

hp StorageWorks

HSG80 ACS Solution Software Version 8.7 for Sun Solaris Installation and Configuration Guide

Part Number: AA-RFBQH-TE

Eighth Edition (August 2002)

Product Version: 8.7

This guide provides installation and configuration instructions and reference material for operation of the HSG80 ACS Solution Software Version 8.7 for Sun Solaris.



i n v e n t

© Hewlett-Packard Company, 2002. All rights reserved.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Microsoft, MS-DOS, Windows, and Windows NT are trademarks of Microsoft Corporation in the U.S. and/or other countries.

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

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

Printed in the U.S.A.

HSG80 ACS Solution Software Version 8.7 for Sun Solaris Installation and
Configuration Guide
Eighth Edition (August 2002)

Part Number: **AA-RFBQH-TE**

Contents

About this Guide

Intended Audience	xv
Related Documentation	xv
Document Conventions	xvii
Configuration Flowchart	xviii
Symbols in Text	xxiii
Symbols on Equipment	xxiii
Rack Stability	xxiv
Getting Help	xxv
Technical Support	xxv
Storage Website	xxv
Authorized Reseller	xxv

1 Planning a Subsystem

Defining Subsystems	1-2
Controller Designations A and B	1-2
Controller Designations “This Controller” and “Other Controller”	1-3
What is Failover Mode?	1-5
Transparent Failover Mode	1-5
Multiple-Bus Failover Mode	1-7
Selecting a Cache Mode	1-9
Read Caching	1-10
Read-Ahead Caching	1-10
Write-Back Caching	1-10
Write-Through Caching	1-10
Enabling Mirrored Caching	1-10
What is the Command Console LUN?	1-11
Determining the Address of the CCL	1-12
Enabling/Disabling the CCL in SCSI-2 Mode	1-12
Enabling the CCL	1-12
Disabling the CCL	1-12
Enabling/Disabling CCL in SCSI-3 Mode	1-12

Determining Connections	1–13
Naming Connections	1–13
Numbers of Connections	1–13
Assigning Unit Numbers	1–16
Matching Units to Host Connections in Transparent Failover Mode	1–17
Matching Units to Host Connections in Multiple-Bus Failover Mode	1–19
Assigning Unit Numbers Depending on SCSI_VERSION	1–19
Assigning Host Connection Offsets and Unit Numbers in SCSI-3 Mode.	1–20
Assigning Host Connection Offsets and Unit Numbers in SCSI-2 Mode.	1–20
What is Selective Storage Presentation?	1–21
Restricting Host Access in Transparent Failover Mode	1–21
Restricting Host Access by Separate Links	1–22
Restricting Host Access by Disabling Access Paths	1–23
Restricting Host Access by Offsets	1–23
Restricting Host Access in Multiple-Bus Failover Mode	1–24
Enable the Access Path of Selected Host Connections	1–24
Restricting Host Access by Offsets	1–26
Worldwide Names (Node IDs and Port IDs)	1–27
Restoring Worldwide Names (Node IDs)	1–28
Unit Worldwide Names (LUN IDs)	1–29

2 Planning Storage Configurations

Where to Start	2–2
Determining Storage Requirements	2–3
Configuration Rules for the Controller	2–3
Addressing Conventions for Device PTL	2–4
Examples - Model 2200 Storage Maps, PTL Addressing	2–7
Choosing a Container Type	2–14
Creating a Storageset Profile	2–16
Planning Considerations for Storageset	2–18
Stripeset Planning Considerations	2–18
Mirrorset Planning Considerations	2–21
Keep these points in mind when planning mirrorsets	2–22
RAIDset Planning Considerations	2–22
Keep these points in mind when planning RAIDsets	2–23
Striped Mirrorset Planning Considerations	2–24
Storageset Expansion Considerations	2–26
Partition Planning Considerations	2–26

Defining a Partition	2–27
Guidelines for Partitioning Storagesets and Disk Drives	2–27
Changing Characteristics through Switches	2–27
Enabling Switches	2–28
Changing Switches	2–28
Specifying Storageset and Partition Switches	2–28
RAIDset Switches	2–28
Mirrorset Switches	2–29
Partition Switches	2–29
Specifying Initialization Switches	2–29
Chunk Size	2–30
Increasing the Request Rate	2–30
Increasing Sequential Data Transfer Performance	2–32
Save Configuration	2–32
Destroy/Nodestroy	2–32
Geometry	2–33
Specifying Unit Switches	2–33
Creating Storage Maps	2–33
Using LOCATE Command to Find Devices	2–34
Example Storage Map - Model 4310R Disk Enclosure	2–35

3 Preparing the Host System

Installing RAID Array Storage System	3–1
Making a Physical Connection	3–6
Preparing to Install Host Bus Adapter	3–6
Installing Host Bus Adapter	3–6
Verifying/Installing Required Versions	3–8
Choosing Automatic or Manual Install Modes	3–8
Installing Solution Software Packages	3–9
Installing with Automatic Mode	3–10
Installing with Manual Mode	3–11
Preparing LUNs for Use by the FileSystem	3–13
Labeling LUNs	3–13
Labeling a LUN Automatically Using the Format Utility	3–13
Creating and Tuning File Systems	3–15
Solution Software Upgrade Procedures	3–16

4 Installing and Configuring HSG Agent

Why Use StorageWorks Command Console (SWCC)?	4-1
Installation and Configuration Overview	4-2
About the Network Connection for the Agent	4-3
Before Installing the Agent.	4-5
Configuring the Agent	4-5
Configuring the Agent with Install.sh (for First-time Configurations)	4-5
Configuring the Agent with RAID Array V8.7 Configuration Menu	4-6
Choosing a Password.	4-7
Adding a Subsystem Entry	4-8
Adding a Client System Entry	4-8
Restarting the SWCC Agent	4-8
Configuring the Agent within FirstWatch	4-9
RAID Manager Mail Messages	4-9

5 FC Configuration Procedures

Establishing a Local Connection	5-2
Establishing Connection with a SPARC System	5-3
Setting Up a Single Controller	5-3
Power On and Establish Communication	5-3
Cabling a Single Controller	5-4
Configuring a Single Controller Using CLI	5-4
Verify the Node ID and Check for Any Previous Connections	5-5
Configure Controller Settings	5-6
Restart the Controller.	5-7
Set Time and Verify all Commands	5-7
Plug in the FC Cable and Verify Connections	5-9
Repeat Procedure for Each Host Adapter	5-9
Verify Installation	5-9
Setting Up a Controller Pair	5-10
Power Up and Establish Communication	5-10
Cabling a Controller Pair	5-10
Configuring a Controller Pair Using CLI	5-11
Configure Controller Settings	5-12
Restart the Controller.	5-13
Set Time and Verify All Commands	5-13
Plug in the FC Cable and Verify Connections	5-15
Repeat Procedure for Each Host Adapter Connection	5-15

