections (data streams) for the replication operation can improve performance in this case.

The following notes give more specific indications of how to choose settings. However, because of the many factors involved these recommendations may not be absolutely optimal. If it is vital to get the very highest performance from your system, some experimentation will be needed to find the most suitable values for the installation.

- The default settings are 4 additional connections and 12 read-ahead processes. These settings should be adequate for a wide range of applications, including files systems with mixed small and medium size files (average file size 250KB), up to file systems containing very large files. Also, these settings are appropriate for file systems with data on a relatively small number of physical disks (such as 32). A single replication operation with these settings should not severely impact other system users.
- If many replication operations are running simultaneously, you may want to reduce the number of read-ahead processes and connections each replication uses.

  For example, if you are running 8 simultaneous replications, the settings might be 1 additional connection and 6 read-ahead processes.
- Where the files in the file systems being replicated are mostly small, increasing the number of connections and read-ahead processes used will usually produce better performance.

  For example:

- For a file system with an average file size of less than 64KB, you may want to set 8 additional connections and 20 read-ahead processes.
- For a file system with an average file size of 32KB, you may want to set 12 additional connections and 24 read-ahead processes.

If the number of files per directory is also very low (less than 4 or 5 files per directory), even more connections and read-ahead processes might improve performance.

- If you are using a high speed cross-continental or inter-continental link, then using multiple connections may help utilize more of the bandwidth of the connection. In cases where latency is high due to the connection in use, it might be useful to increase the number of connections used, even when the average file size is large.

  The default TCP window size used by the server is 256KB. If the latency (round trip time) of the link is 70ms, then the maximum realistic throughput on a single TCP connection is about 3 MB per second.

  For instance, to get TCP connections capable of delivering 30MB/sec from a file system containing mostly large files, 10 additional connections and 12 read-ahead processes might be suitable settings. Note that, in this case, it is not necessary to increase the number of read-ahead processes, because reading from the source file system will not be a limiting factor.

- If the file systems involved have relatively few physical disks, increasing the number of connections and read-ahead processes will gain relatively little performance improvement.

  For instance, for a small source file system with data on only 32 physical disks, there will not be much to gain by increasing the values above the defaults.

Troubleshooting Replication Failures

This section provides troubleshooting tips, in the following sections:

- Manually Restarting a Failed Replication, on page 388
- Rolling Back an Incomplete Replication, on page 388

Listed below are some scenarios where a replication job can fail:

- The destination volume is offline.
- The destination volume was full.
- One of the volumes involved may have been unmounted.
- SMU was rebooted while a replication job was in progress.

Note: Without any further action upon a replication failure, the replication will continue as expected on its next scheduled run. However, this will recopy any changes already copied during the failed replication attempt.

Clicking **Restart** will cause the failed replication to be restarted immediately and will avoid recopying most of the data.

Manually Restarting a Failed Replication

If a replication has failed, the replication will be started normally at its next scheduled run time, rather than "picking up where it left off." To restart the replication from the point of failure (before its next scheduled time), you must restart it manually:

1. **Navigate to the Replication page.**

   From the **Data Protection** page, click **Replication** to display the **Replication** page.

2. **Restart the replication.**

   Click **details** for the failed replication to view its **Replication Schedule** page, then click **restart.**

Rolling Back an Incomplete Replication

Upon successful completion of a replication, the system takes a snapshot to preserve the state of the target file system. With this snapshot, if an offline source leads to failure of a subsequent replication, the target file system can be rolled back to the state of the last successful replication.

To rollback the target file system to the state of the last successful replication:

1. **Navigate to the Replication page.**

   From the **Data Protection** page, click **Replication** to display the **Replication** page.

2. **Roll back the replication.**

   Click **details** for the failed replication to view its **Replication** schedule page, then click **rollback.**

Note: Rollback should only be used when the target will be used as the live file system. If the replication's source file system cannot be used as the live file system (either permanently or temporarily), users can access the latest available data on the replication target (the file system created by the last successful replication). There are two possible approaches:

- **If the target file system will be used as the live file system permanently**, delete the replication policy and all related schedules (since the source will not be used for this replication again). You can then create new replication policies and schedules.
- **If the target file system will be used as the live file system temporarily**, contact SGI Global Services for assistance in synchronizing the "old" (source) and the "new" (target) file systems before transferring access and resuming replication operations as implemented prior to the "rollback".

Transferring Primary Access

A transfer of primary access relocates a file system, with very little down time, while the file system is live and servicing file read requests. For a short period, access is limited to read-only.

Transferring Primary Access

To transfer primary access:

1. **Navigate to the Replication page.**

   From the **Data Protection** page, click **Replication** to display the **Replication** page:

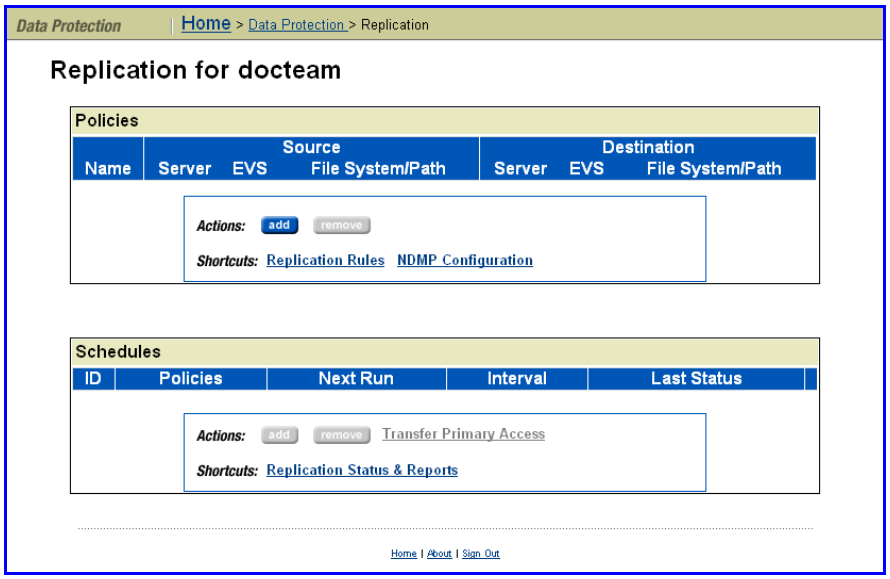

2. **Create a replication policy.**

   Refer to Policy Based Replication, on page 333.

3. **Create a schedule for the replication policy.**

   Refer to Adding a Replication Schedule, on page 377.

4. **Start an initial replication at a convenient time.**

Refer to Understanding Incremental Replications, on page 381 for more information:

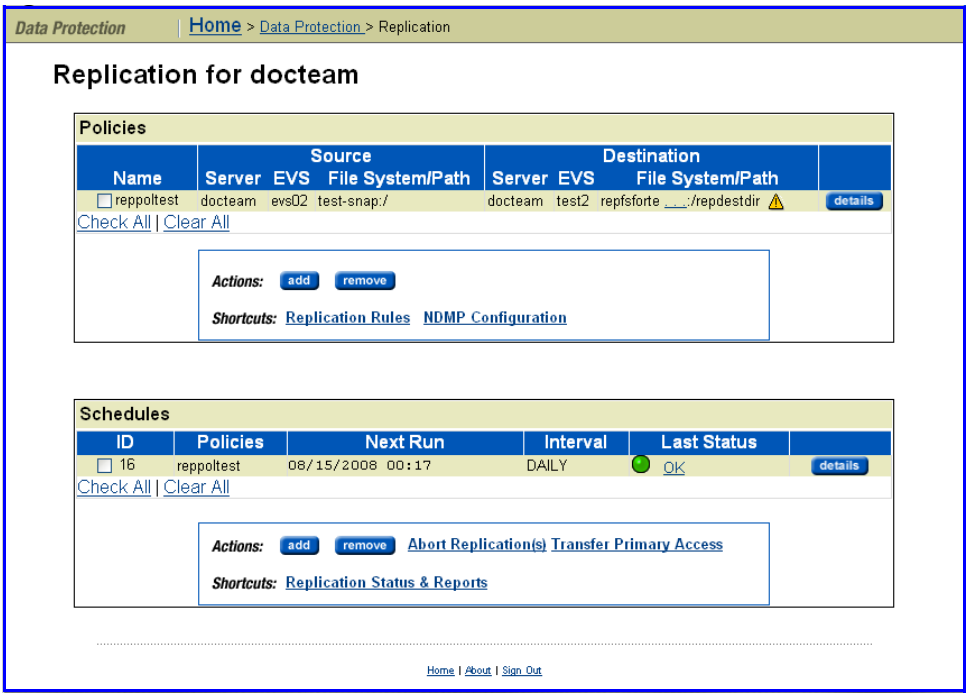

5. **Schedule at least one incremental replication.**

   When the incremental replication is complete, select the schedule and click **Transfer Primary Access** to display the **Transfer Primary Access** page, which displays a summary of the steps required for the transfer of primary access:

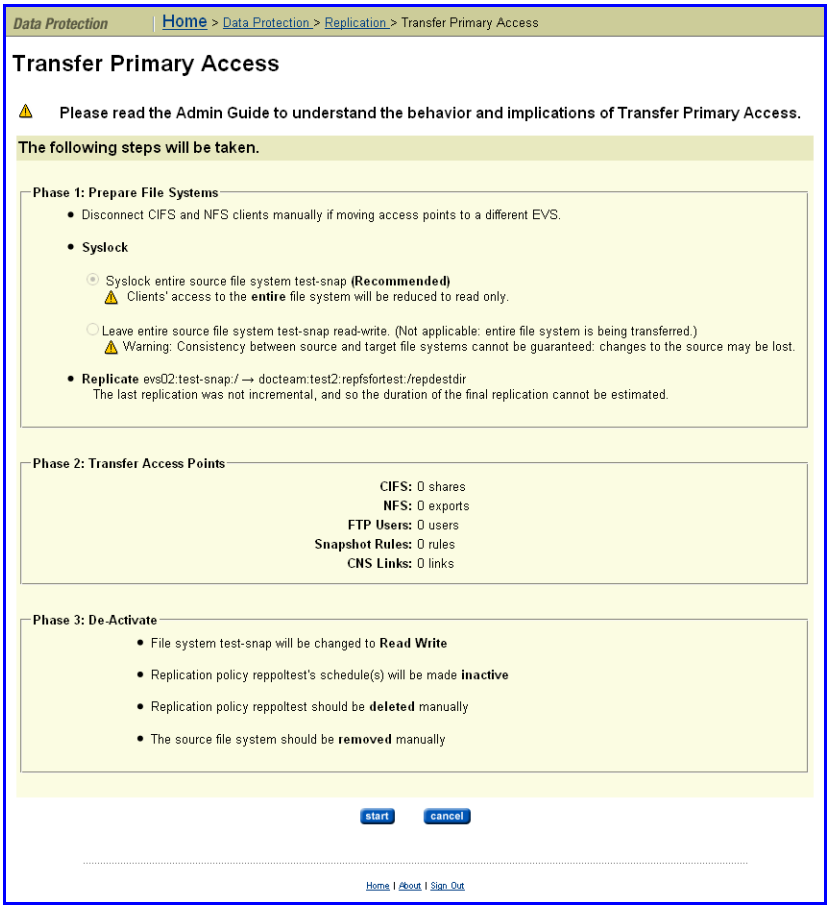

1. **Start the final replication.**

   Click **start** to begin the final replication and to display another **Transfer Primary Access** page, which displays the progress of the transfer of primary access:

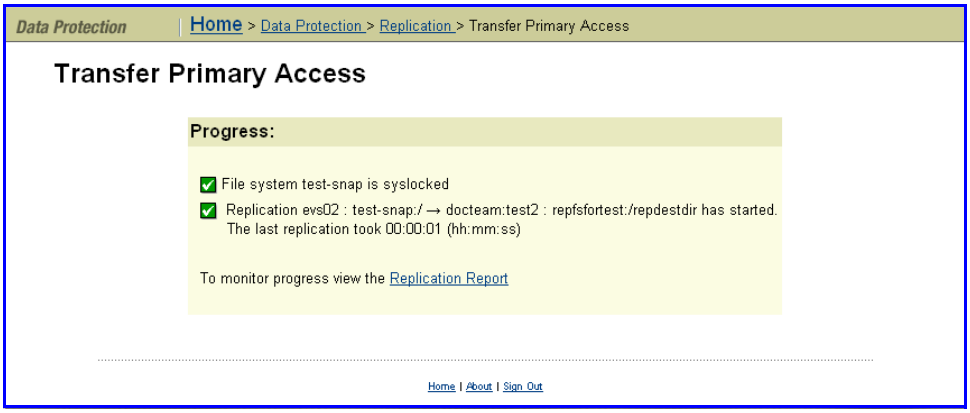

Note: If replication operation fails, the transfer of primary access operation is aborted.

2. **View progress.**

For more detailed information about the progress of the transfer of primary access, click **Replication Report** to view the latest replication report.

Note: If there is a failure during the transfer of primary access, the failure is handled as described in Handling a Failure during a Transfer of Primary Access, on page 392.

Handling a Failure during a Transfer of Primary Access

If a failure occurs during a transfer of primary access:

- The target file system is not brought online in place of the source.
- The source remains accessible and usable to network clients.
- There is no attempt to rollback after a failure.
- The SMU performs as many actions as possible but leaves the replication policy in place.
- A partially failed final replication does not remove the replication policy/schedule.
- The system administrator can usually resolve the issue that caused the failure, then run transfer primary access again.

For example, when replicating several CIFS shares, one share fails to be replicated, but the others are replicated successfully:

- The share that failed is logged to a simple text file (which is viewable from the **Replication Report** page).
- All other shares that were successfully recreated are brought online, and deleted from the source.
- When complete, the system administrator sees the error message and then views the text file using the **Replication Report** page).
- Viewing the text file, the administrator sees that the share could not be created on the target, perhaps because the name is already in use. The system administrator can delete the named share either from the source or the target, and then transfer primary access again, this time successfully.

Using Virus Scanning

Virus Scanning is enabled and configured at the EVS level. Only files accessed using the CIFS protocol are scanned. If a file has not been verified clean by a Virus Scan Engine, it requires scanning before access. As virus scanning can cause a delay when a client requires access, files are automatically queued for scanning as soon as they are closed (after creation or modification). Queued files are scanned promptly, expediting detection of viruses in new or modified files and making it unlikely that a virus infected file will remain dormant on the system for a long period of time.

You can configure multiple Virus Scan Engines to enhance both the performance and to maintain high-availability of the server. If a Virus Scan Engine fails during a virus scan, the storage server automatically redirects the scan to another Virus Scan Engine.

The server maintains a list of file types, the Inclusion List, that allows the administrator to control which files are scanned (for example, .exe, .dll, .doc, etc.). The default Inclusion List includes most file types commonly affected by viruses.

Caution: When virus scanning is enabled, the server must receive notification from a Virus Scan Engine that a file is clean before allowing access to the file. As a result, if virus scanning is enabled and there are no Virus Scan Engines available to service the virus scans, CIFS clients may experience a temporary loss of data access. To ensure maximum accessibility of data, configure multiple Virus Scan Engines to service each EVS on which virus scanning has been enabled.

If virus scanning is temporarily disabled, files continue to be marked as needing to be scanned. In this way, if virus scanning is re-enabled, files that were changed are re-scanned the next time they are accessed by a CIFS client.

Virus Scanning statistics for the storage server (in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

When a virus is detected, a severe event is placed in the Event Log, identifying the path of the infected file and the IP address of the infected machine. For information on accessing the event log, see Event Logging and Notification, on page 497.

Configuring Virus Scanning

To configure virus scanning, the following steps are required:

1. Configure the Virus Scan Engine(s).
2. Enabling anti-virus support on the server.
3. Optionally, disable virus scanning on selected CIFS shares.

Supported Platforms

For information about supported virus scan engines, contact SGI Global Services.

Notes on Installation and Configuration of a Virus Scanning Application

- Multiple Virus Scan Engines should be configured to enhance performance and high-availability of the server.
- When prompted, while installing a Virus Scan Engine, select to use the RPC protocol.

After installation and configuration has been completed, the Virus Scan Engine will automatically self-register with the server.

Enabling Virus Scanning on the Storage Server

To enable virus scanning on the server:

1. **Navigate to the Data Protection page.**

   From the **Data Protection** page, click **Virus Scanning** to display the **Virus Scanning** page:

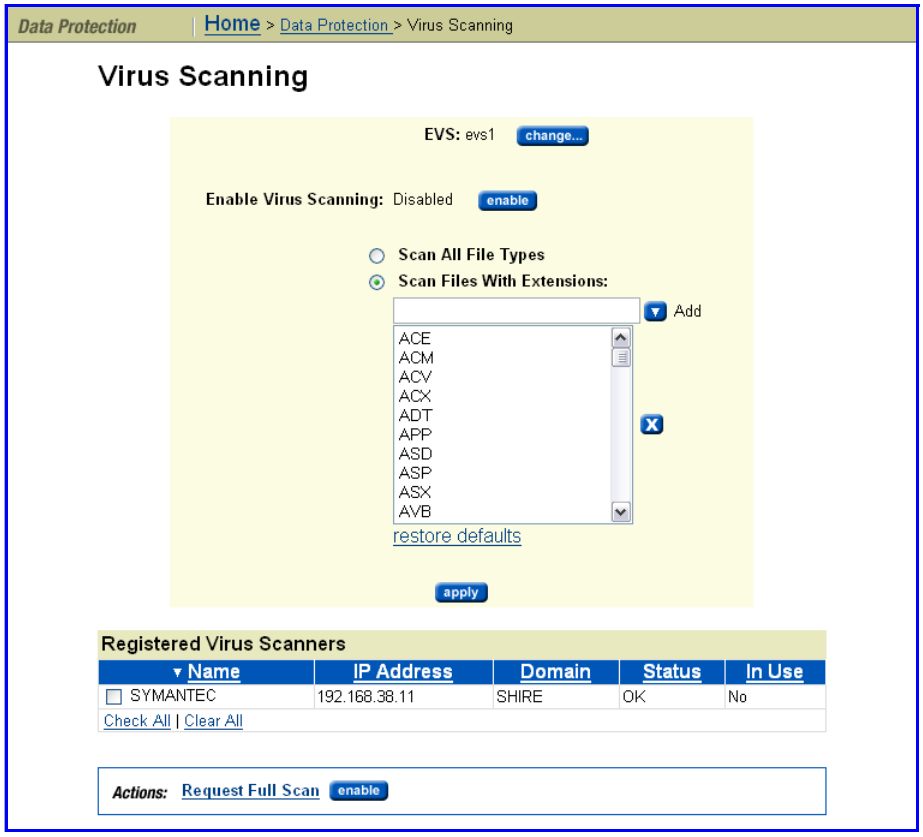

2. **Select the Virtual Server (EVS) on which to enable virus scanning.**

Caution: It is important that at least one virus scan engine is listed in the Registered Virus Scanners table. The account used to start the scanning services on the virus scan engine must be added to the server's Backup Operators Local Group. If the account used to start the antivirus service is not a member of the Backup Operators Local Group, the antivirus engine will not be registered and will not be displayed on the **Virus Scanning** page of Web Manager. If you try to enable virus scanning when no virus scanners have been registered, the SMU restricts the action; virus scanning cannot be enabled when there are no registered virus scanners.

3. **In the Enable Virus Scanning field, ensure that virus scanning is enabled. Click the enable button to enable scanning.**

Tip: Virus Scanning can be disabled on individual CIFS shares by

unchecking the Enable Virus Scanning box in the **Add Shares** page (File Services > CIFS Shares > Add Share).

4. **Optionally, modify the list of files to be scanned:**
   - *To scan all file types regardless of those in the list,* select **Scan All File Types**.It is advisable to select this option while compiling your list of file types to scan.
   - *To add a file type to scan*, click the **Scan Files With Extensions** radio button, enter the file extension in the field below it, then click **Add**.
   - *To delete a file type,* select it from the list, then click the **X**.
   - *To revert back to the original default list of files types to scan,* click **restore defaults**.

Caution: The default list of file extensions contains the most commonly used file types. SGI Global Services recommends that your anti-virus software vendor be contacted for an up-to-date list of file types that should be included for scanning, and to modify the your file extension list accordingly. It is your responsibility to choose the file types you include for scanning. Based on your needs, the antivirus software used, and the recommendations of the antivirus software manufacturer, choose the file types you want to include in the antivirus scanning; *types not listed will not be scanned*.

The default file extension list is as follows:

```
ACE, ACM, ACV, ACX, ADT, APP, ASD, ASP, ASX, AVB, AX, BAT,
BO, BIN, BTM, CDR, CFM, CHM, CLA, CLASS, CMD, CNV, COM, CPL,
CPT, CPY, CSC, CSH, CSS, DAT, DEV, DL, DLL, DOC, DOT, DVB,
DRV, DWG, EML, EXE, FON, GMS, GVB, HLP, HTA, HTM, HTML, HTT,
HTW, HTX, IM, INF, INI, JS, JSE, JTD, LIB, LGP, LNK, MB, MDB,
MHT, MHTM, MHTML, MOD, MPD, MPP, MPT, MRC, MS, MSG, MSO, MP,
NWS, OBD, OBT, OBJ, OBZ, OCX, OFT, OLB, OLE, OTM, OV, PCI,
PDB, PDF, PDR, PHP, PIF, PL, PLG, PM, PNF, PNP, POT, PP, PPA,
PPS, PPT, PRC, PWZ, QLB, QPW, REG, RTF, SBF, SCR, SCT, SH,
SHB, SHS, SHT, SHTML, SHW, SIS, SMM, SWF, SYS, TD0, TLB, TSK,
TSP, TT6, VBA, VBE, VBS, VBX, VOM, VS?, VSD, VSS, VST, VWP,
VXD, VXE, WBT, WBK, WIZ, WK?, WML, WPC, WPD, WS?, WSC, WSF,
WSH, XL?, XML, XTP, 386
```

5. **If necessary, re-enable the use of a disabled virus scanner.**

   If a virus scanner has been disabled for some reason, you can re-enable its usage by filling the check box next to the name of the disabled virus scanner and clicking the **enable** button in the **Actions** area.

6. **Save your changes.**

   Verify your settings, then click **apply** to save.

Forcing Files to be Rescanned

With the appearance of a new virus and release of anti-virus software updates, it is important to re-scan all files, including those that have not changed since the last time they were scanned.

1. **Navigate to the Virus Scanning page.**

   From the **Data Protection** page, click **Virus Scanning** to display the **Virus Scanning** page.

2. **Request a full scan.**

   Click the **Request Full Scan** link. This marks every file as unscanned, so the file will be scanned the next time it is accessed.

9 Scalability and Clustering

The IS-NAS Server/Titan Server architecture provides the following options for scalability and clustering:

Scalability/ Clustering Concept	Conceptual Overview	Associated Tasks
Virtual Servers (EVSs)	Virtual Servers (EVSs), on page 402	Using Virtual Servers (EVSs), on page 416
Secure Virtual Servers (Secure EVSs)	Secure Virtual Servers (Secure EVSs), on page 402 Secure EVS Considerations, on page 403 About Security Contexts, on page 404	Securing an EVS, on page 406 Removing an Individual Security Context From a Secure EVS, on page 408
EVS Name Spaces	EVS Name Spaces, on page 412 EVS Name Spaces Considerations, on page 412	Viewing the Cluster Name Space Tree, on page 439
Clusters and Server Farms	Clusters and Server Farms, on page 398	Using Clusters, on page 427 About Cluster Licensing, on page 427
Cluster Name Spaces	Cluster Name Space (CNS), on page 410	Using Cluster Name Space (CNS), on page 438
Cluster Read Caching	Read Caching, on page 412	Using Read Caching, on page 446

Overview

Administrators can configure a single physical server to act as a standalone server, or as a node in a cluster. The administrator can group several servers into a multi-node *cluster* or into a *server farm*.

- A **cluster** allows multiple physical servers to operate together as a single entity; sharing storage under the centralized management of a single SMU and using a common namespace.
- A **server farm** allows multiple standalone servers and/or clusters to be grouped together under the management of a single SMU, sharing a common pool of storage. Each server/cluster in the server farm operates independently of all other servers/clusters in the farm.

File services within the cluster or server farm are virtualized as Virtual Servers (EVSs), and any file service within the cluster or server farm can reside on, or be migrated to, any node within the cluster/server farm.

Clusters and Server Farms

The key differences between a cluster and a server farm are the behavior when a failure occurs, ease of management, and scalability of operations:

- A **cluster** allows an EVS to be automatically migrated among cluster nodes in the event of a failure. Management of all nodes in the cluster is centralized, and the usage of a single namespace allows clients to mount a single network resource, while having the actual storage virtualized among different devices attached to the cluster.
- A **server farm** allows an EVS to be migrated manually among servers in the server farm, but this is a manual process, and it does not happen automatically in the event of a failure. All servers in the server farm may be managed by a single SMU, but each server must be managed as an independent unit.

Clusters

Clustering provides the following functionality:

- **Simultaneous hosting of multiple EVSs.** Nodes in a cluster can simultaneously host multiple EVSs, allowing all servers to be active at the same time, each providing file services to clients.
- **Redundant monitoring and transparent failover of EVS hosts.** The cluster monitors the health of each server through redundant channels. Should one server fail, one of the other cluster nodes can take over its functions transparently to network clients, so no loss of service will result from the failure.
- **Redundant availability of configuration settings for all nodes.** The cluster provides a cluster-wide replicated registry, containing configuration information for all nodes in the cluster.

The following sections discuss options for configuring server nodes as clusters, in order to expand their functionality.

Cluster Nodes

Multiple EVSs can be hosted simultaneously by all cluster nodes (referring to individual nodes within a cluster). Clustering frees EVSs from a one-to-one relationship with a cluster node. IP addresses are associated with the logical EVS, rather than with physical node. Failover or migration of an EVS from one cluster node to another is transparent to network clients.

N-Way Clustering

Currently, a single cluster can contain up to eight servers. Within clusters, each node may host one or more independent EVSs. Upon failure of one of its nodes, the cluster automatically migrates its EVSs to other cluster nodes,

which then become their hosts. Network clients will not typically be aware of the failover and will not experience a loss of service. After the failed node is restored and is ready for normal operation, previously hosted EVSs can be migrated back. During restoration of the failed node, the cluster may operate with reduced performance.

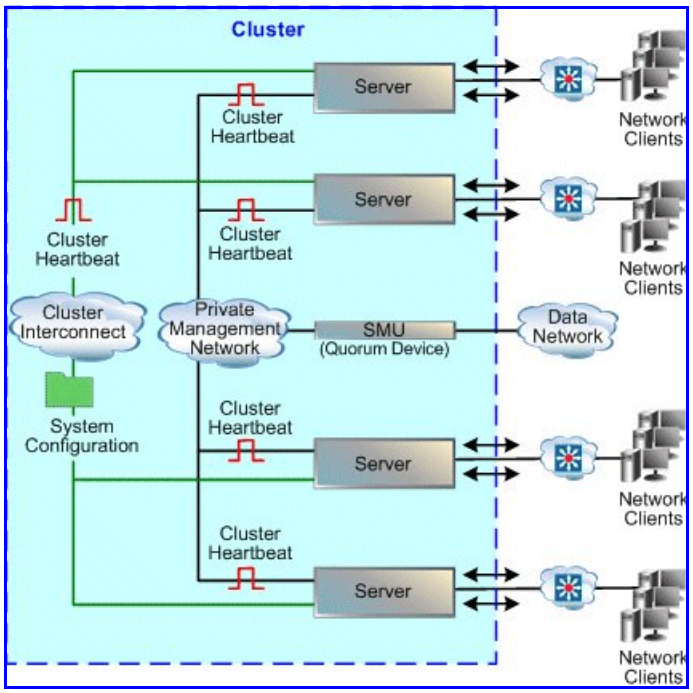

The Quorum Device (QD) in a Cluster Configuration

The Quorum Device (QD) runs on the System Management Unit (SMU), which can provide QD services for up to eight clusters (or up to eight servers in a server farm). The QD enables a cluster to maintain operations following a communications failure between nodes and also to restore the cluster registry (containing the cluster configuration), as follows:

- **Surviving a communication failure between nodes.** Clustering preserves data integrity through a *quorum voting* algorithm that ensures only one node can access a given file system at any time. Under this algorithm, each of the cluster nodes may "vote" regarding file access. When a cluster contains an even number of nodes, the QD also votes. When a cluster node has obtained a *quorum* (a simple majority of the votes available in the entire cluster) it receives exclusive access to the file system. Under certain failure scenarios, cluster nodes can lose communication with each other and may attempt to access the same file system; in this situation, the QD alone "votes" for one of the nodes, establishing the quorum and granting one node exclusive access to the file system.
- **Preserving a copy of the cluster registry.** Although the registry is replicated across cluster nodes, some failure scenarios could result in the loss of recent configuration changes, a condition called *amnesia.* Anticipating the possibility of such a condition, the QD preserves a copy

of the registry, ensuring that configuration changes can always be replicated.

Cluster Topology

Typically, the private management network connects cluster nodes and the QD, keeping cluster traffic off of the public data network, and isolating them from potential congestion due to heavy data access loads.

The high-speed *cluster interconnect* provides an additional, direct connection among the cluster nodes. This dedicated connection consists of dual redundant Gigabit Ethernet (GE) links, and is reserved for clustering traffic and NVRAM mirroring.

Note: Setting up a cluster requires a license. Contact SGI Global Services to purchase a cluster license. For more information on how to enter the license key(s), see Adding a License Key, on page 538.

Server Farms

A typical Server Farm contains at least two standalone servers and/or standalone servers and at least one cluster:

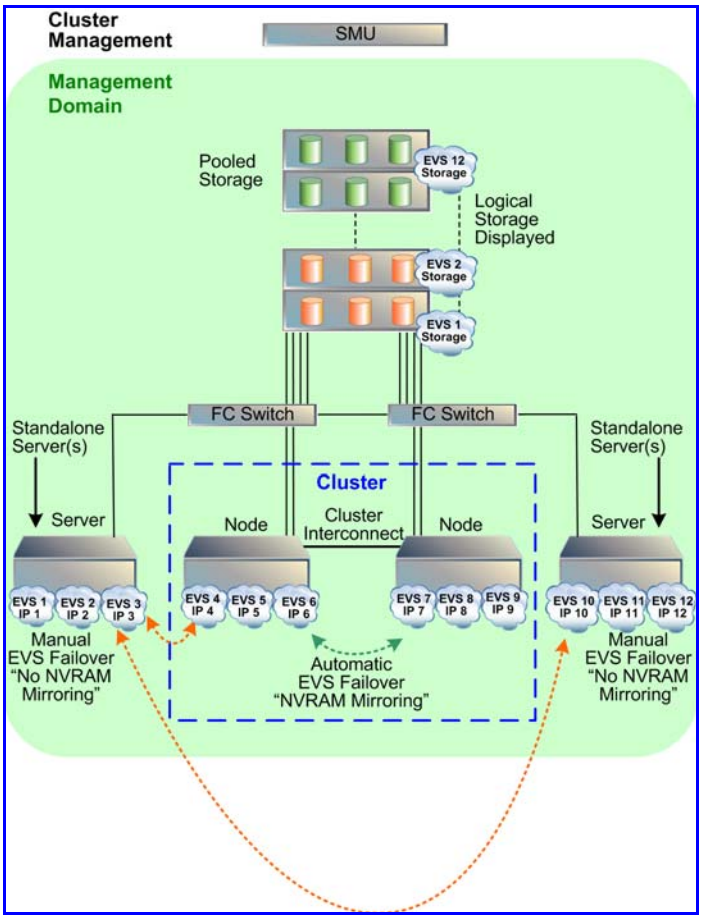

A single SMU (System Management Unit) manages every server and cluster within the server farm. The SMU hosts the management network for the server farm and provides quorum services for up to eight clusters. Managed devices must be located in a single data center, not distributed across a campus or MAN environment.

The server farm offers the following functionality:

- **Optimizing performance.** For maximum throughput, migrate EVSs to a higher-end server, or to a fully dedicated server.
- **Balancing load.** For more efficient use of available resources, migrate heavily used EVSs to less-busy servers, or to higher-end servers that support greater capacity.
- **Redundant failover.** In the event of a catastrophic failure of any standalone server, the EVSs hosted by the failed server can be brought online on any other server or cluster in the server farm.

When configured together as a server farm, standalone and cluster nodes share common access to the same storage subsystem, ensuring that when EVSs move from one node to another, whether due to an automatic failover or manual migration of EVSs among servers, the target server has access to the storage served by the EVS.

Clusters versus Server Farms

The following table distinguishes the properties of a cluster and a Server Farm:

Property	Cluster	Server Farm
Can belong to a Server Farm	Yes	No
EVS migration under server failure	Automatic	Manual
NVRAM mirroring between servers	Yes	No
Maximum number of storage servers	8	No explicit restriction on the number of servers; however, an SMU can manage only 8 Quorum Devices and server farm planning should be adjusted accordingly.
Shared SMU	For central management; cluster quorum	For central management; EVS Migration
Storage Pools	Yes	No
Common Storage Access	Yes	Yes

Virtual Servers (EVSs)

EVSs appear to network clients as actual file servers. Like a physical server, an EVS has IP Addresses, supports CIFS shares and NFS exports for file sharing, and contains file systems. Network clients access EVSs as independent servers, while the administration of an EVS is localized to the host server or cluster. A single physical node supports up to 64 EVSs.

To balance workload and increase availability of file services, EVS can be migrated among servers in a Server Farm or nodes in a cluster. There are two distinct methods of configuring this migration capability:

- In a Server Farm, EVSs hosted on any server in the Server Farm can be migrated to any other server or cluster in the Server Farm. Note, however, that EVS migration in a Server Farm across non-clustered servers requires an EVS Migration license.
- In a Server Farm, EVSs hosted on any server in the Server Farm can be migrated to any other server or cluster in the Server Farm. Note, however, that EVS migration in a Server Farm across non-clustered servers requires an EVS Migration license.

Secure Virtual Servers (Secure EVSs)

A secure Virtual Server is a file serving EVS that has a specifically defined security configuration (called an *individual* security context). When no individual security context is specified for an EVS, it uses the *global* (server or cluster-wide) security configuration settings (the global security context). By defining an individual security context for a particular EVS, you create a secure Virtual Server (secure EVS).

Note: Secure Virtual Servers are a licensed feature, identified as EVS Security. Without an EVS Security license, all EVSs use the global security settings (context). See Adding a License Key, on page 538 for more information about adding a license key.

- When no individual security context is defined for an EVS, the global security settings (the global context) are used by default.

  Note: When an individual security context is added to an EVS, the new individual security context is created using the same settings as are used by the global security context. After adding the individual security context, you can then change settings to make the individual security context settings different than the global settings.

- When using an individual security context, the EVS security context can be configured independently of the global (server or cluster-wide) security settings.

When present, individual security context settings override the global security context settings, allowing a storage server (or cluster) to be shared by multiple groups (departments, customers, or organizations), while maintaining strong security so that no group has access to another group's data.

For example, if a server/cluster has six EVSs, you could define individual security contexts for two of the EVSs, turning them into secure EVSs. Each

secure EVS could then be associated with an NT domain that is different than the one used by the cluster, meaning that each of those secure EVSs could be assigned to its own domain. For network clients, access to the file systems in the secure EVSs can then be restricted or allowed as desired using standard network security policies such as user name or user group membership.

Secure EVS Considerations

When using secure EVSs, keep the following points in mind:

- **Security context defaults.** Unless an individual security context is specified for an EVS (making it a secure EVS), the EVS security context defaults to the global security context.
- **Inherited global settings.** NDMP user name and password settings are not EVS-specific; the same NDMP user name and password settings apply to all EVSs and secure EVSs in a server/cluster.
- **Secure EVS-specific security settings.** Once an EVS has a defined individual security context, it becomes a secure EVS, and each secure EVS is considered to be separate from all other EVSs and secure EVSs in the server/cluster.

  A secure EVS is always treated as an individual unit, regardless of if it uses the same security context settings as another secure EVS or if it uses different security context settings. As a result, different secure EVSs cannot share anything, including an individual EVS name space (see EVS Name Spaces, on page 412 for more information on EVS name spaces).
- **Secure EVS migration.** When a secure EVS migrates to a different cluster, it retains all specified security settings in its individual security context. If, however, a secure EVS is configured to use default settings from the global context, after migrating, the secure EVS switches to use the settings in the global context of the cluster to which it migrates.
- **Moving file systems between secure EVSs.** A system administrator with sufficient privileges can move a file system from one secure EVS to another, but a warning is issued if the security contexts of the source and destination secure EVSs are different.
- **External name server access.** Each secure EVS can be configured to connect to several external name servers, and each secure EVS can connect to different name servers.
- **Secure EVSs and name spaces.** Links from the cluster's CNS tree to a secure EVS are supported, according to the following rules:
  - **Accessing the CNS.** Only a secure EVS that uses the global security context can access links in the CNS.

- **CNS links to a file system hosted by a secure EVS with an individual security context are not allowed.** In the CNS, you cannot add a link to a file system hosted by a secure EVS.

  Similarly, you cannot configure an individual security context for an EVS (turning an EVS into a secure EVS) if there are CNS links to one or more file systems in that EVS.

- **Name space usage and the secure EVS.** If you want to use a name space with a secure EVS that does not use the global configuration settings, you must configure an EVS name space for that secure EVS. An EVS name space is required because file systems hosted by the secure EVS cannot be linked to from the CNS, and file systems hosted by the EVS cannot access links in the CNS. See EVS Name Spaces, on page 412 for more information on EVS name spaces.

  If you want to use a name space with a secure EVS that does use the global configuration settings, you may configure an EVS name space for that secure EVS, but it is not required. If the secure EVS does use the global security settings, the file systems hosted by the EVS can access links in the CNS.

- **Links to a secure EVS individual name space.** In a secure EVS with an individual name space, you can add links between file systems hosted by the same secure EVS (see EVS Name Spaces, on page 412 for more information on EVS name spaces).

- **Configuring a group of EVSs with the same settings.** To create a group of secure EVSs that use the same individual security context settings (that are different from the global settings), you must configure each secure EVS in the group separately.

- **Reconfiguring a secure EVS security context to use the global context.** If a secure EVS is re-configured to use the global security context (reverting it to an EVS), **and** the secure EVS was using a different NT domain than the cluster, CIFS names and CIFS share names become invalid. This occurs because CIFS names (and CIFS share names) are associated with a specific NT domain, and the NT domain name changes.

  If the global security context and the secure EVS and the NT domain are different, after you remove the secure EVS' individual security context (making it an EVS again), you must delete all CIFS names for the EVS and all CIFS shares for the file systems in the EVS. Then, you must recreate the EVS CIFS names and the CIFS shares for the file systems in the EVS.

About Security Contexts

Because EVSs and secure EVSs inherit many of their settings from the cluster's global context, when configuring name services, you must specify if you want to change the global context or the individual context.

For example, on the **NIS/LDAP Configuration** page or the **EVS Details** page, the current security context displayed as follows.

The EVS security context can be any of the following:

- **Global Configuration**, which indicates that the current security context is the global context.
- **Inherits Global Configuration**, which indicates that the current security context is an individual EVS security context that has been set to use the settings that are defined in the global security context.
- **Individual Configuration**, which indicates that the current security context is an individual context that has individually specified settings.

On the **Name Services** and **NIS/LDAP Configuration** pages, click **change** to switch the current context between an individual context for a particular secure EVS and the global context.

If you make changes that affect:

- The global context, those changes apply to all EVSs that have security context settings set to "Inherits Global Configuration."
- An individual context, those changes apply only to the currently selected EVS.

On the **EVS Details** page, click **change** to switch the context used by the EVS:

- If an individual context is being used, you can change the EVS to use the global context, removing the individual security context and changing the secure EVS into a regular EVS.
- If the global context is being used, you can change the EVS to use an individual context (creating a secure EVS).

Security Context Contents

The following security settings make up the security context for all EVSs and secure EVSs:

- Name services (DNS, WINS, and so on)
- Windows NT domain
- NIS domain
- User/group/domain mapping tables
- Security mode (mixed or UNIX)
- Individual groups
- "cifs-auth" setting
- "bypass-permissions-checks" setting

- File/directory umasks
- NFS export and CIFS share access options
- CNS mount points

The parts of the security context that can be configured for a secure EVS include:

- Name services (DNS, WINS, and so on)
- Windows NT domain
- CNS mount points

Securing an EVS

To change an EVS into a secure EVS, you must add an individual security context (see Secure Virtual Servers (Secure EVSs), on page 402 for more information about secure EVSs). To add an individual security context to an EVS:

1. **Make sure the EVS Security license key is installed.**

   Before you can specify an individual security context for an EVS, the secure EVS license must be installed, for information on installing the EVS Security license, see Managing License Keys, on page 535.

2. **Navigate to the EVS Management page.**

   From the **Home** page, click **EVS Management** to display the **EVS Management** page.

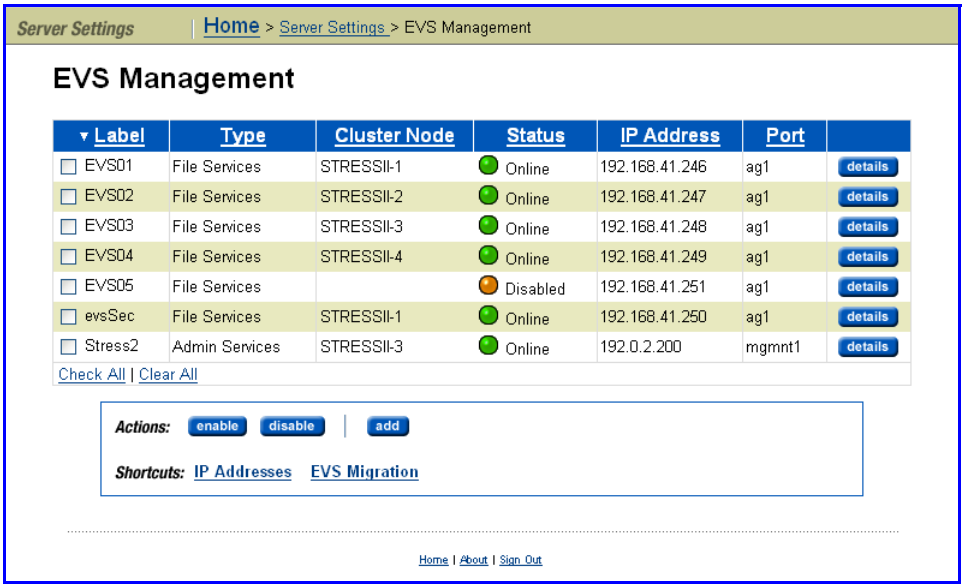

3. **Disable the EVS.**

   Fill the checkbox for the EVS you want to disable and click **disable**.

4. **Navigate to the EVS Details page.**

To display the **EVS Details** page, click the **details** button next to the EVS you want to change into a secure EVS.

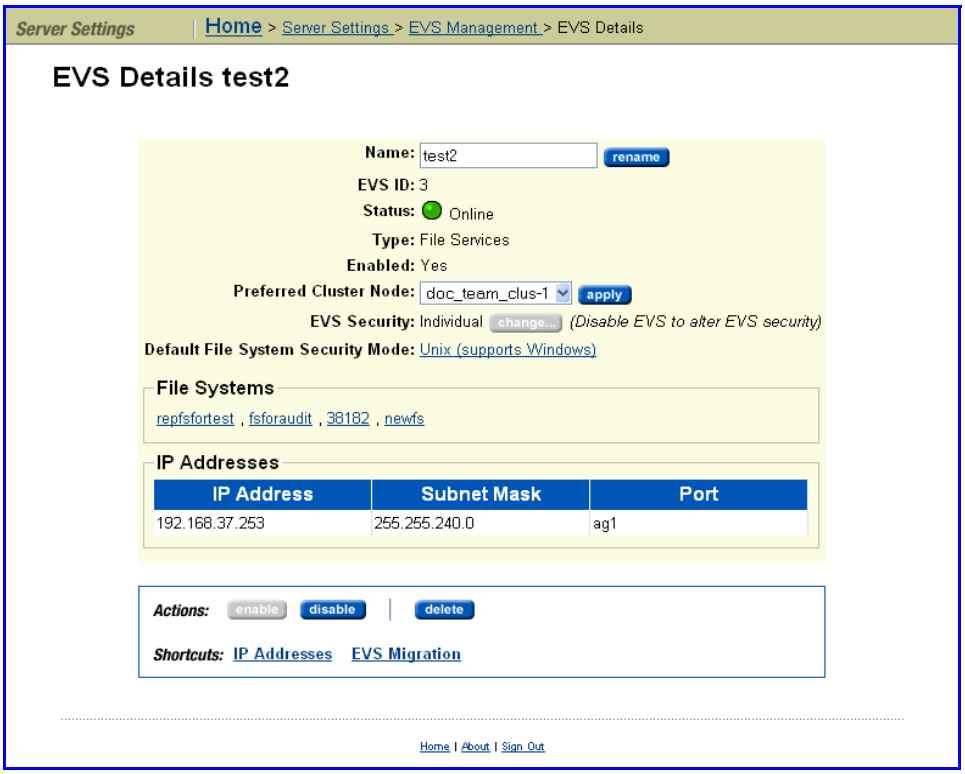

5. **Enable EVS Security.**

   To add an individual security context to the EVS described on the **EVS Details** page, click **change**.

6. **Confirm the change.**

   A warning page appears, reminding you of the consequences of enabling an individual security context and prompting you to confirm that you want to enable an individual security context for this EVS.

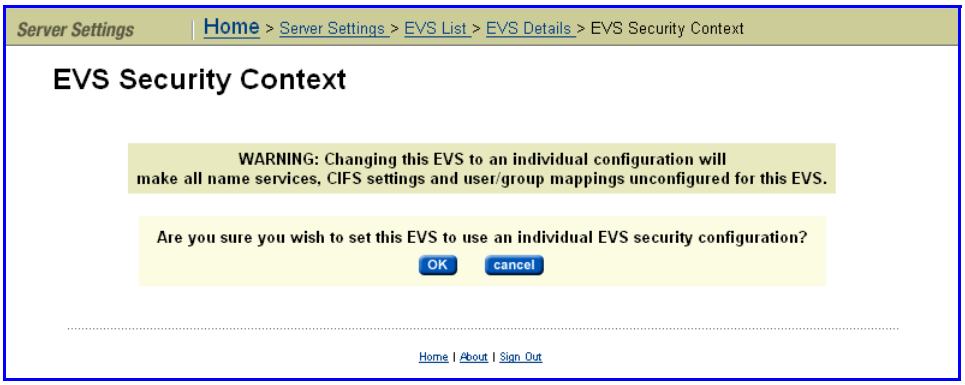

- To confirm the change, click **OK**.
- To cancel without making the change, click **cancel**.

