



SGI® Management Center for
InfiniteStorage Administrator's Guide

007-5652-004

COPYRIGHT

© 2010-2011 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

SGI, the SGI cube, the SGI logo, and XFS are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and other countries.

Active Directory, Internet Explorer, Microsoft, and Windows are registered trademarks of Microsoft Corporation. LSI is a trademark and MegaRAID is a registered trademark of LSI Corporation. Fedora, Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. InfiniBand is a registered trademark and service mark of the InfiniBand Trade Association. Firefox is a registered trademarks of the Mozilla Foundation. Kerberos is a trademark of the Massachusetts Institute of Technology. Linux is a registered trademark of Linus Torvalds in the U.S. and other countries. Novell is a registered trademark, and SUSE is a trademark of Novell, Inc. in the United States and other countries. OpenLDAP is a registered trademark of the OpenLDAP Foundation. UNIX is a registered trademark of The Open Group in the United States and other countries. All other trademarks mentioned herein are the property of their respective owners.

New Features in this Guide

This revision contains a simplification of the Setup Wizard, which helps you set the SMC IS administration password and the `root` password at initial use. See Chapter 2, "Initial System Setup" on page 7.

Record of Revision

Version	Description
001	June 2010 Original publication with ISSP 2.1
002	September 2010 Revision with ISSP 2.2
003	January 2011 Revision with ISSP 2.3
004	April 2011 Revision with ISSP 2.4

Contents

About This Guide	xv
Related Publications	xvi
Obtaining Publications	xvi
Conventions	xvii
Reader Comments	xvii
1. Overview	1
System Management and Monitoring	1
User Interface	2
2. Initial System Setup	7
Requirements	7
Accessing the Setup Wizard	7
Using the Setup Wizard to Set Passwords	8
Passwords	9
Completed	9
Customizing Your Installation	9
3. Server Configuration and Management	11
Network Interface Configuration	13
Management Interface	13
Ethernet Network Interfaces	14
InfiniBand Network Interfaces	16
Bonded Network Interfaces	17
Storage Configuration	21
007-5652-004	vii

Device Failure Notification	21
Listing Filesystems	22
iSCSI	22
Creating iSCSI Targets	24
The iSCSI Initiator	27
Miscellaneous iSCSI Management	28
User and Group Configuration	29
Local Users and Groups	29
Quotas	30
User Quotas	30
Group Quotas	31
NFS Configuration	32
Global Options	32
Export Options	33
Use Export Options	34
Use a Custom Definition	35
CIFS Configuration	36
SNMP Configuration	38
Global Configuration	40
System Name	40
Name Service Client	41
Local Files Only	41
Active Directory	41
LDAP	44
NIS	45
DNS and Hostnames	45
Time and Date	46
Licenses	47

Administrator Password	47
Operations	47
Save/Restore Configuration	47
Support Data	48
Performance Data	48
Shutdown	48
4. Performance Monitoring	49
Monitoring Overview	49
Metrics Collected	51
System Summary	52
Alerts	55
Storage Device Failures	55
Resources	55
Disk Space	56
Disk User Quota and Disk Group Quota	56
Disk Throughput and Disk IOPS	57
CPU Utilization	57
Network Throughput	58
Hardware Inventory	59
Services	59
NFS	59
CIFS	62
Versions	63
Clients	63
5. Troubleshooting	65
Forgotten Password or Corrupt Password File	65
The archives Directory is Too Large	66

Contents

Power Outage and iSCSI	66
Manual System Reboot	67
Network Configuration Issues	67
Reporting Problems to SGI	68
Glossary	69
Index	75

Figures

Figure 1-1	SGI Management Center for InfiniteStorage Interface	3
Figure 3-1	Management Screen	12
Figure 3-2	Bonded Network Interfaces	18
Figure 3-3	iSCSI Storage	23
Figure 4-1	Monitoring Screen	50
Figure 4-2	Color-Coding the Direction of Data Flow	52
Figure 4-3	Summary Screen	54

Tables

Table 4-1	CPU Metrics Reported	58
Table 4-2	Statistics Reported by NFS and CIFS Screens	60
Table 4-3	Additional Information Reported by the NFS Screen	60
Table 4-4	NFS Operation Classes	61
Table 4-5	CIFS Operation Classes	62
Table 4-6	Additional Information Reported by the CIFS Screen	63

About This Guide

This manual describes the operation of SGI® Management Center for InfiniteStorage (SMC IS). It discusses the following:

- Chapter 1, "Overview" on page 1, describes the tasks you can accomplish with SMC IS and introduces the interface
- Chapter 2, "Initial System Setup" on page 7, describes how to use the Setup Wizard to set passwords
- Chapter 3, "Server Configuration and Management" on page 11, describes how to configure the various components of your system and perform general system administration
- Chapter 4, "Performance Monitoring" on page 49, describes the current and historical views of the state and the performance of a storage server
- Chapter 5, "Troubleshooting" on page 65, discusses problems that you might encounter and how to resolve them

In addition, this document includes a glossary of terms.

Related Publications

For information about this release, see the SGI InfiniteStorage Software Platform (ISSP) release notes (`README.txt`).

Note: The external websites referred to in this guide were correct at the time of publication, but are subject to change.

Obtaining Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, man pages, and other information.
- You can also view man pages by typing `man <title>` on a command line.
- The `/docs` directory on the ISSP DVD or in the Supportfolio download directory contains the following:
 - The ISSP release note: `/docs/README.txt`
 - Other release notes: `/docs/README_NAME.txt`
 - The manuals provided with ISSP
 - A complete list of the packages and their location on the media:
`/docs/RPMS.txt`
 - The packages and their respective licenses: `/docs/PACKAGE_LICENSES.txt`
- The ISSP release notes and manuals are installed on the system as part of the `sgi-isspdocs` RPM into the following location:
`/usr/share/doc/packages/sgi-issp-ISSPVERSION/TITLE`

Conventions

The following conventions are used throughout this publication:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	Bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
Menu item	Bold font indicates a menu item or button in the graphical user interface (GUI).
...	Ellipses indicate that a preceding element can be repeated.
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.

Reader Comments

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

Overview

This chapter discusses the following:

- "System Management and Monitoring" on page 1
- "User Interface" on page 2

System Management and Monitoring

SGI® Management Center for InfiniteStorage (SMC IS) is a web-based interface that lets you configure, manage, and monitor a storage server solution. You can use SMC IS to do the following:

- Perform initial system configuration using the Setup Wizard.
- Manage and install SGI software licenses.
- Configure the system components.
- Perform general system administration tasks.
- Monitor the state and performance of the storage server, including the following:
 - Storage device failures (LSI™ MegaRAID® platform only)
 - CPU utilization
 - Disk utilization
 - Network throughput
 - Services
- Review historical data describing the state and performance of the storage server.
- View connected clients and determine how each of these contribute to the current workload.
- Configure iSCSI targets from block devices and logical volumes.
- Configure NFS exports and CIFS shares from existing filesystems.

Note: Reverse lookup for NFS clients must be properly configured in the DNS server because the NFS server will always try to do a reverse lookup on client IP addresses. Improper configuration will cause delays.

User Interface

To access the SGI Manager Center for InfiniteStorage (SMC IS) features, click one of the menu options displayed across the top of the screen. As you page through SMC IS, your location is shown below the menu options. You can also click an item in this path to directly access that location. For example, Figure 1-1 shows the screen you would see if you selected the **CPU Utilization** item from the **Resources** category on the **Monitoring** menu page. The menu path is shown in this guide in the following format:

Monitoring
 > **Resources**
 > **CPU Utilization**

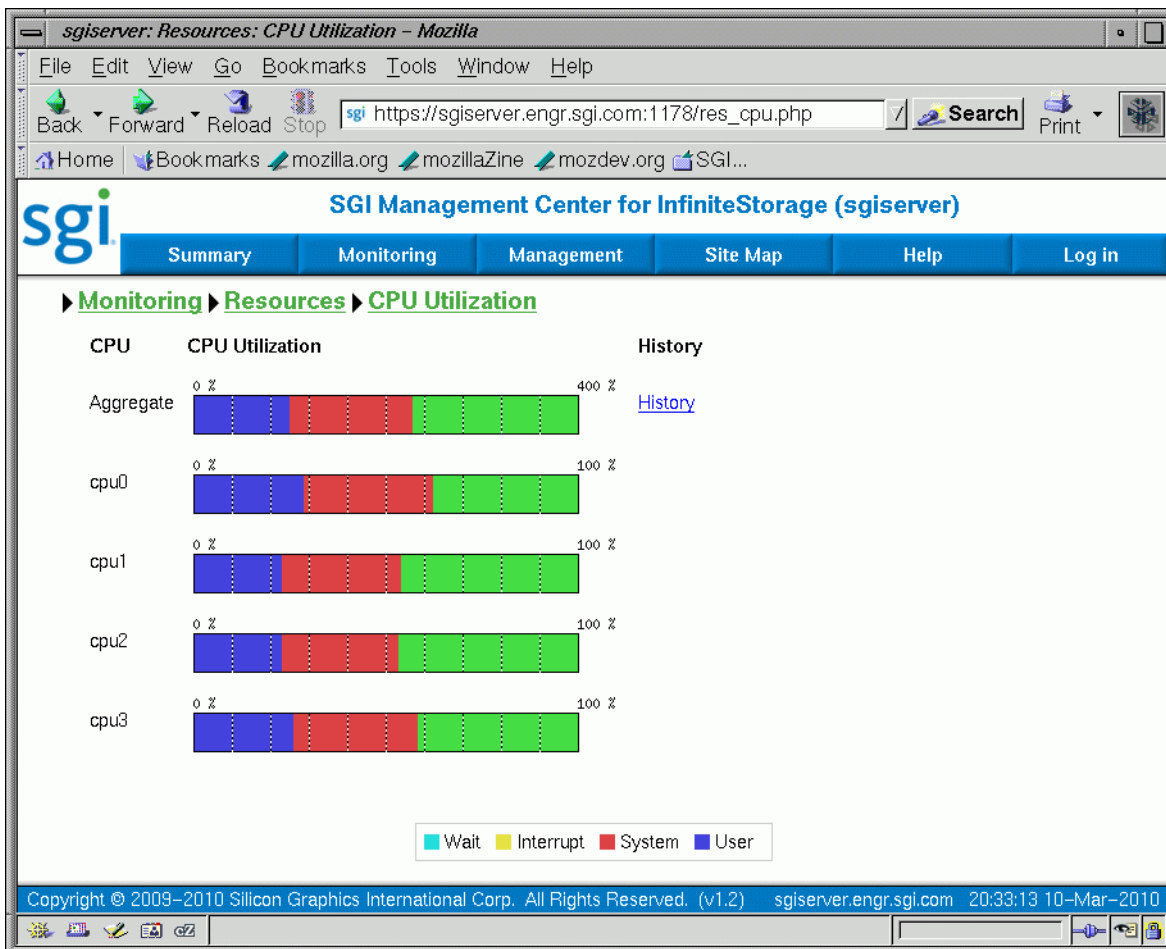


Figure 1-1 SGI Management Center for InfiniteStorage Interface

The menu options are as follows:

Summary

Displays a graphic summary of system utilization, number of alerts, CPU usage, disk space, disk throughput, network throughput, current clients, and uptime. See "System Summary" on page 52.

Monitoring

Lets you monitor features in the following categories:

- **Alerts** displays messages from the system logs. See "Alerts" on page 55.
- **Resources** groups a list of system resources that SMC IS monitors. Select a resource (such as **Disk Space**) to display its status. See "Resources" on page 55.
- **Services** groups a list of services provided by the storage server. Select a service (such as **NFS**) to display its status. You can also display the current versions of installed software. See "Services" on page 59.
- **Clients** displays various I/O criteria by which to display information about the storage server's clients. See "Clients" on page 63.

Management

Lets you perform tasks in the following categories:

- **Resources** groups a list of system resources that you can configure. Select a resource (such as **Network Interfaces**). See:
 - "Network Interface Configuration" on page 13
 - "Storage Configuration" on page 21
 - "User and Group Configuration" on page 29

- **Services** groups a list of services that you can configure using SMC IS. See:
 - "NFS Configuration" on page 32
 - "CIFS Configuration" on page 36
 - "SNMP Configuration" on page 38
- **Global Configuration** groups a list of options for various general system administration tasks. See "Global Configuration" on page 40.
- **Operations** lets you save and restore the system configuration files that you create with SMC IS, gather support and performance data, and shut down or reboot the system. See "Operations" on page 47.