Verify Installation	5-16
Configuring Devices	5-16
Configuring Storage Containers.	5-16
Configuring a Stripeset	5-17
Configuring a Mirrorset.	5-18
Configuring a RAIDset	5-19
Configuring a Striped Mirrorset	5-19
Configuring a Single-Disk Unit (JBOD).	5-20
Configuring a Partition	5-20
Assigning Unit Numbers and Unit Qualifiers	5-22
Assigning a Unit Number to a StorageSet	5-22
Assigning a Unit Number to a Single (JBOD) Disk	5-22
Assigning a Unit Number to a Partition	5-22
Configuration Options	5-23
Changing the CLI Prompt	5-23
Mirroring cache	5-23
Adding Disk Drives.	5-23
Adding a Disk Drive to the Spareset	5-23
Removing a Disk Drive from the Spareset	5-24
Enabling Autospare	5-24
Deleting a StorageSet.	5-25
Changing Switches for a StorageSet or Device	5-25
Displaying the Current Switches.	5-25
Changing RAIDset and Mirrorset Switches	5-26
Changing Device Switches	5-26
Changing Initialize Switches	5-26
Changing Unit Switches	5-26
FC Considerations for Both Loop and Fabric Environments.	5-26
General FC Environmental Considerations.	5-27
Configurable Objects.	5-27
Servers.	5-27
Storage	5-28
Active-Passive vs. Active-Active Controller Pairs	5-29
Hubs	5-30
Switches	5-30
Fibre Channel Access/Transport Paths	5-31
Implementation of Loop and Fabric Bindings.	5-31
Bindings	5-31

Loop Bindings	5-31
Fabric Bindings	5-32
Loop Connections at the Solaris Level	5-32
Loop Mappings	5-33
Analysis	5-33
Fabric Mappings	5-33
Analysis	5-34
The /opt/steam/bin/config.sh Utility	5-34
Adding/Changing Adapters	5-35
Change the Mode from Fabric to Loop (For JNI HBAs only)	5-38
Change the Number of LUNs per Target	5-39
Set the Target IDs	5-39
Accept/Reject Changes	5-40
Configuration Procedures	5-40
Loop Configurations	5-41
Fabric Configurations	5-41
Reconfiguring from Loop to Fabric	5-42
Reconfiguring from Fabric to Loop	5-43

6 Using CLI for Configuration

CLI Configuration Example	6-4
-------------------------------------	-----

7 Backing Up, Cloning, and Moving Data

Backing Up Subsystem Configurations	7-1
Creating Clones for Backup	7-2
Linking WWIDs for Snap and Clone Units	7-5
CLI format	7-6
Implementation Notes	7-6
Moving Stagesets	7-7

A Subsystem Profile Templates

Storageset Profile A-2

Storage Map Template 1 for the BA370 Enclosure A-4

Storage Map Template 2 for the second BA370 Enclosure A-5

Storage Map Template 3 for the third BA370 Enclosure A-6

Storage Map Template 4 for the Model 4214R Disk Enclosure A-7

Storage Map Template 5 for the Model 4254 Disk Enclosure A-9

Storage Map Template 6 for the Model 4310R Disk Enclosure A-11

Storage Map Template 7 for the Model 4350R Disk Enclosure A-14

Storage Map Template 8 for the Model 4314R Disk Enclosure A-16

Storage Map Template 9 for the Model 4354R Disk Enclosure A-19

B Installing, Configuring, and Removing the Client

Why Install the Client? B-2

Before You Install the Client B-2

Installing the Client B-4

Installing the Integration Patch B-5

 Should I Install the Integration Patch? B-5

 How to Install the Integration Patch B-5

 Integrating Controller’s SWCC Storage Window with CIM B-6

 Insight Manager Unable to Find Controller’s Storage Window B-7

 Removing the Integration Patch Will Corrupt Storage Window B-7

Troubleshooting Client Installation B-8

 Invalid Network Port Assignments During Installation B-8

 “There is no disk in the drive” Message B-9

Adding Storage Subsystem and its Host to Navigation Tree B-10

Removing Command Console Client B-12

Where to Find Additional Information B-13

 About the User Guide B-13

 About the Online Help B-13

Glossary

Index

Figures

1	General configuration flowchart (panel 1)	xx
2	General configuration flowchart (panel 2)	xxi
3	Configuring storage with SWCC.	xxii
1-1	Location of controllers and cache modules in a Model 2200 enclosure	1-2
1-2	Location of controllers and cache modules in a BA370 enclosure	1-3
1-3	“This controller” and “other controller” for the Model 2200 enclosure	1-4
1-4	“This controller” and “other controller” for the BA370 enclosure	1-4
1-5	Transparent failover—normal operation	1-6
1-6	Transparent failover—after failover from controller B to controller A	1-7
1-7	Typical multiple-bus configuration	1-9
1-8	Mirrored caching	1-11
1-9	Connections in separate-link, transparent failover mode configurations	1-14
1-10	Connections in single-link, transparent failover mode configurations	1-15
1-11	Connections in multiple-bus failover mode.	1-16
1-12	LUN presentation to hosts, as determined by offset	1-18
1-13	Limiting host access in transparent failover mode	1-22
1-14	Limiting host access in multiple-bus failover mode	1-25
1-15	Placement of the worldwide name label on the Model 2200 enclosure	1-28
1-16	Placement of the worldwide name label on the BA370 enclosure	1-29
2-1	Mapping a unit to physical disk drives	2-4
2-2	PTL naming convention	2-5
2-3	How data is laid out on disks in BA370 enclosure configuration	2-6
2-4	Storage container types	2-14
2-5	3-member RAID 0 stripeset (example 1)	2-18
2-6	3-member RAID 0 stripeset (example 2)	2-19
2-7	Mirrorsets maintain two copies of the same data	2-21
2-8	Mirrorset example 2.	2-21
2-9	5-member RAIDset using parity	2-23
2-10	Striped mirrorset (example 1)	2-25
2-11	Striped mirrorset (example 2)	2-25
2-12	One example of a partitioned single-disk unit.	2-26
2-13	Large chunk size increases request rate.	2-31
3-1	Dual-Bus Enterprise Storage RAID Array Storage System	3-4
3-2	Single-Bus Enterprise Storage RAID Array Storage System	3-5
4-1	An example of a network connection	4-4
4-2	RAID Array V8.7 Configuration Menu	4-7
5-1	Maintenance port connection.	5-2

5-2	Single controller cabling	5-4
5-3	Controller pair failover cabling	5-11
5-4	Storage container types	5-17
6-1	Example storage map for the BA370 Enclosure	6-2
6-2	Example, three non-clustered host systems	6-3
6-3	Example, logical or virtual disks comprised of storagesets	6-4
6-4	Example, virtual system layout from hosts' point of view	6-4
7-1	CLONE utility steps for duplicating unit members	7-2
B-1	Navigation Window	B-10
B-2	Navigation window showing storage host system "Atlanta"	B-11
B-3	Navigation window showing expanded "Atlanta" host icon	B-11

Tables

1	Document Conventions	xvii
2	Summary of Chapter Contents	xviii
1-1	Unit Assignments and SCSI_VERSION	1-21
2-1	PTL addressing, single-bus configuration, six Model 4310R enclosures	2-8
2-2	PTL addressing, dual-bus configuration, three Model 4350R enclosures	2-10
2-3	PTL addressing, single-bus configuration, six Model 4314R enclosures	2-11
2-4	PTL addressing, dual-bus configuration, three Model 4354A enclosures	2-13
2-5	Comparison of Container Types	2-15
2-6	Example of Storageset Profile	2-17
2-7	Example Chunk Sizes	2-32
2-8	Model 4310R disk enclosure, example of storage map	2-35
3-1	<i>StorageWorks</i> Solution Software Packages	3-7
4-1	SWCC Features and Components	4-2
4-2	Installation and Configuration Overview	4-3
5-1	Case 1: Active-Passive (Describes the controller pair)	5-29
5-2	Case 2: Active-Active (Describes the controller pair)	5-29
5-3	Example of Active-Passive with a Single Controller	5-29

About this Guide

This guide describes how to install and configure the HSG80 ACS Solution Software Version 8.7 for Sun Solaris.

