



SGI® Management Center for
InfiniteStorage™ Administrator Guide

007-5652-005

COPYRIGHT

© 2010-2011, 2013 Silicon Graphics International Corp. 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, SGI InfiniteStorage, 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 includes the following changes for SGI[®] Management Center for InfiniteStorage[™] (SMC IS):

- Support for x86_64 servers running either of the following operating systems:
 - Red Hat[®] Enterprise Linux[®] 6 (RHEL 6)
 - SUSE[®] Linux Enterprise Server 11 Service Pack 2 (SLES 11 SP2)

As of ISSP 3.0, SMC IS does not support ia64 servers.

- Internet Protocol version 6 (IPv6) use is documented but is provided only as a **technology preview**



Caution: Technology previews are not supported for production use. For more information, see the ISSP release note.

- Redesigned user interface based on the Dojo Toolkit

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
005	April 2013 Revision with ISSP 3.0

Contents

About This Guide	xv
Related Publications	xv
Obtaining Publications	xv
Conventions	xvi
Reader Comments	xvii
1. Overview	1
Requirements	1
Accessing SMC IS	2
SMC IS User Interface	2
Features	4
Monitoring	4
Management	5
Server	7
Refresh	8
Help	9
Log Out	9
2. Monitoring	11
Monitoring Overview	11
Metrics Collected	12
Storage Monitoring	15
Space	15
Throughput and IOPS	15
Quotas	16
007-5652-005	vii

Set User Quotas	17
Set Group Quotas	18
NFS Exports	18
CIFS Shares	21
NFS & CIFS Clients	22
iSCSI Targets	23
Device Failures	23
System Monitoring	24
Alerts	24
CPU Utilization	26
Network Throughput	27
Hardware Inventory	27
Software Versions	27
3. Management	29
Storage Management	30
Filesystems	31
NFS Exports	31
CIFS Shares	33
iSCSI Targets	36
iSCSI Targets Overview	36
Creating iSCSI Targets	38
Modifying iSCSI Targets	41
The iSCSI Initiator	41
Failure Notification	42
System Management	42
Network Interfaces	43
Ethernet Network Interfaces	43

Bonded Network Interfaces	45
System Name	48
DNS & Name Servers	49
Name Service Client	49
Local Files Only	50
Active Directory	50
LDAP	52
NIS	54
SNMP	54
Time & Date	55
Licenses	56
Administrator Password	56
Local Users and Local Groups	56
Operations Management	57
Save/Restore Configuration	57
Gather Support Data	58
Shut Down System	58
4. Troubleshooting	59
Forgotten Password or Corrupt Password File	59
The archives Directory is Too Large	59
Power Outage and iSCSI	60
Manual System Reboot	61
Network Configuration Issues	61
Reporting Problems to SGI	62
Glossary	63
Index	69

Figures

Figure 1-1	SMC IS Interface	3
Figure 1-2	Summary Page	8
Figure 2-1	Network Throughput Page	12
Figure 2-2	Color-Coding the Direction of Data Flow	14
Figure 2-3	Alerts Page	25
Figure 3-1	Time & Date Page	30
Figure 3-2	iSCSI Storage	37
Figure 3-3	Bonded Network Interfaces	46

Tables

Table 2-1	Statistics Reported by NFS Exports and CIFS Shares	19
Table 2-2	Additional Information Reported by NFS Exports	19
Table 2-3	NFS Operation Classes	20
Table 2-4	CIFS Operation Classes	21
Table 2-5	Additional Information Reported by CIFS Shares	22

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, "Monitoring" on page 11, describes the current and historical views of the state and the performance of a storage server
- Chapter 3, "Management" on page 29, describes how to modify the various components of your system and perform general system administration
- Chapter 4, "Troubleshooting" on page 59, 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

SGI® Management Center for InfiniteStorage™ (SMC IS) is a web-based interface that lets you monitor and manage a NAS storage server. This chapter discusses the following:

- "Requirements" on page 1
- "Accessing SMC IS" on page 2
- "SMC IS User Interface" on page 2
- "Features" on page 4

Requirements

SMC IS requires the following:

- JavaScript™
- One of the following browsers:
 - Firefox® 3.6 or newer (Firefox is the preferred browser)
 - Internet Explorer® 8 or newer

Also note the following:

- If you have a RHEL system with SELinux enabled, the domains for `httpd` and `syslogd` will be set to `permissive` by SMC IS during installation.
- If you are running a browser on Windows, you must install the Arial Unicode MS font (or an equivalent) in order to display check marks for acknowledged alerts.
- If your system depends on `/etc/udev/rules.d/70-persistent-net.rules` for network interface configuration, make sure that an entry for `eth0` exists.