Site Map	Displays an index of direct links to each screen that SMC IS displays.
Help	Displays this guide, the release notes, and legal information about SMC IS.
Log In	Displays the management log-in screen, in which you enter the administration password that enables you to make changes and use the Management screens. (No password is required to use the Monitoring screens.) You must also enable cookies.
Log Out	Allows you to exit from the management function but still access monitoring functions. (After you have logged in, the menu selection changes to Log Out .)

Initial System Setup

This chapter discusses initial system setup for the SGI Management Center for InfiniteStorage (SMC IS):

- "Requirements" on page 7
- "Accessing the Setup Wizard" on page 7
- "Using the Setup Wizard to Set Passwords" on page 8
- "Customizing Your Installation" on page 9

Requirements

Before running the Setup Wizard, you must ensure that the hardware setup instructions have been completed and verified and that the machine has been powered up. For information on system hardware setup, see your system's *Quick Start Guide*.

SMC IS requires the following:

- JavaScript™
- One of the following browsers:
 - Firefox® 2.n through Firefox 3.6 (Firefox is the preferred browser)
 - Internet Explorer® 7.n (7.0 or later) and Internet Explorer 8

Accessing the Setup Wizard

To access the Setup Wizard, do the following:

1. Launch a web browser to the following URL, where *YOUR_SERVER* is the hostname or IP address of your system:

`https://YOUR_SERVER:1178`

Note: If you have just powered the system on, you may need to allow five minutes for the system to boot up before the Setup Wizard will be accessible.

2. Accept the security certificate. In Firefox, this will require a series of steps to add an exception and get the certificate.

Note: SMC IS generates its own SSL certificates, rather than having the SSL certificates signed by a commercial certificate authority. Therefore, the certificate warning is safe to ignore.

The first time you access URL above, the web browser presents the Setup Wizard. After you have completed initial system configuration with the Setup Wizard, the web browser presents the summary screen from which you can access all of the product features.

When using the Setup Wizard, you may see warning or error messages when you click **Next** after filling in the fields. This happens when the system detects a problem in what you have configured. When a warning message appears, the system will still proceed to the next screen. When an error message appears, the system remains on the current screen.

Using the Setup Wizard to Set Passwords

On initial setup, the Setup Wizard leads you through the following screens:

- "Passwords"
- "Completed"

Passwords

The **Passwords** screen displays two sets of text boxes, allowing you to enter and confirm the following:

- SMC IS administration password. This is the password you must enter in order to perform web-based server configuration and management. The password is not required to view the system monitoring screens.
- Command-line configuration password. This is the `root` password for the system

To change the administrator password later on, see "Administrator Password" on page 47.

Completed

The **Completed** screen indicates that the configuration information you entered in the Startup Wizard has been applied. The browser should automatically redirect to the summary screen after a short pause.

Customizing Your Installation

After completing the initial setup, you may customize your system through configuration procedures provided by SMC IS for the following aspects:

- Configuring network interfaces in addition to the management interface. For information on configuring network interfaces, see "Network Interface Configuration" on page 13.
- Configuring *bonded interfaces*, which are virtual network interfaces that consist of real interfaces working in tandem. A virtual interface can provide the aggregated bandwidth of all of the interfaces that you used to create it. For information, see "Bonded Network Interfaces" on page 17.
- Configuring local users and groups, as described in "User and Group Configuration" on page 29.
- Modifying the `/etc/hosts` file. For information on `/etc/hosts` files, see "DNS and Hostnames" on page 45.
- Configuring authentication services. For information on configuring Active Directory[®], LDAP, or NIS for the system, see "Name Service Client" on page 41.

- Setting the time directly, as described in "Time and Date" on page 46.

Server Configuration and Management

This chapter describes how to use SGI Management Center for InfiniteStorage (SMC IS) to configure the various components of your system and perform general system administration:

- "Network Interface Configuration" on page 13 describes how to configure and modify the network interfaces for the system
- "Storage Configuration" on page 21 describes how to list filesystems and iSCSI targets
- "User and Group Configuration" on page 29 describes how to configure a name service client, local users, local groups, and user and group quotas
- "NFS Configuration" on page 32 describes how to configure NFS to share filesystems
- "CIFS Configuration" on page 36 describes how to configure CIFS to share filesystems
- "SNMP Configuration" on page 38 describes how to configure basic Simple Network Management Protocol (SNMP)
- "Global Configuration" on page 40 describes how to perform various general administration functions
- "Operations" on page 47 describes how to save changes to the configuration files and restore them, how to gather support and performance data, and how to shut down the system

Figure 3-1 shows the top level **Management** screen.

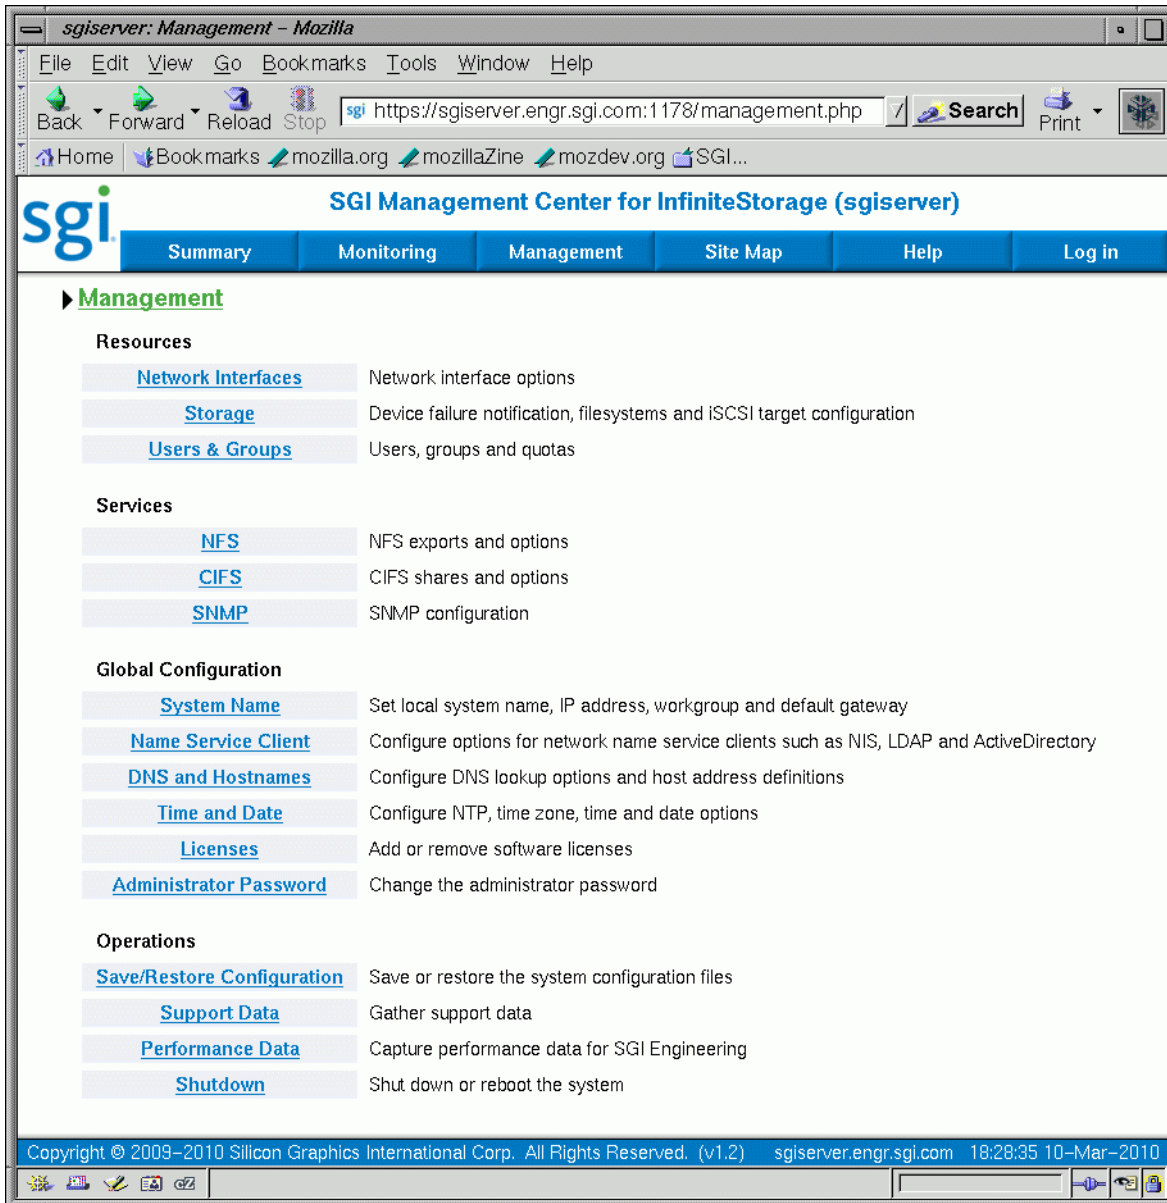


Figure 3-1 Management Screen

Network Interface Configuration

You can use SGI Management Center for InfiniteStorage (SMC IS) to configure and modify the network interfaces for the system. When configuring the system, you must consider the difference between the management interface and the remainder of the interfaces in the system.

The management interface is the first interface in the machine (`eth0`), which is dedicated for use by SMC IS. The remainder of the interfaces in the system are used for fileserving.

You can configure these ports as individual standalone ports or you can group these ports together into a *bonded network interface*.

Bonding interfaces together gives you the aggregated bandwidth for multiple clients of all of the interfaces that constitute the bonded interface. For most systems, this can significantly increase performance over a system in which all of the interfaces are configured as individual network ports.

For further information, see:

- "Management Interface" on page 13
- "Ethernet Network Interfaces" on page 14
- "InfiniBand Network Interfaces" on page 16
- "Bonded Network Interfaces" on page 17



Caution: Ensure that the hardware settings are correct before you configure the network interfaces. For information on hardware setting, see the *Quick Start Guide* for your system.

Management Interface

When the system is shipped from the factory, the management interface has a preconfigured IP address. When using the Setup Wizard, you connect a laptop to that interface in order to perform the initial setup tasks. For information on the Setup Wizard, see Chapter 2, "Initial System Setup" on page 7.

The management interface is always configured as an individual network interface and cannot be part of a bonded interface.

You can modify the management interface by selecting `eth0` from the following screen:

Management
 > **Resources**
 > **Network Interfaces**
 > **Modify**

For information on the network configuration parameters you can modify, see "Ethernet Network Interfaces" on page 14.



Caution: If you configure an incorrect IP address for the management interface, you can make SGI Management Center for InfiniteStorage (SMC IS) inaccessible.

Ethernet Network Interfaces

To see the available Ethernet network interfaces and change their parameters, select the following:

Management
 > **Resources**
 > **Network Interfaces**
 > **Modify**

To modify an interface, select it. You can change the following fields:

Enabled?	Enables the interface. You cannot disable the management interface.
Speed	Displays the port speed of the Ethernet card, which is usually Autonegotiate .
Duplex	Displays the duplex of the Ethernet connection, which is usually Autonegotiate .
Automatic discovery by DHCP	Specifies that dynamic host configuration protocol (DHCP) will be used to configure the Ethernet interface. (Another system must be the DHCP server.)
Static	Specifies that a particular IP address is required for the network interface. If you select this, you must provide the IP address and subnet mask.

Dedicated

Specifies the local and remote IP address for a dedicated network connection between the storage server and another host, for example a dedicated VLAN network or single point-to-point network cable.

A *dedicated network interface* is an interface, such as `eth2`, that has been configured to use a point-to-point connection with a single remote host. All network traffic to and from that server will go via the local dedicated network interface and no other traffic will appear on that interface.

Dedicated network interfaces can be useful when there may be a large amount of network traffic to a specific host and you wish to prevent interference with other network traffic to other hosts.

Note: Dedicated interfaces are an advanced option that may require configuration changes to the network infrastructure and on the remote host. You should only use dedicated interfaces if they are specifically required.

InfiniBand Network Interfaces

To see the available InfiniBand network interfaces and change their parameters, select the following:

Management
 > **Resources**
 > **Network Interfaces**
 > **Modify**

To modify an interface, select it. You can change the following fields:

Enabled	Enables the interface.
Speed	Displays the port speed of the InfiniBand card, which is usually Autonegotiate .
Duplex	Displays the duplex of the InfiniBand connection, which is usually Autonegotiate .
Automatic discovery by DHCP	Specifies that dynamic host configuration protocol (DHCP) will be used to configure the Ethernet interface. (Another system must be the DHCP server.)
Static	Specifies that a particular IP address is required for the network interface. If you select this, you must provide the IP address and subnet mask.
Dedicated	Specifies the local and remote IP address for a dedicated network connection between the storage server and another host, for example a dedicated VLAN network or single point-to-point network cable.