This guide describes:

- How to plan the storage array subsystem; and,
- How to install and configure the storage array subsystem on individual operating system platforms.

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.

Intended Audience

This book is intended for use by system administrators and system technicians who have a basic experience with storage and networking.

Related Documentation

In addition to this guide, corresponding information can be found in:

- ACS v8.7 controller documentation (software delivered via PCMCIA cards)
- HSG80 CLI Reference Guide, EK-G80CL-RA.B01
- HSG80 Maintenance and Service Guide, EK-G80MS-SA.B01
- HSG80 Troubleshooting and Reference Guide, EK-G80TR-SA.B01
- SWCC v2.5 documentation (client software delivered in solutions kits)
- Command Console User Guide, AA-RFA2J-TE
- Command Console Release Notes, AV-RPBKB-TE
- Command Console Help Files, AA-RS20A-TE and AA-RS21A-TE
- Host-specific documentation (SWCC Agent and HBA software delivered in solutions kits)

- Installation and Configuration Guide (platform-specific) - the guide you are reading
- Solution Software Release Notes (platform-specific)
- FC-AL Application Note (AA-RS1ZA-TE) - Solution software host support includes the following platforms:
 - IBM AIX
 - HP-UX
 - Linux (Red Hat x86/Alpha, SuSE x86/Alpha, Caldera x86)
 - Novell NetWare
 - Open VMS
 - Sun Solaris
 - Tru64 UNIX
 - Windows NT/2000

Additional support required by HSG80 ACS Solution Software Version 8.7, but delivered through external programs, include the following:

- Heterogeneous “rules based” SAN configurations
- Host-Bus Adapter (HBA) products
- Applicable Storage Utility Management Suite (SUMS) components
- Vendor-specific switch products
- Secure Path Products (Windows, NetWare, Sun, AIX, HP-UX)
- Data Replication Manager (DRM) under ACS 8.7P
- Enterprise Volume Manager (EVM) under ACS 8.7S
- Enterprise Backup Solution (EBS)
- Additional ACS Variants (W, R)

Document Conventions

The conventions included in [Table 1](#) apply.

Table 1: Document Conventions

Element	Convention
Cross-reference links	Blue text: Figure 1
Key names, menu items, buttons, and dialog box titles	Bold
File names, application names, and text emphasis	<i>Italics</i>
User input, command names, system responses (output and messages)	Monospace font COMMAND NAMES are uppercase unless they are case sensitive
Variables	<i>Monospace, italic font</i>
Website addresses	Sans serif font (http://www.compaq.com)

Configuration Flowchart

A three-part flowchart (Figures 1-3) is shown on the following pages. Refer to these charts while installing and configuring a new storage subsystem. All references in the flowcharts pertain to pages in this guide, unless otherwise indicated.

Table 2 below summarizes the content of the chapters.

Table 2: Summary of Chapter Contents

Chapters	Description
1. Planning a Subsystem	This chapter focuses on technical terms and knowledge needed to plan and implement storage array subsystems.
2. Planning Storage Configurations	Plan the storage configuration of your subsystem, using individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives. This chapter describes addressing conventions, configuration rules, creating storage profiles, and creating storage maps.
3. Preparing the host system	How to prepare your Sun Solaris host computer to accommodate the HSG80 controller storage subsystem.
4. Installing, Configuring the HSG Agent	The Agent for HSG for a specific operating system polls the storage.
5. FC Configuration Procedures	How to configure a subsystem that uses Fibre Channel (FC) fabric topology. In fabric topology, the controller connects to its hosts through switches. A section of this chapter describes the differences and configuration processes for Sun servers and adapters for loop and switch support.
6. Configuration using CLI	How-to example of configuring a storage subsystem using the Command Line Interpreter (CLI).
7. Backup, Clone, Move Data	Description of common procedures that are not mentioned elsewhere in this guide. <ul style="list-style-type: none">• Backing Up Subsystem Configuration• Cloning Data for Backup• Moving Storagesets

Table 2: Summary of Chapter Contents (Continued)

Chapters	Description
Appendix A. Subsystem Profile Templates	This appendix contains storageset profiles to copy and use to create your 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.
Appendix B. Install, Configure, Remove the Client	The Client monitors and manages a storage subsystem. The following information is included in this appendix: <ul style="list-style-type: none">• Why Install the Client?• Before You Install the Client• Installing the Client• Installing Integration Patch• Troubleshooting the Client Installation• Adding Storage Subsystem and its Host to the Navigation Tree• Removing the Command Console Client• Where to Find Additional Information

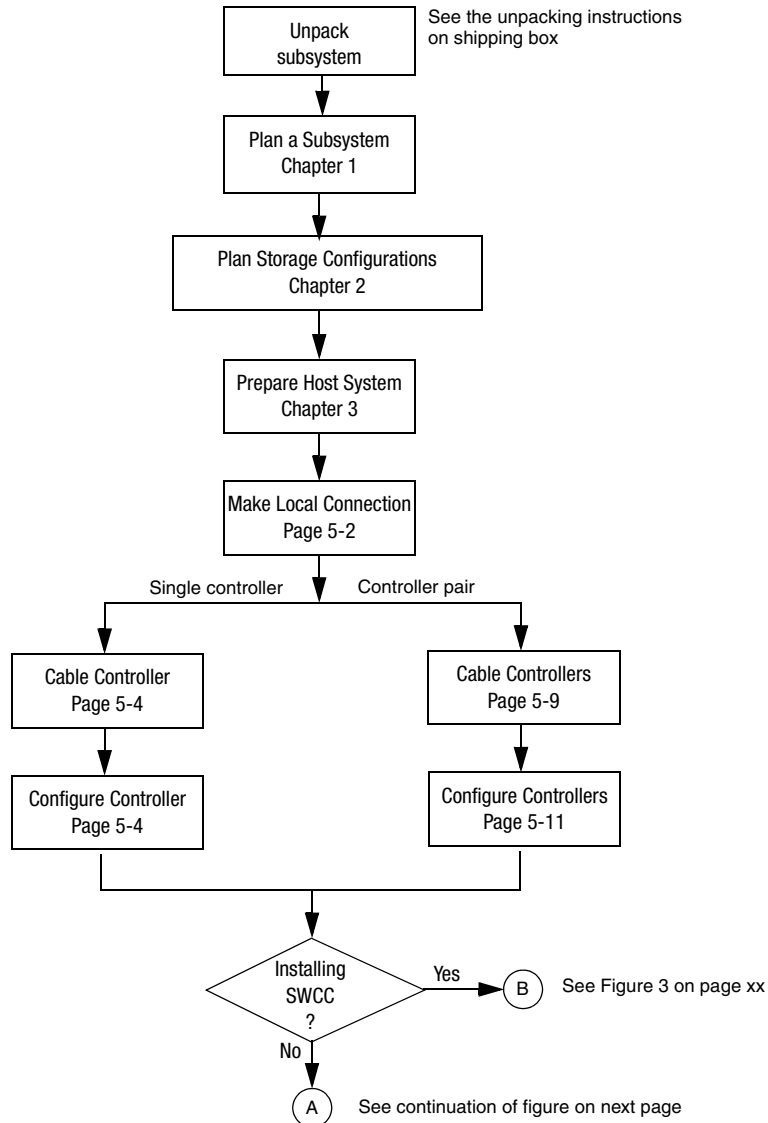


Figure 1: General configuration flowchart (panel 1)