Accessing SMC IS

To access SMC IS, 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
```

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 default password is `INSECURE`. You should change this to something appropriate for your site; see "Administrator Password" on page 56.

SMC IS User Interface

To access the SMC IS features, click one of the buttons displayed across the top of the screen. This will display a set of tabs (such as **Storage**) in the left-hand pane; click a tab to expand it and display a list of features. For example, Figure 1-1 shows the screen you would see if you clicked the **Monitoring** button, the **System** tab in the **Monitoring** pane, and the **CPU Utilization** feature.

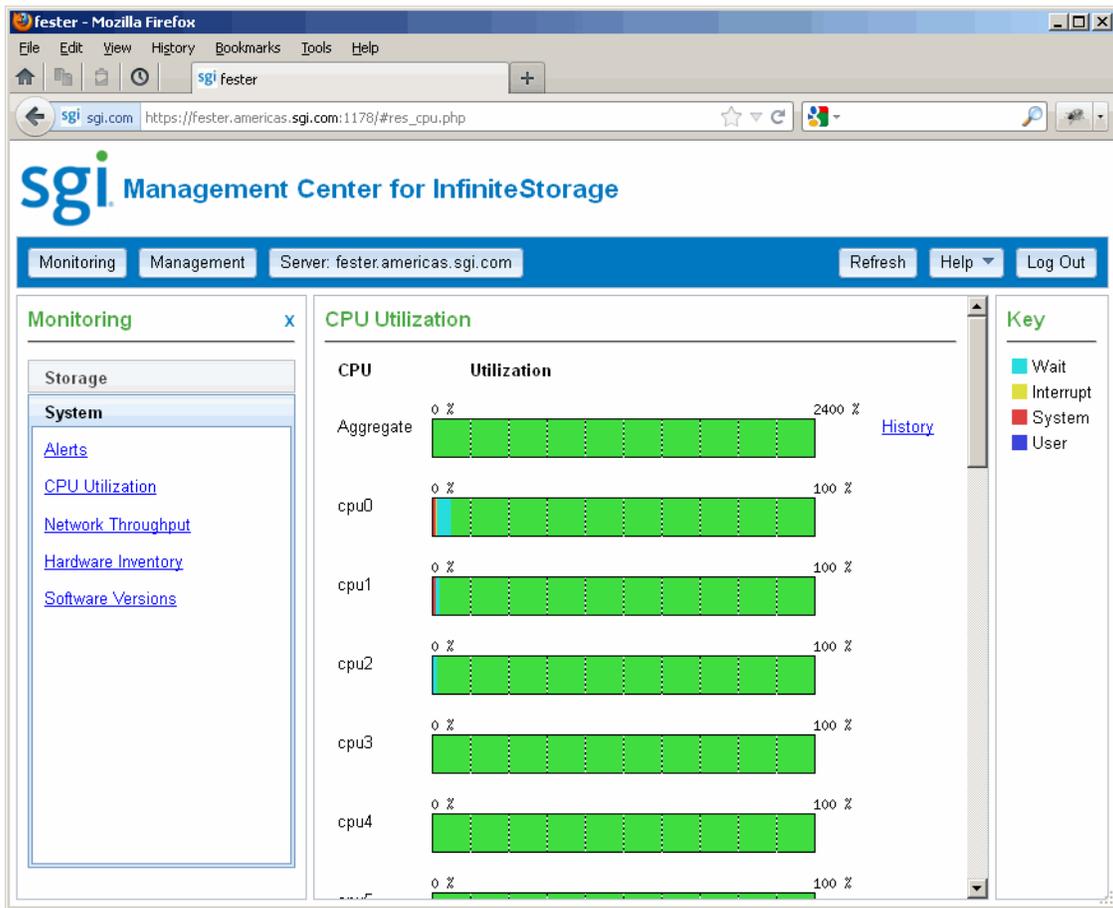


Figure 1-1 SMC IS Interface

To expand the size of the **CPU Utilization** page, click the **X** in the **Monitoring** pane to close that pane. To make the **Monitoring** pane reappear, click the **Monitoring** button.

Features

The following sections summarize the features available with each of the following buttons:

- "Monitoring" on page 4
- "Management" on page 5
- "Server" on page 7
- "Refresh" on page 8
- "Help" on page 9
- "Log Out" on page 9

Monitoring

The **Monitoring** button provides access the following features:

- **Storage** tab:
 - **Space** displays the disk space used on each filesystem. See "Space" on page 15.
 - **Throughput** displays the amount of data that is transferred to and from the disks. See "Throughput and IOPS" on page 15.
 - **IOPS** displays a bar graph of disk I/O per second (IOPS) for each active filesystem. See "Throughput and IOPS" on page 15.
 - **Quotas** displays the disk user/group quotas that provide limits on the number of files and the amount of disk space a user/group is allowed to consume on each filesystem. If you log in, you can modify the settings for user and group quotas. See "Quotas" on page 16.
 - **NFS Exports** displays statistics about Network File System (NFS) exports. See "NFS Exports" on page 18.

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.

- **CIFS Shares** displays statistics about Common Internet File System (CIFS) shared files. See "CIFS Shares" on page 21.
- **NFS & CIFS Clients** displays various I/O criteria by which to display information about the storage server's clients. See "NFS & CIFS Clients" on page 22.
- **iSCSI Targets** displays a list of the connected iSCSI initiators and their targets. See "iSCSI Targets" on page 23.
- **Device Failures** displays details about failed storage devices. See "Device Failures" on page 23.
- **System tab:**
 - **Alerts** displays messages from the system logs. If you log in, you can acknowledge alerts. See "Alerts" on page 24.
 - **CPU Utilization** reports metrics for CPU use. See "CPU Utilization" on page 26.
 - **Network Throughput** displays the amount of data transferred through each network interface card (NIC). See "Network Throughput" on page 27.
 - **Hardware Inventory** shows a summary of the hardware configuration, including the CPUs, memory, network controllers, and SCSI disks. See "Hardware Inventory" on page 27.
 - **Software Versions** displays the installed operating system level and the version numbers of key software packages. See "Software Versions" on page 27.

Management

The **Management** button provides access to the following tasks:

- **Storage tab:**
 - **Filesystems** displays a brief description of the available local filesystems. See "Filesystems" on page 31.
 - **NFS Exports** lets you configure filesystems so that they are available for network clients by means of the NFS network protocol. See "NFS Exports" on page 31.

- **CIFS Shares** lets you configure filesystems so that they are available for network clients by means of the CIFS network protocol. See "CIFS Shares" on page 33.
- **iSCSI Targets** lets you configure the system to access storage across a network just as if the system were accessing a local physical disk. See "iSCSI Targets" on page 36.
- **Failure Notification** lets you configure notification of failed devices. See "Failure Notification" on page 42.
- **System tab:**
 - **Network Interfaces** configures the network interfaces for the system. See "Network Interfaces" on page 43.
 - **System Name** sets the hostname and management IP address for the server. See "System Name" on page 48.