You will return to the **EVS Details** page.

7. **Enable the EVS.**

   Select the EVS and click **enable**.

8. **Recreate CIFS names for the secure EVS.**

   For information on specifying a CIFS name for the secure EVS, see Configuring CIFS Security, on page 255.

9. **Adjust CIFS Shares for the file system(s) in this secure EVS.**

   For information on recreating CIFS shares for the file system(s) in the secure EVS, see Configuring CIFS Shares, on page 268.

10. **Specify user/group access for this secure EVS.**

    For information on configuring user and group access to file system(s) in the secure EVS, see Configuring Local Groups, on page 264.

11. **Configure name services for this secure EVS.**

    For information on setting up name services for the secure EVS, see Configuring Name Services, on page 94.

12. **If necessary, configure the EVS name space.**

    If you want to use a name space with a secure EVS that does not use the global configuration settings, you must configure an EVS name space for that secure EVS. For information on setting up an EVS name space, see EVS Name Spaces, on page 412.

Removing an Individual Security Context From a Secure EVS

Removing an individual security context from an EVS changes the secure EVS back into an EVS (see Secure EVS Considerations, on page 403 for more information about secure EVSs). To remove an individual security context from an EVS:

1. **Navigate to the EVS Management page.**

   From the **Home** page, click **EVS Management** to display the **EVS Management** page.

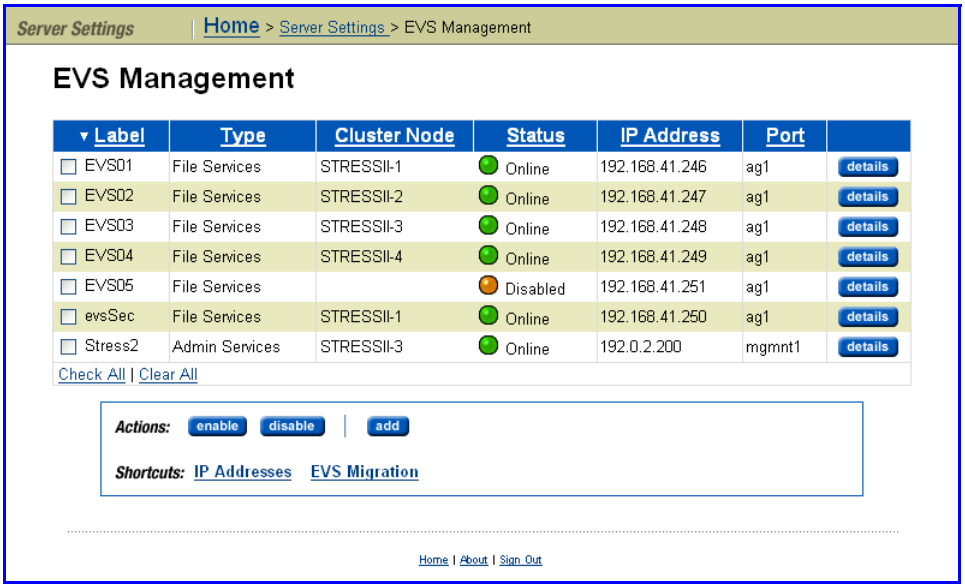

2. **Disable the secure EVS.**

   Fill the checkbox for the secure EVS you want to disable and click **disable**.

3. **Navigate to the EVS Details page.**

   To display the **EVS Details** page, click the **details** button next to the secure EVS you want to change back into an EVS.

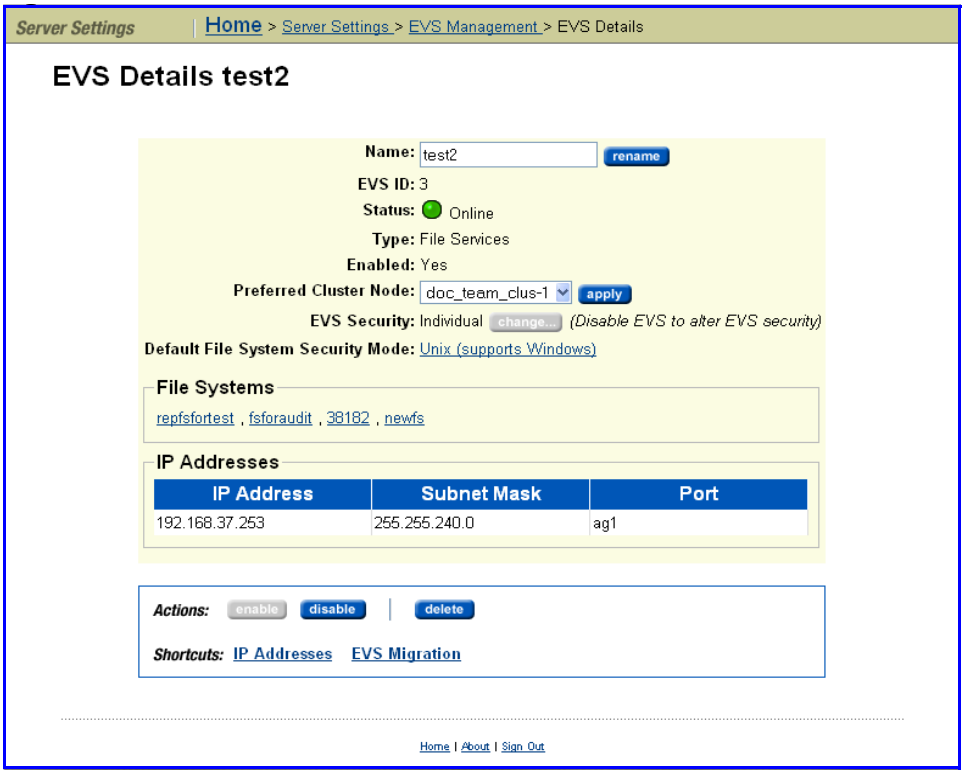

4. **Disable EVS Security.**

To remove an individual security context from the secure EVS described on the **EVS Details** page, click **change**.

5. **Confirm the change.**

   A warning page appears, reminding you of the consequences of removing the individual security context and prompting you to confirm that you want to remove the individual security context for this secure EVS.

   - To confirm the change, click **OK**.
   - To cancel without making the change, click **cancel**.

   You will return to the **EVS Details** page.

6. **Enable the EVS.**

   Click the **enable**.

Cluster Name Space (CNS)

A *Cluster Name Space* (CNS) allows multiple separate file systems on a server to appear as subdirectories of a single logical file system (that is, as one unified file system). They can also make multiple storage elements on that server available to network clients through a single CIFS share or NFS export.

The root directory and subdirectories in the CNS tree are virtual directories. As in a file system, the root occupies the highest position in the CNS tree and subdirectories reside under the root. Access to these virtual directories is read-only. Only the server's physical file systems support read-write access. Physical file systems can be made accessible under any directory in the CNS tree by creating a file system link. File system links associate the virtual directory in the CNS tree with actual physical file systems.

Any or all of the subdirectories in the CNS can be exported or shared, making them (and the underlying physical file systems), accessible to network clients. Creation and configuration of a CNS can be performed through the Web Manager or the CLI.

Once shared or exported, a CNS becomes accessible through any EVS on its server or cluster; therefore, it is not necessary to access a file system through the IP address of its host EVS and, in fact, file systems linked into the CNS can be relocated between EVSs on the server or cluster transparently and without

requiring the client to update its network configuration. This can be useful in distributing load across cluster nodes.

The simplest CNS configuration is also the most common. After creating the root directory of the CNS, create a single CIFS share and NFS export on the CNS root; then, add a file system link for each physical file system under the root directory. Through this configuration, all of the server's storage resources will be accessible to network clients through a single share or export, and each file system will be accessible through its own subdirectory.

Windows and Unix clients can take full advantage of the storage virtualization provided by CNS, because directories in the virtual name space can be shared and exported directly.

Tip: *Best practice to add FTP Mount Points or iSCSI Logical Units to non-CNS file systems.* For the best results, FTP mount points and iSCSI Logical Units should be added to file systems that are not part of a CNS, as CNS does not support FTP mount points or iSCSI Logical Units. Because FTP clients and iSCSI Initiators communicate directly with individual EVSs and their associated file systems, connectivity for any file system containing FTP mount points or iSCSI Logical Units must be reestablished through a new EVS upon relocation.

Note: CNS is a licensed feature. To create a Cluster Name Space, a CNS license must be installed. To purchase a CNS license, please contact SGI.

CNS Usage Considerations

The following recommendations are intended to simplify configuration and maintenance for CNS and for transfers of primary access for the file system:

- A single name space is supported per server or cluster.
- If there is only one CNS link to the file system, and no CIFS shares/NFS exports on the file system, only a single link has to be moved during a transfer of primary access.
- CNS does not support *hard link*s or *move operations* across the individual file systems. These operations are fully supported, but only within a single physical file system; that is, the part of the CNS tree under a file system link.
- Relocating file systems under the CNS may interrupt CIFS access to the file system being relocated. To minimize interruption, relocate file systems when they are idle. For more information, see Relocating a File System, on page 154.
- When using CNS and EVS together:
  - Only one EVS per cluster node is required for all data inside the cluster name space. Having additional EVSs causes unnecessary administrative overhead, and may lead to confusion. Use multiple EVSs on the same cluster node only when you have data that should reside outside the cluster name space.

- Balance loads by moving file systems, instead of migrating EVS. If you migrate an EVS containing a read cache, the files in the read cache become invalidated and, assuming they are still cacheable, they would have to be cached again after the next read request.

  If an EVS containing a read cache is migrated to another cluster node that already has a read cache, the files in the migrated read cache are invalidated, and only the read cache that was not migrated will be used. If the EVS is migrated back to its original cluster node, the read cache will be used again, assuming another read cache has not been created on that cluster node in the interim.

- When using CNS, the recommended configuration is to have a single CIFS share or NFS export at the root of the name space. If that configuration does not suit your needs, the next-best configuration is to have CIFS shares/NFS exports pointing to individual directories in the name space. You should not configure CIFS shares or NFS exports pointing to a path of the real file system unless absolutely necessary.

EVS Name Spaces

An *EVS name space* allows separate file systems within a virtual server (EVS) to appear as subdirectories of a single logical file system (that is, as one unified file system). An EVS name space can also make multiple storage elements on the virtual server available to network clients through a single CIFS share or NFS export.

The EVS name space functions in the same way as the Cluster Name Space (CNS), except that its context is that of the EVS, instead of the cluster.

In order to create an EVS name space, you must have installed a CNS license, and an EVS Security license, and you must have set the EVS to use an individual security context (see Securing an EVS, on page 406 for more information).

EVS Name Spaces Considerations

Linking to and from an EVS name space has the following constraints:

- **Links within an EVS name space.** In an EVS name space tree, you can add links from the EVS name space to file systems hosted by the same secure EVS.
- **Links between the CNS and the EVS name spaces.** The contexts of the CNS Name Space and the EVS name space are mutually exclusive: links from one to the other are not allowed.
- **Links outside the EVS name space.** Links from the individual EVS name space to file systems in other EVSs are not supported.

Read Caching

Read caching allows creation of a *read cache*, a special read-only file system that stores copies of individual files outside of their local file systems, enabling a server or a node to have a cached copy of the file. When NFS v2 or NFS v3

clients submit a read request for a file in the read cache, the server/node can serve the read request from the copy in the *read cache.* Note that a read cache does not benefit CIFS clients.

Note: The Read Caching, N-Way Clustering, and the Cluster Name Space (CNS) features are separately licensed features of the Titan Server storage system.

There are two types of read caching; **local** and **remote**.

- **Local read caching** caches files from a file system on the server/node to which the client has connected.

  With a read cache, when a clients reads a file through an export with local read caching enabled, the file is copied into the node's read cache (if permitted by the local read caching configuration). Subsequent reads of the file are then serviced from the read cache.

  Local read caching is not supported for NFSv4 clients.

  Local read caching provides performance benefits when:

  - Accessing migrated files (through cross file system links). Performance is improved because the data for such files is usually located on lower performance storage, or on an external server.
  - The original copy of a file being cached is stored in a file system that is on slower storage than the read cache.

    For example, if the read cache is on a tier 1 (high-performance) storage device and the original copy of a file in the cache is on a tier 3 (nearline) or tier 4 (archival) storage device. Another example would be if the read cache is on a solid-state storage device and the original copy of the files being cached is on a disk-based storage device. Refer to the *Storage Subsystem Guide* for more information on tiered storage and supported storage devices.
  - File access contention is reduced. Because each EVS serves read requests for the file from its own local copy in the read cache, multiple read requests for the same file can be served.

- **Remote read caching** caches files read through CNS links when the file system being accessed is on a different (remote) cluster node from the server/node to which the client has connected. Note that remote read caching can be enabled only on clusters configured with CNS.

  When remote read caching is not enabled, when an NFS client reads a file that is not hosted by a file system hosted by the node to which the client has connected, the read request is forwarded to the remote server that hosts the file system, which processes the request and sends the response to the original node, which in turn responds to the client. When

subsequent read requests for the same file are received from NFS clients, the same process is followed for every request:

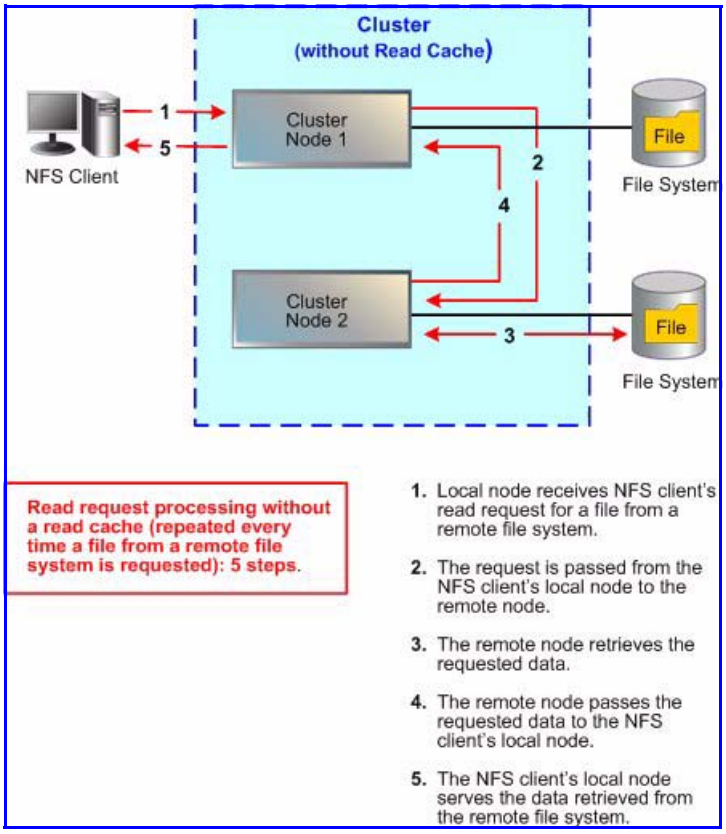

When remote read caching is enabled, on the first request by an NFS client for a non-local file, the request is forwarded to the remote cluster node, which retrieves the data and sends it to the requesting client. If the file can

be cached, the local node stores a copy of the file in the read cache. Upon subsequent read requests, that file is served from the read cache.

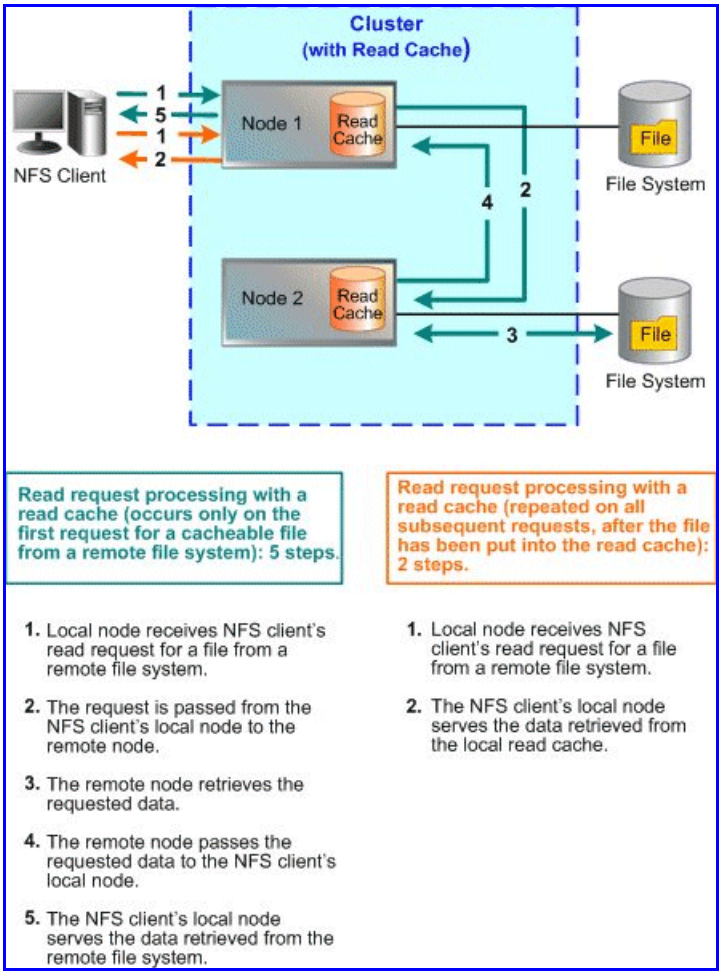

When files from a file system hosted on a remote cluster node are cached, read caching provides the following performance benefits:

- **File access contention is reduced.** Each EVS serves read requests for the file from its own local copy.
- **Interconnect traffic is reduced.** After the first read, subsequent requests do not cause the same data to be transported repeatedly across the cluster interconnect.
- **Files are served more quickly.** The read cache eliminates the time it would take for the local EVS to route a read request to the EVS on the remote cluster node, to receive the reply, and to forward the data to the client.

Read caching is most effective with files that do not change often, such as files containing programming libraries, executable code, graphics, audio, video, or any other files with relatively static content.

If a file has been cached, and the "original" version of that file is modified at its original location, the cached copy is invalidated, and **no further reads of**

the invalidated cache copy are allowed. For a period of time (specified as a file caching option), the file cannot be cached again. Once the time period in the file caching option has expired, the file can be cached. When another read request is received, the invalidated file in the read cache is updated with the changes, and subsequent read requests are served from the updated copy in the read cache.

The *read caching service* provides configurable file caching options, which allow you to control which files can be cached and which files are not eligible for caching, the content of the file caches, the number of files in the cache, and how long a file can remain in the cache without being accessed. For more information on the available file caching options, see Setting File Caching Options, on page 447.

Using Virtual Servers (EVSs)

A server node supports up to 64 EVSs. Likewise, a cluster can have up to 64 EVSs. EVSs can be added, deleted, and changed based on the evolving needs of the network.

EVS Configuration

Before they can be used, EVSs must be created and assigned to an IP address; then, to provide file services, assign one or more file systems.

Creating an EVS

1. **Navigate to the add EVS page.**

   From the **Server Settings** page, select **EVS Management,** then click **Add EVS** to display the **Add EVS** page:

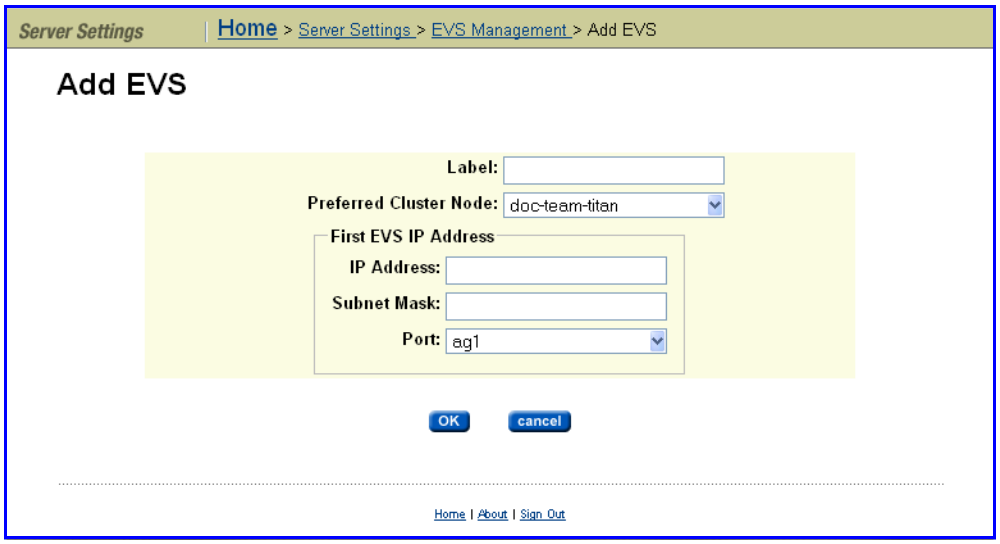

2. **Enter the requested information.**

   All fields required.

3. **Save your changes.**

   Verify your settings, then click **OK** to save or **cancel** to decline.

Assign a File System to an EVS

After the EVS has been created, at least one file system must be assigned to it. You can either create a new file system on the EVS (see Creating a File System, on page 138) or you can assign an existing file system to the EVS. To assign a file system to an EVS, you can relocate a file system currently assigned to another EVS (see Relocating a File System, on page 154) or you can assign a file system that is currently not assigned to an EVS. To assign a currently unassigned file system to an EVS:

1. **Navigate to the Server Settings page.**

   From the **Storage Management** page, click to display the **File Systems** page.

2. **Select the file system.**

   In the file system grid, for the specific file system that will be assigned to the EVS, click **details** to display the **File Systems Detail** page:

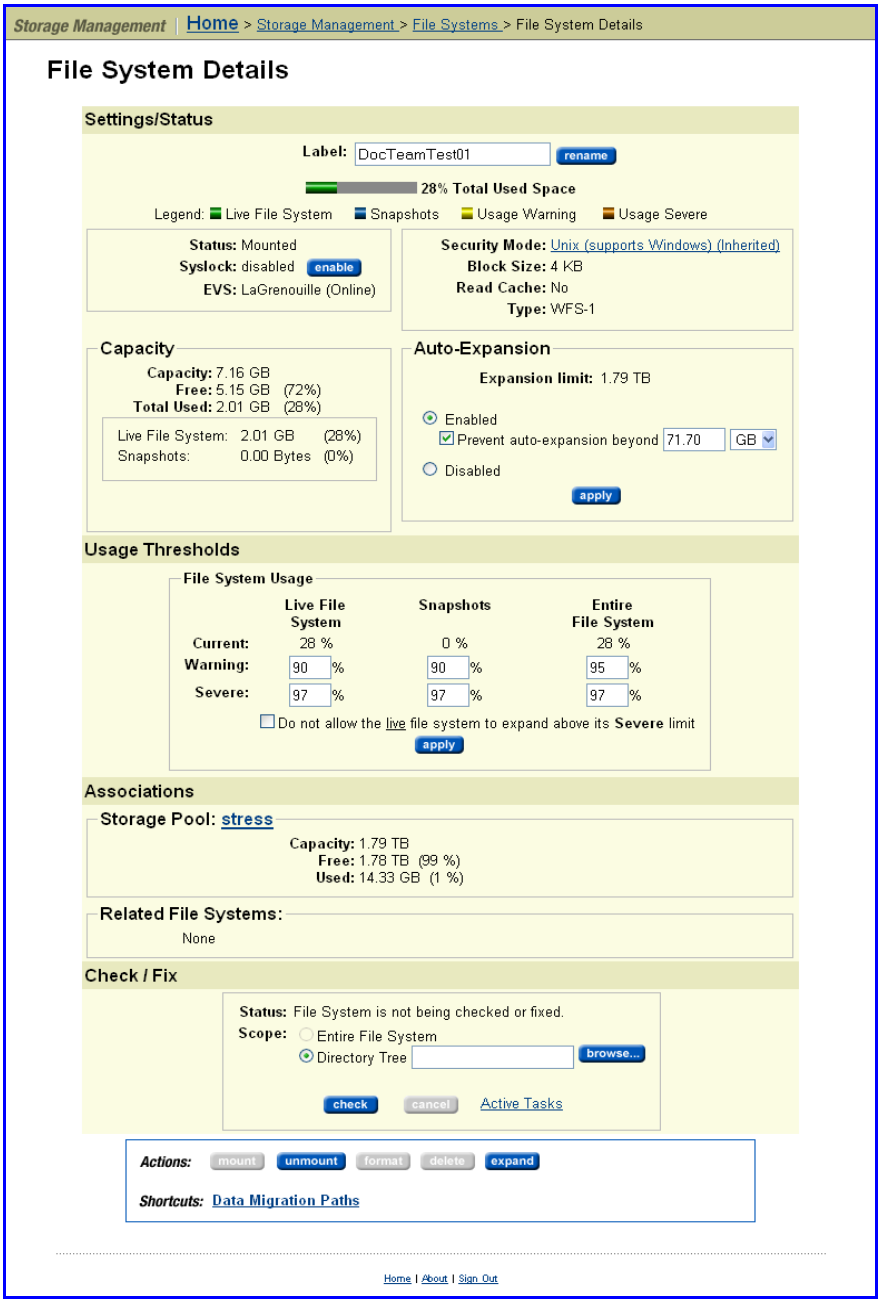

1. **Select an EVS.**

   From the **EVS** drop-down list, select the EVS to which you wan to assign the file system, then click **assign**.

2. **Mount the file system.**

   While the **Server Settings** page returns, in the file system list, verify that the new assignment is displayed in the **EVS** column, then select the file system and click **mount.**

   The file system status changes to *Mounted.*

Virtual Server (EVS) Management

The **EVS Management** page allows EVSs to be added, enabled, and disabled.

To display the **EVS Management** page, navigate to the **Server Settings** page, then click to display the **EVS Management** page:

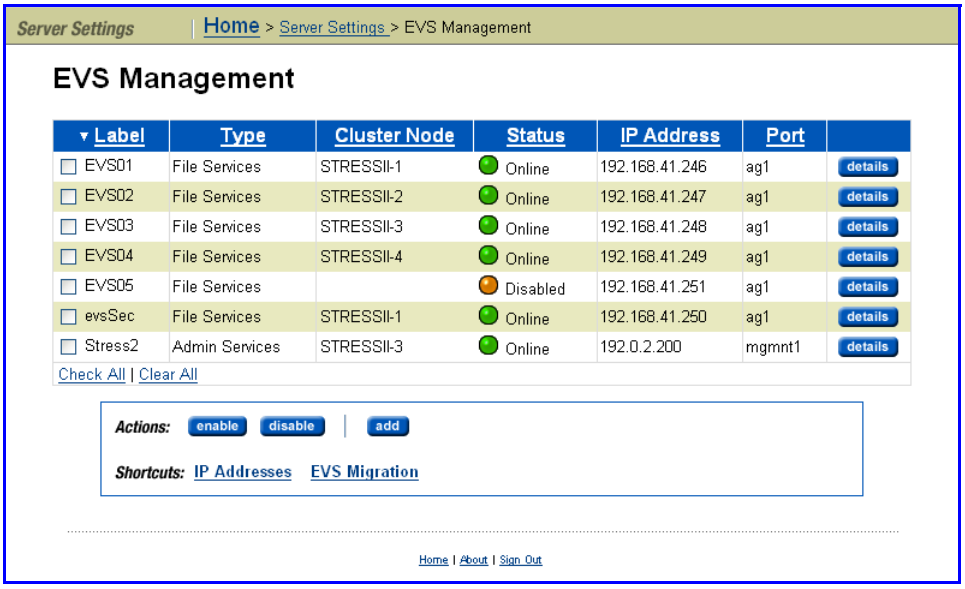

The following table describes the columns in this page:

Item	Description
Label	EVS identifier.
Type	Type of service: *administrative services* or *file services*.
Cluster Node	Assigned cluster node (only displayed for cluster nodes).
Status	Service status: • **Online:** Up and accessible. • **Disabled:** Down and inaccessible.
IP Address	First IP address assigned to an EVS. **Note:** An EVS can have multiple IP addresses.
Port	The cluster node Ethernet group to which the IP address for the EVS is assigned.
details button	Click **details** to display the **Details** page for the EVS.

You can perform actions on one or more EVSs at the same time. Fill the checkbox for each EVS you want to select, then click on one of the following **Actions**:

- Click **enable** to enable a disabled EVS.

- Click **disable** to disable an enabled EVS.
- Click **add** to create a new EVS (refer to Creating an EVS, on page 416 for more information).

Viewing Virtual Server (EVS) Details

To view the **EVS Details** page, navigate to the **EVS Management** page, then click **details** for a selected EVS to display its **EVS Details** page:

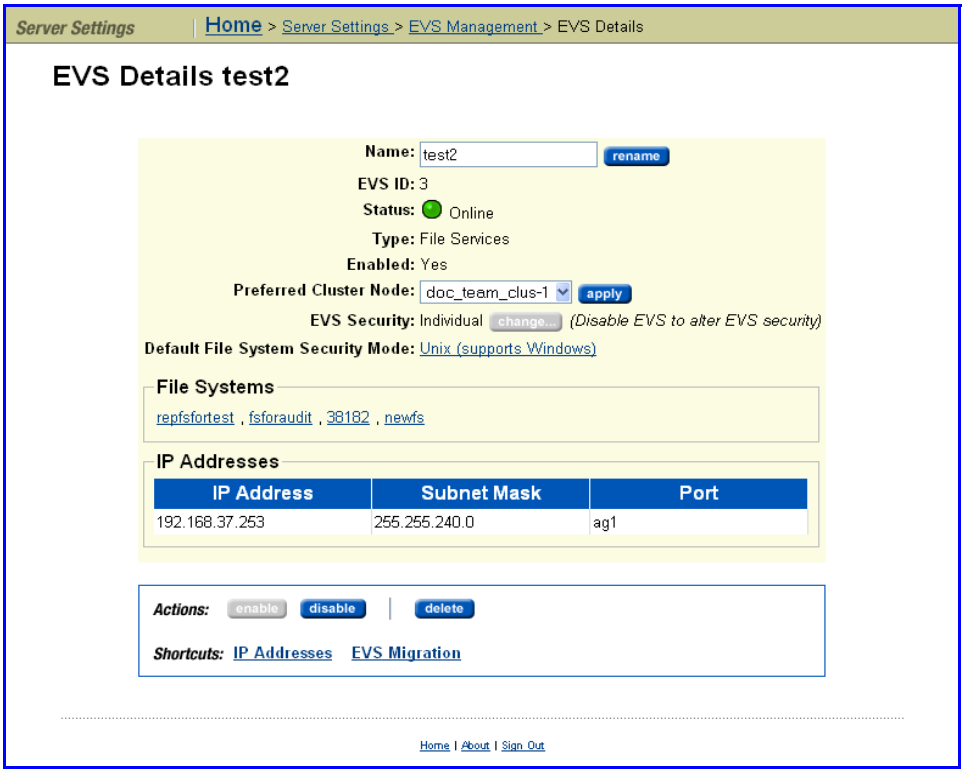

The following table describes the *additional* fields in the **EVS Details** page:

Item	Description
Name	EVS identifier (same as Label in the previous page).
EVS ID	Unique identifier for the EVS within the cluster, generated by the server upon EVS creation. **Note:** If moved to another server in a server farm, the EVS ID may change, but not if the move is within a cluster.
Status	Service status: • **Online**: Up and capable of providing services. • **Offline**: Not running. While offline, EVS are inaccessible.

Item	Description
Type	Type of service provided by the EVS: administration services or file services.
Enabled	*Yes* (enabled) or *No* (disabled).
Preferred Cluster Node	*Preferred* cluster node for the EVS (only displayed for file serving EVSs in a cluster.) Indicates cluster node preference; the EVS may, however, be hosted by a node other than its preferred node after having been migrated, for several reasons, such as node failure, manual migration for load balancing, etc.
EVS Security	Displays the current EVS security context. Click **Change** to select a different EVS security context or to select the global configuration. Selecting a different EVS security context changes how EVS name services, user mappings, group mappings, localgroups, name space (CNS), DNS servers, and NIS/LDAP configuration settings are managed. • If an EVS uses the **Global** configuration, any changes made to the global configuration settings will affect the EVS. • If an EVS uses an **Individual** security context, changes made to the global configuration settings will not affect the EVS. To change the settings of an EVS using an individual security context, you must go to the configuration settings for the EVS' individual security context to make changes, even if those settings are the same as the settings used by the global security context.  **Note:** You must disable the EVS before you can switch the EVS security context being used (between the global security context and an individual security context).
Default File System Security Model	Displays the default security model to be used by file systems in the EVS.
File Systems	List of all file systems hosted by the EVS. For more information about a particular file system, click its name to display its **File System Details** page.  **Note:** Administrative EVSs (EVS Type set to *Admin Services*), the this area is not displayed
IP Addresses	The IP addresses area displays a list of all IP addresses assigned to the EVS. Note that EVS can have multiple IP addresses.
Subnet Mask	Subnet mask for the EVS.
Port	The cluster node Gigabit Ethernet port to which the IP address for the EVS is assigned.

The following **Actions** are available:

- Click **rename** to apply a new label entered in the **Name** field.
- Click **apply** to apply a new preferred cluster node selected from the **Preferred Cluster Node** drop-down list.

- Click **enable/disable** to control status of this EVS.
- Click **delete** to remove the EVS. Do not click **delete** until you first **disable** the EVS.

Note: Deleting an EVS does not affect the file system owned by the EVS. Once the EVS has been deleted, assign the file system to another EVS to make it available for use.

Migrating Virtual Servers (EVSs) within a Cluster

While migration of EVSs occurs automatically as part of the failover resiliency of a cluster, EVSs can be manually migrated to a different node in a cluster, or among servers or clusters within a server farm.

Migrating an EVS within a Cluster

An individual EVS can be migrated to a different node within the same cluster, or *all EVSs* can be migrated to another server or another cluster. The current mapping of *EVSs to cluster nodes* can be preserved, and the saved map is called a preferred mapping.

To migrate EVSs between nodes of the same cluster:

1. **Navigate to the EVS Migrate page.**

   From the **Server Settings** page, click to display the **EVS Migration** page.

   Note: If the currently managed server is in a cluster and the SMU is also managing at least one standalone server, the following page appears (if the SMU is *not* managing a cluster and one or more standalone servers, this page will not appear):

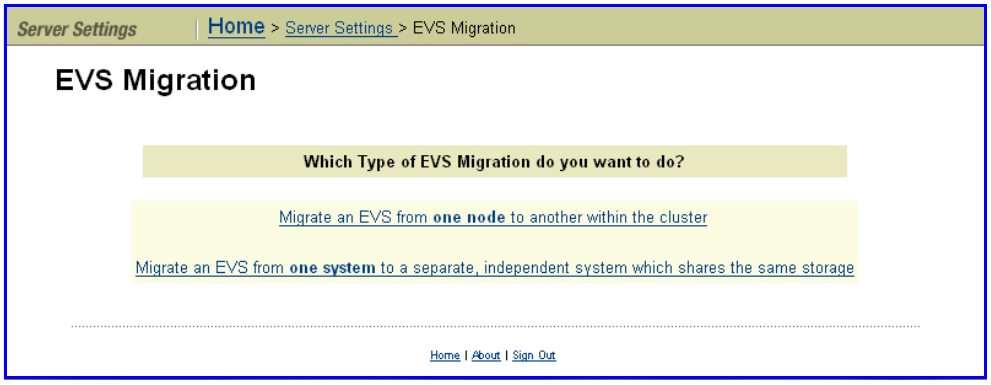

   In the above page, select “Migrate an EVS from one node to another within the cluster” to display the main **EVS Migration** page:

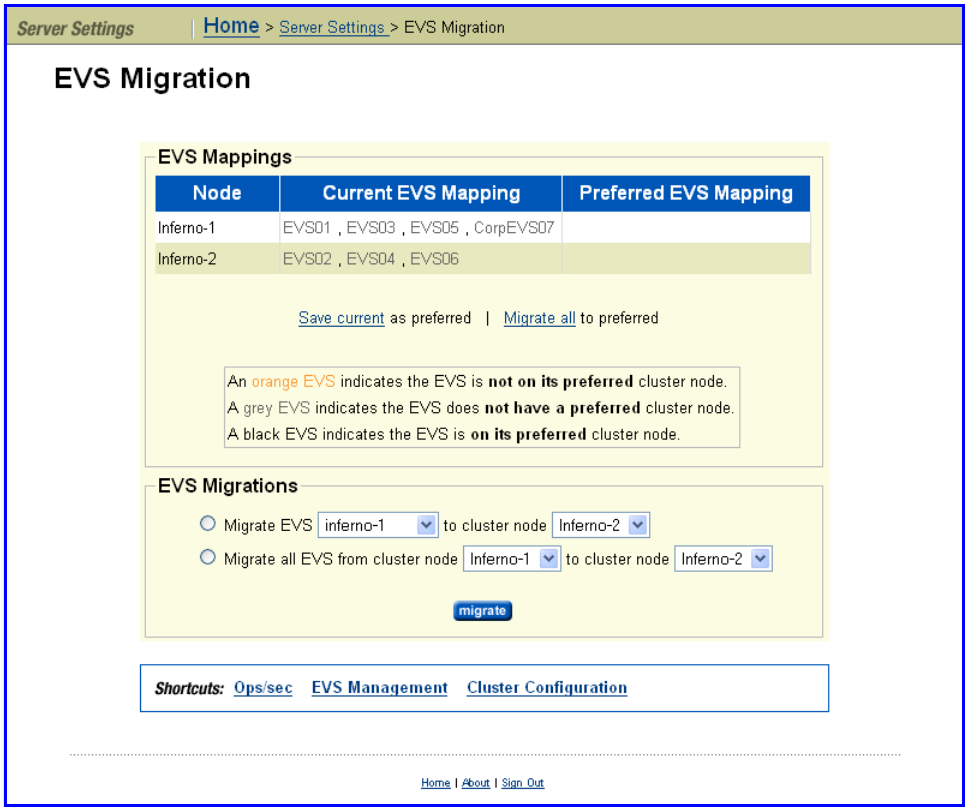

2. **Perform a migration of the type required:**

   a. **To migrate all EVSs between cluster nodes:**

   - Select "Migrate all EVS from cluster node ___ to cluster node ___".
   - From the first drop-down list, select the cluster node from which to migrate all EVS.
   - From the second drop-down list, select the cluster node to which the EVSs will be migrated.
   - Click **Migrate**.

   b. **To migrate an EVS to a cluster node:**

   - Select "Migrate EVS ____ to cluster node ___."
   - From the first drop-down list, select the cluster node to migrate.
   - From the second drop-down list, select the cluster node to which the EVS will be migrated.
   - Click **Migrate**.

   c. **To save a *Preferred EVS to cluster node* mapping.**

Note: Saving the current EVS-to-cluster configuration as the *Preferred Mapping* helps when restoring EVSs to cluster nodes. For example, if a failed cluster node is being restored, the preferred mapping can be used

to easily restore the original cluster configuration.

To perform this operation:

- Migrate the EVS between the cluster nodes until the preferred mapping has been defined. The current mapping will be displayed in the **Current EVS** list box.
- To save current *EVS-to-cluster node* mapping, click the **Save** in "Save current EVS mapping as the preferred mapping." The preferred mapping will then be displayed in the **Current EVS** column.

d. To migrate all EVSs to a *Preferred Mapping:*

To migrate all EVS to their preferred mapping, click **Migrate** in "Migrate all EVS to their preferred mapping."

Migrating an EVS within a Server Farm

Migration within a server farm is supported under the following conditions:

- When both the source and destination server are online.
- If the source server is offline and the destination server is online.
- The EVS does not contain any file systems that are linked into a CNS tree.

Note: After migrating EVS between servers in a Server Farm, the assignment of tape drives and tape autochanger devices to EVS must be manually adjusted:

- Tape devices specifically assigned to a migrated EVS will have become unassigned.
- Tape devices assigned to "any EVS" on the source server will remain assigned to "any EVS" on the source server.

Tape devices must not be assigned to EVSs on more than one server.

While EVSs contain most of the settings required to support client storage access, some settings (including *DNS, Windows NT domain* and *Active Directory*) are functions of the host server or node, not of the EVS itself. Therefore, when preparing to migrate an EVS from one server or cluster to another, verify in advance that the target server's settings can properly support the EVS. To prepare a server or cluster to receive a new EVS, source server settings can be cloned. For information about specific settings that can be cloned, see Cloning Server Settings, on page 424.

Cloning Server Settings

To clone server settings:

1. **Navigate to the Clone Server Settings page.**

   From the **Server Settings** page, select the target of the migration as the SMU's managed server, then click to display **Clone Server Settings** page:

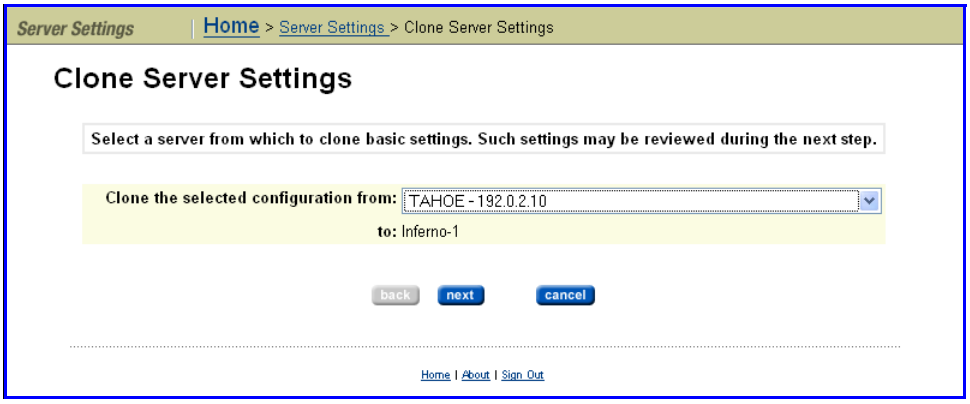

2. **Select the source.**

   From the drop-down menu, select the server/node currently hosting the EVS, then click **next**.

3. **Specify settings to be cloned.**

   Select the settings to clone to the target server/node.

   Caution: Settings selected for cloning will overwrite the currently defined settings on the target server/node. To keep part of the existing configuration, do not clone the configuration item(s) you want to keep.

4. **Start the process.**

   Click **OK** to initiate cloning.

   Once the target server/node has been prepared through Server Cloning, it is ready to receive the EVS migration.

Migrating an EVS Within a Server Farm

To migrate an EVS within a server farm:

1. **Navigate to the EVS Migration page.**

   From the **Server Settings** page, click to display the **EVS Migration** page.

   Note: This page will only appear if the currently managed server is a cluster node. Otherwise, clicking **EVS Migrate** will immediately launch the

EVS Migration page shown in the next step.

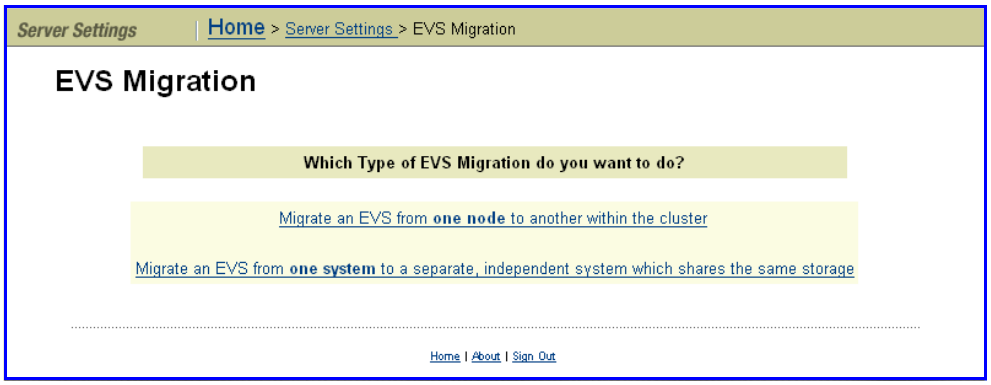

In the above page, click "Migrating an EVS from one system to a separate, independent system" to display the main **EVS Migration** page:

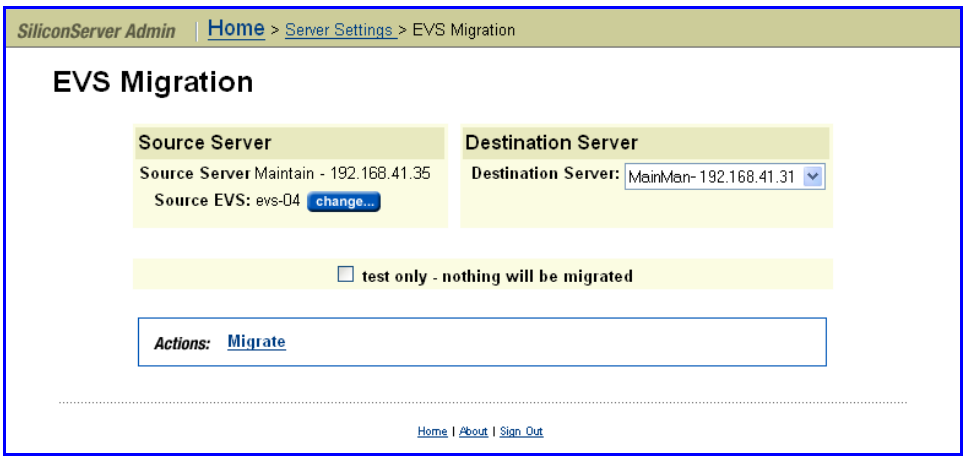

2. **Select Source and Destination servers:**
   - *To select a Source Server and Source EVS,* click **change**.
   - *To select a Target Server,* use the **Destination Server** drop-down menu.
   - *To test the migration before committing the change,* fill the **Test Only** box. This ensures that the EVS migration is possible.

     A message appears telling you that the operation succeeded or failed. If the operation failed, the message includes the reason for the failure.

3. **Start the process.**

   Click **Migrate.**

Note: If the source server is offline or doesn't function, migration will be performed using an existing backup and the following warning will

appear:

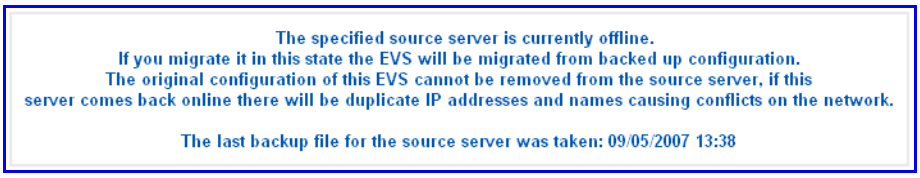

Using Clusters

IS-NAS Servers or Titan Servers can form clusters under the following conditions:

- The cluster to which a node is being added must have a license for at least the currently existing number of nodes.