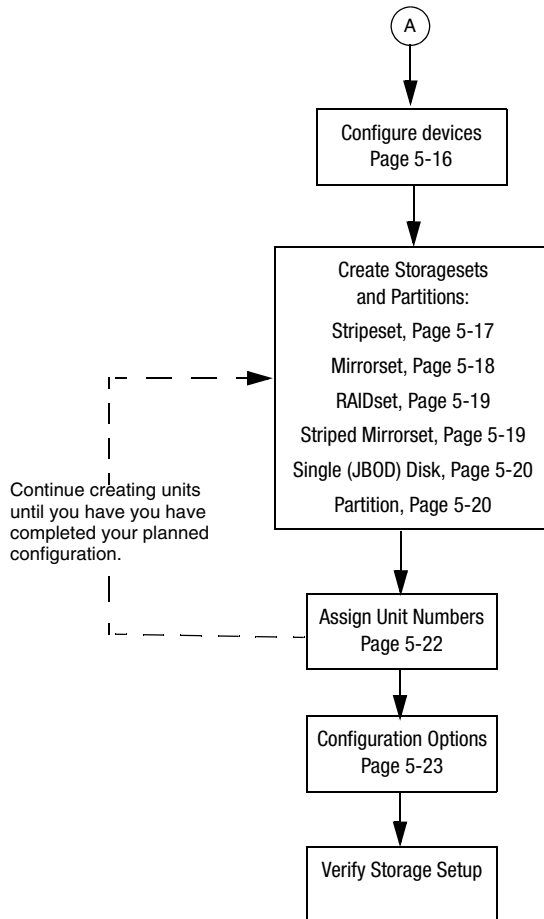


Figure 2: General configuration flowchart (panel 2)

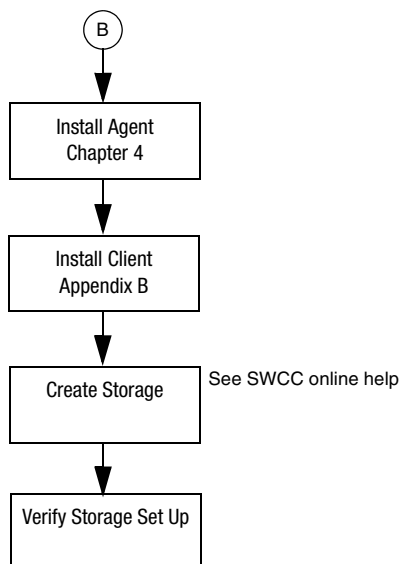


Figure 3: Configuring storage with SWCC

Symbols in Text

These symbols may be found in the text of this guide. They 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 data.

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

NOTE: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Symbols on Equipment



Any enclosed 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.



defined; 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. Contact with this surface could result in injury.

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



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.



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 manually handling material.

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.
 - In single rack installations, the stabilizing feet are attached to the rack.
 - In multiple rack installations, the racks are coupled.
 - Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.
-

Getting Help

If you still have a question after reading this guide, contact an authorized service provider or access our website.

Technical Support

In North America, call technical support at 1-800-OK-COMPAQ, available 24 hours a day, 7 days a week.

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

Outside North America, call technical support at the nearest location. Telephone numbers for worldwide technical support are listed on the following website: <http://www.compaq.com>.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

Storage Website

The Storage website has the latest information on this product, as well as the latest drivers. Access the Storage website at: <http://www.compaq.com/storage>. From this website, select the appropriate product or solution.

Authorized Reseller

For the name of your nearest Authorized Reseller:

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

Planning a Subsystem

This chapter provides information that helps you plan how to configure the storage array subsystem. This chapter focuses on the technical terms and knowledge needed to plan and implement storage subsystems.

NOTE: This chapter frequently references the command line interface (CLI). For the complete syntax and descriptions of the CLI commands, see the *StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide*.

The following information is included in this chapter:

- “Defining Subsystems,” page 1–2
- “What is Failover Mode?,” page 1–5
- “Selecting a Cache Mode,” page 1–9
- “Enabling Mirrored Caching,” page 1–10
- “What is the Command Console LUN?,” page 1–11
- “Determining Connections,” page 1–13
- “Assigning Unit Numbers,” page 1–16
- “What is Selective Storage Presentation?,” page 1–21

IMPORTANT: DILX should be run for ten minutes on all units to delete the 8 MB EISA partition. Refer to *StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide* for details.

Refer to Chapter 2 when planning the types of storage containers you need. Storage containers are individual disk drives (JBOD), storage set types (mirrorsets, stripesets, and so on), and/or partitioned drives.

Defining Subsystems

This section describes the terms *this controller* and *other controller*. It also presents graphics of the Model 2200 and BA370 enclosures.

NOTE: The HSG80 controller uses the BA370 or Model 2200 enclosure.

Controller Designations A and B

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.

Controllers and cache modules are designated either A or B depending on their location in the enclosure, as shown in Figure 1–1 for the Model 2200 enclosure and in Figure 1–2 for the BA370 enclosure.

Model 2200 Enclosure




Figure 1–1: Location of controllers and cache modules in a Model 2200 enclosure




Figure 1–2: Location of controllers and cache modules in a BA370 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 abbreviated 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 abbreviated to “other” in CLI commands.

Figure 1–3 shows the relationship between “this controller” and “other controller” in a Model 2200 enclosure. Figure 1–4 shows the same relationship in a BA370 enclosure.


Model 2200 Enclosure



- ❶ This controller ❷ Other controller

Figure 1-3: “This controller” and “other controller” for the Model 2200 enclosure

BA370 Enclosure



- ❶ Other controller ❷ This controller

Figure 1-4: “This controller” and “other controller” for the BA370 enclosure

What is 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 controller hardware failure or, in multiple-bus mode only, due to a failure of the link between host and controller or host-bus adapter. Failover keeps the storage array available to the hosts by allowing the surviving controller to take over total control of the subsystem.

There are two failover modes:

- **Transparent**, which is handled by the surviving controller and is invisible (transparent) to the hosts.
- **Multiple-bus**, which is handled by the surviving controller and which is handled by the hosts through either additional software or as a feature of the operating system.

Either mode of failover can work with loop or fabric topology.

NOTE: Multiple-bus failover mode is only supported when Secure Path is used. Refer to the Secure Path documentation for more information on using multiple-bus mode.

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 set is a physical or virtual device of the subsystem. It is typically assigned a logical unit number (LUN) and is managed by the HSG80 controller and presented to a server through the Fibre Channel bus and the server's host bus adapter. Disks that are set up as independent disks (JBODs) or RAID sets are referred to as storage sets.