 - **DNS & Name Servers** specifies how to map hostnames to IP addresses for the system. See "DNS & Name Servers" on page 49.
 - **Name Service Client** specifies various directory services that manage information associated with the network users, such as mapping user names with user IDs and group names with group IDs. See "DNS & Name Servers" on page 49.
 - **SNMP** enables Simple Network management Protocol (SNMP) access. See "SNMP" on page 54.
 - **Time & Date** sets the local time and enables automatic time synchronization with NTP. See "Time & Date" on page 55.
 - **Licenses** provides information about the installed licenses and lets you add and delete licenses. See "Licenses" on page 56.
 - **Administrator Password** changes the SMC IS administrator password. See "Administrator Password" on page 56.
 - **Local Users** creates or imports local user accounts. See "Local Users and Local Groups" on page 56.
 - **Local Groups** creates or imports local groups. See "Local Users and Local Groups" on page 56.

- **Operations** tab:
 - **Save/Restore Configuration** saves the current SMC IS configuration or restore a previously saved version. See "Save/Restore Configuration" on page 57.
 - **Gather Support Data** generates an archive containing copies of the storage server's software and hardware configuration and log files, which may be needed by SGI Support for troubleshooting purposes. See "Gather Support Data" on page 58.
 - **Shut Down System** lets you reboot or power down the system in a specified number of minutes. See "Shut Down System" on page 58.

Server

The **Server** button displays a **Summary** page that contains the following:

- Current time as reported by the server
- System uptime
- Number of users
- Load average
- Number of unacknowledged alerts
- 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)

The ticks along the status bars 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. Click other links, such as **Alerts**, to access other pages.

Figure 1-2 shows an example **Summary** page.

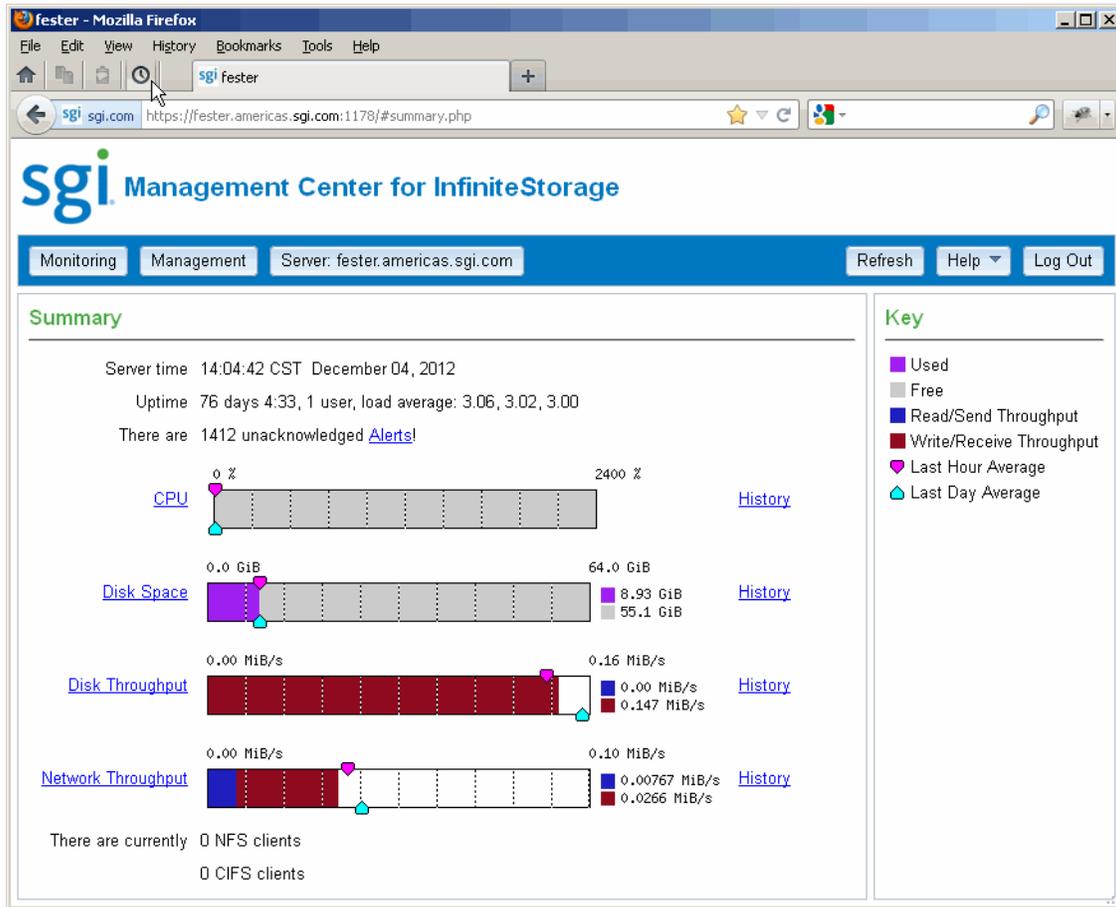


Figure 1-2 Summary Page

In Figure 1-2, the example bar graph for **Network Throughput** shows that more data is being received (the red part of the graph) than sent (the blue 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 2-2 on page 14.

Refresh

The **Refresh** button redraws the currently displayed page.

Help

The **Help** button provides access to information about SMC IS, including the following:

- Copyright, trademark, and open-source statements
- Links for contacting support and providing feedback
- Information about conventions for units of measure, IP addresses, and color-coding of data flow directions
- ISSP release note
- This guide

Log Out

The **Log Out** button exits from management functions. When you log out, SMC IS displays the server **Summary** page. If you are logged out and try to access a management task, you will be prompted to log in. You must enable cookies to log in.

Monitoring

This chapter describes how to use SGI Management Center for InfiniteStorage (SMC IS) to monitor various components of your system. This chapter discusses the following:

- "Monitoring Overview" on page 11
- "Metrics Collected" on page 12
- "Storage Monitoring" on page 15
- "System Monitoring" on page 24

Monitoring Overview

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 lets 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 2-1 shows an example **Network Throughput** page available under the **System** tab in the **Monitoring** pane.

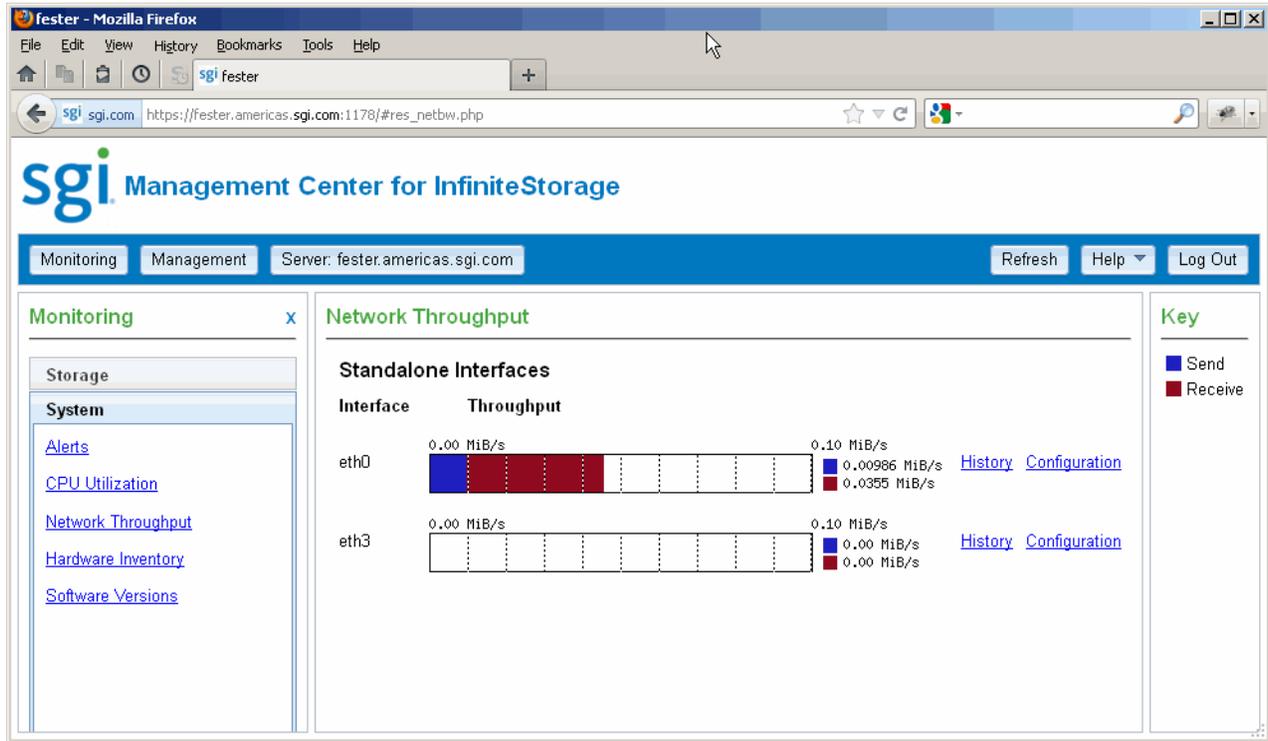


Figure 2-1 Network Throughput Page

Metrics Collected

The information provided by 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 2-2 describes how SMC IS color-codes the direction of data flow graphs. For an example of the result in a graph, see Figure 1-2 on page 8.

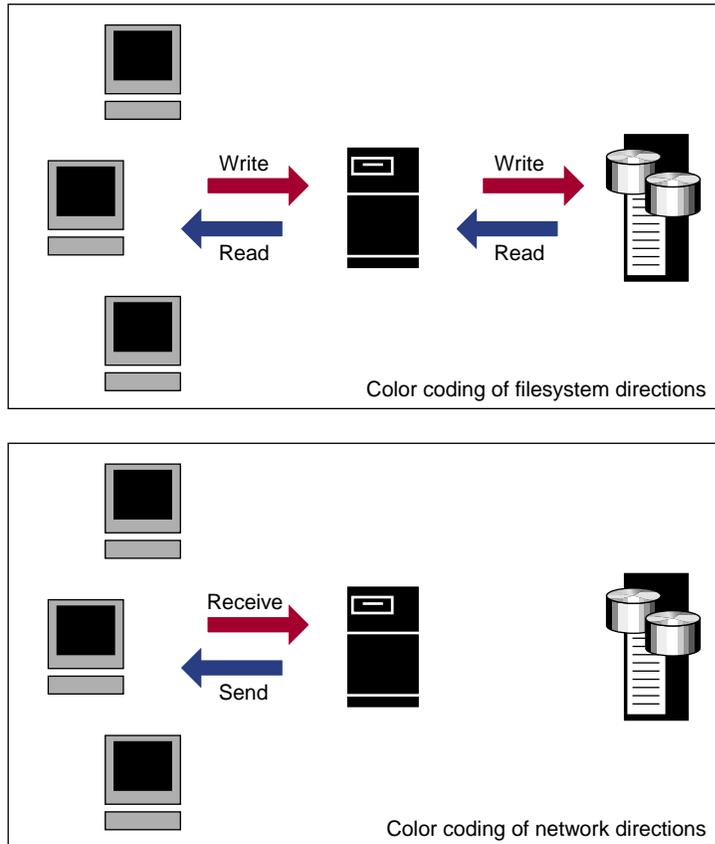


Figure 2-2 Color-Coding the Direction of Data Flow

By default, SMC IS displays **synchronized** scaling, meaning that the vertical scales for historical graphs will be similar; you can also switch to **independent** scaling, meaning that graph-specific scaling may be used, as appropriate.. For more information about historic representation, see "Metrics Collected" on page 12.

Storage Monitoring

The **Storage** tab in the **Monitoring** pane provides access to the following:

- "Space" on page 15