  If the cluster to which a node is being added has a license for a number of nodes that is equal to the currently existing number of nodes, the joining node must also have a cluster license.

  If the cluster to which a node is being added has a license for a number of nodes that is larger than the currently existing number of nodes, the joining node does not need to have a cluster license.
- All nodes in the cluster must have the same hardware configuration. (You cannot form a cluster from a combination of IS-NAS Servers and Titan Servers.)
- The node joining the cluster must be of a compatible software level (within one minor revision level). For example, a server running version 6.0 software can be added to a cluster running version 6.1 software, but not to a cluster running version 6.2 software.

Once the first server has been set to cluster mode, you can:

- Add nodes by "joining" servers to the cluster.
- Add EVSs to the cluster and distribute them among the cluster nodes.

Note: In order to maximize cluster performance, distribute EVSs across nodes to level the network client load between them.

About Cluster Licensing

The maximum number of nodes for a cluster is controlled by several factors, including hardware version, software version, and cluster license(s).

Note: The maximum licensed number of nodes in a cluster will never exceed the maximum number of nodes supported by the hardware and software of the nodes making up the cluster.

A cluster license can be for a single node or for multiple nodes.

- A single node license allows the server/node on which the license is installed to become the first node in a cluster or to join an existing cluster. Using single node cluster licenses, you can form clusters of up to the

maximum number of nodes supported by the hardware/software being used.

Single node cluster licenses can also be used to increase the maximum number of nodes in an already-formed cluster, up to the supported maximum.

- A multi-node license allows the cluster on which the license is installed to form a cluster containing up to the licensed number of nodes, or the supported maximum number of nodes, whichever is lower.

  If a server/node containing a multi-node cluster license joins an existing cluster, the cluster's total licensed number of nodes increases to the higher of the following:

  - The maximum number of nodes licensed by the existing cluster.
  - The maximum number of nodes in the existing cluster's license plus one.

    This happens when the total size of the cluster is already greater than or equal to the licensed maximum number of nodes in the existing cluster.

Note: The only difference between a single-node and a multi-node cluster license is the maximum number of nodes the license permits. After installing the license key, you can see the difference between the number of nodes allowed by the license on the **License Keys** page (see Managing License Keys, on page 535 for more information about managing license keys).

Maximum cluster size can be determined in either of the following ways:

- A cluster containing a multi-node cluster license, for up to "*X*" nodes.

  This method is typically used for new larger-scale installations, where a multi-node cluster is being set up as a new installation and the node containing the multi-node license becomes the first cluster node.

- An additive process, that combines an existing cluster and a node containing a single-node cluster license.

  This method is typically used for installations that are expected to grow over time. The key advantage provided by this additive method is that maximum cluster size need not be determined in advance.

  For example, you can start with a single server without a cluster license. Later, you install a cluster license, configure the server as the first node of the cluster, and then add nodes. In this situation, you could begin with:

  - A multi-node cluster license and then add nodes that don't have cluster licenses into the cluster.

- A single-node cluster license and then install additional nodes (each having their own single-node cluster license) into the cluster.

Another situation where this additive process is used would be if you start with a small cluster, and later add nodes to make a larger cluster. For example, if you start with a two-node cluster that has a four-node license, you can later add two servers (that don't have cluster licenses) to create a four-node cluster. If necessary, you could later grow the cluster by adding individual nodes (each having a single-node cluster license), up to the supported maximum number of nodes.

Assuming that the cluster has fewer nodes than the maximum size supported by the hardware and software, the rules governing the addition of a node to an existing cluster are fairly simple:

- A node may be added if the licensed maximum number of nodes is greater than or equal to the number of existing nodes, plus one.
- A node may be added if the licensed maximum number of nodes is equal to the number of existing nodes, **and** the joining node has a cluster license.

  When joining an existing cluster, if the joining node has a cluster license, that cluster license is transferred to the existing cluster, and the cluster's maximum number of nodes increases by one (1). **The cluster's maximum number of nodes is increased by one, regardless of the maximum number of nodes allowed by the cluster license of the joining node, even if the joining node has a multi-node cluster license**. For this reason, the order of joining nodes into a cluster is important.

When becoming a cluster node, all of its licenses are transferred to the cluster, and different licenses are transferred in different fashions (see License Types, on page 538 for more information).

Configuring New Clusters

Using the Cluster Wizard, you can:

- Create a new cluster by configuring a server as the first cluster node, or
- Join a server to an existing cluster as a new node.

  See About Cluster Licensing, on page 427 for information about adding nodes to an existing cluster.

Configuring the First Cluster Node

If any of the nodes that you are going to use to form the cluster contain a multi-node cluster license, that node is the one that should be configured as the first cluster node. See About Cluster Licensing, on page 427 for more information about adding cluster nodes.

To configure the first cluster node:

1. **Navigate to the Cluster Wizard page.**

From the **Server Settings** page, click **Cluster Wizard** to display the **Cluster Wizard** page:

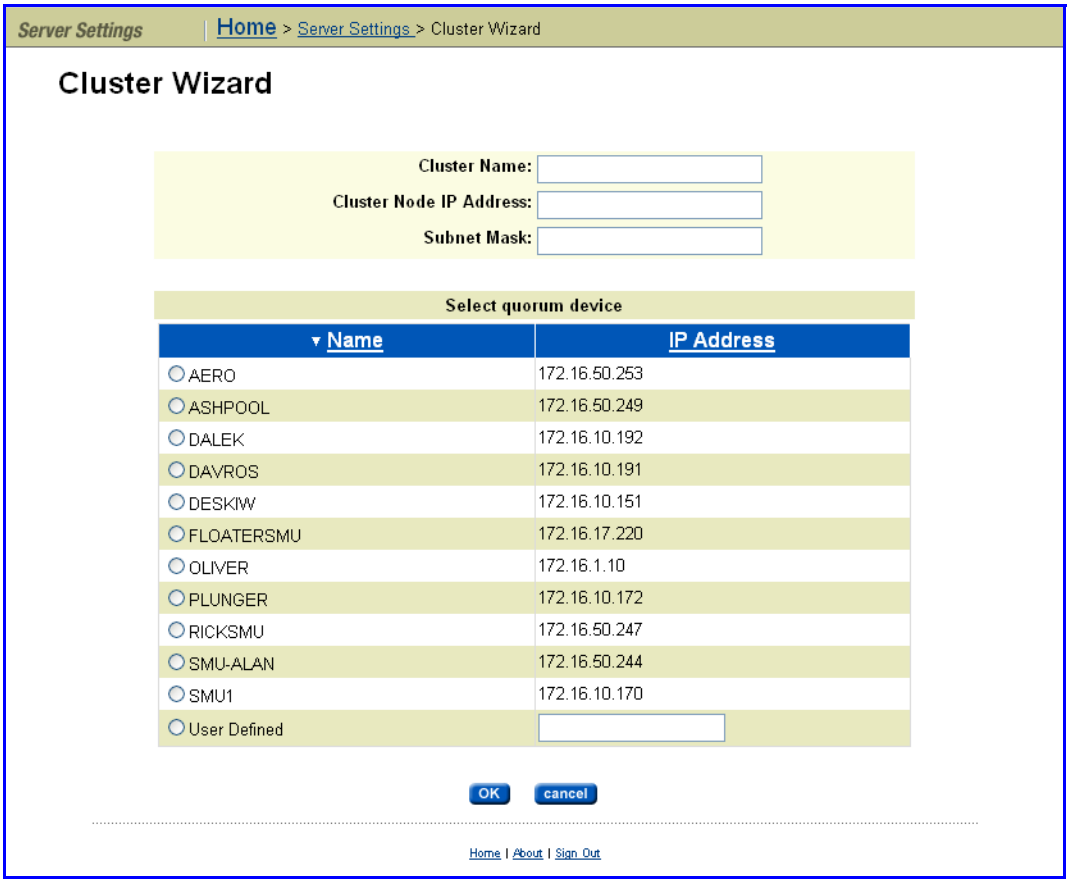

2. **Enter the requested information.**

   Assign a new **Cluster Name**, and associated **Cluster Node IP Address** and **Cluster Subnet Mask**, then select a Quorum Device.

   Note: Whether creating a new cluster or joining a cluster node, a *Cluster Node IP Address* must be defined. This IP address maintains heartbeat communication among cluster nodes and between the cluster nodes and the Quorum Device (QD), which is typically the SMU. Due to the importance of the heartbeat communication, the cluster node IP address should be assigned to the 10/100 management port connected to the private management network, keeping the heartbeats isolated from normal network congestion.

3. **Save the configuration.**

   Click **OK.** The server reboots automatically. On restart, the node joins the cluster.

Joining an Existing Cluster Using the CLI

See About Cluster Licensing, on page 427 for information about adding nodes to an existing cluster.

To join an existing cluster through the CLI:

1. **Connect a terminal to the server.**

   Connect to the built-in RS232 port of the unconfigured server that will join the cluster, as described in Using the Command Line Interface, on page 16.

   Caution: *IP Address and file system bindings overwrite Alert!* When joining a node to a cluster, the cluster's configuration will overwrite the configuration information of the joining node, causing any prior configuration changes to be lost, including IP addresses and file system bindings.

2. **Enter the requested information.**

   When the server boots for the first time, it prompts for cluster membership and requests information about the host cluster:

Prompt	Description
Is this node joining a cluster	To join the cluster, enter y.
Enter cluster node IP address	Enter the cluster node IP address to assign to the joining node. This IP address will be assigned to the 10/100 Ethernet management port (`eth0` or `eth1` for a IS-NAS Server or `mgmnt1` for a Titan Server).
Enter cluster node IP mask	Enter the network mask for the joining node's cluster node IP.
Enter the cluster name	Enter the name of the cluster to join. This is the cluster name and should not be confused with the DNS name (that is, *serverName* should be entered, not 192.0.2.2.)

3. **Allow the system to reboot.**

   The node will automatically join the cluster.

Joining an Existing Cluster Using Web Manager

See About Cluster Licensing, on page 427 for information about adding nodes to an existing cluster.

To join an existing cluster using Web Manager:

1. **Navigate to the Join Cluster Wizard page.**

   From the **Server Settings** page, click to display the **Join Cluster Wizard** page.

2. **Select a server and configure basic settings**

Using the radio buttons, select a server, then check the suggested IP Address for the node (you can change it if necessary), enter a *username* and *password,* then click **next**.

3. **Allow the system to reboot.**

   The selected server will automatically reboot, and join the cluster during the boot process.

Managing a Cluster

The following sections explain how to manage the cluster including cluster services (file services and server administration) and the physical elements which form the cluster (cluster nodes and the quorum device).

Configuring the Cluster

To configure the cluster:

1. **Navigate to the Cluster Configuration page.**

   On the **Server Settings** page, click to display the **Cluster Configuration** page:

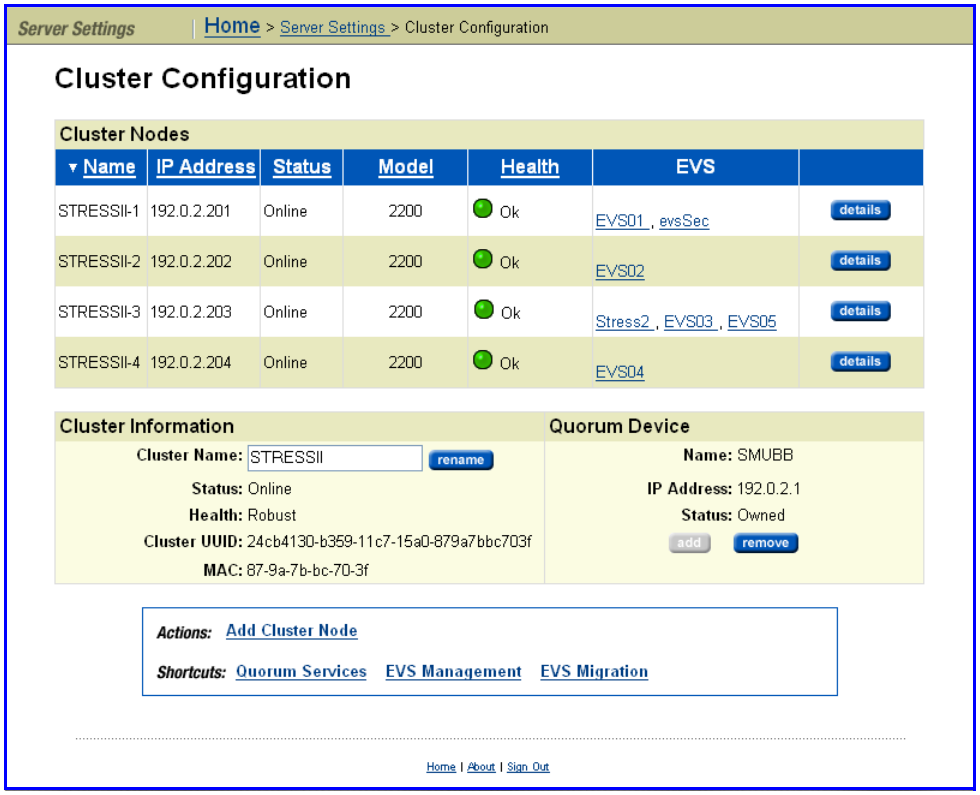

The following table describes the fields in this page:

Item	Description
Cluster Nodes	

Item	Description
Name	Node name.
IP Address	IP address of the cluster node. This IP address is on the private management network, which connects devices within the cluster.
Status	Node status: • **Online.** The cluster node is a part of the cluster and is exchanging heartbeats with the other cluster members. • **Offline.** The cluster node is no longer exchanging heartbeats with the other cluster members. This may be caused by a reboot or a fault condition. Services cannot be migrated to a cluster node in this state. • **Dead.** No heartbeats from this cluster node for a significant period. Services cannot be migrated to a cluster node in this state. • **Unknown.** The node has not been online since the cluster was started.
Model	Server model, if available.
Health	Worst-case status from each node: • **OK.** Operating normally, with no failures. • **Degraded.** Operating, but with one or more failures in connectivity. The problem may be with the cluster interconnect, the management network, or quorum device communication. • **Failed.** Not operating, due to one or more failures in connectivity. The problem may be with the cluster interconnect, the management network, or the quorum device communication.
EVS	EVS currently assigned to the cluster node.
Cluster Information	
Cluster Name	Cluster name. To rename the cluster, enter the new cluster name in the edit box, then click on the **rename** button.
Status	Unless rebooting, the status will be "Online."
Health	Cluster health: • **Robust.** Operating normally, with no failures in the cluster interconnect, the management network, or quorum device communication. • **Degraded.** Operating, but one or more nodes has failed or there has been a failure in connectivity.
Cluster UUID	UUID (unique ID) of the cluster. This string provides a unique identifier for each cluster when there are several clusters on a network.
MAC	MAC address of the cluster.
Quorum Device	

Item	Description
Name	SMU on which the QD resides.
IP Address	IP address of the SMU on which the QD resides.
Status	QD status: • **Configured.** Attached to the cluster, but vote not needed. The QDs vote is not needed when any cluster contains an odd number of operational nodes. • **Owned.** Attached to the cluster and owned by a specific cluster node. • **Not up.** Cannot be contacted. • **Seized.** Taken over by another cluster.

2. **As needed, modify quorum device assignment:**
   - *If a QD has not been specified*, click **add** to assign a QD to the cluster.
   - *To remove the specified QD*, click **remove.**

     If a QD is removed from the cluster, the service will be released back to SMU's pool of available QDs.
3. **As needed, modify cluster node assignment:**

   Note: Services hosted by the cluster node must be migrated to a different cluster node before a node can be removed.

   - *To remove a cluster node*, click its **details** button to display the corresponding **Cluster Node** page. Click **Remove From Cluster**, then **OK** (or **Cancel** to decline) in the confirmation dialog. For additional information about the **Cluster Node Details** page, see Viewing Cluster Node Details, on page 434.

     Upon node removal, any hosted EVSs will automatically be migrated to another cluster node, with details provided in the confirmation dialog.
   - *To add a node to the cluster*, from the **Server Settings** page, select **Cluster Configuration**, then select **Cluster Join Wizard** to display the **Cluster Wizard**. See Configuring the First Cluster Node, on page 429 for information about using the Cluster Wizard.

Viewing Cluster Node Details

The **Cluster Node Details** page displays information about a selected cluster node and allows removal of that node from the cluster.

To access the **Cluster Node Details** page, navigate to the **Server Settings** page, then select **Cluster Configuration**; in this page, select a cluster node, then click **details** to display its **Cluster Node Details** page:

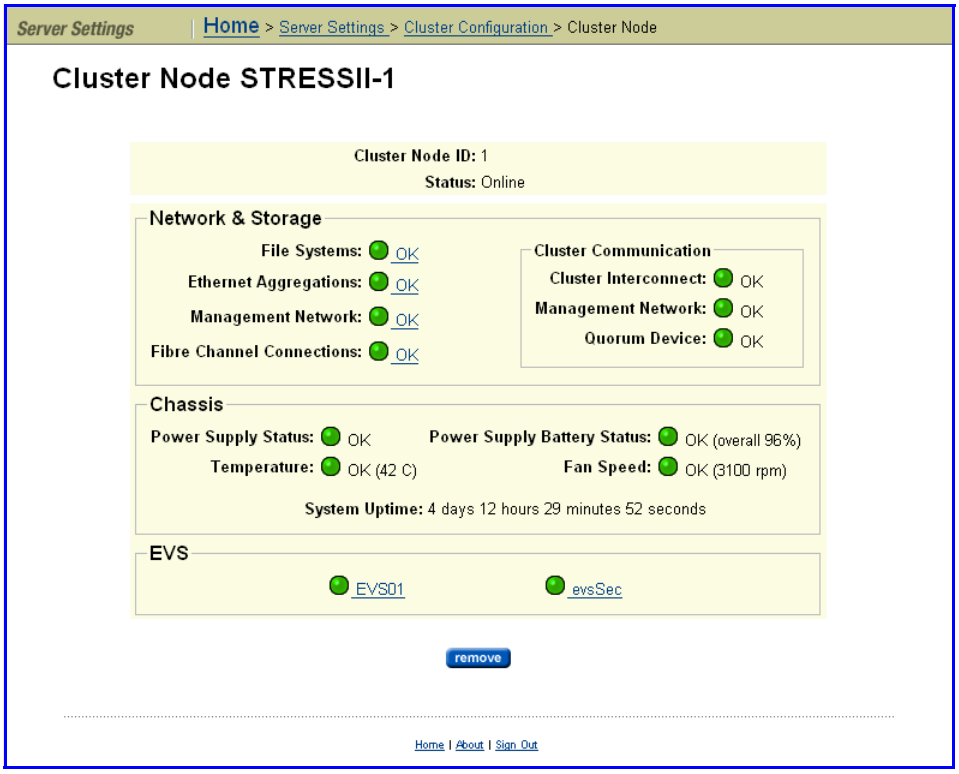

The following table describes the contents of this page:

Item	Description
Cluster Node ID	The ID assigned to the node.
Status	An indicator of node status. • **Online.** Node has completed booting. • **Unknown.** Node has not yet booted. • **Up.** Node is booting (displayed only while the node is booting). • **Not up.** Node is shutting down (displayed only while the node is shutting down). • **Dead.** Node has failed to go online after booting.
Network and Storage	
File Systems	Overall indicator of file system status: • **OK.** All file systems up and operational. • **Failed.** One or more file systems has failed. Click on the status link to display the **File Systems** page, which lists all file systems assigned to the EVS in that cluster node.

Item	Description
Ethernet Aggregations	Overall status of Ethernet aggregations in the cluster node: • **OK.** All aggregated ports are up and linked. • **Degraded.** One or more ports in an aggregation has failed. • **Failed.** All ports in an aggregation have failed. Click on the status link to display the **Link Aggregation** page, which lists all aggregations in the cluster node.
Management Network	Overall status of the management network: • **OK.** Links are up and heartbeats are being received. • **Failed.** No heartbeats are being received, and the links may be up or down. Click on the status link to display the **Ethernet Statistics** page, which lists information about the management port and the aggregated Ethernet ports in the cluster node.
Fibre Channel Connections	An overall status indicator for the Fibre Channel ports in the cluster node: • **OK.** All ports up and operational. • **Degraded.** Some ports up and operational, but one or more has failed. • **Failed.** All ports have failed. Click on the status link to display the **Fibre Channel Statistics Per Port** page, which lists all Fibre Channel ports in use in the cluster node.
Cluster Communication	This section contains status indicators for communications within the cluster node. Cluster Interconnect status: • **OK.** Link is up and heartbeats are being received. • **Standby port down.** The primary link is up and heartbeats are being received, but the secondary link is down. • **Link up, no heartbeating.** At least one link is up, but no heartbeats are being received. • **Link down.** All links are down (and therefore no heartbeats are being received). Management Network status: • **OK.** Both links are up and heartbeats are being received. • **Link up, no heartbeating.** Both links are up, but no heartbeats are being received. • **Link down.** Both links are down (and therefore no heartbeats are being received). Quorum Device status: • **OK.** The Quorum Device is communicating with the cluster node. • **Link up, no quorum communication.** The link to the Quorum Device is up, but the Quorum Device is not communicating with the cluster node. • **Link down.** There is no communication with the Quorum Device.
Chassis	
Power Supply Status	A status indicator for the cluster power supply units (PSUs): • **OK.** Both PSUs are installed and operating normally. • **Degraded.** One PSU not responding to queries, which may mean that it is switched off. • **Failed.** One PSU not responding to queries, and it has failed or is not present.

Item	Description
Temperature	Status indicator for temperature of the cluster node chassis. • **OK.** Within the normal operating range. • **Degraded.** Above normal, but not yet critical. • **Failed.** Critical. When available, the temperature in the chassis is also displayed. The displayed temperature is the highest reported temperature of any of the boards in the chassis.
Power Supply Battery Status	Status of the cluster power supply battery: • **OK.** Capacity and voltage within the normal operating range. • **Degraded.** Capacity and/or voltage below normal. This status should be considered a *warning*; if it continues, the PSU battery should be replaced. • **Failed.** Capacity and/or voltage below acceptable minimum, or the PSU battery is not responding to queries. This status indicates a failure; the PSU battery should be replaced. When available, the level of battery charge is also displayed.
Fan Speed	Status of fans in the cluster chassis: • **OK.** All fans operating normally. • **Degraded.** One or more fans spinning below normal range. • **Failed.** At least one fan has stopped completely, or is not reporting status. When available, the chassis fan speed is also displayed. The displayed fan speed is the slowest reported speed of any of the three fans. An error message may be displayed, even if it does not correspond with the slowest fan.
System Uptime	Duration since last reboot of the cluster node.
EVS	This section displays the names (labels) of EVSs assigned to the node, indicates status for each: • **Green.** Online and operational. • **Amber.** Offline, but listed here because it is hosting the administrative EVS. • **Red.** Failed. Click on the EVS name to display the **EVS Details** page for that EVS.

Quorum Device Management

The SMU provides Quorum services, assigning individual quorum devices (QDs) to clusters during configuration. The SMU can provide quorum services for up to eight clusters, by hosting a pool of eight available QDs, from which it assigns a QD to a cluster upon configuration. Once assigned, the QD is "owned" by its cluster and is no longer available for assignment to another cluster. Removing a QD from a cluster releases its ownership and returns it to the SMU's pool of available QDs.

Viewing and Managing Quorum Devices

Each instance of the QD service can be viewed and managed in the **Quorum Services** page.

To view and manage quorum devices:

1. **Navigate to the Quorum Services page.**

   From the **SMU Administration** page, click to display the **Quorum Services** page:

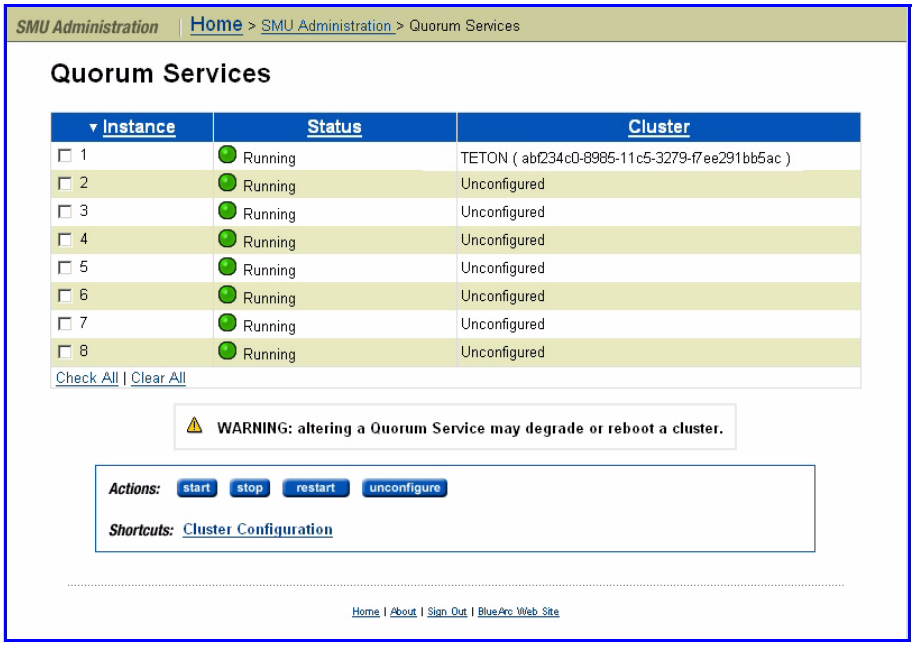

   The following table describes the columns in this page:

Item	Description
Instance	List of available quorum device instances hosted by the SMU.
Status	Service status: • **Running**. Up and capable of providing services. • **Not Running**. Not running. While offline, the QD service is not available to the cluster.
Cluster	Name and cluster UUID of the node.

2. **Start, stop, restart, or unconfigure a selected quorum device.**

   Caution: incorrectly stopping, restarting, or unconfiguring a QD may disrupt QD services provided to clusters.

Using Cluster Name Space (CNS)

The CNS has a tree-like directory structure, much like a real file system. Its virtual root and subdirectories provide access to file systems. The CNS can be

viewed through the CLI or the Web Manager, and shows all of the configured directories and file system links.

Viewing the Cluster Name Space Tree

From **File Services,** click **CNS** to view the **CNS** page:

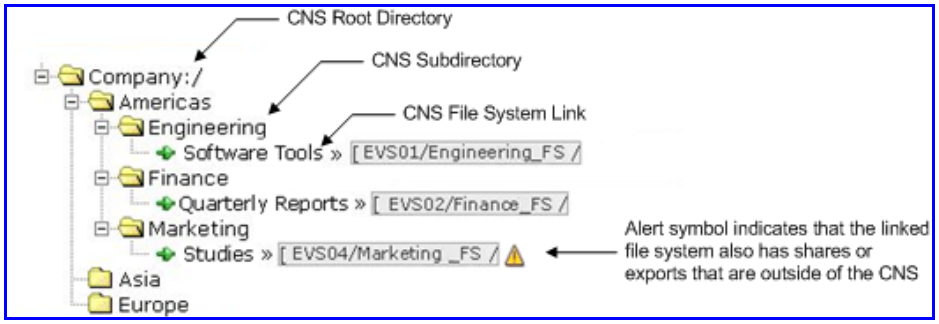

If a Secure EVS has been created and you want to view the EVS Name Space for that Secure EVS, click **change** to select the name space you want to view. Once you have selected a name space, the tree for that context is displayed.

- At the top of the name space is the root directory.
- Under the root directory are a number of subdirectories. In this example topology, one subdirectory has been created for each physical file system.
- Under each subdirectory is a file system link. A file system link associates a directory with a specific file system. The EVS to which the file system is associated is also shown.

Viewing the EVS Name Space Tree

From **File Services,** click **CNS** to display the **CNS** page. The currently selected EVS Security Context and the current name space are displayed at the top of

the page, and the tree for the current name space is displayed below the name space label.

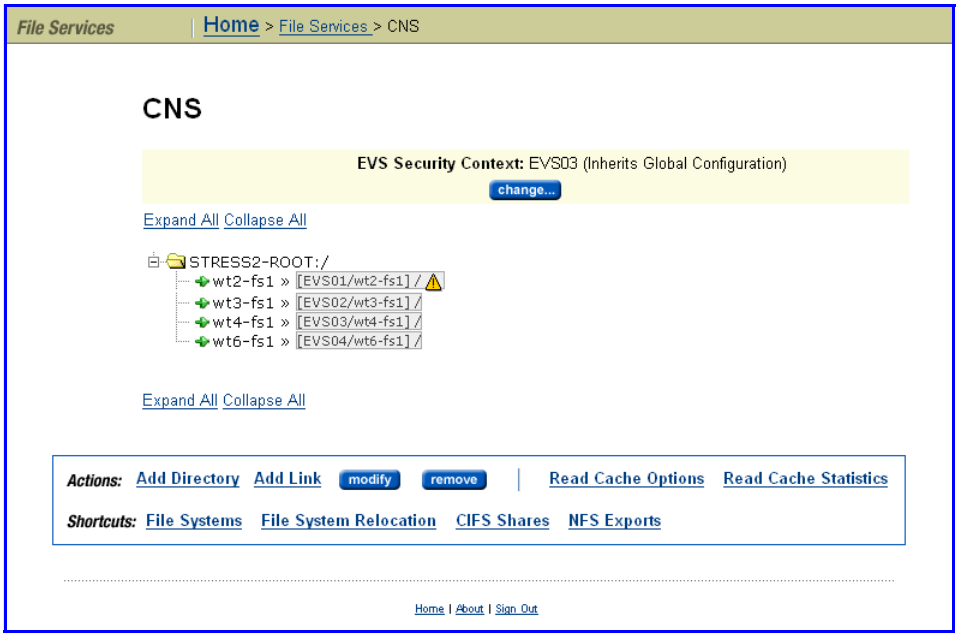

Click **change** to display a list of name spaces (the Global Configuration, and all individual EVS name spaces that have been defined). Click **Global Configuration** or the EVS name space to display the tree for that name space.

- At the top of the name space is the root directory.
- Under the root directory are a number of subdirectories. In this example topology, one subdirectory has been created for each physical file system.
- Under each subdirectory is a file system link. A file system link associates a directory with a specific file system. The EVS to which the file system is associated is also displayed.

Managing Links and Subdirectories in the EVS Name Space

Links and subdirectories in an individual EVS name space are managed in the same way as they are in the CNS. See Using Cluster Name Space (CNS), on page 438 for information on managing links and subdirectories in the CNS.

Creating a Cluster Name Space Tree

A CNS contains a root directory, file system links, and, optionally, subdirectories. The first step required to configure CNS is to create the root directory.

Creating a CNS Root Directory

To create a CNS root directory:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page:

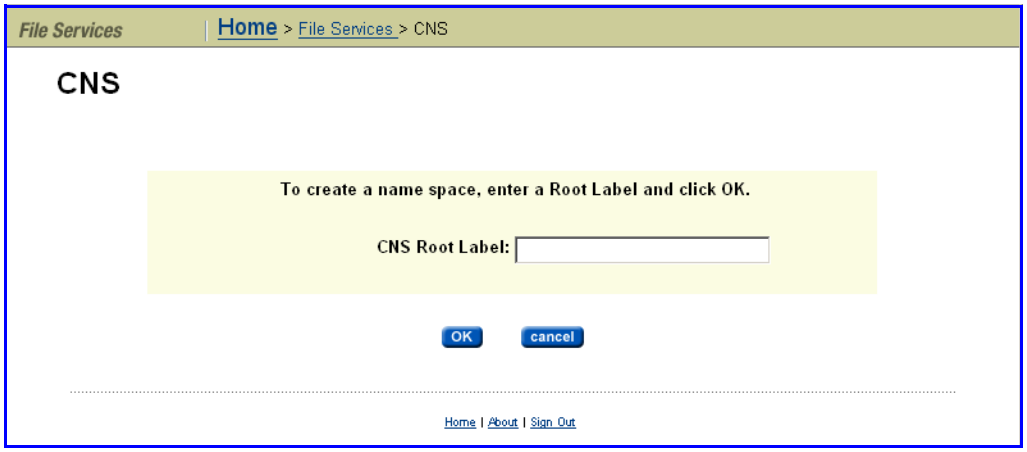

2. **Name the CNS root directory.**

   In the **CNS Root Label** text box, enter a name, then click **OK** to create the CNS.

   Note: For the CNS to be available to clients, a CIFS share and/or an NFS export must be created for it. See Configuring CIFS Shares, on page 268, or Adding an NFS Export, on page 246.

Creating CNS Subdirectories

Subdirectories can be created under the root directory or under other subdirectories in the CNS tree. They are optional, but give structure to the CNS, allowing granular control over the organization of physical file system resources.

To create a CNS subdirectory:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Add subdirectories.**

a. **Navigate to the add CNS Directory page.**

Click **Add Directory** to display the **Add CNS Directory** page:

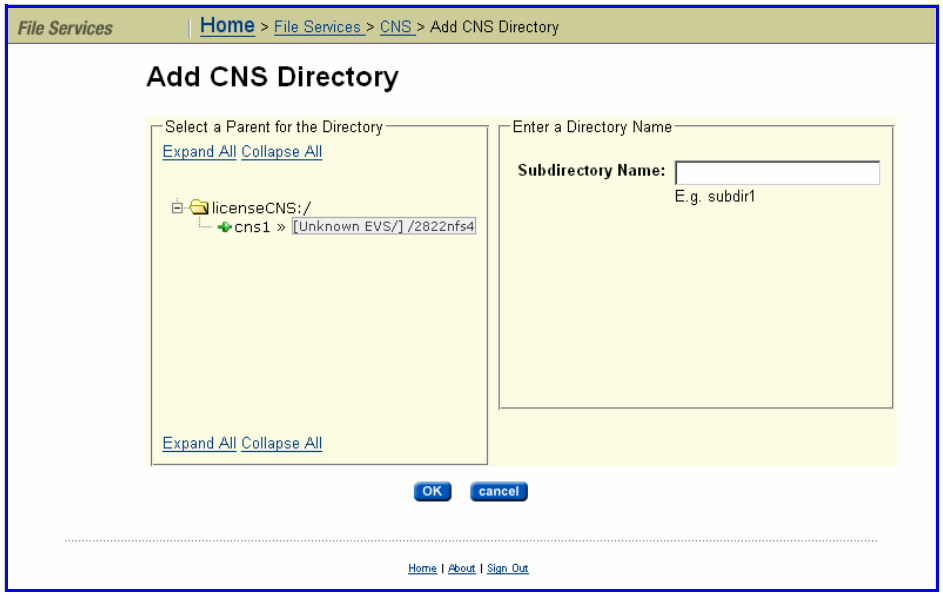

b. **Create subdirectories.**

From the **Select a Parent for the Directory** options box, select a parent directory, then enter a name in the **Subdirectory Name** text box. Click **OK** to create the subdirectory, then repeat to add any additional subdirectories.

Creating a File System Link

File system links make physical file systems accessible through the CNS. A file system link can be associated with either the root directory or a subdirectory in a physical file system. Once created, a file system link will appear as a directory in the CNS. The directory name seen by a network client will be the name given to the file system link. A network client navigating through CNS and into a file system link will see the contents of the directory that was linked.

To create a file system link:

1. **Navigate to the Link File Systems page.**

From the **File Services** page, select **CNS**, then click **Add Link** to display the **Link File System** page:

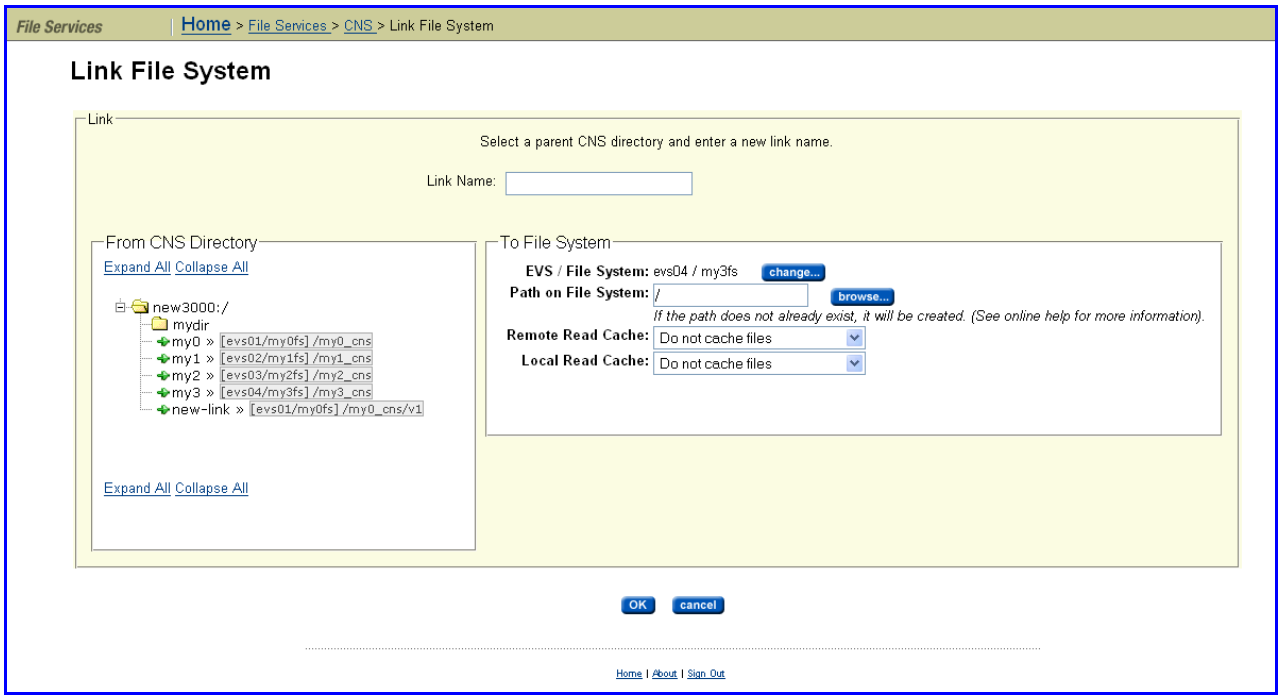

2. **Name and place the link.**
   - In the **Link Name** text box, enter a name for the link.
   - In the **From CNS Directory** options box, select a location in the CNS tree to place the link.
3. **Select a file system target.**

   From the **To File System** options, select **Change**. Then, select a target file system.

   To link a specific directory in the physical file system, rather than the root directory (which will link the entire file system), enter a directory name in the **Path on File System** text box, or click **browse** to search and select.
4. **Enable or disable read caching for the linked file system.**

   Note: For information about *Read Caching*, see Read Caching, on page 412.

   a. **Specify if files and/or links from remote file systems may be read cached.**

   Note: When the link being added is for a file system in an EVS that has an EVS individual namespace, remote read caching is not available.

   To allow files or cross file system links from remote file systems to be read cached, go to the **Remote Read Cache** drop-down list and select one of the following:

- **Cache all files**. Allows caching of files from a file system hosted by an EVS on a remote cluster node, **and** files accessed by local links to a remote file system (cross file system links). A remote cluster node is a node other than the one to which the client is connected.
- **Cache cross file system links**. Allows only cross file system links to be cached.

To disallow read caching of files from remote file systems, do not change the default selection of **Do not cache files**.

b. Specify if files and/or links from local file systems may be read cached.

To allow files or cross file system links from local file systems to be read cached, go to the **Local Read Cache** drop-down list and select one of the following:

- **Cache all files**. Allows caching of files from file systems on the same server/node as the read cache, **and** files accessed by local links to a remote file system (cross file system links). The remote file system may be a remote server or storage device.
- **Cache cross file system links**. Allows only cross file system links to be cached.

To disallow read caching of files from remote file systems, do not change the default selection of **Do not cache files**.

Changing Cluster Name Space Properties

After a CNS has been created, any of its properties can be changed, except the name of the root directory.

Deleting a Cluster Name Space

Deleting a CNS will permanently erase it. Deleting a CNS will not affect the physical file systems accessible through the CNS. However, once the CNS has been deleted, it may be necessary to restore access to the file system by sharing or exporting the file system through its EVS.

To delete a CNS:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Select the CNS to remove.**

   From the CNS directory tree, select the CNS root directory, then click **remove** to open a confirmation dialog.

3. **Apply your changes.**

   Click **OK** to delete the CNS.

Renaming a CNS Subdirectory

To rename a CNS subdirectory:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Select the subdirectory to rename.**

   From the CNS directory tree, select the subdirectory to be renamed, then click **modify** to display the **Modify CNS Directory** page.

3. **Rename the subdirectory.**

   In the **Subdirectory Name** text box, enter a new name for the CNS directory, then click **apply** to open a confirmation message box.

4. **Apply your changes.**

   Click **OK** to rename the directory.

Moving a CNS Directory

Moving a CNS directory from one location in the CNS to another can be done at any time.

To move a CNS directory:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Select the subdirectory to move.**

   From the CNS directory tree, select the subdirectory to be moved, then click **modify** to display the **Modify CNS Directory** page.

3. **Select a new location.**

   From the **Select a Parent for the Directory** options box, select a new location in the CNS tree. From the bottom of the options box, click **apply** to open a confirmation message box.

4. **Apply changes.**

   Click **OK** to move the directory.

Deleting a CNS Directory

Deleting a CNS directory will permanently remove it and all of its subdirectories and file system links. Deleting CNS directories will not affect physical file systems on the server.

To delete a CNS directory:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Select a subdirectory for removal.**

From the CNS directory tree, select a subdirectory. From the box at the bottom of the page, click **remove** to open a confirmation message box.

3. **Apply changes.**

   Verify your settings, then click **OK** to proceed or **Cancel** to decline.

Modifying a File System Link

The name and location of a CNS file system link can be modified.

To modify a file system link:

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Display a link in the Modify Link page.**

   From the CNS tree, select a file system link, then click **modify** to view the **Modify Link** page.

3. **As needed, change the link name or parent directory.**

   To change the name of the file system link, enter the new name in the **Link Name** text box, then click **apply.**

   To change the parent directory, select a new location in the tree from the **Select a New Parent Directory** options, then click **apply.**

4. **As needed, enable or disable caching.**

   If necessary, change the setting to enable or disable the caching of files from this file system, then click **apply.**

Deleting a File System Link

Deleting a file system link will erase the link from the CNS. The actual file system associated with the link will not be deleted.

1. **Navigate to the CNS page.**

   From the **File Services** page, click to display the **CNS** page.

2. **Remove the CNS link.**

   From the CNS tree, select a link, then click **Remove** to display a confirmation dialog, then **OK** to proceed.

Using Read Caching

Prerequisites for Read Caching

A storage server can support read caching under the following conditions:

- License keys to enable the read caching service and the Network File System (NFS) service must be installed. For more information about entering license keys, see Managing License Keys, on page 535.

- Sufficient space must be available in a Storage Pool to create the read cache.

Additionally, to support remote read caching:

- The storage server must be configured as a part of cluster. For more information on the cluster configurations, see Clusters, on page 398.
- The Cluster Name Space (CNS) feature must be properly licensed and configured. For more information on CNS, see Cluster Name Space (CNS), on page 410

Note: After the read cache license key is entered, the server/cluster must be restarted before the read caching service starts.

Configuring Read Caching

Before you can configure the read caching service, you must have already fulfilled the requirements in the section Prerequisites for Read Caching, on page 446. Then you can enable and configure read caching.

To enable and configure read caching, you must:

1. **Enable the read caching service.** To enable read caching, you must add the license key for read caching. Once the key has been added, the service will be enabled upon restart, enabling creation of read caches.

   For more information about adding license keys, see Managing License Keys, on page 535.

2. **Create a read cache on an EVS.** Because a read cache is a kind of file system, the same procedure that creates file systems also creates read caches. For information about creating a read cache or a file system, see Creating a File System, on page 138.

3. **Enable file caching.** The configuration can specify that files from some file systems should be cached, while prohibiting file caching for files from other file systems. To control the caching of files from a file system, select the file caching option when you add the file system link or export. For more information about adding a file system link to the CNS tree, see Creating a File System Link, on page 442.

4. **Set file caching options.**To control which files are eligible for caching, you must configure the file caching options. Once a file system link has been added to the CNS tree, the file system link options can be changed to control whether files from this file system can be cached. For more information on editing file system links, see Modifying a File System Link, on page 446.

Setting File Caching Options

To configure file caching options:

1. **Navigate to the Read Cache Options page.**

From the **Storage Management** page, click **Read Cache Options** to display the **Read Cache Options** page:

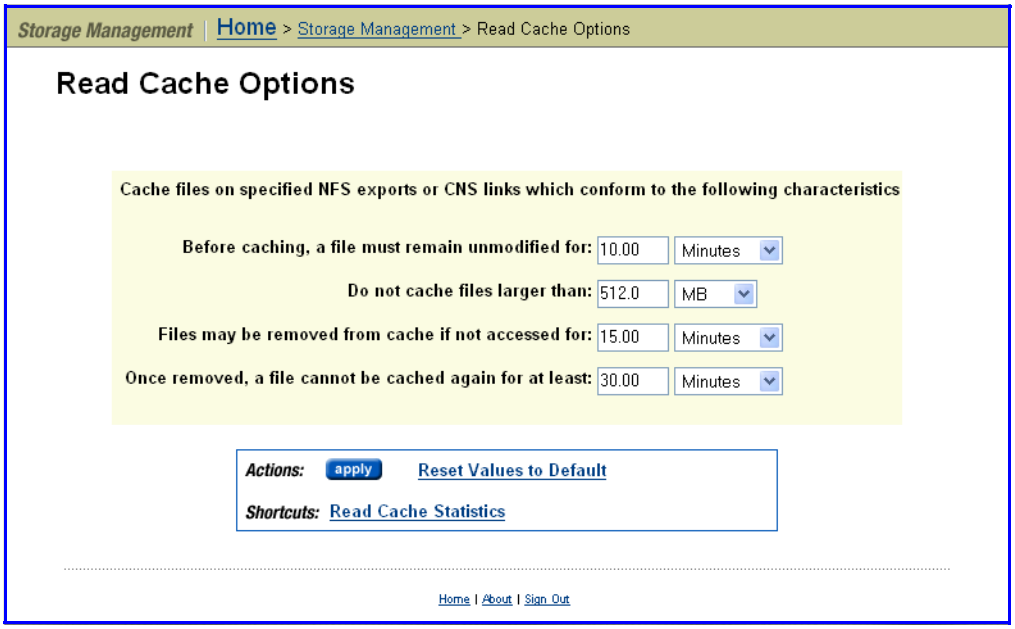

The following table describes the fields in this page:

Field	Description
Before caching, a file must remain unmodified for	Specifies how long a file must be unchanged before it is eligible for caching. The default minimum stable time is 10 minutes. This does *not* indicate how long since the file has been accessed, only that the file may not have been *changed* within this period. You can specify the time in seconds, minutes, hours, or days.
Do not cache files larger than	Limits the maximum size of a file eligible for caching. The default maximum size is 512 megabytes. Caching large files may limit the number of files that the read cache can contain. If necessary, you can expand the read cache as described in Expanding a File System, on page 153.
Files may be removed from cache if not accessed for	Which indicates the amount of time that a file will remain in the read cache without being accessed before it is designated as inactive. Inactive files are eligible for removal from the cache, and are deleted on an "as space is needed" basis, with the oldest inactive files being deleted until there is enough space for a new file to be added to the read cache. The default duration is 15 minutes. You can specify the amount of time in seconds, minutes, hours, or days.