In transparent failover mode, host port 1 of controller A and host port 1 of controller B must be on the same Fibre Channel link. Host port 2 of controller A and host port 2 of controller B 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 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, both host port 1 on controller A and host port 2 on controller B are active. A representative

configuration is shown in Figure 1–5. 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 1–6.

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).

Transparent failover only compensates for a controller failure, and not for failures of either the Fibre Channel link or host Fibre Channel adapters.


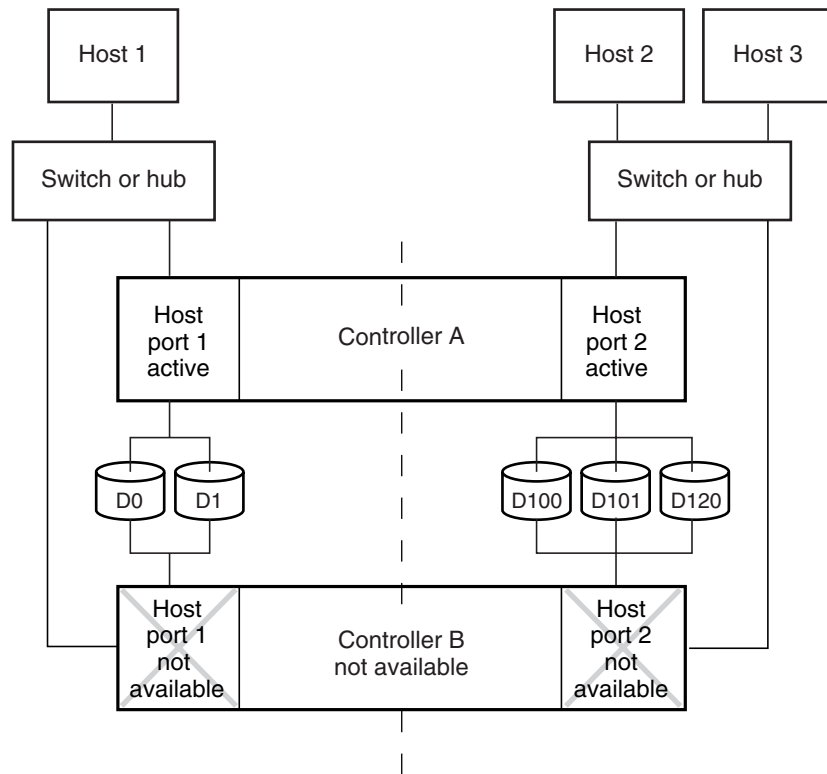


Figure 1–5: Transparent failover—normal operation



CXO7035A

Figure 1–6: Transparent failover—after failover from controller B to controller A

Multiple-Bus Failover Mode

Multiple-bus failover mode has the following characteristics:

- Host controls the failover process by moving the units from one controller to another
- All units (0 through 199) are visible at all host ports
- Each host has two or more paths to the units

All hosts must have operating system software that supports multiple-bus failover mode. With this software, the host sees the same units visible through two (or more) paths. When one path fails, the host can issue commands to move the units from one path to another. A typical multiple-bus failover configuration is shown in Figure 1–7.


In multiple-bus failover mode, you can specify which units are normally serviced by a specific controller of a controller pair. Units can be preferred to one controller or the other by the *PREFERRED_PATH* switch of the *ADD UNIT* (or *SET unit*) command. For example, use the following command to prefer unit D101 to “this controller”:

```
SET D101 PREFERRED_PATH=THIS_CONTROLLER
```

NOTE: This is an initial preference, which can be overridden by the hosts.

Keep the following points in mind when configuring controllers for multiple-bus failover:

- Multiple-bus failover can compensate for a failure in any of the following:
 - Controller
 - Switch or hub
 - Fibre Channel link
 - Host Fibre Channel adapter
- A host can redistribute the I/O load between the controllers
- All hosts must have operating system software that supports multiple-bus failover mode



NOTE: FCA = Fibre Channel Adapter

CXO7094B

Figure 1–7: Typical multiple-bus configuration

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 allowing 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 1-8.

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




Figure 1–8: 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.

What is the Command Console LUN?

StorageWorks Command Console (SWCC) software communicates with the HSG80 controllers through an existing storage unit, or logical unit number (LUN). The dedicated LUN that SWCC uses is called the Command Console LUN (CCL). The 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 HSG80 controller on host port 1. The HSG80 uses both SCSI-2 or SCSI-3 with your platform.

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 sections alter the setting of the CCL on both controllers. The CCL is enabled only on host port 1. At least one storage device of any type must be configured on host port 2 before installing the Agent on a host connected to host port 2.

Select a storageset that you plan to configure and that is not likely to change. This storageset can be used by the Agent to communicate with the RAID Array. Deleting this storageset (LUN) later breaks the connection between the Agent and the RAID Array.

Determining the Address of the CCL

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

```
HSG80 > SHOW THIS_CONTROLLER
```

Enabling/Disabling the CCL in SCSI-2 Mode

Enabling the CCL

To enable the CCL, enter the following CLI command:

```
HSG80 > SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

Disabling the CCL

To disable the CCL, enter the following CLI command:

```
HSG80 > 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 CCL in SCSI-3 Mode

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


Determining 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 Version 8.7, the maximum number of supported connections is 96.

Naming Connections

It is highly recommended 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 the host named 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.

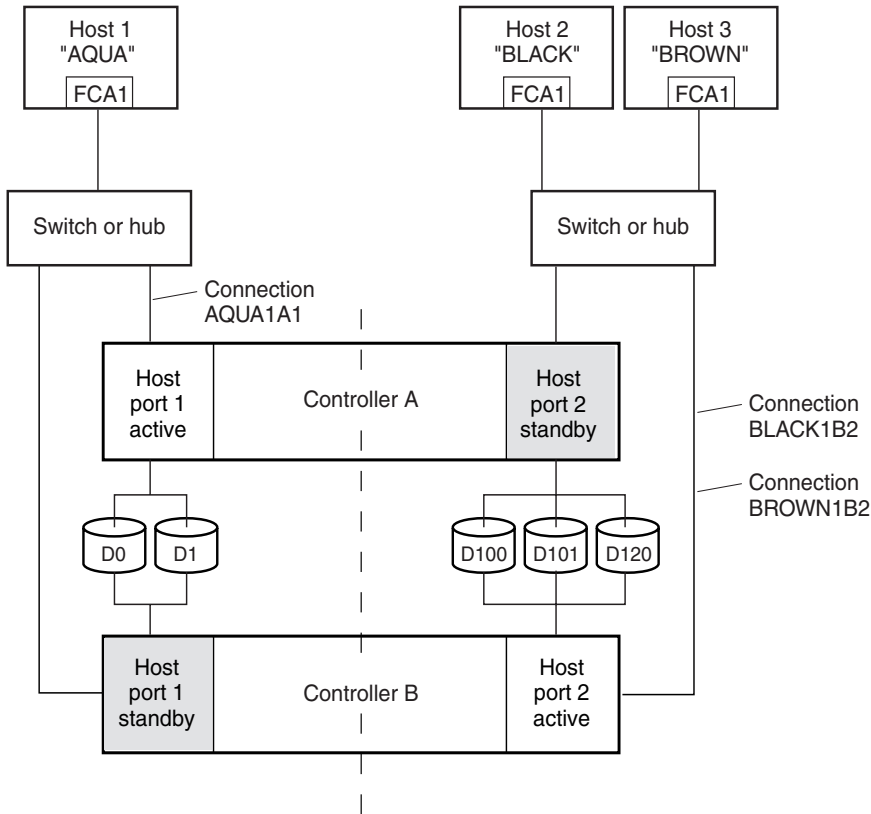
NOTE: Connection names can have a maximum of 9 characters.