- "Throughput and IOPS" on page 15
- "Quotas" on page 16
- "NFS Exports" on page 18
- "CIFS Shares" on page 21
- "NFS & CIFS Clients" on page 22
- "iSCSI Targets" on page 23
- "Device Failures" on page 23

Space

The **Disk Space** page 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 Quotas** page to find out who is using the most disk space; see "Quotas" on page 16.

To see historical output, for a given filesystem, click its **History** link.

Throughput and 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 **IOPS** page displays 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 Exports** page, **CIFS Shares** page, and the various client pages available from the **NFS & CIFS Clients** page. See:

- "NFS Exports" on page 18
- "CIFS Shares" on page 21
- "NFS & CIFS Clients" on page 22

To see historical output, for a given disk, click its **History** link.

Quotas

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 when you choose **All Filesystems** 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 for **All Filesystems** is additive, so you can use this information to determine the user/groups who are currently consuming the most disk space. The **All Filesystems** display 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.

To manage quotas, you must log in. See:

- "Set User Quotas" on page 17
- "Set Group Quotas" on page 18

Set User Quotas

To set a user quota, select **User** on the **Disk Quotas** page and choose a specific filesystem name or else **All Filesystems**, then click **Update**.

Click the user name to access the user's disk and file limits fields. Enter the desired values and click **Apply changes**.

The fields are as follows:

- 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.)

Set Group Quotas

To set a user quota, select **Group** on the **Disk Quotas** page and choose a specific filesystem name or else **All Filesystems**, then click **Update**.

Click the name of the group to access the disk and file limits fields. Enter the desired values and click **Apply changes**.

The fields are as follows:

- 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 Exports

Note: The **NFS Exports** page 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 2-1 describes the statistics reported by both the **NFS Exports** and **CIFS Shares** pages. Table 2-2 and Table 2-3 describe additional information that is reported.

NFS services gather like-operations into a smaller number of operation classes. Table 2-3 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.

To see historical output for either NFS throughput or IOPS for the selected NFS export, click the associated **History** link.

Table 2-1 Statistics Reported by NFS Exports and CIFS Shares

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 2-2 Additional Information Reported by NFS Exports

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 2-3 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 Shares

Note: The **CIFS Shares** page 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 2-1 describes the statistics reported by both the **NFS Exports** and **CIFS Shares** pages.

CIFS services gather like operations into a smaller number of operation classes. While these classes are largely similar, there are some differences. Table 2-4 summarizes these classes. Table 2-5 describes additional information that is reported.

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.

To see historical output for either NFS throughput or IOPS for the selected CIFS share, click the associated **History** link.

Table 2-4 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

Class	Description
getsecurity	Operations requesting file access permissions
ioctl	Operations performing special filesystem features, such as sparse file handling
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 2-5 Additional Information Reported by CIFS Shares

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

NFS & CIFS Clients

A *NAS client* is a computer running a program that accesses the storage server. NAS clients are known to 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 **NFS & CIFS Clients** page will not be available if neither SGI Samba nor SGI Enhanced NFS are installed.

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

You can change the sorted display by selecting different values and clicking **Update**.