A *dedicated network interface* is an interface that has been configured to use a point-to-point connection with a single remote host. All network traffic to and from that server will go via the local dedicated network interface and no other traffic will appear on that interface.

Dedicated network interfaces can be useful when there may be a large amount of network traffic to a specific host and you wish to prevent interference with other network traffic to other hosts.

Note: Dedicated interfaces are an advanced option that may require configuration changes to the network infrastructure and the remote host. You should only use dedicated interfaces if they are specifically required.

Bonded Network Interfaces

A bonded interface is a virtual network interface that consists of real interfaces working in tandem. You use bonded interfaces on NAS systems to increase bandwidth to NFS and CIFS clients.

A virtual interface can provide the aggregated bandwidth of all of the interfaces that you used to create it.

Note: Any single client can achieve the bandwidth of only a single interface at a time. A bonded interface increases the aggregate bandwidth for multiple clients.

For example, if you have three interfaces each with a bandwidth of 10, the aggregate bandwidth is 30. For an individual client, however, the maximum bandwidth remains 10. When additional clients access the bonded interface, the clients are assigned to the subinterfaces, and up to three clients can use a bandwidth of 10 at the same time. Thus multiple clients accessing the system increase the aggregate bandwidth, improving the performance to a maximum bandwidth of 30.

For example, Figure 3-2 shows a configuration in which all clients connect to a single IP address (192.168.0.3). The switch is responsible for sharing the load across four bonded interfaces (eth1-eth4). Therefore, four times as many clients can communicate with the same server without a loss in overall performance.

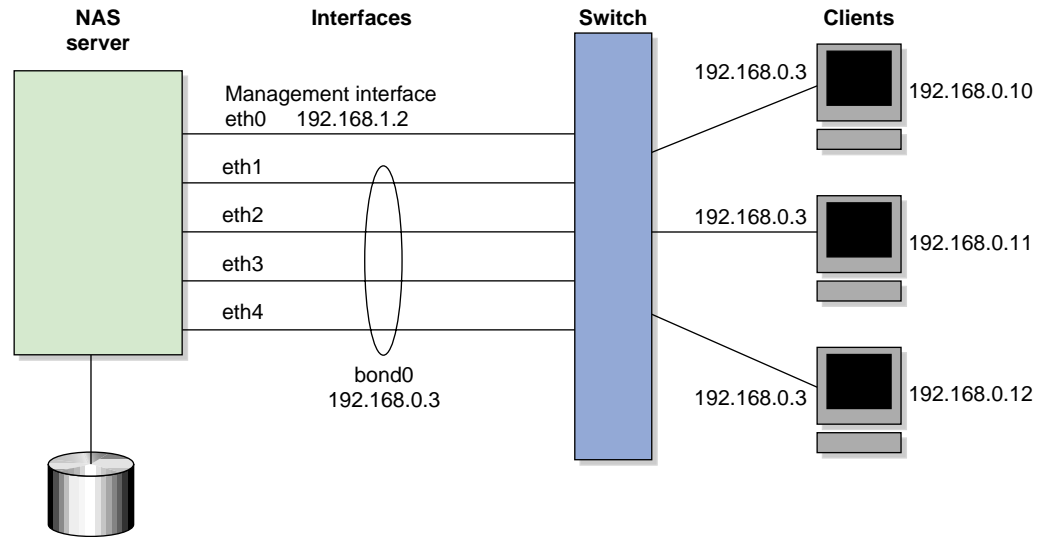


Figure 3-2 Bonded Network Interfaces

Output load balancing controls how the server chooses which subinterface will send replies. *Input load balancing* controls how clients are assigned to subinterfaces, and how and when clients are moved from one subinterface to another. Load balancing happens on a per-packet basis. When a client sends a packet, it traverses a switch, which determines at which subinterface the packet arrives. Input load balancing ensures that each client arrives at a different subinterface. The clients see only one interface because the balancing is done by the system.

In addition to configuring a bonded interface in SGI Management Center for InfiniteStorage (SMC IS), you must configure the ports on the switch so that they use either static trunking or 802.3ad dynamic trunking. For more information, refer to the user manual for your switch.

To create a bonded interface, select the following:

- Management**
 - > **Resources**
 - > **Network Interfaces**
 - > **Create a bonded interface**

The available interfaces are displayed for selection.

When you configure a bonded interface, you specify the following:

Available interfaces

Specifies the interfaces to be used.

Bonding mode

Selects a bonding mode that governs the relation of the subinterfaces to a switch and defines the protocol that is used for assigning network switch ports to a bonded interface:

- **Dynamic 802.3ad** uses the 802.3ad protocol to communicate with the switch and automatically bond the appropriate switch ports together. You may need to configure your switch to enable the 802.3ad protocol on a range of switch ports or for the switch as a whole.
- **Static** requires that the switch be manually configured to bond specific switch ports together.

Your choice depends upon what your switch supports:

- If your switch supports the 802.3ad protocol, choose dynamic bonding.
- If your switch only supports manually grouping ports together in a bond, choose static bonding.
- If your switch does not support any bonding, you must configure all your network interfaces as separate individual interfaces.

Output Load Balancing

Specifies how the server chooses which subinterface will send replies:

- **Layer 3 (IP header)** specifies that the server and client are on different subnets.
- **Layer 2 (MAC address)** specifies that all packets sent to the clients use separate MAC addresses. This option is more efficient than **Layer 3 (IP header)**. Use this option only if the clients are in the same broadcast domain as the server.

Note: Do not select this option if the switch immediately upstream of the server is acting as a router rather than a switch (that is, making packet routing decisions at Layer 3 rather than Layer 2) or if the clients are in a different subnet and you have another router between the server and clients.

IP address

Specifies the IP address of the new bonded interface. The IP address for a bonded interface must be configured statically. SMC IS does not support DHCP and dedicated IP addresses for bonded interfaces.

Subnet mask

Specifies the subnet mask of the new bonded interface. All configured network interfaces should be on different subnets.

Click **Apply Changes** to create the bond.

Storage Configuration

This section discusses the following:

- "Device Failure Notification" on page 21
- "Listing Filesystems" on page 22
- "iSCSI" on page 22

These features are available under the following menu selection:

Management
 > **Resources**
 > **Storage**

Device Failure Notification

To configure notification of failed devices, use the **Failure Notification** option:

Management
 > **Resources**
 > **Storage**
 > **Devices**
 > **Failure Notification**

For each available platform, enter the interval at which to scan for for device failures (0-59 minutes, 0 to disable), enter one or more email addresses to send notifications to, and check the **SNMP trap** box to enable that as desired.

Each SNMP trap is sent using the `SGI-SSMC-SMI::devstatDriveFailure` object identifier. To allow the SNMP monitoring application to fully parse the object, see the procedure in "SNMP Configuration" on page 38.

For the LSI MegaRAID platform, any drive whose state is `Failed`, `Unconfigured`, `Bad`, or `Offline` is deemed to have failed.

Listing Filesystems

To display a brief description of the available local filesystems, use the **List** option:

```
Management
  > Resources
    > Storage
      > Filesystems
        > List
```

For filesystems that are mounted, storage capacity and usage are displayed, along with related NFS exports and CIFS shares.

Filesystems that are listed in `/etc/fstab` will be presented with mount and unmount check boxes (except for the root filesystem). Those selected will be put into effect when the **Apply** button is clicked.

iSCSI

Internet Small Computer Systems Interface (iSCSI) is a protocol that is used to transport SCSI commands across a TCP/IP network. This allows a system to access storage across a network just as if the system were accessing a local physical disk. To a client accessing the iSCSI storage, the storage appears as a disk drive would appear if the storage were local.

In an iSCSI network, the client accessing the storage is called the *initiator* and runs iSCSI Initiator software. The remote storage that the client accesses is called the *target*, which is what appears to the initiator as a disk drive.

A common application of an iSCSI network is to configure an Exchange Server as an iSCSI initiator that uses an iSCSI target as its mail store.

Figure 3-3 illustrates iSCSI storage. Each client (initiator) is configured to connect to a specific iSCSI target (an area allocated in the RAID iSCSI storage pool), and views this target as if it were a local disk. The lines in Figure 3-3 indicate data flow.




Figure 3-3 iSCSI Storage

You can use SGI Management Center for InfiniteStorage (SMC IS) to create iSCSI targets on the RAID storage. An iSCSI initiator will be able to connect to the system and access those targets, format them, and use the targets as it would use a disk drive.

You cannot configure SMC IS itself as an initiator, and you cannot re-export iSCSI targets with NFS or CIFS. In addition, you cannot export existing filesystems that you have created with SMC IS as iSCSI targets; you can create filesystems and configure them to be exported by NFS or CIFS, but you must configure iSCSI targets separately on the RAID device.

SMC IS supports the `iscsitarget` and `tgt` packages in SUSE® Linux® Enterprise Server (SLES) 11 for creating targets. They are mutually exclusive; installing `tgt` removes `iscsitarget`.

Note: Due to the nature of iSCSI as a block-level protocol (as distinct from file-level protocols such as NFS and CIFS), particular care must be taken in the event of a system crash, power failure, or extended network outage. See "Power Outage and iSCSI" on page 66.

This section discusses the following:

- "Creating iSCSI Targets" on page 24
- "The iSCSI Initiator" on page 27
- "Miscellaneous iSCSI Management" on page 28

Creating iSCSI Targets

Perform the following steps to create an iSCSI target:

1. Select the **Create Target** option:

```
Management
  > Resources
    > Storage
      > iSCSI
        > Create Target
```

2. The **Target Name** screen lets you specify the target information. Enter the domain and optional identifier for the iSCSI name and the LUNs for the target in the following fields:

Domain	Specifies an iSCSI qualified name (which is a unique name that starts with <code>iqn</code>), then a year and month, then an internet domain name in reverse order. A default name appears based on the current system configuration. If in doubt, leave this field as is.
---------------	---

Identifier	Specifies a string that will be used to uniquely identify the target. If you create only one target, this is optional. If you create more than one target, each must have a unique identifier. By default, a unique target identifier is provided for you.
-------------------	--

LUNs Specifies logical units (LUNs) to be used for the target. Enter the path to a block device to add to the list. Applicable block devices and logical volumes will be listed in pulldown menus if they are available. Use the buttons on the right to reorder or remove entries in the list.

Click **Next**.

3. The **Target Options** screen defines access to the target. You must specify at least one authentication option:

Note: If more than one initiator were to write to the same target at the same time, there is a high risk of data loss. By using one or more authentication options, you ensure that only one client (initiator) can access an individual target at a time.

- Authentication:

Initiator IP Address Specifies the IP addresses of the initiators that will be allowed access to this target

- Challenge Handshake Authentication Protocol (CHAP) authentication, in which the initiator will supply the following information to the target:

Target Username Specifies the username that the initiator must supply to connect to the target using CHAP authentication. (This is not the username with which you logged in to SMC IS; it is specific to the iSCSI target that you are defining.)

Target CHAP Secret Specifies the password that the initiator must supply to connect to the target using CHAP authentication. It must be in the range from 12 through 16 characters. (This is not the password with which you logged in to SMC IS; it is specific to the iSCSI target you are defining.)

Re-enter Target CHAP Secret Verifies the CHAP secret.

- Mutual CHAP authentication, in which the target will supply the following information to the initiator:

Mutual Username Specifies the target username for mutual CHAP authentication. With mutual CHAP authentication, after the initiator supplies a username, the target must supply a username and password back to the initiator. If you leave the **Mutual Username** field blank, it defaults to the target username.

The mutual name is usually ignored by initiators, which only care about the mutual secret. When the client connects to a target, the iSCSI initiator software verifies that the mutual secret specified in SMC IS matches the secret specified in the initiator.

Mutual CHAP Secret Specifies the mutual CHAP secret.

Note: This secret should be different from the target CHAP secret.

Re-enter Mutual CHAP Secret Verifies the mutual CHAP secret.

You must enter the CHAP username and secret specified on this screen in the iSCSI initiator software on the client in order for the initiator to be able to authenticate with and connect to the target. For a Windows client, this is the username and secret you enter in the Microsoft™ iSCSI Initiator program.

4. The **Confirm** screen summarizes the target options you have selected. Click **Next** to confirm your choices and create the iSCSI target.
5. The **Finished** screen indicates that the iSCSI target has been created. Select **Done**.

After you have created iSCSI targets, select the following to see what initiators are connected to what targets:

Monitoring
 > **Clients**
 > **iSCSI**

The iSCSI Initiator

SMC IS lets you configure iSCSI targets for use by an iSCSI initiator, such as the Microsoft iSCSI Software Initiator or the iSCSI initiator included with various Linux[®] and UNIX[®] distributions.

