

Compaq StorageWorks

HSG60 ACS Solution Software Version 8.6 for SGI Irix

Installation and Configuration Guide

First Edition (June 2001)
Part Number: AA-RMCEB-TE
Compaq Computer Corporation

© 2001 Compaq Computer Corporation.

Compaq, the Compaq logo, and StorageWorks Registered in U.S. Patent and Trademark Office.

Microsoft, MS-DOS, Windows, Windows NT and Windows 2000 are trademarks of Microsoft Corporation.

UNIX is a trademark of The Open Group.

All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Compaq service tool software, including associated documentation, is the property of and contains confidential technology of Compaq Computer Corporation. Service customer is hereby licensed to use the software only for activities directly relating to the delivery of, and only during the term of, the applicable services delivered by Compaq or its authorized service provider. Customer may not modify or reverse engineer, remove, or transfer the software or make the software or any resultant diagnosis or system management data available to other parties without Compaq's or its authorized service provider's consent. Upon termination of the services, customer will, at Compaq's or its service provider's option, destroy or return the software and associated documentation in its possession.

Printed in the U.S.A.

HSG60 ACS Solution Software Version 8.6 for SGI Irix
Installation and Configuration Guide
First Edition (June 2001)
Part Number: AA-RMCEB-TE

Contents

About this Guide

Text Conventions	xiv
Symbols in Text	xv
Symbols on Equipment	xv
Rack Stability	xvi
Installation and Setup of the CD-ROM	xvi
Configuration Flowchart	xvii
Getting Help	xx
Compaq Technical Support	xx
Compaq Website	xxi
Compaq Authorized Reseller	xxi

Chapter 1

Planning a Subsystem

Defining the Subsystems	1-1
Controller Designations A and B	1-2
Controller Designations “This Controller” and “Other Controller”	1-3
Failover Mode	1-4
Transparent Failover Mode	1-4
Selecting a Cache Mode	1-7
Read Caching	1-8
Read-Ahead Caching	1-8
Write-Back Caching	1-8
Write-Through Caching	1-8
Enabling Mirrored Caching	1-9
The Command Console LUN	1-9
Determining the Address of the CCL	1-10
Enabling/Disabling the CCL in SCSI-2 Mode	1-10
Enabling the CCL	1-10

Disabling the CCL	1-10
Enabling/Disabling the CCL in SCSI-3 Mode	1-10
Connections	1-11
Naming Connections	1-11
Numbers of Connections	1-11
Assigning Unit Numbers	1-14
Matching Units to Host Connections in Transparent Failover Mode	1-14
Assigning Unit Numbers Depending on SCSI_VERSION	1-16
Assigning Unit Numbers and Host Connections in SCSI-2 Mode	1-16
Selective Storage Presentation	1-17
Restricting Host Access in Transparent Failover Mode	1-17
Restricting Host Access by Separate Links	1-17
Worldwide Names (Node IDs and Port IDs)	1-17
Restoring Worldwide Names (Node IDs)	1-18
Unit Worldwide Names (LUN IDs)	1-19
HSG60 Setup	1-19
Arbitrated Loop Setup	1-19
Fabric Setup	1-20
Host System Setup	1-22

Chapter 2

Planning Storage

Where to Start	2-2
Configuration Rules	2-3
Device PTL Addressing Convention	2-3
Determining Storage Requirements	2-6
Choosing a Container Type	2-6
Creating a Storageset Profile	2-8
Storageset Planning Considerations	2-10
Stripeset Planning Considerations	2-10
Mirrorset Planning Considerations	2-13
RAIDset Planning Considerations	2-15
Striped Mirrorset Planning Considerations	2-17
Storageset Expansion Considerations	2-19
Partition Planning Considerations	2-19
Defining a Partition	2-20
Guidelines for Partitioning Storagesets and Disk Drives	2-20
Changing Characteristics through Switches	2-21
Enabling Switches	2-21
Changing Switches	2-21
Storageset and Partition Switches	2-22

RAIDset Switches	2-22
Mirrorset Switches	2-22
Partition Switches	2-23
Initialization Switches	2-23
Chunk Size	2-23
Increasing the Request Rate	2-24
Increasing the Data Transfer Rate	2-25
Increasing Sequential Write Performance	2-26
Save Configuration	2-26
Destroy/Nodestry	2-26
Geometry	2-27
Unit Switches	2-27
Storage Maps	2-27
Creating a Storage Map	2-27
Using the LOCATE Command to Find Devices	2-29

Chapter 3

Preparing the Host System

Making a Physical Connection	3-1
Installing the Host Bus Adapter	3-1
Loading the StorageWorks Software onto the Hard Disk	3-2
Copying the File Archive from CD-ROM to Temporary Storage	3-2
Executing the Installation Script	3-3
Host Device Initialization	3-4
Arbitrated Loop Device Name Convention	3-5
Fabric Device Name Convention	3-6
Creating the Disk Label	3-6
Partitioning the Storage	3-8
Creating Filesystems	3-10

Chapter 4

Installing and Configuring the HS-Series Agent

Why Use StorageWorks Command Console (SWCC)?	4-1
About the Network Connection for the Agent	4-3
Installing and Configuring the Agent	4-5
Running the Agent	4-8
Reconfiguring the Agent	4-8
Removing the Agent	4-8

Chapter 5

Configuration Procedures

Establishing a Local Connection	5-2
Setting Up a Single Controller	5-3
Power Up and Establish Communication	5-3
Cabling a Single controller	5-3
Configuring a Single Controller Using CLI	5-4
Verify the Node ID and Check for Any Previous Connections	5-4
Configure Controller Settings	5-5
Restart the Controllers	5-6
Set Time and Verify All Commands	5-6
Plug in the FC Cable and Verify Connections	5-7
Repeat Procedure for Each Host Adapter Connections	5-8
Verify Installation	5-8
Setting Up a Controller Pair	5-8
Power Up and Establish Communication	5-8
Cabling a Controller Pair	5-9
Configuring a Controller Pair Using CLI	5-10
Verify the Node ID and Check for Any Previous Connections	5-10
Configure Controller Settings	5-11
Restart the Controller	5-11
Set Time and Verify All Commands	5-12
Plug in the FC Cable and Verify Connections	5-13
Repeat Procedure for Each Host Adapter	5-14
Verify Installation	5-14
Configuring Devices	5-14
Configuring a Stripeset	5-15
Configuring a Mirrorset	5-15
Configuring a RAIDset	5-16
Configuring a Striped Mirrorset	5-17
Configuring a Single-Disk (JBOD) Unit	5-17
Configuring a Partition	5-18
Assigning Unit Numbers and Unit Qualifiers	5-19
Assigning a Unit Number to a Storageset	5-19
Assigning a Unit Number to a Single (JBOD) Disk	5-19
Assigning a Unit Number to a Partition	5-20
Configuration Options	5-20
Changing the CLI Prompt	5-20
Mirroring cache	5-20
Adding Disk Drives	5-20
Adding a Disk Drive to the Spareset	5-21

Removing a Disk Drive from the Spareset.	5-21
Enabling Autospare	5-22
Deleting a Storageset	5-22
Changing Switches for a Storageset or Device	5-23
Displaying the Current Switches.	5-23
Changing RAIDset and Mirrorset Switches	5-23
Changing Device Switches	5-23
Changing Initialize Switches.	5-23
Changing Unit Switches	5-23

Chapter 6

Configuration Example Using CLI

CLI Configuration Example.	6-5
------------------------------------	-----

Chapter 7

Backing Up the Subsystem, Cloning Data for Backup, and Moving Storagesets

Backing Up the Subsystem Configuration	7-1
Cloning Data for Backup.	7-2
Moving Storagesets	7-5

Appendix A

Subsystem Profile Templates

Storageset Profile.	A-2
Storage Map Template 1 for the Model 4214R Disk Enclosure	A-4
Storage Map Template 2 for the Model 4254 Disk Enclosure	A-5
Storage Map Template 3 for two Model 4214R Disk Enclosures.	A-6

Appendix B

Installing, Configuring, and Removing the Client

Why Install the Client?	B-1
Before You Install the Client.	B-2
Installing the Client	B-2
Troubleshooting the Client Installation	B-3
Invalid Network Port Assignments	
During Installation	B-3
“There is no disk in the drive” Message.	B-4
Adding the Storage Subsystem and its Host to the Navigation Tree.	B-4
Removing the Command Console Client	B-6
Where to Find Additional Information	B-7
About the User Guide.	B-7
About the Online Help	B-8

Glossary

Index

Tables

Table 1-1 Unit Assignments and SCSI_VERSION	1-16
Table 1-2 Controller Parameters and Recommended Settings for Arbitrated Loop	1-19
Table 1-3 Controller Parameters and Recommended Settings for Fabric	1-21
Table 2-1 A Comparison of Container Types	2-7
Table 2-2 Example Chunk Sizes	2-25
Table 4-1 SWCC Features and Components	4-2
Table 4-2 Installation and Configuration Overview	4-2
Table 4-3 Information Required for Install Program Prompts	4-5

Figures

Figure 1	General configuration flowchart	xviii
Figure 2	Configuring storage with the CLI	xix
Figure 3	Configuring storage with SWCC	xx
Figure 1-1	Location of controllers and cache modules in a Model 2100 enclosure	1-2
Figure 1-2	“This controller” and “other controller” for the Model 2100 enclosure	1-3
Figure 1-3	Transparent failover—normal operation.	1-5
Figure 1-4	Transparent failover—after failover from controller B to controller A	1-6
Figure 1-5	Mirrored caching	1-9
Figure 1-6	Connections in separate-link, transparent failover mode configurations	1-12
Figure 1-7	Connections in single-link, transparent failover mode configurations	1-13
Figure 1-8	LUN presentation to hosts, as determined by offset	1-15
Figure 1-9	Limiting host access in transparent failover mode	1-17
Figure 1-10	Placement of the worldwide name label on the Model 2100 enclosure	1-18
Figure 2-1	PTL naming convention	2-4
Figure 2-2	PTL addressing in a single-bus configuration, using two shelves	2-4
Figure 2-3	PTL addressing in a dual-bus configuration	2-5
Figure 2-4	Mapping a unit to physical disk drives	2-5
Figure 2-5	Container types	2-7
Figure 2-6	An example storage set profile.	2-9
Figure 2-7	A 3-member RAID 0 stripeset (example 1)	2-11
Figure 2-8	A 3-member RAID 0 stripeset (example 2)	2-12
Figure 2-9	Mirrorsets maintain two copies of the same data	2-14
Figure 2-10	Mirrorset example 2	2-14
Figure 2-11	A five member RAIDset using parity	2-16
Figure 2-12	Striped mirrorset (example 1)	2-18
Figure 2-13	Striped mirrorset (example 2)	2-18
Figure 2-14	One example of a partitioned single-disk unit.	2-19
Figure 2-15	Chunk size larger than the request size	2-24
Figure 2-16	Chunk size smaller than the request size.	2-25
Figure 2-17	Example storage map-single bus	2-28
Figure 2-18	Example storage map-dual bus	2-29
Figure 4-1	An example of a network connection	4-4

Figure 5-1	Maintenance port connection	5-2
Figure 5-2	Single controller cabling	5-4
Figure 5-3	Controller pair failover cabling.	5-9
Figure 6-1	Example storage map	6-2
Figure 6-2	Example system	6-3
Figure 6-3	Example system	6-4
Figure 6-4	Example virtual layout from the hosts' point of view	6-4
Figure 7-1	Steps the CLONE utility follows for duplicating unit members.	7-3
Figure B-1	Navigation window	B-5
Figure B-2	Navigation window showing storage host system "Atlanta".	B-5
Figure B-3	Navigation window showing expanded "Atlanta" host icon	B-6

About this Guide

This guide provides installation and configuration instructions for operation of the HSG60 ACS Solution Software Version 8.6 for SGI Irix.

Thank You for selecting a *Compaq StorageWorks* RAID Array subsystem for your growing storage needs. StorageWorks RAID subsystems are designed to support the most popular computer platforms in the industry. The solution software that accompanies this software kit enables the storage subsystem to operate with your chosen platform.

Please fill out and return the Registration Card included in this kit. This information is used by Compaq to provide notification services to its customers. You can also register online at:

<http://www.compaq.com/products/registration>

This guide describes:

- Considerations while planning a configuration
- Configuration procedures

This book does not contain information about the operating environments to which the controller may be connected; nor does it contain detailed information about subsystem enclosures or their components. See the documentation that accompanied these peripherals for information about them.

Text Conventions

This document uses the following conventions:

Keys	Keys appear in boldface. A plus sign (+) between two keys indicates that they should be pressed simultaneously.
USER INPUT, COMMANDS	* User input and commands appear in this typeface and in uppercase.
<i>type of user input</i>	The type of user input, such as <i>device-name</i> appear in italics.
<i>FILENAMES</i>	* File names appear in uppercase italics.
Menu Options, Dialog Box Names	These elements appear in initial capital letters.
DIRECTORY NAMES, and DRIVE NAMES	* These elements appear in uppercase.
Enter	When you are instructed to enter information, type the information and press the Enter key.
<i>Script</i>	Script names appear in upper and lower case in <i>this typeface</i> .
<i>SWITCHES</i>	Elements designated as switches appear in uppercase italics.

* UNIX commands are case sensitive and will not appear in uppercase.

Symbols in Text

The symbols found in this guide have the following meanings:



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life.



CAUTION: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or loss of information.

IMPORTANT: Text set off in this manner presents clarifying information or specific instructions.

Symbols on Equipment



Any surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of injury from electrical shock hazards, do not open this enclosure.



Any RJ-45 receptacle marked with these symbols indicates a Network Interface Connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. If this surface is contacted, the potential for injury exists.

WARNING: To reduce the risk of injury from a hot component, allow the surface to cool before working with it.



Power Supplies or Systems marked with these symbols indicate that the equipment is supplied by multiple sources of power.
WARNING: To reduce the risk of injury from electrical shock, remove all power cords to completely disconnect power from the system.



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.
WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manual material handling.

Rack Stability



WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
 - The full weight of the rack rests on the leveling jacks.
 - The stabilizing feet are attached to the rack if it is a single rack installation.
 - The racks are coupled together in multiple rack installations.
 - A rack may become unstable if more than one component is extended for any reason. Extend only one component at a time.
-

Installation and Setup of the CD-ROM

Contents of CD-ROM:

- Solution Software
- SWCC Agent and Client software
- Documentation

NOTE: Refer to the platform specific release notes for detailed solution kit contents and release specific information.

Configuration Flowchart

A three-part flowchart is shown on the following pages. Refer to these charts while configuring a new storage subsystem:

- Figure 1 on page xviii shows the start of the configuration process.
- Figure 2 on page xix shows how to configure storage with the command line interpreter (CLI), which is the low-level interface to the controller.
- Figure 3 on page xx shows how to configure storage using StorageWorks Command Console (SWCC), which is the graphical user interface to the controller.

All references in the flowcharts pertain to pages in this guide, unless otherwise indicated.

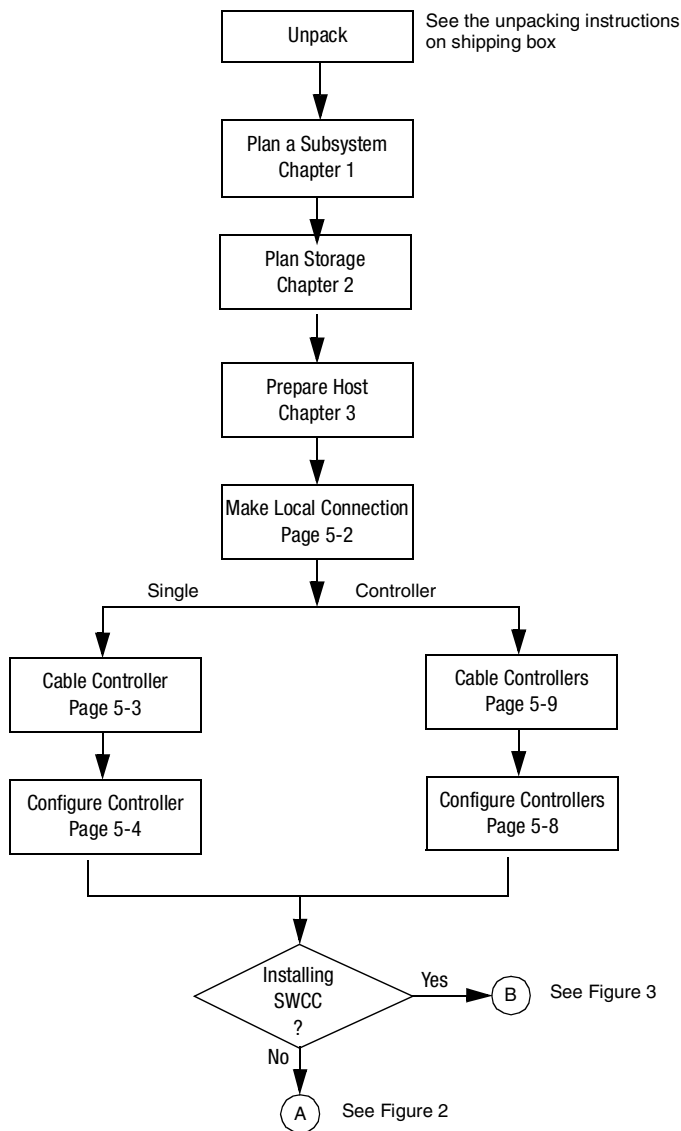


Figure 1. General configuration flowchart

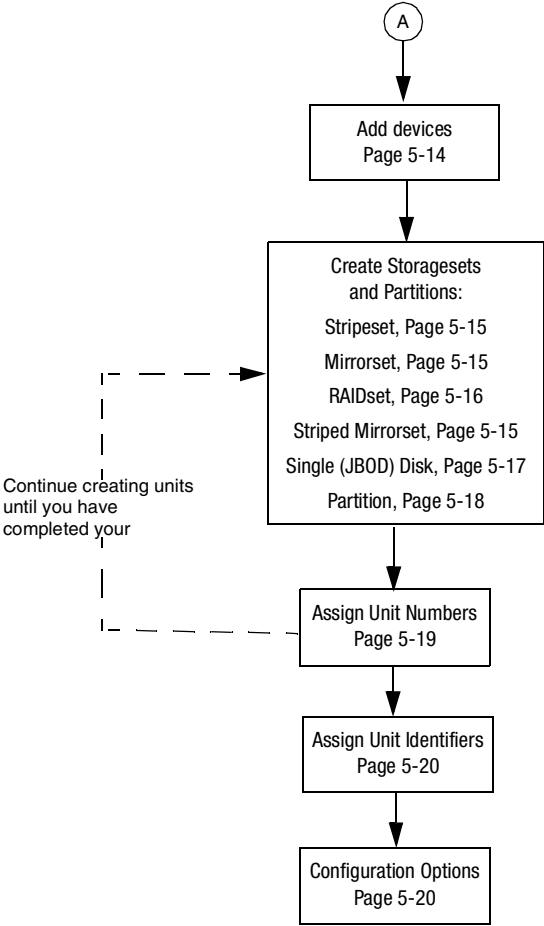


Figure 2. Configuring storage with the CLI

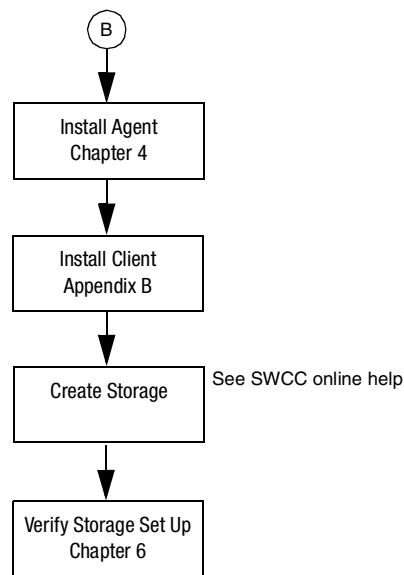


Figure 3. Configuring storage with SWCC

Getting Help

If you have a problem and have exhausted the information in this guide, you can receive further information and other help in the following locations.

Compaq Technical Support

A technical support specialist can help you diagnose a problem or guide you to the next step in the warranty process.

In North America, call the Compaq Technical Phone Support Center at 1-800-OK-COMPAQ. This service is available 24 hours a day, 7 days a week.

NOTE: For continuous quality improvement, calls may be recorded or monitored.

Outside North America, call the nearest Compaq Technical Support Phone Center. Telephone numbers for worldwide Technical Support Centers are listed on the Compaq website. Access the Compaq website:

<http://www.compaq.com>

Be sure to have the following information available before you call Compaq:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level
- Detailed, specific questions

Compaq Website

The Compaq website has latest information on this product as well as the latest drivers. Access the Compaq website at:

<http://www.compaq.com/storage>

Compaq Authorized Reseller

For the name of the nearest Compaq Authorized Reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the Compaq website for locations and telephone numbers.

Chapter 1

Planning a Subsystem

The information in this chapter will help you plan how to configure the subsystem. Refer Chapter 2 to plan the types of storage containers you need.

IMPORTANT: This chapter frequently references the command line interface (CLI). For the complete syntax and descriptions of the CLI commands, see the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

The following information is included in this chapter:

- “Defining the Subsystems,” page 1–1
- “Failover Mode,” page 1–4
- “Selecting a Cache Mode,” page 1–7
- “Enabling Mirrored Caching,” page 1–9
- “The Command Console LUN,” page 1–9
- “Connections,” page 1–11
- “Assigning Unit Numbers,” page 1–14
- “Selective Storage Presentation,” page 1–17
- “Worldwide Names (Node IDs and Port IDs),” page 1–17

Defining the Subsystems

The HSG60 controller uses the Model 2100 enclosure.

Controller Designations A and B

A hardware device that, with proprietary software, facilitates communications between a host and one or more devices organized in an array. HSG60 family controllers are examples of array controllers.

The terms A, B, “this controller,” and “other controller,” respectively are used to distinguish one controller from another in a two-controller (also called dual-redundant) subsystem. These terms are described more thoroughly in the following sections.

Controllers and cache modules are designated either A or B depending on their location in the enclosure, as shown in Figure 1-1.




Figure 1-1. Location of controllers and cache modules in a Model 2100 enclosure

Controller Designations “This Controller” and “Other Controller”

Some CLI commands use the terms “this” and “other” to identify one controller or the other in a dual-redundant pair. These designations are a shortened form of “this controller” and “other controller.” These terms are defined as follows:

- **“this controller”**—the controller that is the focus of the CLI session. “This controller” is the controller to which the maintenance terminal is attached and through which the CLI commands are being entered. “This controller” can be shortened to “this” in CLI commands.
- **“other controller”**—the controller that is not the focus of the CLI session and through which CLI commands are not being entered. “Other controller” can be shortened to “other” in CLI commands.

Figure 1-2 shows the relationship between “this controller” and “other controller.”




Figure 1-2. “This controller” and “other controller” for the Model 2100 enclosure

Failover Mode

Failover is a way to keep the storage array available to the host if one of the controllers becomes unresponsive. A controller can become unresponsive because of a hardware failure, such as a controller. Failover keeps the storage array available to the hosts by enabling the surviving controller to take over total control of the subsystem.

Consult the release notes for supported failover modes on SGI Irix.

Transparent Failover Mode

Transparent failover mode has the following characteristics:

- Hosts do not know failover has taken place.
- Units are divided between host ports 1 and 2.

A unit or storage unit is a physical or virtual device of the subsystem. It is assigned a logical unit number (LUN) and is managed by the HSG60 controller. The unit is presented to a server through the Fibre Channel bus and servers host bus adapter. Disks that are set up as independent disks (JBODs) or RAIDsets are referred to as storagesets. Storagesets are units.

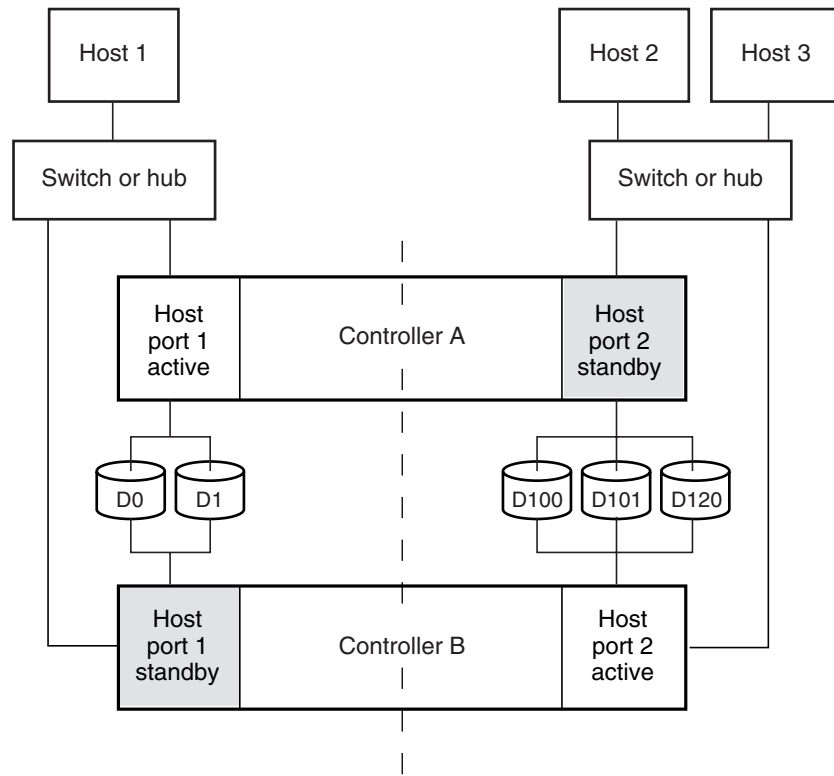
In transparent failover mode, the controller A host port 1 and the controller B host port 1 must be on the same Fibre Channel link. Controller A host port 2 and controller B host port 2 must also be on the same Fibre Channel link. Depending on operating system restrictions and requirements, the port 1 link and the port 2 link can be separate links, or they can be the same link.

At any one time, host port 1 is active on only one controller, and host port 2 is active on only one controller. The other ports are in standby mode. In normal operation, controller A port 1 is active, and controller B port 2 is active. A representative configuration is shown in Figure 1-3. The active and standby ports share port identity, enabling the standby port to take over for the active one. If one controller fails, its companion controller (known as the surviving controller) takes control by making both its host ports active, as shown in Figure 1-4.