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 1–9.

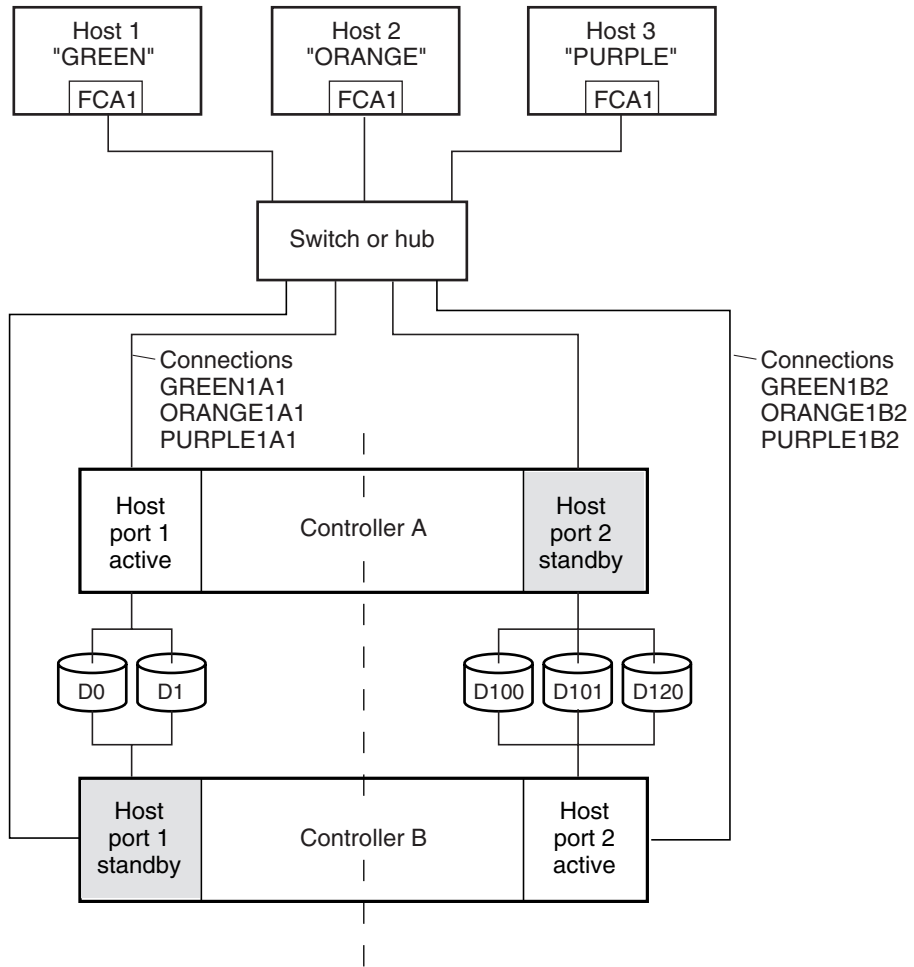
- If a controller pair is in transparent failover mode and port 1 and port 2 are on the same link (that is, all ports are on the same loop or fabric), each adapter will have two connections, as shown in Figure 1–10.
- If a controller pair is in multiple-bus failover mode, each adapter has two connections, as shown in Figure 1–11.



NOTE: FCA = Fibre Channel Adapter

CXO7081B


Figure 1–9: Connections in separate-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CX07079B

Figure 1–10: Connections in single-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CXO7080B

Figure 1–11: Connections in multiple-bus failover mode

Assigning Unit Numbers

The controller keeps track of the unit with the unit number. The 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 connection to access and assigns it to either port 1 or 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 available through 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 is shown in the following equation:

$$\text{LUN number} = \text{unit number} - \text{offset}$$

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 D102 minus offset of 100).

Figure 1–12 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 1–12: 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 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 1–21.

Matching Units to Host Connections in Multiple-Bus Failover Mode

In multiple-bus failover mode, the `ADD UNIT` command creates a unit for host connections to access. All unit numbers (0 through 199) are potentially visible on all four controller ports, but are accessible only to those host connections for which access path is enabled and which have offsets in the unit's range.

The LUN number a host connection assigns to a unit is a function of the `UNIT_OFFSET` switch of the `ADD` (or `SET`) `CONNECTIONS` command. The default offset is 0. The relationship of offset, LUN number, and unit number is shown in the following equation:

$$\text{LUN number} = \text{unit number} - \text{offset}$$

Where...

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

For example, unit D7 would be visible to a host connection with an offset of 0 as LUN 7 (unit number of 7 minus offset of 0). Unit D17 would be visible to a host connection with an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible at all to a host connection with a unit offset of 18 or greater, because that offset is not within the units range (unit number of 17 minus offset of 18 is a negative number).

In addition, the access path to the host connection must be enabled for the connection to access the unit. This is done through the `ENABLE_ACCESS_PATH` switch of the `ADD UNIT` (or `SET unit`) command.

The `PREFERRED_PATH` switch of the `ADD UNIT` (or `SET unit`) command determines which controller of a dual-redundant pair initially accesses the unit. Initially, `PREFERRED_PATH` determines which controller presents the unit as Ready. The other controller presents the unit as Not Ready. Hosts can issue a SCSI Start Unit command to move the unit from one controller to the other.

Assigning Unit Numbers Depending on SCSI_VERSION

The `SCSI_VERSION` switch of the `SET THIS_CONTROLLER/OTHER_CONTROLLER` command determines how the 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 Host Connection Offsets and Unit Numbers in SCSI-3 Mode

If `SCSI_VERSION` is set to SCSI-3, the CCL is presented as LUN 0 to all connections. The CCL supersedes any other unit assignment. Therefore, in SCSI-3 mode, a unit that would normally be presented to a connection as LUN 0 is not visible to that connection at all.

The following methods are recommended for assigning host connection offsets and unit numbers in SCSI-3 mode:

- Offsets should be divisible by 10 (for consistency and simplicity).
- Unit numbers should not be assigned at connection offsets (to avoid being masked by the CCL at LUN 0).

For example, if a host connection has an offset of 20 and SCSI-3 mode is selected, the connection will see LUNs as follows:

LUN 0 - CCL

LUN 1 - unit 21

LUN 2 - unit 22, etc.

In this example, if a unit 20 is defined, it will be superseded by the CCL and invisible to the connection.

Assigning Host Connection Offsets and Unit Numbers 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.

StorageWorks recommends to use the following conventions 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 SCSI_VERSION

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

What is Selective Storage Presentation?

Selective Storage presentation is a feature of the HSG80 controller that enables the user to control the allocation of storage space and shared access to storage across multiple hosts. This is also known as *Restricting Host Access*.

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.

NOTE: The default condition is ENABLE_ACCESS_PATH=ALL. This specifies that access paths to ALL hosts are enabled. StorageWorks recommends that the user restrict host access and that the access path be carefully specified to avoid providing undesired host connections access to the unit.