Field	Description
Once removed, a file cannot be cached again for at least	Specifies the minimum time that must elapse before a file is re-evaluated for read caching after: • Having been read cached, and then having been flushed from the read cache for any reason. • Being evaluated for read caching, and not being cached for some reason. Reasons for not caching a file may include file size, too recent modification, or insufficient space available in the read cache. The default retry time is 30 minutes. You can specify the amount of time that must elapse before a file can be cached again in seconds, minutes, hours, or days.

2. **Set and invoke the displayed options.**

   Once you have set the options, click **apply** to apply changes, or **Reset Values to Default** to restore defaults.

Note: You can prohibit read caching of files from a particular file system when the link to that file system is added to the CNS tree. To disallow the caching of files from particular file systems, see Creating a File System Link, on page 442.

Reviewing Read Cache Statistics

Read cache statistics provide information about a read cache, including:

- **Successfully Cached Files:** The number of successfully read cached files.
- **Candidate Files Encountered:** For remote read caching, the number of read cacheable files that have been read by a remote node. For local read caching, the number of read cacheable files that have been read by the local node.
- **Files Rejected: Has Named Streams:** The number of read cacheable files that were not cached because they have associated named streams.
- **Files Rejected: Not Stable:** The number of read-cacheable files were not cached because they were modified within the window of time specified by the “Before caching, a file must remain unmodified for” setting.
- **Files Rejected: Too Large:** The number of read-cacheable files that were not cached because they exceed the size specified in the “Do not cache files larger than” setting.
- **Flushes Due To Active Set Limit:** The number of times a file was flushed from the read cache because the read cache reached its maximum number of active files. By default, a maximum of 250,000 files may be in the read cache at any one time.
- **Flushes To Reclaim Space:** The number of times a file was flushed from the read cache to free space in the read cache file system.

- **Flushes Aborted:** The number of times an unaccessed file in the read cache was not flushed because it was still considered active according to the "Files may be removed from cache if not accessed for" setting.
- **File Lock Revoked:** The number of times all files in the read cache were invalidated at the same time. This statistic also counts the times a file lock is revoked because the "real" file has been modified, which causes the cached copy to be removed from the read cache.

  Certain situations will cause the simultaneous invalidation of files in the read cache; some are the result of normal operations (like the unmounting of a file system), while others are due to error conditions.

  For local read caching, this situation occurs whenever the local file system is unmounted, or when the EVS hosting the file system is migrated to another cluster node.

  For remote read caching, this situation occurs whenever there is a loss of communication with the remote file system. The number of times all files in the read cache from a particular remote file system were invalidated at the same time. This situation occurs whenever there is a loss of communication with the remote file system.

  For example, all files from a particular remote file system are invalidated simultaneously when:

  - The remote file system is unmounted.
  - The cluster node on which the remote file system is located crashes.
  - The cluster interconnect fails.
- **Average Cached File Size:** The average size of files stored in the read cache.
- **Average Cached File Lifetime:** The average time a read cached file stays valid and can therefore service read requests. This statistic provides a very good indicator of the efficiency of the read cache. If this value is small, it can have several causes:

  - Files are being flushed from the cache too often because too many files are being cached.

    If files are being flushed from the cache too often, consider reducing the number of CNS links marked as read cacheable, or increasing the maximum number of files allowed in the read cache. (Contact SGI Global Services for assistance if you want to increase the maximum number of files allowed in the read cache.)
  - Files are being flushed from the cache because the read cache is running out of space.

    If the read cache is running out of space, you can increase the size of the read cache file system, or you can decrease the number of files that are cached (either by decreasing the maximum number of files allowed in the read cache, or by reducing the number of CNS links marked as read cacheable).

- Files that are identified as read cacheable are actually being modified too often.

  If files identified as read cacheable are being modified too often, increase the value of the "Before caching, a file must remain unmodified for" option.

To display statistics about a read cache:

1. **Navigate to the Read Cache Statistics page.**

   From the **Status and Monitoring** page, select **Read Cache Statistics** to display the **Read Cache Statistics** page:

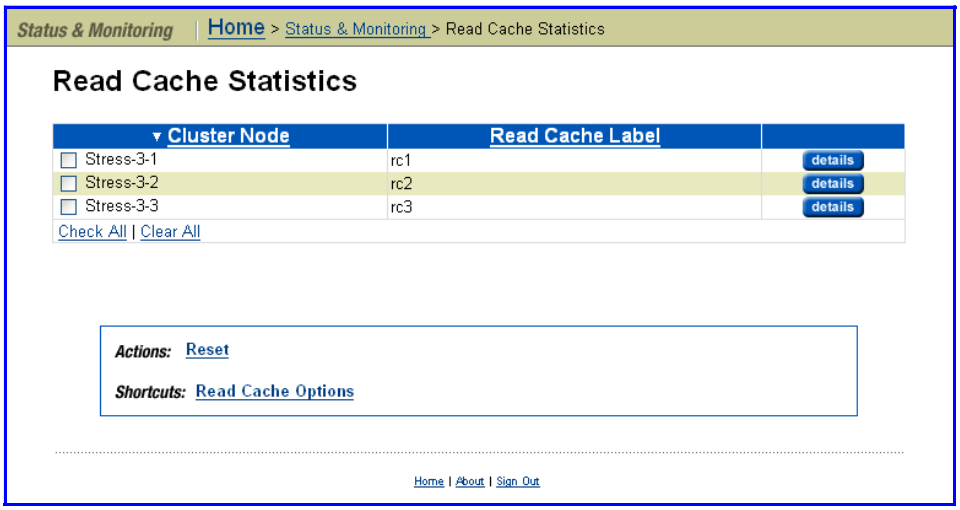

2. **Display read cache information.**

   Select a read cache, then click **details** to display its statistics page.

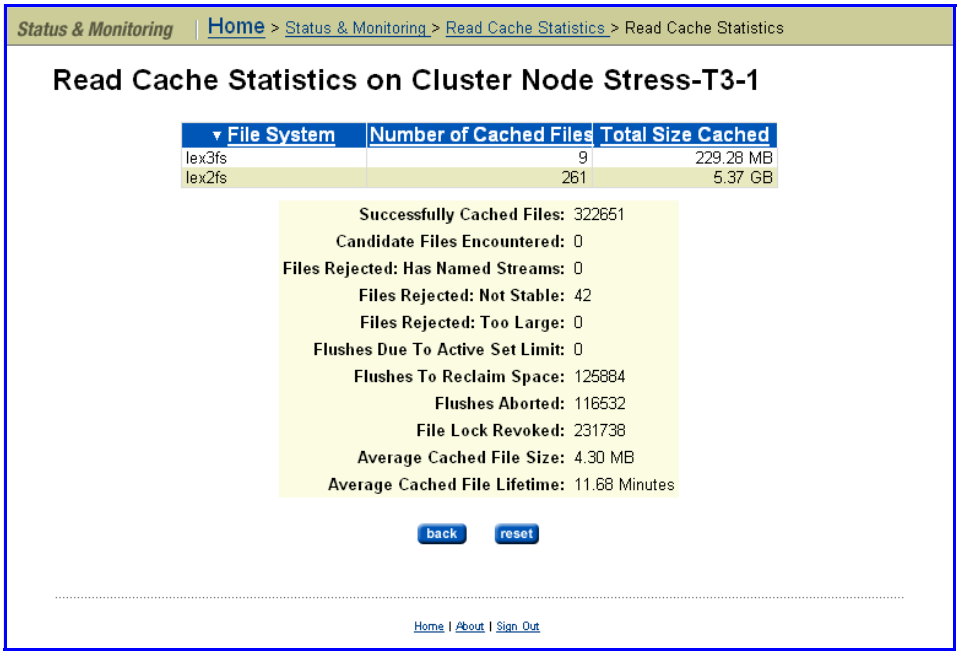

Note: The table at the top of the **Read Cache Statistics** page lists the name of each file system that currently has files in the read cache. For each file system that currently has cached files, the table lists the number of files and their total (aggregated) size.

3. **Reset or return.**

   Once you have reviewed the available statistics, you can:

   - Click **reset** to restart the gathering of statistics (you will lose previously gathered statistics for the read cache).
   - Click on the **back** button to return to the **Read Cache Statistics** page.

Deleting a Read Cache

For information on how to completely delete a read cache, see Deleting a File System, on page 150.

Read Caching Considerations

The following recommendations are intended to take full advantage of read caching:

- Because remote read caching requires CNS, you should review the Cluster Name Space considerations, see CNS Usage Considerations, on page 411.
- **In a cluster configuration, define one EVS per cluster node, and assign a read cache to each EVS.**
- **Balance loads by moving file systems, instead of migrating EVS.** If you migrate an EVS containing a read cache, the files in the read cache become invalidated and, assuming they are still cacheable, they would have to be cached again after the next read request.

  If an EVS containing a read cache is migrated to another cluster node that already has a read cache, the files in the migrated read cache are invalidated, and only the read cache that was not migrated will be used. If the EVS is migrated back to its original cluster node, the read cache will be used again, assuming another read cache has not been created on that cluster node in the interim.
- **Do not relocate read caches**. If you relocate a read cache, the files in the read cache become invalidated and, assuming they are still cacheable, the previously cached files would have to be cached again after the next read request.

10 Status and Monitoring

The IS-NAS Server/Titan Server architecture includes the following status and monitoring tools:

Topic	Overview and Instructions
Storage System Status	SGI Storage System Status, on page 453
Storage Server Statistics	Storage Server Statistics, on page 465
Event Logging and Notification	Event Logging and Notification, on page 497
File System Auditing	File System Auditing, on page 514
FTP Auditing	FTP Auditing, on page 524
Monitoring Fibre Channel Switches	Monitoring Fibre Channel Switches, on page 528

Status & Monitoring Overview

Web Manager provides comprehensive and integrated management of the storage server and its storage subsystem. Its management pages provide color-coded information about the status of the various installed devices. Web Manager also provides a comprehensive event logging and alerting mechanism, which can notify the system administrator, as well as SGI Global Services, as soon as a problem occurs. Alerts are issued through email, SNMP, Syslog, and Windows pop-ups.

File system auditing monitors and records file system operations performed through the CIFS protocol. File system operations such as file access and deletion are recorded in the server's file system audit log. You can then view the file system audit log through a remote Windows Event Viewer, and save the log entries in ".evt" format for later review. See File System Auditing, on page 514 for more information.

SGI Storage System Status

Web Manager provides a flexible, customizable, and easy-to-use interface, displaying the status of each managed device in the storage system. Ethernet-connected auxiliary devices can be added to the System Monitor as managed objects, so that the status of these devices will be displayed. The **System Monitor** page provides a central management console for the management and status monitoring of all devices that comprise the network storage system.

Using the Server Status Console

Summary status information for the currently managed server can be viewed from the Web Manager's **Server Status Console,** located in the server's **Home** page:

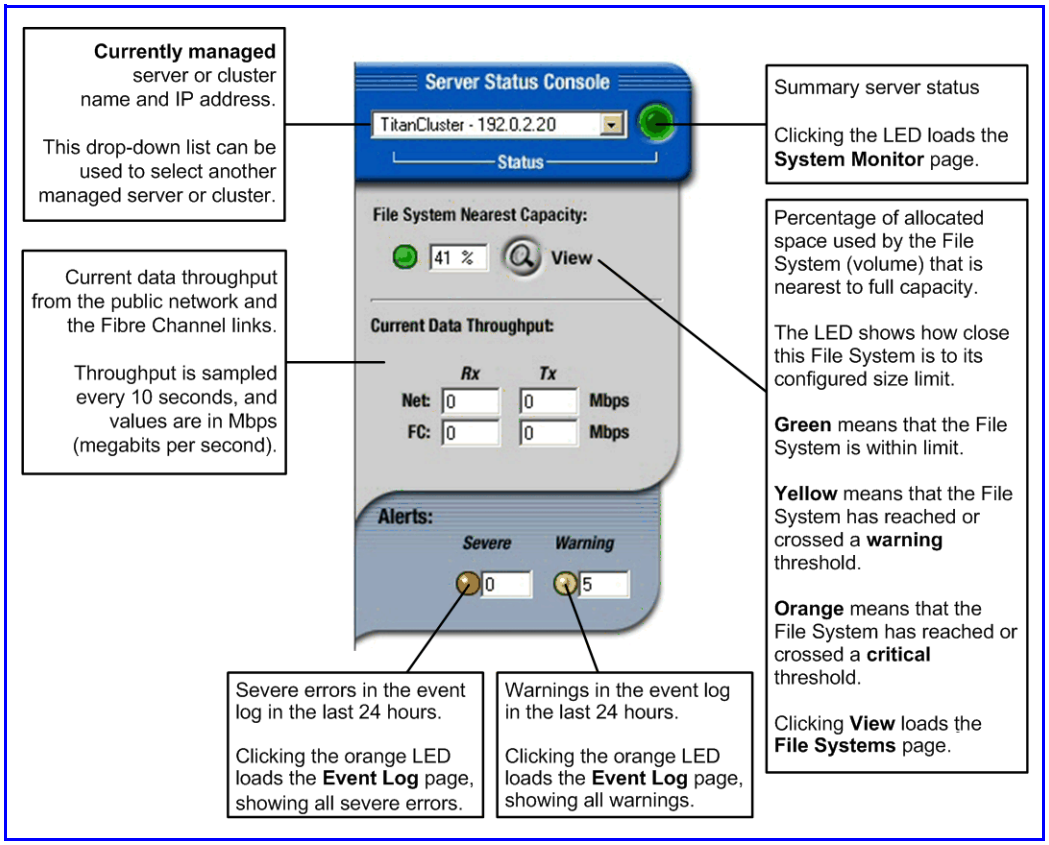

Checking the System Status

To display the components contained within the system and view the status of the currently managed server, navigate to the **Status & Monitoring** page, then click to display the **System Monitor** page:

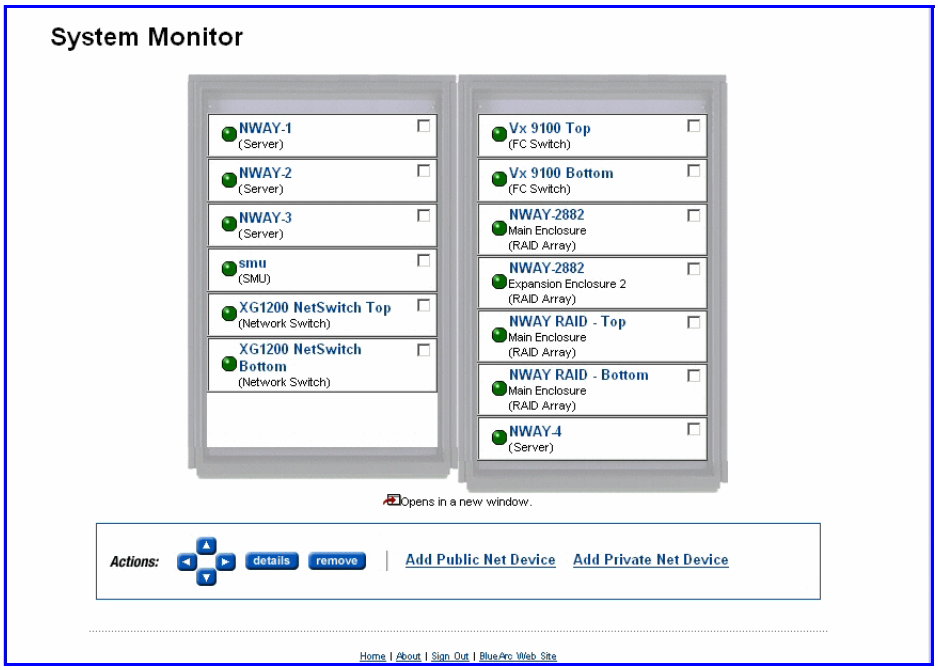

This page provides a graphical representation of the main components that make up the system, displaying status and links to more information (see table below).

To change the position of any of the items in the System Monitor, select the item's checkbox and use the arrows in the **Action** section.

Note: The System Monitor reflects a 60 second delay for status information cached by the SMU.

When displaying status of a device using a colored LED, the following conventions apply:

Color	Status	Means that the item is
	Information	Is operating normally and not displaying an alarm condition
	Warning	Needs attention, but does not necessarily represent an immediate threat to the operation of the system
	Severe Warning	Has failed in a way that poses a significant threat to the operation of the system

A system can contain the following basic components:

Component	Description	Action when clicking the component	Action when clicking the details button
Server	The server provides multiple Gigabit Ethernet interfaces to the network and multiple Fibre Channel interfaces to the main enclosure. In a cluster configuration, there are up to four nodes (servers).	Loads the **Server Status** page.	
Storage Enclosure	A storage enclosure contains disk drives, power supplies, and RAID controllers.	Loads the **Enclosure Status** page.	Loads the **System Drives** page.
Expansion Enclosure	An expansion enclosure contains disk drives and power supplies, but does not contain RAID controllers.	Loads the **Enclosure Status** page.	Loads the **System Drives** page.
SMU	The System Management Unit	Loads the **SMU System Status page.**	
System Power Unit	This component is also known as an uninterruptible power supply (UPS).	Loads the **UPS Status** page.	Loads the **UPS Configuration** page.
NDMP Backup Devices	The server monitors its FC links every 60 seconds, automatically detecting the presence of backup devices and adding them to the System Monitor. Since the server can be connected into a FC network shared with other servers, it does not automatically make use of backup devices found on its FC links. Usage of backup devices is enabled/disabled through the **NDMP Device List** page.	Loads the **NDMP Device Access Details** page.	Loads the **NDMP Details** page for the device if the device can be contacted, or loads the **NDMP Device List** page if the device cannot be contacted.
FC Switches	FC switches (and cables) connect FC devices, generally storage arrays, to the server(s).  **Note:** FC Switches are added to the System Monitor automatically, after being added through the **FC Switches** page.	Loads either the embedded management utility for the switch, or the **FC Switch Details** page for the switch, depending on the protocol specified when the switch was added. (See Adding FC Switches, on page 531 for more information.)	Loads the **FC Switch Details** page.

Component	Description	Action when clicking the component	Action when clicking the details button
Other Components	Any component can be added to the system monitor. If the device supports a web-based management interface, the management interface can be launched directly from the server management interface.	Loads the embedded management utility for the device. For example, for a storage enclosure, it loads the **Home Page** for the device.	Loads either the add **Public Net Device** or the add **Private Net Device** page. Settings for the component can be changed from this page.

Checking the Status of a Server Unit

To check the status of a server, navigate to the **Status & Monitoring** page, then click to display the **Server Status** page.

- For a stand-alone server that is not part of a cluster, the **Server Status** page appears:

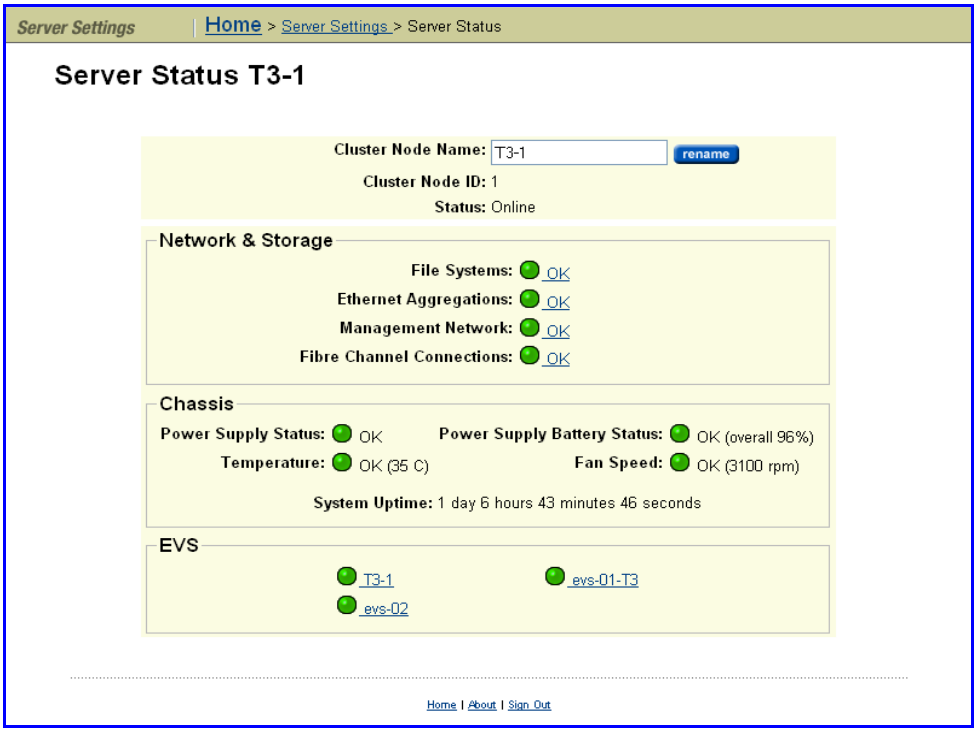

The table below describes the items on this page:

Item	Description
Cluster Node Name	The server/node name (label). To change the server/node name, enter the new name in the field, then click **rename**.
Cluster Node ID	The ID assigned to the node.
Status	An indicator of node status. • **Online**: Node has completed booting. • **Unknown**: Node has not yet booted. • **Up**: Node is booting (displayed only while the node is booting). • **Not up**: node is shutting down (displayed only while the node is shutting down). • **Dead**: Node has failed to go online after booting.
Network and Storage	
File Systems	Overall indicator of file system status: • **OK.** All file systems up and operational. • **Failed.** One or more file systems has failed. Click on the status link to display the **File Systems** page, which lists all file systems assigned to the EVS in that cluster node.
Ethernet Aggregations	Overall status of Ethernet aggregations in the server/node: • **OK.** All aggregated ports are up and linked. • **Degraded.** One or more ports in an aggregation has failed. • **Failed.** All ports in an aggregation have failed. Click on the status link to display the **Link Aggregation** page, which lists all aggregations (trunks) in the server/node.
Management Network	Overall status of the management network: • **OK.** Links are up and heartbeats are being received. • **Failed.** No heartbeats are being received, and the links may be up or down. Click on the status link to display the **Ethernet Statistics** page, which lists information about the management port and the aggregated Ethernet ports in the server/node.
Fibre Channel Connections	An overall status indicator for the Fibre Channel ports in the server/node: • **OK.** All ports up and operational. • **Degraded.** Some ports up and operational, but one or more has failed. • **Failed.** All ports have failed. Click on the status link to display the **Fibre Channel Statistics Per Port** page, which lists all Fibre Channel ports in use in the server/node.

Item	Description
Chassis	
Power Supply Status	A status indicator for the power supply units (PSUs): • **OK.** Both PSUs are installed and operating normally. • **Degraded.** One PSU not responding to queries, which may mean that it is switched off. • **Failed.** One PSU not responding to queries, and it has failed or is not present.
Temperature	Status indicator for temperature of the server/node chassis. • **OK.** Within the normal operating range. • **Degraded.** Above normal, but not yet critical. • **Failed.** Critical. When available, the temperature in the chassis is also displayed. The displayed temperature is the highest reported temperature of any of the boards in the chassis.
Power Supply Battery Status	Status of the power supply battery: • **OK.** Capacity and voltage within the normal operating range. • **Degraded.** Capacity and/or voltage below normal. This status should be considered a *warning*; if it continues, the PSU battery should be replaced. • **Failed.** Capacity and/or voltage below acceptable minimum, or the PSU battery is not responding to queries. This status indicates a failure; the PSU battery should be replaced. When available, the level of battery charge is also displayed.
Fan Speed	Status of fans in the server/node chassis: • **OK.** All fans operating normally. • **Degraded.** One or more fans spinning below normal range. • **Failed.** At least one fan has stopped completely, or is not reporting status. When available, the chassis fan speed is also displayed. The displayed fan speed is the slowest reported speed of any of the three fans. An error message may be displayed, even if it does not correspond with the slowest fan.
System Uptime	Duration since last reboot of the server/node.
EVS	
EVS	This section displays the names (labels) of EVSs assigned to the node, and displays a status indicator for each EVS: • **Green.** Online and operational. • **Amber.** Offline, but listed here because the server/node is hosting the administrative EVS. • **Red.** Failed. Click on the EVS name to display the **EVS Details** page for that EVS.

- For a cluster node, the **Cluster Configuration** page appears:

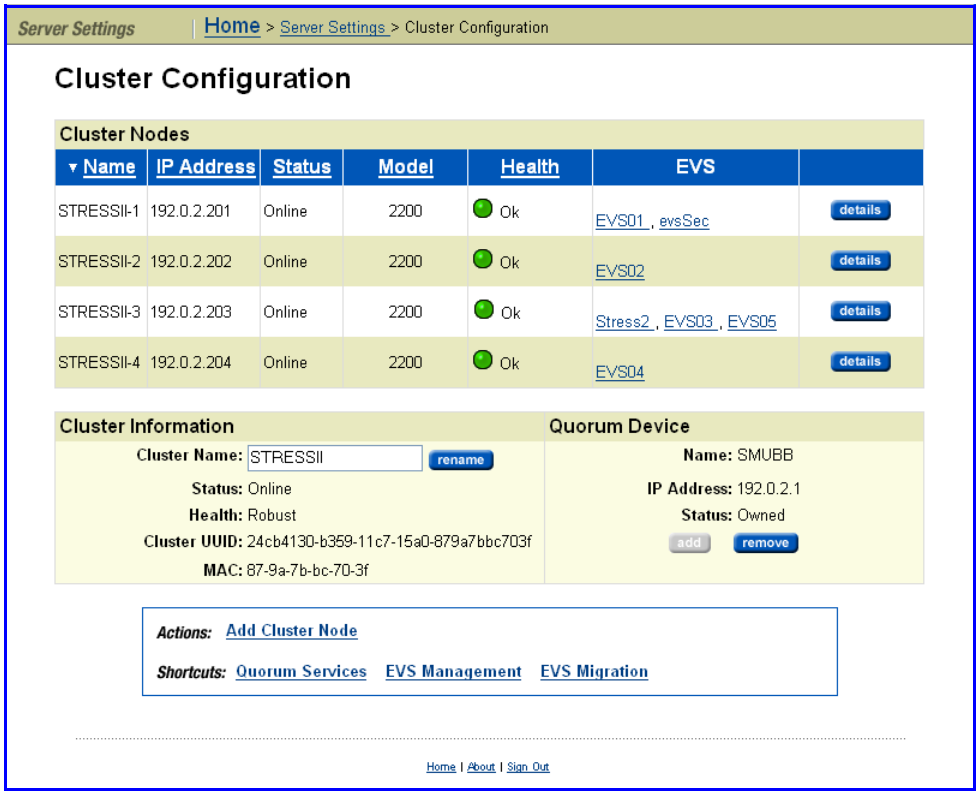

The table below describes the items on this page:

Item	Description
Cluster Nodes	
Name	Node name.
IP Address	IP address of the cluster node. This IP address is on the private management network, which connects devices within the cluster.
Status	Node status: • **Online.** The cluster node is a part of the cluster and is exchanging heartbeats with the other cluster members. • **Offline.** The cluster node is no longer exchanging heartbeats with the other cluster members. This may be caused by a reboot or a fault condition. Services cannot be migrated to a cluster node in this state. • **Dead.** No heartbeats from this cluster node for a significant period. Services cannot be migrated to a cluster node in this state. • **Unknown.** The node has not been online since the cluster was started.
Model	Server model, if available.
Health	Worst-case status from each node: • **OK.** Operating normally, with no failures. • **Degraded.** Operating, but with one or more failures in connectivity. The problem may be with the cluster interconnect, the management network, or quorum device communication. • **Failed.** Not operating, due to one or more failures in connectivity. The problem may be with the cluster interconnect, the management network, or the quorum device communication.
EVS	This section displays the names (labels) of EVSs hosted on each cluster node. Click on the EVS name to display the **EVS Details** page for that EVS.
Cluster Information	
Cluster Name	Cluster name. To rename the cluster, enter the new cluster name in the edit box, then click on the **rename** button.
Status	Unless rebooting, the status will be "Online."
Health	Cluster health: • **Robust.** Operating normally, with no failures in the cluster interconnect, the management network, or quorum device communication. • **Degraded.** Operating, but one or more nodes has failed or there has been a failure in connectivity.
Cluster UUID	UUID (unique ID) of the cluster. This string provides a unique identifier for each cluster when there are several clusters on a network.
MAC	MAC address of the cluster.

Item	Description
Quorum Device	
Name	SMU on which the QD resides.
IP Address	IP address of the SMU on which the QD resides.
Status	QD status: • **Configured.** Attached to the cluster, but vote not needed. The QDs vote is not needed when any cluster contains an odd number of operational nodes. • **Owned.** Attached to the cluster and owned by a specific cluster node. • **Not up.** Cannot be contacted. • **Seized.** Taken over by another cluster.

Checking the Status of a UPS

To check the status of a power unit, navigate to the **Server Settings** page, then click **UPS Configuration** to display the **UPS Configuration** page.

At the bottom of the page is the **UPS Devices** table listing the configured UPS devices. For each UPS device, the table lists its IP address, percentage charged, runtime remaining on the UPS, and online status (whether currently supplying power).

The **UPS Configuration** page is also where you configure ethernet-connected UPS devices and specify how the NAS server/cluster reacts to power failures. For more information on the **UPS Configuration** page, see Managing Uninterruptible Power Supply Usage (Titan Server only), on page 60.

Checking the Status of the SMU

To check the status of the SMU:

1. **To view the current status of the SMU, from the SMU Administration page, click SMU Status.**

The **SMU Status** page displays the services available, the current status of each service and the option to restart each service if necessary:

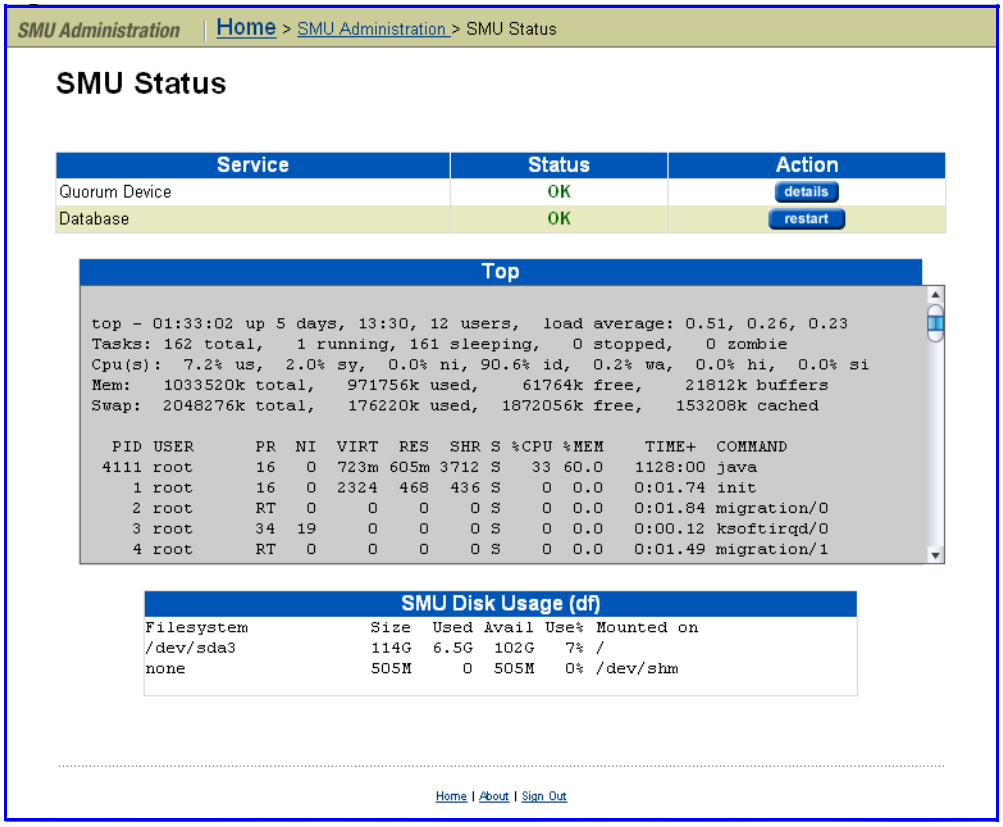

The following table describes the components of this page:

Services	Description
Services	**Quorum Device.** Used by a cluster which has become partitioned by a network failure, to determine which partition is allowed to talk to the storage. **Database.** Allows communication between the SMU and the servers.
Status	The desired state of these services is *OK*, and if a service is not running correctly, an error message will be displayed.
Action	If the Quorum Device service is not running correctly, an error message will be displayed in the **Status** column. To see more information about the status, click **details**. If the Database service status is not running correctly an error message will be displayed in the **Status** column. You can restart the Database service by clicking **restart**.
Top	Displays the status of the SMU Operating System. This is the actual output gathered from the Unix `top` command and indicates the current running status of the SMU's internal processes.

Services	Description
SMU Disk Usage (df)	Displays the details of the space used in each of the partitions of the SMU hard disk. This is the actual output gathered from the Unix `df` command.

Monitoring Multiple Servers

Servers can be managed from the SMU. To display all of the currently managed servers, navigate from the **SMU Administrator** page to the **Managed Servers** page:

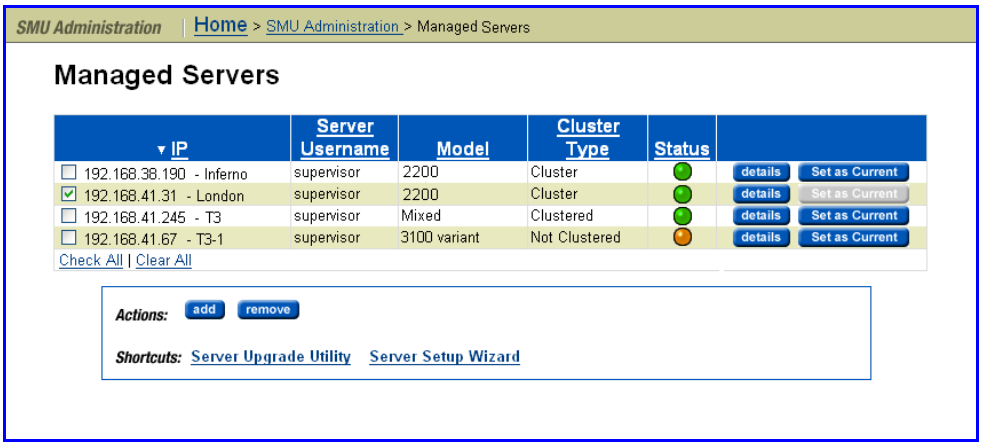

The following table describes the components of this page:

Item/Field	Description
IP Address	IP address of the server. This should be the Administration Services IP address, as used on the private management network (for example, 192.0.2.x).
Username	Username of server.
Model	Displays the storage server model number. For a cluster with different server models, this field displays "mixed", and the specific server models can be displayed in the Cluster Configuration page.
Cluster Type	Cluster type (for example, *Single Node* or *Clustered*).
Status	Current status of the server: Green indicates that the server is operating normally (i.e., not showing an alert condition) Amber indicates a warning (i.e., operating normally, however, action should be taken to maintain normal operation) Red indicates a critical condition (i.e., the server is no longer functioning properly).
Details	Link to a page displaying detailed information about contacting or managing the server.
Set as Current	Select the server as the *currently managed server.*

The following **Actions** are available:

- Click **add** to add a managed server.
- Click **remove** to remove one or more selected servers.

When a server is removed:

- Replication policies and schedules are deleted.
- Data migration policies and schedules are deleted.
- The system monitor for that server is deleted.
- Racks managed by that server are deleted.

Storage Server Statistics

Servers provide extensive statistics that can be used to monitor their operation. These include:

- Networking (Ethernet and TCP/IP)
- Fibre Channel
- File access protocols (CIFS, NFS, and FTP)
- Block access protocols (iSCSI)
- Management access (SNMP and SSC for a IS-NAS Server, and Telnet, SSC, and SNMP for a Titan Server)
- Virus scanning

Ethernet Statistics

Ethernet statistics (per port in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Viewing Ethernet Statistics

To view Ethernet statistics, navigate from the **Status & Monitoring** page to the **Ethernet Statistics** page:

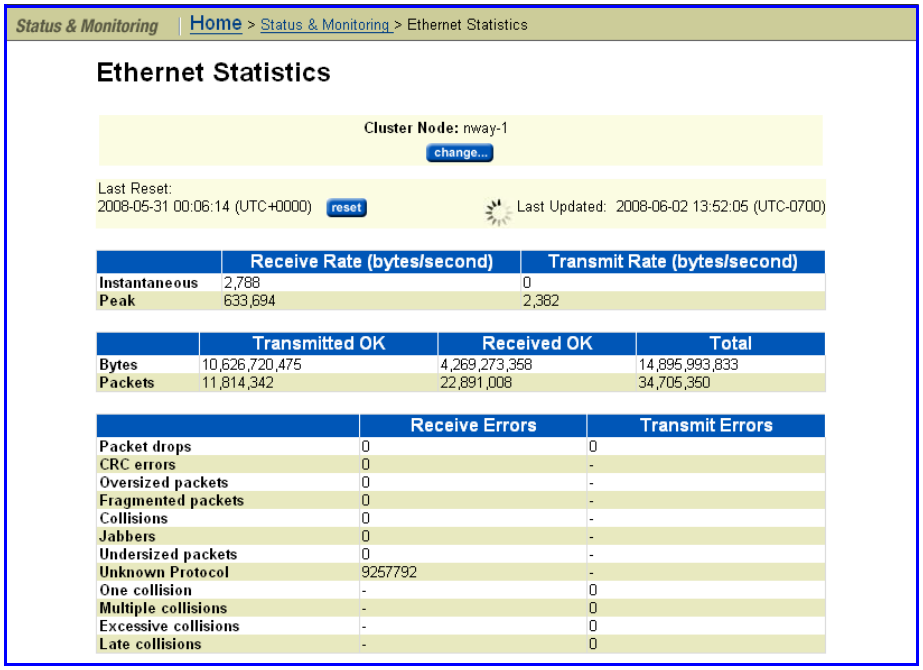

The following table describes the fields and buttons in this page:

Item/Field	Description
Cluster Nodes	When connected to a cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click the **change** button.
change	Clicking the **change** button opens the Select a Cluster Node page where you can select a different node for which to view statistics.
Last Reset	Displays the date and time the statistics on this page were reset. To reset the statistics to zero, click the **reset** button.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Receive Rate	The amount of data received in bytes per second. Includes the **Instantaneous** (current) and **Peak** throughput.
Transmit Rate	The amount of data transmitted in bytes per second. Includes the **Instantaneous** (current) and **Peak** throughput.
Transmitted OK	The total number of **Bytes** and **Packets** successfully transmitted.
Received OK	The total number of **Bytes** and **Packets** successfully received.
Errors	Lists the number of **Receive Errors** and **Transmit Errors** logged on the server/node. The following types of errors are reported (For the five **collision** errors, the values will always be 0 for the file-serving ports --ge, tg, and ag-- as only full duplex point-point connections are supported): *Packet drops, CRC Errors, Oversized packets, Fragmented Packets, Collisions, Jabbers, Undersized packets, Unknown Protocol, One collision, Multiple collisions, Excessive collisions, Late collisions.*

Viewing Aggregated Ports or Per Port Ethernet Statistics

To view Ethernet statistics for individual ports or aggregated for all ports, navigate from the **Status & Monitoring** page and click **Ethernet Statistics (per port)**. This opens the **Aggregated Ports Ethernet Statistics** page. To view these statistics for individual ports, from this page, click the link **Physical Ports Ethernet Statistics** which opens the **Physical Ports Ethernet Statistics** page.

The following table describes the fields in both of these pages:

Item/Field	Description
Cluster Nodes	When connected to a NAS server cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click the **change** button.
change	Clicking the **change** button opens the Select a Cluster Node page where you can select a different node for which to view statistics.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Bytes	Displays the number of bytes **Transmitted OK** and **Received OK**, and the **Total** number of bytes.
Packets	Displays the number of packets **Transmitted OK** and **Received OK**, and the **Total** number of packets.
Receive Throughput Rate	The receive rate in bytes/second for the **Instantaneous** (current) and **Peak** throughput.
Transmit Throughput Rate	The transmit rate in bytes/second for the **Instantaneous** (current) and **Peak** throughput.
Receive Errors	Lists the number of **Receive Errors** logged on the server/node. The following types of errors are reported: *Packet drops, CRC Errors, Oversized packets, Fragmented packets, Collisions, Jabbers, Undersized packets, Unknown Protocol.*
Transmit Errors	Lists the number of **Transmit Errors** logged on the server/node. The following types of errors are reported: *Packet drops, one collision, multiple collisions, excessive collisions, and late collisions.*
Link Status	Indicates the condition of the link. The possible values are either Up or Down.
MAC Addresses	Displays the MAC address of each port. For the aggregated statistics, the MAC address is for the port to which the aggregation is linked.
Last Reset Time	Displays the date and time when the statistics for this port were last reset to zero.
Select to Reset Statistics	Fill this checkbox for each port whose statistics you want to reset to zero. The statistics are reset when you click the **reset** button.
reset	Click **reset** to reset all statistics of the selected ports to zero.

TCP/IP Statistics

The TCP/IP statistics display activity since the last server reboot or since the TCP/IP statistics were last reset. Both per-port and overall statistics are available. The statistics are updated every ten seconds.

Viewing TCP/IP Statistics

To view TCP/IP statistics, navigate from the **Status & Monitoring** page to the **TCP/IP Statistics** page:

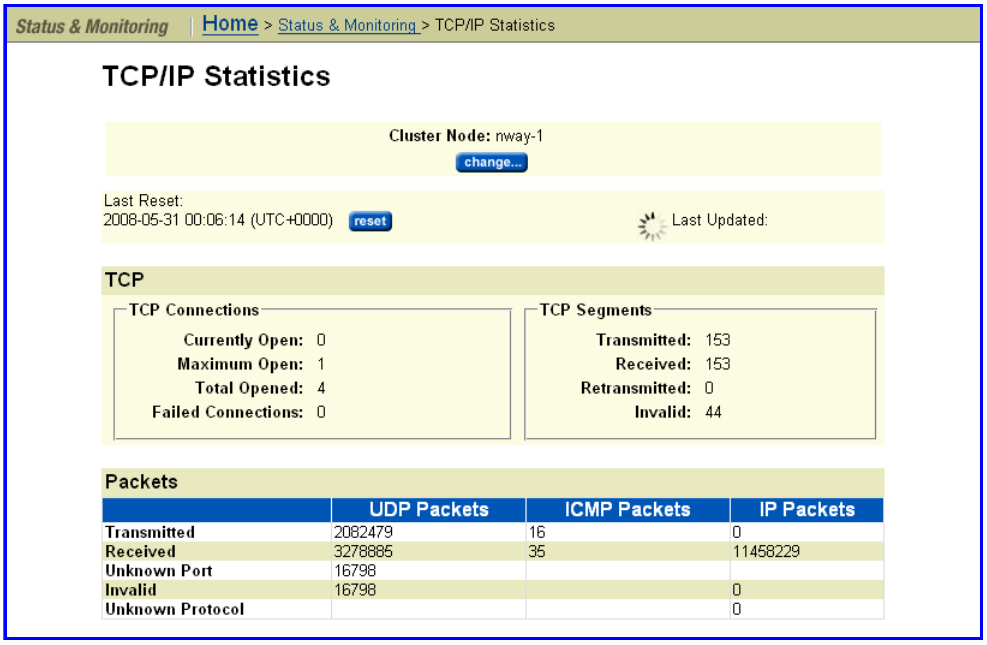

The following table describes the components of this page:

Item/Field	Description
Cluster Node	When connected to a NAS server cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click the **change** button.
change	Clicking the **change** button opens the Select a Cluster Node page where you can select a different node for which to view statistics.
Last Reset	Displays the date and time the statistics on this page were reset. To reset the statistics to zero, click the **reset** button.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.

Item/Field	Description
TCP Connections	Displays statistics about the TCP connections. **Currently Open** is the number of currently open connections. **Maximum Open** is the maximum number of connections opened at one time since the last reset. **Total Opened** is the number of connections that have been opened since the last reset. **Failed Connections** is the number of failed incoming and outgoing connections.
TCP Segments	Displays the number of **Transmitted**, **Received** and **Retransmitted** TCP segments since the last reset. Also displays the number of segments received with **Invalid** TCP checksums.
UDP Packets	Displays the number of **Transmitted** and **Received** UDP packets; the number of packets received on an unknown port with no UDP listener (**Retransmitted**); the number of packets received with **Invalid** UDP checksums
ICMP Packets	Displays the number of **Transmitted** and **Received** ICMP packets.
IP Packets	Displays the number of **Transmitted** and **Received** IP packets; the number of **Unknown Protocol** packets; the number of **Invalid** IP packets. An IP packet is invalid when any of the following is invalid: • The header checksum • The length field (too long for the packet) • The source address • The destination address (this is the most common cause)

Viewing Aggregated Ports or Per Port TCP/IP Statistics

To view TCP/IP statistics for individual ports or aggregated for all ports, navigate from the **Status & Monitoring** page and click **TCP/IP Statistics (per port)**. This opens the **Aggregated Ports TCP/IP Statistics** page. To view these statistics for individual ports, from this page, click the link **Physical Ports TCP/IP Statistics** which opens the **Physical Ports TCP/IP Statistics** page.

The following table describes the fields in both of these pages:

Item/Field	Description
Cluster Node	When connected to a NAS server cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click the **change** button.
change	Clicking the **change** button opens the Select a Cluster Node page where you can select a different node for which to view statistics.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.