Units are divided between the host ports:

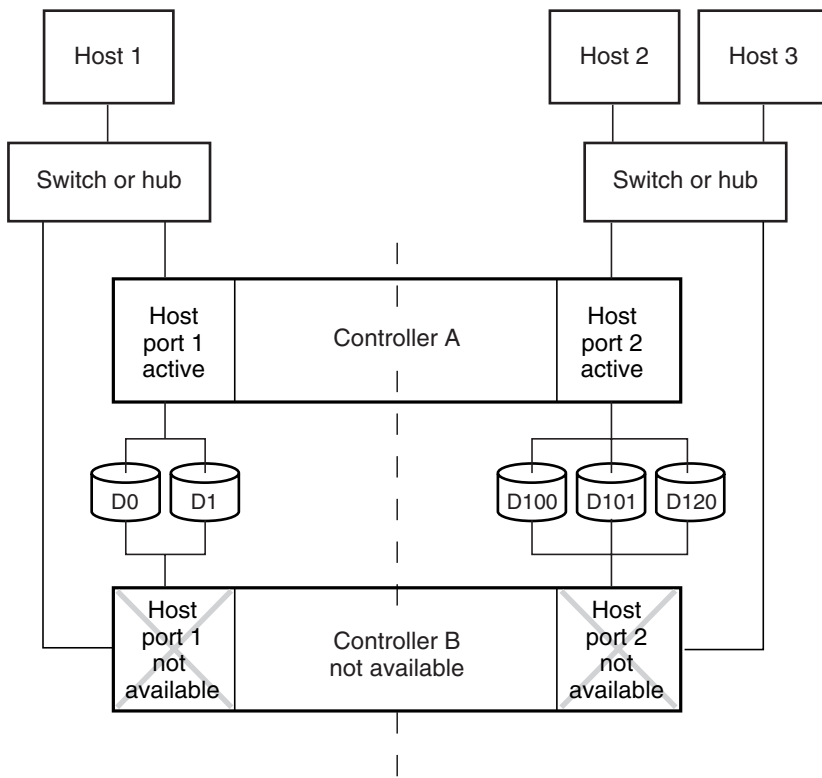
- Units 0-99 are on host port 1 of both controllers (but accessible only through the active port).
- Units 100-199 are on host port 2 of both controllers (but accessible only through the active port).

The limitation of transparent failover is that it only compensates for a controller failure, and not for failures of either the Fibre Channel link or host Fibre Channel adapters.



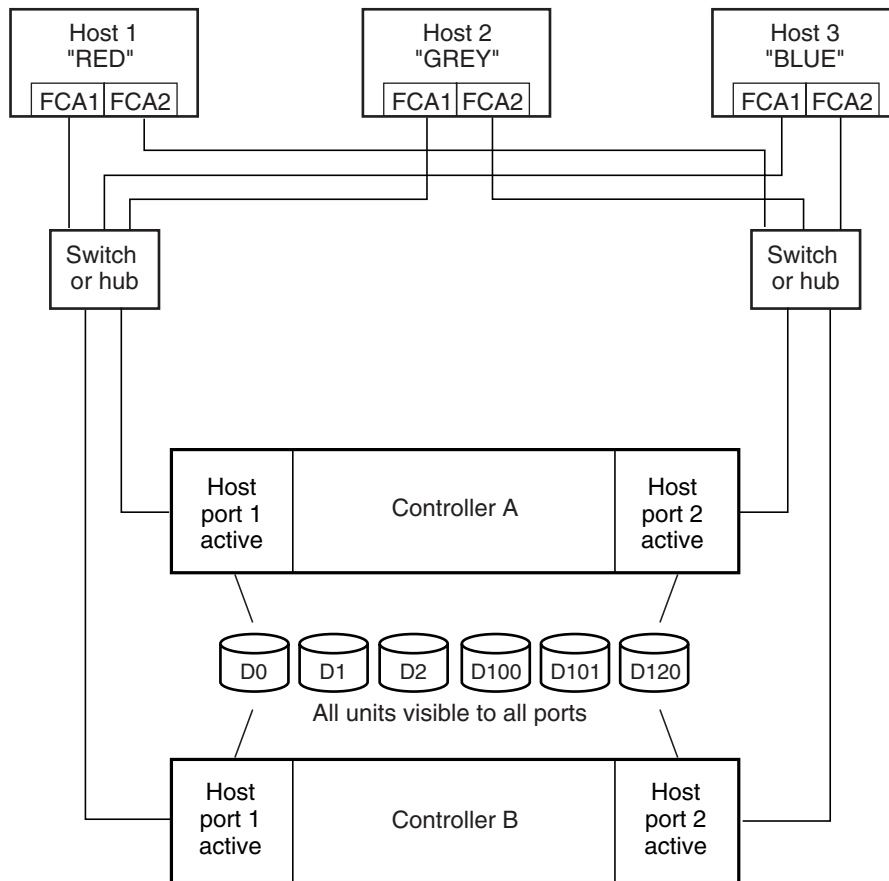
CXO7036A

Figure 1-3. Transparent failover—normal operation



CXO7035A

Figure 1-4. Transparent failover—after failover from controller B to controller A



NOTE: FCA = Fibre Channel Adapter

CXO7094B

Selecting a Cache Mode

The cache module supports read, read-ahead, write-through, and write-back caching techniques. The cache technique is selected separately for each unit. For example, you can enable only read and write-through caching for some units while enabling only write-back caching for other units.

Read Caching

When the controller receives a read request from the host, it reads the data from the disk drives, delivers it to the host, and stores the data in its cache module. Subsequent reads for the same data will take the data from cache rather than accessing the data from the disks. This process is called read caching.

Read caching can improve response time to many of the host's read requests. By default, read caching is enabled for all units.

Read-Ahead Caching

During read-ahead caching, the controller anticipates subsequent read requests and begins to prefetch the next blocks of data from the disks as it sends the requested read data to the host. This is a parallel action. The controller notifies the host of the read completion, and subsequent sequential read requests are satisfied from the cache memory. By default, read-ahead caching is enabled for all units.

Write-Back Caching

Write-back caching improves the subsystem's response time to write requests by enabling the controller to declare the write operation complete as soon as the data reaches cache memory. The controller performs the slower operation of writing the data to the disk drives at a later time.

By default, write-back caching is enabled for all units, but only if there is a backup power source for the cache modules (either batteries or an uninterruptable power supply).

Write-Through Caching

Write-through caching is enabled when write-back caching is disabled. When the controller receives a write request from the host, it places the data in its cache module, writes the data to the disk drives, then notifies the host when the write operation is complete. This process is called write-through caching because the data actually passes through—and is stored in—the cache memory on its way to the disk drives.

Enabling Mirrored Caching

In mirrored caching, half of each controller's cache mirrors the companion controller's cache, as shown in Figure 1-5.

The total memory available for cached data is reduced by half, but the level of protection is greater.




Figure 1-5. Mirrored caching

Before enabling mirrored caching, make sure the following conditions are met:

- Both controllers support the same size cache.
- Diagnostics indicate that both caches are good.
- No unit errors are outstanding, for example, lost data or data that cannot be written to devices.
- Both controllers are started and configured in failover mode.

The Command Console LUN

StorageWorks Command Console (SWCC) software communicates with HSG60 controllers through an existing storage unit, or LUN. The dedicated LUN that SWCC uses is called the Command Console LUN (CCL). CCL serves as the communication device for the HS-Series Agent and identifies itself to the host by a unique identification string. By default, a CCL device is enabled within the HSG60 controller on host port 1. The HSG60 uses SCSI-2 and SCSI-3. The use of the CCL is recommended on SGI Irix in SCSI-2 mode and SCSI-3. Consult the release notes for version restrictions.

The CCL does the following:

- Allows the RAID Array to be recognized by the host as soon as it is attached to the SCSI bus and configured into the operating system.
- Serves as a communications device for the HS-Series Agent. The CCL identifies itself to the host by a unique identification string.

In dual-redundant controller configurations, the commands described in the following section alter the setting of the CCL on both controllers. The CCL is enabled only on host port 1.

Determining the Address of the CCL

The CCL is enabled by default. Its address can be determined by entering the following CLI command:

```
HSG60 > SHOW THIS_CONTROLLER
```

Enabling/Disabling the CCL in SCSI-2 Mode

Enabling the CCL

To enable the CCL, enter the following CLI command:

```
HSG60 > SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

Disabling the CCL

To disable the CCL, enter the following CLI command:

```
HSG60> SET THIS_CONTROLLER NOCOMMAND_CONSOLE_LUN
```

To see the state of the CCL, use the `SHOW THIS CONTROLLER/ OTHER CONTROLLER` command. Because the CCL is not an actual LUN, the `SHOW UNITS` command will not display the CCL location.

Enabling/Disabling the CCL in SCSI-3 Mode

The CCL is enabled all the time. There is no option to enable/disable.

Connections


The term “connection” applies to every path between a Fibre Channel adapter in a host computer and an active host port on a controller.

NOTE: In ACS V8.6 Solution Software the maximum number of supported hosts is 96.

Naming Connections

Compaq highly recommends that you assign names to connections that have meaning in the context of your particular configuration. One system that works well is to name each connection after its host, its adapter, its controller, and its controller host port, as follows:

Examples:



A connection from the first adapter in host RED that goes to port 1 of controller A would be called RED1A1.

A connection from the third adapter in host GREEN that goes to port 2 of controller B would be called GREEN3B2.

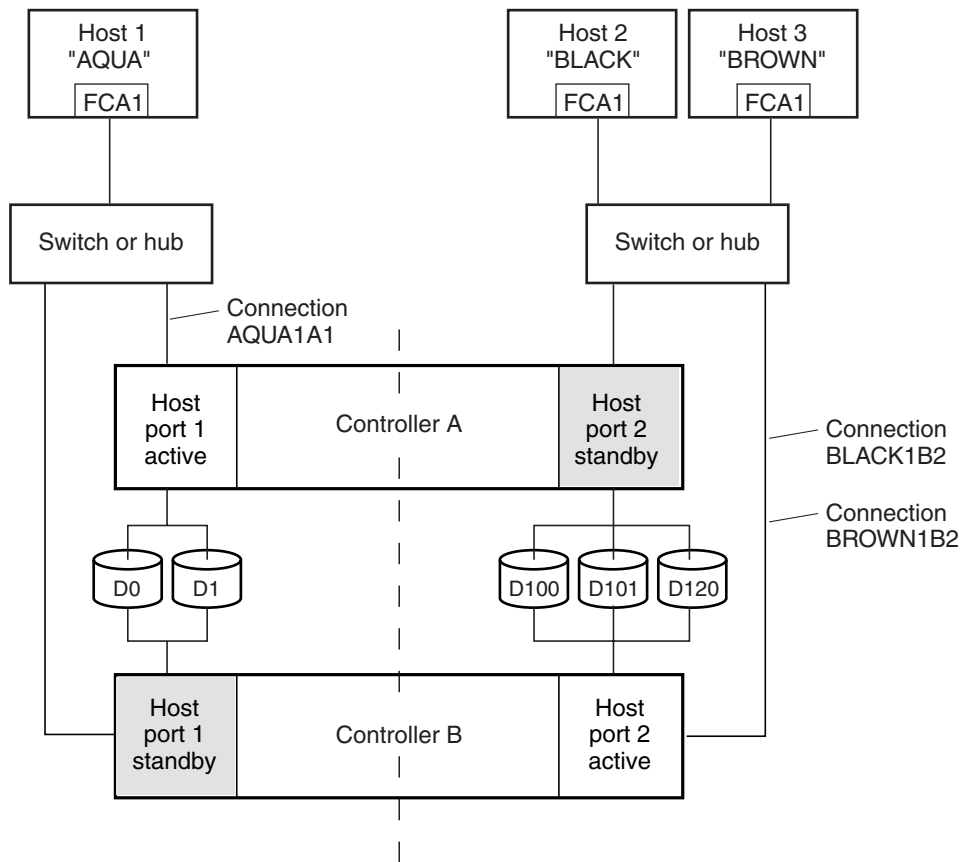
NOTE: Connection names can have a maximum of 9 characters.

Numbers of Connections

The number of connections resulting from cabling one adapter into a switch or hub depends on failover mode and how many links the configuration has:

- If a controller pair is in transparent failover mode and the port 1 link is separate from the port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another), each adapter will have one connection, as shown in Figure 1-6.

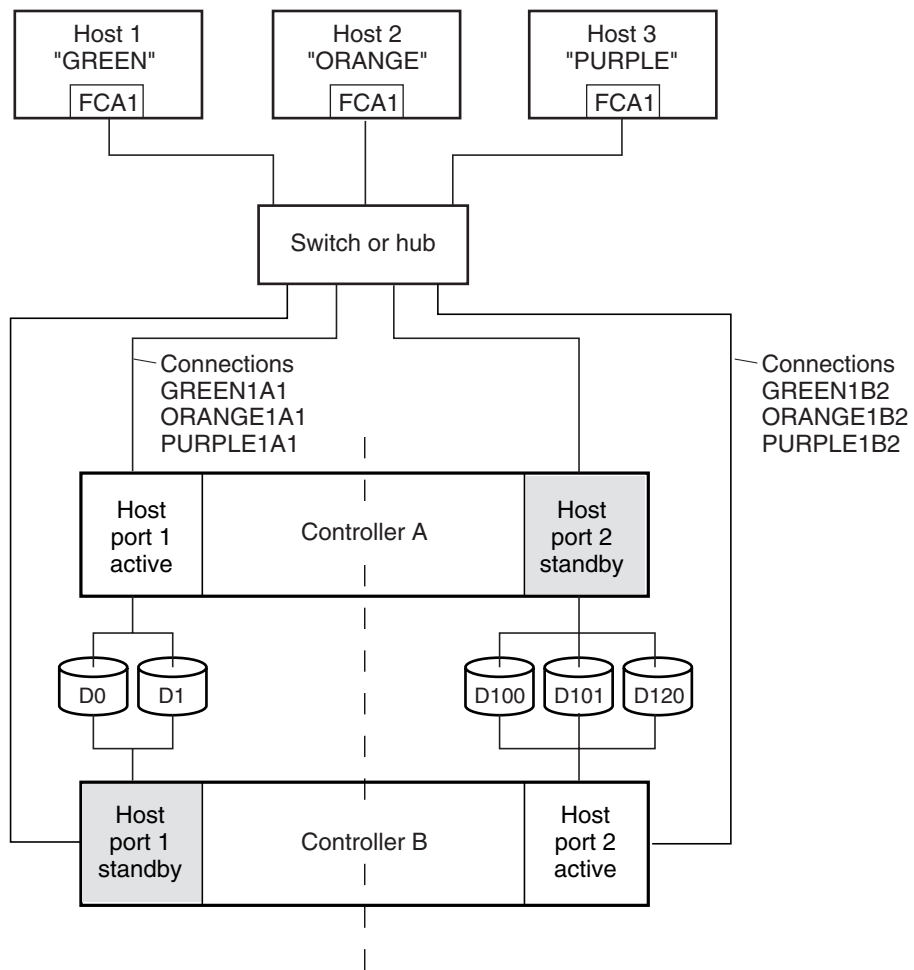
- If a controller pair is in transparent failover mode and port 1 and port 2 are on the same link (that is, all ports are on the same loop or fabric), each adapter will have two connections, as shown in Figure 1-7.



NOTE: FCA = Fibre Channel Adapter

CXO7081B

Figure 1-6. Connections in separate-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CX07079B

Figure 1-7. Connections in single-link, transparent failover mode configurations

Assigning Unit Numbers

The controller keeps track of the unit with the unit number. Unit number can be from 0 - 199 prefixed by a D, which stands for disk drive. A unit can be presented as different LUNs to different connections. The interaction of a unit and a connection is determined by several factors:

- Failover mode of the controller pair
- The *ENABLE_ACCESS_PATH* and *PREFERRED_PATH* switches in the *ADD UNIT* (or *SET unit*) commands
- The *UNIT_OFFSET* switch in the *ADD CONNECTIONS* (or *SET connections*) commands
- The controller port to which the connection is attached
- The *SCSI_VERSION* switch of the *SET this_controller/other_controller* command

The considerations for assigning unit numbers are discussed in the following sections.

Matching Units to Host Connections in Transparent Failover Mode

In transparent failover mode, the *ADD UNIT* command creates a unit for host connections to access and assigns it to either port 1 of both controllers or to port 2 of both controllers.

Unit numbers are assigned to ports as follows:

- 0 - 99 are assigned to host port 1 of both controllers.
- 100 - 199 are assigned to host port 2 of both controllers.

For example, unit D2 is on port 1, and unit D102 is on port 2.

The LUN number that a host connection assigns to a unit, is a function of the *UNIT_OFFSET* switch of the *ADD* (or *SET*) *CONNECTIONS* command. The relationship of offset, LUN number, and unit number follows:

LUN number = unit number – offset

Where...

- LUN number is relative to the host (what the host sees the unit as)
- unit number is relative to the controller (what the controller sees the unit as)

If no value is specified for offset, then connections on port 1 have a default offset of 0 and connections on port 2 have a default offset of 100.

For example, if all host connections use the default offset values, unit D2 will be presented to a port 1 host connection as LUN 2 (unit number of 2 minus offset of 0). Unit D102 will be presented to a port 2 host connection as LUN 2 (unit number of 102 minus offset of 100).

Figure 1-8 shows how units are presented as different LUNs, depending on the offset of the host. In this illustration, host connection 1 and host connection 2 would need to be on host port 1; host connection 3 would need to be on host port 2.

Controller units	Host connection 1 Offset: 0	Host connection 2 Offset: 20	Host connection 3 Offset: 100
D0	LUN 0		
D1	LUN 1		
D2	LUN 2		
D3	LUN 3		
D20	LUN 20	LUN 0	
D21	LUN 21	LUN 1	
D100			LUN 0
D101			LUN 1
D102			LUN 2
D130			LUN 30
D131			LUN 31

CXO6455B

Figure 1-8. LUN presentation to hosts, as determined by offset

Offsets other than the default values can be specified. For example, unit D17 would be visible to a host connection on port 1 that had an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible to a host connection with a unit offset of 18 or greater, because that offset is not within the unit's range (unit number of 17 minus offset of 18 is a negative number).

Similarly, unit D127 would be visible to a host connection on port 2 that had an offset of 120 as LUN 7 (unit number of 127 minus offset of 120). The unit would not be visible to a host connection with a unit offset of 128 or greater, because that offset is not within the unit's range (unit number of 127 minus offset of 128 is a negative number).

An additional factor to consider when assigning unit numbers and offsets is SCSI version. If the *SCSI_VERSION* switch of the SET *this_controller/other_controller* command is set to SCSI-3, the Command Console LUN (CCL) is presented as LUN 0 to every connection, superseding any unit assignments. The interaction between SCSI version and unit numbers is explained further in the next section.

In addition, the access path to the host connection must be enabled for the connection to access the unit. See “Restricting Host Access in Transparent Failover Mode,” page 1-17.

Assigning Unit Numbers Depending on SCSI_VERSION

The *SCSI_VERSION* switch of the SET THIS/OTHER CONTROLLER command determines how the Command Console LUN (CCL) is presented. There are two choices: SCSI-2 and SCSI-3. The choice for *SCSI_VERSION* affects how certain unit numbers and certain host connection offsets interact.

Assigning Unit Numbers and Host Connections in SCSI-2 Mode

Some operating systems expect or require a disk unit to be at LUN 0. In this case, it is necessary to specify SCSI-2 mode.

If *SCSI_VERSION* is set to SCSI-2 mode, the CCL floats, moving to the first available LUN location, depending on the configuration. Because the SWCC Agent depends on the exact location of the CCL, keep the CCL at LUN 0 by assigning units starting with D1 or D101. The following conventions are recommended when assigning host connection offsets and unit numbers in SCSI-2 mode:

- Offsets should be divisible by 10 (for consistency and simplicity).
- Unit numbers should be assigned at connection offsets (so that every host connection has a unit presented at LUN 0).

Table 1-1 summarizes the recommendations for unit assignments based on the *SCSI_VERSION* switch.

Table 1-1 Unit Assignments and <i>SCSI_VERSION</i>			
SCSI_VERSION	Offset	Unit Assignment	What the connection sees LUN 0 as
SCSI-2	Divisible by 10	At offsets	Unit whose number matches offset
SCSI-3	Divisible by 10	Not at offsets	CCL

Selective Storage Presentation

In a subsystem that is attached to more than one host or if the hosts have more than one adapter, it is possible to reserve certain units for the exclusive use of certain host connections.

For a controller pair, the method used to restrict host access depends on whether the controllers are in transparent or multiple-bus failover mode. For a single controller, the methods are the same as for a controller pair in transparent failover.

Restricting Host Access in Transparent Failover Mode

The following methods can be used to restrict host access to storage units in transparent failover mode:

- Using separate Fibre Channel links (either loop or fabric)
- Enabling the access path of selected host connections on a shared loop or fabric
- Setting offsets

NOTE: These techniques also work for a single controller.

Restricting Host Access by Separate Links

In transparent failover mode, controller A host port 1 and the controller B host port 1 share a common Fibre Channel link. Controller A host port 2 and the controller B host port 2 also share a common Fibre Channel link. If the host 1 link is separate from the host 2 link, the simplest way of limiting host access is to have one host or set of hosts on the port 1 link, and another host or set of hosts on the port 2 link. Each host can then see only units assigned to its respective controller port. This separation of host buses is shown in Figure 1-9. This method applies only if the host 1 link and host 2 link are separate links.

Figure 1-9. Limiting host access in transparent failover mode

Worldwide Names (Node IDs and Port IDs)

A worldwide name—also called a node ID—is a unique, 64-bit number assigned to a subsystem prior to shipping. The node ID belongs to the subsystem itself and never changes.

Each subsystem's node ID ends in zero, for example 5000-1FE1-FF0C-EE00. The controller port IDs are derived from the node ID.

In a subsystem with two controllers in transparent failover mode, the controller port IDs are incremented as follows:

- Controller A and controller B, port 1—worldwide name + 1, for example 5000-1FE1-FF0C-EE01
- Controller A and controller B, port 2—worldwide name + 2, for example 5000-1FE1-FF0C-EE02

Use the CLI command, `SHOW this_controller/other_controller` to display the subsystem's worldwide name.

Restoring Worldwide Names (Node IDs)

If a situation occurs that requires you to restore the worldwide name, you can restore it using the worldwide name and checksum printed on the sticker on the frame into which the controller is inserted.

Figure 1-10 shows the placement of the worldwide name label for the Model 2100 enclosure.




Figure 1-10. Placement of the worldwide name label on the Model 2100 enclosure

Unit Worldwide Names (LUN IDs)

In addition, each unit has its own worldwide name, or LUN ID. This is a unique, 128-bit value that the controller assigns at the time of unit initialization. It cannot be altered by the user but does change when the unit is reinitialized. Use the SHOW command to list the LUN ID.

HSG60 Setup

Before the HSG60 RAID Array can be utilized by the Host System, the array must be properly configured. Use the HSG60 Command Line Interpreter (CLI) to complete the following tasks.

NOTE: The exact command syntax can be found in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

Ensure that the HSG60 controllers have been set correctly for arbitrated loop or fabric. Compare the settings displayed by the CLI with those in Table 1-2 for arbitrated loop and Table 1-3 for fabric. The current settings can be obtained with the command:

```
HSG60>SHOW THIS_CONTROLLER
```

Arbitrated Loop Setup

Table 1-2 Controller Parameters and Recommended Settings for Arbitrated Loop

Controller Parameter	Recommended Settings
PORT_1_AL_PA	71
PORT_2_AL_PA	72
PORT_1_Topology	Loop_Hard
PORT_2_Topology	Loop_Hard
This_Controller SCSI_Version	SCSI-2
Other_Controller SCSI_Version	SCSI-2

If the controller parameter settings displayed do not match those in the Recommended Settings column in Table 1-2, the following CLI commands can be used to alter their values.

NOTE: HSG60> is included in the CLI commands to show the prompt that will appear and is not to be typed as part of the command line.

Setting the ALPA:

```
HSG60 >SET THIS_CONTROLLER PORT_1_AL_PA=71
```

```
HSG60 >SET THIS_CONTROLLER PORT_2_AL_PA=72
```

Setting the Port Topology:

```
HSG60>SET OTHER_CONTROLLER PORT_1_TOPOLOGY=LOOP_HARD
```

```
HSG60>SET OTHER_CONTROLLER PORT_2_TOPOLOGY=LOOP_HARD
```

Setting the SCSI Version:

```
HSG60>SET THIS_CONTROLLER SCSI_VERSION=SCSI-2
```

```
HSG60>SET THIS_CONTROLLER SCSI_VERSION=SCSI-2
```

If the StorageWorks Command Console (SWCC) will be used, the HSG60 Command Console LUN must be enabled.

```
HSG60>SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

Once all of the changes have been made, restart both controllers as follows:

```
HSG60>RESTART OTHER_CONTROLLER
```

```
HSG60>RESTART THIS_CONTROLLER
```

Once both controllers have been restarted, the connection type must be set. First examine the connections by typing:

```
HSG60>SHOW CONNECTIONS
```

In the following command substitute the **[CONNECTIONNAME]** with the names shown in the output of the SHOW CONNECTIONS command. Set the connections to SGI by typing:

```
HSG60>SET [CONNECTIONNAME] OPERATING __SYSTEM =SGI
```

Repeat the command for each connection shown.

Fabric Setup

Ensure that the HSG60 controllers have been set correctly for fabric. Compare the settings displayed by the CLI with those in Table 1-3. The current settings can be obtained with the command:

```
HSG60>SHOW THIS_CONTROLLER
```

Table 1-3 Controller Parameters and Recommended Settings for Fabric

Controller Parameter	Recommended Setting
PORT_1_Topology	Fabric
PORT_2_Topology	Fabric
This_Controller SCSI_Version	SCSI-2
Other_Controller SCSI_Version	SCSI-2

If the HSG60 controller parameter settings displayed do not match those in the Recommended Settings column in Table 1-3, the following CLI commands can be used to alter their values.

NOTE: HSG60> is included in the CLI commands to show the prompt that will appear and is not to be typed as part of the command line.

Setting the Port Topology:

```
HSG60 >SET OTHER_CONTROLLER PORT_1_TOPOLOGY=FABRIC
HSG60 >SET OTHER_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

Setting the SCSI Version:

```
HSG60 >SET THIS_CONTROLLER SCSI_VERSION=SCSI-2
HSG60 >SET THIS_CONTROLLER SCSI_VERSION=SCSI-2
```

If the StorageWorks Command Console (SWCC) will be used, the HSG60 Command Console LUN must be enabled.