Restricting Host Access in Transparent Failover Mode

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

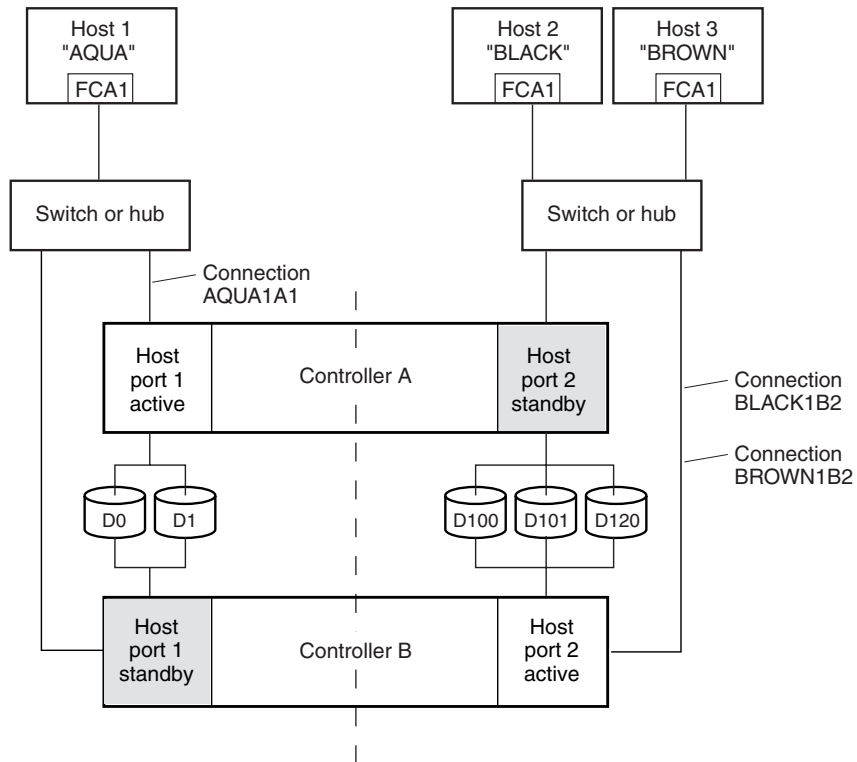
- Using separate Fibre Channel links (either loop or fabric)
- Enabling the access path of selected host connections on a shared loop or fabric
- Setting offsets

NOTE: These techniques also work for a single controller.

Restricting Host Access by Separate Links

In transparent failover mode, host port 1 of controller A and host port 1 of controller B share a common Fibre Channel link. Host port 2 of controller A and host port 2 of controller B also share a common Fibre Channel link. If the host 1 link is separate from the host 2 link, the simplest way to limit 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 1–13. This method applies only if the host 1 link and host 2 link are separate links.

NOTE: It is highly recommended that you provide access to only specific connections. This way, if new connections are added, they will not have automatic access to all units. See the following section Restricting Host Access by Disabling Access Paths.



NOTE: FCA = Fibre Channel Adapter

CXO7081B

Figure 1–13: Limiting host access in transparent failover mode

Restricting Host Access by Disabling Access Paths

If more than one host is on a link (that is, attached to the same port), host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the *ENABLE_ACCESS_PATH* and *DISABLE_ACCESS_PATH* switches of the *ADD UNIT* (or *SET unit*) commands. The access path is a unit switch, meaning it must be specified for each unit. Default access enables the unit to be accessible to all hosts.

For example:

In Figure 1–14, restricting the access of unit D101 to host 3, the host named BROWN can be done by enabling only the connection to host 3. Enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=BROWN1B2
```

If the storage subsystem has more than one host connection, carefully specify the access path to avoid providing undesired host connections access to the unit. The default condition for a unit is that access paths to all host connections are enabled. To restrict host access to a set of host connections, specify *DISABLE_ACCESS_PATH=ALL* for the unit, then specify the set of host connections that are to have access to the unit.

Enabling the access path to a particular host connection does not override previously enabled access paths. All access paths previously enabled are still valid; the new host connection is simply added to the list of connections that can access the unit.

NOTE: The procedure of restricting access by enabling all access paths then disabling selected paths is **not recommended** because of the potential data/security breach that occurs when a new host connection is added.

Restricting Host Access by Offsets

Offsets establish the start of the range of units that a host connection can access.

For example:

In Figure 1–13, assume both host connections on port 2 (connections BLACK1B2 and BROWN1B2) initially have the default port 2 offset of 100. Setting the offset of connection BROWN1B2 to 120 will present unit D120 to host BROWN as LUN 0.

```
SET BROWN1B2 UNIT_OFFSET=120
```

Host BROWN cannot see units lower than its offset, so it cannot access units D100 and D101. However, host BLACK can still access D120 as LUN 20 if the operating system permits. To restrict access of D120 to only host BROWN, enable only host BROWN's access, as follows:

```
SET D120 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=BROWN1B2
```

NOTE: StorageWorks recommends that you provide access to only specific connections, even if there is just one connection on the link. This way, if new connections are added, they will not have automatic access to all units.

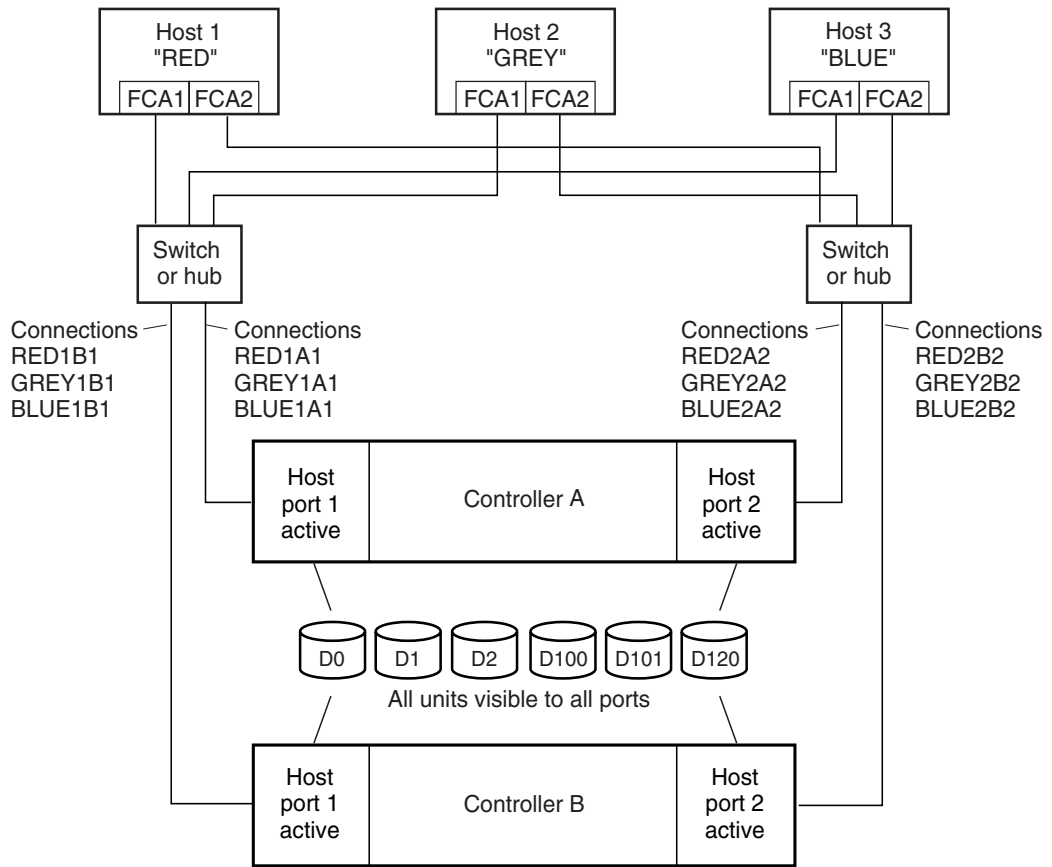
Restricting Host Access in Multiple-Bus Failover Mode

In multiple-bus mode, the units assigned to any port are visible to all ports. There are two ways to limit host access in multiple-bus failover mode:

- Enabling the access path of selected host connections
- Setting offsets

Enable the Access Path of Selected Host Connections

Host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the *ENABLE_ACCESS_PATH* and *DISABLE_ACCESS_PATH* switches of the *ADD UNIT* (or *SET unit*) commands. Access path is a unit switch, meaning it must be specified for each unit. Default access means that the unit is accessible to all hosts. It is important to remember that at least two paths between the unit and the host must be enabled in order for multiple-bus failover to work.



NOTE: FCA = Fibre Channel Adapter

CXO7078B

Figure 1–14: Limiting host access in multiple-bus failover mode

For example:

Figure 1–14 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host BLUE can be done by enabling only the connections to host BLUE. At least two connections must be enabled for multiple-bus failover to work. For most operating systems, it is desirable to have all connections to the host enabled. To enable all connections for host BLUE, enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=BLUE1A1 , BLUE1B1 , BLUE2A2 , BLUE2B2
```