After you have created an iSCSI target, you must configure the initiator on the client system that will connect to the target. You must specify the following:

- Hostname of the storage server
- Target identifier
- Any CHAP authentication details you configured when creating the target (for specific instructions, see the documentation supplied with your iSCSI initiator)

After the iSCSI initiator has connected to the target, the target will appear as a disk drive on the client system and can then be formatted using the tools supplied with the client operating system.

The following is an example of configuring a Windows client (it assumes that you have already created a target or targets):

1. Download the iSCSI Initiator from Microsoft's web site (<http://www.microsoft.com/>) and install it on the Windows client.
2. Open the **iSCSI Initiator Control Panel** applet.
3. Add the storage server to the list of **Target Portals**.
4. Select the iSCSI target to connect to from the **Targets** list and click **Log On**.
5. Specify CHAP authentication details in the **Advanced** settings.

6. Use the following tool to partition and format the target and assign a drive letter:

Start Menu
 > **Administrative Tools**
 > **Computer Management**
 > **Disk Management**

Miscellaneous iSCSI Management

The **iSCSI** menu also provides the following management options:

List Targets	Lists the existing iSCSI targets.
Create Target	Creates an iSCSI target.
Modify Target	Modifies the authentication settings you defined on the Target Options screen when you created an iSCSI target.
Destroy Target	Destroys an existing iSCSI target.
Stop/Start	Stops or starts the iSCSI service. If you are backing up the system, taking the iSCSI service offline ensures that the data is in a consistent state.

User and Group Configuration

SGI Management Center for InfiniteStorage (SMC IS) lets you configure the following:

- "Local Users and Groups" on page 29
- "Quotas" on page 30

Local Users and Groups

SGI Management Center for InfiniteStorage (SMC IS) can create and add user and group accounts to access the storage server locally. This is a local database only; these users and groups do not interact with the users and groups provided by the name server. If you search the site directory and do not find the user or group data you are looking for, the system searches this local database. The local user accounts will be used for authentication for CIFS shares if you are not using LDAP or Active Directory authentication.



Caution: If you create a local user and subsequently add that user in the sitewide directory, access problems may result. For example, if you create local user `Fred` with a UID of 26, `Fred` will be able to create local files. But if you subsequently add a user `Fred` on a sitewide name services directory with a different UID, user `Fred` will be unable to access those local files because the system will use the sitewide name and UID first.

If you are using LDAP or Active Directory as a name service client, a user must be present in LDAP or Active Directory and you will not be able to authenticate local users and groups. In this case, adding local users and groups may be useful for ID mapping, but authentication does not use the local password files.

When you select the **Import** option for either **Local Users** or **Local Groups**, you can choose among the following actions:

- Merge the imported new users or groups with the current list, ignoring any accounts or groups with the same name. (That is, if there is an existing user or group, keep it rather than the new imported user or group.)
- Merge the imported new users and groups with the current list, overwriting any exists in accounts or groups of the same name. (That is, if there is an existing user or group, replace it with the new imported user or group.)

- Replace all current unrestricted users or groups with the new imported users or groups.

Accounts with a UID or GID of less than 1000 are considered restricted and are not imported or replaced.

If you use a *shadow file*, which is a file that is protected from all access by non-root users and stores the encrypted passwords, then you can use the **Import Users** screen to import this file as well as the password file itself.

Quotas

SGI Management Center for InfiniteStorage (SMC IS) will create new filesystems with both user and group quotas enabled by default.

This section discusses the following:

- "User Quotas" on page 30
- "Group Quotas" on page 31

User Quotas

You can use the following screen to specify the user for whom you want to modify quotas:

```
Management
  > Resources
    > Users & Groups
      > User Quotas
```

Enter the name of the user and click **Submit**. (To modify the default for user quotas, leave the field blank.) The following screen displays the current amount of disk space that can be used (disk limits, in KiB) and the number of files that can be owned (file limits):

- The **soft limit** is the number of 1-KiB blocks or the number of files that the user is expected to remain below. If a user hits the soft limit, a grace period of 7 days will begin. If the user still exceeds the soft limit after the grace period expires, the user will not be able to write to that filesystem until he or she removes files in order to reduce usage.

- The **hard** limit is the number of 1-KiB blocks or the number of files that the user cannot exceed. If a user's usage reaches the hard limit, he or she will be immediately unable to write any more data.

Note: The administrator can set quotas for the `root` user. However, instead of enforcing these quotas against the `root` user specifically, they will apply to all users that do not have their own quotas set. In other words, setting quotas for the `root` user will set the default quotas for all normal users and groups. (The actual `root` user is exempt from quota limits.)

Group Quotas

You can use the following screen to specify the group for which you want to modify quotas:

```
Management
  > Resources
    > Users & Groups
      > Group Quotas
```

Enter the name of the group and click **Submit**. (To modify the default for group quotas, leave the field blank.) The following screen displays the current amount of disk space that can be used (disk limits, in KiB) and the number of files that can be owned (file limits):

- The **soft** limit is the number of 1-KiB blocks or the number of files that the group is expected to remain below. If any user in that group hits the soft limit, a grace period of 7 days will begin. If the user still exceeds the soft limit after the grace period expires, the user will not be able to write to that filesystem until he or she removes files in order to reduce usage.
- The **hard** limit is the number of 1-KiB blocks or the number of files that the group cannot exceed. If the usage for a user in that group reaches the hard limit, he or she will be immediately unable to write any more data.

Note: The administrator can set quotas for the `root` group. However, instead of enforcing these quotas against the `root` group specifically, they will apply to all groups that do not have their own quotas set. In other words, setting quotas for the `root` group will set the default quotas for all normal groups. (The actual `root` user is exempt from quota limits.)

NFS Configuration

To configure filesystems so that they are available for network clients by means of the NFS network protocol, select the following:

Management
 > **Services**
 > **NFS**

This screen displays a link for **Global Options** and all of the filesystems that have been created with SMC IS, whether or not they have been enabled for export.

To specify NFSv4 options, select **Global Options**. To change the export options, select an individual filesystem name or **All Filesystems**. See:

- "Global Options" on page 32
- "Export Options" on page 33

Note: Reverse lookup for NFS clients must be properly configured in the DNS server.

Global Options

The **Global Options** screen lets you specify the following:

Enable NFSv4	Specifies whether NFSv4 is enabled (checked) or not. If enabled, an NFS exported filesystem will be accessible via both NFSv3 and NFSv4. The following fields are only relevant if you have enabled NFSv4.
NFS serving domain	Specifies the serving domain. If NFSv4 is enabled, the mapping of user/group IDs between the client and server requires both to belong to the same NFS serving domain.
Enable Kerberos	Specifies whether Kerberos™ is enabled (checked) or not. Enabling Kerberos forces encrypted authentication between the NFS client and server. Furthermore, the NFS exported filesystems will only be accessible to a Kerberos enabled client via NFSv4. The following fields are only relevant if you have enabled Kerberos.

Note: SGI Management Center for InfiniteStorage (SMC IS) supports Kerberos 5. You must use a mechanism to synchronize the time between all systems.

Realm	Specifies the Kerberos realm in which the NFSv4 server operates.
Domain	Specifies the DNS domain name that corresponds to the realm.
KDC	Specifies the key distribution center (KDC). In most cases, the KDC will be the same system as the Kerberos admin server. However, if the admin server in your Kerberos environment is not used for granting tickets, then set the KDC to the system that grants tickets.
Admin Server	Specifies the server containing the master copy of the realm database.
Keep Existing Keytab	Select this radio button to keep the existing keytab without changes.
Update Keytab	Select this radio button to change the principal user and password for the existing keytab.
Principal	Specifies a user that belongs to the Kerberos server with sufficient privileges to generate a keytab for the NFS server.
Password	Specifies the principal's password.
Upload Keytab	Copies the selected file to <code>/etc/krb5.keytab</code> on the NFS server. Click Browse to see a list of available files.
Verify Keytab	Specifies that the keytab should be verified. This is not supported by Active Directory.

Export Options

You can choose to export or not export a filesystem by clicking the check box. The current export point is shown following the **Directory** label. To export a subdirectory under the export point, enter its path in the text field (a leading “/” may be omitted.)

When you enable a filesystem for export, you can do one of the following:

- "Use Export Options" on page 34
- "Use a Custom Definition" on page 35

After specifying the configuration parameters, click **Apply changes**.

Use Export Options

If you select **Use export options**, you must specify the following:

Read-only	Specifies that the client has access to the filesystem but cannot modify files or create new files.
Asynchronous writes	<p>Specifies whether or not to use asynchronous writes.</p> <p>Data that is written by the client can be buffered on the server before it is written to disk. This allows the client to continue to do other work as the server continues to write the data to the disk.</p> <p>By default, writes are performed synchronously, which ensures that activity on the client is suspended when a write occurs until all outstanding data has been safely stored onto stable storage.</p>
Allow access from unprivileged ports	Allows access for Mac OS X clients or other NFS clients that initiate mounts from port numbers greater than 1024. If there are no such clients on your network, leave this option unchecked.
All hosts	Allows connections from anywhere on a network.
Local subnet	Allows connections from the indicated subnet. You can select any subnet from those that have been defined for the network interfaces.
Kerberos aware clients (krb5)	Allows connections only from those systems that are Kerberos aware (if Kerberos is enabled in "Global Options" on page 32) over NFSv4.
Kerberos with Integrity support aware clients (krb5i).	Allows connections only from those systems that are Kerberos with Integrity support aware (if Kerberos is enabled in "Global Options" on page 32) over NFSv4.

Restrict to hosts

Specifies the set of hosts that are permitted to access the NFS filesystem. You can specify the hosts by hostname or IP address; separate values with a space or tab. For example, you could restrict access to only the hosts on a Class C subnet by specifying something like the following:

```
150.203.5
```

To allow hosts of IP address 150.203.5.* and myhost.mynet.edu.au, specify the following:

```
150.203.5. myhost.mynet.edu.au
```

You can also specify hosts by network/subnet mask pairs and by netgroup names if the system supports netgroups.

To allow hosts that match the network/subnet mask of 150.203.15.0/255.255.255.0, you would specify the following:

```
150.203.15.0/255.255.255.0
```

To allow two hosts, hostA and hostB, specify the following:

```
hostA hostB
```

Note: Access still requires suitable user-level passwords. The localhost address 127.0.0.1 will always be allowed.

Use a Custom Definition

If you select **Use custom definition**, you can enter any NFS export options that are supported in the Linux `/etc/exports` file.

For example, the following entry gives 192.168.10.1 read-write access, but read-only access to all other IP addresses:

```
192.168.10.1(rw) *(ro)
```

Note: There cannot be a space between the IP address and the export option.

For information on the `/etc/exports` file, see the `exports(5)` man page. ¹

CIFS Configuration

To configure filesystems so that they are available for network clients by means of the CIFS network protocol, select the following:

Management
 > **Services**
 > **CIFS**

All of the filesystems created with SGI Management Center for InfiniteStorage (SMC IS) are displayed on this screen, whether or not they have been enabled for sharing. To share a file, select it and click the **Shared?** box. The current share path is shown following the **Directory** label. To share a subdirectory under the share path, enter its path in the text field (a leading “/” may be omitted.)

Specify the following **Share Options**:

Share name	Specifies the name under which the filesystem will appear to a Windows client, as displayed in its Network Neighborhood.
Comment	Specifies an arbitrary string to describe the share.
Read-only	Specifies that the client has access to the filesystem but cannot modify files or create new files.
Allow guest users	Specifies that users can gain access to the CIFS filesystem without authenticating. Uncheck this option to allow connections only to valid users. By default, the CIFS protocol requires a password for authentication. If you are configured as an Active Directory client, then the authentication is distributed. See "Active Directory" on page 41.

¹ You can access man pages from the SGI Technical Publications Library at <http://docs.sgi.com>.

Always synchronize writes

Ensures that write activity on the client is suspended when a write occurs until all outstanding data has been safely stored onto stable storage. If you do not check this box, data that is written by the client can be buffered on the server before it is written to disk. This allows the client to continue to do other writing as the server continues to write the data to the disk. This is the faster write option and is recommended.

Allow symbolic linking outside of the share

Specifies that symbolic links made by NFS users that point outside of the Samba share will be followed.



Caution: This feature is a performance/security tradeoff that is only interesting for sites running both CIFS and NFS from the same filesystem. Allowing linking could be a security risk if, for example, an NFS user created a symbolic link to `/etc/passwd`. However, unchecking the box will cause a decrease in performance.