```
HSG60>SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

Once all of the changes have been made, restart both controllers as follows:

```
HSG60>RESTART OTHER_CONTROLLER
HSG60>RESTART THIS_CONTROLLER
```

Once both controllers have been restarted, the connection type must be set. First examine the connections by typing:

```
HSG60>SHOW CONNECTIONS
```

In the following command substitute the [CONNECTIONNAME] with the names shown in the output of the SHOW CONNECTIONS command.

Set the connections to SGI by typing:

```
HSG60>SET [CONNECTIONNAME] OPERATING __SYSTEM =SGI
```

Repeat the command for each connection shown.

Host System Setup

Once the RAID Array has been configured, the Host System must be set up.

1. Shut down the host operating system, using your site-specific procedures. Once the operating system has halted, power off the SGI system enclosure.
2. If necessary, install the SGI Fibre Channel host adapter, and the optical to electrical Media Interface Adapter (MIA), using specific SGI FC-AL adapters only. The MIA should be secured to the host adapter with its two attached fastening screws. The MIA is available from SGI using the following part number X-F-OE-KIT
3. Before starting the Host System, connect the optical cables from the RAID Controller to the hub or switch. Then connect the Fibre Channel Host adapter in the SGI enclosure to the hub or switch with an optical cable.
4. Once the cables are securely connected, power on the SGI enclosure and start the operating system.
5. Once the bootstrap process has completed, log in to the system as superuser (*root*).

Chapter 2

Planning Storage

The information in this chapter will help you plan the storage configuration of your subsystem. Use the guidelines found in this section to plan the various types of storage containers needed.

The following information is included in this chapter:

- “Where to Start,” page 2–2
- “Configuration Rules,” page 2–3
- “Determining Storage Requirements,” page 2–6
- “Choosing a Container Type,” page 2–6
- “Creating a Stageset Profile,” page 2–8
- “Stageset Planning Considerations,” page 2–10
- “Mirrorset Planning Considerations,” page 2–13
- “Stageset Expansion Considerations,” page 2–19
- “Partition Planning Considerations,” page 2–19
- “Changing Characteristics through Switches,” page 2–21
- “Stageset and Partition Switches,” page 2–22
- “Initialization Switches,” page 2–23
- “Unit Switches,” page 2–27
- “Storage Maps,” page 2–27

Where to Start

The following procedure outlines the steps to take to plan your storage configuration. Containers are individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on) and/or partitioned drives. See Appendix A to locate the blank templates for keeping track of the containers being configured.

1. Review configuration rules. See “Configuration Rules,” page 2-3.
2. Familiarize yourself with the current physical layout of the devices and their addressing scheme. See “Device PTL Addressing Convention,” page 2-3.
3. Determine your storage requirements. Use the questions in “Determining Storage Requirements,” page 2-6 to help you.
4. Choose the type of storage container(s) you need to use in your subsystem. See “Choosing a Container Type,” page 2-6 for a comparison and description of each type of storageset.
5. Create a storageset profile (described in “Creating a Storageset Profile,” page 2-8). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:
 - “Storageset Planning Considerations,” page 2-10
 - “Mirrorset Planning Considerations,” page 2-13
 - “RAIDset Planning Considerations,” page 2-15
 - “Partition Planning Considerations,” page 2-19
 - “Striped Mirrorset Planning Considerations,” page 2-17
6. Decide which switches you will need for your subsystem. General information on switches is detailed in “Storageset and Partition Switches,” page 2-22.
 - Determine what unit switches you want for your units (“Unit Switches,” page 2-27)
 - Determine what initialization switches you want for your planned storage containers (“Initialization Switches,” page 2-23)
7. Create a storage map (“Storage Maps,” page 2-27).
8. Configure the storage you have now planned using one of the following methods:
 - Use SWCC. See the SWCC documentation for details.
 - Use the Command Line Interpreter (CLI) commands. This method allows you flexibility in defining and naming your storage containers. See the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

Configuration Rules

Review these configuration rules to ensure your planned configuration meets the following requirements and conditions:

- Maximum of 128 LUNs: if Command Console LUN is enabled, the result is 127 visible LUNs and one CCL
- Maximum 1.024 TB LUN capacity
- Maximum 72 physical drives with 2 optional expansion enclosures per storage system
- Maximum 24 physical devices
- Maximum 20 RAID-5 storagesets
- Maximum 8 partitions per storageset or individual disk
- Maximum 6 physical devices per RAID 1 (mirrorset)
- Maximum 14 physical devices per RAID-5 storageset
- Maximum 24 physical devices per RAID 0 (stripeset)
- Maximum 24 physical devices per striped mirrorset

Device PTL Addressing Convention

The controller has two SCSI I/O device ports, each of which connects to a SCSI bus. In dual-controller subsystems, these device buses are shared between the two controllers. (SWCC calls the device ports “channels.”) The standard Model 2100-Series UltraSCSI Controller enclosure provides a maximum of 12 SCSI target IDs for each device port.

The controller identifies devices based on a Port-Target-LUN (PTL) numbering scheme, shown in Figure 2-1. The physical location of a device in its enclosure determines its PTL.

- P—Designates the controller's SCSI device port number (1 or 2).
- T—Designates the target identification (ID) number of the device. Valid target ID numbers for a single-controller configuration and dual-redundant controller configuration are 0 - 5 and 8 - 13.
- L—Designates the logical unit (LUN) of the device. For disk devices the LUN is always 0.




Figure 2-1. PTL naming convention

The controller operates with a Model 2100 UltraSCSI controller enclosure combined with either Model 4214R (maximum of two) or Model 4254 (maximum of one) disk enclosures. Physical drives are assigned SCSI target ID numbers 0 through 5 and 8 through 13 (SCSI target IDs 6 and 7 are reserved for the controllers, and SCSI target IDs 14 and 15 are not supported). Figure 2-2 shows the addresses for each device using two Model 4214R disk enclosure shelves in a single-bus configuration. Figure 2-3 shows the addresses for each drive using a single Model 4254 disk enclosure shelf in a dual-bus configuration.

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

Model 4214R Disk Enclosure Shelf 2 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

Figure 2-2. PTL addressing in a single-bus configuration, using two shelves

Model 4254 Disk Enclosure Shelf 1 (dual-bus)														
Bus A							Bus B							
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800

Figure 2-3. PTL addressing in a dual-bus configuration

When the controller receives an I/O request, it identifies the storage set unit number for the request, then correlates the unit number to the storage set name. From the storage set name, the controller locates the appropriate device(s) for the I/O request (For example, the RAIDset “RAID1” might contain DISK10000, DISK20000 and DISK10100). The controller generates the read or write request to the appropriate device(s) using the PTL addressing convention. Figure 2-4 illustrates the concept of mapping unit numbers to PTL addresses.

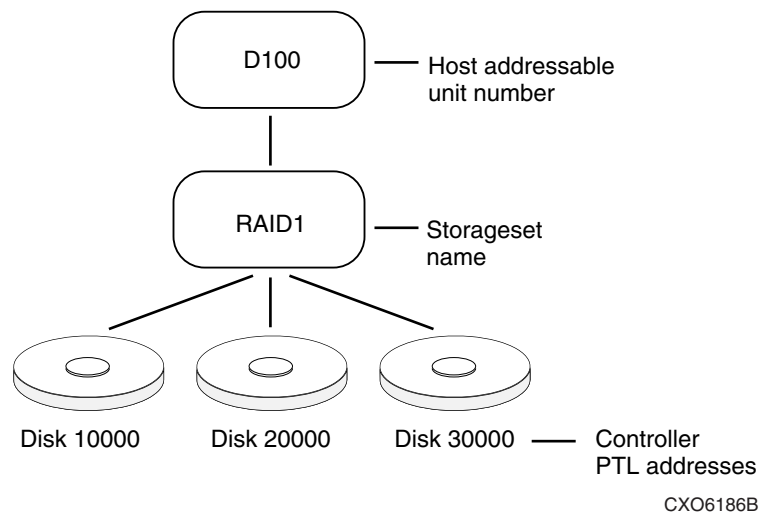


Figure 2-4. Mapping a unit to physical disk drives

Determining Storage Requirements

It is important to determine your storage requirements. Here are a few of the questions you should ask yourself regarding the subsystem usage:


- What applications or user groups will access the subsystem? How much capacity do they need?
- What are the I/O requirements? If an application is data-transfer intensive, what is the required transfer rate? If it is I/O-request intensive, what is the required response time? What is the read/write ratio for a typical request?
- Are most I/O requests directed to a small percentage of the disk drives? Do you want to keep it that way or balance the I/O load?
- Do you store mission-critical data? Is availability the highest priority or would standard backup procedures suffice?

Choosing a Container Type

Different applications may have different storage requirements. You will probably want to configure more than one kind of container within the subsystem.

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in Figure 2-5. The independent disks and the selected storageset may also be partitioned.

The storagesets implement RAID (Redundant Array of Independent Disks) technology. Consequently, they all share one important feature: each storageset, whether it contains two disk drives or ten, looks like one large, virtual disk drive to the host.



CXO6677A

Figure 2-5. Container types

Table 2-1 compares the different kinds of containers to help you determine which ones satisfy your requirements.

Container Name	Relative Availability	Request Rate (Read/Write) I/O per second	Transfer Rate (Read/Write) MB per second	Applications
Independent disk drives (JBOD)	Equal to number of JBOD disk drives	Comparable to single disk drive	Comparable to single disk drive	—
Stripeset (RAID 0)	Proportionate to number of disk drives; worse than single disk drive	Excellent if used with large chunk size	Excellent if used with small chunk size	High performance for non-critical data
Mirrorset (RAID 1)	Excellent	Good/Fair	Good/Fair	System drives; critical files

Table 2-1 A Comparison of Container Types (Continued)

Container Name	Relative Availability	Request Rate (Read/Write) I/O per second	Transfer Rate (Read/Write) MB per second	Applications
RAIDset (RAID 3/5)	Excellent	Excellent/good	Read: excellent (if used with small chunk sizes) Write: good (if used with small chunk sizes)	High request rates, read-intensive, data lookup
Striped Mirrorset (RAID 0+1)	Excellent	Excellent if used with large chunk size	Excellent if used with small chunk size	Any critical response-time application

For a comprehensive discussion of RAID, refer to *The RAIDBOOK—A Source Book for Disk Array Technology*.

Creating a Stageset Profile

Creating a profile for your stagesets, partitions, and devices can simplify the configuration process. Filling out a stageset profile helps you choose the stagesets that best suit your needs and to make informed decisions about the switches you can enable for each stageset or storage device that you configure in your subsystem.

The example stageset profile shown in Figure 2-6. Appendix A contains blank profiles that you can copy and use to record the details for your stagesets. Use the information in this chapter to help you make decisions when creating stageset profiles.

Type of StorageSet:

<input type="checkbox"/> Mirrorset	<input checked="" type="checkbox"/> RAIDset	<input type="checkbox"/> Stripeset	<input type="checkbox"/> Striped	<input type="checkbox"/> JBOD
Mirrorset				

StorageSet Name *R1*

Disk Drives *D10300, D20300, D10400, D20400*

Unit Number *D101*

Partitions:

Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #
%	%	%	%	%	%	%	%

RAIDset Switches:

Reconstruction Policy	Reduced Membership	Replacement Policy
<input checked="" type="checkbox"/> Normal (default)	<input checked="" type="checkbox"/> No (default)	<input checked="" type="checkbox"/> Best performance (default)
<input type="checkbox"/> Fast	<input type="checkbox"/> Yes, missing:	<input type="checkbox"/> Best fit
		<input type="checkbox"/> None

Mirrorset Switches:

Replacement Policy	Copy Policy	Read Source
<input type="checkbox"/> Best performance (default)	<input type="checkbox"/> Normal (default)	<input type="checkbox"/> Least busy (default)
<input type="checkbox"/> Best fit	<input type="checkbox"/> Fast	<input type="checkbox"/> Round robin
<input type="checkbox"/> None		<input type="checkbox"/> Disk drive:

Initialize Switches:

Chunk size	Save Configuration	Metadata
<input checked="" type="checkbox"/> Automatic (default)	<input type="checkbox"/> No (default)	<input checked="" type="checkbox"/> Destroy (default)
<input type="checkbox"/> 64 blocks	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> Retain
<input type="checkbox"/> 128 blocks		
<input type="checkbox"/> 256 blocks		
<input type="checkbox"/> Other:		

Unit Switches:

Caching	Access by following hosts enabled
Read caching <input checked="" type="checkbox"/>	<u>ALL</u>
Read-ahead caching <input type="checkbox"/>	_____
Write-back caching <input checked="" type="checkbox"/>	_____
Write-through caching <input type="checkbox"/>	_____

Figure 2-6. An example storageSet profile

Storageset Planning Considerations

This section contains the guidelines for choosing the storageset type that would best work in your subsystem:

- “Stripeset Planning Considerations” on page 2-10
- “Mirrorset Planning Considerations” on page 2-13
- “RAIDset Planning Considerations” on page 2-15
- “Striped Mirrorset Planning Considerations” on page 2-17
- “Storageset Expansion Considerations” on page 2-19
- “Partition Planning Considerations” on page 2-19

NOTE: Storageset expansion is synonymous with Concat Set.

Stripeset Planning Considerations

Stripesets (RAID 0) enhance I/O performance by spreading the data across multiple disk drives. Each I/O request is broken into small segments called “chunks.” These chunks are then simultaneously “striped” across the disk drives in the storageset, thereby enabling several disk drives to participate in one I/O request.

For example, in a three-member stripeset that contains disk drives Disk10000, Disk20000, and Disk10100, the first chunk of an I/O request is written to Disk10000, the second to Disk 20000, the third to Disk 10100, the fourth to Disk10000, and so forth until all of the data has been written to the drives (Figure 2-7).




Figure 2-7. A 3-member RAID 0 stripeset (example 1)

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rate. You can set the chunk size or use the default setting (see “Chunk Size,” page 2-23, for information about setting the chunk size). Figure 2-8 shows another example of a three-member RAID 0 Stripeset.

A major benefit of striping is that it balances the I/O load across all of the disk drives in the storageset. This can increase the subsystem performance by eliminating the hot spots (high localities of reference), that occur when frequently accessed data becomes concentrated on a single disk drive.




Figure 2-8. A 3-member RAID 0 stripeset (example 2)

Keep the following points in mind when planning stripesets:

- Reporting methods and size limitations prevent certain operating systems from working with large stripesets.
- A storageset should only contain disk drives of the same capacity. The controller limits the effective capacity of each member to the capacity of the smallest member in the storageset (base member size) when the storageset is initialized. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you will waste 5 GB of capacity on each 9 GB member.

If you need high performance and high availability, consider using a RAIDset, striped-mirrorset, or a host-based shadow of a stripeset.

- Striping does not protect against data loss. In fact, because the failure of one member is equivalent to the failure of the entire stripeset, the likelihood of losing data is higher for a stripeset than for a single disk drive.

For example, if the mean time between failures (MTBF) for a single disk is 1 hour, then the MTBF for a stripeset that comprises N such disks is 1/N hours. As another example, if the MTBF of a single disk is 150,000 hours (about 17 years), a stripeset comprising four of these disks would only have an MTBF of slightly more than 4 years.


For this reason, you should avoid using a stripeset to store critical data. Stripesets are more suitable for storing data that can be reproduced easily or whose loss does not prevent the system from supporting its critical mission.

- Evenly distribute the members across the device ports to balance load and provide multiple paths.
- Stripesets may contain between two and 24 members.
- Stripesets are well-suited for the following applications:
 - Storing program image libraries or run-time libraries for rapid loading.
 - Storing large tables or other structures of read-only data for rapid application access.
 - Collecting data from external sources at very high data transfer rates.
- Stripesets are not well-suited for the following applications:
 - A storage solution for data that cannot be easily reproduced or for data that must be available for system operation.
 - Applications that make requests for small amounts of sequentially located data.
 - Applications that make synchronous random requests for small amounts of data.

Spread the member drives as evenly as possible across the two I/O device ports.

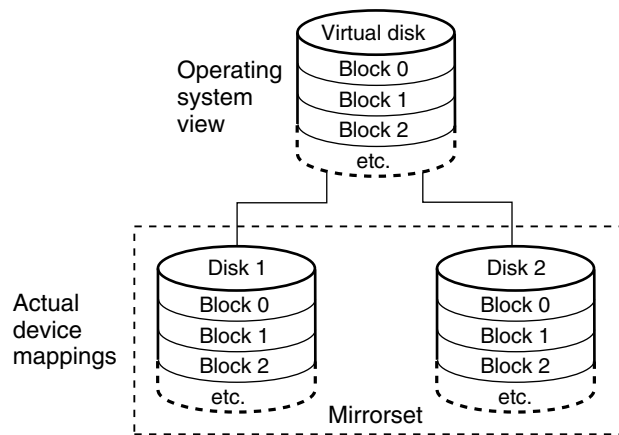
Mirrorset Planning Considerations

Mirrorsets (RAID 1) use redundancy to ensure availability, as illustrated in Figure 2-9. For each primary disk drive, there is at least one mirror disk drive. Thus, if a primary disk drive fails, its mirror drive immediately provides an exact copy of the data. Figure 2-10 shows a second example of a Mirrorset.



CXO7288A

Figure 2-9. Mirrorsets maintain two copies of the same data



CXO4594B

Figure 2-10. Mirrorset example 2

Keep these points in mind as when planning mirrorsets:

- Data availability with a mirrorset is excellent but comes with a higher cost—you need twice as many disk drives to satisfy a given capacity requirement. If availability is your top priority, consider using dual-redundant controllers and redundant power supplies.

- You can configure up to 20 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset may contain up to 6 members.
- Both write-back cache modules must be the same size.
- A mirrorset should only contain disk drives of the same capacity.
- Mirrorsets are well-suited for the following:
 - Any data for which reliability requirements are extremely high
 - Data to which high-performance access is required
 - Applications for which cost is a secondary issue
- Mirrorsets are not well-suited for the following applications:
 - Write-intensive applications (a performance hit of 10% will occur)
 - Applications for which cost is a primary issue

RAIDset Planning Considerations

RAIDsets (RAID 3/5) are enhanced stripesets—they use striping to increase I/O performance and distributed-parity data to ensure data availability. Figure 2-11 shows an example of a RAIDset that uses five members.




Figure 2-11. A five member RAIDset using parity

RAIDsets are similar to stripesets in that the I/O requests are broken into smaller “chunks” and striped across the disk drives. RAIDsets also create chunks of parity data and stripe them across all the members of the RAIDset. This parity data is derived mathematically from the I/O data and enables the controller to reconstruct the I/O data if a single disk drive fails. Thus, it becomes possible to lose a disk drive without losing access to the data it contained. Data could be lost, however, if a second disk drive fails before the controller replaces the first failed disk drive and reconstructs the data.

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rates. You can set the chunk size or use the default setting. See “Chunk Size,” page 2-23, for information about setting the chunk size.

Keep these points in mind when planning RAIDsets:

- Reporting methods and size limitations prevent certain operating systems from working with large RAIDsets.
- Both cache modules must be the same size.

- A RAIDset must include at least 3 disk drives, but no more than 14.
- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you will waste 5 GB of capacity on each 9 GB member.
- RAIDsets are particularly well-suited for the following:
 - Small to medium I/O requests
 - Applications requiring high availability
 - High read request rates
 - Inquiry-type transaction processing
- RAIDsets are not particularly well-suited for the following:
 - Write-intensive applications
 - Database applications in which fields are continually updated
 - Transaction processing

Striped Mirrorset Planning Considerations

Striped mirrorsets (RAID 0+1) are a configuration of stripesets whose members are also mirrorsets (Figure 2-12). Consequently, this kind of storageset combines the performance of striping with the reliability of mirroring. The result is a storageset with very high I/O performance and high data availability. Figure 2-13 shows a second example of a striped mirrorset using six members.





Figure 2-12. Striped mirrorset (example 1)

The failure of a single disk drive has no effect on the ability of the storageset to deliver data to the host. Under normal circumstances, it also has very little effect on performance. Because striped mirrorsets do not require any more disk drives than mirrorsets, this storageset is an excellent choice for data that warrants mirroring.



CXO6462A

Figure 2-13. Striped mirrorset (example 2)

Plan the mirrorset members, then plan the stripeset that will contain them. Review the recommendations in “Storageset Planning Considerations,” page 2-10, and “Mirrorset Planning Considerations,” page 2-13.

Storageset Expansion Considerations

Storageset Expansion allows for the joining of two of the same kind of storage containers by concatenating RAIDsets, Stripesets, or individual disks, thereby forming a larger virtual disk which is presented as a single unit. The *Compaq StorageWorks HSG60 Array Controller (ACS) Version 8.6 CLI Reference Guide* describes the CLI command: ADD CONCATSETS which is used to perform concatenation.



CAUTION: The ADD CONCATSETS command should only be executed with host operating systems that support dynamic volume expansion. If the operating system cannot handle one of its disks increasing in size, use of this command could make data inaccessible.

Partition Planning Considerations

Use partitions to divide a container (storageset or individual disk drive) into smaller pieces, each of which can be presented to the host as its own storage unit. Figure 2-14 shows the conceptual effects of partitioning a single-disk container.




Figure 2-14. One example of a partitioned single-disk unit

You can create up to eight partitions per storageset (disk drive, RAIDset, mirrorset, stripeset, or striped mirrorset). Each partition has its own unit number so that the host can send I/O requests to the partition just as it would to any unpartitioned storageset or device. Partitions are separately addressable storage units, therefore, you can partition a single storageset to service more than one user group or application.

Defining a Partition

Partitions are expressed as a percentage of the storageset or single disk unit that contains them:

- Mirrorsets and single disk units—the controller allocates the largest whole number of blocks that are equal to or less than the percentage you specify.
- RAIDsets and stripesets—the controller allocates the largest whole number of stripes that are less than or equal to the percentage you specify.
 - Stripesets—the stripe size = chunk size × number of members.
 - RAIDsets—the stripe size = chunk size × (number of members minus 1)

An unpartitioned storage unit has more capacity than a partition that uses the whole unit because each partition requires a small amount of disk space for metadata.

Guidelines for Partitioning Storagesets and Disk Drives

Keep these points in mind when planning partitions for storagesets and disk drives:

- A storageset or disk drive can contain up to eight partitions.
- In transparent failover mode, all partitions of a particular container must be on the same host port. Partitions cannot be split across host ports.
- Partitions cannot be combined into storagesets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAIDset.
- Just as with storagesets, you do not have to assign unit numbers to partitions until you are ready to use them.
- The CLONE utility cannot be used with partitioned mirrorsets or partitioned stripesets (See “Cloning Data for Backup,” page 7-2 for details about cloning.)

Changing Characteristics through Switches

CLI command switches allow the user another level of command options. There are three types of switches that modify the storageset and unit characteristics:

- Storageset switches
- Initialization switches
- Unit Switches

The following sections describe how to enable/modify switches. They contain a description of the major CLI command switches.

Enabling Switches

If you use SWCC to configure the device or storageset, you can set switches from the SWCC screens during the configuration process, and SWCC automatically applies them to the storageset or device. See the SWCC help online for information about using SWCC.

If you use CLI commands to configure the storageset or device manually, the configuration procedure found in Chapter 5 of this guide indicates when and how to enable each switch. The *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* contains the details of the CLI command and their switches.

Changing Switches

You can change the RAIDset, mirrorset, device, and unit switches at any time. You cannot change the initialize switches without destroying the data on the storageset or device. These switches are integral to the formatting and can only be changed by re-initializing the storageset.



CAUTION: Initializing a storageset is similar to formatting a disk drive; all of the data is destroyed during this procedure.

Storageset and Partition Switches

The characteristics of a particular storageset can be set by specifying switches when the storageset is added to the controllers' configuration. Once a storageset has been added, the switches can be changed by using a SET command. Switches can be set for partitions and the following types of storagesets:

- RAIDset
- Mirrorset

Stripesets have no specific switches associated with their ADD and SET commands.

RAIDset Switches

You can enable the following kinds of switches to control how a RAIDset behaves to ensure data availability:

- Replacement policy
- Reconstruction policy
- Remove/replace policy

For details on the use of these switches refer to SET RAIDSET and SET *RAIDset-name* commands in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

Mirrorset Switches

You can enable the following switches to control how a mirrorset behaves to ensure data availability:

- Replacement policy
- Copy speed
- Read source
- Membership

For details on the use of these switches refer to ADD MIRRORSET and SET *mirrorset-name* commands in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

Partition Switches

You can specify the following switches when creating a partition:

- Size
- Geometry

For details on the use of these switches refer to CREATE_PARTITION command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

Initialization Switches

Initialization switches set characteristics for established storagesets before they are made into units. The following kinds of switches affect the format of a disk drive or storageset:

- Chunk Size (for stripesets and RAIDsets only)
- Save Configuration
- Destroy/Nodeestroy
- Geometry

Each of these is described in the following sections.

NOTE: After you initialize the storageset or disk drive, you cannot change these switches without reinitializing the storageset or disk drive.

Chunk Size

Specify the chunk size of the data to be stored to control the stripesize used in RAIDsets and stripesets:

- *CHUNKSIZE=DEFAULT* lets the controller set the chunk size based on the number of disk drives (d) in a stripeset or RAIDset. If $d \leq 9$ then chunk size = 256. If $d > 9$ then chunk size = 128.
- *CHUNKSIZE=n* lets you specify a chunk size in blocks. The relationship between chunk size and request size determines whether striping increases the request rate or the data-transfer rate.

Increasing the Request Rate

A large chunk size (relative to the average request size) increases the request rate by enabling multiple disk drives to respond to multiple requests. If one disk drive contains all of the data for one request, then the other disk drives in the storageset are available to handle other requests. Thus, in principle, separate I/O requests can be handled in parallel, thereby increasing the request rate. This concept is shown in Figure 2-15.




Figure 2-15. Chunk size larger than the request size

Large chunk sizes also tend to increase the performance of random reads and writes. Compaq recommends that you use a chunk size of 10 to 20 times the average request size, rounded to the closest prime number.

To calculate the chunk size that should be used for the subsystem, you first must analyze the types of requests that are being made to the subsystem:

- Many parallel I/Os that use a small area of disk should use a chunk size of 10 times the average transfer request rate.
- Random I/Os that are scattered over all the areas of the disks should use a chunk size of 20 times the average transfer request rate.
- If you do not know, then you should use a chunk size of 15 times the average transfer request rate.