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** page, changing to display the clients in order of aggregate throughput will quickly identify the contributing clients.

To display a detailed view of the NFS and CIFS traffic generated by a particular client, click on the client IP address. 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 page can help diagnose problems, such as hung NFS mounts.

To see historical output for either throughput or IOPS for the selected client, click the associated **History** link.

iSCSI Targets

The **iSCSI Targets** page displays a list of the connected iSCSI initiators, their path targets, the size in GiB, and the transport mechanism. To manage a target, create a new target, or stop the iSCSI service, you must log in. See "iSCSI Targets" on page 36 in Chapter 3, "Management".

Device Failures

The **Storage Device Failures** page displays details of about failed storage 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.

System Monitoring

The **System** tab in the **Monitoring** pane provides access to the following:

- "Alerts" on page 24
- "CPU Utilization" on page 26
- "Network Throughput" on page 27
- "Hardware Inventory" on page 27
- "Software Versions" on page 27

Alerts

The **Alerts** page displays messages from the system logs. These provide informative messages, notifications of unusual events, and error conditions. An icon denotes the priority; hover the mouse pointer over the icon for more information.

Only unacknowledged alerts are displayed unless you make a different selection from the **Options** menu. You must log in to acknowledge alerts.

By default, similar messages are grouped and their number is displayed in the **Count** column. To display the messages individually, select **Show repeats** from the **Options** menu.

You can sort the display by clicking on a column head. For example, to sort the display by alert priority in ascending order of importance (from informative to warning), click on the **Priority** header. To resort in descending order (from warning to informative), click the **Priority** header again.

By default, 25 messages are displayed on the page. To display a different amount, select a different number at the top center of the table. To jump to a specific page, click the page number. For example, Figure 2-3 shows that 10 items will be shown per page, that the alerts are sorted according to those that have appeared most often. The `host avalanche` alert has occurred 112 times and is a critical message that requires action, where as the `SuSEfirewall12` alert has occurred 17 times and is a warning that might require action. None of the message on this page have been acknowledged.

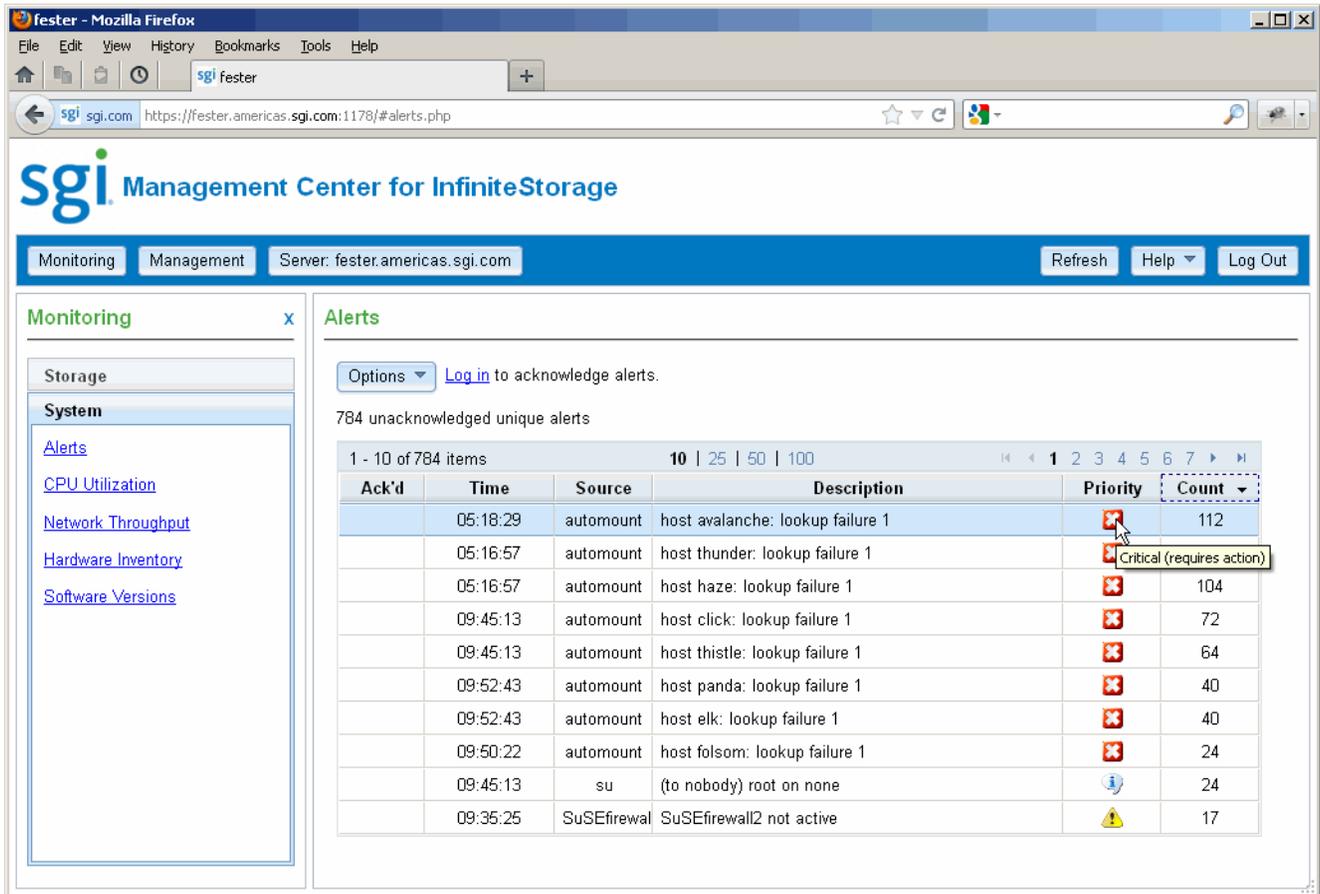


Figure 2-3 Alerts Page

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.

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.

The key is as follows:

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

You can also display historical output for the aggregate CPU utilization by clicking on the **History** link. For more information about historical representation, see "Metrics Collected" on page 12.

Network Throughput

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

If an interface is load-balanced, SMC IS displays throughput for both the bonded interface and its constituent interfaces.

Note: 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.

The throughput displayed is total network throughput (which includes protocol headers), so real data transfer will be somewhat lower than this value.

To see historical data for a given interface, select its **History** link. For more information about historical representation, see "Metrics Collected" on page 12.

To see details about the configuration of a given interface, select its **Configuration** link.

The **NFS Exports** and **CIFS Shares** pages show the amount of real data transferred from a variety of perspectives. See:

- "NFS Exports" on page 18
- "CIFS Shares" on page 21

Hardware Inventory

The **Hardware Inventory** page shows a summary of the hardware configuration, including the CPUs, memory, network controllers, and SCSI disks. The list of SCSI disks includes both the system `root` disk and the configured RAID logical units (LUNs).

To view the contents of the `/proc/cpuinfo` file, click the **Raw output** link.

Software Versions