All hosts

Allows connections from anywhere on a network.

Local subnets

Allows connections from the indicated subnet. You can select one subnet in this field and you must choose it from the available interfaces as set in the **Network Interfaces** screen.

Restrict to hosts

Specifies the set of hosts that are permitted to access the CIFS filesystem. You can specify the hosts by name or IP number; separate values by a space or tab. For example, you could restrict access to only the hosts on a Class C subnet by specifying something like the following:

```
150.203.5
```

To allow hosts of IP address `150.203.5.*` and `myhost.mynet.edu.au`, specify the following:

```
150.203.5 myhost.mynet.edu.au
```

You can also specify hosts by network/subnet mask pairs and by netgroup names if the system supports

netgroups. You can use the `EXCEPT` keyword to limit a wildcard list.

For example, to allow all IP address in `150.203.*.*` except one address (`150.203.6.66`), you would specify the following:

```
150.203. EXCEPT 150.203.6.66
```

To allow hosts that match the network/subnet mask of `150.203.15.0/255.255.255.0`, you would specify the following:

```
150.203.15.0/255.255.255.0
```

To allow two hosts, `hostA` and `hostB`, specify the following:

```
hostA, hostB
```

Note: Access still requires suitable user-level passwords. The localhost address `127.0.0.1` will always be allowed.

After specifying the configuration parameters, select **Apply changes**.

SNMP Configuration

SGI Management Center for InfiniteStorage (SMC IS) lets you configure basic SNMP monitoring support on your storage server. In order to query the SNMP service and receive SNMP traps, you will require an external management station with appropriately configured monitoring software.

To configure the SNMP service, select the following:

```
Management
  > Services
    > SNMP
```

The **SNMP** screen lets you configure the following parameters:

Enable SNMP Enables or disables the SNMP service.

Allow SNMP access from	Specifies the IP address of the Network Monitoring Station (NMS) or the network segment that is allowed to access the SNMP service.
Trap destination	Specifies the IP address of your NMS for receiving default SNMP traps.
Community string	Specifies the SNMP community string to use when sending SNMP traps and when querying the SNMP service. The default is <code>public</code> .
System name	Specifies the system name. This field is automatically set by SMC IS to the hostname of the server. However, you may change this to something more appropriate to your environment.
System location	Specifies the physical location of the storage server (optional).
System contact	Specifies the contact details (such as the name and email address) of one or more persons responsible for administration of the server (optional).
System description	Provides addition descriptive information for identifying the server (optional).

After applying your configuration changes to the SNMP service, you should receive start/stop SNMP v2 traps notifying you that the SNMP service has been restarted.

To allow the SNMP monitoring application to fully parse trap objects, do the following:

1. On the trap destination system, install the `sgi-snmpagent-mibs` package from the SGI Foundation Software media.
2. Copy the following file from the storage server system to the same directory on the trap destination system:

```
/opt/sgi/snmpagents/mibs/sgi-ssmc-smi.mib
```

3. Make the management information bases (MIBs) in `/opt/sgi/snmpagents/mibs` known to the SNMP monitoring application.

Global Configuration

The following sections describe the following aspects of system administration that you can perform with SGI Management Center for InfiniteStorage (SMC IS):

- "System Name" on page 40
- "Name Service Client" on page 41
- "DNS and Hostnames" on page 45
- "Time and Date" on page 46
- "Licenses" on page 47
- "Administrator Password" on page 47

System Name

Use the **System Name** screen to set the following system components:

System name	Specifies the fully qualified domain name (FQDN) for this storage server. The default system name is <code>sgiserver</code> .
<hr/>	
	Note: After changing the system name, the various SGI Management Center for InfiniteStorage (SMC IS) screens will still display the old hostname. SGI recommends that you reboot the system to complete the name change.
<hr/>	
Workgroup	Specifies the NetBIOS workgroup to which the machine should belong. The default is <code>WORKGROUP</code> . If you are not using CIFS, you can ignore this setting.
Default network gateway	Specifies the IP address of the router that this system should use to communicate with machines that are outside of its subnet.
Management IP address	Specifies the IP address of the management interface.
Subnet mask	Specifies the subnet mask of the management interface.

Use DHCP Specifies whether or not to use dynamic host configuration protocol (DHCP).

You can also use the **Network Interfaces** screen for `eth0` to configure or modify the management interface. For information on these options, see "Ethernet Network Interfaces" on page 14.

Name Service Client

The **Name Service Client** screen lets you specify a *name service* (or *directory service*) for the system. A name service is the application that manages the information associated with the network users. For example, it maps user names with user IDs and group names with group IDs. It allows for centralized administration of these management tasks.

You can specify whether you are using local files (if you have no sitewide protocol and names and IDs are kept locally on server), Active Directory services, lightweight directory access protocol (LDAP), or the sitewide network information service (NIS).

Note: When specifying servers on the **Name Service Client** screen, you must use IP addresses rather than hostnames, because the system may require a name service client to determine the IP address from the hostname.

Local Files Only

The **Local Files Only** selection specifies that an external name server will not be used. All user and group name to ID mapping will be done using local users and groups. See "Local Users and Groups" on page 29.

Active Directory

Active Directory is a directory service that implements LDAP in a Windows environment. It provides a hierarchical structure for organizing access to data. CIFS authentication will automatically use the Active Directory service.

Note: The **Active Directory** section is disabled if there are no Active Directory DNS servers specified. See "DNS and Hostnames" on page 45.

The following Active Directory components appear on the **Name Service Client** screen:

Active Directory domain	Specifies the full domain name of the Active Directory. <hr/> Note: If you later change the server hostname on which SGI Management Center for InfiniteStorage (SMC IS) runs, you must rejoin the Active Directory domain because the Active Directory Security Identifier (SID) will be changed. <hr/>
Domain Controller	Specifies a domain controller.
Administrative user	Specifies the user with administrator privileges.
Allow this user to remotely manage CIFS share permissions	Specifies whether or not the Administrative user shown will be able to use the Windows MMC Computer Management GUI to manipulate CIFS share permissions remotely when you join the Active Directory domain.
Password	Specifies the password for the administrator user. For security reasons, the Active Directory password cannot contain the following characters: <code> ; * & \ ' < > ? []</code>
Re-enter password	Verifies the password for the administrative user.
UID/GID Mapping	Lets you manage UNIX user ID (UID) and group ID (GID) mapping on the Active Directory server, using one of the following: <ul style="list-style-type: none">• RFC 2307 (Microsoft Windows Server 2003 R2). In order for this to function correctly:<ul style="list-style-type: none">– The Active Directory domain controller must be running Microsoft Windows Server 2003 R2.– The Identity Management for UNIX service must be installed on the domain controller.– You must use the UNIX Attributes tab in Active Directory user management to set up UIDs and GIDs for all users requiring access to this system.

- **Microsoft Windows Services For UNIX.** In order for this to function correctly:
 - Microsoft Windows Services for UNIX must be installed on the Active Directory domain controller.
 - You must use the **UNIX Attributes** tab in Active Directory user management to set up UIDs and GIDs for all users requiring access to this system.
- **Automatic assignment based on Windows SID.** In this mode, UIDs and GIDs are automatically based on the Windows SID and are set to be in the range 16777216 through 33554431.

Note: This method can only be used within a single Active Directory domain and is incompatible with trusted domains.

- **Automatic assignment in range 10000-20000.** In this mode, UIDs and GIDs in the range 10000 through 20000 will be automatically assigned to Active Directory users on a first-come, first-served basis.

Note: The default is **Automatic assignment based on Windows SID**. For best interoperability, SGI recommends that you choose either **RFC 2307 (Microsoft Windows Server 2003 R2)** or **Microsoft Windows Services For UNIX** when applicable, as appropriate for your environment.



Caution: Depending on your environment, making changes to the UID/GID mapping may result in ownership changes of user files.

LDAP

Lightweight directory access protocol (LDAP) is a networking protocol that organizes access to data in a directory tree structure. Each entry in the tree has a unique identifier called the *distinguished name*.

The default LDAP server IP address is the local host. You will probably need to specify a different IP address.

Fields:

LDAP server	Specifies the IP address of the LDAP server.
Base	Specifies the distinguished name of the base of the subtree you will be searching.
Root binddn	Specifies the distinguished name of the user to whom you are assigning <code>root</code> privileges for administration. This is expressed as a node in the directory tree that refers to a user account.
Password	Specifies the password that will be required to authenticate against the LDAP server. For security reasons, the LDAP password cannot contain the following characters: <code>; * & \ < > ? []</code>
Re-enter password	Verifies the password that will be required to authenticate against the LDAP server.

To use LDAP for CIFS authentication, you must configure the LDAP server to use the RFC2307bis or NIS schema to supply POSIX account information. In addition, you must add a Samba schema to the LDAP database. These schemas specify how the user and group data is organized in the database. The database must be organized using these particular schemas so that the CIFS authentication mechanism is able to extract the data it needs.

For a description of how to add the Samba schema to a Fedora® Directory Server, see:

<http://directory.fedora.redhat.com/wiki/Howto:Samba>

For a description of how to add the samba schema to an OpenLDAP® Server, see:

<http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/passdb.html#id327194>

The following website provides another description of an OpenLDAP configuration:

<http://www.unav.es/cti/ldap-smb/ldap-smb-3-howto.html>

For other LDAP servers (such as the Sun Directory Server, Novell's eDirectory, and IBM's Tivoli Directory Server) the above information may be useful; however, please refer to the relevant documentation for your server product for more information.

NIS

Network information service (NIS) is a network lookup service that provides a centralized database of information about the network to systems participating in the service. The NIS database is fully replicated on selected systems and can be queried by participating systems on an as-needed basis. Maintenance of the database is performed on a central system.

Note: NIS cannot be used for CIFS authentication.

Specify the following:

Domain name	Specifies the NIS domain name for this system.
NIS server IP address	Specifies the IP address of the NIS server. If the NIS server is on the same subnet as SGI Management Center for InfiniteStorage (SMC IS), SMC IS finds the NIS server IP address and provides it as a default. If you are not on the same subnet, you must enter the address in this field.

Click **Apply changes**. You will then be presented with a confirmation screen that allows you to verify whether or not you want to commit the changes.

DNS and Hostnames

You can use the **DNS and Hostnames** screen to specify how to map hostnames to IP addresses for the system. Click **Edit local hosts table** to access the **Hosts** screen, where you can edit the `/etc/hosts` file that contains local mappings or import the contents of a file you specify. For information on the `/etc/hosts` file, see the `hosts(5)` man page.

You can also specify the DNS servers to map hostnames to IP addresses and to resolve hostnames that are incomplete.

Domain Search	<p>Specifies the domain name or names of the DNS servers that the system uses to provide hostname-to-IP-address translation.</p> <p>If you have multiple domains, list them in the order you want to use for lookup. This is important in cases where you have two machines with the same name, each on a different domain, to establish the lookup priority.</p>
Nameserver #	<p>You can specify up to three IP addresses for the DNS name servers to use. If an address you specify is down, the system will use the next one.</p>

Time and Date

Use the **Time and Date** screen to set the following:

Time Zone	<p>Sets the local time zone. You can choose a time zone from a drop-down list of options or you can set a custom time zone. For example, the following specifies what the name of the time zone is for both standard and daylight savings periods, and when the change-over is from daylight to standard and back again (going from standard to daylight on the 10th month and the 5th Sunday, and back again on the 4th month and the first Sunday):</p> <pre>AEST-10AEDT,M10.5.0,M4.1.0</pre> <p>For more information about custom time-zone format, see the <code>tzfile</code> man page.</p>
NTP Time Synchronization	<p>Enables automatic time synchronization with Network Time Protocol (NTP). The NTP protocol is used to synchronize clocks on computer systems over a network. Select Apply NTP changes keep the system's time in synchronization with an NTP server or Set time from NTP server to go off and synchronize it now once only.</p>

If the server has Internet access, see the following website for information about using the public NTP timeserver:

<http://www.pool.ntp.org/>

Set Current Time and Date

Sets the system date (in the format *year/month/day*) and time directly instead of using NTP time synchronization.

Licenses

The **Licenses** screen provides information required to request licenses and a text box in which you can type in or paste permanent licenses obtained from SGI.

Administrator Password

The **Administrator Password** screen changes the SGI Management Center for InfiniteStorage (SMC IS) *administration password*, which is the password required to perform server configuration and management. This password is not required to view the monitoring screens.