- If you have mostly sequential reads or writes (like those needed to work with large graphic files), then make the chunk size a small number (that is 17 sectors).

Table 2-2 shows a few examples of chunk size selection.

Table 2-2 Example Chunk Sizes			
Transfer Size (KB)	Small Area of I/O Transfers	Unknown	Random Areas of I/O Transfers
2	41	59	79
4	79	113	163
8	157	239	317

Increasing the Data Transfer Rate

A small chunk size relative to the average request size increases the data transfer rate by enabling multiple disk drives to participate in one I/O request. This concept is shown in Figure 2-16.



CXO-5172A-MC

Figure 2-16. Chunk size smaller than the request size

Increasing Sequential Write Performance

Sequential write and read requests on stripesets (or striped mirrorsets), should use a small chunk size relative to the I/O size to increase the performance. A chunk size of 17 generally works well.

Save Configuration

This switch is for a single-controller configuration only. This switch reserves an area on each of the disks that constitute the container being initialized. The controller can write subsystem configuration data on this area. If the controller is replaced, the new controller can read the subsystem configuration from the reserved area of disks.

If you specify `SAVE_CONFIGURATION` for a multi-device storageset, such as a stripeset, the complete subsystem configuration is periodically written on each disk in the storageset.

The `SHOW DEVICES FULL` command shows which disks are used to backup configuration information.

Destroy/Nodestroy

Specify whether to destroy or retain the user data and metadata when a disk is initialized that has been previously used in a mirrorset or as a single-disk unit.

NOTE: The `DESTROY` and `NODESTROY` switches are only valid for mirrorsets and striped mirrorsets.

- `DESTROY` (default) overwrites the user data and forced-error metadata on a disk drive when it is initialized.
- `NODESTROY` preserves the user data and forced-error metadata when a disk drive is initialized. Use `NODESTROY` to create a single-disk unit from any disk drive that has been used as a member of a mirrorset. See the `REDUCED` command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* for information on removing disk drives from a mirrorset.

`NODESTROY` is ignored for members of a RAIDset.

Geometry

The geometry parameters of a storage set can be specified. The geometry switches are:

- *CAPACITY*—the number of logical blocks. The range is from 1 to the maximum container size.
- *CYLINDERS*—the number of cylinders used. The range is from 1 to 16777215.
- *HEADS*—the number of disk heads used. The range is from 1 to 255.
- *SECTORS_PER_TRACK*—the number of sectors per track used. The range is from 1 to 255.

Unit Switches

There are several switches that control the characteristics of units. The unit switches are described under the SET *unit-number* command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

One unit switch, *ENABLE/DISABLE_ACCESS_PATH*, determines which host connections can access the unit, and is part of the larger topic of matching units to specific hosts. This complex topic is covered in Chapter 1 under the following headings:

- “Assigning Unit Numbers,” page 1-14
- “Selective Storage Presentation,” page 1-17

Storage Maps

Configuring a subsystem will be easier if you know how the storage sets, partitions, and JBODs correspond to the disk drives in your subsystem. You can more easily see this relationship by creating a hardcopy representation (a storage map).

Creating a Storage Map

To make a storage map, fill out a blank storage map as you add storage sets, partitions, and JBOD disks to your configuration and assign them unit numbers. Appendix A contains blank templates you may use in the creation of your subsystem storage map. Label each disk drive in the map with the higher levels it is associated with, up to the unit level.

The following example (see Figure 2-17) is of a single-shelf disk enclosure using a single-bus I/O module (Model 4214R):

- Unit D100 is a 3-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk10100, and Disk10200.
- Unit D101 is a 2-member striped mirror set named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirror set consisting of Disk10300 and Disk10400.
 - M2 is a 2-member mirror set consisting of Disk10500 and Disk10800.
- Unit D102 is a 2-member mirror set named M3. M3 consists of Disk10900 and Disk11000.
- Unit D103 is a 3-member stripeset named S2. S2 consists of Disk11100, Disk11200, and Disk11300.
- Unit D104 is a JBOD consisting of Disk11400.
- Disk11500 is a spareset member.

The I/O module assigns a SCSI address to each drive bay in the disk enclosure using a SCSI address map.

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D100 R1 DISK10000	D100 R1 DISK10100	D100 R1 DISK10200	D101 S1 M1 DISK10300	D101 S1 M1 DISK10400	D101 S1 M2 DISK10500	D101 S1 M2 DISK10800	D102 M3 DISK10900	D102 M3 DISK11000	D103 S2 DISK11100	D103 S2 DISK11200	D103 S2 DISK11300	D104 DISK11400	spareset member DISK11500

Figure 2-17. Example storage map-single bus

The following example (see Figure 2-18) is of a single-shelf disk enclosure using a dual-bus I/O module (Model 4254):

- Unit D100 is a 3-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, and Disk10100.

- Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirrorset consisting of Disk10200 and Disk20200.
 - M2 is a 2-member mirrorset consisting of Disk20300 and Disk10300.
- Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10400 and Disk20400.
- Unit D103 is a 3-member stripeset named S2. S2 consists of Disk 10500, Disk20500, and Disk10800.
- Unit D104 is a single (JBOD) disk named Disk20100.
- Disk20800 is a spareset member.

Model 4254 Disk Enclosure Shelf 1 (dual-bus)														
Bus A								Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	DISK10000	DISK10100	DISK10200	DISK10300	DISK10400	DISK10500	DISK10800	DISK20000	DISK20100	DISK20200	DISK20300	DISK20400	DISK20500	DISK20800
	D100 R1	D100 R1	D101 S1 M1	D101 S1 M2	D102 M3	D103 S2	D103 S2	D100 R1	D104	D101 S1 M1	D101 S1 M2	D102 M3	D103 S2	spareset member

Figure 2-18. Example storage map-dual bus

Using the LOCATE Command to Find Devices

If you want to complete a storage map at a later time but do not remember where everything is, use the CLI command LOCATE. The LOCATE command flashes the (fault) LED on the drives associated with the specific storageset or unit. To turn off the flashing LEDs, enter the CLI command LOCATE CANCEL.

The following is an example of the commands needed to locate all the disk drives that make up unit D104:

1. Enter the following command:

```
LOCATE D104
```

The LEDs on the disk drives that make up unit D104 will flash.

2. Note the position of all the drives contained within D104.
3. Enter the following command to turn off the flashing LEDs:

LOCATE CANCEL

The following is an example command needed to locate all the drives that make up RAIDset R1:

1. Enter the following command:
- LOCATE R1
2. Note the position of all the drives contained within R1.
 3. Enter the following command to turn off the flashing LEDs:

LOCATE CANCEL

Chapter 3

Preparing the Host System

This chapter describes how to prepare the SGI Irix host computer to accommodate the HSG60 controller storage subsystem.

The following information is included in this chapter:

- “Making a Physical Connection,” page 3–1
- “Loading the StorageWorks Software onto the Hard Disk,” page 3–2
- “Host Device Initialization,” page 3–4

Making a Physical Connection

To attach the host computer to the storage subsystem, install one or more host bus adapters into the computer. A Fibre Channel (FC) cable goes from the host bus adapter to an FC switch or hub.

Installing the Host Bus Adapter

Once the RAID Array has been configured, perform the following procedure to install the host system.

1. Shut down the host operating system using your site-specific procedures. Once the operating system has halted, power off the SGI system enclosure.
2. If necessary, install the SGI Fibre Channel host adapter, and the optical to electrical Media Interface Adapter (MIA), FC-AL only. The MIA should be secured to the host adapter with its two attached fastening screws. The MIA is available from SGI using the part number X-F-OE-KIT

3. Before starting the host system, connect the optical cable(s) from the RAID controller to the hub or switch. Then connect the Fibre Channel host adapter in the SGI enclosure to the hub or switch with an optical cable.
4. Once the cables are securely connected, power on the SGI enclosure and start the operating system.
5. Once the bootstrap process has completed, log in to the system as superuser (*root*).

Loading the StorageWorks Software onto the Hard Disk

The StorageWorks (SWCC) software consists of the RAID Array software installation scripts, the support files, the Client, and the Agent. Installation and operation of the software requires temporary and permanent space on your hard disk. You need approximately 4 M-bytes of permanent space on your hard disk to use the StorageWorks software and an additional 2 Mbytes of temporary space on your hard disk for installation.

Loading the StorageWorks software on an SGI-Irix system requires the following major steps:

1. Mounting the CD-ROM which contains the software.
2. Extracting the files from the file archive to a temporary storage area.
3. Executing the installation and configuration script.

Copying the File Archive from CD-ROM to Temporary Storage

Copy the file archive that contains the StorageWorks software for SGI-Irix from the CD-ROM to a temporary directory, before running the installation script.

1. Insert the StorageWorks software distribution CD-ROM into the drive.
2. Log in as root (superuser).
3. Change the directory to a temporary storage area (/tmp). For example:

```
# cd /tmp
```
4. Create a directory into which the file archive can be loaded.

```
# mkdir HSG60
```

5. Change to that directory.

```
# cd HSG60
```

6. If the SGI-Irix automounter does not mount the CD-ROM, use the following command:

```
# mount -t iso9660 -r /dev/rdisk/dks0d7vol /mnt
```

NOTE: The name of *dks0d7vol* may be different on your system.

7. Extract the file archive from the CD-ROM by entering the following command:

```
# tar xvf /mnt/agent/sgi/stm_sgi.tar
```

All of the files necessary to complete the StorageWorks Software installation have now been copied to this directory.

Executing the Installation Script

Perform the following procedure to install the StorageWorks software onto the system:

1. Enter the following command and press **Enter**.

```
# ./stgwks_sgi.sh
```

The installation options menu appears.

2. Choose option 1 to install StorageWorks Software. The Copyright Notice and the following information will be displayed:

```
Copyright Compaq Computer Corp. 1994 - 2001. All rights reserved.
Restricted Rights: Use, duplication by the U.S. Government is subject to restriction as set forth in
subparagraph (c) (1) (ii) of DFARS 252.227-7013, or in FAR 52.227-19, or in FAR 52.227-14 Alt.III, as
applicable
Portions of this software are proprietary to and embody the confidential technology of Compaq Computer
Corporation. Possession, use, or copying of this software and media is authorized only pursuant to a valid
written license from Compaq or an authorized sublicensor.
Checking User...OK
Checking Irix Revision Level...OK
Checking that Host Adapter is installed and correctly configured...OK
Copying installation files to /usr/stgwks directory...OK
StorageWorks Software Installed
-- Hit RETURN to continue --
```

Once the StorageWorks software has been installed, logical units may be created on the RAID Array and configured into the system. These logical units can be created using the SWCC Client or the CLI. In order to utilize Client to create logical units, Agent must first be installed on the host SGI system. Details on Agent are in Chapter 4 and details on Client are in Appendix B. Once the Agent and Client have been installed, the following sections will guide you through the remaining steps in the configuration and initialization process.

Host Device Initialization

New logical units created on the RAID Array with SWCC or the CLI must be initialized and partitioned before they can be used by the host operating system. Once the logical unit has been created, the system must be restarted for the logical unit to be recognized by the operating system. The following steps outline the process.

1. Using the CLI or the SWCC Client, create the logical units on the RAID Array.
2. Restart Irix according to the established local system administration procedures.
3. Once the bootstrap process has completed, log in to the system as the superuser (root).
4. Verify that the host has located and configured the new RAID Array devices into the operating system with the command:

```
# hinv -v
```

The output of this command should show the logical units that were created on the RAID Array. In the following portion of the command output, an HSG60 with logical units at IDs 1,2,3,4,5 are connected to controller 10 of an XIO Arbitrated Loop Fibre Channel host adapter. In this example, the first disk drive, unit 64, logical unit 0, is the Command Console LUN (CCL).

```
Integral SCSI controller 10: Version Fibre Channel AIC-1160, revision 2
Disk drive: unit 64 on SCSI controller 10
Disk drive: unit 64 lun 1 on SCSI controller 10 1
Disk drive: unit 64 lun 2 on SCSI controller 10 2
Disk drive: unit 64 lun 3 on SCSI controller 10 3
Disk drive: unit 64 lun 4 on SCSI controller 10 4
Disk drive: unit 64 lun 5 on SCSI controller 10 5
Integral SCSI controller 3: Version QL1040B (rev. 2)
Integral SCSI controller 4: Version QL1040B (rev. 2)
Integral SCSI controller 5: Version QL1040B (rev. 2)
Integral SCSI controller 6: Version QL1040B (rev. 2)
Integral SCSI controller 7: Version QL1040B (rev. 2)
Integral SCSI controller 8: Version QL1040B (rev. 2)
Integral SCSI controller 9: Version QL1040B (rev. 2)
```

5. If the devices have been correctly recognized by the operating system, they can be labeled and partitioned. If the devices were not recognized, recheck the status of the RAID Controller, logical units, and Fibre Channel connections.
6. If the RAID controller is properly configured and the optical cables are connected, then use the following command sequence to re-initialize the loop and configure the devices into Irix that are attached to controller #10:

```
# scsiha -r -p 10
# cd /hw
# ioconfig -f./
# hinv -v
```

If necessary, repeat the sequence for each host controller with Fibre Channel RAID devices attached.

7. The logical units that were recognized by the operating system must be initialized before they can be used for data storage. This initialization includes:
 - Creating and writing the disk label
 - Partitioning the storage
 - Creating the file system, if required

Arbitrated Loop Device Name Convention

SGI Irix creates unique device names for devices attached to the Fibre Channel Arbitrated Loop host adapters on the system. Understanding the device nomenclature will help in providing device name-related answers in the following section.

Format used for LUN 0: *dk s x d y s a*

Format used for other LUNs: *dk s x d y l z s a*

NOTE: The “l” in the example is the letter “el”)

Where ...

x = controller (host Fibre Channel adapter) number assigned by host.

y = Loop ID in decimal

z = Logical Unit Number, LUN

a = Slice or partition number

The fourth slice (partition) on RAID Array logical unit D4 on Fibre Channel controller 10 would be seen as:

dk s 10d64l4s3

Fabric Device Name Convention

SGI Irix creates unique device names for devices attached to the Fabric Fibre Channel host adapters on the system. Understanding the device nomenclature will help in providing device name-related answers in the following section.

Format used for LUNs: `/dev/rdisk/<nodeWWN>/lunzvol/cxp<portID>`

NOTE: The “l” in the example is the letter “el”)

Where:

z (in lunzvol) = Logical Unit Number, LUN

x (in cxp<portID>) = controller (host Fibre Channel adapter) number assigned by host.

<nodeWWN> = worldwide name (Node ID) of HSG60 controllers (THIS and OTHER)

<portID> = Port ID of the HSG60 controller (THIS or OTHER)

Creating the Disk Label

The disk label contains a variety of information about the storage device. Included in the label are settings for a performance enhancement feature called Command Tagged Queuing (CTQ). In order to utilize CTQ, it must be enabled in the disk label for each RAID Array logical unit that is visible to the host. In the following example, CTQ is enabled and the queue depth is set.

NOTE: To partition the device, follow the prompts and enter the answers displayed in the large ***bold face italic*** type. The default answers are shown in () braces. Press the **Enter** key to utilize the default answer.

The controller number, drive number and LUN number of your drive will probably be different for Arbitrated Loop or Fabric.

The following entries are for Arbitrated Loop only.

```
# fx -x
fx version 6.5, Dec 31, 2000
fx: "device-name" = (dksc) dksc
fx: ctrl# = (0) 10
fx: drive# = (1) 64
fx: lun# = (0) 4
...opening dksc(10,64,4)
...drive selftest...OK
```


NOTE: The following entries are for Fabric only.

For Irix V6.5.9 or earlier, use:

```
# fx - x - d /dev/rdisk/50001fe1000007e0/lun0vol/c5p50001fe1000007e1
```

For Irix V6.5.10 or earlier, use:

```
# fx - x - d /dev/rdisk/50001fe1000007e0/lun0vol/c5p1
```

```
fx version 6.5, Dec 31, 2000
...opening c5p50001fe1000007e1
...drive selftest...OK
```

NOTE: The following entries are for both Arbitrated Loop and Fabric.

Scsi drive type == DEC HSG60 V86

fx: Warning: no sgilabel on disk

fx: Warning: can't read sgilabel on disk

creating new sgilabel

----- please choose one (? for help, .. to quit this menu)-----

```
[exi]t [d]ebug/ [l]abel/ [a]uto
[b]adblock/ [exe]rcise/ [r]epartition/ [f]ormat
fx> label
```

----- please choose one (? for help, .. to quit this menu)-----

```
[sh]ow/ [sy]nc [se]t/ [c]reate/
fx/label> set
```

----- please choose one (? for help, .. to quit this menu)-----

```
[para]meters [part]itions [s]giinfo
[g]eometry [m]anufacturer_ params [b]ootinfo
fx/label/set> parameters
```

```
fx/label/set/parameters: Error correction = (enabled)
fx/label/set/parameters: Data transfer on error = (enabled)
fx/label/set/parameters: Report recovered errors = (enabled) disabled
fx/label/set/parameters: Delay for error recovery = (enabled)
fx/label/set/parameters: Err retry count = (4)
fx/label/set/parameters: Transfer of bad data blocks = (enabled)
fx/label/set/parameters: Auto bad block reallocation (write) = (enabled)
fx/label/set/parameters: Auto bad block reallocation (read) = (enabled)
fx/label/set/parameters: Read ahead caching = (enabled)
fx/label/set/parameters: Write buffering = (disabled)
fx/label/set/parameters: Drive disable prefetch = (0)
fx/label/set/parameters: Drive minimum prefetch = (0)
fx/label/set/parameters: Drive maximum prefetch = (0)
fx/label/set/parameters: Drive prefetch ceiling = (0)
fx/label/set/parameters: Enable CTQ = (disabled) enabled
fx/label/set/parameters: CTQ depth = (2) 3
```

```
fx/label/set/parameters: Read buffer ratio = (0/256)
fx/label/set/parameters: Write buffer ratio = (0/256)
***** W A R N I N G * * * * *
about to modify drive parameters on disk C5P1 ok? yes

----- please choose one (? for help, .. to quit this menu)-----
[para]meters [part]itions [s]giinfo
[g]eometry [m]anufacturer_params [b]ootinfo
fx/label/set>

----- please choose one (? for help, .. to quit this menu)-----
[sh]ow/ [sy]nc [se]t/ [c]reate/
fx/label> ..

----- please choose one (? for help, .. to quit this menu)-----
[exi]t [d]ebug/ [l]abel/ [a]uto
[b]adblock/ [exe]rcise/ [r]epartition/ [f]ormat
fx> exit

label info has changed for disk C5P1. write out changes? (yes)
```

As Irix writes the disk label to the RAID Array, an error message is usually displayed on the system console. The error message reads:

```
dks10d64: [Alert] Illegal Request: Illegal Field in CDB.....
```

This message is caused by the Host adapter sending an illegal command to the RAID Array. The error message can be ignored.

Partitioning the Storage

After the disk label has been successfully written, the storage can be divided into the slices or partitions that are required by the user application. Since it is not possible to alter the partition structure of the storage device once it contains user data, this step should be carefully planned. A few factors to consider when deciding on the partition/slice size include:

- Total storage space required for application and user data, now and in the future
- Maximum amount of storage space in one partition that can be addressed by the application
- Backup time, strategy, and tool capabilities

The following simple example, shows ID 64, LUN 4 being partitioned with the default partition allocation used by the **fx** utility.

NOTE: To partition the device, follow the prompts and enter the answers displayed in the large **bold face italic** type. The default answers are shown in () braces. Press the **Enter** key to utilize the default answer

The controller number, drive number and LUN number of your drive will probably be different for arbitrated loop or fabric.

NOTE: The following entries are for Arbitrated Loop only.

```
# fx -x
fx version 6.5, Jan 20, 2000
fx: "device-name" = (dksc) dksc
fx: ctlr# = (0) 10
fx: drive# = (1) 64

fx: lun# = (0) 4
...opening dksc(10,64,4)
...drive selftest...OK
```

NOTE: The following entries are for Fabric only.

```
# fx -x -d /dev/rdisk/50001fe1000007e0/lun0vol/c5p50001fe1000007e1
fx version 6.5, Jan 20,, 2000
...opening c5p50001fe1000007e1
...drive selftest...OK
```

NOTE: The following entries are for both Arbitrated Loop and Fabric.

```
Scsi drive type == DEC HSG60 V86
----- please choose one (? for help, .. to quit this menu)-----
[ex]it [d]ebug/ [l]abel/
[b]adblock/ [ex]rcise/ [r]epartition/
fx> repartition

----- partitions-----
part type blocks Megabytes (base+size)
7: xfs 4096 + 20538224 2 + 10028
8: volhdr 0 + 4096 0 + 2
10: volume 0 + 20542320 0 + 10030
capacity is 20542320 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive [u]srootdrive [o]ptiondrive [re]size
fx/repartition> optiondrive

fx/repartition/optiondrive: type of data partition = (xfs) xfs
Warning: you will need to re-install all software and restore user data
from backups after changing the partition layout. Changing partitions will cause all
```

```
data on the drive to be lost. Be sure you have the drive
backed up if it contains any user data. Continue? yes

----- partitions-----
part type blocks Megabytes (base+size)
7: xfs 4096 + 20538224 2 + 10028
8: volhdr 0 + 4096 0 + 2
10: volume 0 + 20542320 0 + 10030
capacity is 20542320 blocks

----- please choose one (? for help, .. to quit this menu)-----
[ro]otdrive [u]srootdrive [o]ptiondrive [re]size
fx/repartition> ..

----- please choose one (? for help, .. to quit this menu)-----
[exi]t [d]ebug/ [l]abel/
[b]adblock/ [exe]rcise/ [r]epartition/
fx> exit
```

Creating Filesystems

Once the storage space has been divided into partitions, filesystems can be created. Filesystems can be constructed from one storage device or multiple storage devices. To create a filesystem using multiple storage devices, use the SGI XLV Volume Manager. The XLV utilities are easily accessed through the graphical system management tool, **sysmgr**.

To create and mount an XFS filesystem on an arbitrated loop host adapter, use:

```
# mkfs /dev/dsk/dks10d64l4s7
# mount /dev/dsk/dks10d64l4s7 /mnt
```

To create and mount an XFS filesystem on a fabric host adapter, use:

```
# mkfs /dev/dsk/50001fe1000007e0/lun0s7/c5p50001fe1000007e1
# mount /dev/dsk/50001fe1000007e0/lun0s7/c5p50001fe1000007e1 /mnt
```

The logical units on the StorageWorks RAID Array are now ready to use.

Chapter 4

Installing and Configuring the HS-Series Agent

The following information is included in this chapter:

- “Why Use StorageWorks Command Console (SWCC)?,” page 4–1
- “Installation and Configuration Overview,” page 4–2
- “About the Network Connection for the Agent,” page 4–3
- “Installing and Configuring the Agent,” page 4–5
- “Running the Agent,” page 4–8
- “Reconfiguring the Agent,” page 4–8
- “Removing the Agent,” page 4–8

Why Use StorageWorks Command Console (SWCC)?

StorageWorks Command Console (SWCC) enables you to monitor and configure the storage connected to the HSG60 controller. SWCC consists of Client and Agent.

- Client provides pager notification and lets you manage your virtual disks. Client runs on Windows 2000 with Service Pack 1 and Windows NT 4.0 (Intel) with Service Pack 6A.
- Agent obtains the status of the storage connected to the controller. It also passes the status of the devices connected to the controller to other computers and provides email notification and error logging.

4-2 HSG80 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

To receive information about the devices connected to your HSG60 controller over a TCP/IP network, you must install the Agent on a computer that is connected to a controller.

The Agent can also be used as a standalone application without Client. In this mode, which is referred to as Agent only, Agent monitors the status of the subsystem and provides local and remote notification in the event of a failure. A subsystem includes the HSG60 controller and its devices. Remote and local notification can be made by email, System Error Log and/or SNMP messages to an SNMP monitoring program.

Table 4-1 SWCC Features and Components

Features	Agent Required?	Client Required?
Able to create the following: <ul style="list-style-type: none">■ Striped device group (RAID 0)■ Mirrored device group (RAID 1)■ Striped mirrored device group (RAID 0+1)■ Striped parity device group (RAID 3/5)■ Individual device (JBOD)	Yes	Yes
Able to monitor many subsystems at once	Yes	Yes
Event logging	Yes	No
Email notification	Yes	No
Pager notification	Yes	Yes

NOTE: For serial and SCSI connections, the Agent is not required for creating virtual disks.

Table 4-2 provides an overview of the installation.

Table 4-2 Installation and Configuration Overview

Step	Procedure
1	Verify that your hardware has been set up correctly. See the previous chapters in this guide.
2	Verify that you have a network connection for the Client and Agent systems. See "About the Network Connection for the Agent" on page 4-3.

Table 4-2 Installation and Configuration Overview (Continued)


Step	Procedure
3	Verify that there is a LUN to communicate through. This can be either the CCL or a LUN that was created with the CLI. See "The Command Console LUN" described in Chapter 1.
4	Install the Agent (TCP/IP network connections) on a system connected to the HSG60 controller. See "Installing and Configuring the Agent," page 4-5.
5	Add the name of the Client system to the Agent's list of Client system entries (TCP/IP network connections). This can be done during installation or when reconfiguring the Agent.
6	Install the Client software on Windows 2000 with Service Pack 1 or Windows NT 4.0 (Intel) with Service Pack 6A. See Appendix B.
7	Add the name of the agent system to the Navigation Tree of each Client system that is on the Agent's list of Client system entries (TCP/IP network connections). See Appendix B.
8	Set up pager notification (TCP/IP network connections). Refer to "Setting Up Pager Notification" in the <i>Compaq StorageWorks Command Console Version 2.4 User Guide</i> .

About the Network Connection for the Agent

The network connection, shown in Figure 4-1, displays the subsystem connected to a hub or a switch. SWCC can consist of any number of Clients and Agents in a network. However, it is suggested that you install only one Agent (except for the HSZ22 Agent) on a computer.

By using a network connection, you can configure and monitor the subsystem from anywhere on the LAN. If you have a WAN or a connection to the Internet, monitor the subsystem with TCP/IP.

IMPORTANT: SWCC does not support the dynamic host configuration protocol (DHCP) or the Windows Internet Name Service (WINS).