Item/Field	Description
TCP Segments	Displays the number of **Transmitted**, **Received** and **Retransmitted** TCP segments since the last reset. Also displays the number of segments received with **Invalid** TCP checksums.
UDP Packets	Displays the number of **Transmitted** and **Received** UDP packets; the number of packets received on a port with no UDP listener (**Retransmitted**); the number received with **Invalid** UDP checksums.
ICMP Packets	Displays the number of **Transmitted** and **Received** ICMP packets.
IP Packets	Displays the number of **Transmitted** and **Received** IP packets; the number of **Unknown Protocol** packets; the number of **Invalid** IP packets. An IP packet is invalid when any of the following is invalid: • The header checksum • The length field (too long for the packet) • The source address • The destination address (this is the most common cause)
Last Reset Time	Displays the date and time when the statistics for this port were last reset to zero.
Select to Reset Statistics	Fill this checkbox for each port whose statistics you want to reset to zero. The statistics are reset when you click the **reset** button.
reset	Click **reset** to reset all statistics of the selected ports to zero.

Viewing TCP/IP Detailed Statistics

To view TCP/IP detailed statistics, navigate from the **Status & Monitoring** page to the **Detailed TCP/IP Statistics** page:

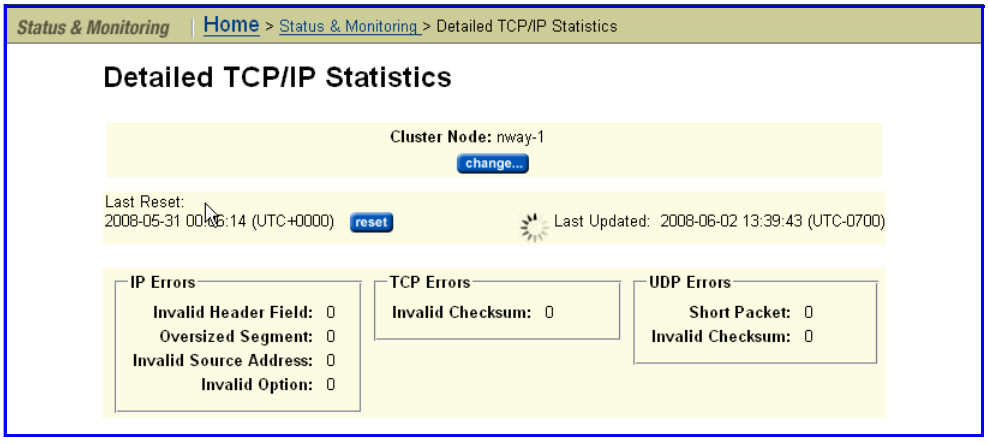

The following table describes the components of this page:

This field	Shows the number of
Cluster Nodes	When connected to a cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click the **change** button.
change	Clicking the **change** button opens the **Select a Cluster Node** page where you can select a different node for which to view statistics.
Last Reset	Displays the date and time the statistics on this page were reset. To reset the statistics to zero, click the **reset** button.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
IP Errors	
Invalid Header Field	Displays the number of IP errors arising from an invalid header field.
Oversized Segment	Displays the number of fragmented TCP packets greater than the Maximum Transmission Unit (MTU) size when re-assembled. The transmitting source made an error or the packet was corrupted in transit.
Invalid Source Address	Displays the number of IP packets with an invalid source address (often caused by DHCP broadcast requests using the source address 0).
Invalid Option	Displays the number of IP packets that were not decoded because the IP option length was invalid. The transmitting source made an error or the packet was corrupted in transit.
TCP Errors	
Invalid Checksum	Displays the number of invalid TCP packet checksums. The transmitting source made an error or the packet was corrupted in transit.
UDP Errors	
Short Packet	Displays the number of UDP packets that were too short for the UDP header or length. The transmitting source made an error or the packet was corrupted in transit.
Invalid Checksum	Displays the number of invalid UDP packet checksums. The transmitting source made an error or the packet was corrupted in transit.

Fibre Channel Statistics

The Fibre Channel (FC) statistics for the server (per port in ten-second time slices) are available for activity since the previous reboot or since the point when statistics were last reset.

Viewing the Fibre Channel Statistics

To view the Fibre Channel statistics, navigate from the **Status & Monitoring** page to the **Fibre Channel Statistics** page:

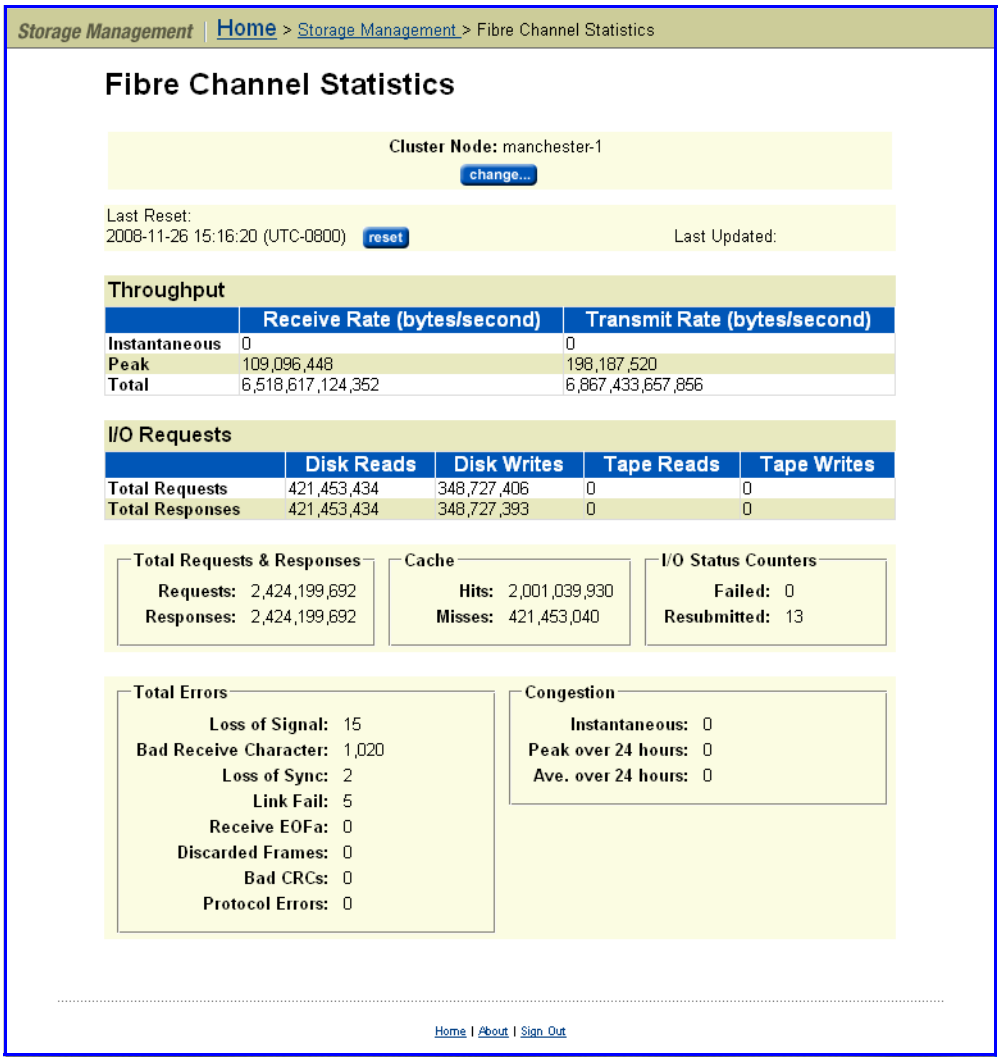

The following table describes the components in this page:

Item/Field	Description
Cluster Node	When connected to a NAS server cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click **change** and select the node for which you want to see the statistics. This field will not appear when connected to a stand-alone server.
change	Clicking the **change** button opens the **Select a Cluster Node** page where you can select a different node for which to view statistics.

Item/Field	Description
Last Reset	Displays the date and time the statistics on this page were reset. To reset the statistics to zero, click the **reset** button.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Receive Rate	The amount of data received in bytes per second. Includes the Instantaneous (current), Peak, and Total throughput.
Transmit Rate	The amount of data transmitted in bytes per second. Includes the Instantaneous (current), Peak, and Total throughput.
Disk Reads	The number of read requests that the attached disk devices have received, and the number of responses sent.
Disk Writes	The number of write requests that the attached disk devices have received, and the number of responses sent.
Tape Reads	The number of read requests that the attached tape devices have received, and the number of responses sent.
Tape Writes	The number of write requests that the attached tape devices have received, and the number of responses sent.
Total Requests and Responses	The number of data requests that the server has received, and the number of responses it has sent out. These include both requests that have been sent to the storage devices and requests that the cache has served internally.
Cache	Number of hits (requests that the cache has served) and misses (requests not served by the cache and passed to the storage subsystem).
I/O Status Counters	Numbers of failed and resubmitted input and output requests.
Total Errors	Number of errors logged at the Fibre Channel interface. The following types of errors are reported: *Loss of Signal*, *Bad Receive Character*, *Loss of Sync, Link Fail*, *Receive EOF*, *Discarded Frames*, *BAD CRCs*, and *Protocol Errors*.
Congestion	Displays congestion rates. Includes the Instantaneous (current), Peak (over the past 24hrs), and Average (Ave. over the past 24hrs) rates.

To View Per Port Fibre Channel Statistics

To view per port Fibre Channel statistics, navigate from the **Status & Monitoring** page to the **Fibre Channel Statistics (per port)** page, which displays statistics for each of the defined ports:

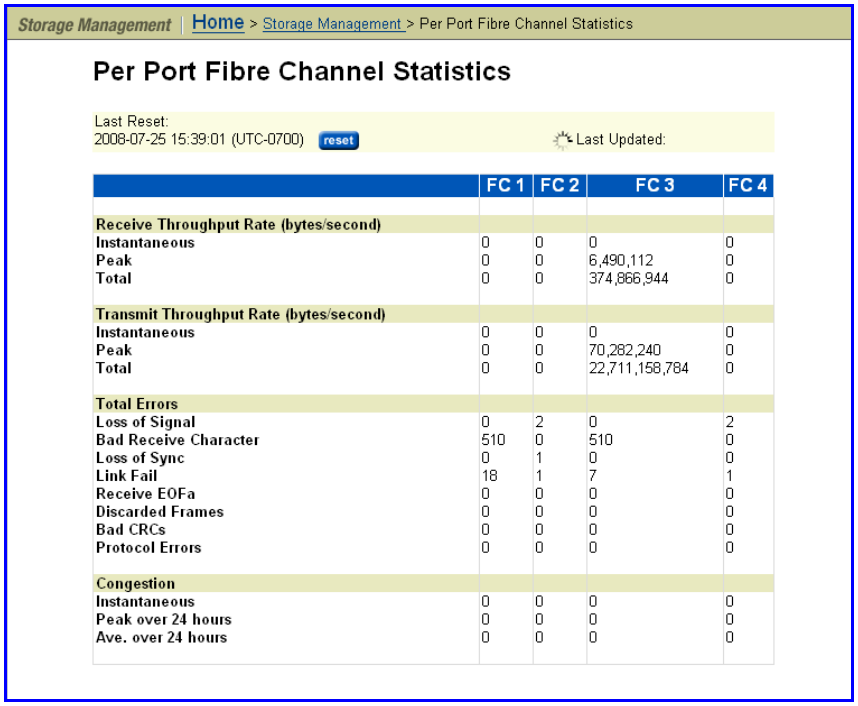

	FC 1	FC 2	FC 3	FC 4
Receive Throughput Rate (bytes/second)				
Instantaneous	0	0	0	0
Peak	0	0	6,490,112	0
Total	0	0	374,866,944	0
Transmit Throughput Rate (bytes/second)				
Instantaneous	0	0	0	0
Peak	0	0	70,282,240	0
Total	0	0	22,711,158,784	0
Total Errors				
Loss of Signal	0	2	0	2
Bad Receive Character	510	0	510	0
Loss of Sync	0	1	0	0
Link Fail	18	1	7	1
Receive EOFa	0	0	0	0
Discarded Frames	0	0	0	0
Bad CRCs	0	0	0	0
Protocol Errors	0	0	0	0
Congestion				
Instantaneous	0	0	0	0
Peak over 24 hours	0	0	0	0
Ave. over 24 hours	0	0	0	0

The following table describes the components of this page:

Item/Field	Description
Cluster Nodes	When connected to a NAS server cluster, this field indicates the node for which the statistics are displayed. To view statistics for another node, click the **change** button.
change	Clicking the **change** button opens the Select a Cluster Node page where you can select a different node for which to view statistics.
Last Reset	Displays the date and time the statistics on this page were last reset. To reset the statistics to zero, click the **reset** button.
Last Updated	Displays the date and time this page was refreshed. The page automatically refreshes every 10 seconds.
Receive Throughput Rate	The receive rate in bytes/second for the **Instantaneous** (current), **Peak**, and **Total** throughput.
Transmit Throughput Rate	The transmit rate in bytes/second for the **Instantaneous** (current), **Peak**, and **Total** throughput.

Item/Field	Description
Total Errors	Lists the number of errors logged on the Fibre Channel ports. The following types of errors are reported: *Loss of Signal, Bad Receive Character, Loss of Sync, Link Fail, Receive EOF, Discarded Frames, BAD CRCs*, and *Protocol Errors*.
Congestion	Displays congestion rates. Includes the **Instantaneous** (current), **Peak** (over the past 24hrs), and **Average** (over the past 24hrs) rates.

File and Block Protocol Statistics

The server provides statistics to monitor data access via the following network protocols:

- Network File System (NFS)
- Common Internet File System (CIFS)
- File Transfer Protocol (FTP)
- Internet Small Computer System Interface (iSCSI)

Viewing NFS Statistics

NFS statistics display activity since the last server reboot or since NFS statistics were last reset. They are updated every ten seconds. To view NFS Statistics, navigate from the **Status & Monitoring** page to the **NFS Statistics** page.

This page displays the current number of RPC calls of different types that clients have issued to the NAS server/cluster node. For a cluster node, the node is shown in the **Cluster Node** field, and you can change nodes by clicking the **change** button. These statistics are updated every ten seconds. You can reset all the values displayed on this page to zero by clicking **reset**.

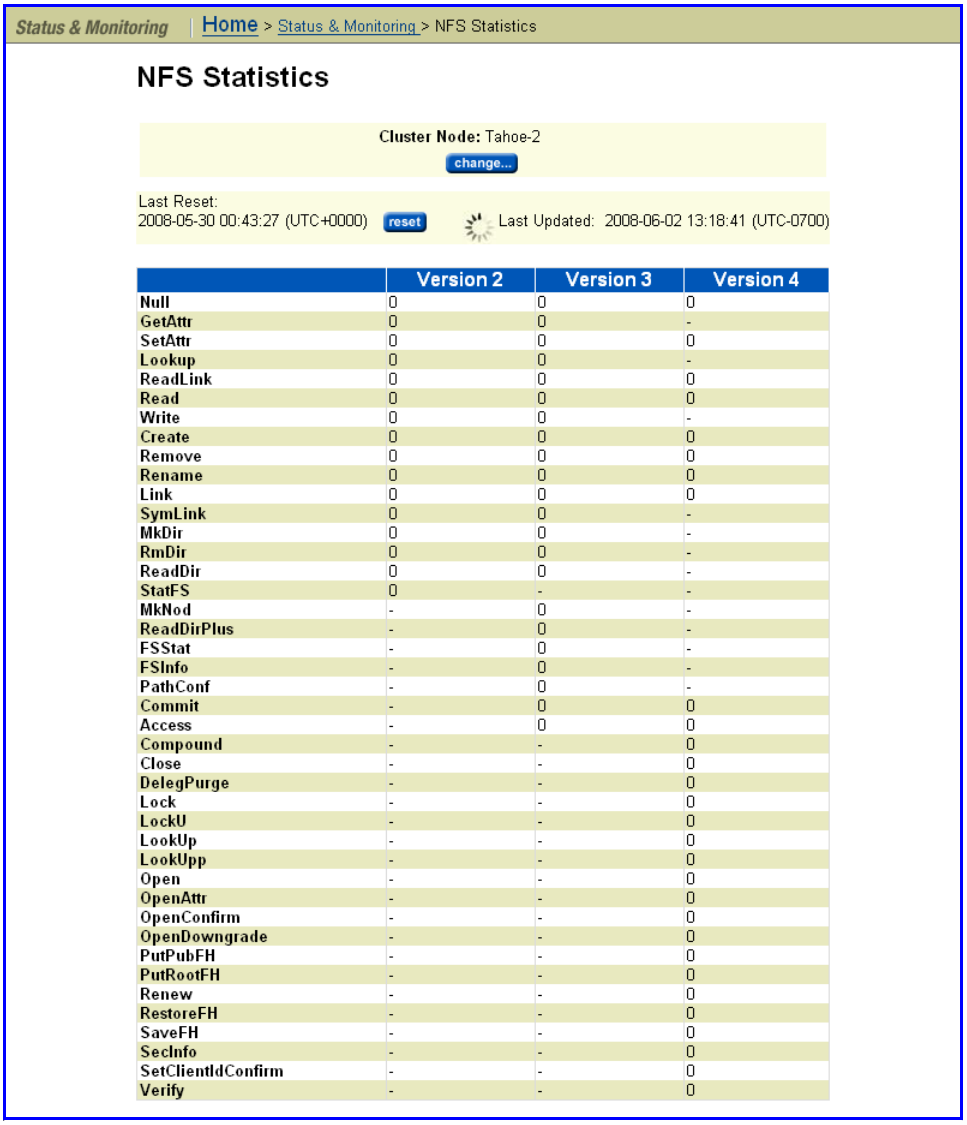

	Version 2	Version 3	Version 4
Null	0	0	0
GetAttr	0	0	-
SetAttr	0	0	0
Lookup	0	0	-
ReadLink	0	0	0
Read	0	0	0
Write	0	0	-
Create	0	0	0
Remove	0	0	0
Rename	0	0	0
Link	0	0	0
SymLink	0	0	-
MkDir	0	0	-
RmDir	0	0	-
ReadDir	0	0	-
StatFS	0	-	-
MkNod	-	0	-
ReadDirPlus	-	0	-
FSStat	-	0	-
FSInfo	-	0	-
PathConf	-	0	-
Commit	-	0	0
Access	-	0	0
Compound	-	-	0
Close	-	-	0
DelegPurge	-	-	0
Lock	-	-	0
LockU	-	-	0
LookUp	-	-	0
LookUpp	-	-	0
Open	-	-	0
OpenAttr	-	-	0
OpenConfirm	-	-	0
OpenDowngrade	-	-	0
PutPubFH	-	-	0
PutRootFH	-	-	0
Renew	-	-	0
RestoreFH	-	-	0
SaveFH	-	-	0
SecInfo	-	-	0
SetClientIdConfirm	-	-	0
Verify	-	-	0

The following table describes the components of this page:

Request	Description
Access	Gets the file security accesses for a file.
Close	Closes a file.
Commit	Commits the cached data on the server to stable storage.
Compound	Compound operations.
Create	Creates a file or symbolic link.
DelegPurge	Purge delegations awaiting recovery.
FSInfo	Gets static file system state information.

Request	Description
FSStat	Gets dynamic file system state information.
GetAttr	Retrieves the attributes of a file or directory.
Link	Creates a hard link to an object.
Lock	Creates a lock.
LockU	Unlocks a file.
Lookup	Looks up a file name in a directory.
LookUpp	Looks up a parent directory.
MkDir	Creates a directory.
MkNod	Creates a special device node (device file or named pipe).
Null	Does nothing, except to make sure the connection is up.
Open	Opens a regular file.
OpenAttr	Opens the named attribute directory.
OpenConfirm	Confirms open.
OpenDowngrade	Reduces open file access.
PathConf	Retrieves POSIX information for the file system.
PutPubFH	Sets public filehandle.
PutRootFH	Sets root filehandle.
Read	Reads data from a file.
ReadDir	Reads from a directory.
ReadDirPlus	Performs an expanded read from a directory.
ReadLink	Reads the data associated with a symbolic link.
Remove	Removes a file.
Rename	Renames a file or directory.
Renew	Renews a lease.
RestoreFH	Restores saved filehandle.
RmDir	Removes a directory.
SaveFH	Saves current filehandle.
SecInfo	Obtains available security.
SetAttr	Sets the attributes of a file or directory.

Request	Description
SetClientIdConfirm	Confirms client ID.
StatFS	Gets dynamic file system state information.
SymLink	Creates a symbolic link.
Verify	Verifies same attributes.
Write	Writes data to a file.

Viewing CIFS Statistics

CIFS statistics display SMB1 and SMB2 activity since the last server reboot or since CIFS statistics were last reset. They are updated every ten seconds. To view CIFS Statistics, navigate from the **Status & Monitoring** page to the **CIFS Statistics** page, which displays number of current clients and the number of CIFS calls that clients have sent to the server:

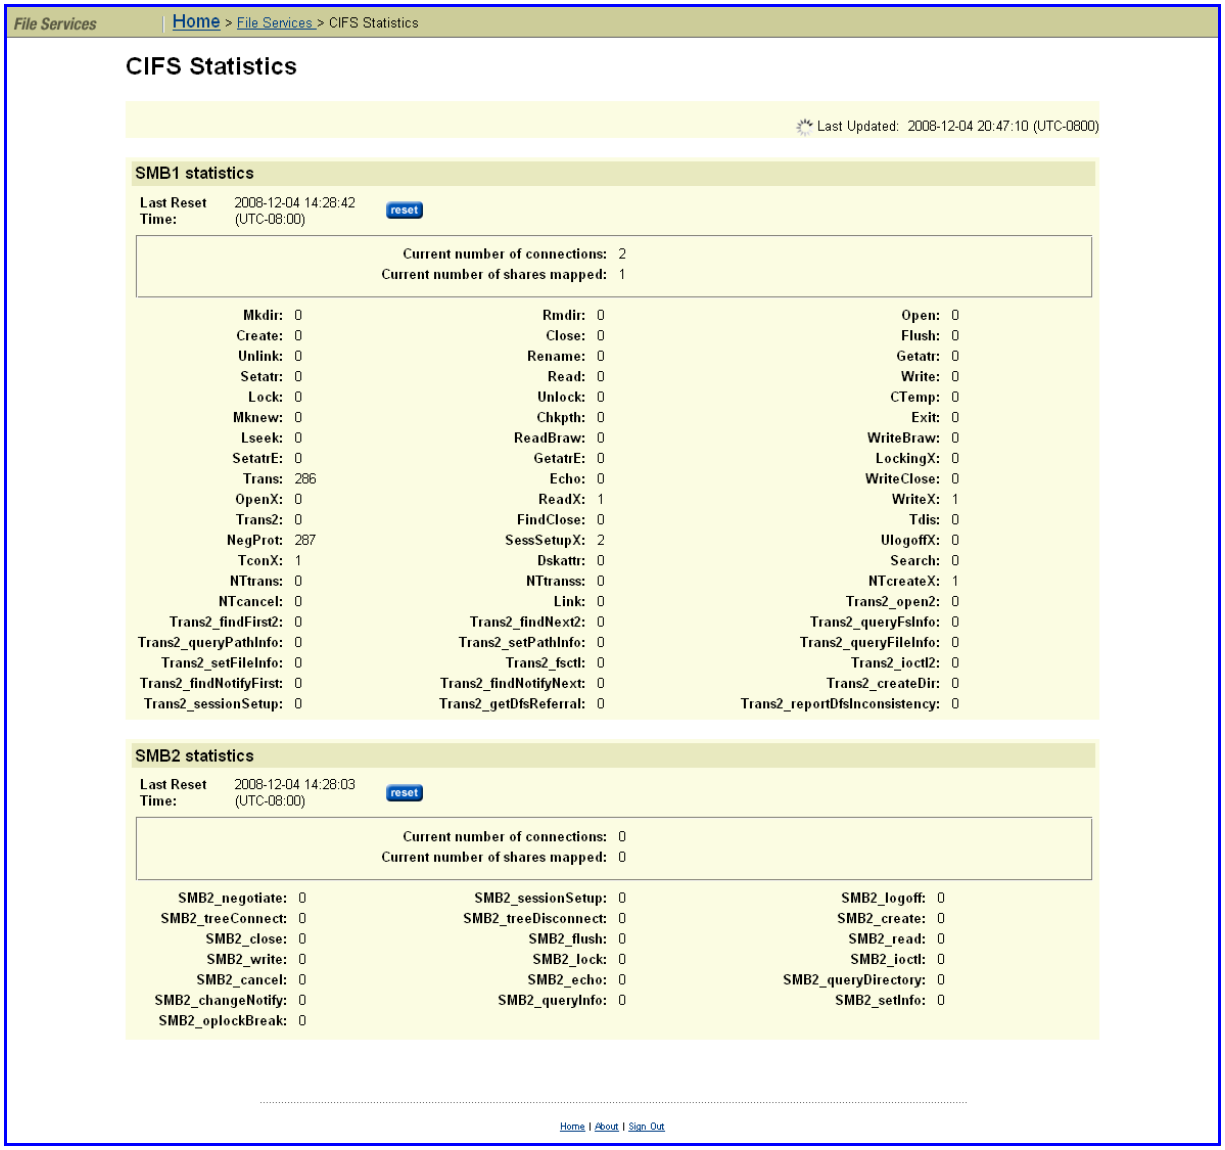

The following tables describe the components of the SMB1 section of this page:

Item/Field	Description
Last Updated	Displays the date, time, and UTC offset for when the statistics on this page were last updated.
Last Reset	Displays the date and time the statistics on this page were last reset.
Time	Displays the UTC offset of the date and time the statistics on this page were last reset.
reset	To reset the statistics to zero, click the **reset** button.

Item/Field	Description
Current number of connections	Displays the current number of SMB1 (CIFS v1) connections.
Current number of shares mapped	Displays the current number of CIFS shares being accessed through the current connections.

SMB1 calls:

Call	Description
Chkpth	Checks that the specified directory path exists.
Close	Closes a file.
Create	Creates a new file or opens an existing one.
CTemp	Creates a temporary file with a random server-generated name.
Dskattr	Retrieves file system attributes.
Echo	Pings the server.
Exit	Used by a process when it exits. Currently unsupported.
FindClose	Closes a CIFS FindFirst subfunction.
Flush	Instructs the server to flush cached information on a file.
Getattr	Retrieves the attributes of a file or directory.
GetatrE	Retrieves the expanded attributes of a file or directory.
Link	Creates a hard link to an object.
Lock	Takes out a byte-range lock on a file.
LockingX	Locks or unlocks a range of bytes in a file.
Lseek	Sets the file pointer to a given offset in the file.
Mkdir	Creates a new directory.
Mknew	Creates a new file.
NegProt	Negotiates the protocol with which the client and server will communicate.
NTcancel	Cancels an outstanding operation.
NTcreateX	Creates a new file or opens an existing one.
NTtrans	Multifunction command for operating subfunctions.
NTtranss	Multifunction command for operating subfunctions.

Call	Description
Open	Creates a new file or opens an existing one.
OpenX	Creates a new file or opens an existing one.
Read	Reads data from a file.
ReadBraw	Reads a block of data with no CIFS header.
ReadX	Reads data from a file.
Rename	Renames a file or directory.
Rmdir	Removes a directory.
Search	Lists the files in a directory.
SessSetupX	Logs the client in to a CIFS session.
Setattr	Sets the attributes of a file or directory.
SetattrE	Sets the attributes of a file or directory.
TconX	Connects the client to a file system resource.
Tdis	Breaks a connection that a TconX call previously established.
Trans	Multifunction command for operating subfunctions.
Trans2	Multifunction command for operating subfunctions.
Trans2_creatDir	Create a directory that has expanded attributes.
Trans2_findFirst2	Begin a search for files.
Trans2_findNext2	Resume a search for files.
Trans2_findNotifyFirst	Commence monitoring changes on a file or directory.
Trans2_findNotifyNext	Continue monitoring changes on a file or directory.
Trans2_fsctl	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
Trans2_GetDfsReferral	Get a DFS referral.
Trans2_ioctl2	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
Trans2_open2	Create a file that has expanded attributes.
Trans2_queryFileInfo	Get information about a file handle.
Trans2_queryFsInfo	Get information about a file system.
Trans2_queryPathInfo	Get information about the named file or directory.

Call	Description
Trans2_reportDfsInconsistency	Report an inconsistency in DFS knowledge.
Trans2_sessionSetup	Set up a session with expanded security.
Trans2_setFileInfo	Set file information by handle.
Trans2_setPathInfo	Set information about a named file or directory.
UlogoffX	Breaks a connection that a SessSetupX call previously established.
Unlink	Deletes a file.
Unlock	Releases a byte-range lock on a file.
Write	Writes data to a file.
WriteBraw	Write a block of data with no CIFS header.
WriteClose	Writes data to a file and then closes the file.
WriteX	Writes data to a file.

Click **reset** to set all values on this page to zero.

The following tables describe the components of the SMB2 section of this page:

Item/Field	Description
Last Reset	Displays the date and time the statistics on this page were last reset.
Time	Displays the UTC offset of the date and time the statistics on this page were last reset.
reset	To reset the statistics to zero, click the **reset** button.
Current number of connections	Displays the current number of SMB2 (CIFS v2) connections.
Current number of shares mapped	Displays the current number of CIFS shares being accessed through the current connections.

The following table describes the components of the SMB2 section of this page:

Call	Description
SMB2_cancel	Cancels a previously sent message on the same SMB2 transport connection.

Call	Description
SMB2_changeNotify	Change notifications on a directory.
SMB2_close	Close the named resource (pipe or file).
SMB2_create	Either create a file or access an existing file.
SMB2_echo	Determine if a server is processing requests.
SMB2_flush	Flush all cached file information for a specified open of a file to the persistent store that backs the file.
SMB2_ioctl	Issue an implementation-specific file system control or device control (FSCTL/IOCTL) command across the network.
SMB2_lock	Lock or unlock portions of a file.
SMB2_logoff	Terminate the named session.
SMB2_negotiate	Notify the server what dialects of the SMB 2.0 Protocol the client can process.
SMB2_oplockBreak	Server notification that the underlying object store indicates that an oplock is being broken, meaning that there is (or will be) a change in the oplock level.
SMB2_queryDirectory	Get a directory enumeration on an open directory.
SMB2_queryInfo	A request for information on a file, named pipe, or underlying volume.
SMB2_read	Request for a read operation on a specified file.
SMB2_sessionSetup	Request for a new authenticated session within a new or existing SMB 2.0 Protocol transport connection to the server.
SMB2_setInfo	Set information on a file or underlying file system.
SMB2_treeConnect	Request to access to a particular share on the server.
SMB2_treeDisconnect	Request to terminate the access to the specified tree.
SMB2_write	Write data to the file or named pipe on the server.

Click **reset** to set all values on this page to zero.

FTP Statistics

FTP statistics display activity since the last server reboot or since FTP statistics were last reset. They are updated every ten seconds. You can reset the values displayed on this page to zero by clicking **reset**. All values are reset except the ones concerning active sessions; for example, the number of active sessions, number of files incoming/outgoing for active sessions and number of bytes incoming/outgoing for active sessions.

To view FTP Statistics, navigate from the **Status & Monitoring** page to the **FTP Statistics** page:

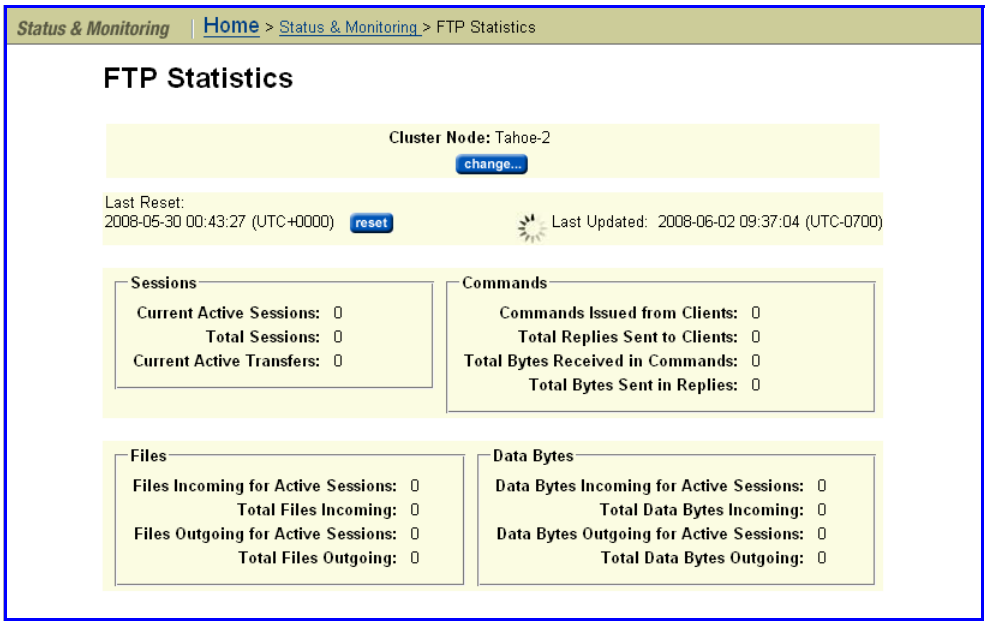

The following table describes the components of this page:

Item/Field	Description
Sessions	
Current Active Sessions	Currently active FTP sessions.
Total Sessions	Total FTP sessions since last server restart or statistics reset.
Current Active Transfers	Currently active FTP transfers.
Commands	
Commands Issued from Clients	Number of commands sent by clients.
Total Replies Sent to Clients	Number of replies sent to clients.
Total Bytes Received in Commands	Bytes in commands that clients have sent to the FTP server.
Total Bytes Sent in Replies	Bytes in replies that the FTP server has sent to clients.
Files	
Files Incoming for Active Sessions	Files that clients have transferred to the FTP server in currently active sessions.
Total Files Incoming	Files that clients have transferred to the FTP server since last server restart or statistics reset.

Item/Field	Description
File Outgoing for Active Sessions	Files that the FTP server has transferred to clients in currently active sessions.
Total Files Outgoing	Files that the FTP server has transferred to clients since last server restart or statistics reset.
Data Bytes	
Data Bytes Incoming for Active Sessions	Bytes of data that clients have transferred to the FTP server in currently active sessions.
Total Data Bytes Incoming	Bytes of data that clients have transferred to the server since last server restart or statistics reset.
Data Bytes Outgoing for Active Sessions	Bytes of data that the FTP server has transferred to clients in currently active sessions.
Total Data Bytes Outgoing	Bytes of data that the server has transferred to clients since last server restart or statistics reset.

iSCSI Statistics

The **iSCSI Statistics** page provides a summary of the iSCSI and SCSI requests on a NAS server/cluster node.

This page displays the current number of connections, the current number of sessions, and the number of iSCSI requests that initiators have sent to NAS server/cluster node. For a cluster node, the node is shown in the **Cluster Node** field and you can change nodes by clicking the **change** button. These statistics are updated every ten seconds. You can reset all the values displayed on this page to zero by clicking **reset**.

To view iSCSI statistics, navigate from the **File Services** page to the **iSCSI Statistics** page:

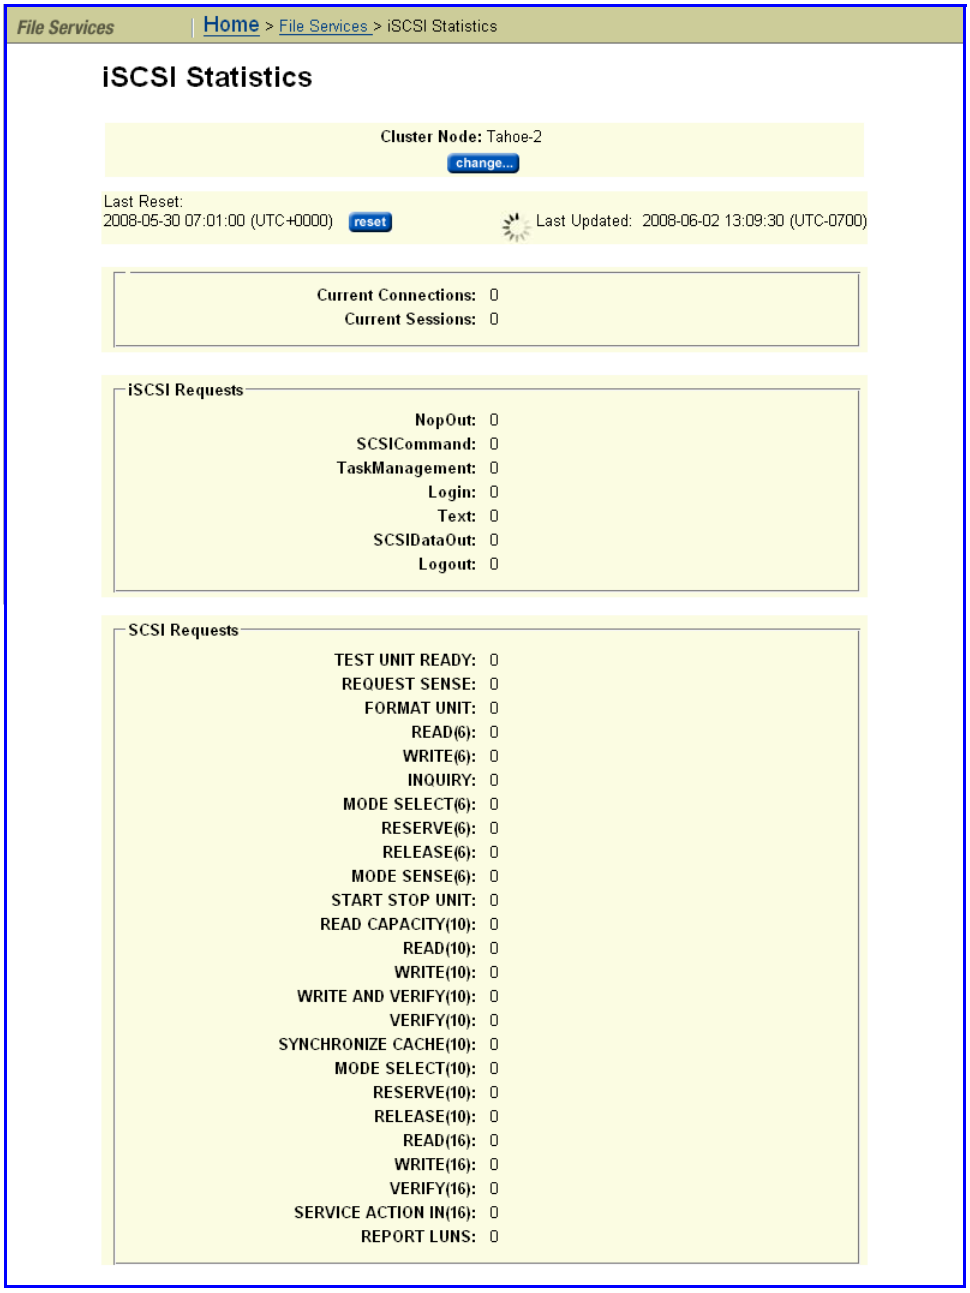

The following table describes the contents of this page:

Item / Field	Description
Current Connections	The current number of iSCSI connections to the server.
Current Number of Session	The number of iSCSI sessions currently hosted by the server.
iSCSI Requests	

Item / Field	Description
NopOut	No operation.
Task Management	Requests used for task management functions.
Text	Requests used to negotiate behavior.
Logout	Logout requests
SCSICommand	Carries a SCSI Command.
Login	Login requests.
SCSIDataOut	Requests containing SCSI data.
iSCSI Requests	
TestUnitReady	Tests that the target is ready to receive commands.
Read(6)	Reads data.
ModeSelect(6)	Configure SCSI behavior.
Release(6)	Releases (unlocks) a Logical Unit reservation.
StartStopUnit	Warm reboots the target.
Read(10)	Reads data.
Verify(10)	Verifies data.
ModeSelect(10)	Configure SCSI behavior.
Release(10)	Releases (unlocks) a Logical Unit reservation.
RequestSense	Requests state information.
Inquiry	Requests device information.
Reserve(6)	Reserves (locks) a Logical Unit for exclusive access.
ModeSense(6)	Requests SCSI configuration information.
ReadCapacity	Reads the size of Logical Unit.
Write(10)	Write data.
SynchronizeCache	Flushes cached data to disk.
Reserve(10)	Reserves (locks) a Logical Unit for exclusive access.
ReportLuns	Retrieves a list of available Logical Units.
Format Unit	Formats a logical unit.
Write(6)	Writes data.
WriteAndVerify(10)	Writes then verifies data.

Item / Field	Description
Read(16)	Reads data.
Write(16)	Writes data.
Verify(16)	Verifies data.
ServiceActionIn(16)	Performs an extended SCSI command, such as ReadCapacity(16).

Data Access and Performance Statistics

The server provides measures and tools for monitoring the impact of network clients on internal resources. In particular, the server provides:

- Server and file system load statistics
- File System NVRAM usage statistics

Server and File System Load (Ops per second)

In addition to Ethernet and Fibre Channel throughput statistics, server performance can also be measured in operations per second (ops/sec). The Web Manager provides a graphic representation of ops/sec, at two levels:

- Total operations per server.
- Total operations per individual file system.

The *total operations on a server* is an aggregate of the operations performed by all file systems hosted by that server.

Understanding the performance profile of servers and individual file systems is especially useful in environments where more than one server is installed, as it enables intelligent relocation of EVSs or file systems to more equally distribute the load among the available servers.

Viewing Ops/Sec Statistics

To view ops/sec statistics, navigate from the **Status & Monitoring** page to the **File System Ops/Sec** (or **Node Ops/Sec**) page:

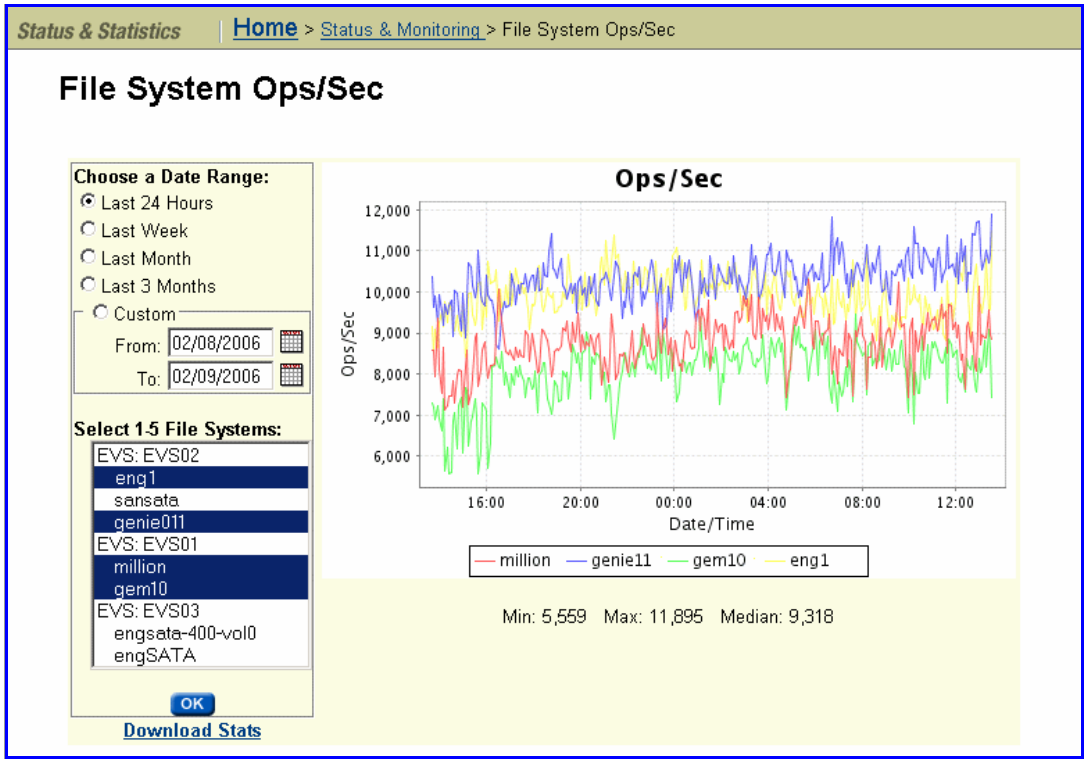

The following functionality is provided:

- *If File System Ops/sec was selected,* select between one and five file systems to view under **Select 1-5 File Systems.**Hold **Ctrl** while clicking to select more than one file system.
- Statistics can be viewed based on a specified date range. Customize the date range by selecting an option under **Choose a Date Range**.
- The statistics can be downloaded into `.csv` format by clicking the **Download Stats** link.

File System NVRAM Statistics

The **File System NVRAM Statistics** page displays NVRAM activity. To display File System NVRAM Statistics, navigate from the **Status & Monitoring** page to the **File System NVRAM Statistics** page.

Note: When an EVS has a Read Cache file system, no NVRAM statistics are presented.

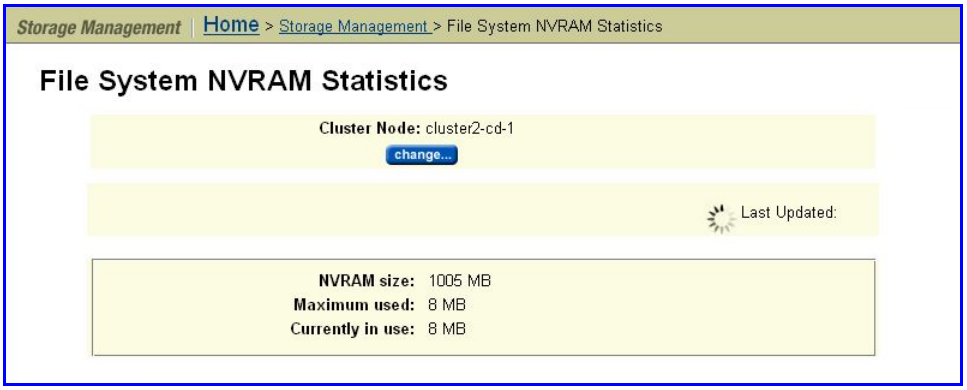

The following table describes the fields in this page:

Item/Field	Description
Cluster Node	When connected to a cluster, this field indicates the node for which NVRAM statistics are displayed. To view statistics for another node, click the **change** button.
change	Click the **change** button to open the **Select a Cluster Node** page where you can select a different node for which to view statistics.
Last Updated	Displays the date and time this page was refreshed.
NVRAM size	Size of NVRAM buffer, used to preserve data for disk-modifying operations until written to disk. The default is 2GB.
Maximum used	Maximum amount of the NVRAM buffer that has been used since the node was last started.
Currently in use	Current NVRAM buffer usage.