Operations

The following sections describe other operations you can perform with SGI Management Center for InfiniteStorage (SMC IS):

- "Save/Restore Configuration" on page 47
- "Support Data" on page 48
- "Performance Data" on page 48
- "Shutdown" on page 48

Save/Restore Configuration

The **Save/Restore Configuration** screen lets you save the current SGI Management Center for InfiniteStorage (SMC IS) configuration or restore a previously saved version. The configuration information saved includes how the interfaces are

configured and what filesystems should be mounted. You may find this useful if you have made an error in the present configuration and you wish to return to a previously configured state.



Caution: This procedure does not provide a system backup and specifically does not save or restore user data; it provides a snapshot record of the configuration.

This screen lists previously saved configurations, labeled by date. After restoring a configuration, you should restart the system.

Support Data

If there is a problem with the system, SGI Call Center Support may request support data in order to find and resolve the problem. The **Gather Support Data** screen lets you generate an archive containing copies of the storage server's software and hardware configuration and log files.

To collect the data, select **Yes, gather information**. This process can take more than 30 seconds on large RAID configurations and requires at least 200 MB of free space in `/tmp`.

Performance Data

This screen lets you capture and download archives of performance data from the server on which SGI Management Center for InfiniteStorage (SMC IS) is running. SGI may request such an archive for performance-analysis purposes, but please be aware that it may contain potentially sensitive information such as network traces.

Note: The **Performance Data** screen in SMC IS is only available if you have installed the `oprofile` and `ethereal` packages.

Shutdown

From the **Shutdown** screen, you can specify to reboot or shut down the system in a specified number of minutes.

Performance Monitoring

This chapter discusses the following:

- "Monitoring Overview" on page 49
- "Metrics Collected" on page 51
- "System Summary" on page 52
- "Alerts" on page 55
- "Storage Device Failures" on page 55
- "Resources" on page 55
- "Services" on page 59
- "Clients" on page 63

Monitoring Overview

SGI Management Center for InfiniteStorage (SMC IS) provides current and historical views of the state and the performance of a storage server. This includes CPU usage, disk and network throughput, and many other metrics. It also allows you to view connected clients and determine how each of these contribute to the current workload.

Note: The monitoring capability is provided by the `pcp-storage` package, which in turn relies on several optional packages that provide specific metrics: `pcp-pmda-infiniband`, `pmdanfs`, and `sgi-samba-pmda`. If any of the packages are installed after `pcp-storage`, you must do the following:

1. Change (`cd`) to the appropriate directory under `/var/lib/pcp/pmdas` and run the `Install` script.
 2. Run `/usr/lib/pcp-storage/setup-pcp-storage`.
-

Figure 4-1 shows the top-level **Monitoring** screen.

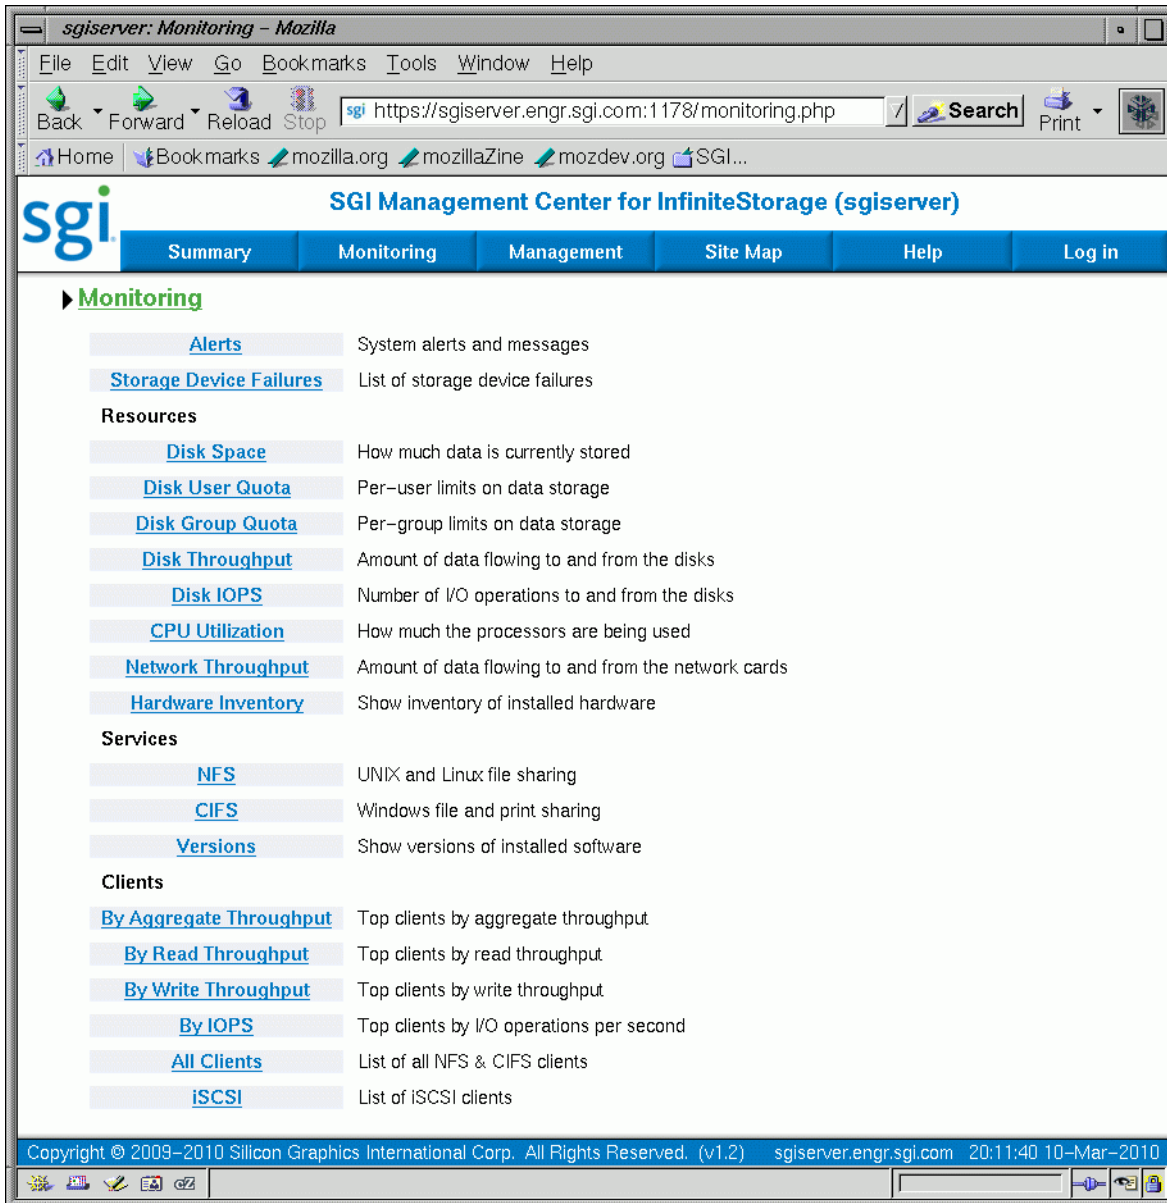


Figure 4-1 Monitoring Screen

Metrics Collected

The information provided by SGI Management Center for InfiniteStorage (SMC IS) can be roughly broken down into “who” and “how much.” SMC IS continuously gathers performance metrics and stores them in archives in `/var/lib/pcp-storage/archives`. Each month, a data reduction process is performed on the metric gathered for the month. This reduces the size of the archives while retaining a consistent amount of information.

Although the size of metric archives has a bounded maximum, this can still be quite large depending on the configuration of the server and how many clients access it. For example, a server with a large number of filesystems could generate up to 100 Mbytes of archives per day. You should initially allow around 2 Gbytes of space for archive storage and monitor the actual usage for the first few weeks of operation.

Note: SMC IS uses the International Electrotechnical Commission's International Standard names and symbols for binary multiples of units. In particular, this means that 1 MiB/s is $2^{20} = 1048576$ Bytes per second. For more information on this standard, see the National Institute of Standards & Technology information about prefixes for binary multiples at:

<http://physics.nist.gov/cuu/Units/binary.html>

SMC IS distinguishes between *current* and *historic* time. Current metrics are either drawn live from the server or are taken from the last few minutes of the metric archives. Historic metrics are taken exclusively from the metric archives. SMC IS displays this historical information for three time periods:

- Last hour
- Last day (the previous 24 hours)
- Last month (the previous 30 days)

Within bar graphs, SMC IS uses color-coding to display the direction of data flow:

- Red represents write and receive data flow
- Blue represents read and send data flow

Figure 4-2 describes how SMC IS color-codes the direction of data flow graphs. For an example of the result in a graph, see Figure 4-3 on page 54.




Figure 4-2 Color-Coding the Direction of Data Flow

System Summary

SGI Management Center for InfiniteStorage (SMC IS) provides a **Summary** menu selection at the top of the screen. This screen displays the following on a per-node basis:

- CPU utilization
- Disk space
- Disk throughput

- Network throughput
- InfiniBand throughput (if installed)
- The number of NFS, CIFS, and iSCSI clients (if iSCSI targets have been created)
- System uptime
- Number of users
- Load average

The screen displays ticks along the status bars that represent the average value over the past day or hour, rather than the immediate value that is shown by the graph. You can drill down to more detailed status by clicking the headings to the left of the graphs. Click **History** to view the historical status of a parameter.

Figure 4-3 shows an example **Summary** screen.




Figure 4-3 Summary Screen

In Figure 4-3, the bar graph for **Disk Throughput** shows 0.0539 MiB/s of data read/sent (the blue part of the graph) and 0.121 MiB/s of data written/received (the red part of the graph). If you were sending and receiving data at the same rate, there

would be equal amounts of red and blue in the graph. For more information, see Figure 4-2 on page 52.

Alerts

The **Alerts** screen displays messages from the system logs. These provide informative messages, notifications of unusual events, and error conditions.

Only unacknowledged alerts are displayed unless you click **Show Acknowledged**. You must log in in order to acknowledge alerts.

After a period of time, alerts are archived and will not be redisplayed. Acknowledged alerts are archived after 2 days and unacknowledged alerts are archived after 7 days. The `/var/lib/ssmc/alerts/archive` file contains all the archived alert messages.

Storage Device Failures

The **Storage Device Failures** screen displays details of any failed devices. You can choose to view the data in the default table format or switch to a tree (hierarchical) listing.

For the LSI MegaRAID platform, the information includes adapter firmware version and device location within its enclosure.

Resources

SGI Management Center for InfiniteStorage (SMC IS) contains a separate screen to display the utilization of each resource.

The following sections provide details about the resources:

- "Disk Space" on page 56
- "Disk User Quota and Disk Group Quota" on page 56
- "Disk Throughput and Disk IOPS" on page 57
- "CPU Utilization" on page 57
- "Network Throughput" on page 58

- "Hardware Inventory" on page 59

Where multiple physical resources are bonded into a single logical resource (for example, load-balanced NICs and RAID volumes in a filesystem), SMC IS shows the structure of the aggregated resource, and (where possible) shows metrics for both the aggregate and the component resources.

Disk Space

The **Disk Space** screen shows the GiB used on each filesystem. If the amount of disk space appears low on a filesystem on which disk quotas are enabled, you can use the **Disk User Quota** screen to find out who is using the most disk space.

Disk User Quota and Disk Group Quota

Disk user/group quotas provide limits on the number of files and the amount of disk space a user/group is allowed to consume on each filesystem. A side effect of this is that they make it possible to see how much each user/group is currently consuming.

Because quotas are applied on a per-filesystem basis, the limits reported in the **All Filesystems** screen are not additive. This means that if a user/group has a 500-MiB disk space limit on filesystem A and a 500-MiB limit on filesystem B, the user/group cannot store a 1-GiB file because there is no single filesystem with a large-enough space allowance.

However the current usage shown in the **used** column on the **All Filesystems** screen is additive, so you can use this screen to determine the user/groups who are currently consuming the most disk space. The **All Filesystems** screen highlights user/groups who have exceeded the quota on any filesystem on which they have been allocated a quota.

Note: Users/groups that do not have quotas explicitly assigned to them are not listed in the monitoring pages.

Disk Throughput and Disk IOPS

Disk operations occur when the result of a file operation is committed to disk. The most common types of disk operation are data reads and writes, but in some types of workload, metadata operations can be significant. *Metadata operations* include the following:

- Truncating and removing files
- Looking up filenames
- Determining the size and types of files

Disk operations are measured in I/O per second (IOPS).

Disk throughput is the amount of data that is transferred to and from the disks. This is predominantly the result of reading and writing data.

The **Disk Throughput** and **Disk IOPS** screens display a bar graph for each active filesystem. For RAID filesystems, a separate graph is displayed for each volume element.