CXO7240A

Figure 4-1. An example of a network connection

- | | |
|---|---|
| ① Agent system (has the Agent software) | ⑤ Hub or switch |
| ② TCP/IP Network | ⑥ HSG60 controller and its device subsystem |
| ③ Client system (has the Client software) | ⑦ Servers |
| ④ Fibre Channel cable | |

Installing and Configuring the Agent

IMPORTANT: Copy the installation and configuration script to a temporary directory before beginning this installation. For instructions, see Chapter 3.

The Agent will automatically start when you install it and will restart automatically if it stops for any reason.

1. Log on to the host system as root (superuser). You must have access to the superuser (root) account to perform this installation.
2. Change to the directory to which you copied the installation and configuration script from the CD-ROM. For example, if you installed the installation and configuration script to the /tmp/HSG directory, you would enter the following:

```
# cd /tmp/HSG80
```

3. To start the StorageWorks installation and configuration script, enter the following:

```
# ./stgwks_sgi.sh
```

The StorageWorks and Command Console Agent Installation and Configuration menu appears.

4. To install the Agent, select option **3**.
5. Follow the text-based prompts and enter the appropriate information.

NOTE: If you entered Y to start the Agent during the installation, the Agent starts. Its default is to restart automatically after a system restart or if it stops for any reason.

6. To remove the temporary directory, enter the following:

```
# rm -r /tmp/HSG80
```

Table 4-3 Information Required for Install Program Prompts

Information Required to Complete Installation	How to Respond
Target directory for installation	/usr/stgwks is the directory where StorageWorks will be installed. It contains the software copied in Chapter 3. The directory must have a minimum of 4 MB of space available.
*Port that the Command Console Client uses to communicate with the Command Console Agent for RAID subsystem configuration and maintenance.	*Default is 4998. If the default port is already used, press N and specify a new port.

Table 4-3 Information Required for Install Program Prompts (Continued)

Information Required to Complete Installation	How to Respond
*Port number that the Command Console Agent uses to send notifications to the Command Console Client.	*Default is 4999. If the default port is already used, press N and specify a new port.
The directory which contains the device special files.	Irix device special files are located in the /dev/scsi directory.
A password to gain access to the agent to configure the RAID subsystem with the Client.	Enter a 4-16 character password. The install program does not echo what you enter.
*The name of CC Agent manufacturer and the Enterprise object identifier for use with its MIB, so that it can operate with other SNMP applications.	*Enter Compaq Computer Corporation as the CC Agent manufacturer.
*The Enterprise Object Identifier (OID) is a vendor specific value used to identify SNMP messages.	*Enter: 36.2.15.21
*Names of the Host Systems on which you will run the Client.	*Enter the network names for the hosts from which you will run the Client.
Level of access a client has for communicating with the storage systems. The possible options are: 0 = Navigation Tree monitoring only 1 = Navigation Tree and Storage Window monitoring only 2 = Navigation Tree monitoring, Storage Window monitoring and configuration	You can give a client configuration level access only if you specified an access password.

Table 4-3 Information Required for Install Program Prompts (Continued)

Information Required to Complete Installation	How to Respond
<p>How do you want a CC Agent server to notify a client when an error condition occurs?</p> <p>Notification schemes available are:</p> <p>0 = No Error Notification</p> <p>1 = Notification via a TCP/IP Socket</p> <p>2 = Notification via the SNMP protocol</p> <p>3 = Notification via both TCP/IP and SNMP</p>	<p>Selecting this option results in:</p> <p>0 = No notification over network, but local notification through email and an entry in the CC Agent error log file.</p> <p>1 = Notification messages sent to the GUI and local notification sent through email and an entry in the CC Agent error log file.</p> <p>2 = Notification messages sent to the SNMP monitor and local notification sent through email and an entry in the CC Agent error log file.</p> <p>3 = Notification messages sent to the GUI, the SNMP monitor, and local notification sent through email and an entry in the CC Agent error log file.</p>
<p>Names of your subsystems.</p>	<p>Specify unique names for each subsystem in your configuration.</p>
<p>*The character special file name for accessing the storage system.</p>	<p>Enter the device name that your operating system assigned to the communications logical unit that you enabled in Appendix B.</p> <p>From the choices displayed select the device that corresponds to the "HSGCCL"</p>
<p>Do you want email notification of errors or failures on the RAID subsystem?</p> <p>*If yes to email notification, specify email address and notification level.</p> <p>1 = Fatal error only</p> <p>2 = Warning and fatal errors</p> <p>3 = Info, warning and fatal errors</p>	<p>Enter Y to receive email and syslog notification. Enter N for syslog notification only.</p> <p>*Enter email address and notification level for each user.</p> <p>1 = Notification for serious errors, which require immediate attention.</p> <p>2 = Notification for serious errors and warning messages.</p> <p>3 = Notification for serious errors, warnings and informational messages.</p>

The Agent program is installed. It was automatically started if you typed **Y** to start the Agent during the installation. It will restart automatically after a system restart or if it stops for any reason.

Running the Agent

The Agent program was automatically started when you installed it, and it was installed to restart automatically if the system is restarted or it stops for any reason.

The installation script places an entry in the *“/etc/inittab”* file to implement automatic execution of the Agent. The tag field in the file is *“steam”*.

Reconfiguring the Agent

You can reconfigure previously-installed Agents using the configuration and installation program. To reconfigure an Agent execute the *.Stgwks_sgi.sh* script in */usr/stgwks* the StorageWorks Installation and Configuration directory (*/usr/stgwks*) and follow the prompts to Agent maintenance.

Removing the Agent

You can remove the Agent program by using the Agent Maintenance option from the *Stgwks_sgi.sh* Bourne-shell script.

1. Run the *stgwks_sgi.sh* script from the */usr/stgwks* directory.
2. Select option 4, Agent Maintenance.
3. Select option 5, Uninstall Agent, from the Agent Configuration menu.

Chapter 5

Configuration Procedures

This chapter describes procedures to configure a subsystem that uses Fibre Channel fabric topology. In fabric topology, the controller connects to its hosts through switches.

The following information is included in this chapter:

- “Establishing a Local Connection,” page 5–2
- “Setting Up a Single Controller,” page 5–3
- “Setting Up a Controller Pair,” page 5–8
- “Configuring Devices,” page 5–14
- “Configuring a Stripeset,” page 5–15
- “Configuring a Mirrorset,” page 5–15
- “Configuring a RAIDset,” page 5–16
- “Configuring a Striped Mirrorset,” page 5–17
- “Configuring a Single-Disk (JBOD) Unit,” page 5–17
- “Configuring a Partition,” page 5–18
- “Assigning Unit Numbers and Unit Qualifiers,” page 5–19
- “Configuration Options,” page 5–20

Use the command line interpreter (CLI) or StorageWorks Command Console (SWCC) to configure a subsystem. This chapter uses CLI, which is the low-level interface to the controller. If you want to use SWCC for configuration, see the SWCC online help for assistance.


The configuration procedures assume that controllers and cache modules are installed in a fully functional and populated enclosure and that the PCMCIA cards are installed.

To install a controller or cache module and the PCMCIA card, see the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 Maintenance and Service Guide*.

Establishing a Local Connection

A local connection is required to configure the controller until a command console LUN (CCL) is established using the CLI. Communication with the controller can be through the CLI or SWCC.

The maintenance port, as shown in Figure 5-1, provides a way to connect a maintenance terminal. The maintenance terminal can be an EIA-423 compatible terminal or a computer running a terminal emulator program. The maintenance port accepts a standard RS-232 jack. The maintenance port cable shown in Figure 5-1 has a 9-pin connector molded onto the end for a PC connection. If you need a terminal connection or a 25-pin connection, you can order optional cabling.



① Maintenance Port Cable

② Maintenance Port

Figure 5-1. Maintenance port connection



CAUTION: The maintenance port generates, uses, and can radiate radio-frequency energy through its cables. This energy may interfere with radio and television reception. Disconnect all maintenance port cables when not communicating with the controller through the local connection.

Setting Up a Single Controller


Power Up and Establish Communication

1. Connect the computer or terminal to the controller as shown in Figure 5-1. The connection to the computer is through the COM1 or COM2 ports.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Configure the computer or terminal as follows:
 - 9600 baud
 - 8 data bits
 - 1 stop bit
 - no parity
5. Press **Enter**. A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

Cabling a Single controller

The cabling for a single controller is shown in Figure 5-2.

NOTE: It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure ("Configuring a Single Controller Using CLI," page 5-4).



- ❶ Controller
- ❷ Host port 1
- ❸ Host port 2
- ❹ Cable from the switch to the host Fibre Channel adapter
- ❺ FC switch

Figure 5-2. Single controller cabling

Configuring a Single Controller Using CLI

To configure a single controller using CLI involves the following processes:

- Verify the Node ID and Check for Any Previous Connections.
- Configure Controller Settings.
- Restart the Controllers.
- Set Time and Verify All Commands.
- Plug in the FC Cable and Verify Connections.
- Repeat Procedure for Each Host Adapter Connections
- Verify Installation

Verify the Node ID and Check for Any Previous Connections

1. Enter a SHOW THIS command to verify the node ID:
SHOW THIS

See “Worldwide Names (Node IDs and Port IDs),” page 1–17, for the location of the sticker.

The node ID is located in the third line of the SHOW THIS result:

```
HSG60> show this
```

Controller:

```
HSG60 ZG80900583 Software V8.6L-0, Hardware E11
```

```
NODE_ID      = 5000-1FE1-0001-3F00
```

```
ALLOCATION_CLASS = 0
```

- ❑ If the node ID is present, go to step 6.

If the node ID is all zeroes, enter the node ID and checksum, which are located on a sticker on the controller enclosure. Use the following syntax to enter the node ID:

```
SET THIS NODE_ID=NNNN-NNNN-NNNN-NNNN nn
```

NNNN-NNNN-NNNN-NNNN is the node ID, and nn is the checksum.

2. If the controller you are working with is not new from the factory, enter the following command to take it out of any failover mode that may have been configured previously:

```
SET NOFAILOVER
```

If the controller did have a failover mode previously set, the CLI may report an error. Clear the error with this command:

```
CLEAR_ERRORS CLI
```

3. Enter the following command to remove any previously configured connections:

```
SHOW CONNECTIONS
```

A list of named connections, if any, is displayed.

4. Delete these connections by entering the following command:

```
DELETE !NEWCON01
```

Repeat this delete command for each of the listed connections. When completed, no connections will be displayed.

Configure Controller Settings

5. Set the SCSI version to SCSI-2 , using the following command.

```
SET THIS SCSI_VERSION=SCSI-2
```

Setting the SCSI version to SCSI-3 does not make the controller fully compliant with the SCSI-3 standards.

6. Set the topology for the controller. If both ports are used, set topology for both ports:

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

If the controller is not new from the factory, it may have another topology set, in which case these commands will result in an error message. If this happens, take both ports offline first, then reset the topology:

```
SET THIS PORT_1_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_2_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

Restart the Controllers

7. Restart the controller, using the following command:

```
RESTART THIS
```

It takes about one minute for the CLI prompt to come back after a RESTART command.

Set Time and Verify All Commands

8. Set the time on the controller using the following syntax:

```
SET THIS TIME=DD-MMM-YYYY:HH:MM:SS
```

9. Use the FRUTIL utility to set up the battery discharge timer. Enter the following command to start FRUTIL:

```
RUN FRUTIL
```

When FRUTIL asks if you intend to replace the battery, answer “Y”:

```
Do you intend to replace this controller's cache battery? Y/N [N] Y
```

FRUTIL will print out a procedure, but won't give you a prompt. Ignore the procedure and just press enter.

10. Set up any additional optional controller settings, such as changing the CLI prompt. See the SET *this_controller/other controller* command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* for the format of optional settings.
11. Verify that all commands have taken effect. Use the following command:

```
SHOW THIS
```

Verify node ID, allocation class, SCSI version, failover mode, identifier, and port topology. Following is a sample result of a SHOW THIS command, with the areas of interest in bold.

Controller:

HSG60 (C) DEC ZG09030200 Software V86L, Hardware 0000

NODE_ID = 5000-1FE1-0000-0000

ALLOCATION_CLASS = 01

SCSI_VERSION = SCSI-2

Not configured for dual-redundancy

Device Port SCSI address 7

Time: 10-Mar-1999:12:30:34

Command Console LUN is lun 0

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0000-0001

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 210313

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0000-0002

PORT_2_TOPOLOGY = FABRIC (fabric up)

Address = 210513

NOREMOTE_COPY

.....

12. Turn on the switches, if not done previously.

If you want to communicate with the FC switches through Telnet, set an IP address for each switch. See the manuals that came with the switches for details.

Plug in the FC Cable and Verify Connections

13. Plug in the Fibre Channel cable from the first adapter in the first host into the switch.

Enter a *SHOW connections* command to view the connection table:

SHOW CONNECTIONS

The first connection will have one or more entries in the connection table. Each connection will have a default name of the form !NEWCONxx, where xx is a number representing the order in which the connection was added to the connection table.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON00	WNNT	OTHER	1	200613	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

For an explanation about why there may be multiple connections shown for one adapter, see page 1-11.

14. Rename the connection(s) to something meaningful to the system and easy to remember. For example, to assign the name ANGEL1A1 to connection !NEWCON01, enter:

```
RENAME !NEWCON01 ANGEL1A1
```

For a recommended naming convention, see “Naming Connections,” page 1-11.

15. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=SGI
```

16. Verify the changes:

```
SHOW CONNECTIONS
```

Mark or tag all Fibre Channel cables at both ends for ease of maintenance.

Repeat Procedure for Each Host Adapter Connections

17. Repeat step 14, 15, and 16 for each of that adapter’s host connections, or delete the unused connections from the table.
18. For each host adapter, repeat steps 13 through 17.

Verify Installation

To verify installation for your SGI Irix host, enter the following command:

```
SHOW DEVICES
```

Setting Up a Controller Pair

Power Up and Establish Communication


1. Connect the computer or terminal to the controller as shown in Figure 5-1. The connection to the computer is through the COM1 or COM2 ports.

2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Configure the computer or terminal as follows:
 - 9600 baud
 - 8 data bits
 - 1 stop bit
 - no parity
5. Press **Enter**. A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

Cabling a Controller Pair

The cabling for a controller pair failover showing one HBA per server with HSG60 controller in transparent failover mode is shown in Figure 5-3.

NOTE: It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure ("Configuring a Controller Pair Using CLI," page 5-10).



- | | |
|----------------|--|
| ❶ Controller A | ❷ Host port 2 |
| ❸ Controller B | ❸ Cable from the switch to the host FC adapter |
| ❹ Host port 1 | ❹ FC switch |

Figure 5-3. Controller pair failover cabling

Configuring a Controller Pair Using CLI

To configure a controller pair using CLI involves the following processes:

- Verify the Node ID and Check for Any Previous Connections.
- Configure Controller Settings.
- Restart the Controller.
- Plug in the FC Cable and Verify Connections.
- Repeat Procedure for Each Host Adapter.
- Verify Installation.

Verify the Node ID and Check for Any Previous Connections

1. Enter a SHOW THIS command to verify the node ID:

```
SHOW THIS
```

See “Worldwide Names (Node IDs and Port IDs),” page 1-17, for the location of the sticker.

The node ID is located in the third line of the SHOW THIS result:

```
HSG60> show this
```

```
Controller:
```

```
  HSG60 ZG80900583 Software V8.6L-0, Hardware E11
```

```
  NODE_ID      = 5000-1FE1-0001-3F00
```

```
  ALLOCATION_CLASS = 0
```

If the node ID is present, go to step 6.

If the node ID is all zeroes, enter the node ID and checksum, which are located on a sticker on the controller enclosure. Use the following syntax to enter the node ID:

```
SET THIS NODE_ID=NNNN-NNNN-NNNN-NNNN nn
```

NNNN-NNNN-NNNN-NNNN is the node ID, and nn is the checksum.

2. If working with anything but a factory-new controller, enter the following command to take it out of any failover mode that may have been previously configured:

```
SET NOFAILOVER
```

If the controller did have a failover mode previously set, the CLI may report an error. Clear the error with this command:

```
CLEAR_ERRORS CLI
```

3. Enter the following command to remove any previously configured connections:

```
SHOW CONNECTIONS
```

A list of named connections, if any, is displayed.

4. Delete these connections by entering the following command:

```
DELETE !NEWCON01
```

Repeat this delete command for each of the listed connections. When completed, no connections will be displayed.

Configure Controller Settings

5. Set the version to SCSI-2 using the following command:

```
SET THIS SCSI_VERSION=SCSI-2
```

Setting the SCSI version to SCSI-3 does not make the controller fully compliant with the SCSI-3 standards.

6. Set the topology for the controller. If both ports are used, set topology for both ports:

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

If the controller is not factory-new, it may have another topology set, in which case these commands will result in an error message. If this happens, take both ports offline first, then reset the topology:

```
SET THIS PORT_1_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_2_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

Restart the Controller

7. Restart the controller, using the following command:

```
RESTART THIS
```

It takes about one minute for the CLI prompt to come back after a RESTART command.

Set Time and Verify All Commands

8. Set the time on the controller using the following syntax:

```
SET THIS TIME=DD-MMM-YYYY:HH:MM:SS
```

9. Use the FRUTIL utility to set up the battery discharge timer. Enter the following command to start FRUTIL:

```
RUN FRUTIL
```

When FRUTIL asks if you intend to replace the battery, answer “**Y**”:

```
Do you intend to replace this controller's cache battery? Y/N [N] Y
```

FRUTIL will print out a procedure, but will not give you a prompt. Ignore the procedure and just press **Enter**.

10. Set up any additional optional controller settings, such as changing the CLI prompt. See the SET *this_controller/other controller* command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* for the format of optional settings.
11. Verify that all commands have taken effect. Use the following command:

```
SHOW THIS
```


Verify node ID, allocation class, SCSI version, failover mode, identifier, and port topology. The following display shows the result of a SHOW THIS command, with the areas of interest in bold.

Controller:

HSG60 (C) DEC ZG09030200 Software V8.6L, Hardware E11

NODE_ID = 5000-1FE1-0000-0000

ALLOCATION_CLASS = 1

SCSI_VERSION = SCSI-2

configured for dual-redundancy

Device Port SCSI address 7

Time: 10-Mar-2001:12:30:34

Command Console LUN is lun 0

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0000-0001

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 210313

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0000-0002

PORT_2_TOPOLOGY = FABRIC (fabric up)

Address = 210513

NOREMOTE_COPY

.....

12. Turn on the switches, if not done previously.

If you want to communicate with the FC switches through Telnet, set an IP address for each switch. See the manuals that came with the switches for details.

Plug in the FC Cable and Verify Connections

13. Plug in the Fibre Channel cable from the first adapter on the first host into the switch.

Enter a SHOW connections command to view the connection table:

SHOW CONNECTIONS

The first connection will have one or more entries in the connection table. Each connection will have a default name of the form !NEWCONxx, where xx is a number representing the order in which the connection was added to the connection table.

For a description of why plugging in one adapter can result in multiple connections, see “Numbers of Connections,” page 1-11.

14. Rename the connection(s) to something meaningful to the system and easy to remember. For example, to assign the name ANGEL1A1 to connection !NEWCON01, enter:

```
RENAME !NEWCON01 ANGEL1A1
```

Compaq recommends using a naming convention, see “Naming Connections,” page 1-11.

15. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=SGI
```

16. Verify the changes:

```
SHOW CONNECTIONS
```

Mark or tag all Fibre Channel cables at both ends for ease of maintenance.

Repeat Procedure for Each Host Adapter

17. Repeat steps 14, 15, and 16 for each of that adapter’s host connections or delete the unwanted connections from the table.

18. For each host adapter, repeat steps 13 through 17.

Verify Installation

To verify installation for your SGI Irix host, enter the following command:

```
SHOW DEVICES
```

Configuring Devices

The disks on the device bus of the HSG60 can be configured manually or with the CONFIG utility. The CONFIG utility is easier. Invoke CONFIG with the following command:

```
RUN CONFIG
```

CONFIG takes about 2 minutes to discover and to map the configuration of a completely populated storage system.

Configuring a Stripese

1. Create the stripese by adding its name to the controller's list of storagesets and by specifying the disk drives it contains. Use the following syntax:
ADD STRIPESSET STRIPESSET-NAME DISKNNNNN DISKNNNNN
2. Initialize the stripese, specifying any desired switches:
INITIALIZE STRIPESSET-NAME SWITCHES
See "Initialization Switches" on page 2-23 for a description of the initialization switches.
3. Verify the stripese configuration:
SHOW STRIPESSET-NAME
4. Assign the stripese a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5-19.

For example:

The commands to create Stripe1, a stripese consisting of three disks (DISK10000, DISK20000, and DISK10100) and having a chunksize of 128:

```
ADD STRIPESSET STRIPE1 DISK10000 DISK20000 DISK10100
INITIALIZE STRIPE1 CHUNKSIZE=128
SHOW STRIPE1
```

Configuring a Mirrorset

1. Create the mirrorset by adding its name to the controller's list of storagesets and by specifying the disk drives it contains. Optionally, you can append mirrorset switch values:
ADD MIRRORSET MIRRORSET-NAME DISKNNNNN DISKNNNNN SWITCHES
NOTE: See the ADD MIRRORSET command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* for a description of the mirrorset switches.
2. Initialize the mirrorset, specifying any desired switches:
INITIALIZE MIRRORSET-NAME SWITCHES
See "Initialization Switches" on page 2-23 for a description of the initialization switches.
3. Verify the mirrorset configuration:
SHOW MIRRORSET-NAME

4. Assign the mirrorset a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5-19.

For example:

The commands to create Mirr1, a mirrorset with two members (DISK10000 and DISK20000), and to initialize it using default switch settings:

```
ADD MIRRORSET MIRR1 DISK10000 DISK20000
INITIALIZE MIRR1
SHOW MIRR1
```

Configuring a RAIDset

1. Create the RAIDset by adding its name to the controller's list of storagesets and specifying the disk drives it contains. Optionally, you can specify RAIDset switch values:

```
ADD RAIDSET RAIDSET-NAME DISKNNNNN DISKNNNNN DISKNNNNN SWITCHES
```

NOTE: See the ADD RAIDSET command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* for a description of the RAIDset switches.

2. Initialize the RAIDset, specifying any desired switches:

```
INITIALIZE RAIDSET-NAME SWITCH
```

NOTE: Compaq recommends that you allow initial reconstruct to complete before allowing I/O to the RAIDset. Not doing so may generate forced errors at the host level. To determine whether initial reconstruct has completed, enter SHOW RAIDSET FULL.

See “Initialization Switches” on page 2-23 for a description of the initialization switches.

3. Verify the RAIDset configuration:

```
SHOW RAIDSET-NAME
```

4. Assign the RAIDset a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5-19.

For example:

The commands to create RAID1, a RAIDset with three members (DISK10000, DISK20000, and DISK10100) and to initialize it with default values:

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK10100
INITIALIZE RAID1
SHOW RAID1
```

Configuring a Striped Mirrorset

1. Create, but do not initialize, at least two mirrorsets.
See “Configuring a Mirrorset” on page 5-15.
2. Create a stripeset and specify the mirrorsets it contains:
ADD STRIPESET STRIPESET-NAME MIRRORSET-1 MIRRORSET-2...MIRRORSET-N
3. Initialize the striped mirrorset, specifying any desired switches:
INITIALIZE STRIPESET-NAME SWITCH
See “Initialization Switches” on page 2-23 for a description of the initialization switches.
4. Verify the striped mirrorset configuration:
SHOW STRIPESET-NAME
5. Assign the stripeset mirrorset a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5-19.

For example:

The commands to create Stripe1, a striped mirrorset that comprises Mirr1, Mirr2, and Mirr3, each of which is a two-member mirrorset:

```
ADD MIRRORSET MIRR1 DISK10000 DISK20000
ADD MIRRORSET MIRR2 DISK20100 DISK10100
ADD MIRRORSET MIRR3 DISK10200 DISK20200
ADD STRIPESET STRIPE1 MIRR1 MIRR2 MIRR3
INITIALIZE STRIPE1
SHOW STRIPE1
```

Configuring a Single-Disk (JBOD) Unit

1. Initialize the disk drive, specifying any desired switches:
INITIALIZE DISK-NAME SWITCHES
See “Initialization Switches” on page 2-23 for a description of the initialization switches.
2. Verify the configuration using the following command:
SHOW DISK-NAME

3. Assign the disk a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5-19.

Configuring a Partition

1. Initialize the storageset or disk drive, specifying any desired switches:

```
INITIALIZE STORAGESET-NAME SWITCHES
```

or

```
INITIALIZE DISK-NAME SWITCHES
```

See “Initialization Switches” on page 2-23 for a description of the initialization switches.

2. Create each partition in the storageset or disk drive by indicating the partition's size. Also specify any desired switch settings:

```
CREATE_PARTITION STORAGESET-NAME SIZE=N SWITCHES
```

or

```
CREATE_PARTITION DISK-NAME SIZE=N SWITCHES
```

where N is the percentage of the disk drive or storageset that will be assigned to the partition. Enter SIZE=LARGEST to let the controller assign the largest free space available to the partition.

NOTE: See the CREATE_PARTITION command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide* for a description of the partition switches.

3. Verify the partitions:

```
SHOW STORAGESET-NAME
```

or

```
SHOW DISK-NAME
```

The partition number appears in the first column, followed by the size and starting block of each partition.

4. Assign the partition a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5-19.

For example:

The commands to create RAID1, a three-member RAIDset, then partition it into two storage units:

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK10100
INITIALIZE RAID1
CREATE_PARTITION RAID1 SIZE=25
CREATE_PARTITION RAID1 SIZE=LARGEST
SHOW RAID1
```

Assigning Unit Numbers and Unit Qualifiers

Each storageset, partition, or single (JBOD) disk must be assigned a unit number for the host to access. As the units are added, their properties can be specified through the use of command qualifiers, which are discussed in detail under the ADD UNIT command in the *Compaq StorageWorks HSG60 Array Controller ACS Version 8.6 CLI Reference Guide*.

Each unit can be reserved for the exclusive use of a host or group of hosts. See “Restricting Host Access in Transparent Failover Mode,” page 1-17.