Management Statistics

A IS-NAS Server/cluster provides the following management statistics:

- Access Management Statistics for SSC, SNMP, HTTPS, and VSS
- Virus Scanning Statistics

A Titan Server/cluster also provides the following management statistics:

- Access Management Statistics for Telnet, SSC, SSH, SNMP, HTTPS, and VSS
- Virus Scanning Statistics

Access Management Statistics

Management statistics are displayed since the server was last started or when it was last reset. The system updates the management statistics every ten seconds.

To view Access Management Statistics, navigate from the **Status & Monitoring** page, then select one of the items in the **Management Access Statistics** section to display its **Management Statistics** page:

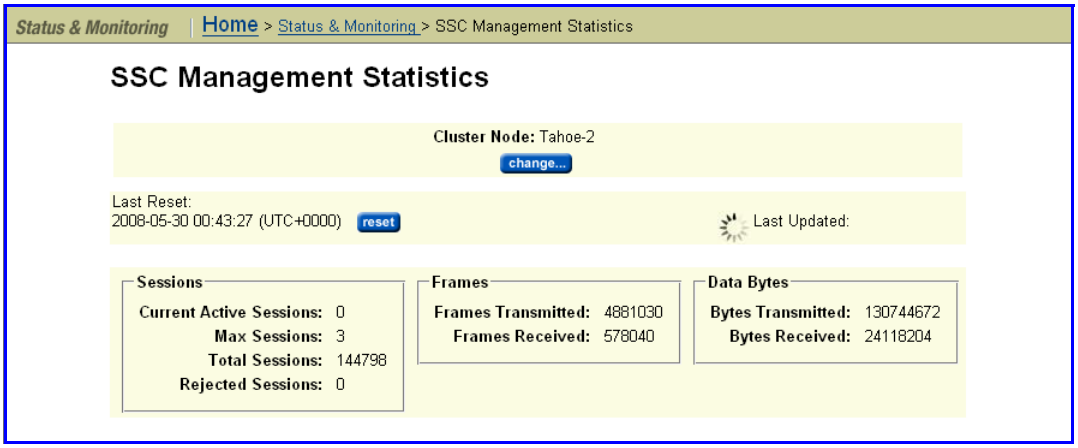

The following table describes the components of this page:

This statistic	Shows
Sessions	
Current Active Sessions	Number of sessions that are currently in progress.
Max Sessions	Peak number of concurrent sessions.
Total Sessions	Total number of sessions.
Rejected Sessions	Number of failed attempts to establish a connection. A connection may fail because the client does not have required permissions or because the maximum number of concurrent sessions is already in progress.
Frames	
Frames Transmitted	Number of sections (packages of information transmitted as single units) that the system has sent to clients.
Frames Received	Number of sections that clients have sent to the system.
Data Bytes	
Bytes Transmitted	Number of data bytes that the system has sent to clients.
Bytes Received	Number of data bytes that clients have sent to the system.

SNMP Management Statistics

The SNMP management statistics page shows the SNMP statistics for the server since the server was last reset. It displays statistics regarding Input, Output, and Drops. These statistics are updated every ten seconds. You can reset all the values displayed on this page to zero by clicking **reset**.

To view SNMP Management Statistics, navigate from the **Status & Monitoring** page to the **SNMP Management Statistics** page. When the server is part of a cluster, the **Cluster Node** field identifies the node, and the **change** button allows you to change nodes, and view statistics for that node.

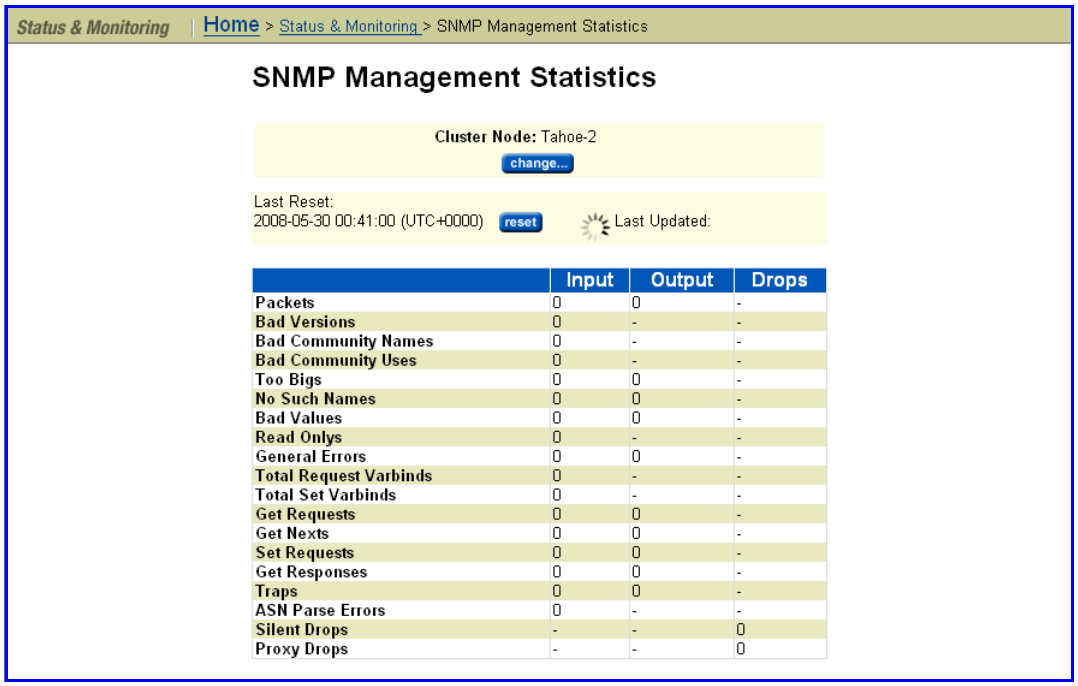

Item/Field	Shows the number of
Input	
Packets	SNMP packets the agent has received.
Bad Community Names	SNMP messages received using an unknown community name.
Too Bigs	Protocol Data Units (PDUs) received containing an error-status field value of tooBig.
Bad Values	PDUs received containing an error-status field value of badValue.
General Errors	PDUs received containing an error-status field value of genErr.
Total Set Varbinds	MIB objects successfully altered because of valid SNMP Set-Request PDUs.
Get Nexts	Get-Next PDUs received and processed.
Get Responses	Get-Response PDUs received and processed.
ASN Parse Errors	Abstract Syntax Notation (ASN) errors found in SNMP messages received.
Bad Versions	Packets received for an unsupported SNMP version.

Item/Field	Shows the number of
Bad Community Uses	SNMP messages received representing an operation not allowed by the SNMP community named in the message.
No Such Names	PDUs received containing an error-status field value of nosuchName.
Read Onlys	PDUs received containing an error-status field value of ReadOnly. This value is used to detect incorrect SNMP implementations.
Total Request Varbinds	MIB objects successfully retrieved because of valid SNMP Get-Request and Get-Next PDUs.
Get Requests	Get-Request PDUs received and processed.
Set Requests	Set-Request PDUs received and processed.
Traps	Trap PDUs received and processed.
Output	
Packets	SNMP packets the agent has sent.
No Such Names	Sent PDUs receiving an error-status field value of noSuchName.
General Errors	Sent PDUs receiving an error-status field value of genErr.
Get Nexts	Get-Next PDUs sent.
Get Responses	Get-Response PDUs sent.
Too Bigs	Sent PDUs receiving an error-status field value of tooBig.
Bad Values	Sent PDUs receiving an error-status field value of badValue.
Get Request	Get-Request PDUs sent.
Set Requests	Set-Request PDUs sent.
Traps	Trap PDUs sent.
Drops	
Silent Drops	PDUs delivered but silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests.
Proxy Drops	PDUs delivered but silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned.

HTTPS Management Statistics

The HTTPS management statistics page shows the HTTPS statistics for the server since the server was last reset. It displays statistics regarding sessions, data sent/received, and connections. These statistics are updated every ten

seconds. You can reset all the values displayed on this page to zero by clicking **reset**.

To view HTTPS Management Statistics, navigate from the **Status & Monitoring** page to the **HTTPS Management Statistics** page. When the server is part of a cluster, the **Cluster Node** field identifies the node, and the **change** button allows you to change nodes, and view statistics for that node.

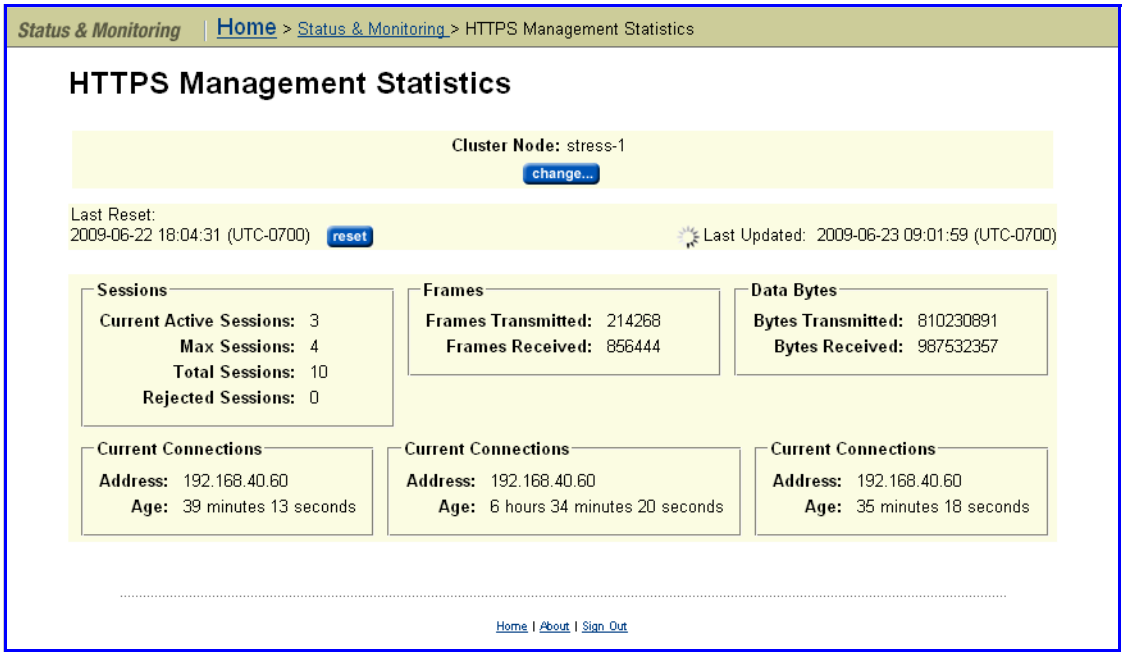

Item/Field	Shows the number of
Sessions	
Current Active Sessions	The number of HTTPS sessions that are currently in progress.
Max Sessions	The peak number of concurrent HTTPS sessions.
Total Sessions	The total number of HTTPS sessions.
Rejected Sessions	The total number of failed attempts to establish an HTTPS connection. A connection may fail because the client does not have the required permissions or because the maximum number of concurrent sessions are already in progress.
Frames	
Frames Transmitted	The total number of frames that the system has sent to clients over an HTTPS connection.
Frames Received	The total number of frames that clients have sent to the system over an HTTPS connection.
Data Bytes	

Item/Field	Shows the number of
Bytes Transmitted	The number of data bytes that the system has sent to clients over an HTTPS connection.
Bytes Received	The number of data bytes that clients have sent to the system over an HTTPS connection.
Current Connections (for each active connection)	
Address	IP address of the connected client.
Age	The duration of the connection.

To reset all the values to zero, click **reset**.

VSS Management Statistics

The VSS management statistics page shows the VSS statistics for the server since the server was last reset. It displays statistics regarding sessions, frames sent/received, and data sent/received. These statistics are updated every ten seconds. You can reset all the values displayed on this page to zero by clicking **reset**.

To view HTTPS Management Statistics, navigate from the **Status & Monitoring** page to the **VSS Management Statistics** page. When the server is part of a cluster, the **Cluster Node** field identifies the node, and the **change** button allows you to change nodes, and view statistics for that node.

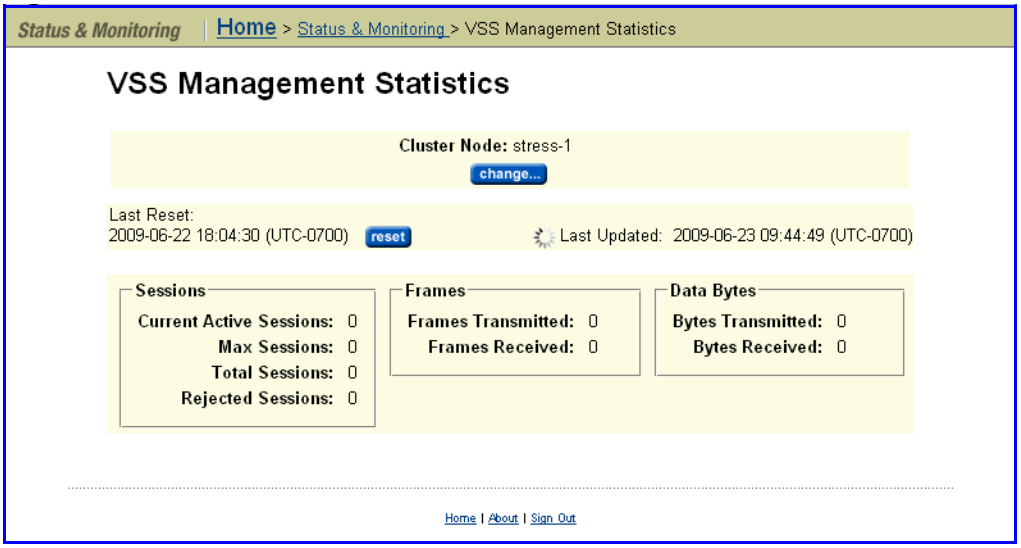

Item/Field	Shows the number of
Sessions	
Current Active Sessions	The number of HTTPS sessions that are currently in progress.

Item/Field	Shows the number of
Max Sessions	The peak number of concurrent HTTPS sessions.
Total Sessions	The total number of HTTPS sessions.
Rejected Sessions	The total number of failed attempts to establish an HTTPS connection. A connection may fail because the client does not have the required permissions or because the maximum number of concurrent sessions are already in progress.
Frames	
Frames Transmitted	The total number of frames that the system has sent to clients over an HTTPS connection.
Frames Received	The total number of frames that clients have sent to the system over an HTTPS connection.
Data Bytes	
Bytes Transmitted	The number of data bytes that the system has sent to clients over an HTTPS connection.
Bytes Received	The number of data bytes that clients have sent to the system over an HTTPS connection.

To reset all the values to zero, click **reset**.

Virus Scanning Statistics

The **Virus Statistics** page summarizes virus scanning activity. To display the Virus Statistics, navigate from the **Data Protection** page to the **Virus Statistics** page:

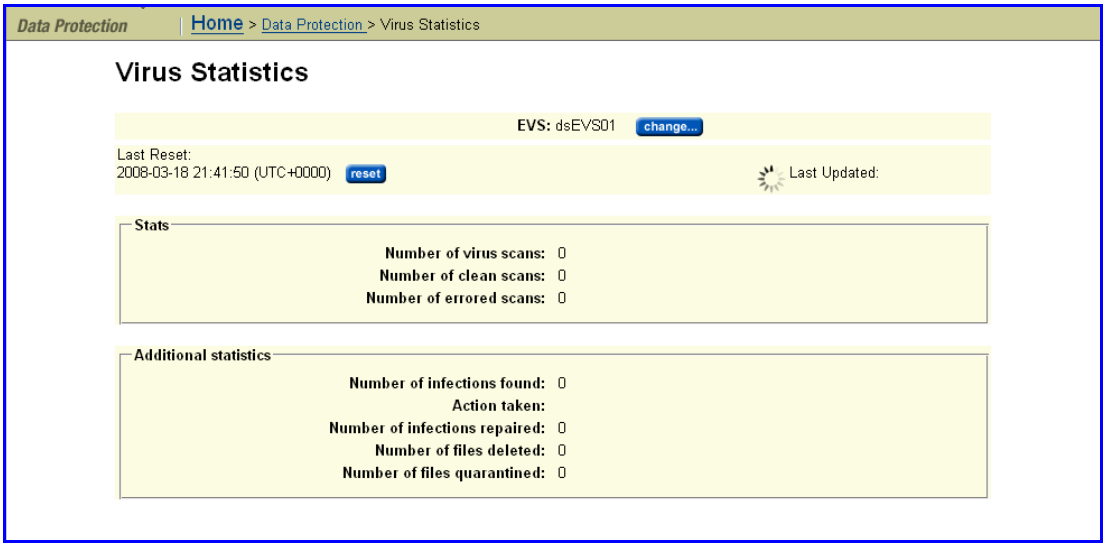

Note: When a virus is detected, a severe event is placed in the Event Log, identifying the path of the infected file and the IP address of the client machine

that wrote the file.

Note: Files will only be deleted, repaired or quarantined if the virus scan engine has been configured to do so.

If the server is configured as a cluster, the node is displayed in the **Cluster Node** field. You can click the **change** button to select a different node for which the statistic are displayed.

To change the EVS for which the statistic are displayed, click the **change** button next to the EVS field. The **Last Reset** field shows the last time the statistics on this page were reset to zero. Click the **reset** button to reset all values to zero. The **Last Updated** field displays the date and time the statistics on this page were last updated.

The following table describes the components of this page:

Item	Description
Number of virus scans	Number of times files have been scanned for viruses.
Number of clean scans	Number of times files have been scanned with no viruses detected.
Number of errored scans	Number of times a failure occurred while scanning a file.
Additional statistics (not supported on every Virus Scan Engine):	
Number of infections found	Number of times files have been scanned and detected infections are found.
Number of infections repaired	Number of times the Virus Scan Engine has been able to repair infections found.
Number of files deleted	Number of deleted files because they contain irreparable infections.
Number of files quarantined	Number of files quarantined because they contain irreparable infections.

Event Logging and Notification

The server provides a comprehensive event logging and alert mechanism and auxiliary devices in the storage subsystem automatically direct any events and SNMP traps to the server (or can be configured to do so).

All event messages generated by the server (including those issued by its auxiliary devices) are logged into an event log, which can be downloaded and cleared by the system administrator. The event log provides a record of past events that have occurred on the server, for use in trend/fault analysis.

Event message severity can be changed, and messages can be suppressed entirely, using the CLI command `event-log-filter`. Using the CLI

command `event-log-filter`, you can specify that a command is to be run whenever a specified message is logged. For more information on the `event-log-filter` command, refer to the *Command Line Reference*.

The server can also be configured for automated notification according to pre-defined severity categories, including *daily summary* and *status notification.* With automated notification enabled, the system will notify selected personnel when an event is generated, based on the level of severity of the event. 24x7 automated notifications allow SGI Global Services personnel to proactively monitor the health of the system and address any issues that may arise.

Using the Event Log

The server continuously monitors temperature, fans, power supply units, and disk drives in the cabinet. Each time an event occurs (for example, a disk failure or a possible breach of security, the system records it in an event log). The event log can be viewed, filtered, and saved as a permanent record.

The log can contain a maximum of 10,000 events. Once the event log limit has been reached, each new event replaces the oldest event in the log.

Viewing and Filtering the Event Log

1. **Navigate to the Event Log Management page.**

   From the **Status & Monitoring** page, click **Event Log** to display the **Event Log Management** page:

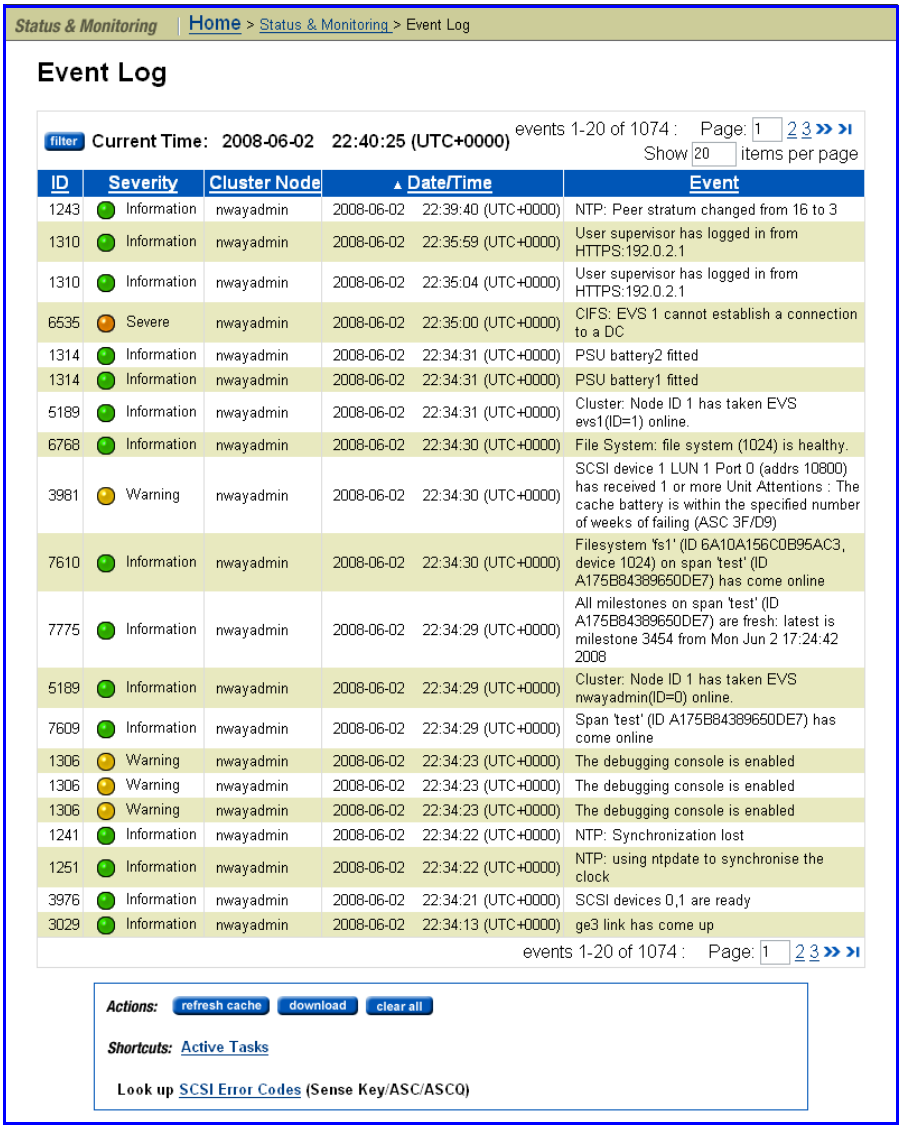

2. **Set up filtering.**

   Click the **filter** button to open the **Filter** dialog.

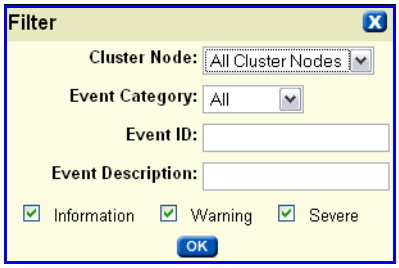

   Specify your filtering criteria:

   - In a cluster, specify the cluster node for which to display the log. In the **Cluster Node** field, you can select the specific node or *All Cluster Nodes*.

- In the **Event Category** field, select the type of events to be included in the log: *All* events, *System* events, or *Security* events.

  System events are events that the system components have logged, such as the failure of a drive. Security events track changes to the security system and identify possible breaches of security.
- You can specify an **Event ID** that you want included in the log.
- You can specify an **Event Description** that you want included in the log.
- Select the severity level of the events you want included in the log by filling one or more of the boxes: **Information**, **Warning**, or **Severe**.

Click **OK** to filter the log events being displayed according to the filter criteria you specified.

3. **As needed, view event details**

   Click an event to display a dialog box displaying the cause and resolution:

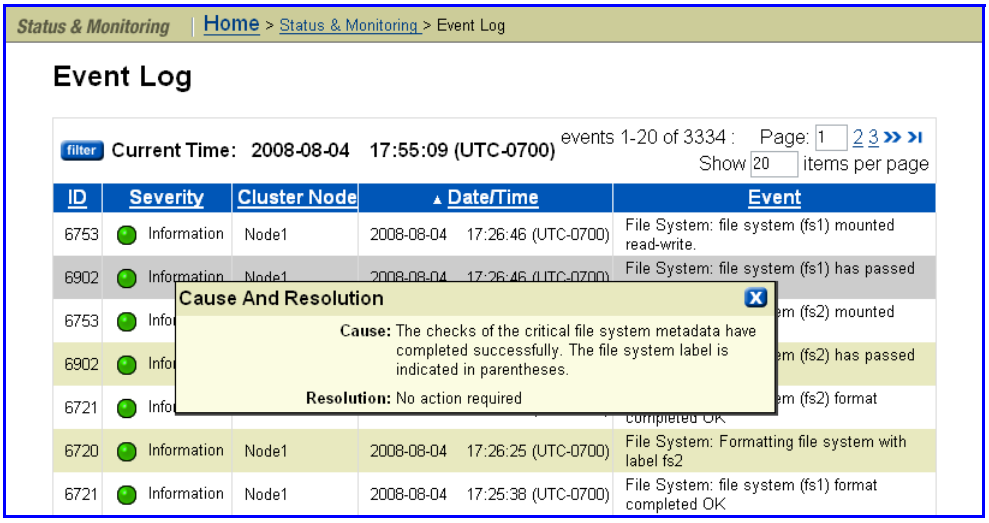

4. **As needed, refresh the cache.**

   Click **refresh cache** to clear the SMU's cache, and then repopulate the cache with the relevant objects. Note that this is different than clicking the browser refresh button, which picks up any recent updates without clearing the cache.

5. **As needed, download or empty the log.**

   To download the log to your computer, click **Download Log,** then print or save to a text file.

   To empty the log, click **Clear Event Log.**

Setting up Event Notification

The server can be configured for automatic notification of selected users when particular types of system events occur. Once warned of an event, these users

can run Web Manager to diagnose the problem remotely, with a direct connection or virtual private link to the network.

The event notification may take three forms:

- An **Email** message, which the system sends through an SMTP server. See Managing Email Alerts and Profiles, on page 507 for more information.
- An **SNMP trap**, to notify a central Network Management Station (NMS) of any events generated by the server; for example, HP OpenView. See Sending SNMP Traps, on page 511 for more information.
- A **Syslog** alert enables the user to send alerts from a server to a UNIX system log (the UNIX system must have its syslog daemon configured to receive remote syslog messages). See Setting Up Syslog Notification, on page 513 for more information.

Note: With any of the event notifications, SGI Global Services recommends setting a notification frequency of *Immediately* for the most serious alert type (*Severe*) and to send these alerts to at least two users.

Using Email Alerts

The server can be configured to send emails to specified recipients to alert them on system events. Setting up email alerts requires configuring:

- **SMTP Servers.** The servers on the network to which the reporting server should email alerts.
- **Email Profiles.** Email profiles allow distribution groups to be created, so that email recipients are properly notified based on alert threshold criteria.

  The server allows classification of email recipients into specific profiles, so that they can receive customized alerts with the depth of focus they require.

  For instance, profiles can define different tiers of user responsibility for the server, such that recipients in one profile will only receive alerts on *Severe* events, while recipients in a second profile receive alerts on *Warning and Severe* events, and recipients in a third profile get summary emails on *all events.* In a large user group, dividing these users into separate profiles saves time and simplifies event notification.

  Note: A special profile, SupportProfile, is intended for support use and sends an alert email to SGI Global Services. You cannot modify, add recipients to, or remove recipients from the SupportProfile. You can, however, disable this profile. If this profile was changed in an earlier version of the firmware, the only way to reset this profile to its default state is to navigate to the **SMTP Email Profile** page for the SupportProfile. To display this page, go to the **Email Alerts Setup** page (Home > Status & Monitoring > Email Alerts Setup), click **details** for the SupportProfile to display the **SMTP Email Profile** page for the SupportProfile, then click **Restore Default Support Profile**.

To configure email alerts:

1. **Navigate to the Email Alerts Setup page.**

   From the **Status & Monitoring** page, click to display the **Email Alerts Setup** page:

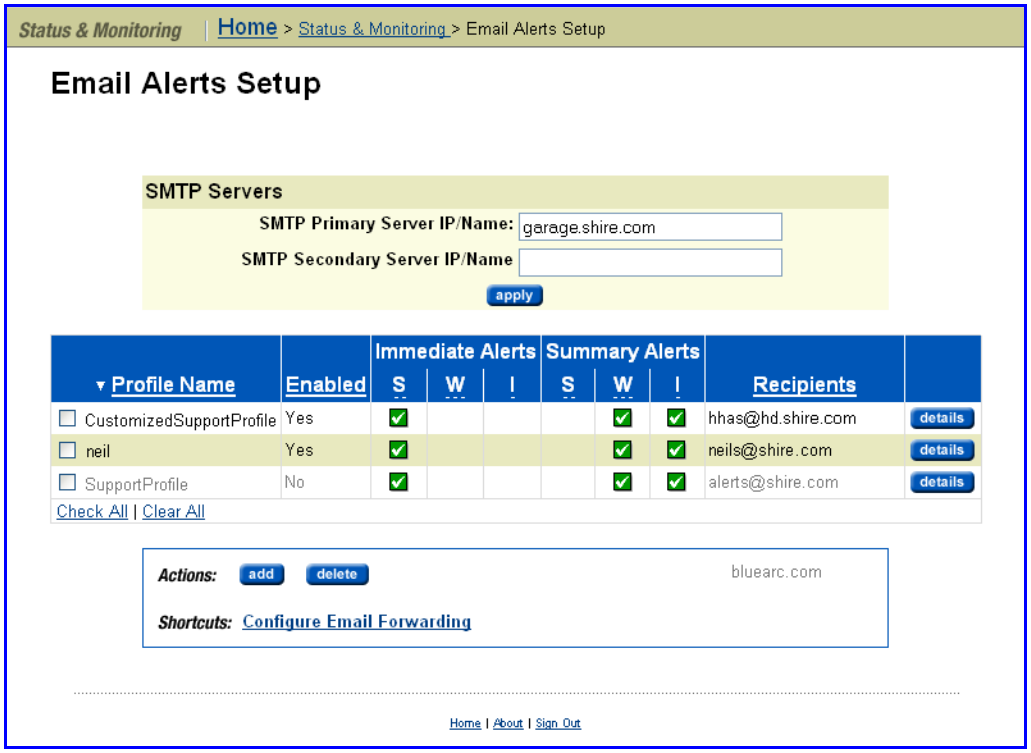

2. **Specify SMTP Server information.**

   The fields and items required to specify SMTP server information are described in the following table:

Field	Description
SMTP Primary Server IP/Name	Enter the host name or IP address of the primary mail server. The server specified as the SMTP Server will be used for email alert notification. If the Primary SMTP Server is offline, the Server will re-direct email notifications to the defined SMTP Secondary Server. **Tip:** As the server should always be in contact with the SMU, SGI Global Services recommends that the SMU's `eth1` IP address be defined as the Primary SMTP server. The SMU can be configured for email forwarding and relay any messages to the public mail server.
SMTP Secondary Server IP/Name	Enter the host name or IP address of the secondary mail server. Email alerts are redirected to this server if the Primary SMTP Server is unresponsive.
apply	Click **apply** to save the SMTP server information specified above.

3. **Manage existing email profiles.**

   The fields and items required to manage email profiles are described in the following table:

Field	Description
Profile Name	A descriptive name for the profile. A special profile, SupportProfile, is intended for support use and sends an alert email to technical support. You cannot modify, add recipients to, or remove recipients from, the SupportProfile.
Enabled	Indicates whether the profile is enabled.
Applies to SMU	Indicates whether the profile applies to the SMU. Note that this column only appears when an internal SMU is used (when using an external SMU, this column does not appear).
Immediate Alerts	A check mark in one of these columns indicates that an immediate email alert is to be sent to the recipients in the named profile when an event of the indicated type occurs. • **S** indicates that an alert will be sent for any severe event. • **W** indicates that an alert will be sent for any warning event. • **I** indicates that an alert will be sent for any informational events. It is recommended that immediate email alerts be sent on any severe event.
Summary Alerts	A check mark in one of these columns indicates that a summary email alert is to be sent to the recipients in the named profile when an event of the indicated type occurs. • **S** indicates that an alert will be sent for any severe event. • **W** indicates that an alert will be sent for any warning event. • **I** indicates that an alert will be sent for any informational events. It is recommended that summary email alerts be sent on any warning event.
Recipients	Displays the email addresses that will receive alert emails based on the profile definition.
details	Click **details** to display the **SMTP Email Profile** page, where you can enable, disable, or edit the email profile. See Managing Email Alerts and Profiles, on page 507 for more information.
add	Click **add** to display the **Add Email Profile** page, where you can create a new email profile. See Adding an Email Profile, on page 505 for more information.
delete	Fill the checkbox for the email profile you want to delete, and click **delete** to remove the selected email profile.

Field	Description
configure email forwarding	Click **configure email forwarding** to display the **SMTP Configuration** page, which allows you to specify the host name of the email server to which the SMU can send and relay event notification emails.

Daily Status Emails

A IS-NAS Server system/Titan Server system is made up of multiple components. To get an accurate description of the overall status of the various components of the storage system, two daily status emails are generated:

- **Daily Status Email from the server.** The server's Daily Status Email contains logs of server performance and battery health, descriptive information regarding the health of the server and storage subsystem, and current space utilization by the file systems.

  This email is sent to all recipients in all mail profiles in which the **Send a Daily Status Email at midnight** option has been selected.

- **Daily Status Email from the SMU.** The SMU's Daily Status Email contains a list of the SMU's managed servers and their current firmware versions. It also contains the SMU's current software version. The SMU and server names are links that can be clicked to manage the specified server. The email also provides the ID, Model, Type (for example, *single node* or c*luster node),* and Status information about servers.

- **SMU Diagnostic Emails.** The SMU sends all of its configured email recipients a diagnostic email when any of the following events occur:

  - The server has unexpectedly rebooted.
  - If enabled, once per day at a specified time.

  These diagnostic emails contain details regarding the servers, storage, and FC switches managed by the SMU. As the details in these diagnostic mails can be useful to Silicon Graphics International Corporation (should its assistance be required). SGI Global Services strongly recommends including `is-nas_call_home@sgi.com` as one of the email recipients.

  SGI also recommends enabling **Monthly Call Home Emails**. When enabled, the SMU sends a full set of server, SMU, and storage diagnostics to SGI on the first of every month. These provide an archive of the complete configuration of the storage system, which can aid in the detection of problems, provide a wealth of background diagnostic information to Silicon Graphics International Corporation should a problem occur and, if necessary, access to a known good configuration for restoration.

Note: When the monthly diagnostic email is first enabled, an initial email is sent at midnight that night, allowing you to verify that the email configuration is set up correctly.

For information on configuring the SMU's diagnostic emails, see Adding an Email Profile, on page 505.

Adding an Email Profile

To add an email profile:

1. **Navigate to the add Email Profile page.**

   From the **Status & Monitoring** page, select **Email Alert Configuration,** then click **add** to display the **Add Email Profile** page:

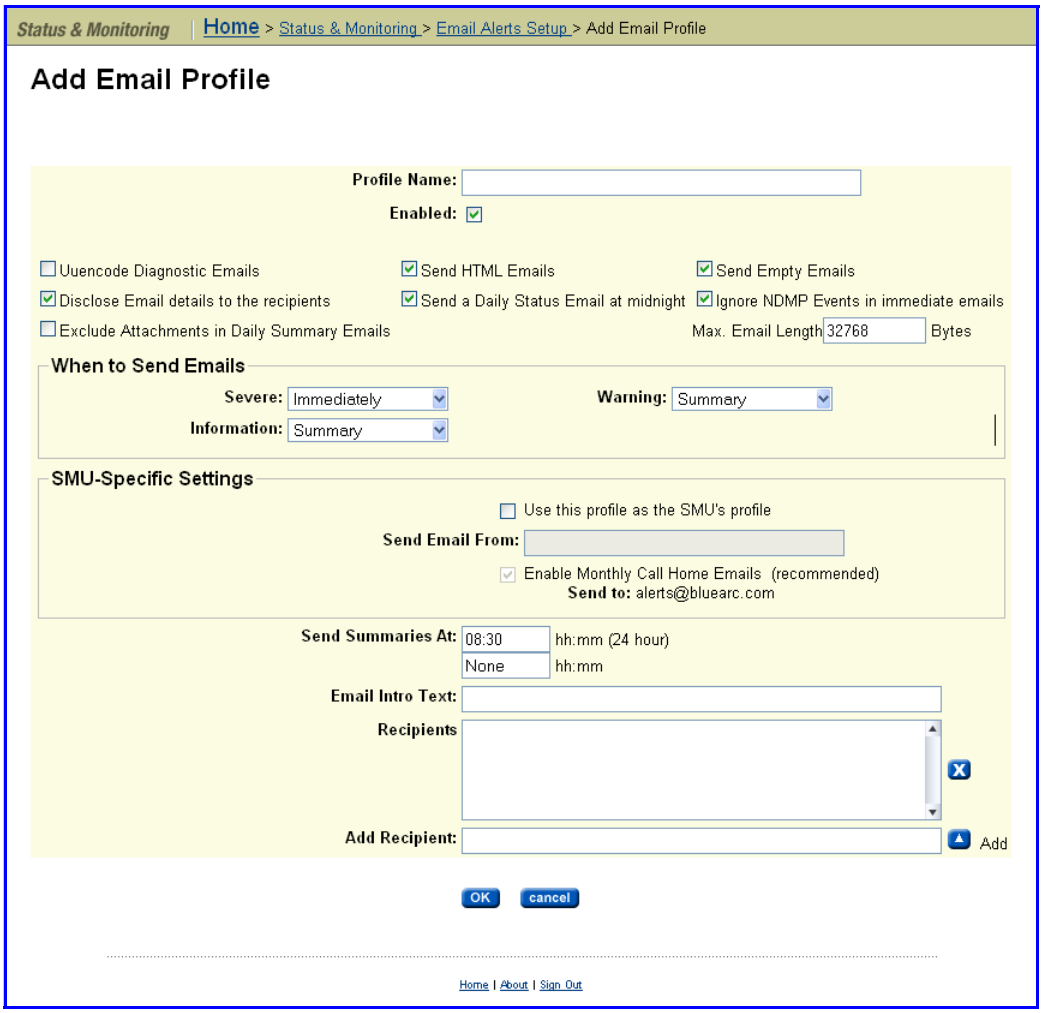

2. **Complete the requested information.**

   The fields on this page are described in the following table:

Field	Description
Profile Name	Select a name for the profile being created.
Enabled	Fill the checkbox to enable the profile, or leave it inactive.
Uuencode Diagnostic Emails	Fill this checkbox to uuencode email attachments. By uuencoding the message, any virus scanning software at the recipient's site will be bypassed.
Send HTML Emails	Fill this checkbox to receive emails in HTML format. HTML emails are easier to read compared to plain text mails, and this provides easy access to the web UI, since the server name in the email is clickable.
Send Empty Emails	By default, the Send Empty Emails checkbox will be filled. Empty summary emails will be sent to the specified recipient when this is selected. To avoid sending empty summarized Emails, clear the checkbox.
Disclose Email Details to the recipient	By default, the Disclose Email details to the recipient checkbox will be filled. Detailed emails containing restricted or confidential information (account names, IP addresses, portions of user data, etc.) will be sent to the specified recipient. To avoid sending detailed emails, clear the checkbox.
Send a Daily Status Email	By default, the Send a Daily Status Email checkbox will be filled. Detailed emails containing logs of server performance and battery health, descriptive information regarding the health of the server and storage subsystem, and the current space utilization of the file systems will be sent to the specified recipient. To avoid sending Daily Status Emails, clear the checkbox.
Ignore NDMP events in immediate emails	Fill this checkbox to prevent emails from being sent when events are generated by the NDMP backup system.
Exclude Attachments in Daily Summary Emails	Fill this checkbox to prevent attachments from being sent when daily summary emails are sent.
Max. Email Length	Limit the size of the email by specifying the maximum number of bytes it can contain. It must be stated numerically, such as: 32768.
When to Send Emails	
Severe/ Warning/ Information	Select the preferred option for the chosen recipient from the drop-down menu: • Immediately • Summary • Never
SMU-Specific Settings	

Field	Description
Use this profile as the SMU's profile	Fill the check box to use this profile as the SMU's profile. (Note that this is only valid for an internal SMU.)
Send Email From	For emails that will be sent by this SMU, enter the address that you want listed as the sender's email address. Note that this field is not available for internal SMUs.
Enable Monthly Call Home Emails	Fill the check box to enable monthly "call home" emails. Clear the check box if you do not to receive monthly "call home" emails.
Send Summaries At	Set the time when summary emails should be sent. Set the exact time (hh:mm) in a 24-hour format (i.e. 2 PM will be set as 14:00). A second summary can also be sent by entering a time in the second box.
Email Intro Text	Custom text to add to the body of the email. You can use this text field to add information or comments to the body of the email. If you are sending HTML emails, you can add basic HTML formatting (italics, bold, new lines and paragraphs, etc.) to the email, and the additional text will be displayed according to the formatting you entered.
Recipients	Displays the current recipient's Email Address.
add Recipient	Enter the Email Address of the recipient about to be added to the profile. Click **Add** to add the specified recipient to the current profile. Click **X** to delete the selected recipient from the current profile.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Managing Email Alerts and Profiles

The **Email Alerts Setup** page can be used to delete a profile or modify its properties.

To modify email alerts and profiles:

1. **Navigate to the SMTP Email Profile page.**

From the **Status & Monitoring** page, select **Email Alerts Setup.** Select a profile, then click **details** to display the **SMTP Email Profile** page:

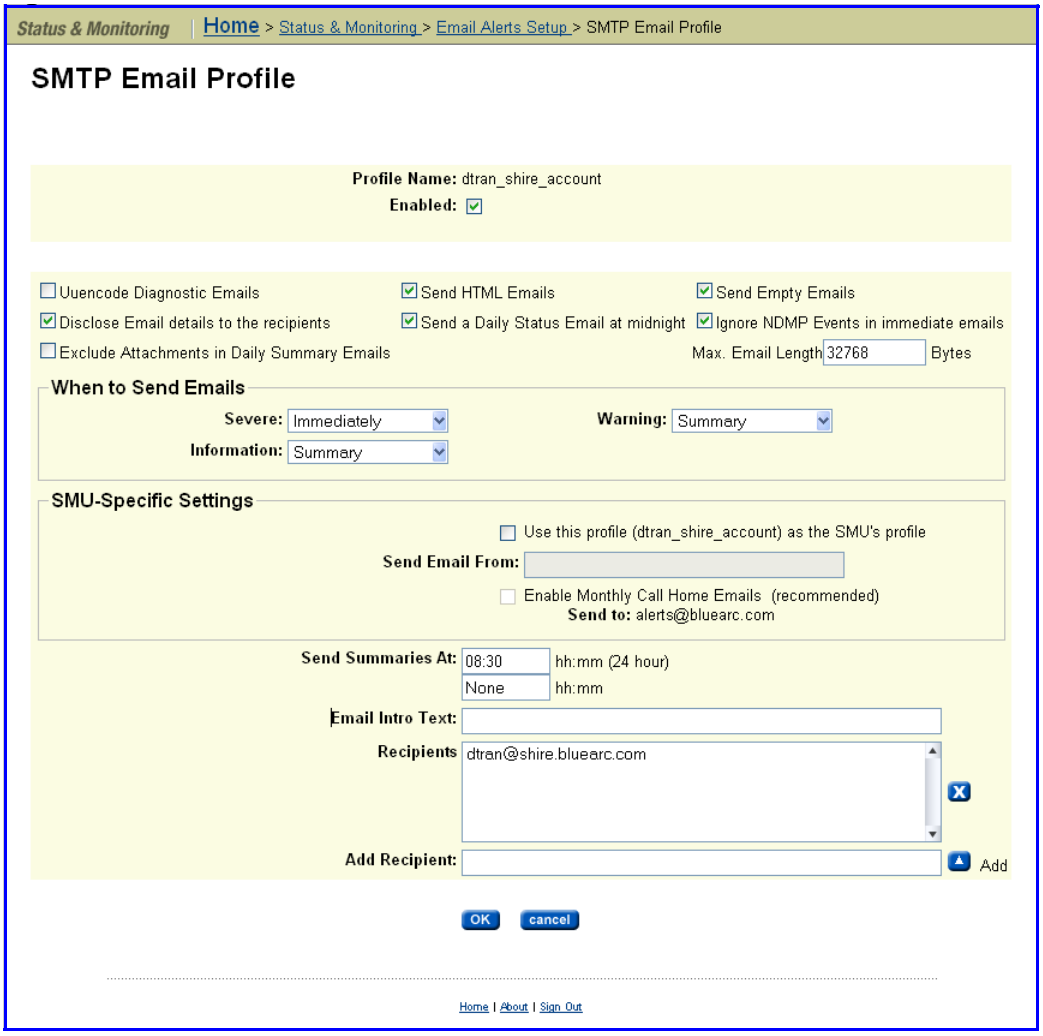

2. **Modify as needed.**

   Modify the profile by selecting the desired alert options from the drop-down menus or the checkboxes. See "Adding an Email Profile" for a description of the fields and options on this page.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Note: The **SMTP Primary Server IP/Name** should always point to SMU's Private IP address, while the **SMTP Secondary Server IP/Name** should always point to the company's main SMTP server.

Setting Up an SNMP Agent

The Simple Network Management Protocol (SNMP) is a standard protocol for managing connected network devices. An SNMP agent can be set up so that

Network Management Stations (NMS) or SNMP managers can access its management information.

The server supports SNMP versions 1 and 2c.

SNMP Statistics

SNMP statistics (per port and overall in ten-second timeslices) are available for activity since the previous reboot or since the point when statistics were last reset.

The Management Information Base

The SNMP agent maintains a Management Information Base (MIB) that is organized in a treelike structure, with each item of data having a unique object identifier (OID) that is written as a series of numbers separated by dots.

The storage server SNMP agent not only supports the MIB-II specification as described in RFC1213, but also provides an Enterprise MIB module, making management facilities available beyond those in the MIB-II specification. Download the Enterprise MIB module from the SMU, or contact SGI Global Services for the latest Enterprise MIB module. The Enterprise MIB module is compiled for SNMP v2c, and is defined in two modules, BLUEARC-SERVER-MIB and BLUEARC-TITAN-MIB.