If the cache hit rate is low and the network throughput is high, the disk throughput should be high. Usually, the disk throughput would be steady somewhere a little under the maximum bandwidth of the disk subsystem. If the disk throughput is consistently too high relative to the network throughput, this might indicate that the server has too little memory for the workload.

Under heavy loads, a storage server must be able to sustain a high rate of disk operations. You can use the disk operations metrics in conjunction with other metrics to determine the characteristics of a workload so that you can tune the server can be tuned. For example, a high utilization of NICs but few IOPS could indicate that a workload is coming straight from the cache. A large number of IOPS but low throughput (either disk or network) indicates a metadata-dominated load. You can determine the contributing operations or clients from the **NFS** screen, **CIFS** screen, and the various screens under the **Clients** category.

CPU Utilization

Serving files places demands on the storage server CPU as well as the I/O subsystem. The CPU helps with copying data to and from disks, calculating checksums, and other tasks. Table 4-1 shows the CPU metrics reported.

Table 4-1 CPU Metrics Reported

CPU Metric	Description
Wait time	Time when a CPU was forced to do nothing while waiting for an event to occur. Typical causes of wait time are filesystem I/O and memory swapping.
Interrupt time	Time the CPU spent processing requests from I/O devices. In a storage server context, these are almost exclusively generated by disk operations or network packets and by switching between processes.
System time	Time the CPU spent executing kernel code. This is usually dominated by NFS file serving and accessing data from disks.
User time	Time when the CPU is devoted to running ordinary programs. The biggest consumers of user time in a storage server would usually be the CIFS server, HTTP server, or FTP server.

CPU time is displayed as a percentage, where 100% is the total time available from a single CPU. This means that for an 8-CPU server, the total available CPU time is 800%.

In general, NFS workloads consume more system time, whereas CIFS, HTTP, and FTP workloads consume more user time. SGI Management Center for InfiniteStorage (SMC IS) performance monitoring infrastructure consumes only a small amount of user time.

The most useful problem indicator is consistently having little or no idle time. This can mean that the server is underpowered compared to the workload that is expected of it.

Network Throughput

The **Network Throughput** screen displays the amount of data transferred through each network interface card (NIC).

If an interface is load-balanced, SGI Management Center for InfiniteStorage (SMC IS) displays throughput for both the bonded interface and its constituent interfaces.

Note: The throughput displayed is total network throughput (which includes protocol headers), so real data transfer will be somewhat lower than this value. The **Services** category screens show the amount of real data transferred from a variety of perspectives.

Hardware Inventory

The hardware inventory is a summary of the hardware configuration, including the CPUs, I/O controllers, memory, network controllers, and SCSI disks. The list of SCSI disks includes both the system `root` disk and the configured RAID logical units (LUNs).

Services

A *service* is a task that is performed by the storage server. While the primary service is fileserving, SGI Management Center for InfiniteStorage (SMC IS) breaks this down by the different methods of accessing the server. The services known to SMC IS are NFS and CIFS. This screen also provides access to the software versions installed.

This section discusses the following screens available under the **Services** category:

- "NFS" on page 59
- "CIFS" on page 62
- "Versions" on page 63

NFS

Note: The **NFS** screen is available only if SGI Enhanced NFS is installed.

NFS traffic is a major contributor to storage server utilization. NFS services report statistics aggregated across all exports/shares as well as statistics for each export/share.

Table 4-2 describes the statistics reported by both the **NFS** and **CIFS** screens. Table 4-3 and Table 4-6 describe additional information that is reported.

NFS services gather like-operations into a smaller number of operation classes. Table 4-4 summarizes these classes. (The NFS operation statistics measure classes of NFS protocol operations sent by clients.)

Note: There is not a one-to-one correspondence between CIFS and NFS IOPS. The former measures operations that are received from a network client, the latter measures operations that are sent to a local filesystem.

Table 4-2 Statistics Reported by NFS and CIFS Screens

Graph	Description
Throughput	Current incoming and outgoing traffic for the export/share (the NFS service Throughput graph includes all types of operations, whereas the CIFS graph only shows actual data transfer)
Operations by Type	Export/share operations by class
Read Block Sizes	Reads by size
Write Block Sizes	Writes by size

Table 4-3 Additional Information Reported by the NFS Screen

Category	Description
IOPS	I/O per second for TCP and for UDP
Service Times	Number of operations falling into each service time band as tracked by the NFS server for each operation

Table 4-4 NFS Operation Classes

Class	Description
access	File accessibility tests; checks whether a client can open a particular file
commit	Commit request; requests that the server flush asynchronously written data to stable storage
fsinfo	Filesystem statistics and information requests, <code>pathconf</code> calls, and service availability tests
getattr	File attribute retrieval operations
inode_mods	New file or directory creation, hard and symbolic link creation, file renaming, and device file creation operations
lockd	General lock operations not covered by other classes
lockd_granted	Number of lock granting operations
lockd_share	Number of export/share reservation operations
lookup	Operations that result in filename translations; that is, operations that are applied to a filename rather than to a file handle, such as <code>open</code>
read	File read operations and symbolic link resolution operations
readdir	Directory entry listing operations
readdirplus	Extended directory entry listing operations; returns the attributes of the directory entries as well as their names
remove	File deletion operations
setattr	File attribute setting operations, which include file truncations and changing permissions and ownership
write_async	Asynchronous writes; the written data may be cached and scheduled for writing at a later time
write_sync	Synchronous write; these do not complete until the data is written to stable storage
xattr	Operations that manipulate XFS [®] extended attributes

CIFS

Note: The **CIFS** screen is available only if SGI Samba packages are installed.

CIFS traffic is a major contributor to storage server utilization. CIFS services report statistics aggregated across all exports/shares as well as statistics for each export/share.

Table 4-2 describes the statistics reported by both the **NFS** and **CIFS** screens.

CIFS services gather like operations into a smaller number of operation classes. While these classes are largely similar, there are some differences. Table 4-5 summarizes these classes.

Note: Clients can perform file operations in a variety of different ways, which can result in similar logical operations being recorded as differing sets of CIFS operations depending on the application.

Table 4-5 CIFS Operation Classes

Class	Description
cancel	Cancel current activity operations
change/notify	Operations requesting notification of changes to a file or in a directory
close	File close operations
create/open	File and directory create and open operations
delete/remove	File deletion and directory removal operations
findfirst/next	Operations searching for files or scanning a directory
flush	Operations requesting a flush to disk of buffered data
getattr	Operations requesting file and directory attributes, such as access times
getsecurity	Operations requesting file access permissions
ioctl	Operations performing special filesystem features, such as sparse file handling

Class	Description
lock/unlock	File locking and unlocking operations
misc	All other operations, including infrequent filesystem features
move	File and directory move and rename operations
read	File read operations
setattr	Operations setting file and directory attributes, such as hidden file status
setsecurity	Operations setting file access permissions
write	File write operations

Table 4-6 Additional Information Reported by the **CIFS** Screen

Category	Description
IOPS	Number of SMB operations per second
Latencies	Number of SMB operations falling into each service time band

Versions

The **Versions** screen displays the version numbers of key software packages that have been installed.

Clients

A *NAS client* is a computer running a program that accesses the storage server. NAS clients are known to SGI Management Center for InfiniteStorage (SMC IS) by their IP address; if multiple accessing programs are running on the same computer, they are all counted as a single client.

Note: Detailed client information is gathered only for CIFS and NFS protocols.

The **All Clients** screen will not be available if neither SGI Samba nor SGI Enhanced NFS are installed.

The **All Clients** screen displays the NAS clients sorted according to hostname. The other selections sort according to the chosen selection (such as by aggregate throughput).

From each of these screens, you can change the sorted display of the data without returning to the **Monitoring** screen.

Displaying the NAS clients in this fashion is useful for pinpointing how the current set of clients are contributing the workload profile. For example, upon noticing an unusually large amount of network traffic on the **Network Throughput** screen, changing to display the clients in order of aggregate throughput will quickly identify the contributing clients.

From the list of clients, you can display a detailed view of the NFS and CIFS traffic generated by a particular client. This is useful when trying to diagnose problems that affect only a single client or type of client. For example, by viewing the client detail, it may be obvious that throughput is limited by the client using very small read and write sizes. Continuing from the client details to the client history screen can help diagnose problems, such as hung NFS mounts.

The **iSCSI** screen displays a list of the connected iSCSI initiators are connected and their targets.

Troubleshooting

This section discusses the following:

- "Forgotten Password or Corrupt Password File" on page 65
- "The archives Directory is Too Large" on page 66
- "Power Outage and iSCSI" on page 66
- "Manual System Reboot" on page 67
- "Network Configuration Issues" on page 67
- "Reporting Problems to SGI" on page 68

Forgotten Password or Corrupt Password File

If you forget the administrator password or if the **Alerts Page** reports that the `/etc/ssmc/passwd` file is corrupt (preventing administrator login) run the following to set a new password of your choice (*NEWPASSWORD*):

```
# echo "ssmc_admin:`echo -n $NEWPASSWORD | md5sum | cut -d' ' -f1`" > /etc/ssmc/passwd
```

The archives Directory is Too Large

SGI Management Center for InfiniteStorage (SMC IS) stores historical information in the directory `/var/lib/pcp-storage/archives`. On a large machine, this directory may require too much disk space to fit in the `/` or `/var` filesystem. This directory can be moved to any other filesystem (assuming the new filesystem always remains mounted) using the following procedure:

1. Stop services related to SMC IS, `pcp-storage`, and Performance Co-Pilot:

```
# service ssmc stop
# service pcp-storage stop
# service pcp stop
```

2. Change to the `pcp-storage` directory:

```
# cd /var/lib/pcp-storage
```

3. Move the `archives` directory to a different filesystem:

```
# mv archives /some/other/filesystem/
```

4. Link the other filesystem to the `archives` location:

```
# ln -s /some/other/filesystem/archives
```

5. Restart services:

```
# service pcp start
# service pcp-storage start
# service ssmc start
```

Power Outage and iSCSI

Due to the nature of iSCSI as a block-level protocol (as distinct from file-level protocols such as NFS and CIFS), particular care must be taken in the event of a system crash, power failure, or extended network outage.

If power is lost to the server while an iSCSI initiator is performing a write to an iSCSI target, the write will not be completed and the filesystem created on that particular target may then be in an inconsistent state. The iSCSI initiator should be made to perform a filesystem check on the iSCSI target immediately after power is restored, and before trying to access that target for normal usage.

For example, on a Windows client:

1. Use the iSCSI Initiator program to connect to the iSCSI target.
2. Open **My Computer**.
3. Right-click the iSCSI target drive and select **Properties**.
4. In the **Properties** window, select the **Tools** tab and click the **Check Now** button.
5. In the **Check Disk** window, select both **Automatically fix file system errors** and **Scan for and attempt recovery of bad sectors**.
6. Click **Start** to verify the filesystem and attempt recovery of any errors.

Manual System Reboot

If you must reboot the system but SGI Management Center for InfiniteStorage (SMC IS) is inaccessible, do the following:

1. Log in via the system console as `root`, such as via the L2 on an SGI ia64 system or via IPMI or a monitor/keyboard on an SGI x86_64 system.
2. Reboot the system:

```
# reboot
```

Network Configuration Issues

If the network configuration is damaged or if the system running SGI Management Center for InfiniteStorage (SMC IS) becomes inaccessible via the network, do the following:

1. Log in via the system console as `root`, such as via the L2 on an SGI ia64 system or via IPMI or a monitor/keyboard on an SGI x86_64 system.
2. Reconfigure the management interface (`eth0`) by using the following commands, as appropriate for your site:

- Static IP address:

```
# /usr/lib/ssmc/ssmc-cli -c "network if-enable-static eth0 IPaddress 255.255.255.0"
```

For example, for a static IP address of 192.168.9.9:

```
# /usr/lib/ssmc/ssmc-cli -c "network if-enable-static eth0 192.168.9.9 255.255.255.0"
```

- DHCP:

```
# /usr/lib/ssmc/ssmc-cli -c "if-enable-dhcp eth0"
```

3. To set the default gateway (such as if the system must communicate with other systems outside the local network or if the default gateway is not supplied by a DHCP server), enter the following:

```
# /usr/lib/ssmc/ssmc-cli -c "network default-gateway-set Default_Gateway_IPaddress"
```

For example, for a default gateway of 192.168.9.254:

```
# /usr/lib/ssmc/ssmc-cli -c "network default-gateway-set 192.168.9.254"
```

4. Reset eth0:

```
# /usr/lib/ssmc/ssmc-cli -c "network if-reset eth0"
```

5. Restart SMC IS:

```
# service ssmc restart
```

Reporting Problems to SGI

See "Support Data" on page 48 for information about gathering the information that SGI Support will require when diagnosing problems.