Assigning a Unit Number to a Storageset

To assign a unit number to a storageset, use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESET-NAME
```

For example:

To assign unit D102 to RAIDset R1, use the following command:

```
ADD UNIT D102 R1
```

Assigning a Unit Number to a Single (JBOD) Disk

To assign a unit number to a single (JBOD) disk, use the following syntax:

```
ADD UNIT UNIT-NUMBER DISK-NAME
```

For example:

To assign unit D4 to DISK20300, use the following command:

```
ADD UNIT D4 DISK20300
```

Assigning a Unit Number to a Partition

To assign a unit number to a partition, use the following syntax:

```
ADD UNIT UNIT-NUMBER STORAGESET-NAME PARTITION=PARTITION-NUMBER
```

For example:

To assign unit D100 to partition 3 of mirrorset mirr1, use the following command:

```
ADD UNIT D100 MIRR1 PARTITION=3
```

Configuration Options

Changing the CLI Prompt

To change the CLI prompt, enter a 1- to 16- character string as the new prompt, as follows:

```
SET THIS_CONTROLLER PROMPT = "NEW PROMPT"
```

If you are configuring dual-redundant controllers, also change the CLI prompt on the "other controller." Use the following command:

```
SET OTHER_CONTROLLER PROMPT = "NEW PROMPT"
```

It is suggested that the prompts reflect something about the controllers. For example, if the subsystem is the third one in a lab, give the top controller a prompt like LAB3A and the bottom controller, LAB3B.

Mirroring cache

To specify mirrored cache:

```
SET THIS MIRRORED_CACHE
```

Adding Disk Drives

If you add new disk drives to the subsystem, the disk drives must be added to the controllers' list of known devices:

- To add one new disk drive to the list of known devices, enter the following command:

```
ADD DISK DISKNNNNN P T L
```


- To add several new disk drives to the list of known devices, enter the following command:

```
RUN CONFIG
```

Adding a Disk Drive to the Spareset

The spareset is a collection of spare disk drives that are available to the controller should it need to replace a failed member of a RAIDset or mirrorset.

NOTE: This procedure assumes that the disks that you are adding to the spareset have already been added to the controller's list of known devices.

To add the disk drive to the controller's spareset list:

```
ADD SPARESET DISKNNNNN
```

Repeat this step for each disk drive you want to add to the spareset:

For example:

The following example shows the commands for adding DISK11300 and DISK21300 to the spareset.

```
ADD SPARESET DISK11300
```

```
ADD SPARESET DISK21300
```

Removing a Disk Drive from the Spareset

You can delete disks in the spareset if you need to use them elsewhere in your subsystem.

NOTE: The RUN CONFIG command does not delete disks from the controllers' device table if a disk has been physically removed or replaced. In this case, use the command DELETE DISKNNNNN.

1. Show the contents of the spareset using the following command:

```
SHOW SPARESET
```

2. Delete the desired disk drive using the following command:

```
DELETE SPARESET DISKNNNNN
```

3. Verify the contents of the spareset using the following command:

```
SHOW SPARESET
```

Enabling Autospare

With AUTOSPARE enabled on the failedset, any new disk drive that is inserted into the PTL location of a failed disk drive is automatically initialized and placed into the spareset. If initialization fails, the disk drive remains in the failedset until you manually delete it from the failedset.

To enable autospare, enter the following command:

```
SET FAILEDSET AUTOSPARE
```

To disable autospare, enter the following command:

```
SET FAILEDSET NOAUTOSPARE
```

During initialization, AUTOSPARE checks to see if the new disk drive contains metadata. Metadata is information the controller writes on the disk drive when that the disk drive is configured into a storageset. Therefore, the presence of metadata indicates the disk drive belongs to, or has been used by, a storageset. If the disk drive contains metadata, initialization stops. (A new disk drive will not contain metadata but a repaired or reused disk drive might. To erase metadata from a disk drive, add it to the controller's list of devices, then set it to be nottransportable and initialize it.)

Deleting a Storageset

NOTE: If the storageset you are deleting is partitioned, delete each partitioned unit before you can delete the storageset.

1. Show the storageset's configuration:

```
SHOW STORAGESET-NAME
```

2. Delete the unit number that uses the storageset. Use the following command:

```
DELETE UNIT-NUMBER
```

3. Delete the storageset. Use the following command:

```
DELETE STORAGESET-NAME
```

4. Verify the configuration:

```
SHOW STORAGESET-NAME
```

Changing Switches for a StorageSet or Device

You can optimize a storageSet or device at any time by changing the switches that are associated with it. Remember to update the storageSet's profile when you change its switches.

Displaying the Current Switches

To display the current switches for a storageSet or single-disk unit, enter a `SHOW` command, specifying the `FULL` switch:

```
SHOW STORAGESET-NAME  
or  
SHOW DEVICE-NAME FULL
```

Changing RAIDset and Mirrorset Switches

Use the `SET storageSet-name` command to change the RAIDset and Mirrorset switches associated with an existing storageSet. For example, the following command changes the replacement policy for RAIDset RAID1 to `BEST_FIT`:

```
SET RAID1 POLICY=BEST_FIT
```

Changing Device Switches

Use the `SET device-name` command to change the device switches. For example, to request a data transfer rate of 20 MHz for DISK10000:

```
SET DISK10000 TRANSFER_RATE_REQUESTED=20MHZ
```

Changing Initialize Switches

The initialization switches cannot be changed without destroying the data on the storageSet or device. They can only be changed by reinitializing the storageSet. This is similar to formatting a disk drive; all data is destroyed during this procedure.

Changing Unit Switches

Use the `SET unit-name` command to change the characteristics of a unit. For example, the following command enables write protection for unit D100:

```
SET D100 WRITE_PROTECT
```


Chapter 6

Configuration Example Using CLI

This chapter presents an example of configuring a particular storage subsystem using the CLI. The CLI configuration example shown assumes:

- A normal, new controller pair, meaning:
 - ❑ NODE ID set
 - ❑ No previous failover mode
 - ❑ No previous topology set
- Two single-bus model 4214R disk enclosure shelves
- PCMCIA cards installed in both controllers

A storage subsystem example is shown in Figure 6–1. The example system contains three non-clustered Irix hosts, as shown in Figure 6–2. The resulting virtual system, from the hosts' point of view, is shown in Figure 6–4.

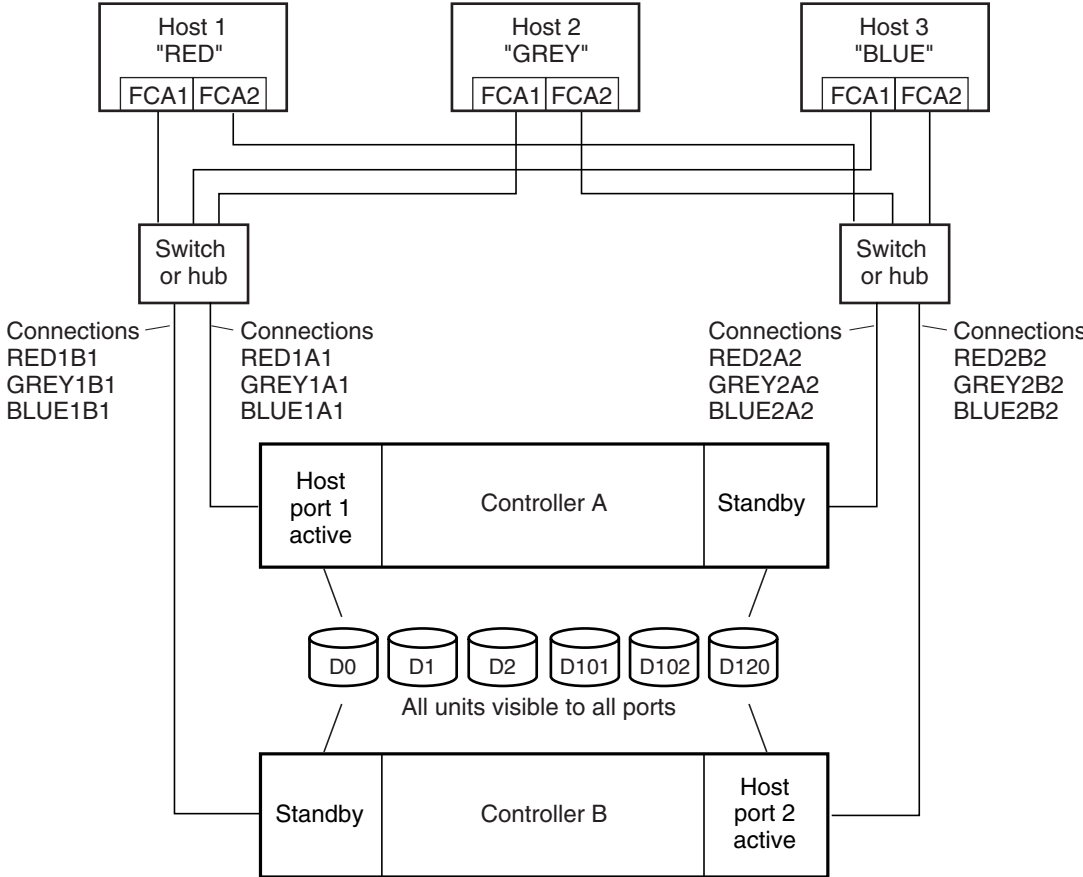
Figure 6–1 shows an example storage system map for the 4214R enclosure. Details on building your own map are described in Chapter 2. Templates to help you build your storage map are supplied in Appendix A.

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D0 R1 DISK10000	D0 R1 DISK10100	D0 R1 DISK10200	D1 S1 M1 DISK10300	D1 S1 M1 DISK10400	D1 S1 M2 DISK10500	D1 S1 M2 DISK10800	D100 M3 DISK10900	D100 M3 DISK11000	D120 S2 DISK11100	D120 S2 DISK112000	Spare DISK11300	not supported	not supported

Model 4214R Disk Enclosure Shelf 2 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D0 R1 DISK20000	D0 R1 DISK20100	D0 R1 DISK20200	D1 S1 M1 DISK20300	D1 S1 M1 DISK20400	D1 S1 M2 DISK20500	D1 S1 M2 DISK20800	D100 M3 DISK20900	D100 M3 DISK21000	D120 S2 DISK21100	D120 S2 DISK21200	Spare DISK21300	not supported	not supported

Figure 6-1. Example storage map

The example system shown in Figure 6-2, contains three non-clustered Irix hosts. The port 1 link is separate from the port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another) therefore, each adapter has two connections.

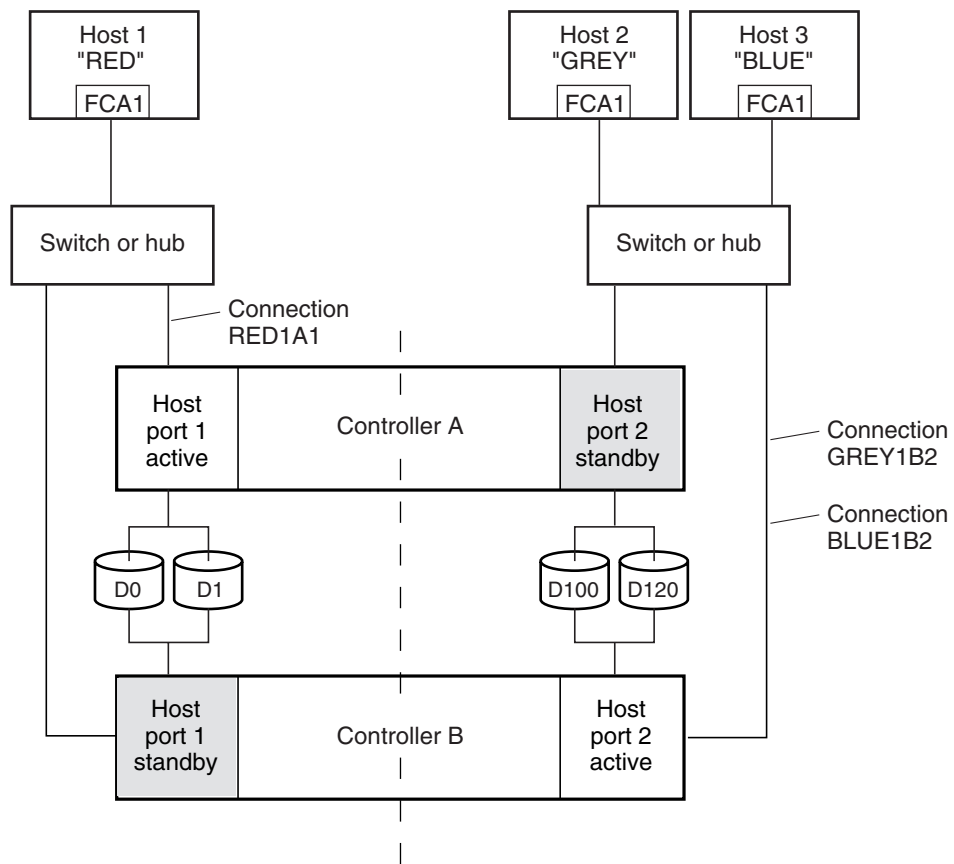


NOTE: FCA = Fibre Channel Adapter

CXO7547A

Figure 6-2. Example system


The following graphic represents units that are logical or virtual disks comprised of storagesets configured from physical disks.



FCA = Fibre Channel Adapter

CXO7296B

Figure 6-3. Example system



CXO7297B

Figure 6-4. Example virtual layout from the hosts' point of view

CLI Configuration Example

The series of commands and information presented in this section provides a CLI configuration example. Text conventions used in this section are listed below:

- Text in *italics* indicates an action you take.
- Text in THIS FORMAT, indicates a command you type. Be certain to press **Enter** after each command.
- Text enclosed within a box, indicates information that is displayed by the CLI interpreter.
- **NOTE** text provides additional helpful information.

NOTE: "This" controller is top controller (A).

Plug serial cable from maintenance terminal into top controller.

```
CLEAR CLI
SET FAILOVER COPY=THIS
CLEAR CLI
SET THIS SCSI_VERSION=SCSI-2
SET THIS PORT_1_TOPOLOGY=FABRIC
SET THIS PORT_2_TOPOLOGY=FABRIC
SET OTHER PORT_1_TOPOLOGY=FABRIC
SET OTHER PORT_2_TOPOLOGY=FABRIC
RESTART OTHER
RESTART THIS
SET THIS TIME=10-Mar-2001:12:30:34
RUN FRUTIL
```

Do you intend to replace this controller's cache battery? Y/N [N]

Y

NOTE: If "N" is selected, there is no warning when the battery is low.

Plug serial cable from maintenance terminal into bottom controller.

NOTE: Bottom controller (B) becomes “this” controller.

RUN FRUTIL

Do you intend to replace this controller's cache battery? Y/N [N]

Y

SET THIS MIRRORED_CACHE

NOTE: This command causes the controllers to restart.

SET THIS PROMPT=“HLNDR BOTTOM”

SET OTHER PROMPT=“HLNDR TOP”

SHOW THIS

SHOW OTHER

Plug in the Fibre Channel cable from the adapter in host “PURPLE.”

SHOW CONNECTIONS

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON00	WNNT	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				

RENAME !NEWCON00 PURPLE1A1

SET PURPLE1A1 OPERATING_SYSTEM=SGI

SHOW CONNECTIONS

NOTE: Connection table sorts alphabetically.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
RED1A1		OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				

Mark or tag both ends of Fibre Channel cables.

Plug in the Fibre Channel cable from the second adapter in host “WHITE.”

SHOW CONNECTIONS

NOTE: The offset for !NEWCON01 is automatically set to 100 because it is plugged into host port 2.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON01	WNNT	THIS	2	XXXXXX	OL this	100
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
PURPLE1A1	SGI	OTHER	1	XXXXXX	OL other	0
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		

```

RENAME !NEWCON01 WHITE1B2
SET WHITE1B2 OPERATING_SYSTEM=SGI
SHOW CONNECTIONS
    
```

Mark or tag both end of Fibre Channel cables.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
PURPLE1A1	SGI	OTHER	1	XXXXXX	OL other	0
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
WHITE1B2	SGI	THIS	2	XXXXXX	OL this	100
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		

Plug in the Fibre Channel cable from the adapter in host "TAN."

```
SHOW CONNECTIONS
```

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON02	SGI	THIS	2	XXXXXX	OL this	100
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
PURPLE1A1	SGI	OTHER	1	XXXXXX	OL other	0
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
WHITE1B2	SGI	THIS	2	XXXXXX	OL this	100
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		

```

RENAME !NEWCON02 TAN1B2
SET TAN1B2 OPERATING_SYSTEM=SGI
    
```

6-8 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

```
SET TAN1B2 UNIT_OFFSET=120
```

```
SHOW CONNECTIONS
```

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
PURPLE1A1	SGI	OTHER	1	XXXXXX	OL other	0
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
WHITE1B2	SGI	THIS	2	XXXXXX	OL this	100
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
TAN1B2	SGI	THIS	2	XXXXXX	OL this	120
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		

Mark or tag both end of Fibre Channel cables.

```
RUN CONFIG
```

```
ADD RAIDSET R1 DISK10000 DISK20000 DISK10100 DISK20100 DISK10200 DISK20200
```

```
INITIALIZE R1
```

```
ADD UNIT D0R1 DISABLE_ACCESS_PATH=ALL
```

```
SET D0 ENABLE_ACCESS_PATH=PURPLE1A1
```

```
ADD MIRRORSET MI DISK10300 DISK20300 DISK10400 DISK20400
```

```
ADD MIRRORSET M2 DISK10500 DISK20500 DISK10800 DISK20800
```

```
ADD STRIPESET S1 M1 M2
```

```
INITIALIZE S1
```

```
ADD UNIT D1 S1 DISABLE_ACCESS_PATH=ALL
```

```
SET D1 ENABLE_ACCESS_PATH=PURPLE1A1
```

```
ADD MIRRORSET M3 DISK10900 DISK20900 DISK11000 DISK21000
```

```
INITIALIZE M3
```

```
ADD UNIT D100 M3 DISABLE_ACCESS_PATH=ALL
```

```
SET D100 ENABLE_ACCESS_PATH=(WHITE1B2)
```

```
ADD STRIPESET S2 DISK11100 DISK11200 DISK21100 DISK21200
```

```
INITIALIZE S2
```

```
ADD UNIT D120 S2 DISABLE_ACCESS_PATH=ALL
```

```
SET D120 ENABLE_ACCESS_PATH=(TAN1B2)
```

```
ADD SPARESET DISK11300  
ADD SPARESET DISK21300  
SHOW UNITS ALL
```


Chapter 7

Backing Up the Subsystem, Cloning Data for Backup, and Moving Stagesets

This chapter describes some common procedures that are not mentioned previously in this guide. The following information is included in this chapter:

- “Backing Up the Subsystem Configuration,” page 7-1
- “Cloning Data for Backup,” page 7-2
- “Moving Stagesets,” page 7-5

Backing Up the Subsystem Configuration

The controller stores information about the subsystem configuration in its nonvolatile memory. This information could be lost if the controller fails or when you replace a module in the subsystem.

You can use the following command to produce a display that shows if the save configuration feature is active and which devices are being used to store the configuration.

```
SHOW THIS_CONTROLLER FULL
```

The resulting display includes a line that indicates status and how many devices have copies of the configuration. The last line shows on how many devices the configuration is backed up.

The `SHOW devices FULL` command shows which disk drive are set up to back up the configuration. The syntax for this command is shown below:


```
SHOW DEVICES FULL
```

Cloning Data for Backup

Use the CLONE utility to duplicate the data on any unpartitioned single-disk unit, stripeset, mirrorset, or striped mirrorset in preparation for backup. When the cloning operation is complete, you can back up the clones rather than the storageset or single-disk unit, which can continue to service its I/O load. When you are cloning a mirrorset, CLONE does not need to create a temporary mirrorset. Instead, it adds a temporary member to the mirrorset and copies the data onto this new member.

The CLONE utility creates a temporary, two-member mirrorset for each member in a single-disk unit or stripeset. Each temporary mirrorset contains one disk drive from the unit you are cloning and one disk drive onto which CLONE copies the data. During the copy operation, the unit remains online and active so that the clones contain the most up-to-date data.

After the CLONE utility copies the data from the members to the clones, it restores the unit to its original configuration and creates a clone unit you can backup. The CLONE utility uses steps shown in Figure 7-1 to duplicate each member of a unit.



CXO5510A

Figure 7-1. Steps the CLONE utility follows for duplicating unit members

Use the following steps to clone a single-disk unit, stripeset, or mirrorset:

1. Establish a connection to the controller that accesses the unit you want to clone.
2. Start CLONE using the following command:
RUN CLONE
3. When prompted, enter the unit number of the unit you want to clone.
4. When prompted, enter a unit number for the clone unit that CLONE will create.
5. When prompted, indicate how you would like the clone unit to be brought online:
either automatically or only after approval.
6. When prompted, enter the disk drives you want to use for the clone units.

7. Back up the clone unit.

The following example shows the commands you would use to clone storage unit D98. The clone command terminates after it creates storage unit D99, a clone or copy of D98.

```
RUN CLONE
CLONE LOCAL PROGRAM INVOKED
UNITS AVAILABLE FOR CLONING:
    98
ENTER UNIT TO CLONE? 98
CLONE WILL CREATE A NEW UNIT WHICH IS A COPY OF UNIT 98.
ENTER THE UNIT NUMBER WHICH YOU WANT ASSIGNED TO THE NEW UNIT? 99
THE NEW UNIT MAY BE ADDED USING ONE OF THE FOLLOWING METHODS:
1. CLONE WILL PAUSE AFTER ALL MEMBERS HAVE BEEN COPIED. THE USER MUST THEN
   PRESS RETURN TO CAUSE THE NEW UNIT TO BE ADDED.
2. AFTER ALL MEMBERS HAVE BEEN COPIED, THE UNIT WILL BE ADDED AUTOMATICALLY.
UNDER WHICH ABOVE METHOD SHOULD THE NEW UNIT BE ADDED[]?1
DEVICES AVAILABLE FOR CLONE TARGETS:
DISK20200 (SIZE=832317)
DISK20300 (SIZE=832317)
USE AVAILABLE DEVICE DISK20200(SIZE=832317) FOR MEMBER DISK10300(SIZE=832317)
(Y,N) [Y]? Y
MIRROR DISK10300 C_MA
SET C_MA NOPOLICY
SET C_MA MEMBERS=2
SET C_MA REPLACE=DISK20200
DEVICES AVAILABLE FOR CLONE TARGETS:
DISK20300 (SIZE=832317)
```

```
USE AVAILABLE DEVICE DISK20300(SIZE=832317) FOR MEMBER DISK10000(SIZE=832317)
(Y,N) [Y]? Y
MIRROR DISK10000 C_MB
SET C_MB NOPOLICY
SET C_MB MEMBERS=2
SET C_MB REPLACE=DISK20300
COPY IN PROGRESS FOR EACH NEW MEMBER. PLEASE BE PATIENT...
.
.
COPY FROM DISK10300 TO DISK20200 IS 100% COMPLETE
COPY FROM DISK10000 TO DISK20300 IS 100% COMPLETE
PRESS RETURN WHEN YOU WANT THE NEW UNIT TO BE CREATED
REDUCE DISK20200 DISK20300
UNMIRROR DISK10300
UNMIRROR DISK10000
ADD MIRRORSET C_MA DISK20200
ADD MIRRORSET C_MB DISK20300
ADD STRIPESET C_ST1 C_MA C_MB
INIT C_ST1 NODESTROY
ADD UNIT D99 C_ST1
D99 HAS BEEN CREATED. IT IS A CLONE OF D98.
CLONE - NORMAL TERMINATION
```

Moving Stagesets

You can move a stageset from one subsystem to another without destroying its data. You also can follow the steps in this section to move a stageset to a new location within the same subsystem.



CAUTION: Move only normal stagesets. Do not move stagesets that are reconstructing or reduced, or data corruption will result.

See the release notes for the version of the controller's software for information on which drives can be supported.



CAUTION: Never initialize any container or this procedure will not protect data.

Use the following procedure to move a storageset, while maintaining the data the storageset contains:

1. Show the details for the storageset you want to move. Use the following command:

```
SHOW STORAGESET-NAME
```

2. Label each member with its name and PTL location.

If you do not have a storageset map for the subsystem, you can enter the LOCATE command for each member to find its PTL location. Use the following command:

```
LOCATE DISK-NAME
```

To cancel the locate command, enter the following:

```
LOCATE CANCEL
```

3. Delete the unit number shown in the “Used by” column of the SHOW *storageset-name* command. Use the following syntax:

```
DELETE UNIT-NUMBER
```

4. Delete the storageset shown in the “Name” column of the SHOW *storageset-name* command. Use the following syntax:

```
DELETE STORAGESET-NAME
```

5. Delete each disk drive, one at a time, that the storageset contained. Use the following syntax:

```
DELETE DISK-NAME
```

```
DELETE DISK-NAME
```

```
DELETE DISK-NAME
```

6. Remove the disk drives and move them to their new PTL locations.

7. Again add each disk drive to the controller's list of valid devices. Use the following syntax:

```
ADD DISK DISK-NAME PTL-LOCATION
```

```
ADD DISK DISK-NAME PTL-LOCATION
```

```
ADD DISK DISK-NAME PTL-LOCATION
```

8. Recreate the storageSet by adding its name to the controller's list of valid storageSets and by specifying the disk drives it contains. (Although you have to recreate the storageSet from its original disks, you do not have to add the storageSets in their original order.) Use the following syntax to recreate the storageSet:

ADD STORAGESET-NAME DISK-NAME DISK-NAME

9. Represent the storageSet to the host by giving it a unit number the host can recognize. You can use the original unit number or create a new one. Use the following syntax:

ADD UNIT UNIT-NUMBER STORAGESET-NAME

For example:

The following example moves unit D100 to another cabinet. D100 is the RAIDset RAID99 that consists of members DISK10000, DISK20000, and DISK10100.

DELETE D100

DELETE RAID99

DELETE DISK10000

DELETE DISK10100

DELETE DISK20000

DELETE DISK20100

ADD DISK DISK10000

ADD DISK DISK10100

ADD DISK DISK20000

ADD DISK DISK20100

ADD RAIDSET RAID99 DISK10000 DISK10100 DISK20000 DISK20100

ADD UNIT D100 RAID99

Appendix **A**

Subsystem Profile Templates

This appendix contains storageset profiles you can copy and use to create system profiles. It also contains an enclosure template to use to help keep track of the location of devices and storagesets in your shelves. Four (4) templates will be needed for the subsystem.

NOTE: The storage map templates for the Model 4310R and Model 4214R or 4314R reflect the disk enclosures physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.

- “Storageset Profile,” page A-2
- “Storage Map Template 1 for the Model 4214R Disk Enclosure,” page A-4
 - “Storage Map Template 1 for the Model 4214R Disk Enclosure,” page A-4
 - “Storage Map Template 2 for the Model 4254 Disk Enclosure,” page A-5
 - “Storage Map Template 3 for two Model 4214R Disk Enclosures,” page A-6

Storageset Profile

Type of Storageset:

<input type="checkbox"/> Mirrorset	<input checked="" type="checkbox"/> RAIDset	<input type="checkbox"/> Stripeset	<input type="checkbox"/> Striped	<input type="checkbox"/> JBOD
Mirrorset				

Storageset Name

Disk Drives

Unit Number

Partitions:

Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #	Unit #

RAIDset Switches:

Reconstruction Policy	Reduced Membership	Replacement Policy
<input type="checkbox"/> Normal (default)	<input type="checkbox"/> No (default)	<input type="checkbox"/> Best performance (default)
<input type="checkbox"/> Fast	<input type="checkbox"/> Yes, missing:	<input type="checkbox"/> Best fit
		<input type="checkbox"/> None

Mirrorset Switches:

Replacement Policy	Copy Policy	Read Source
<input type="checkbox"/> Best performance (default)	<input type="checkbox"/> Normal (default)	<input type="checkbox"/> Least busy (default)
<input type="checkbox"/> Best fit	<input type="checkbox"/> Fast	<input type="checkbox"/> Round robin
<input type="checkbox"/> None		<input type="checkbox"/> Disk drive:

Initialize Switches:

Chunk size	Save Configuration	Metadata
<input type="checkbox"/> Automatic (default)	<input type="checkbox"/> No (default)	<input type="checkbox"/> Destroy (default)
<input type="checkbox"/> 64 blocks	<input type="checkbox"/> Yes	<input type="checkbox"/> Retain
<input type="checkbox"/> 128 blocks		
<input type="checkbox"/> 256 blocks		
<input type="checkbox"/> Other:		

Unit Switches:

Caching	Access by following hosts enabled
Read caching_____	_____
Read-ahead caching_____	-
Write-back caching_____	_____
Write-through caching_____	-
	-

	-

Storage Map Template 1 for the Model 4214R Disk Enclosure

Use this template for an MA6000 subsystem with a single Model 4214R disk enclosure (single-bus).

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

Storage Map Template 2 for the Model 4254 Disk Enclosure

Use this template for an MA6000 subsystem with a single Model 4254 disk enclosure (dual-bus).

Model 4254 Disk Enclosure Shelf 1 (dual-bus)														
Bus A								Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800

Storage Map Template 3 for two Model 4214R Disk Enclosures

Use this template for an MA6000 subsystem with two Model 4214R disk enclosures (single-bus)

NOTE: Disk enclosure shelf 1 is a single-bus shelf connected to device port 1 and disk enclosure shelf 2 is a single-bus shelf connected to device port 2.

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	Disk11400	Disk11500

Model 4214R Disk Enclosure Shelf 2 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	Disk21400	Disk21500

Appendix B

Installing, Configuring, and Removing the Client

The following information is included in this appendix:

- Why Install the Client?
- Before You Install the Client
- Installing the Client
- Troubleshooting the Client Installation
- Adding the Storage Subsystem and its Host to the Navigation Tree
- Removing the Command Console Client
- Where to Find Additional Information

Why Install the Client?

Client enables you to monitor and manage your storage subsystem by performing the following tasks:

- Create mirrored device group (RAID 1)
- Create striped device group (RAID 0)
- Create striped mirrored device group (RAID 0+1)
- Create striped parity device group (3/5)
- Create an individual device (JBOD)
- Monitor many subsystems at once
- Set up pager notification

Before You Install the Client

1. Verify that you are logged into an account that is a member of the administrator group.
2. Check the software product description that came with the software for a list of supported hardware.
3. Verify that you have the SNMP service installed on the computer. You must have SNMP installed on the computer for this software to work properly. The Client software uses SNMP to receive traps from the Agent. The SNMP service is available on the Windows NT or Windows 2000 installation CD-ROM. To verify that you have the SNMP service:
 - For Windows NT, double click Services in *Start > Settings > Control Panel*. The entry for SNMP is shown in this window. If you install the SNMP service and you already have a Windows NT service pack on the computer, reinstall the service pack after installing the SNMP service.
 - For Windows 2000, click *Start > Settings > Control Panel > Administrative Tools > Component Services*. The entry for SNMP is shown in the Component Services window.
4. Read the release notes.
5. Read "Troubleshooting the Client Installation" in this appendix.
6. If you have the Command Console Client open, exit the Command Console Client.
7. If you have Command Console Client version 1.1b or earlier, remove the program by using the Windows Add/Remove Programs utility.
8. If you have a previous version of Command Console, you can save the Navigation Tree configuration by copying the *SWCC2.MDB* file to another directory. After you have installed the product, move *SWCC2.MDB* to the directory to which you installed SWCC.
9. Install the HS-Series Agent. For more information, see Chapter 4.

Installing the Client

1. Insert the CD-ROM into a computer running Windows 2000 with Service Pack 1 and Windows NT 4.0 (Intel) with Service Pack 6A.
2. Using Microsoft Windows Explorer, navigate to:
`\swcc\client\intel`
Double click *SETUP.EXE*

3. Select HSG60 Controller and click **Next**.

NOTE: If the computer does not find a previous installation, it will install the SWCC Navigation Window and the CLI Window.

4. Follow the instructions on the screen. After you install the software, the Asynchronous Event Service (AES) starts. AES is a service that runs in the background. It collects and passes traps from the subsystems to the Navigation Tree and to individual pagers (for example, to show that a disk has failed). AES needs to be running for the client system to receive updates.

NOTE: For more information on AES, see *Compaq StorageWorks Command Console Version 2.4 User Guide*.

Troubleshooting the Client Installation

This section provides the following information on how to resolve some of the problems that may appear during installation of the Client software:

- Invalid Network Port Assignments During Installation
- “There is no disk in the drive” Message

Invalid Network Port Assignments During Installation

SWCC Clients and Agents communicate by using sockets. The SWCC installation attempts to add entries into each system list of services (services file or for UCX, the local services database). If the SWCC installation finds an entry in the local services file with the same name as the one it wants to add, it assumes the one in the file is correct.

The SWCC installation may display a message, stating that it cannot upgrade the services file. This happens if it finds an entry in the local services file with the same number as the one it wants to add, but with a different name. In that case, appropriate port numbers must be obtained for the network and added manually to the services file.

There are two default port numbers, one for Command Console (4998) and the other for the device-specific Agent and Client software, such as the Fibre Channel Interconnect Client and Agent (4989). There are two exceptions. The following software has two default port numbers:

- The KZPCC Agent and Client, (4991 and 4985)
- The RA200 Agent and Client, (4997 and 4995)

If the Network Information Services (NIS) are being used to provide named port lookup services, contact the network administrator to add the correct ports.

The following shows how the network port assignments appear in the services file:

spgui	4998/tcp	#Command Console
ccdevmgt	4993/tcp	#Device Management Client and Agent
kzpccconnectport	4991/tcp	#KZPCC Client and Agent
kzpccdcoveryport	4985/tcp	#KZPCC Client and Agent
ccfabric	4989/tcp	#Fibre Channel Interconnect Agent
spagent	4999/tcp	#HS-Series Client and Agent
spagent3	4994/tcp	#HSZ22 Client and Agent
ccagent	4997/tcp	#RA200 Client and Agent
spagent2	4995/tcp	#RA200 Client and Agent

“There is no disk in the drive” Message

When you install the Command Console Client, the software checks the shortcuts on the desktop and in the Start menu. The installation will check the shortcuts of all users for that computer, even if they are not currently logged on. You may receive an error message if any of these shortcuts point to empty floppy drives, empty CD-ROM drives, or missing removable disks. Do one of the following:

- Ignore the error message by clicking **Ignore**.
- Replace the removable disks, and place a disk in the floppy drive and a CD-ROM in the CD-ROM drive. Then, click **Retry**.

Adding the Storage Subsystem and its Host to the Navigation Tree

The Navigation Tree allows you to manage storage over the network by using the Storage Window. If planning to use pager notification, add the storage subsystem to the Navigation Tree.

1. Verify that you have properly installed and configured the HS-Series Agent on the storage subsystem host.
2. Click *Start > Programs > Command Console > StorageWorks Command Console*. Client displays the Navigation Window. The Navigation Window lets you monitor and manage many storage subsystems over the network.



Figure B-1. Navigation window

3. Click *File > Add System*. The Add System window appears.
4. Type the host name or its TCP/IP address and click *Apply*.
5. Click *Close*.



Figure B-2. Navigation window showing storage host system "Atlanta"

6. Click the plus sign to expand the host icon. When expanded, the Navigation Window displays an icon for the storage subsystem. To access the Storage Window for the subsystem, double-click the Storage Window icon.



Figure B-3. Navigation window showing expanded "Atlanta" host icon

NOTE: You can create virtual disks by using the Storage Window. For more information on the Storage Window, refer to *Compaq StorageWorks Command Console Version 2.4 User Guide*.

Removing the Command Console Client

Before you remove the Command Console Client, remove AES. This will prevent the system from reporting that a service failed to start every time the system is started. Steps 2 through 5 describe how to remove the Command Console Client.

NOTE: When removing the Command Console Client, the SWCC2.MDB file is deleted. This file contains the Navigation Tree configuration. If you want to save this information, move the file to another directory.

1. Click *Start > Programs > Command Prompt* and change to the directory to which you installed the Command Console Client.
2. Enter the following command:

```
C:\Program Files\Compaq\SWCC> AsyncEventService -remove
```

3. Do one of the following:
 - On Windows NT 4.0, click *Start > Settings > Control Panel*, and then double click the Add/Remove Programs icon in the Control Panel. The Add/Remove Program Properties window appears.
 - On Windows 2000, click *Start > Settings > Control Panel > Add/Remove Programs*. The Add/Remove Program window appears.
4. Select Command Console in the window.
5. Do one of the following:
 - On Windows NT 4.0, click *Add/Remove*.
 - On Windows 2000, click *Change/Remove*.
6. Follow the instructions on the screen.

NOTE: This procedure removes only the Command Console Client (SWCC Navigation Window). You can remove the HSG60 Client by using the Add/Remove program.

Where to Find Additional Information

You can find additional information about SWCC by referring to the online Help and to *Compaq StorageWorks Command Console Version 2.4 User Guide*.

About the User Guide

Compaq StorageWorks Command Console Version 2.4 User Guide contains additional information on how to use SWCC. Some of the topics in the user guide are the following:

- About AES
- Adding Devices
- Adding Virtual Disks
- Setting Up Pager Notification
- How to Integrate SWCC with Compaq Insight Manager
- Troubleshooting Information

About the Online Help

Most of the information about the Client is provided in the online Help. Online Help is provided in two places:

- **Navigation Window** – Online Help provides information on pager notification and a tour of the Command Console Client, in addition to information on how to add a system to the Navigation Tree.
- **Storage Window** – Online Help provides detailed information about the Storage Window, such as how to create virtual disks.

Glossary

This glossary defines terms pertaining to the ACS solution software. It is not a comprehensive glossary of computer terms.

adapter	A device that converts the protocol and hardware interface of one bus type into another without changing the function of the bus.
ACS	<i>See</i> array controller software.
AL_PA	<i>See</i> arbitrated loop physical address.
ANSI	Pronounced “ann-see.” Acronym for the American National Standards Institute. An organization who develops standards used voluntarily by many manufacturers within the USA. ANSI is not a government agency.
arbitrated loop physical address	Abbreviated AL_PA. A one-byte value used to identify a port in an Arbitrated Loop topology.
array controller	<i>See</i> controller.
array controller software	Abbreviated ACS. Software contained on a removable ROM program card that provides the operating system for the array controller.
asynchronous	Pertaining to events that are scheduled as the result of a signal asking for the event; pertaining to that which is without any specified time relation. <i>See also</i> synchronous.
autospare	A controller feature that automatically replaces a failed disk drive. To aid the controller in automatically replacing failed disk drives, you can enable the AUTOSPARE switch for the failedset causing physically replaced disk drives to be automatically placed into the spareset. Also called “autonewspare.”
bad block	A data block that contains a physical defect.

Glossary–2 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

bad block replacement	Abbreviated BBR. A replacement routine that substitutes defect-free disk blocks for those found to have defects. This process takes place in the controller, transparent to the host.
backplane	The electronic printed circuit board into which you plug subsystem devices—for example, the SBB or power supply.
BBR	<i>See</i> bad block replacement.
BIST	<i>See</i> built-in self-test.
bit	A single binary digit having a value of either 0 or 1. A bit is the smallest unit of data a computer can process.
block	Also called a sector. The smallest collection of consecutive bytes addressable on a disk drive. In integrated storage elements, a block contains 512 bytes of data, error codes, flags, and the block address header.
bootstrapping	A method used to bring a system or device into a defined state by means of its own action. For example, a machine routine whose first few instructions are enough to bring the rest of the routine into the computer from an input device.
built-in self-test	A diagnostic test performed by the array controller software on the controller policy processor.
byte	A binary character string made up of 8 bits operated on as a unit.
cache memory	A portion of memory used to accelerate read and write operations.
CDU	Cable distribution unit. The power entry device for StorageWorks cabinets. The CDU provides the connections necessary to distribute power to the cabinet shelves and fans.
channel	An interface that allows high speed transfer of large amounts of data. Another term for a SCSI bus. <i>See also</i> SCSI.
chunk	A block of data written by the host.
chunk size	The number of data blocks, assigned by a system administrator, written to the primary RAIDset or stripeset member before the remaining data blocks are written to the next RAIDset or stripeset member.
CLCP	An abbreviation for code-load code-patch utility. This utility is used to upgrade the controller and EMU software. It can also be used to patch the controller software.

coax	A two-conductor wire in which one conductor completely wraps the other with the two separated by insulation.
cold swap	A method of device replacement that requires the entire subsystem to be turned off before the device can be replaced. <i>See also</i> hot swap and warm swap.
command line interpreter	(CLI) The configuration interface to operate the controller software.
concat commands	Concat commands implement storageset expansion features.
configuration file	A file that contains a representation of a storage subsystem configuration.
container	1) Any entity that is capable of storing data, whether it is a physical device or a group of physical devices. (2) A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers the controller uses to create units.
controller	A hardware device that, with proprietary software, facilitates communications between a host and one or more devices organized in an array. The HSG60 family controllers are examples of array controllers.
copying	A state in which data to be copied to the mirrorset is inconsistent with other members of the mirrorset. <i>See also</i> normalizing.
copying member	Any member that joins the mirrorset after the mirrorset is created is regarded as a copying member. Once all the data from the normal member (or members) is copied to a normalizing or copying member, the copying member then becomes a normal member. <i>See also</i> normalizing member.
CSR	An acronym for control and status register.
DAEMON	Pronounced “demon.” A program usually associated with a UNIX systems that performs a utility (housekeeping or maintenance) function without being requested or even known of by the user. A daemon is a diagnostic and execution monitor.
data center cabinet	A generic reference to large Compaq subsystem cabinets, such as the cabinets in which StorageWorks components can be mounted.

data striping	The process of segmenting logically sequential data, such as a single file, so that segments can be written to multiple physical devices (usually disk drives) in a round-robin fashion. This technique is useful if the processor is capable of reading or writing data faster than a single disk can supply or accept the data. While data is being transferred from the first disk, the second disk can locate the next segment.
device	<i>See</i> node and peripheral device.
differential I/O module	A 16-bit I/O module with SCSI bus converter circuitry for extending a differential SCSI bus. <i>See also</i> I/O module.
differential SCSI bus	A bus in which a signal level is determined by the potential difference between two wires. A differential bus is more robust and less subject to electrical noise than is a single-ended bus.
DIMM	Dual inline Memory Module.
dirty data	The write-back cached data that has not been written to storage media, even though the host operation processing the data has completed.
DMA	Direct Memory Access.
DOC	DWZZA-On-a-Chip. ASCSI bus extender chip used to connect a SCSI bus in an expansion cabinet to the corresponding SCSI bus in another cabinet.
driver	A hardware device or a program that controls or regulates another device. For example, a device driver is a driver developed for a specific device that allows a computer to operate with the device, such as a printer or a disk drive.
dual-redundant configuration	A controller configuration consisting of two active controllers operating as a single controller. If one controller fails, the other controller assumes control of the failing controller devices.
dual-simplex	A communications protocol that allows simultaneous transmission in both directions in a link, usually with no flow control.
DUART	Dual universal asynchronous receiver and transmitter. An integrated circuit containing two serial, asynchronous transceiver circuits.
ECB	External cache battery. The unit that supplies backup power to the cache module in the event the primary power source fails or is interrupted.
ECC	Error checking and correction.
EDC	Error detection code.

EIA	The abbreviation for Electronic Industries Association. EIA is a standards organization specializing in the electrical and functional characteristics of interface equipment.
EMU	Environmental monitoring unit. A unit that provides increased protection against catastrophic failures. Some subsystem enclosures include an EMU which works with the controller to detect conditions such as failed power supplies, failed blowers, elevated temperatures, and external air sense faults. The EMU also controls certain cabinet hardware including DOC chips, alarms, and fan speeds.
ESD	Electrostatic discharge. The discharge of potentially harmful static electrical voltage as a result of improper grounding.
extended subsystem	A subsystem in which two cabinets are connected to the primary cabinet.
external cache battery	<i>See</i> ECB.
F_Port	A port in a fabric where an N_Port or NL_Port may attach.
fabric	A group of interconnections between ports that includes a fabric element.
failedset	A group of failed mirrorset or RAIDset devices automatically created by the controller.
failover	The process that takes place when one controller in a dual-redundant configuration assumes the workload of a failed companion controller. Failover continues until the failed controller is repaired or replaced.
FC-AL	The Fibre Channel Arbitrated Loop standard. <i>See Fibre Channel.</i>
FC-ATM	ATM AAL5 over Fibre Channel
FC-FG	Fibre Channel Fabric Generic Requirements
FG-FP	Fibre Channel Framing Protocol (HIPPI on FC)
FC-GS-1	Fibre Channel Generic Services-1
FC-GS-2	Fibre Channel Generic Services-2
FC-IG	Fibre Channel Implementation Guide
FC-LE	Fibre Channel Link Encapsulation (ISO 8802.2)

FC–PH	The Fibre Channel Physical and Signaling standard.
FC–SB	Fibre Channel Single Byte Command Code Set
FC–SW	Fibre Channel Switched Topology and Switch Controls
FCC	Federal Communications Commission. The federal agency responsible for establishing standards and approving electronic devices within the United States.
FCC Class A	This certification label appears on electronic devices that can only be used in a commercial environment within the United States.
FCC Class B	This certification label appears on electronic devices that can be used in either a home or a commercial environment within the United States.
FCP	The mapping of SCSI-3 operations to Fibre Channel.
FDDI	Fiber Distributed Data Interface. An ANSI standard for 100 megabaud transmission over fiber optic cable.
FD SCSI	The fast, narrow, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s. <i>See also</i> FWD SCSI and SCSI.
fiber	A fiber or optical strand. Spelled <i>fib</i> re in Fibre Channel.
fiber optic cable	A transmission medium designed to transmit digital signals in the form of pulses of light. Fiber optic cable is noted for its properties of electrical isolation and resistance to electrostatic contamination.
Fibre Channel	A high speed, high-bandwidth serial protocol for channels and networks that interconnect over twisted pair wires, coaxial cable or fiber optic cable. The Fibre Channel Switched (FC-SW) (fabric) offers up to 16 million ports with cable lengths of up to 10 kilometers. The Fibre Channel Arbitrated Loop (FC-AL) topology offers speeds of up to 100 Mbytes/seconds and up to 127 nodes, all connected in serial. In contrast to SCSI technology, Fibre Channel does not require ID switches or terminators. The FC-AL loop may be connected to a Fibre Channel fabric for connection to other nodes.
fib re channel topology	An interconnection scheme that allows multiple Fibre Channel ports to communicate with each other. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.
FL_Port	A port in a fabric where N_Port or an NL_Port may be connected.
flush	The act of writing dirty data from cache to a storage media.

FMU	Fault management utility.
forced errors	A data bit indicating a corresponding logical data block contains unrecoverable data.
frame	An invisible unit used to transfer information in Fibre Channel.
FRU	Field replaceable unit. A hardware component that can be replaced at the customer location by Compaq service personnel or qualified customer service personnel.
FRUTIL	Field Replacement utility.
full duplex (n)	A communications system in which there is a capability for 2-way transmission and acceptance between two sites at the same time.
full duplex (adj)	Pertaining to a communications method in which data can be transmitted and received at the same time.
FWD SCSI	A fast, wide, differential SCSI bus with a maximum 16-bit data transfer rate of 20 MB/s. <i>See also</i> SCSI and FD SCSI.
GBIC	Gigabit Interface Converter. GBICs convert electrical signals to optical signals (and vice-versa.) They are inserted into the ports of the Fibre Channel switch and hold the Fibre Channel cables.
GLM	Gigabit link module
giga	A prefix indicating a billion (10^9) units, as in gigabaud or gigabyte.
gigabaud	An encoded bit transmission rate of one billion (10^9) bits per second.
gigabyte	A value normally associated with a disk drives storage capacity, meaning a billion (10^9) bytes. The decimal value 1024 is usually used for one thousand.
half-duplex (adj)	Pertaining to a communications system in which data can be either transmitted or received but only in one direction at one time.
hard address	The AL_PA which an NL_Port attempts to acquire during loop initialization.
heterogeneous host support	Also called <i>noncooperating host support</i> .
HIPPI-FC	Fibre Channel over HIPPI
host	The primary or controlling computer to which a storage subsystem is attached.

host adapter	A device that connects a host system to a SCSI bus. The host adapter usually performs the lowest layers of the SCSI protocol. This function may be logically and physically integrated into the host system.
hot disks	A disk containing multiple hot spots. Hot disks occur when the workload is poorly distributed across storage devices which prevents optimum subsystem performance. <i>See also</i> hot spots.
hot spots	A portion of a disk drive frequently accessed by the host. Because the data being accessed is concentrated in one area, rather than spread across an array of disks providing parallel access, I/O performance is significantly reduced. <i>See also</i> hot disks.
hot swap	A method of device replacement that allows normal I/O activity on a device bus to remain active during device removal and insertion. The device being removed or inserted is the only device that cannot perform operations during this process. <i>See also</i> cold swap and warm swap.
hub	A device (concentrator) which performs some or all of the following functions: <ul style="list-style-type: none">■ Automatic insertion of operational loop devices without disrupting the existing configuration.■ Automatic removal of failed loop devices without impacting the existing configuration.■ Provides a centralized (star) wiring configuration and maintenance point.■ Provides central monitoring and management.
IBR	Initial Boot Record.
ILF	Illegal function.
INIT	Initialize input and output.
initiator	A SCSI device that requests an I/O process to be performed by another SCSI device, namely, the SCSI target. The controller is the initiator on the device bus. The host is the initiator on the host bus.
instance code	A four-byte value displayed in most text error messages and issued by the controller when a subsystem error occurs. The instance code indicates when during software processing the error was detected.
interface	A set of protocols used between components, such as cables, connectors, and signal levels.

I/O	Refers to input and output functions.
I/O driver	The set of code in the kernel that handles the physical I/O to a device. This is implemented as a fork process. Same as driver.
I/O interface	<i>See</i> interface.
I/O module	A 16-bit SBB shelf device that integrates the SBB shelf with either an 8-bit single ended, 16-bit single-ended, or 16-bit differential SCSI bus.
I/O operation	The process of requesting a transfer of data from a peripheral device to memory (or visa versa), the actual transfer of the data, and the processing and overlaying activity to make both of those happen.
IPI	Intelligent Peripheral Interface. An ANSI standard for controlling peripheral devices by a host computer.
IPI-3 Disk	Intelligent Peripheral Interface Level 3 for Disk
IPI-3 Tape	Intelligent Peripheral Interface Level 3 for Tape
JBOD	Just a bunch of disks. A term used to describe a group of single-device logical units.
kernel	The most privileged processor access mode.
LBN	Logical Block Number.
L_port	A node or fabric port capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
LED	Light Emitting Diode.
link	A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.
local connection	A connection to the subsystem using either its serial maintenance port or the host SCSI bus. A local connection enables you to connect to one subsystem controller within the physical range of the serial or host SCSI cable.
local terminal	A terminal plugged into the EIA-423 maintenance port located on the front bezel of the controller. <i>See also</i> maintenance terminal.
logical bus	A single-ended bus connected to a differential bus by a SCSI bus signal converter.

logical unit	A physical or virtual device addressable through a target ID number. LUNs use their target bus connection to communicate on the SCSI bus.
logical unit number	A value that identifies a specific logical unit belonging to a SCSI target ID number. A number associated with a physical device unit during a task I/O operations. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.
logon	Also called login. A procedure whereby a participant, either a person or network connection, is identified as being an authorized network participant.
loop	<i>See</i> arbitrated loop.
loop_ID	A seven-bit value numbered contiguously from zero to 126-decimal and represent the 127 legal AL_PA values on a loop (not all of the 256 hex values are allowed as AL_PA values per FC-AL.)
loop tenancy	The period of time between the following events: when a port wins loop arbitration and when the port returns to a monitoring state.
L_Port	A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
LRU	Least recently used. A cache term used to describe the block replacement policy for read cache.
Mbps	Approximately one million (10^6) bits per second—that is, megabits per second.
MBps	Approximately one million (10^6) bytes per second—that is, megabytes per second.
maintenance terminal	An EIA-423-compatible terminal used with the controller. This terminal is used to identify the controller, enable host paths, enter configuration information, and check the controller status. The maintenance terminal is not required for normal operations. <i>See also</i> local terminal.
member	A container that is a storage element in a RAID array.

metadata	The data written to a disk for the purposes of controller administration. Metadata improves error detection and media defect management for the disk drive. It is also used to support storage set configuration and partitioning. Nontransportable disks also contain metadata to indicate they are uniquely configured for StorageWorks environments. Metadata can be thought of as “data about data.”
mirroring	The act of creating an exact copy or image of data.
mirrored write-back caching	A method of caching data that maintains two copies of the cached data. The copy is available if either cache module fails.
mirrorset	<i>See</i> RAID level 1.
MIST	Module Integrity Self-Test.
N_port	A port attached to a node for use with point-to-point topology or fabric topology.
NL_port	A port attached to a node for use in all topologies.
network	In data communication, a configuration in which two or more terminals or devices are connected to enable information transfer.
node	In data communications, the point at which one or more functional units connect transmission lines.
Non-L_Port	A Node of Fabric port that is not capable of performing the Arbitrated Loop functions and protocols. N_Ports and F_Ports loop-capable ports.
non-participating mode	A mode within an L_Port that inhibits the port from participating in loop activities. L_Ports in this mode continue to retransmit received transmission words but are not permitted to arbitrate or originate frames. An L_Port in non-participating mode may or may not have an AL_PA. <i>See also</i> participating mode.
nominal membership	The desired number of mirrorset members when the mirrorset is fully populated with active devices. If a member is removed from a mirrorset, the actual number of members may fall below the “nominal” membership.
node	In data communications, the point at which one or more functional units connect transmission lines. In Fibre Channel, a device that has at least one N_Port or NL_Port.

nonredundant controller configuration	(1) A single controller configuration. (2) A controller configuration that does not include a second controller.
normal member	A mirrorset member that, block-for-block, contains the same data as other normal members within the mirrorset. Read requests from the host are always satisfied by normal members.
normalizing	Normalizing is a state in which, block-for-block, data written by the host to a mirrorset member is consistent with the data on other normal and normalizing members. The normalizing state exists only after a mirrorset is initialized. Therefore, no customer data is on the mirrorset.
normalizing member	A mirrorset member whose contents are the same as all other normal and normalizing members for data that has been written since the mirrorset was created or lost cache data was cleared. A normalizing member is created by a normal member when either all of the normal members fail or all of the normal members are removed from the mirrorset. <i>See also</i> copying member.
NVM	Non-Volatile Memory. A type of memory where the contents survive power loss. Also sometimes referred to as NVMEM.
OCP	Operator control panel. The control or indicator panel associated with a device. The OCP is usually mounted on the device and is accessible to the operator.
other controller	The controller in a dual-redundant pair that is connected to the controller serving the current CLI session. <i>See also</i> this controller.
outbound fiber	One fiber in a link that carries information away from a port.
parallel data transmission	A data communication technique in which more than one code element (for example, bit) of each byte is sent or received simultaneously.
parity	A method of checking if binary numbers or characters are correct by counting the ONE bits. In odd parity, the total number of ONE bits must be odd; in even parity, the total number of ONE bits must be even.
parity bit	A binary digit added to a group of bits that checks to see if errors exist in the transmission.
parity check	A method of detecting errors when data is sent over a communications line. With even parity, the number of ones in a set of binary data should be even. With odd parity, the number of ones should be odd.

participating mode	A mode within an L_Port that allows the port to participate in loop activities. A port must have a valid AL_PA to be in participating mode.
PCM	Polycenter Console Manager.
PCMCIA	Personal Computer Memory Card Industry Association. An international association formed to promote a common standard for PC card-based peripherals to be plugged into notebook computers. The card commonly known as a PCMCIA card is about the size of a credit card.
parity	A method of checking if binary numbers or characters are correct by counting the ONE bits. In odd parity, the total number of ONE bits must be odd; in even parity, the total number of ONE bits must be even. Parity information can be used to correct corrupted data. RAIDsets use parity to improve the availability of data.
parity bit	A binary digit added to a group of bits that checks to see if there are errors in the transmission.
parity RAID	<i>See</i> RAIDset.
partition	A logical division of a container, represented to the host as a logical unit.
peripheral device	Any unit, distinct from the CPU and physical memory, that can provide the system with input or accept any output from it. Terminals, printers, tape drives, and disks are peripheral devices.
point-to-point connection	A network configuration in which a connection is established between two, and only two, terminal installations. The connection may include switching facilities.
port	(1) In general terms, a logical channel in a communications system. (2) The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus. Regarding the controller, the port is (1) the logical route for data in and out of a controller that can contain one or more channels, all of which contain the same type of data. (2) The hardware and software that connects a controller to a SCSI device.
port_name	A 64-bit unique identifier assigned to each Fibre Channel port. The Port_Name is communicated during the logon and port discovery process.
preferred address	The AL_PA which an NL_Port attempts to acquire first during initialization.

Glossary–14 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

primary cabinet	The primary cabinet is the subsystem enclosure that contains the controllers, cache modules, external cache batteries, and the PVA module.
private NL_Port	An NL_Port which does not attempt login with the fabric and only communicates with NL_Ports on the same loop.
program card	The PCMCIA card containing the controller operating software.
protocol	The conventions or rules for the format and timing of messages sent and received.
PTL	Port-Target-LUN. The controller method of locating a device on the controller device bus.
PVA module	Power Verification and Addressing module.
quiesce	The act of rendering bus activity inactive or dormant. For example, “quiesce the SCSI bus operations during a device warm-swap.”
RAID	Redundant Array of Independent Disks. Represents multiple levels of storage access developed to improve performance or availability or both.
RAID level 0	A RAID storage set that stripes data across an array of disk drives. A single logical disk spans multiple physical disks, enabling parallel data processing for increased I/O performance. While the performance characteristics of RAID level 0 is excellent, this RAID level is the only one that does not provide redundancy. Raid level 0 storage sets are sometimes referred to as stripesets.
RAID level 0+1	A RAID storage set that stripes data across an array of disks (RAID level 0) and mirrors the striped data (RAID level 1) to provide high I/O performance and high availability. This RAID level is alternatively called a striped mirrorset.
RAID level 1	A RAID storage set of two or more physical disks that maintain a complete and independent copy of the entire virtual disk's data. This type of storage set has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storage sets are sometimes referred to as mirrorsets.
RAID level 3	A RAID storage set that transfers data parallel across the array disk drives a byte at a time, causing individual blocks of data to be spread over several disks serving as one enormous virtual disk. A separate redundant check disk for the entire array stores parity on a dedicated disk drive within the storage set. <i>See also</i> RAID level 5.

RAID level 5	A RAID storageset that, unlike RAID level 3, stores the parity information across all of the disk drives within the storageset. <i>See also</i> RAID level 3.
RAID level 3/5	A RAID storageset that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID.
RAIDset	<i>See</i> RAID level 3/5.
RAM	Random access memory.
read ahead caching	A caching technique for improving performance of synchronous sequential reads by prefetching data from disk.
read caching	A cache management method used to decrease the subsystem response time to a read request by allowing the controller to satisfy the request from the cache memory rather than from the disk drives.
reconstruction	The process of regenerating the contents of a failed member data. The reconstruct process writes the data to a spareset disk and incorporates the spareset disk into the mirrorset, striped mirrorset, or RAIDset from which the failed member came. <i>See also</i> regeneration.
reduced	Indicates that a mirrorset or RAIDset is missing one member because the member has failed or has been physically removed.
redundancy	The provision of multiple interchangeable components to perform a single function in order to cope with failures and errors. A RAIDset is considered to be redundant when user data is recorded directly to one member and all of the other members include associated parity information.
regeneration	(1) The process of calculating missing data from redundant data. (2) The process of recreating a portion of the data from a failing or failed drive using the data and parity information from the other members within the storageset. The regeneration of an entire RAIDset member is called reconstruction. <i>See also</i> reconstruction.
request rate	The rate at which requests are arriving at a servicing entity.
RFI	Radio frequency interference. The disturbance of a signal by an unwanted radio signal or frequency.

replacement policy	The policy specified by a switch with the SET FAILEDSET command indicating whether a failed disk from a mirrorset or RAIDset is to be automatically replaced with a disk from the spareset. The two switch choices are AUTOSPARE and NOAUTOSPARE.
SBB	StorageWorks building block. (1) A modular carrier plus the interface required to mount the carrier into a standard StorageWorks shelf. (2) any device conforming to shelf mechanical and electrical standards installed in a 3.5-inch or 5.25-inch carrier, whether it is a storage device or power supply.
SCSI	Small computer system interface. (1) An ANSI interface standard defining the physical and electrical parameters of a parallel I/O bus used to connect initiators to devices. (2) a processor-independent standard protocol for system-level interfacing between a computer and intelligent devices including hard drives, floppy disks, CD-ROMs, printers, scanners, and others.
SCSI-A cable	A 50-conductor (25 twisted-pair) cable generally used for single-ended, SCSI-bus connections.
SCSI bus signal converter	Sometimes referred to as an adapter. (1) A device used to interface between the subsystem and a peripheral device unable to be mounted directly into the SBB shelf of the subsystem. (2) a device used to connect a differential SCSI bus to a single-ended SCSI bus. (3) A device used to extend the length of a differential or single-ended SCSI bus. <i>See also</i> DWZZA, DWZZB, DWZZC, and I/O module.
SCSI device	(1) A host computer adapter, a peripheral controller, or an intelligent peripheral that can be attached to the SCSI bus. (2) Any physical unit that can communicate on a SCSI bus.
SCSI device ID number	A bit-significant representation of the SCSI address referring to one of the signal lines, numbered 0 through 7 for an 8-bit bus, or 0 through 15 for a 16-bit bus. <i>See also</i> target ID number.
SCSI ID number	The representation of the SCSI address that refers to one of the signal lines numbered 0 through 15.
SCSI-P cable	A 68-conductor (34 twisted-pair) cable generally used for differential bus connections.
SCSI port	(1) Software: The channel controlling communications to and from a specific SCSI bus in the system. (2) Hardware: The name of the logical socket at the back of the system unit to which a SCSI device is connected.

serial transmission	A method transmission in which each bit of information is sent sequentially on a single channel rather than simultaneously as in parallel transmission.
signal converter	<i>See</i> SCSI bus signal converter.
single ended I/O module	A 16-bit I/O module. <i>See also</i> I/O module.
single-ended SCSI bus	An electrical connection where one wire carries the signal and another wire or shield is connected to electrical ground. Each signal logic level is determined by the voltage of a single wire in relation to ground. This is in contrast to a differential connection where the second wire carries an inverted signal.
spareset	A collection of disk drives made ready by the controller to replace failed members of a storageset.
storage array	An integrated set of storage devices.
storage array subsystem	<i>See</i> storage subsystem.
storageset	(1) A group of devices configured with RAID techniques to operate as a single container. (2) Any collection of containers, such as stripesets, mirrorsets, striped mirrorsets, and RAIDsets.
storageset expansion	Concat commands implement storageset expansion features.
storage subsystem	The controllers, storage devices, shelves, cables, and power supplies used to form a mass storage subsystem.
storage unit	The general term that refers to storagesets, single-disk units, and all other storage devices that are installed in your subsystem and accessed by the host. A storage unit can be any entity that is capable of storing data, whether it is a physical device or a group of physical devices.
StorageWorks	<p>A family of Compaq modular data storage products that allow customers to design and configure their own storage subsystems. Components include power, packaging, cabling, devices, controllers, and software. Customers can integrate devices and array controllers in StorageWorks enclosures to form storage subsystems.</p> <p>StorageWorks systems include integrated SBBs and array controllers to form storage subsystems. System-level enclosures to house the shelves and standard mounting devices for SBBs are also included.</p>

Glossary–18 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

stripe	The data divided into blocks and written across two or more member disks in an array.
striped mirrorset	<i>See</i> RAID level 0+1.
stripeset	<i>See</i> RAID level 0.
stripe size	The stripe capacity as determined by $n-1$ times the chunksize, where n is the number of RAIDset members.
striping	<p>The technique used to divide data into segments, also called chunks. The segments are striped, or distributed, across members of the stripeset. This technique helps to distribute hot spots across the array of physical devices to prevent hot spots and hot disks.</p> <p>Each stripeset member receives an equal share of the I/O request load, improving performance.</p>
surviving controller	The controller in a dual-redundant configuration pair that serves its companion devices when the companion controller fails.
switch	A method that controls the flow of functions and operations in software.
synchronous	Pertaining to a method of data transmission which allows each event to operate in relation to a timing signal. <i>See also</i> asynchronous.
tape	A storage device supporting sequential access to variable sized data records.
target	(1) A SCSI device that performs an operation requested by an initiator. (2) Designates the target identification (ID) number of the device.
target ID number	The address a bus initiator uses to connect with a bus target. Each bus target is assigned a unique target address.
this controller	The controller that is serving your current CLI session through a local or remote terminal. <i>See also</i> other controller.
tape inline exerciser	(TILX) The controller diagnostic software to test the data transfer capabilities of tape drives in a way that simulates a high level of user activity.
transfer data rate	The speed at which data may be exchanged with the central processor, expressed in thousands of bytes per second.
ULP	Upper Layer Protocol.

ULP process	A function executing within a Fibre Channel node which conforms to the Upper Layer Protocol (ULP) requirements when interacting with other ULP processes.
Ultra SCSI	A Fast-20 SCSI bus. <i>See also</i> Wide Ultra SCSI.
unit	A container made accessible to a host. A unit may be created from a single disk drive or tape drive. A unit may also be created from a more complex container such as a RAIDset. The controller supports a maximum of eight units on each target. <i>See also</i> target and target ID number.
unwritten cached data	Sometimes called unflushed data. <i>See</i> dirty data.
UPS	Uninterruptible power supply. A battery-powered power supply guaranteed to provide power to an electrical device in the event of an unexpected interruption to the primary power supply. Uninterruptible power supplies are usually rated by the amount of voltage supplied and the length of time the voltage is supplied.
VHDCI	Very high-density-cable interface. A 68-pin interface. Required for Ultra-SCSI connections.
virtual terminal	A software path from an operator terminal on the host to the controller's CLI interface, sometimes called a host console. The path can be established via the host port on the controller or via the maintenance port through an intermediary host.
VTDPY	An abbreviation for Virtual Terminal Display Utility.
warm swap	A device replacement method that allows the complete system to remain online during device removal or insertion. The system bus may be halted, or quiesced, for a brief period of time during the warm-swap procedure.
Wide Ultra SCSI	Fast/20 on a Wide SCSI bus.
Worldwide name	A unique 64-bit number assigned to a subsystem by the Institute of Electrical and Electronics Engineers (IEEE) and set by Compaq manufacturing prior to shipping. This name is referred to as the node ID within the CLI.
write-back caching	A cache management method used to decrease the subsystem response time to write requests by allowing the controller to declare the write operation “complete” as soon as the data reaches its cache memory. The controller performs the slower operation of writing the data to the disk drives at a later time.

**write-through
caching**

A cache management method used to decrease the subsystem response time to a read. This method allows the controller to satisfy the request from the cache memory rather than from the disk drives.

write hole

The period of time in a RAID level 1 or RAID level 5 write operation when an opportunity emerges for undetectable RAIDset data corruption. Write holes occur under conditions such as power outages, where the writing of multiple members can be abruptly interrupted. A battery backed-up cache design eliminates the write hole because data is preserved in cache and unsuccessful write operations can be retried.

**write-through
cache**

A cache management technique for retaining host write requests in read cache. When the host requests a write operation, the controller writes data directly to the storage device. This technique allows the controller to complete some read requests from the cache, greatly improving the response time to retrieve data. The operation is complete only after the data to be written is received by the target storage device.

This cache management method may update, invalidate, or delete data from the cache memory accordingly, to ensure that the cache contains the most current data.

Index

A

ADD CONNECTIONS

- transparent failover 1-14

ADD UNIT

- transparent failover 1-14

adding

- virtual disks B-7

adding a disk drive to the spareset

- configuration options 5-21

adding disk drives

- configuration options 5-20

Agent

- functions 4-1

- installing 4-5

- reconfiguring 4-8

- running 4-8

- uninstalling 4-8

arbitrated loop

- naming convention 3-5

array of disk drives 2-7

assigning unit numbers 1-14

assignment

- unit numbers

 - fabric topology 5-19

- unit qualifiers

 - fabric topology 5-19

assignment of unit numbers

- fabric topology

 - partition 5-20

 - single disk 5-19

asynchronous event service B-7

autospare

- enabling

 - fabric topology 5-22

availability 2-15

B

backup

- cloning data 7-2

- subsystem configuration 7-1

C

cabling

- controller pair 5-9

- multiple-bus failover

 - fabric topology configuration 5-9

- single controller 5-4

cache modules

- location 1-2

- read caching 1-8

- write-back caching 1-8

- write-through caching 1-8

caching techniques

- mirrored 1-9

- read caching 1-8

- read-ahead caching 1-8

- write-back caching 1-8

- write-through caching 1-8

changing switches

- configuration options 5-23

chunk size

Index-2 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

- choosing for RAIDsets and stripesets 2-23
 - controlling stripesize 2-23
 - using to increase data transfer rate 2-25
 - using to increase request rate 2-24
 - using to increase write performance 2-26
 - CHUNKSIZE 2-23
 - CLI commands
 - installation verification 5-8
 - CLI prompt
 - changing
 - fabric topology 5-20
 - Client
 - removing B-6
 - uninstalling B-6
 - CLONE utility
 - backup 7-2
 - cloning
 - backup 7-2
 - command console LUN 1-9
 - SCSI-2 mode 1-16
 - Compaq Insight Manager B-7
 - comparison of container types 2-7
 - Configuration
 - fabric topology
 - procedure flowchart
 - illustrated 0-xviii
 - configuration
 - backup 7-1
 - fabric topology
 - devices 5-14
 - multiple-bus failover
 - cabling 5-9
 - multiple-bus failover using CLI 6-5
 - single controller cabling 5-3
 - fibre channel fabric topology
 - using CLI 5-1
 - using SWCC 5-1
 - restoring 2-26
 - rules 2-3
 - configuration options
 - fabric topology
 - adding a disk drive to the spareset 5-21
 - adding disk drives 5-20
 - changing switches
 - device 5-23
 - displaying the current switches 5-23
 - initialize 5-23
 - RAIDset and mirrorset 5-23
 - unit 5-23
 - changing switches for a storageset or device 5-23
 - changing the CLI prompt 5-20
 - deleting a storageset 5-22
 - enabling autospare 5-22
 - removing a disk drive from the spareset 5-21
 - configuring
 - pager notification B-7
 - configuring devices
 - fabric topology 5-14
 - connections 1-11
 - naming 1-11
 - containers
 - attributes 2-6
 - comparison 2-7
 - illustrated 2-7
 - mirrorsets 2-13
 - planning storage 2-6
 - stripesets 2-12
 - controller verification
 - installation 5-8
 - controllers
 - cabling 5-4, 5-9
 - location 1-2
 - node IDs 1-17
 - verification of installation 5-8
 - worldwide names 1-17
 - creating
 - storageset and device profiles 2-8
 - creating disk label 3-6
 - creating filesystems 3-10
- D**
- data transfer rate 2-25
 - DESTROY 2-26

- device switches
 - changing
 - fabric topology 5-23
- devices
 - changing switches
 - fabric topology 5-23
 - configuration
 - fabric topology 5-14
 - creating a profile 2-8
- disk drives
 - adding
 - fabric topology 5-20
 - adding to the spareset
 - fabric topology 5-21
 - array 2-7
 - corresponding storagesets 2-27
 - dividing 2-19
 - removing from the spareset
 - fabric topology 5-21
- disk label
 - creation 3-6
- displaying the current switches
 - fabric topology 5-23
- dividing storagesets 2-19

E

- enabling
 - switches 2-21
- erasing metadata 2-26
- establishing a local connection 5-2
- executing the installation script 3-3

F

- fabric device
 - naming convention 3-6
- Fabric topology
 - procedure flowchart
 - illustrated 0-xviii
- fabric topology
 - configuration
 - single controller cabling 5-3
- failover 1-4

- transparent 1-4
- file archive
 - temporary storage 3-2
- filesystems
 - creation 3-10
- first enclosure of multiple-enclosure subsystem
 - storage map
 - template 1 A-4
- functions
 - Agent 4-1

G

- geometry
 - initialize switches 2-27

H

- host access
 - restricting in transparent failover mode 1-17
 - separate links 1-17
- host adapter
 - installation 3-1
 - preparation 3-1
- host connections 1-11
 - naming 1-11
- host device
 - initialization 3-4

I

- initialization
 - host device 3-4
- initialize switches
 - changing
 - fabric topology 5-23
 - CHUNKSIZE 2-23
 - DESTROY 2-26
 - geometry 2-27
 - NODESTROY 2-26
 - NOSAVE_CONFIGURATION 2-26
 - SAVE_CONFIGURATION 2-26
- installation
 - controller verification 5-8
 - host adapter 3-1
 - invalid network port assignments B-3

Index-4 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

- solution software 3-2
- there is no disk in the drive message B-4
- installation script
 - execution 3-3
- installation verification
 - CLI commands 5-8
- installing
 - Agent 4-5
- integrating SWCC B-7
- invalid network port assignments B-3

J

- JBOD 2-7

L

- LOCATE
 - find devices 2-29
- location
 - cache module 1-2
 - controller 1-2
- LUN IDs
 - general description 1-19
- LUN presentation 1-15

M

- maintenance port connection
 - establishing a local connection 5-2
 - illustrated 5-2
- mapping storagesets 2-27
- messages
 - there is no disk in the drive B-4
- mirrored caching
 - enabling 1-9
 - illustrated 1-9
- mirrorset switches
 - changing
 - fabric topology 5-23
- mirrorsets
 - planning considerations 2-13
 - important points 2-14
 - switches 2-22
- moving storagesets 7-5
- multiple-bus failover

- CLI configuration procedure
 - fabric topology 6-5
- fabric topology configuration
 - cabling 5-9

N

- naming convention
 - arbitrated loop 3-5
 - fabric device 3-6
- network port assignments B-3
- node IDs 1-17
 - restoring 1-18
- NODE_ID
 - worldwide name 1-17
- NODESTROY 2-26
- NOSAVE_CONFIGURATION 2-26

O

- offset
 - LUN presentation 1-15
 - SCSI version factor 1-16
- online help
 - SWCC B-7
- options
 - for mirrorsets 2-22
 - for RAIDsets 2-22
 - initialize 2-23
- other controller 1-3
- overwriting data 2-26

P

- pager notification B-7
 - configuring B-7
- partitioning
 - storage 3-8
- partitions
 - assigning a unit number
 - fabric topology 5-20
 - defining 2-20
 - planning considerations 2-19
 - guidelines 2-20
- performance 2-15
- planning

- overview 2-8
- striped mirrorsets 2-18
- stripesets 2-12
- planning considerations 2-15
- planning storage
 - containers 2-6
 - where to start 2-2
- planning storage sets
 - characteristics
 - changing switches 2-21
 - enabling switches 2-21
 - initialization switch 2-21
 - storage set switch 2-21
 - unit switch 2-21
 - switches
 - initialization 2-23
 - storage set 2-22
 - unit 2-27
- profiles
 - creating 2-8
 - description 2-8
 - storage set A-1
 - example A-2

R

- RAID set switches
 - changing
 - fabric topology 5-23
- RAID sets
 - choosing chunk size 2-23
 - maximum membership 2-17
 - planning considerations 2-15
 - important points 2-16
 - switches 2-22
- read caching
 - enabled for all storage units 1-8
 - general description 1-8
- read requests
 - decreasing the subsystem response time with
 - read caching 1-8
- read-ahead caching 1-8
 - enabled for all disk units 1-8

- reconfiguring
 - Agent 4-8
- removing
 - Client B-6
- request rate 2-24
- restricting host access
 - separate links
 - transparent failover 1-17
 - transparent failover 1-17
- running
 - Agent 4-8

S

- SAVE_CONFIGURATION 2-26
- saving configuration 2-26
- SCSI version
 - offset 1-16
- SCSI-2
 - assigning unit numbers 1-16
 - command console lun 1-16
- SCSI-3
 - assigning unit numbers 1-16
- selective storage presentation 1-17
- SET CONNECTIONS
 - transparent failover 1-14
- setting
 - controller configuration handling 2-26
- single disk (JBOD)
 - assigning a unit number
 - fabric topology 5-19
- single-enclosure subsystem
 - storage map
 - template 1
 - A-4
- solution software
 - installation 3-2
- storage
 - creating map 2-27
 - partitioning 3-8
 - profile
 - example A-2
 - storage map 2-27

Index-6 HSG60 ACS Solution Software Version 8.6 for SGI Irix Installation and Configuration Guide

- example 2-29
- template 1 A-4
 - first enclosure of multiple-enclosure subsystem A-4
 - single enclosure subsystem A-4
- storageset
 - deleting
 - fabric topology 5-22
 - fabric topology
 - changing switches 5-23
 - planning considerations 2-10
 - mirrorsets 2-13
 - partitions 2-19
 - RAIDsets 2-15
 - striped mirrorsets 2-17
 - stripesets 2-10
- profile 2-8
 - example 2-9
- profiles A-1
- storageset profile 2-8
- storageset switches
 - SET command 2-22
- storagesets
 - creating a profile 2-8
 - moving 7-5
- striped mirrorsets
 - planning 2-18
 - planning considerations 2-17
- stripesets
 - distributing members across buses 2-13
 - planning 2-12
 - planning considerations 2-10
 - important points 2-12
- subsystem
 - saving configuration 2-26
- subsystem configuration
 - backup 7-1
- SWCC 4-1
 - additional information B-7
 - integrating B-7
 - online help B-7
- switches

- changing 2-21
- changing characteristics 2-21
- CHUNKSIZE 2-23
- DESTROY 2-26
- enabling 2-21
- mirrorsets 2-22
- NODESTROY 2-26
- NOSAVE_CONFIGURATION 2-26
- RAIDset 2-22
- SAVE_CONFIGURATION 2-26
- switches for storagesets
 - overview 2-21

T

- templates
 - subsystem profile A-1
- terminology
 - other controller 1-3
 - this controller 1-3
- this controller 1-3
- transparent failover 1-4
 - ADD CONNECTIONS 1-14
 - ADD UNIT 1-14
 - matching units to host connections 1-14
 - restricting host access 1-17
 - separate links 1-17
 - SET CONNECTIONS 1-14
- troubleshooting
 - invalid network port assignments B-3
 - there is no disk in the drive message B-4

U

- uninstalling
 - Agent 4-8
 - Client B-6
- unit numbers
 - assigning 1-14
 - fabric topology 5-19
 - assigning depending on SCSI version 1-16
 - assigning in fabric topology
 - partition 5-20
 - single disk 5-19
- unit qualifiers

- assigning
 - fabric topology 5-19
- unit switches
 - changing
 - fabric topology 5-23
- units
 - LUN IDs 1-19

V

- verification
 - controller installation 5-8
- verification of installation
 - controller 5-8
- virtual disks
 - adding B-7

W

- worldwide names 1-17
 - NODE_ID 1-17
 - REPORTED PORT_ID 1-17
 - restoring 1-18
- write performance 2-26
- write requests
 - improving the subsystem response time with
 - write-back caching 1-8
 - placing data with with write-through caching 1-8
- write-back caching
 - general description 1-8
- write-through caching
 - general description 1-8