Implementing SNMP Security

The SNMP agent is provided for monitoring purposes only; it provides Read Only access. By default, the SNMP agent does not permit access to the Management Information Base (MIB). Access is enabled by specifying:

- The version of the SNMP protocol with which requests must comply.
- The community names of the SNMP hosts and their associated access levels.
- The IP address or name of hosts from which requests may be accepted (or just choose to accept requests from any host).

To configure SNMP access to the MIB:

1. **Navigate to the SNMP Access Configuration page.**

   From the **Server Settings** page, click **SNMP Access Configuration** to display the page:

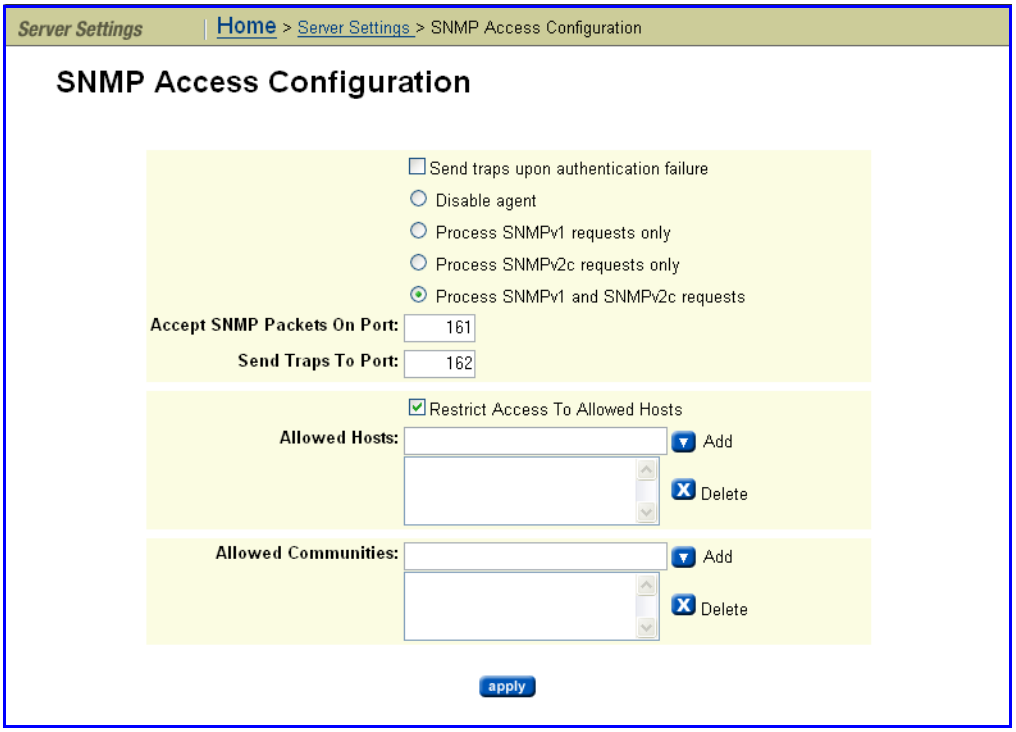

2. Complete the requested information.

The following table describes the components in this page:

Item/Field	Description
Send traps upon authentication failure	Fill this checkbox if the SNMP agent is to send a trap in the event of an authentication failure (caused, for example, by the SNMP host using an incorrect community string when formulating a request).
SNMP Protocol Support	Using the four radio buttons at the top of the page, select the version of the SNMP protocol with which hosts must comply when sending requests to the agent, or alternatively, disable the SNMP agent.
Accept SNMP Packets On Port	Enter the port number that the server monitors for communication through the SNMP protocol.
Send Traps to Port	Enter the port number that the server uses to send traps.
Restrict Access To Allowed Hosts	Fill this checkbox to restrict protocol access to the hosts specified on this page. Make sure the checkbox is empty to enable the protocol to access any host.

Item/Field	Description
Allowed Hosts	To permit requests from authorized hosts only, type the IP address of a host in this field, then click **Add** to include it in the list. (If the system has been set up to work with a name server, you can type the name of the SNMP manager host rather than its address.) You can delete an entry in the list by selecting it and clicking the **X**.  **Note:** If access is restricted to specified hosts, you should add the SMU as an allowed host.
Allowed Communities	Type the name of a community (a password) that will provide authentication into the MIB, then click **Add** to include it in the list. Community names are case-sensitive. You can delete an entry in the list by selecting it and clicking the **X**.  **Note:** SGI Global Services recommends that at least one community entry be defined.

3. **Save your changes.**

   Verify your settings, then click **apply** to save your changes.

Sending SNMP Traps

A trap is unsolicited information that the SNMP agent sends to a manager. It enables the agent to alert the manager to some unusual system event. The SNMP agent supports the following limited set of traps:

- **AuthenticationFailure.** Indicates that the SNMP agent received a request from an unauthorized manager. Either the manager used an incorrect community name or the agent has been set up to deny access to the manager.
- **ColdStart.** Indicates that the SNMP agent has started or been restarted.
- **LinkUp.**Indicates that the status of an Ethernet link has changed from *Down* to *Up*.

To set up SNMP notification:

1. **Navigate to the SNMP Traps Setup page.**

   From the **Status & Monitoring** page, click **SNMP Traps Setup** to display the page:

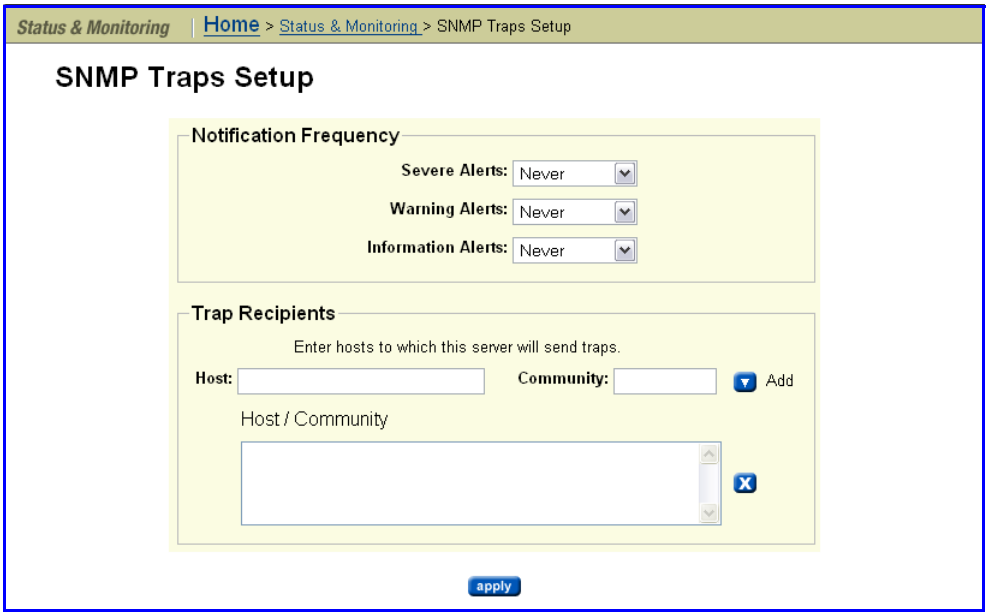

2. **Enter the necessary information.**

   The following table describes the components of this page:

Item/Field	Description
Notification Frequency	Using the drop-down list, select the notification frequency for each type of alert: • **Severe Alerts**: The specified component has failed in a way that poses a significant threat to the continued operation of the system. • **Warning Alerts**: The specified component needs attention but does not necessarily represent an immediate threat to the continued operation of the system. • **Information Alerts**: The specified component is operating normally and is not displaying an alarm condition.
Trap Recipients	In this area, enter the hosts to which this server will send traps. In the **Host** box, enter the IP address of an SNMP host to associate with each community. (If the system has been set up to work with a name server, you can type the name of the SNMP manager host rather than its address.) In the **Community** field, type the name of the SNMP community (community names are case-sensitive). Click **Add** to save the information in the list. You can delete an entry in the list by selecting it and clicking the **X**.
apply	Click **apply** to save the settings.

Setting Up Syslog Notification

You can use Syslog notification to send a Syslog alert from the server to a UNIX system log when three types of events occur. The UNIX system must have its syslog daemon configured to receive remote syslog messages.

To set up Syslog notification:

1. **Navigate to the Syslog Alerts Setup page.**

   From the **Status & Monitoring** page, click **Syslog Alerts Setup** to display the page:

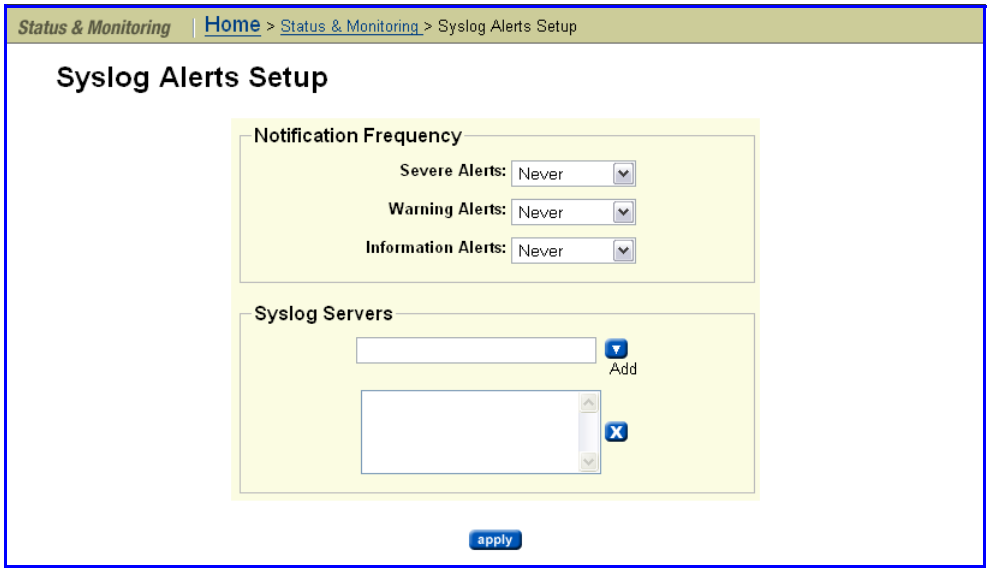

2. **Enter the necessary information.**

   The fields on this page are described in the following table:

Field	Description
Notification Frequency	Using the drop-down list, select the notification frequency for each type of alert: • **Severe Alerts**: The specified component has failed in a way that poses a significant threat to the continued operation of the system. • **Warning Alerts**: The specified component needs attention but does not necessarily represent an immediate threat to the continued operation of the system. • **Information Alerts**: The specified component is operating normally and is not displaying an alarm condition.
New Syslog Recipient	In this area, enter the syslog servers to which this server will send alerts. In the first box, enter the IP address or host name of the syslog server. Click **Add** to save the address in the list. You can delete an entry in the list by selecting it and clicking the **X**.

Testing the Alert Configuration

After setting up the alert configuration, send a test alert to all selected recipients.

1. **Navigate to the Send Test Alerts page.**

   From the **Status & Monitoring** page, **click Send Test Event** to display the page:

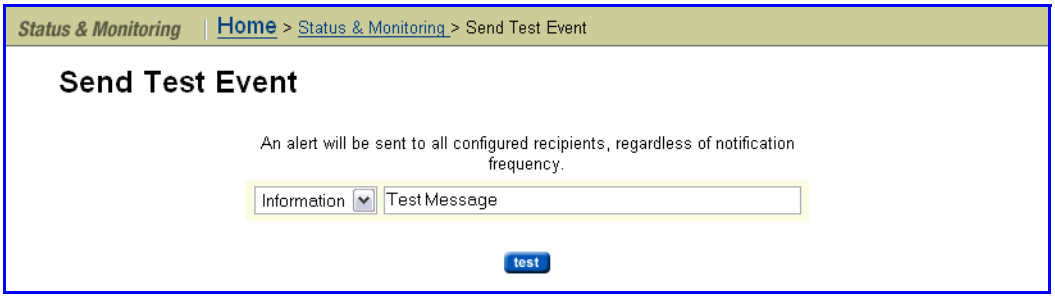

2. **Create and send a test message.**

   Select a type of message to send from the drop down list (*information*, *warning*, or *severe*), then enter a test message in the empty box, and click **test**.

File System Auditing

File system auditing monitors and records file access and deletion operations performed through the CIFS protocol. These operations are recorded in the file system's audit log. You can then view the file system's audit log, and use a remote Windows Event Viewer to save the log entries for later review. File system audit logging is performed and controlled on a per file system basis.

Because CIFS defines open and close operations, auditing file system object access performed by clients using other protocols would be costly in terms of system performance, because each I/O operation would have to be audited as an open operation. Therefore, when file system auditing is enabled, **by default, only clients connecting through the CIFS protocol are allowed access to the file system**. Access by clients using other protocols, like NFS, can, however, be allowed. When such access is allowed, access to file system objects through these protocols is not audited.

Note: You can configure file system auditing to deny access to clients connecting with protocols that cannot be audited (NFS).

After a file has been externally migrated (migrated to an external server), subsequent access to the file through the NAS server is audited as if the file were still local.

For known users (users with a Windows user mapping), the NAS server logs Object Access events 560, 562, 563 and 564. As with the Windows operating

system, auditable events for objects are specified by SACLs (system access control lists). Auditing events are logged under the following conditions:

- 560 – open handle

  This event is logged when a network client asks for access to an object. An access check is performed against the DACL (discretionary access control list) and an audit check is performed against the SACL. If the result of the access check matches the result of the audit check, an audit record is generated.

  For successful accesses, the audit records the accesses that were granted, and for failures the audit records the accesses that were requested.

- 562 – close handle

  This event is logged when an application closes (disposes of) an existing handle, and is logged in conjunction with event 560.

- 563 – open handle for delete

  This event is logged when a network client asks for access to a file with the CreateFile API is called, and the delete-on-close flag is specified. An access check is performed against the DACL and an audit check is performed against the SACL. If the result of the access check matches the result of the audit check, an audit record is generated.

  For successful deletions, the audit records the accesses that were granted, and for failures the audit records the accesses that were requested.

- 564 – delete

  This event is logged when an application closes (disposes of) an existing handle, and is logged in conjunction with event 563.

About File System Audit Logs

The file system audit log is buffered in memory, and may be permanently stored in a file in the file system being audited. Active audit log files are stored in Windows event log file format (.evt) so that standard tools can access them. The name, location, size of the active audit log file, log file retention, and active log file backup settings are defined when enabling auditing for a file system.

Note: File System Audit logs are saved in Windows XP format. An effect of this is that, depending upon how the saved .evt file is opened, a Windows Vista or Windows 2008 Server event viewer may report the file as corrupted, or may not be able to fully interpret the events. Note that the same situation occurs when a Windows Vista event viewer is used to view saved logs from an XP system. To view the logs correctly, use a Windows XP event viewer.

Audit log files are limited in size, and the retention behavior when a log fills is configurable. When an audit log reaches its maximum size, log entries (file system events) can be overwritten, or the full audit log can be saved, and a new log started.

Note: All file system audit log parameters are specified on a per file system

basis.

You can specify a backup policy, which backs up the active log at regular intervals, and starts a new active log file. Backup log files are created in the same directory as the active audit log file.

In the event of a server crash, active file system audit logs are recovered only if a rollback is performed on restart. Note that a rollback may reset the audit log file to a time when it can be recovered, thus saving some records that would otherwise be lost.

Controlling File System Auditing

File system auditing requires that a file system audit policy be defined for the file system to be monitored, and that auditing is enabled for the specific file system. File system auditing is performed and controlled on a per file system basis.

Enabling Auditing for a File System

File system auditing is a licensed feature. Once the FS Auditing license is installed, you can enable file system auditing on a per-file system basis. See "Managing License Keys" for information about adding license keys and viewing installed license keys.

Note: By default, when file system auditing is enabled, access to the file system will be limited to the CIFS protocol. Access by clients using other protocols, like NFS, can, however, be allowed. When such access is allowed, access to file system objects through these protocols is not audited.

To enable file system auditing for a particular file system, the file system must be added to the file system audit list. To add the file system to the file system audit list:

1. **Navigate to the File System Audit Policies page.**

   From the **File Services** page, click **File System Audit Policies** to display the page:

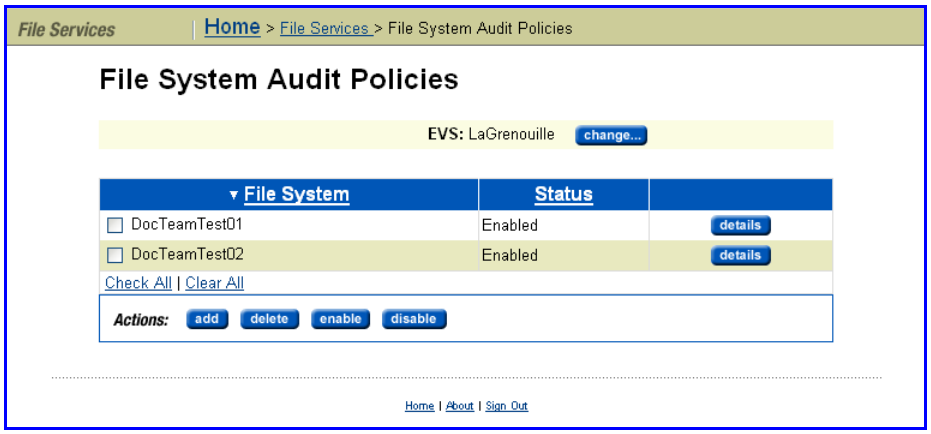

   This page displays a list of all file systems in the currently selected EVS that have file system auditing enabled.

The fields on this page are described in the following table:

Field/Item	Description
EVS	Lists the EVS to which host the file system is assigned. Click **change** to go to the **Select an EVS** page, where you can select a different EVS.
File System	Lists all file systems in the specified EVS that have an audit policy.
Audit Policy Status	Indicates whether file system auditing is enabled or disabled.
details	Click to open the **File System Audit Policy Details** page, where you can change the auditing options for a file system.
add	Click to open the **Add File System Audit Policy** page, where you can set the auditing options for a file system. Only one audit policy is allowed per file system.
delete	Click to delete the audit policy for a selected file system.
enable	Click to enable file system auditing for the selected file system.
disable	Click to disable file system auditing for the selected file system.

2. **Select the file system and enable logging.**

   If the file system on which you want to enable auditing is listed, an audit policy has already been defined for that file system.

   - If the **Audit Policy Status** is "Enabled," logging is already enabled for the file system, and no further actions are required.
   - If the **Audit Policy Status** is "Disabled," fill the check box next to the file system name, and click **enable**.

   You can now review the audit policy, and make changes if you want to do so. See Modifying a File System Audit Policy, on page 521 for more information on modifying file system audit policies.

   If the file system on which you want to enable auditing is not displayed, a file system audit policy may not have been defined for that file system, or the file system may have an audit policy defined, but the file system is not in the currently selected EVS.

   To select a different EVS, click **change** to go to the **Select an EVS** page, where you can select a different EVS.

- If, after selecting the EVS that hosts the file system, the file system on which you want to enable auditing is now listed on the **File System Audit Policies** page, fill the check box next to the file system name, and click **enable**.

  You can now review the audit policy, and make changes if you want to do so. See Modifying a File System Audit Policy, on page 521 for more information on modifying file system audit policies.
- If, after selecting the EVS that hosts the file system, the file system on which you want to enable auditing is still not displayed, you must define a file system audit policy for that file system. Click **add** to display the **Add File System Audit Policy** page, where you can set the auditing options for a file system. See Creating a File System Audit Policy, on page 518 for information on creating the file system audit policy.

Creating a File System Audit Policy

The file system audit policy specifies access restrictions for clients connecting through unauditable protocols (if access is allowed or denied), and specifies audit log details. The audit log policy specifies naming, location in the file system, size, the log roll over policy, and the backup policy. To create a file system audit policy:

1. **Navigate to the Add File System Audit Policy page.**

From the **File Services** page, click **File System Audit Policies** to display the **File System Audit Policies** page, then click **add** to display the **Add File System Audit Policy** page:

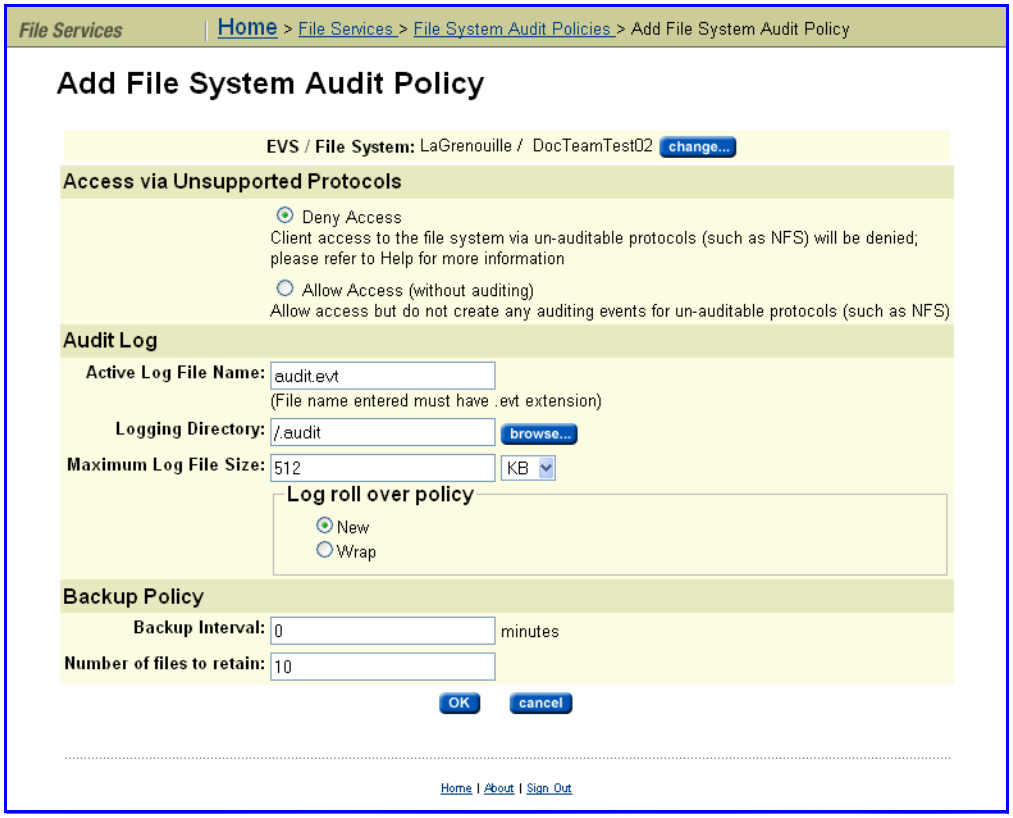

This page displays the file system audit policy settings and options. The fields on this page are described in the following table:

Field/Item	Description
EVS/File System	Lists the currently selected EVS and file system, to which the audit policy will apply. Click **change** to go to the **Select a File System** page, where you can select a different EVS and file system.
Access via Unsupported Protocols	Allows you to: • **Deny Access**. Client access to the file system via unauditable protocols (such as NFS) is denied. • **Allow Access**. Allows client access to the file system via unauditable protocols (such as NFS), but does not create any auditing events.
Active Log File Name	Specify the file name for the file system audit log. The file name must have an ".evt" extension. The default file name is audit.evt.
Logging Directory	Specify the directory within the file system where the file system audit log files are saved. You can use the **browse** button to search for an existing directory, or enter the name of a directory to be created.

Field/Item	Description
Maximum Log File Size	Specify the maximum size of the active audit log file in KB, MB, or GB.
Log roll over policy	Determines what the system does once the active audit log file is full (when it reaches the Maximum Log File Size). You can select either: • **Wrap**, which causes the system to delete the oldest existing audit entry to allow room for a new entry. • **New**, which causes the system to create a new active audit log file. The default is New.
Backup Interval	Specify the time (in minutes) between automatic backups of the active audit log. A value of 0 disables the automatic backups. The default is 0.
Number of files to retain	Specify the number of backup audit log files to retain. The default is 10.
OK	Click **OK** to save the file system audit policy.
cancel	Click **cancel** to exit this page without creating the file system audit policy.

2. **Specify the access settings for unsupported (unauditable) protocols.**

   When clients attempt to access the file system through a protocol that does not support auditing (such as NFS), this setting determines if those clients are permitted to access the file system. You can select either:

   - **Deny Access**. Client access to the file system via unauditable protocols (such as NFS) is denied.

     Specifying **Deny Access** will generate an error if there is an NFS export for the file system or the file system has a FTP user that has a directory available. To make sure this error is not generated, you can remove the NFS export for the file system, remove the FTP user, or select the **Allow Access** option.
   - **Allow Access**. Allows client access to the file system via unauditable protocols (such as NFS), but does not create any auditing events.

3. **Specify the active audit log file name.**

   Specify the name for the active audit log file. The file type suffix must be "evt."

4. **Specify the logging directory.**

   You can choose an existing directory by clicking **browse**, or you can enter the name of a directory, and that directory will be created.

Note: For ease of access to the audit log files, the logging directory should be within in a CIFS share that can be accessed by those who need to review the access log.

5. **Specify the maximum log file size.**

6. **Specify the roll over (retention) policy.**

   When the log file reaches the size specified by the **Maximum Log File Size**, this setting determines what happens to the audit log file. You can select either:

   - **Wrap**, which causes the system to delete the oldest existing audit entry in the active audit log file, so that a new entry can be added.
   - **New**, which causes the system to create a new active audit log file.

   The default is **New**.

7. **Specify the backup interval.**

   This is the time (in minutes) between saving (backing up) copies of the active audit log.

8. **Specify the number of files to retain.**

   The number of backup files to keep. A value of 0 disables automatic backups. The default is 0 (backups are disabled).

9. **Save the policy.**

   Click **OK** to save the policy as specified.

Modifying a File System Audit Policy

To modify an existing file system audit policy:

1. **Navigate to the File System Audit Policies page.**

   From the **File Services** page, click **File System Audit Policies** to display the **File System Audit Policies** page.

2. **Navigate to the File System Audit Policies page containing the policy you want to modify.**

   If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, where you can select a different EVS.

3. **Select the policy to modify.**

   Click the **details** button on the file system with the audit policy you want to modify to display the **File System Audit Policy Details** page:

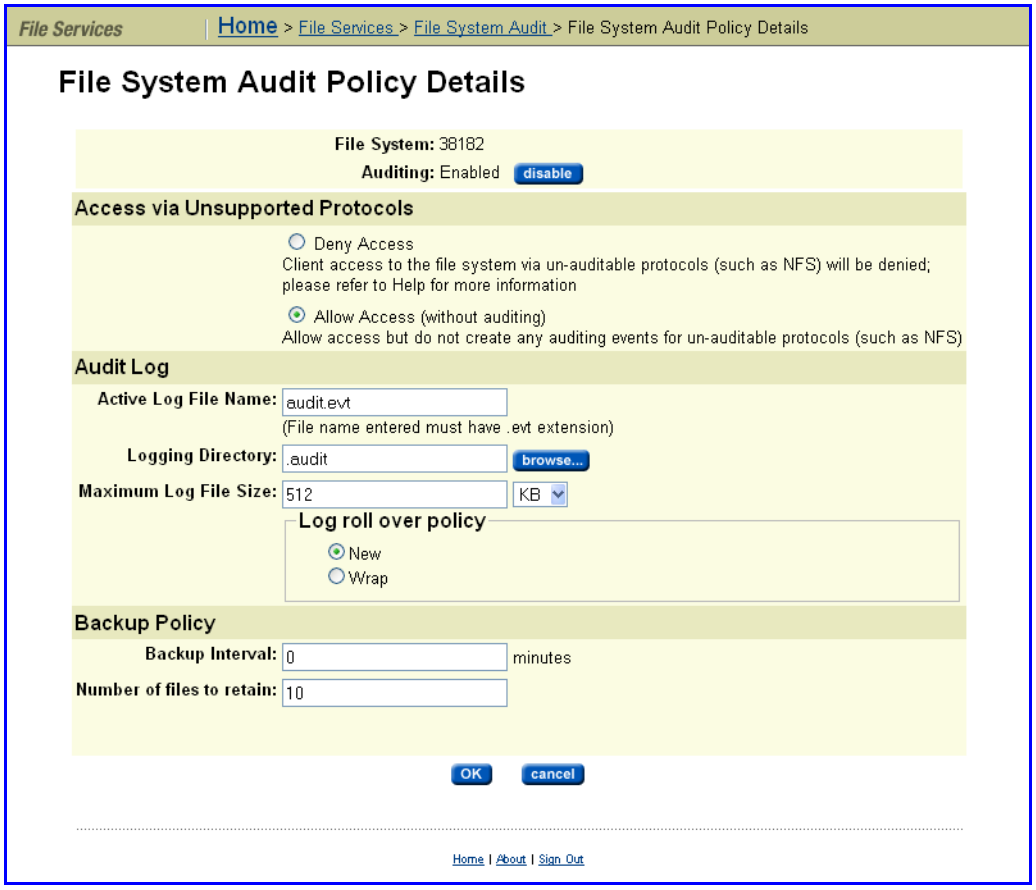

This page displays the file system audit policy settings and options. The fields on this page are the same as those described in Creating a File System Audit Policy, on page 518.

4. **Modify the policy as required.**

   Change the policy as required. For information about the fields and options on this page, see Creating a File System Audit Policy, on page 518.

5. **Save the modified policy.**

   Click **OK** to save the policy as specified.

   To return to the **File System Audit Policies** page without saving changes to the policy, click **cancel**.

Enabling/Disabling Auditing for a File System

To enable/disable file system auditing:

1. **Navigate to the File System Audit Policies page.**

   From the **File Services** page, click **File System Audit Policies** to display the **File System Audit Policies** page.

2. **Navigate to the File System Audit Policies page containing the policy you want to enable/disable.**

If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to disable. To select a different EVS, click **change** to go to the **Select an EVS** page, where you can select a different EVS.

3. **Select the policy to enable/disable.**

   Fill the check box next to the name of the file system with the audit policy you want to disable.

4. **Enable/disable the policy.**

   Click:

   - **Enable** to allow a disabled policy to function again.
   - **Disable** to stop the policy from functioning (note, however, that the policy is not removed).

Note: When disabled, file system access operations are not logged, and protocol restrictions are not enforced.

Deleting a File System Audit Policy

To delete a file system audit policy:

1. **Navigate to the File System Audit page.**

   From the **File Services** page, click **File System Audit Policies** to display the **File System Audit Policies** page.

2. **Navigate to the File System Audit page containing the policy you want to modify.**

   If the file system with the audit policy you want to change is not displayed, change the currently selected EVS to display the EVS hosting the file system with the audit policy you want to change. To select a different EVS, click **change** to go to the **Select an EVS** page, where you can select a different EVS.

3. **Select the policy to delete.**

   Fill the check box next to the name of the file system with the audit policy you want to delete.

4. **Delete the policy.**

   Click **delete** to completely remove the policy.

Note: Existing log files are not deleted automatically when a policy is deleted. If you want to delete these logs, you must do so manually,

Viewing File System Audit Logs

The NAS server supports using a remote Windows Event Viewer to display file system audit log events. Assuming that the logging directory is within a CIFS share, the audit log files can be accessed by the Windows Event Viewer. Using the Windows Event Viewer, you can view, save and clear the local

event logs, or those on a remote computer. Audit logs can be saved in several formats, including a ".evt" event format or a plain text file. The Windows Event Viewer can only save in ".evt" format to a file on the same computer as the event log, because it is the computer being viewed that does the copy (meaning the Event Viewer doesn't just read the event log and write it to a file). The Event Viewer can also be used to open and view saved audit log files.

FTP Auditing

FTP Audit Logging is controlled on a per-EVS basis. When enabled, the system maintains an audit log which tracks user activity performed through the FTP protocol for all file systems in the EVS. Each time a user takes any of the following actions, the system records the event:

- Logging in or out (including when a session timeout occurs).
- Renaming or deleting a file.
- Retrieving, appending or storing a file.

  In this case, the system records the success or otherwise of the action at both its start and end.
- Creating or removing a directory.

The **FTP Audit Logs** page displays the FTP audit logging status for each EVS in the server/cluster. Using this page, you can view FTP logging status, enable/disable FTP audit logging, and you can also display the **FTP Audit Log Details** page, which allows you to configure log file details.

From the **File Services** page, click **FTP Audit Logs** to display the **FTP Audit Logs** page:

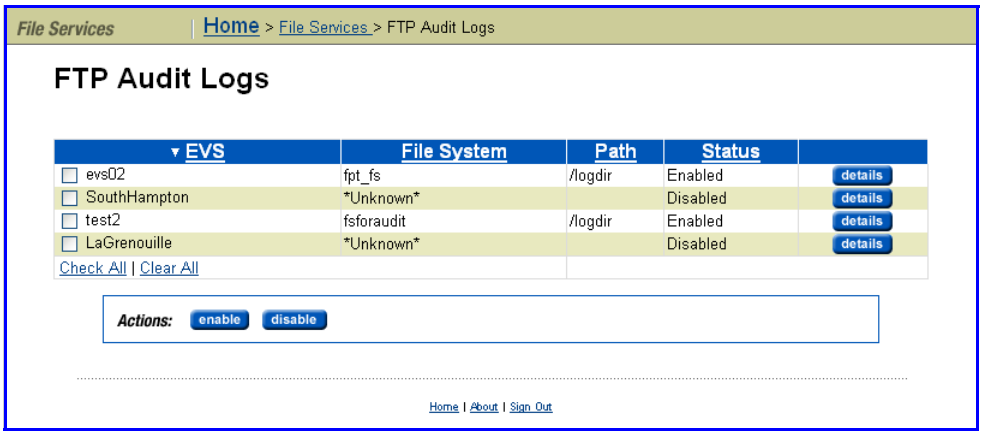

For each EVS in the server/cluster, this page lists the status of FTP audit logging, and displays the file systems being monitored, as well as the path to the FTP audit logs for each monitored file system. The fields on this page are described in the following table:

Field/Item	Description
EVS	Lists the file serving EVS in the server/cluster.
File System	Lists the file systems in the server/cluster.
Path	Displays the directory path in the file system where the FTP audit log is located.
Status	Indicates whether FTP auditing is enabled or disabled. Click **enable** or **disable** to enable/disable auditing.
details	Click details to display the **FTP Audit Log** page, which allows you to configure FTP audit logging for the file system.

Managing FTP Audit Logging

This section provides information on:

- Enabling and Disabling FTP Error Logging
- Configuring FTP Auditing

Enabling and Disabling FTP Audit Logging for an EVS

FTP Audit Logging is enabled or disabled on a per-EVS basis, meaning that is is enabled or disabled for all file systems served by the EVS and accessed through the FTP protocol. To enable or disable FTP Audit Logging:

1. **Navigate to the FTP Audit Logs page.**

   From the **File Services** page, click **FTP Audit Logs** to display the **FTP Audit Logs** page.

2. **Select the EVS for which you want to enable/disable FTP Audit Logging.**

   Fill the check box for the EVS for which you want to enable/disable FTP Audit Logging.

   - If FTP Audit Logging is disabled, you can enable it by clicking **enable**.
   - If FTP Audit Logging is enabled, you can disable it by clicking **disable**.

   Use the information in the following sections to specify the FTP Audit Logging configuration.

Configuring FTP Audit Logging

To configure FTP Auditing for the file systems in an EVS:

1. **Navigate to the FTP Audit Logs page.**

   From the **File Services** page, click **FTP Audit Logs** to display the **FTP Audit Logs** page.

2. **Select the EVS for which you want to configure FTP Audit Logging.**

For the EVS for which you want to configure FTP Audit Logging, click details to display the **FTP Audit Log Details** page.

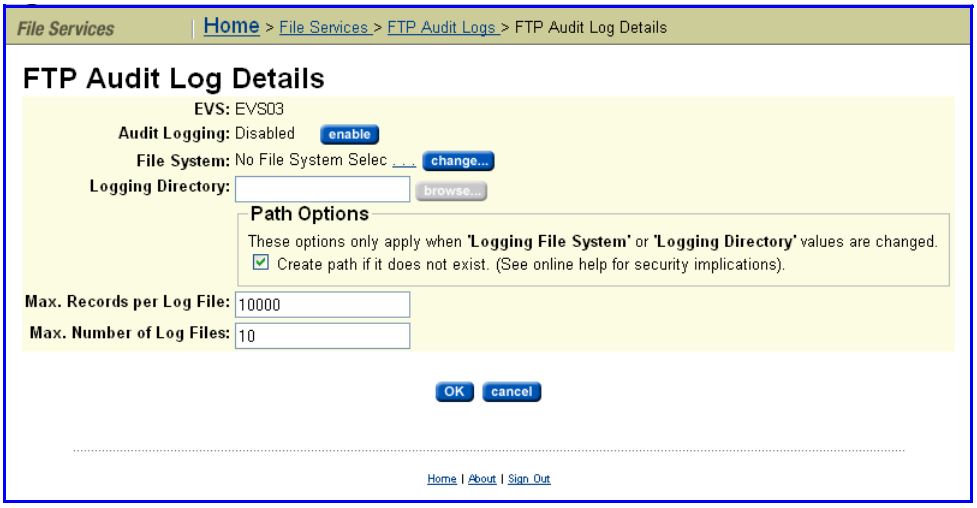

This page displays the FTP Audit Logging settings and options and is used to configure FTP Audit logging. This configuration includes:

- Enabling/disabling FTP audit logging for an EVS.
- Specifying the file system where the FTP Audit Log file(s) are stored.
- Specifying the maximum number of records in an FTP Audit Log file.
- Specifying the number of FTP Audit Log files that are stored.

The fields on the **FTP Audit Log Details** page are described in the following table:

Field/Item	Description
EVS	Lists the currently selected EVS and file system, to which the audit configuration will apply.
Audit Logging	Indicates if FTP Audit Logging is enabled or disabled for the EVS. Click **enable**/**disable** to enable/disable FTP Audit Logging for the EVS.
File System	Displays the name of the file system that will contain the FTP Audit Log file(s). Click **change** to select a different file system.

Field/Item	Description
Logging Directory	Displays the directory path in the file system where the FTP audit log files are stored. The Path Options allow you to select an existing directory, or to create the directory if it does not already exist.  **Note:** The **browse**... button only exists if the path being created is the path in a file system, not a namespace.  **Note:** Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users, i.e. the permissions are set to rwxrwxrwx. It is recommended that such directories are created via CIFS or NFS or that such directories are given the desired permissions explicitly after being created via this option.
Max. Records per Log File	Specifies the maximum number of records per log file. Once the maximum number of records per file is reached, a new log file is started. Each log file is a tab-delimited text file containing one line per FTP event. Besides logging the date and time at which an event occurs, the system logs the user name and IP address of the client and a description of the executed command.
Max. Number of Log Files	Specifies the maximum number of log files to be kept. Once the maximum number of log files is reached, when the current log file becomes full, the oldest log file is deleted. The newest log file is called ftp.log, and the older files are called ftpn.log (the larger the value of n, the older the file).
OK	Click **OK** to save the FTP Audit Logging configuration described above.
cancel	Click **cancel** to exit the **FTP Audit Logs** page without saving the FTP Audit Logging configuration described above.

3. **Enable/Disable FTP Auditing for the EVS.**

   For the currently selected EVS, you can enable/disable FTP Audit Logging.

   - If FTP Audit Logging is disabled, you can enable it by clicking **enable**.
   - If FTP Audit Logging is enabled, you can disable it by clicking **disable**.

4. **Specify the FTP Audit Logging Configuration**

   a. **Specify the file system for the logging directory.**

      In the File System field, choose a file system in which to keep the log files. Click change to see a list of file systems in the EVS.

      For optimum performance, keep the log files on a different system drive than the files that users will access over FTP.

b. **Specify the logging directory.**

The logging directory specifies the location where the FTP Audit Logs are kept. In the **Logging Directory** field, specify the directory in which to keep the log files. Click **browse** to choose an existing directory, or specify a path to be created. To create the path automatically when it does not already exist, fill the checkbox **Create path if it does not exist**.

Note: Automatically created directories will be owned by the root user and group (UID:0 / GID:0) and will be accessible to all users (that is, the permissions are set to rwxrwxrwx). It is recommended that such directories are created via CIFS or NFS, or that such directories are given the desired permissions explicitly after being created via this option.

c. **Specify the maximum number of records per log file.**

In the **Max. Number of Records per Log File** field, specify the maximum number of records to store in each log file.

For optimum performance, produce a small number of large files instead of a large number of small files.

d. **Specify the maximum number of log files to keep.**

In the **Max. Number of Log Files** field, specify the maximum number of log files to keep. Once it has reached this limit, the server deletes the oldest log file each time it creates a new one.

5. **Save the configuration.**

Click **OK** to save the FTP Audit Logging configuration described above.

Click **cancel** to exit the **FTP Audit Logs** page without saving the FTP Audit Logging configuration described above.

Viewing FTP Audit Logs

FTP Audit logs can be viewed with a text editor. If the logging directory is within an NFS Export or a CIFS share, simply access the directory and open the log file. If the logging directory is available through FTP, you can download the file, then open it with a text editor.

Monitoring Fibre Channel Switches

The server allows you to add Fibre Channel (FC) switches to the System Monitor, so you can easily check FC Switch connectivity status, which indicates whether the SMU received a response to an Ethernet ping of its last-known IP address. The connectivity status does **not** indicate whether the FC switch has connectivity with the storage subsystem.

When adding an FC Switch to the System Monitor, you can associate it with one or more servers. For information on adding FC Switches to the System Monitor, see Adding FC Switches, on page 531. Once an FC switch has been

associated with a server, you can monitor switch connectivity status, view log events and SNMP traps, download FC switch diagnostic information, and configure emailing of switch-related diagnostic information.

Displaying the Connectivity Status of Fibre Channel Switches

The System Monitor displays FC switch connectivity status at a glance, and also lists FC switches, which can be selected to display detailed switch information.

Using System Monitor to Display Switch Connectivity Status

To use System Monitor to display Switch Connectivity Status, navigate from the **Server Status Console** (click the indicator to the right of the server) to display the **System Monitor** page.

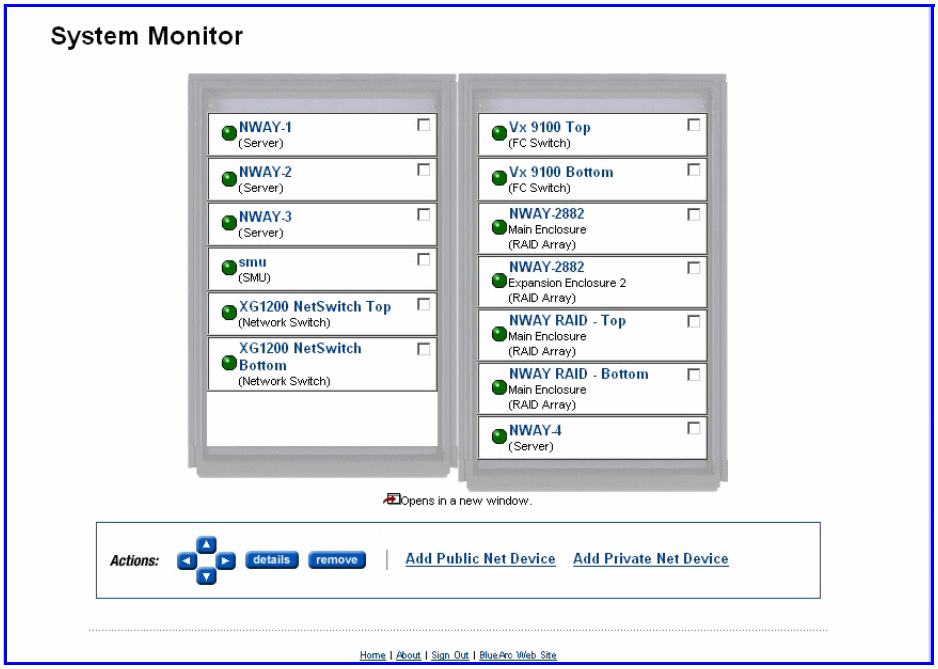

The status indicator next to the FC switch indicates its connectivity status.

For more information about the System Monitor and the device status indicators, see Checking the System Status, on page 455.

Using Web Manager to Display Switch Connectivity Status

To use Web Manager to display Switch Connectivity Status, navigate from the **Storage Management** page to the **FC Switches** page:

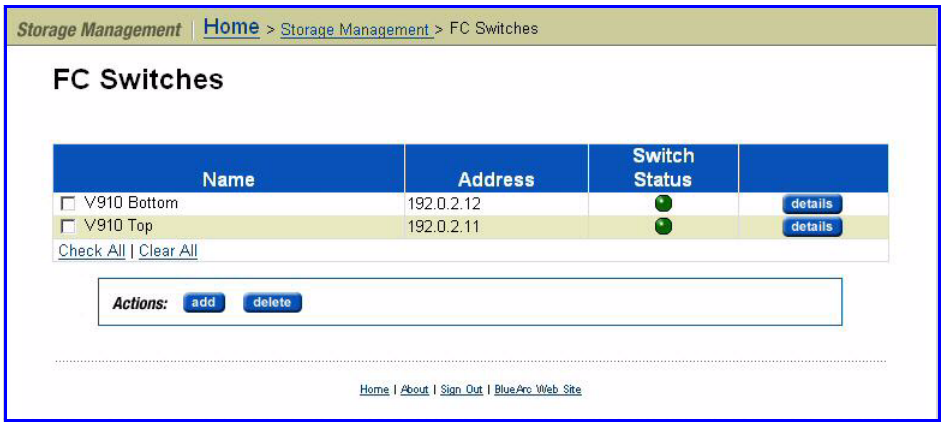

The following table describes the components of this page:

Item/Field	Description
Name	The name of the switch, defined when the switch was added. This name should be sufficiently descriptive as to be able to identify the switch.
Address	The IP address or DNS name of the switch, defined when the switch was added.
Switch Status	An indicator of the connectivity status of the switch. Connectivity status indicators are: • **Green – OK.** A response was received from a ping of the last-known IP address of the switch. • **Gray – Determining state.** A FC switch will appear as gray for up to 60 seconds, immediately after it is added. After a ping of the switch IP address, the status will change to OK or severe (green or red), depending on whether there was a response to the ping. • **Red – Severe.** No response was received from a ping of the IP address of the switch. See Checking the System Status, on page 455 for more information about status indicators.
Actions	
Details	Click **details** to view the **FC Switch Details** page for the switch. From the **FC Switch Details** page, you can open the embedded management interface for the switch (if available) and change the switch name or address.
add	Click **add** to add an FC switch.
delete	Click **delete** to delete an FC switch.