Glossary

Active Directory

A directory service that implements *LDAP* in a Windows environment. It provides a hierarchical structure for organizing access to data.

administration password

The password required to log into the **Management** screens

bonded network interface

Virtual network interface that consists of real interfaces working in tandem. A virtual interface can provide the aggregated bandwidth of all of the interfaces that you used to create it.

CHAP

Challenge Handshake Authentication Protocol is a means of authentication used between a client and server where the password is sent over the wire in a form that is impossible to discover and impossible to replay. Both client and server must know what the original password is, but someone snooping on wire traffic cannot recover the password and cannot later send the original (snooped upon) authentication packet to the server in an attempt to try to trick it into letting them authenticate as a valid client.

CIFS

Common internet filesystem. This protocol is usually used by Microsoft Windows clients.

current metric

Metric drawn live from the server or taken from the last few minutes of the metric archives.

default network gateway

The IP address of the router that this system should use to communicate with machines that are outside of its subnet.

DHCP

Dynamic host configuration protocol (DHCP) allows one or more server systems to dynamically distribute network IP addresses and site configuration parameters to new or requesting client systems. By using DHCP, a site with only a few available addresses can serve a large number of hosts that connect to the network only occasionally, or a large site can manage the permanent assignment of addresses with a minimum of administrative attention. The NAS server can be configured as a DHCP client.

directory service

See *name service*.

disk IOPS

Disk I/O per second.

disk throughput

The amount of data that is transferred to and from disks.

distinguished name

A unique identifier for an entry in an LDAP directory tree structure.

DNS

Domain name system.

FC

Fibre Channel storage interface connection.

FQDN

Fully qualified domain name.

historic metric

Metric taken exclusively from the metric archives.

idle time

Time that remained when the CPU could not find any tasks to run.

initiator

The client accessing the storage in an iSCSI network.

interrupt time

Time the CPU spent processing requests from I/O devices. In a storage server context, these are almost exclusively generated by disk operations or network packets and by switching between processes.

IOPS

I/O per second.

IPMI

Intelligent Platform Management Interface, a system reset method.

IPoIB

IP over InfiniBand.

iSCSI

Internet Small Computers Systems Interface is a protocol that is used to transport SCSI commands across a TCP/IP network. This allows a system to access storage across a network just as if the system were accessing a local physical disk. In an iSCSI network, the client access the storage is called the *initiator*. The remote storage that the client accesses is called the *target*.

LDAP

Lightweight directory access protocol (LDAP) is a networking protocol that organizes access to data in a directory tree structure.

KDC

Key distribution center.

metadata

Information that describes a file, such as the file's name, size, location, and permissions.

name service

Application that manages the information associated with network users.

NAS client

Computer running a program that accesses the storage server.

NFS

Network file system.

NIC

Network interface card.

NIS

Network information service (NIS) is a network lookup service that provides a centralized database of information about the network to systems participating in the service.

NTP

Network Time Protocol.

physical volume element

The combination of multiple RAID disk drives into a single logical unit.

RAID

Redundant array of independent disks.

RAID 5

A level of RAID that uses block-level striping and distributed parity.

resource

Something that is monitored and managed by SGI Management Center for InfiniteStorage (such as network interfaces).

Serial ATA (SATA)

Serial advanced technology attachment storage interface connection.

service

Task performed by the storage server.

shadow file

A file that is protected from all access by non-root users and stores the encrypted passwords.

smart host

The gateway server where email should be delivered.

SMC IS

SGI Management Center for InfiniteStorage

ssmc

Name of the service that provides SGI Management Center for InfiniteStorage

system time

Time the CPU spent executing kernel code. This is usually dominated by NFS file serving and accessing data from disks.

target

The storage that appears to the initiator as a disk drive in an iSCSI network.

wait time

Time when a CPU was forced to do nothing while waiting for an event to occur. Typical causes of wait time are filesystem I/O and memory swapping.

Index

802.3ad standard, 19

A

- About menu selection, 5
- access operation, 61
- Active Directory, 10, 29, 41
- admin server, 33
- administration password, 9
- Administrator Password screen, Global
 - Configuration, 47
- aggregate interface
 - See "bonded interface", 9
- alerts, 55
- Alerts menu selection, 4
- All Clients screen, 64
- All Filesystems screen, 56
- archives, 51
- archives directory size, 66
- Asynchronous Writes, NFS export option, 34
- authentication services, 10
- autonegotiate, 14, 16

B

- backup of the SMC IS configuration, 5
- blue color in graphs, 51
- bonded interface, 9
- bonded network interface, 13
- bonded network interfaces, 17
- bonding mode, 18, 19

C

- cancel operation, 62
- change/notify operation, 62
- CHAP authentication, 25
- CIFS, 62
 - client number, 53
 - configuration, 36
 - iSCSI and, 23
- CIFS authentication, 41, 45
- CIFS screen, 57, 62
- clients, 63
- Clients category, 57
- Clients menu selection, 4
- close operation, 62
- colors in graphs, 51
- commit operation, 61
- community string, 39
- Completed screen, 9
- corrupt password file, 65
- CPU utilization, 52, 57
- create/open operation, 62
- current time, 51
- custom installation, 9

D

- data flow color-coding in graphs, 52
- data reduction process, 51
- dedicated network interface, 15, 16
- default network gateway, 40
- device failure notification, 21
- DHCP, 14, 16
- disk
 - IOPS, 57
 - operations, 57

- quotas, 56
 - space, 52, 56
 - throughput, 53, 57
 - throughput, monitoring, 57
- Disk IOPS screen, 57
- Disk Quota screen, 56
- DNS and Hostnames screen, Global Configuration, 45
- domain, 24, 33
- Domain Search, DNS and Hostnames screen, 46
- duplex option, 14, 16
- dynamic bonding mode, 19

E

- /etc/hosts, 9
- /etc/krb5.keytab, 33
- eth0, 14
- ethereal, 48
- Exchange Server as an iSCSI initiator , 22
- export options, 33

F

- Fedora Directory Server, 44
- filesystem listing, 22
- findfirst/next operation, 62
- Firefox, 7
- flush operation, 62
- fsinfo operation, 61
- full-duplex, 14, 16

G

- gateway, 40
- Gather Support Data screen, Global Configuration, 48
- getattr operation
 - CIFS, 62

- NFS, 61
- getsecurity operation, 62
- global configuration, 40
- Global Configuration menu selection, 5
- global operations, 32
- group quotas, 31

H

- half-duplex , 14, 16
- hard limit, 31
- hardware inventory, 59
- historic time, 51
- historical status of a parameter, 53
- History menu selection, 53

I

- identifier for target, 24
- IEEE 802.3ad standard, 19
- Import Users option, Local Users screen, 29
- InfiniBand network interface, 16
- InfiniBand throughput, 53
- initial system setup, 7
- initiator for iSCSI, 22, 27
- inode_mods operation, 61
- input load balancing, 18
- installation customization, 9
- interface overview, 2
- Internet Explorer, 7
- Internet Small Computer Systems Interface
 - See "iSCSI", 22
- interrupt time, 58
- ioctl operation, 62
- IOPS, 57, 60, 63
 - CIFS, 60
 - NFS, 60
- IP address, 20
- IP header, 20

iqn, 24

iSCSI

- client number, 53
- create target, 28
- destroy targets, 28
- domain, 24
- identifier, 24
- initiator, 22, 27
- list targets, 28
- modify targets, 28
- network, 22
- NFS and CIFS, 23
- protocol, 22
- qualified name, 24
- re-exporting targets, 23
- start/stop, 28
- target, 22
- targets, 22, 24

iSCSI and power outage, 66

iSCSI Initiator program, 26

J

JavaScript, 7

K

KDC, 33

Kerberos, 33

Kerberos aware clients, 34

Kerberos with Integrity support aware clients, 35

key distribution center, 33

keytab, 33

krb5 and krb5i, 34

L

Layer 2 (MAC address), 20

Layer 3 (IP header), 20

LDAP, 10, 29

LDAP (lightweight directory access protocol), 44

Licenses screen, Global Configuration, 47

load average, 53

load balancing, 18, 20

local subnet, NFS access, 34

local users and groups, 9, 29

lockd operation, 61

lockd_granted operation, 61

lockd_share operation, 61

Log In menu selection, 5

Log Out menu selection, 5

lookup operation, 61

LSI MegaRAID, 21, 55

M

MAC address header, 20

mail store and iSCSI, 22

main menu, 4

management interface, 13

Management menu selection, 4

management password, 9

menu path, 2

metadata operations, 57

metrics

 CPU, 58

 type collected, 51

MiB vs MB, 51

misc operation, 63

Modify option, 14, 16

modify the installation, 9

Monitoring menu selection, 4

monitoring performance, 49

monitoring screen example, 50

move operation, 63

mutual CHAP authentication, 26

N

- name service client, 41
- nameserver, 46
- network configuration issues, 67
- Network Information Service (NIS), 45
- network interface
 - bonded, 17
 - InfiniBand, 16
 - management, 13
 - standalone, 14
- network interface configuration, 13
- network interfaces, 9
- network throughput, 53, 58, 64
- Network Time Protocol (NTP), 46
- NFS, 32, 59
 - client number, 53
 - custom definition, 35
 - export options, 33
 - iSCSI and, 23
 - restrict to hosts, 35
- NFS screen, 57, 59
- NFS serving domain, 32
- NFSv4 enabling, 32
- NIS, 10, 45
- NTP, 46
- NTP Time Synchronization, Date and Time screen, 46
- number of users, 53

O

- OpenLDAP Server, 44
- operation
 - CIFS, 62, 63
- operation classes, 60
- Operations, 5
- operations by type, 60
- oprofile, 48
- output load balancing, 18, 20
- overview, 1

P

- password default, 9
- password problems, 65
- Passwords screen, 9
- performance archives, 51
- performance data, 5
- Performance Data screen, 48
- performance monitoring, 49
- port speed, 14, 16
- power outage and iSCSI, 66
- principal user (Kerberos) , 33

Q

- Quick Start guide, 7
- quotas
 - disk, 56
 - group, 31
 - user, 30
 - user and group, 30

R

- re-exporting iSCSI targets with NFS or CIFS, 23
- read block sizes, 60
- read operation
 - CIFS, 63
 - NFS, 61
- Read-only NFS Export option, 34
- readdir operation, 61
- readdirplus operation, 61
- realm, 33
- reboot, 5
- red color in graphs, 51
- remove operation, 61
- reporting problems to SGI, 69
- resources, 4, 55
- Resources menu selection, 4

restrict to hosts, NFS option, 35

S

Samba schema, 44
Save/Restore Configuration screen, 47
secret for CHAP authentication, 25
server configuration and management, 11
service times, 60, 63
services, 59
Services menu selection, 4, 5
serving domain for NFS, 32
setattr operation
 CIFS, 63
 NFS, 61
Setup Wizard, 7
 Completed screen, 9
 Passwords screen, 9
Share Options, CIFS configuration, 36
shutdown, 5
Shutdown screen, Global Configuration, 48
Site Map menu selection, 5
SNMP, 38
soft limit, 31
ssmc_admin, 65
standalone network interface, 14
static bonding mode, 19
Static option, 14, 16
storage configuration, 21
storage device failures, 55
subnet mask, 20
Summary menu selection, 4, 52
Summary screen example, 54
system alerts, 55
system console, 67
system logs, 55
system name, 40
System Name, Global Configuration, 40
system setup, 7
system time, 58
system uptime, 53

T

target for iSCSI, 22
 CHAP authentication, 25
 creating, 24
 identifier, 24
 re-exporting with NFS or CIFS, 23
 username, 25
target name, 24
throughput, 57
 CIFS, 60
 network, 58
 NFS, 60
Time and Date screen, 46
time zone, 11
Time Zone, Date and Time screen, 46
trap destination, 39
troubleshooting, 65
 archives directory size, 66
 password issues, 65
 power outage and iSCSI, 66
 reporting problems, 69
 SGI Management Center for InfiniteStorage is
 inaccessible
 rebooting the system, 67
 SMC IS is inaccessible
 network configuration issues, 67

U

unit measures, 51
uptime of system, 53
Use custom definition option, NFS screen, 35
user and group configuration, 29
user numbers, 53
user quotas, 30
user time, 58
users and groups, 9

V

/var/lib/pcp-storage/archives, 66
/var/lib/pcp-storage/archives directory, 51
/var/lib/ssmc/alerts/archive, 55
versions, 63

W

wait time, 58
web browsers, 7

write block sizes, 60
write operation/primary>, 63
write_async operation, 61
write_sync operation, 61

X

xattr operation
NFS, 61