The **Software Versions** page displays the installed operating system level and the version numbers of key software packages.

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:

- "Storage Management" on page 30
- "System Management" on page 42
- "Operations Management" on page 57

Figure 3-1 shows an example **Management** feature, the **Time & Date** page available under the **System** tab.

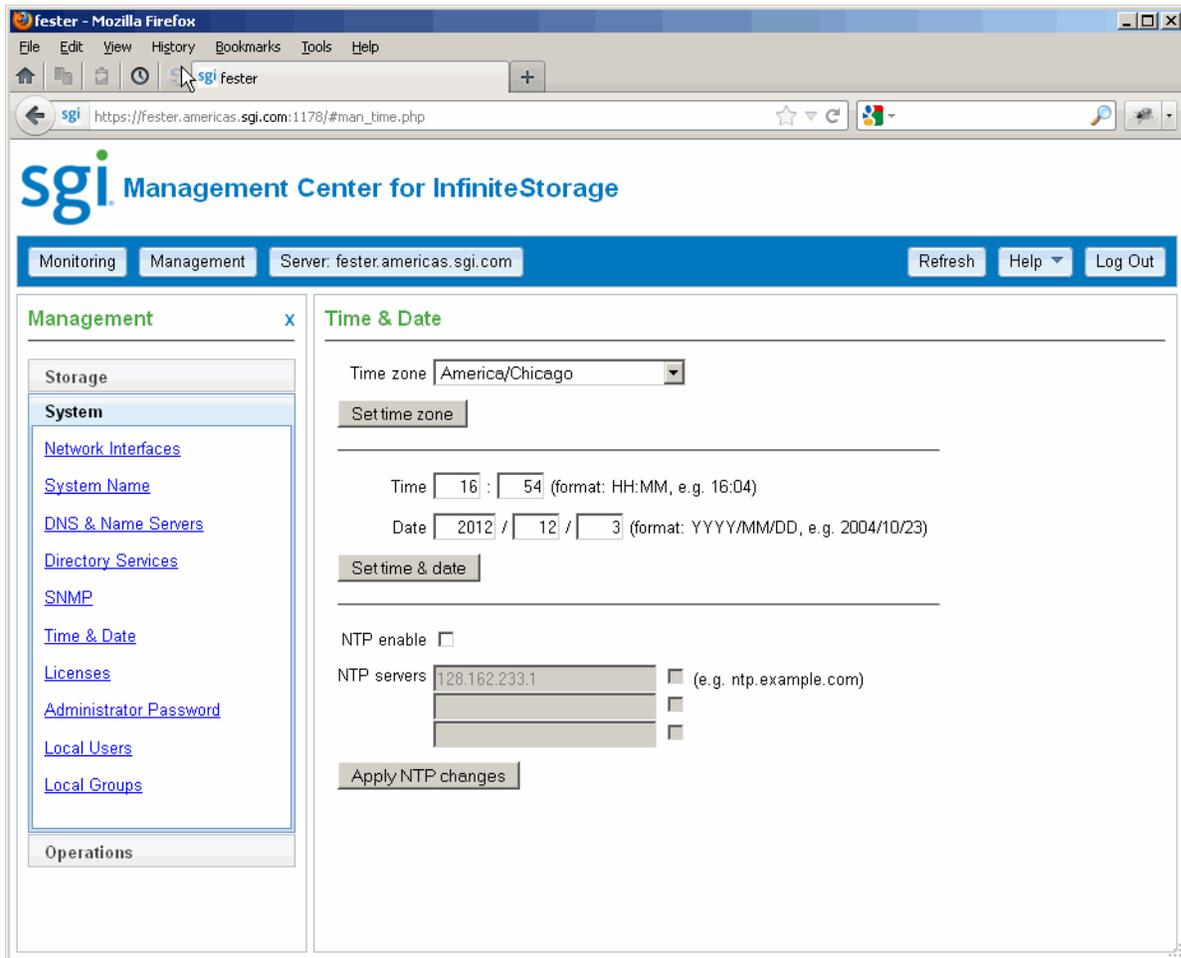


Figure 3-1 Time & Date Page

Storage Management

The **Storage** tab in the **Management** pane lets you manage the following:

- "Filesystems" on page 31
- "NFS Exports" on page 31

- "CIFS Shares" on page 33
- "iSCSI Targets" on page 36
- "Failure Notification" on page 42

Filesystems

To display a brief description of the available local filesystems, select **Filesystems**.

Applicable filesystems found in `/etc/fstab` will be listed along with storage capacity, usage, NFS exports, and CIFS shares.

To unmount a given filesystem, select it and click **Unmount selected**. To select all available filesystems, click the box at the top of the table.

Note: If a directory is currently being exported or shared, then you cannot unmount it.

Unmounted filesystems are listed separately. Devices that are currently in use as backing stores for iSCSI targets cannot be mounted. To mount a given filesystem, select it and click **Mount selected**.

NFS Exports

To configure directories so that they are available for network clients by means of the NFS network protocol, select **NFS Exports**.

This page lists all of the directories that may be exported through NFS.

To specify NFSv4 options, select the **NFSv4** button to display the following fields:

Enable	Specifies whether NFSv4 is enabled (checked) or not. If enabled, an NFS exported directory will be accessible via both NFSv3 and NFSv4. The following fields are only relevant if you have enabled NFSv4.
NFS 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.

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.
	<hr/> Note: SMC IS supports Kerberos 5. You must use a mechanism to synchronize the time between all systems. <hr/>
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	Keeps the existing keytab without changes.
	<hr/> Note: If no keytab is present, a note appears. <hr/>
Update keytab	Changes 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.

To change the export options, select an individual directory name.

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

To export a directory, click its **Export** check box. The current export point is shown next to the **Directory** label. Enter a subdirectory in the text field to specify a new export point and select the desired export options.

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

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

CIFS Shares

To configure directories so that they are available for network clients by means of the CIFS network protocol, select **CIFS Shares**.

This page lists all of the directories that may be shared through CIFS. You can also stop/start the corresponding SMB and NMB services.

To share a directory, select it and click the **Share** box. The current share path is shown next to 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**:

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

Share name	Specifies the name under which the directory 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 directory but cannot modify files or create new files.
Allow guest users	<p>Specifies that users can gain access to the CIFS filesystem without authenticating. Uncheck this option to allow connections only to valid users.</p> <p>By default, the CIFS protocol requires a password for authentication. If configured as an Active Directory client, then the authentication is distributed. See "Active Directory" on page 50.</p>
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	<p>Specifies that symbolic links made by NFS users that point outside of the Samba share will be followed.</p> <hr/> <p>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 <code>/etc/passwd</code>. However, unchecking the box will cause a decrease in performance.</p> <hr/>
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 page; see "Network Interfaces" on page 43.



Restrict to hosts

Specifies the set of hosts that are permitted to access the CIFS share. 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:

```
50.203.15.0/255.255.255.0
```

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**.

iSCSI Targets

This section discusses the following:

- "iSCSI Targets Overview" on page 36
- "Creating iSCSI Targets" on page 38
- "Modifying iSCSI Targets" on page 41
- "The iSCSI Initiator" on page 41

iSCSI Targets Overview

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 3-2 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 3-2 indicate data flow.

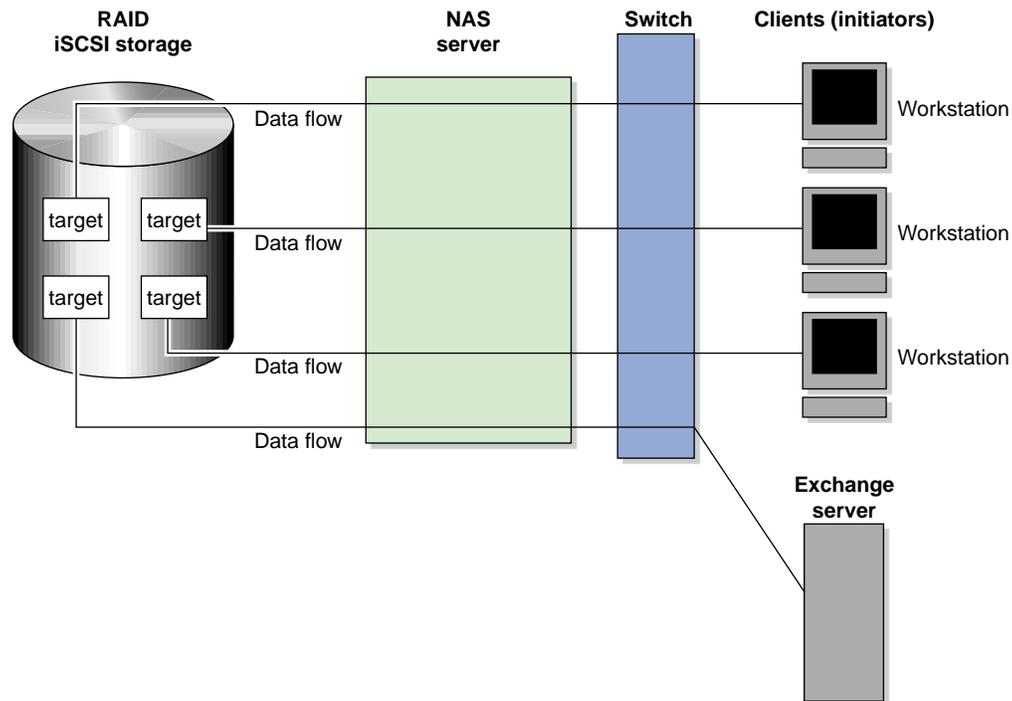


Figure 3-2 iSCSI Storage

You can use 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 following packages for creating targets:

- Red Hat Enterprise Linux (RHEL) 6: `scsi-target-utils`
- SUSE Linux Enterprise Server (SLES) 11: `iscsitarget` or `tgt` (these 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 60.

Creating iSCSI Targets

Perform the following steps to create an iSCSI target:

1. Select the **iSCSI Targets** item from the **Storage** tab.
2. Click **Create target** to access the **Create iSCSI Target** page, which provides a series of pages.
3. On the **Target Name** page, 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**.

4. On the **Target Options** page, 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 page 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.

5. On the **Confirm** page, click **Next** to confirm your choices and create the iSCSI target.
6. The **Finished** page indicates that the iSCSI target has been created. Select **Done**.

To see the initiators and their connected targets, select the **iSCSI Targets** feature from the **Storage** tab in the **Monitoring** pane.

Modifying iSCSI Targets

The **iSCSI Targets** page displays the identifier, path, size, transport mechanism, and client for each configured iSCSI target. To change a given target, click the **Modify** button, which will invoke a series of pages similar to those described in "Creating iSCSI Targets" on page 38. To remove the target, click the **Destroy** button

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 Windows tool to partition and format the target and assign a drive letter:

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

Failure Notification

To configure notification of failed devices, select **Failure Notification**.

For each available platform, enter the interval at which to scan 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" on page 54.

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

System Management

The **System** tab in the **Management** pane lets you manage the following:

- "Network Interfaces" on page 43
- "System Name" on page 48
- "DNS & Name Servers" on page 49
- "Name Service Client" on page 49
- "SNMP" on page 54
- "Time & Date" on page 55
- "Licenses" on page 56
- "Administrator Password" on page 56

- "Local Users and Local Groups" on page 56

Network Interfaces

You can use SMC IS to modify the network interfaces for the system and create a bonded interface. 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 more information, see "Bonded Network Interfaces" on page 45.



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.

Ethernet Network Interfaces

To see the available Ethernet network interfaces and change their parameters, select **Network Interfaces**.

You can change an interface by clicking the **Modify** button for the interface on the **Network Interfaces** page.



Caution: If you configure an incorrect IP address for the management interface, you can make SMC IS inaccessible.

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

To change an interface, click its **Modify** button. You can change the following fields:

- Enable** Enables the interface. You cannot disable the management interface.
- 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.

- 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**.

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 3-3 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.

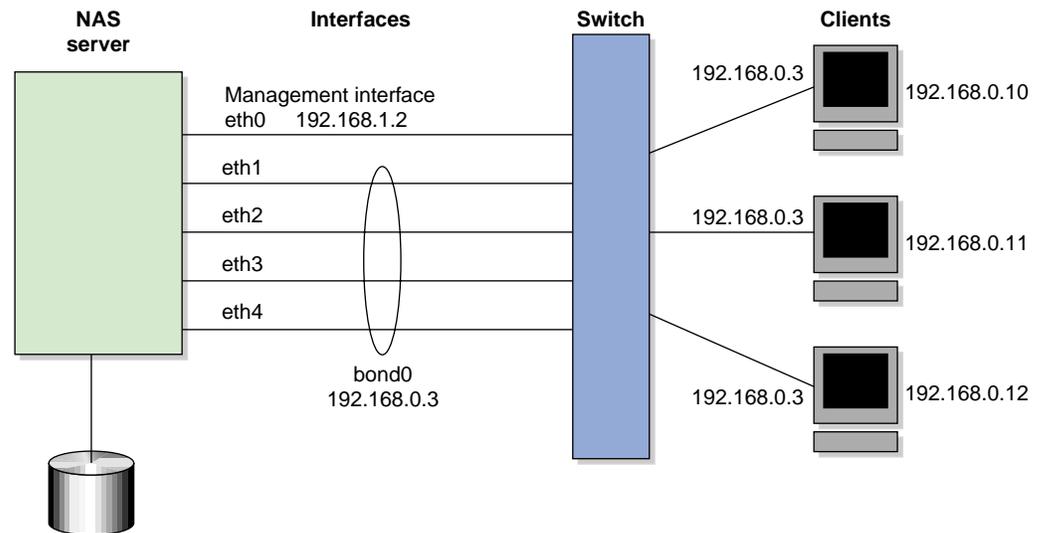


Figure 3-3 Bonded Network Interfaces

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 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 **Create bonded interface** on the **Network Interfaces** page.

The available interfaces are displayed for selection.

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

- | | |
|------------------|---|
| Interface | Specifies the name of the bonded interface. |
| Enable | Enables the bonded interface. |

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.
Available interfaces	Specifies the interfaces to be used.
Bonding mode	<p>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:</p> <ul style="list-style-type: none"> • 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. <p>Your choice depends upon what your switch supports:</p> <ul style="list-style-type: none"> • 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	<p>Specifies how the server chooses which subinterface will send replies:</p> <ul style="list-style-type: none"> • 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.

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

System Name

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

System name Specifies the fully qualified domain name (FQDN) for this storage server. The default system name is `sgiserver`.

Note: After changing the system name, the various SMC IS pages will still display the old hostname. SGI recommends that you reboot the system to complete the name change.

Workgroup Specifies the NetBIOS workgroup to which the machine should belong. The default is `WORKGROUP`. 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** page for `eth0` to configure or modify the management interface. For information on these options, see "Ethernet Network Interfaces" on page 43.

DNS & Name Servers

You can use the **DNS & Name Servers** page to specify how to map hostnames to IP addresses for the system. Click **Edit local hosts file** to access the **Hosts** page, 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 Specifies the domain name or names of the DNS servers that the system uses to provide hostname-to-IP-address translation.

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.

Nameserver # Specifies the IP address for a name server. You can specify up to three IP addresses; if an address you specify is down, the system will use the next one.

Name Service Client

The **Name Service Client** page lets you specify various directory services that manage information associated with the network users, such as mapping user names with user IDs and group names with group IDs.

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** page, 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.

The directory services are:

- "Local Files Only" on page 50
- "Active Directory" on page 50
- "LDAP" on page 52
- "NIS" on page 54

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 Local Groups" on page 56.

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 & Name Servers" on page 49.

The following Active Directory components appear:

Active Directory domain

Specifies the full domain name of the Active Directory.

Note: If you later change the server hostname on which SMC IS runs, you must rejoin the Active Directory domain because the Active Directory Security Identifier (SID) will be changed.

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: ; * & \ ' < > ? []
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:



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

- **RFC 2307 (Microsoft Windows Server 2003 R2).** In order for this to function correctly:
 - 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.

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.

LDAP

Note: This selection requires that the NSD OpenLDAP module is installed.

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 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 page that allows you to verify whether or not you want to commit the changes.

SNMP

SMC IS lets you configure basic Simple Network Management Protocol (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.

The **SNMP** page lets you specify the following information:

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 additional 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.

Time & Date

Use the **Time & Date** page to set the following:

Time zone	Sets the time zone from a drop-down list of options.
Time	Sets the time in hours and minutes, using a 24-hour clock. For example, use <code>16:04</code> for 4:04 PM.
Date	Sets the date by year, month, and day. Use four characters for the year, such as <code>2013</code> .
NTP enable	Enables automatic time synchronization with Network Time Protocol (NTP) using specific NTP servers. 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.

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

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

NTP servers

Specify the servers to be used for the NTP service.

Licenses

The **Licenses** page provides information required to request licenses and lets you add and delete licenses.

Administrator Password

The **Administrator Password** page changes the SMC IS administrator password, which is required to perform server configuration and management. This password is not required to view the pages available from the **Monitoring** pane.

Local Users and Local Groups

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** page to import this file as well as the password file itself.

Operations Management

The **Operations** tab in the **Management** pane lets you do the following:

- "Save/Restore Configuration" on page 57
- "Gather Support Data" on page 58
- "Shut Down System" on page 58

Save/Restore Configuration

The **Save/Restore Configuration** page lets you save the current 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 page lists previously saved configurations, labeled by date. After restoring a configuration, you should restart the system.

Gather Support Data

If there is a problem with the system, SGI Support may request support data in order to find and resolve the problem. The **Gather Support Data** page 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`.

Shut Down System

From the **Shut Down System** page, you can specify to reboot or power down the system in a specified number of minutes.

Troubleshooting

This section discusses the following:

- "Forgotten Password or Corrupt Password File" on page 59
- "The archives Directory is Too Large" on page 59
- "Power Outage and iSCSI" on page 60
- "Manual System Reboot" on page 61
- "Network Configuration Issues" on page 61
- "Reporting Problems to SGI" on page 62

Forgotten Password or Corrupt Password File

If you forget the administrator password or if the **Alerts** 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

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 SMC IS is inaccessible, do the following:

1. Log in via the system console as `root`, such as via IPMI or a monitor/keyboard.
2. Reboot the system:

```
# reboot
```

Network Configuration Issues

If the network configuration is damaged or if the system running SMC IS becomes inaccessible via the network, do the following:

1. Log in via the system console as `root`, such as via IPMI or a monitor/keyboard.
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 "Gather Support Data" on page 58 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** features.

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 and is also known as *Samba*.

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 method of remotely controlling a computer system

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.

Serial ATA (SATA)

Serial advanced technology attachment storage interface connection.

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 SMC IS.

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, 47

A

About menu selection, 9
access operation, 20
Active Directory, 50, 56, 57
admin server, 32
Administrator Password page, 56
alert numbers, 7
alerts, 24
Alerts menu selection, 5
archives, 12
archives directory size, 59
autonegotiate, 45

B

blue color in graphs, 13
bonded network interface, 43
bonded network interfaces, 45
bonding mode, 46

C

cancel operation, 21
change/notify operation, 21
CHAP authentication, 39
CIFS, 21
 client number, 7
 configuration, 33
 iSCSI and, 37
CIFS authentication, 50, 54
CIFS page, 21

CIFS Shares menu selection, 5
CIFS Shares page, 16
clients, 22
close operation, 21
colors in graphs, 13
commit operation, 20
community string, 55
corrupt password file, 59
CPU utilization, 7, 26
create/open operation, 21
current time, 13

D

data flow color-coding in graphs, 14
data reduction process, 12
dedicated network interface, 44
default network gateway, 48
device failure notification, 42
Device Failures menu selection, 5, 23
DHCP, 44
disk
 IOPS, 16
 operations, 15
 quotas, 16
 space, 7, 15
 throughput, 7, 16
 throughput, monitoring, 16
Disk IOPS page, 15
Disk menu selection, 4
Disk Quotas page, 15
DNS & Name Servers, 49
domain, 32, 38
Domain Search, DNS and Hostnames page, 49
duplex option, 45
dynamic bonding mode, 46

E

- /etc/krb5.keytab, 33
- Exchange Server as an iSCSI initiator , 36
- export options, 33

F

- Fedora Directory Server, 53
- filesystem listing, 31
- findfirst/next operation, 21
- Firefox, 1
- flush operation, 21
- fsinfo operation, 20
- full-duplex, 45

G

- gateway, 48
- Gather Support Data page, 58
- getattr operation
 - CIFS, 21
 - NFS, 20
- getsecurity operation, 22
- group quotas, 18

H

- half-duplex , 45
- hard limit, 17
- hardware inventory, 27
- historic time, 13
- historical status of a parameter, 7
- History menu selection, 7

I

- identifier for target, 38

70

- IEEE 802.3ad standard, 47
- Import Users option, Local Users page, 57
- InfiniBand throughput, 7
- initiator for iSCSI, 36, 41
- inode_mods operation, 20
- input load balancing, 46
- interface overview, 2
- Internet Explorer, 1
- Internet Small Computer Systems Interface
 - See "iSCSI", 36
- interrupt time, 26
- ioctl operation, 22
- IOPS, 15, 19, 22
 - CIFS, 19
 - NFS, 19
- IOPS menu selection, 4
- IP address, 47
- IP header, 47
- iqn, 38
- iSCSI
 - client number, 7
 - domain, 38
 - identifier, 38
 - initiator, 36, 41
 - network, 36
 - NFS and CIFS, 37
 - protocol, 36
 - qualified name, 38
 - re-exporting targets, 37
 - target, 36
 - targets, 36, 38
- iSCSI and power outage, 60
- iSCSI Initiator program, 40
- iSCSI Targets menu selection, 5, 23

J

- JavaScript, 1

K

- KDC, 32
- Kerberos, 32
- key distribution center, 32
- keytab, 32

L

- Layer 2 (MAC address), 47
- Layer 3 (IP header), 47
- LDAP, 56
- LDAP (lightweight directory access protocol), 52
- Licenses page, 56
- load average, 7
- load balancing, 46, 47
- local users and groups, 56
- lockd operation, 20
- lockd_granted operation, 20
- lockd_share operation, 20
- Log Out menu selection, 11
- lookup operation, 20
- LSI MegaRAID, 23, 42

M

- MAC address header, 47
- mail store and iSCSI, 36
- main menu, 3
- management interface, 43
- Management menu selection, 5
- metadata operations, 15
- metrics
 - type collected, 12
- MiB vs MB, 13
- misc operation, 22
- Modify option, 43
- Monitoring menu selection, 4
- Monitoring pane example, 12
- monitoring performance, 11

- move operation, 22
- mutual CHAP authentication, 40

N

- name service client, 49
- nameserver, 49
- network configuration issues, 61
- Network Information Service (NIS), 54
- network interface
 - bonded, 45
 - management, 43
 - standalone, 43
- network interface configuration, 43
- network throughput, 7, 23, 27
- Network Time Protocol (NTP), 56
- NFS, 18, 31
 - client number, 7
 - custom definition, 33
 - export options, 33
 - iSCSI and, 37
- NFS & CIFS Clients category, 16
- NFS & CIFS Clients menu selection, 5
- NFS Exports menu selection, 4
- NFS Exports page, 16
- NFS page, 18
- NFS serving domain, 32
- NFSv4 enabling, 31
- NIS, 54
- NTP, 56
- NTP Time Synchronization, Date and Time page, 56
- number of alerts, 7
- number of users, 7

O

- OpenLDAP Server, 53
- operation
 - CIFS, 21, 22

- operation classes, 18
- operations by type, 19
- output load balancing, 46, 47
- overview, 1

P

- password
 - initial, 2
- password problems, 59
- performance archives, 12
- performance monitoring, 11
- port speed, 44
- power outage and iSCSI, 60
- principal user (Kerberos) , 32

Q

- quotas
 - disk, 16
 - set group, 18
 - set user, 17, 18

R

- re-exporting iSCSI targets with NFS or CIFS, 37
- read block sizes, 19
- read operation
 - CIFS, 22
 - NFS, 20
- readdir operation, 20
- readdirplus operation, 20
- realm, 32
- red color in graphs, 13
- remove operation, 20
- reporting problems to SGI, 63
- resources, 4

S

- Samba schema, 53
- Save/Restore Configuration page, 57
- secret for CHAP authentication, 39
- server configuration and management, 29
- service times, 19, 22
- serving domain for NFS, 32
- setattr operation
 - CIFS, 22
 - NFS, 20
- Share Options, CIFS configuration, 33
- Shut Down System page, 58
- Site Map menu selection, 9
- SNMP, 54
- soft limit, 17
- ssmc_admin, 59
- standalone network interface, 43
- static bonding mode, 46
- Static option, 44
- Storage Device Failures page, 23
- Storage menu selection, 4
- subnet mask, 47
- Summary menu selection, 7
- Summary page, 8
- system alerts, 24
- system console, 61
- system logs, 24
- System Name, 48
- system name, 48
- system time, 26
- system uptime, 7

T

- target for iSCSI, 36
 - CHAP authentication, 39
 - creating, 38
 - identifier, 38
 - re-exporting with NFS or CIFS, 37

- username, 39
- target name, 38
- throughput, 16
 - CIFS, 19
 - network, 27
 - NFS, 19
- Throughput menu selection, 4
- Time & Date page, 55
- trap destination, 54
- troubleshooting, 59
 - archives directory size, 59
 - password issues, 59
 - power outage and iSCSI, 60
 - reporting problems, 63
 - SMC IS is inaccessible
 - network configuration issues, 61
 - rebooting the system, 61

U

- unit measures, 13
- uptime of system, 7
- user numbers, 7
- user quotas, 17, 18
- user time, 26

V

- /var/lib/pcp-storage/archives, 59
- /var/lib/pcp-storage/archives directory, 12
- /var/lib/ssmc/alerts/archive, 26
- versions, 27

W

- wait time, 26
- web browsers, 1
- write block sizes, 19
- write operation/primary>, 22
- write_async operation, 20
- write_sync operation, 20

X

- xattr operation
 - NFS, 20