Adding FC Switches

Upon adding an FC Switch, the SMU displays it in the System Monitor, with connectivity status. Because multiple servers or clusters may use the storage connected to an FC switch, it can be associated with multiple servers or clusters managed by an SMU, thereby appearing in the System Monitor for all servers and cluster to which it has been associated.

To add an FC switch:

1. **Navigate to the add FC Switch page.**

   From the **Storage Management** page, select **FC Switches,** then click **add** to display **Add FC Switch** page:

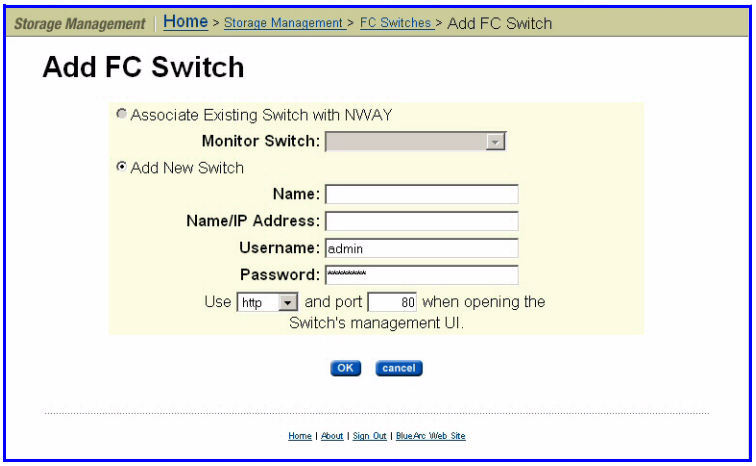

2. **Enter the requested information.**

   The following table describes the fields in this page:

Item/Field	Description
Associate Existing Switch with *name* (currently managed server)	
	Select an existing switch to associate with the named server or cluster. When you associate a switch with a managed server or a cluster, the switch is added to the system monitor of that server/ cluster.
Monitor Switch	Use the drop-down list to select the switch you want to associate with the named server/cluster.
Add New Switch	
	Select to add a new FC switch. After the switch has been added, you can associate it with a managed server or a cluster.
Name	The name you want to use to refer to the switch. This name should be sufficiently descriptive as to be able to identify the switch.
Address	The IP address or DNS name of the switch.
Username	Enter the user login name for the embedded management interface of the FC switch.

Item/Field	Description
Password	Enter the password associated with the user name for the embedded management interface of the FC switch.
Use http/https/ Telnet/other on port...	From the drop-down list, select the protocol and port for connecting with the embedded management interface of the FC switch. Defaults are *http* protocol and *port 80*. **Note:** If *http*, *https*, or *Telnet*, clicking the switch in the System Monitor displays the embedded management interface. If *other*, the **FC Switch Details** page is displayed instead of the management interface.

3. **Save your settings.**

   Verify your settings, then click **OK** to save or **Cancel** to decline.

Displaying or Changing Details about an FC Switch

You can display a list of the FC Switches that have been added to the System Monitor of any server or cluster managed by an SMU on the **FC Switches** page. Once you have displayed this list, you can view and/or change details about a switch.

To display or change FC switch details:

1. **Navigate to the FC Switches page.**

   From the **Storage Management** page, select **FC Switches**, then click **detail** for a selected Switch to display the **FC Switch Details** page, which lists all FC switches that have been added to the System Monitor of the server/ cluster:

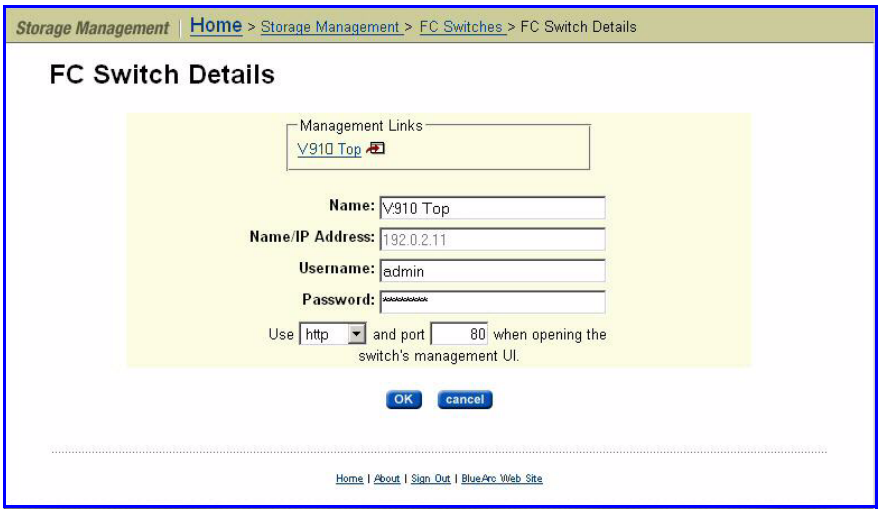

2. **As needed, view or modify switch information.**

   The following table describes the fields in this page:

Item/Field	Description
Management Links	This area provides links to the embedded management interface(s) for the FC switch. Click a link to open the interface.  **Note:** The FC switch management interface may or may not support multiple concurrent logins. Refer to the documentation for the switch regarding use of the embedded management interface.
Name	Name of the switch, specified when the switch was added. This name should be sufficiently descriptive as to be able to identify the switch.
Name/IP Address	The IP address or DNS name of the switch, specified when the switch was added.
Username	User login name for the embedded management interface of the FC switch.
Password	Password associated with the user name for the embedded management interface of the FC switch.
Use http/https/ Telnet/other on port...	Protocol and port for connecting with the embedded management interface of the FC switch. Defaults are *http* protocol and *port 80.*  **Note:** If *http*, *https*, or *Telnet*, clicking the switch in the System Monitor displays the embedded management interface. If *other,* the FC Switch Details page is displayed instead of the management interface.

3. **Save your settings.**

   If you have made changes, verify the new settings, then click **OK** to save or **Cancel** to decline.

11 Maintenance Tasks

The IS-NAS Server/Titan Server architecture includes the following maintenance tools:

Tool	Overview and Instructions
Managing License Keys	Managing License Keys, on page 535
Checking Version Information	Checking Version Information, on page 539
Installing and Managing Certificates	Providing an SSL Certificate, on page 541

System Software and Firmware Upgrade

The System Management Unit (SMU) software and storage server firmware can be upgraded to newer releases. For information on upgrading software and firmware, refer to the BlueArc Titan Server *Software Installation Guide*.

Managing License Keys

License keys add powerful services to the IS-NAS Server/Titan Server and can be purchased and added whenever needed. A License Certificate identifies all of the purchased services and should be kept in a safe place. The License Certificate is included in the User Documentation Wallet that was shipped with the system.

System Administrators manage keys for licensed services from the **License Keys** page, which displays the status (and features enabled by) each key and provides controls for adding and deleting keys.

The **License Keys** page displays a table listing:

- Each installed license key.
- Only displayed for a cluster, this is the maximum number of nodes licensed. This indicates the maximum number of servers that can be configured as nodes of a cluster.

  Cluster licenses are handled somewhat differently than other licenses. See About Cluster Licensing, on page 427 for more information about cluster licenses.
- The maximum number EVSs allowed on the server/cluster.

- The of the maximum amount of storage allowed for the server/cluster.

  The amount of licensed storage in terabytes. Note that the amount of licensed storage must be equal to or greater than the total amount of storage in **all** subsystems connected to the storage server or cluster.

- The expiration date for each key, if the key expires (expired license keys are shown in grey).

Following this table, a section displays a list of all services enabled by the installed keys.

Note: The **Total Licensed on All Keys** section of the **License Keys** page displays each of the services enabled by all installed license keys. To see the services licensed by a particular key, fill the checkbox next to the key and click the **Show licensed services**. A checkmark will appear next to the any service enabled by the selected key. To display more details about a particular license key, click **details**.

The following table lists all services that can be licensed:

Service	Description
CIFS	**Common Internet File System.** This is a message format used by Windows and MS-DOS to share files, directories, and devices.
NFS	**Network File System.** This is Sun's distributed file system that enables users of UNIX workstations (including Windows NT systems running an NFS emulation program) to access remote files and directories on a network as if they were local.
iSCSI	**Internet Small Computer System Interface.** This license enables iSCSI Initiators to communicate at block level with the servers' iSCSI targets.
Data Migrator	**Data Migrator.** Enables more efficient use of primary storage space by transferring older, less performance-critical data to secondary storage.
WORM	**Write Once Read Many file systems.** Used to store crucial company data in an unalterable state for a specific duration.
SFM	**Server Farm Migration of Virtual Servers.** Enables migration of Virtual Servers (EVSs) between servers in a Server Farm.
CNS	**Cluster Name Space.** Creates a virtual name space through which multiple file systems can be made accessible using a single mount point. If the EVS Security license is also installed, you can also create individual EVS Name Spaces.
EVS Security	**EVS Security**. Enables the creation of Secure Virtual Servers (Secure EVSs).
FS Roll Back	**File System Rollback.** A tool for restoring a file system to the state of its last successful replication.
Snapshot Restore	A tool for restoring a WFS-1 file system to the state it was in at the time a specific snapshot was taken. (For file systems formatted using WFS-1 only.)
FSRS	**File System Recovery from Snapshot.** A tool for rolling back one or more files in a WFS-2 file system to a previous version without actually copying the data from a snapshot. (For file systems formatted using WFS-2 only.)

Service	Description
Read Cache	**Cluster Read Caching.** Enables Read Caching service, which allows one cached read-only file system per EVS.  **Note:** After adding the Read Cache license, you must restart the server/cluster before you can use the read cache.
Replication	**Replication.** Enables replication to external servers (other storage servers, clusters, or an NFS server). This license is required to enable any form of replication outside the server or cluster, including ADC copies. This means that, without a replication license, you can use replication within a server/cluster (you can replicate within an EVS or to a different EVS hosted by the same server/cluster), but you cannot replicate to an external server/cluster.  **Note:** The replication license is enforced at the replication source. However, in order to reverse a replication, the source and the target must each have a replication license.
FS Auditing	**File System Auditing.** Enables the auditing of file system operations performed through the CIFS and NFS v4 protocols (which are licensed separately). When enabled, file system operations are monitored and recorded in the NAS server's file system audit log. You can then view the log through a remote Windows Event Viewer, and save the log entries for later review. See File System Auditing, on page 514 for more information.  **Note:** A CIFS license is required, because the server must be able to connect to a remote Windows Event Viewer, which can then save the server's file system audit log entries in ".evt" format.
Storage Pools	**Storage Pools.** Allows Storage Pools to host more than one file system.
WFS2	Allows file systems to be formatted using the WFS-2 file system format. The WFS-2 file system format provides an alternative file system format to the WFS-1 format (the original file system supported by the Titan Server). Note that the WFS-2 file system format is supported only on the Series 3000 Titan Server and later hardware platforms (including the IS-NAS Server), and that the WFS-2 license is provided automatically. For more information on file system formats, see File System Formats, on page 111.
XVL	**External Volume Links**. Indicates that cross volume links to data migrated to storage devices attached to a remote server (not necessarily a IS-NAS Server or a Titan Server) are enabled. See Cross Volume Links in Data Migrator, on page 124 for information on cross volume links.
Storage	Displays the supported storage subsystem options.
SGI	Enables the use of supported storage subsystems manufactured by Silicon Graphics International.

Note: *Expiration of License Keys!* License keys that have been purchased do not expire. Trial licenses, which enable features for use on a trial basis, have a predefined expiration date. Five days before the expiration of a trial license, the server's event log begins receiving a daily warning event, indicating imminent expiration; then, two days before expiration, the warning events escalate to

"severe." When a trial license has expired, the features that enabled by the license become disabled.

License Types

Licenses can be grouped into three types:

- **Boolean** licenses simply enable features/protocols, and when the license is installed the feature/protocol is enabled (for example external volume links, CIFS, or NFS). These licenses operate in a boolean fashion; if the license is present the feature/function is enabled, if not present, the feature/function is disabled.
- **Limit-based** licenses specify a limit that cannot be exceeded. These licenses limit your system to a certain total numerical upper limit of the licensed feature/function. For example, the EVS (virtual server) license is a limit-based license.

  Limit-based licenses are not cumulative. For example, if your existing cluster has an EVS license for up to 9 EVSs, and you install another EVS license for up to 8 EVSs, you still cannot have more than 9 EVSs (the highest licensed amount).For more information, contact SGI Global Services.

  Note: A cluster license is a special kind of limit-based license. When a node joins an existing cluster, its cluster license is transferred to the cluster (if necessary). See About Cluster Licensing, on page 427 for more information about cluster licensing.
- **Cumulative** licenses. Only the Storage capacity license is cumulative, and several storage capacity licenses can be used to increase the capacity or capability of the system. For example, if you have one storage capacity license for 40 terabytes and another storage capacity license for 60 terabytes, your system could manage up to 100 terabytes of storage.

If a node is removed from a cluster, you must restore its license keys for it to function properly as a standalone server. You should retain the licensing information, in case a node needs to be removed from the cluster.

Adding a License Key

To add a license key:

1. **Navigate to the License Keys page.**

   From the Web Manager **Home** page, click **Server Settings**, then click **License Keys**.

2. **Add the key:**

   Click **add**:

   - If you are entering the key manually, enter the key number in the **License Key** field, then click **add**. If you have received the license key electronically, to avoid errors, you should copy the license key from the file and paste it into the **License Key** field.
   - If you have a file that contains the license key, click **Browse** to select it, then click **Import**.

After all the keys have been entered, follow the instructions to reboot the system (if necessary).

Note: Rebooting is only necessary after adding certain license keys (the Read Cache license key for example). After adding a license key, if a reboot is required in order to start a service/protocol or enable a feature, you will be instructed to reboot or restart the system.

Deleting a License Key

To delete a license key:

1. **Navigate to the License Keys page.**

   From the Web Manager **Home** page, click **Server Admin**, then **License Keys**.

2. **Select a key.**

   Fill the checkbox next to the key you want to delete (scroll down if necessary).

3. **Commit the deletion.**

   Click **delete**.

Checking Version Information

When requesting technical support, it is important to have version information about storage server firmware and hardware. The following sections explain how to retrieve storage server firmware version information for clusters, stand-alone servers, and the SMU.

Displaying Storage Server Version Information

To display version information, navigate to the **Server Settings** page, then click **Version Information** to display either the **Version Information** page (for a cluster), or the **Version for Node** page (for a server). Both are shown below:

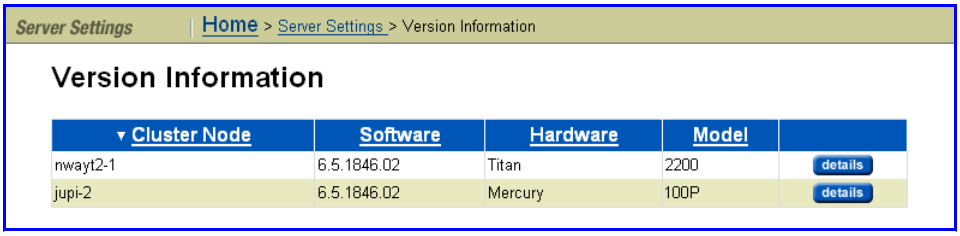

The **Version Information** page lists the nodes of the cluster along with information about the software version, hardware type, and model number. For more information on a node, click the **details** button to view the **Version for Node** page for that node.

For a IS-NAS Server, the Version Information page looks like the following:

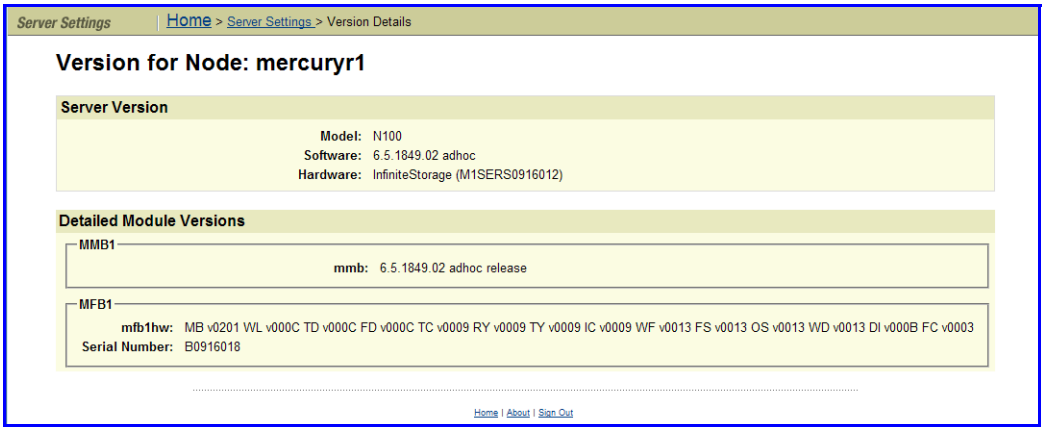

For a IS-NAS Server, the **Version for Node** page displays detailed version information of the hardware and software of the node, including information about the server model, software version, and main boards.

For a Titan Server, the Version Information page looks like the following:

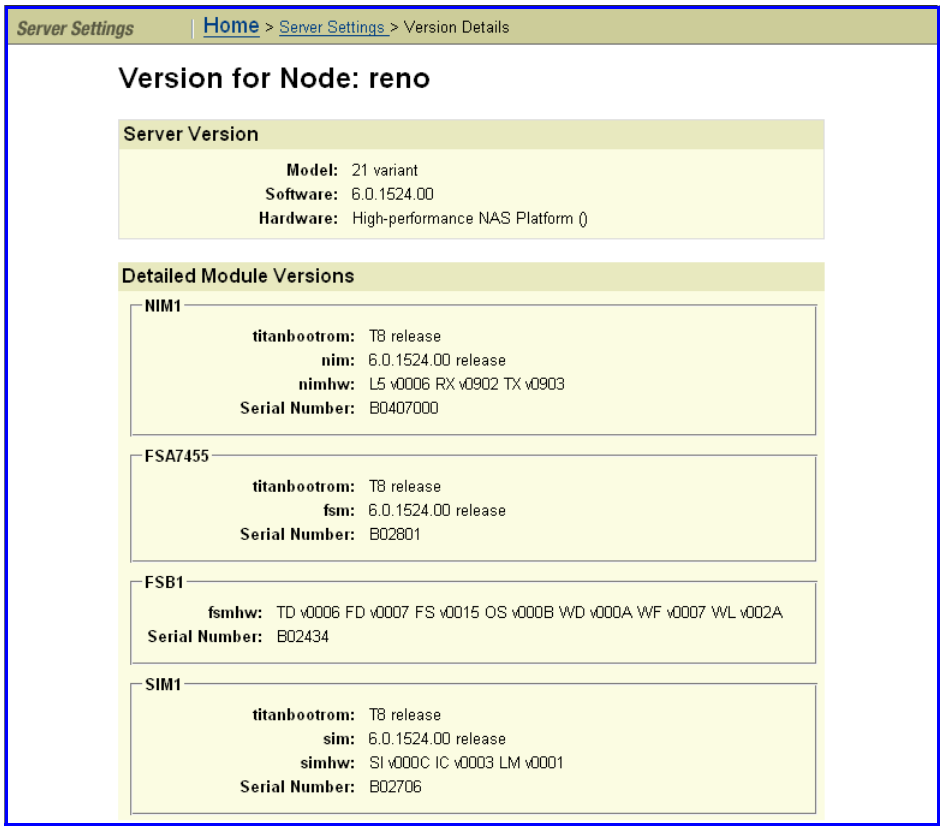

For a Titan Server, the **Version for Node** page displays detailed version information of the hardware and software of the node, including information about the four main boards (SIM, NIM, FSA, and FSB).

Displaying Version Information for the SMU

To display SMU version information, navigate to the Web Manager **Home** page, then click **About** to display the **About SMU** page:

Providing an SSL Certificate

Both the server and the SMU are pre-configured with default SSL certificates. These default certificates should provide an acceptable level of security for most users. For added security, this certificate may be replaced with a certificate signed by a Certificate Authority (for example, Verisign).

Requesting and Generating Certificates

To request a certificate from a certificate authority (CA):

- Generate a custom private key (optional).
- Generate a Certificate Signing Request (CSR).

Generating a Custom Private Key and SSL Certificate

The SMU already contains a default private key from which a CSR may be generated. It uses default values:

- Common Name (CN) uses the SMU's hostname but other values are static (for example: *OU=., O=SGI, L=San Jose, ST=CA, C=US*)
- Valid for 3650 days (10 years).
- Key length of 2048 bits.

To view these values by displaying the SMU's default certificate, type the following at the SMU CLI:

```
cert-showall.sh
```

If other values **must** be used, a custom private key may be generated via the following steps:

1. **Log in.**

   Log onto the SMU (through ssh or through its serial port) as the user "*manager*", then type:

```
sudo cert-gencustom.sh
```

Enter the password when prompted.

2. **Enter the requested information.**

   As prompts appear, provide details of the following (`enter` accepts the defaults):

   ```
   Organizational Unit (OU)

   Organization (O)

   Location (L)

   State (ST)

   Country (C)

   Valid Period (in days)

   Key Size (e.g. 1024, 2048 – must be divisible by 64).
   ```

   After the system confirms the input, it generates a new private key and self-signed certificate.

3. **Restart the web server and any browsers that connect to the SMU.**

   - Restart the web server when prompted so that it may pick up the new SSL certificate.
   - Close and restart any browsers used to connect to the SMU. This is required to purge the browser of any previously negotiated SSL session values. When logging into the SMU Web Manager, the new SSL Certificate should be provided.

4. **Propagate the new certificate to all managed servers.**

   Navigate from the **SMU Administration** page to the **Managed Servers** page. For each server, click **details,** then **OK.**

5. **Back up the private key and certificate.**

   A backup of this private key and certificate (i.e. the whole keystore) may be made for safekeeping:

   a. **Log onto the Web Manager.**

   b. **Navigate from the SMU Administration page to the SMU Backup page.**

   c. **Click Backup and save the resulting zip file to a safe and secure location.**

      The zip file contains a full backup of the SMU's configuration information. The file "smu.keystore" within the zip file contains the SMU's private key.

Generating a Certificate Signing Request (CSR)

A Certificate Signing Request is a file that contains the encoded information needed to request a certificate from an authority. After generating the Certificate Signing Request, it can be submitted to the authority. For example, on Verisign's Web site at `http://www.verisign.com/`, paste the Certificate Signing Request into the Web page.

Generating a CSR

To generate a CSR:

1. **Complete into steps in Generating a Custom Private Key and SSL Certificate, on page 541.**
2. **Log in.**

   Log onto the SMU (through ssh or through its serial port) as the user "*manager*", then type:

   ```
   sudo cert-gencsr.sh
   ```

   Enter the password when prompted.
3. **Copy and paste the CSR.**

   Copy the CSR that is displayed after step 1. Paste into the Web site of the Certificate Authority.

   Alternatively, the same information may be copied off the SMU via the file:

   ```
   /var/opt/smu/conf/ssl/certreq.csr
   ```

Acquiring a SSL Certificate from a Certificate Authority (CA)

At this point, the CSR can be submitted to a Certificate Authority such as Verisign, Thwate, etc. The details of how to do this are beyond the scope of this document.

When acquisition of a certificate from a Certificate Authority has been completed, move ahead to the section Installing and Managing Certificates, on page 543.

Installing and Managing Certificates

Once a certificate has been obtained from the Certificate Authority, follow these instructions to install it.

Installing a Certificate

1. **Copy the certificate to the SMU.**

   Copy the certificate provided by the Certificate Authority to the SMU (for example, scp to `/home/manager/server.cer`).

   If necessary, provide the Certificate Authority's Trusted Certificate Chain as a file (for example, `/home/manager/veritas.pem`). The SMU already includes popular Certificate Authority Trust Chains, so this step can typically be skipped. To view these popular Certificate Authorities, see Sun's documentation:

`http://java.sun.com/j2se/1.5.0/docs/tooldocs/solaris/keytool.html#cacerts`

Note: The content of the certificate and trust chain files should only start with "`-----BEGIN`" and end with "`-----END CERTIFICATE-----`".

2. **Log in.**

   Log onto the SMU (through ssh or through its serial port) as the user *"manager"*.

3. **Import the Certificate Authority's Trusted Certificate Chain (optional).**

   This may require multiple files/chains, so repeat as necessary:

   ```
   sudu cert-importtrustchain.sh <path to trust chain file>
   <unique alias>
   ```

   When prompted, enter the password for user "*manager*".

Note: Any unique alias may be used. If the alias already exists in the SMU's key store, you will be prompted to replace the old certificate or cancel the import.

An example Intermediate Certificate Authority trust chain may be found at: `http://www.verisign.com/support/install2/intermediate.html`

4. **Import the signed Certificate Reply.**

   The signed "Certificate Reply" from the CA may imported (replacing the default SMU SSL certificate):

   ```
   sudo cert-import.sh <path to cert file>
   ```

5. **Restart the web server and any browsers that connect to the SMU.**

   - Restart the web server when prompted so that it may pick up the new SSL certificate. When prompted to overwrite the existing certificate, enter `y`.
   - Close and restart any browsers used to connect to the SMU. This is required to purge the browser of any previously negotiated SSL session values.
   - When logging into the SMU Web Manager, the new SSL Certificate should be provided.

6. **As needed, verify the contents of the keystore.**

   To view and verify the contents (SSL certificate and Trust Chain) of the keystore, type:

   ```
   sudo cert-showall.sh
   ```

7. **Propagate the new certificate to all managed servers.**

   Navigate from the **SMU Administration** page to the **Managed Servers** page. For each server, click **details,** then **OK.**

Restoring the Default SMU Certificate

If troubles are encountered when trying to create/import an SSL certificate, the SMU's default certificate may be restored.

To restore the default certificate:

1. **Log in.**

   Log onto the SMU (through ssh or through its serial port) as the user "*manager*", then type:

   ```
   sudo cert-gendefault.sh
   ```

   Enter the password when prompted.

2. **Restart the web server and any browsers that connect to the SMU.**
   - Restart the web server when prompted so that it may pick up the new SSL certificate. When prompted to overwrite the existing certificate, enter y.
   - Close and restart any browsers used to connect to the SMU. This is required to purge the browser of any previously negotiated SSL session values.
   - When logging into the SMU Web Manager, the new SSL Certificate should be provided.

3. **Propagate the new certificate to all managed servers.**

   Navigate from the **SMU Administration** page to the **Managed Servers** page. For each server, click **details,** then **OK.**

Accepting Self-Signed Certificates

If a self-signed certificate has been installed, users receive a security alert similar to the following when they first access the Web Manager over a secure connection:

Although users can click **Yes** to proceed, the alert reappears when they next run the Web Manager. To suppress the alert, users must choose to trust the certifying authority.

In Internet Explorer, from the Security Alert dialog box, click **View Certificate** to display the certificate:

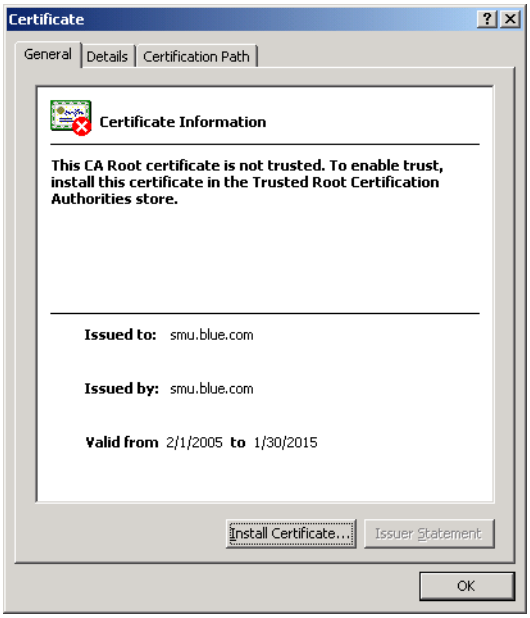

Click **Install Certificate**, then follow the on-screen instructions to install the certificate in the Trusted Root Certification Authorities store.

Mozilla-based browsers will see an alert message similar to the following. Accepting the certificate permanently suppresses the alert in future sessions.

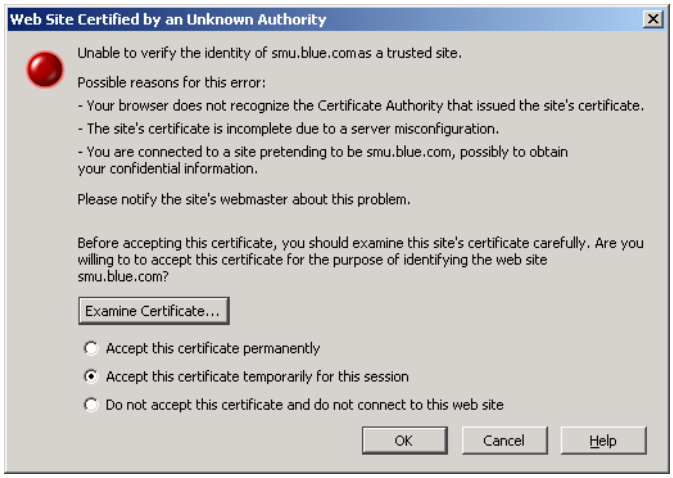

A Using Storage Management Applications

The storage server acts as an NDMP host, operating with leading storage management applications. It supports NDMP Version 2, 3 and 4. A conceptual overview of Storage Management Applications can be found at Storage Management Applications, on page 333.

This section lists environmental variables and provides naming guidelines, in the following sections:

Supported Environment Variables

Environment variables can be used to modify actions taken by applications. The following sections discuss environment variables.

NDMP

NDMP environment variables can be used to modify backup actions. The storage management application generates most of these variables and supports configuration of additional variables. They are invoked from the **Replication Rules: Add Rules** page.

DIRECT

Possible value	Notes
y or **n**	Used on recovery to request Direct Access Recovery (DAR). May be used to recover a subset of a full backup. If the storage management application supports DAR, the recovery will position the tape to the start of the required data rather than reading a complete backup image to find the data. This saves time in recovery of single files and similar operations. The Storage Management Application may control the setting of this variable, based on either the setting of a user interface option, or on an assessment of the likely efficiency of using DAR; however, in some cases, it may be necessary to explicitly set `DIRECT=y`.

EXCLUDE

Possible value	Notes
Comma-separated list of files or directories	Specifies files or directories to exclude from a backup. By default, none are excluded. When specifying a file or directory, type either: • A full path name, relative to the top-level directory specified in the backup path. The path name must start with a forward slash (/). An asterisk (*) can be typed at the end as a wildcard character. • A terminal file or directory, which is simply the last element in the path. The name must not contain any forward slash (/) characters, but it may start or end with the wildcard character *. For example: `ENVIRONMENT EXCLUDE "/dir1/tmp*,core,*.o"` This command excludes all files and directories that: • Start with the letters `tmp` in the directory `/dir1` • Are called `core` • End with the characters `.o` The command is case-sensitive if backing up an NFS export but not if backing up a CIFS share.

EXTRACT

Possible value	Notes
y or **n**	The default value **y** causes a recovery operation to extract files from a file list rather than recovering the whole backup.

FILESYSTEM

Possible value	Notes
Name of directory to back up	The Storage Management Application sets the `FILESYSTEM` variable to the name of the path to be backed up.

FUTURE_FILES

Possible value	Notes
y or **n**	Enables back up of files created after the start of the current backup. With NDMP version 2, the inode number that identifies a file can be reused during a backup, thereby causing the backup to fail. By default, therefore, only files created before the start of the backup are backed up. To override this behavior, set `FUTURE_FILES=y`.

HIST

Possible value	Notes
y or **n**	The default value **y** causes file history information to be sent to the storage management application. This enables the display and recovery of the contents of a backup.

LEVEL

Possible value	Notes
0 – 9, or **i**	The default value is **0** (full backup). If the value is set to **4**, an incremental backup is taken based on the most recent previous backup of the same `FILESYSTEM` with level **0**, **1**, **2**, or **3**. If the value is set to **i**, an incremental backup is taken based on the most recent previous backup of the same `FILESYSTEM` of any level.

NDMP_BLUEARC_AWAIT_IDLE

Possible value	Notes
y or **n** (Default **y**)	By default, the data management engine imposes an interlock to prevent *NDMP backups* and *ADC copies* from a replication destination while a replication copy is actively writing data. This is intended for installations that replicate to a particular volume, then back up from that same volume. However, as the lock is held at a volume level, it may be desirable in the case of directory-level replication to override this action. To make use of this replication interlock, specify **y** on both the replication that is intended to do the waiting and the replication that is waited upon.

NDMP_BLUEARC_BB_COMPATIBLE

Possible value	Notes
y or n	Used to request that the server produces backups compatible with legacy Si7500/8000 series servers.

NDMP_BLUEARC_EMBEDDED_HARDLINKS

Possible value	Notes
y or **n**	Used to enable or disable inline hard linked file support. Set the value to y to enable or n to disable. For backups, inline hard linked file support is set to **n** (disabled) by default, but for multi-stream operations, such as replications and ADC copies between servers, the default is overridden and inline hard linked file support is enabled. By default, replication and ADC copy operations use multiple data streams, so for those operations, inline hard linked file support is used by default. When enabled, inline hard linked file support causes NDMP to back up hard linked files with both file data and file metadata inline (in a single data stream), which reduces the amount of memory the server needs to manage the data. Set to **n** to disable inline hard linked file support, which causes file metadata and file data to be sent in two data streams. Disabling inline hard link file support maintains backup compatibility with older systems or releases. Inline hard linked file support may not be enabled using the `ndmp-option` command. Rather, the command used to invoke NDMP must request inline hard linked file support.  **Note:** Existing programs that can read NDMP data streams for releases prior to Release 6.1 will not be able to read backups or recover from backups created using inline hard linked file support. If a replication fails part way through, it will not be possible to restart replication if the server is downgraded to a release prior to Release 6.1.

Using this Option with Replications and ADC Copies

When multi-streamed replication or ADC copy operations are started, this option is enabled. Starting in release 6.1, replication and ADC copy operations are multi-streamed by default, meaning that this option will be enabled by default for those operations.

Using this Option with Backups

When backing up a file system:

- When the embedded hard link option is enabled, the data for each hard linked file is included in the data stream wherever a path to that file is included.

  When enabled, the embedded hard link option increases the amount of data backed up, because multiple copies of the hard linked file data are included. However, it reduces the complexity of managing the backup.

  Also, note that enabling the embedded hard link option reduces the memory requirements needed to keep track of all the hard links.

- When the embedded hard link option is disabled, paths to hard linked files are included without any data in the main part of the backup and a single copy of the hard link file data is included at the end of the backup.

  This reduces the amount of data backed up, because only a single copy of the hard linked file data is included.

Recommendations for Usage with Backups

- If the backup contains many (more than a few hundred thousand) hard linked files, you should enable this option, because it reduces the memory overhead. Note that, where the backup includes many millions of hard linked files, enabling this option may allow the backup to complete where it would not complete if the option is disabled.
- If the backup contains a relatively small number of hard linked files each containing a large amount of data, you should disable the option.
- If there is a chance that the backup may need to be restored on an older version of software, you should disable this option.

NDMP_BLUEARC_EXCLUDE_MIGRATED

Possible value	Notes
y or **n**	Indicates whether backups or replications will include files whose data has been migrated to secondary storage. If set to **y**, the backup or copy will not include files whose data has been migrated to another volume using the Data Migrator facility. The default **n** specifies that migrated files and their data will be backed up as normal files. The backup/copy retains the information that these files had originally been migrated.

NDMP_BLUEARC_EXTERNAL_LINKS

Possible values	Notes
ignore, **recreate_link**, **remigrate**, or **unmigrate**	Controls what happens when a replication operation encounters a cross volume link (a link to a file that has been migrated to an external server). • If set to *ignore*, the replication operation copies only the files on the primary (migrated files are not copied). Use this setting when files have been migrated because they are less useful, so they are not replicated in order to save time. • If set to *recreate_link*, the replication operation copies only the details of the cross volume link. The cross volume link is recreated on the destination if the relevant external migration data path is in place and the migrated file is accessible. Use this setting when the replication is between storage servers or clusters on the same site, and there is a single external migration target server. • If set to *remigrate*, the replication operation copies the file contents but marks the file as having been externally migrated. The destination re-migrates to secondary storage if there is an existing data migration path. Use this setting when the replication is between a main site and a disaster recovery site, where the disaster recovery site includes a similar data migration configuration. This is the default. • If set to *unmigrate*, the replication operation copies all of the files to the primary without re-migrating. Use this setting when data is migrated from expensive storage to more cost effective storage and then replicated to cost effective storage (so no data migration is necessary on the replication target). The default is **remigrate**.

NDMP_BLUEARC_FH_CHARSET

Possible value	Notes
ASCII, **ISO8859** or **UTF8**	Specifies the character set to use when sending file history information to the storage management application. Most file and directory names use characters in the standard ASCII set. If the names of directories and files contain national variant characters outside the ASCII set, they must be encoded when they are sent to the storage management application. Consult the storage management application provider for advice on setting the `NDMP_BLUEARC_FH_CHARSET` variable. UTF-8 is the most wide-ranging option, as it maps the full Unicode character set, covering the alphabets of most of the world's languages. ISO8859 (which can also be specified as ISO8859-1) refers to the 8-bit ISO Latin-1 character set and covers all Western European languages. If path names include characters unavailable in the chosen character set, they will be encoded as a hexadecimal representation of the value of the Unicode character enclosed in caret (^) symbols. For instance if using ASCII, the "£" character will be "^a3^". To avoid confusion, the caret symbol itself is doubled in names in ASCII or ISO8859.

NDMP_BLUEARC_FH_NAMETYPE

Possible value	Notes
UNIX or **NT**	Name type passed by the system to the storage management application in the file history information. NDMP allows files to be described as either UNIX files or NT files. By default, the server's NDMP implementation describes files under an NFS export as UNIX files and those under a CIFS share as NT files. If the storage management application can handle UNIX-style names only, use the `NDMP_BLUEARC_FH_NAMETYPE` variable to request UNIX-style names when backing up a CIFS share.

NDMP_BLUEARC_INCLUDE_ONLY_MIGRATED

Possible value	Notes
y or **n**	Indicates if backups or ADC copies will exclude files whose data has not been migrated to secondary storage. If set to **y**, the backup or copy includes only those files whose data has been migrated to another volume using the Data Migrator facility. The default **n** indicates that files whose data has not been migrated be backed up like normal files.

NDMP_BLUEARC_MULTI_CONNECTION

Possible value	Notes
0 to 30	The number of additional server connections to start during a replication or ADC copy operation to enable multi-stream replication (multiple parallel data streams). The default is four additional connections. When using additional server connections, you may also want to increase the number of read ahead processes (see NDMP_BLUEARC_READAHEAD_PROCESSES, on page 555 or Setting NDMP Performance Options, on page 385 for more information).

NDMP_BLUEARC_OVERWRITE

Possible value	Notes
ALWAYS or **OLDER or NEVER**	Used on recovery to indicate whether an existing file should be replaced by a file of the same name. The default setting is `ALWAYS`, meaning that the backup file will always replace the existing file. `OLDER` means the existing file will be replaced if it is older than the recovered file. If either the existing file or the backup file is actually a directory, this option does not affect behavior. In this case, no overwriting will occur, unless both are directories, in which case the recovered directory contents are merged into the existing directory.

NDMP_BLUEARC_QUOTAS

Possible value	Notes
y or **n**	Default **y** causes NDMP to back up and recover virtual volume and quota information. Set to **n** to disable quota backup or recovery.

NDMP_BLUEARC_READAHEAD_PROCESSES

Possible value	Notes
0 to **10** (In exceptional cases could be increased to 30.)	This variable controls the number of read-ahead processes used when reading directory entries in the backup or copy operation. As each additional readahead process takes up resources, limit the number of additional processes unless it makes a significant difference in performance. The default for this value can be set using `ndmp-option readahead_procs` CLI command. It defaults to **1** if not set explicitly. **0** disables directory readahead, a reasonable option where file sizes are large. **1** to **10** applies to file systems with smaller files; where most of the files are very small (16 KB or less), 10 processes may be optimal. In extreme cases, where most of the deepest-level directories have only one or two files and those files are very small, it may be useful to increase the amount of second-level readahead, using the CLI command `ndmp-option ext_readahead`. If this second level readahead option is set to a higher value such as 10, then setting `NDMP_BLUEARC_READAHEAD_PROCESSES` up to a value of **30** might be advisable.

NDMP_BLUEARC_REMIGRATE

Possible value	Notes
y or **n**	Controls recovery/replication for a file that had been migrated within the original source file system: • *If set to **y**,* the file will be re-migrated on recovery provided the volume or virtual volume has a Data Migrator path to indicate the target volume. • *If set to **n**,* all the files and their data will be written directly to the recovery or replication destination volume.

NDMP_BLUEARC_SNAPSHOT_DELETE

Possible value	Notes
IMMEDIATELY LAST or OBSOLETE	Overrides the **Automated Snapshot Deletion** field in the **Backup > Snapshot Options** page of the Web Manager. `IMMEDIATELY` has the same effect as **Delete snapshot after use**, `LAST` is the same as **Delete snapshot after next backup** and `OBSOLETE` is the same as **Delete snapshot when obsolete**.

NDMP_BLUEARC_SPARSE_DATA

Possible value	Notes
NONE, BASE, or UPDATE	Controls the omission of unset or unchanged data when backing up or transferring files. To enable this feature, a *Block Level Replication license* is required. If `NONE` (default), files are always sent in their entirety, including any unset data areas. If `BASE`, non-initialized areas of files are omitted from the data stream. This is particularly useful for *iSCSI LUNs* or *Database files*, where significant parts of the files may not have been written. `UPDATE` includes the `BASE` setting features but sends only changed blocks of a file with incremental copies or backups. Therefore, this setting is not recommended, as files backed up in this way are only partially included in the backup and require a full correct sequence of backups to be recovered.

NDMP_BLUEARC_SPARSE_FILE

Possible value	Notes
Comma-separated list of files or directories. A list of files similar in format to that specified by the EXCLUDE variable.	Controls sparse file/block-level incremental processing. If set, only files on the list will be considered for sparse transfer.

NDMP_BLUEARC_SPARSE_LIMIT

Possible value	Notes
Numeric value followed by K, M or G signifying Kilobytes, Megabytes or Gigabytes respectively. (For instance, 32M for 32 Megabytes).	Files smaller than the specified value specified will not be considered for sparse transfer. The default value is 32 MB.

NDMP_BLUEARC_TAKE_SNAPSHOT

Possible value	Notes
y or **n**	Overrides the **Automatic Snapshot Creation** backup configuration option.

NDMP_BLUEARC_USE_CHANGE_LIST

Possible value	Notes
y or **n**	Indicates whether incremental backups or replications will use a *changed object list* to direct the search for changed files; otherwise, it will have to search the entire directory tree looking for changed files. When using the changed object list, the search only passes through those directories that contain changed files. Where a relatively small proportion of the file system includes directories containing changed files, the use of *changed object lists* may significantly reduce incremental backup/replication time; however, processing of the *changed object list* itself may take considerable time. Therefore, where file changes exist in many directories, its use is not recommended. The default setting for this option can be set using the CLI command `ndmp-option change_list_incr.`

NDMP_BLUEARC_USE_SNAPSHOT_RULE

Possible value	Notes
Snapshot rule name	Causes NDMP to back up the latest snapshot created under a specified snapshot rule. This can be used to backup a snapshot taken at a specific time; for example, for databases. If set, NDMP does not create or delete snapshots.  **Note:** Following a successful backup, the snapshot should not be deleted until after the operation has completed. In addition, the snapshot should be kept around long enough to support incremental backups.

TYPE

Possible value	Notes
dump or **tar**	Use **dump** in favor of **tar**. The two backup types use exactly the same environment variables and produce the same backup data on tape. However, the format of the information sent to the storage management application differs: • **dump** produces NDMP **add directory entry** and **add node** file history information; • **tar** produces NDMP **add path** file history information.

UPDATE

Possible value	Notes
y or **n**	Default **y** causes a record of the backup time to be kept. Future incremental backups can be carried out using this backup as a base.

Specifying File Names

A file name is made up of the name of an NFS export or CIFS share, followed by a sub-path. File and directory names in an NFS export are case-sensitive, but those in a CIFS share are not. Here are some examples:

- `/nfsroot/dir1/dir2/file1`

  specifies the file `dir1/dir2/file1` under the NFS export `/nfsroot`.

- `/ntroot/nt_dir` (or `/NTROOT/NT_DIR`)

  specifies the file **nt_dir** under the CIFS share **/ntroot**.

Special prefixes can be used to further define the file or directory path. Where no NFS export or CIFS share is available, the path can be based on the volume name:

```
/__VOLUME__/volname/subpath
```

If the name of an NFS export is the same as that of a CIFS share, a file name clash may occur. This can be avoided by adding a unique prefix to the name. For example, if `/root` is both an export and a share, differentiate between the two as follows:

```
/__EXPORT__/root
```

```
/__SHARE__/root
```

In each case, there are two underscores before and after the prefix keyword, which must be in uppercase characters.

Important Notes

- NDMP does not specify the format of the backup data on tape. The format of data on tape is unique to the NDMP server that created the backup image. As a result, it is not possible to use NDMP backups to exchange data with other types of servers.
- An incremental or differential backup backs up changes made since a previous base backup. When asked to do an incremental or differential backup the NDMP code refers to the record of backups to check for such a base backup to compare against. If there is such a backup, and it was a backup of a snapshot, and that snapshot still exists on the system, then the NDMP code executes a "Comparative Incremental" backup, using the original snapshot to identify changes. If the base backup was not of a snapshot, or its snapshot has been deleted, then the only information the code has is the date/time of the backup, and so a "Date-Based Incremental" backup is done.
- Since the Date-Based incremental backup has no record of the files backed up in the original backup, it cannot identify files that have been deleted in the intervening period. Similarly, if a directory has been moved there is no way of knowing that the contents of the moved directory have changed. Therefore contents of moved directories will not be backed up unless the individual files have themselves changed."
- Adding any new equipment to a FC-AL causes the FC-AL to reset, which in turn can leave any attached tape library in an indeterminate state. Similarly, if a failover takes place during a tape backup operation, the tape status may become unknown. While a backup is running, it is therefore advisable not to:
    - Add devices to the FC-AL loop, or remove them from it.
    - Reboot another node in the cluster.
    - Load any firmware on another node in the cluster.

Taking any of these actions causes the FC-AL to reset.

- In recovery operations the storage management application sends a list of files to recover. If it returns file history information with each file on the list then the list is of practically unlimited length. However, if the storage management application does not include the file history information the list is limited to 1024 names.
- The maximum tape block size supported is 256 KB.