To enable only two connections for host BLUE (if it is a restriction of the operating system), select two connections that use different adapters, different switches or hubs, and different controllers:

```
SET D101 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=( BLUE1A1 , BLUE2B2 )
```

or

```
SET D101 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=( BLUE1B1 , BLUE2A2 )
```

If the storage subsystem has more than one host connection, the access path must be specified carefully to avoid giving undesirable host connections access to the unit. The default condition for a unit is that access paths to all host connections are enabled. To restrict host access to a set of host connections, specify *DISABLE_ACCESS_PATH=ALL* when the unit is added, then use the SET *unit* command to specify the set of host connections that are to have access to the unit.

Enabling the access path to a particular host connection does not override previously enabled access paths. All access paths previously enabled are still valid; the new host connection is simply added to the list of connections that can access the unit.

IMPORTANT: The procedure of restricting access by enabling all access paths then disabling selected paths is **not recommended** because of the potential data/security breach that occurs when a new host connection is added.

Restricting Host Access by Offsets

Offsets establish the start of the range of units that a host connection can access. However, depending on the operating system, hosts that have lower offsets may be able to access the units in the specified range.

NOTE: All host connections to the same host computer must be set to the same offset.

For example:

In Figure 1–14, assume all host connections initially have the default offset of 0. Giving all connections access to host BLUE, an offset of 120 will present unit D120 to host BLUE as LUN 0. Enter the following commands:

```
SET BLUE1A1 UNIT_OFFSET=120
SET BLUE1B1 UNIT_OFFSET=120
SET BLUE2A2 UNIT_OFFSET=120
SET BLUE2B2 UNIT_OFFSET=120
```

Host BLUE cannot see units lower than its offset, so it cannot access any other units. However, the other two hosts can still access D120 as LUN 20 if their operating system permits. To restrict access of D120 to only host BLUE, enable only host BLUE's access, as follows:

```
SET D120 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(BLUE1A1, BLUE1B1, BLUE12A2, BLUE2B2)
```

NOTE: StorageWorks recommends that you always provide access to only specific connections. This way, if new connections are added, they will not have automatic access to all units.

NOTE: See “Restricting Host Access by Disabling Access Paths,” page 1–23.

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

In multiple-bus failover mode, each of the host ports has its own port ID:

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

- Controller A, port 1—worldwide name + 3, for example 5000-1FE1-FF0C-EE03
- Controller A, port 2—worldwide name + 4, for example 5000-1FE1-FF0C-EE04

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 1–15 shows the placement of the worldwide name label for the Model 2200 enclosure, and Figure 1–16 for the BA370 enclosure.




Figure 1–15: Placement of the worldwide name label on the Model 2200 enclosure




Figure 1–16: Placement of the worldwide name label on the BA370 enclosure



CAUTION: Each subsystem has its own unique worldwide name (node ID). If you attempt to set the subsystem worldwide name to a name other than the one that came with the subsystem, the data on the subsystem will not be accessible. Never set two subsystems to the same worldwide name, or data corruption will occur.

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.

Planning Storage Configurations

This chapter provides information to help you plan the storage configuration of your subsystem. Storage containers are individual disk drives (JBOD), storageset types (mirrorsets, stripesets, and so on), and/or partitioned drives. Use the guidelines found in this section to plan the various types of storage containers needed.

This chapter also focuses on the required design and implementation aspects, such as addressing conventions, configuration rules, creating storage profiles, and creating storage maps.

The following storage configuration information can be found in this chapter:

- “Where to Start,” page 2–2
- “Determining Storage Requirements,” page 2–3
- “Configuration Rules for the Controller,” page 2–3
- “Addressing Conventions for Device PTL,” page 2–4
- “Choosing a Container Type,” page 2–14
- “Creating a Storageset Profile,” page 2–16
- “Planning Considerations for Storageset,” page 2–18
- “Changing Characteristics through Switches,” page 2–27
- “Specifying Storageset and Partition Switches,” page 2–28
- “Specifying Initialization Switches,” page 2–29
- “Specifying Unit Switches,” page 2–33
- “Creating Storage Maps,” page 2–33

Refer to Chapter 3 for instructions on how to prepare your host computer to accommodate the HSG80 controller storage subsystem.

Where to Start

The following procedure outlines the steps to follow when planning your storage configuration.

See Appendix A to locate the blank templates for keeping track of the containers being configured.

1. Determine your storage requirements. Use the questions in “Determining Storage Requirements,” page 2–3, to help you.
2. Review configuration rules. See “Configuration Rules for the Controller,” page 2–3.
3. Familiarize yourself with the current physical layout of the devices and their addressing scheme. See “Addressing Conventions for Device PTL,” page 2–4.
4. Choose the type of storage containers you need to use in your subsystem. See “Choosing a Container Type,” page 2–14, for a comparison and description of each type of storageset.
5. Create a storageset profile (described in “Creating a Storageset Profile,” page 2–16). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:
 - “Planning Considerations for Storageset,” page 2–18
 - “Mirrorset Planning Considerations,” page 2–21
 - “RAIDset Planning Considerations,” page 2–22
 - “Partition Planning Considerations,” page 2–26
 - “Striped Mirrorset Planning Considerations,” page 2–24
6. Decide which switches you need for your subsystem. General information on switches is detailed in “Specifying Storageset and Partition Switches,” page 2–28.
 - Determine the unit switches you want for your units (“Specifying Unit Switches,” page 2–33).
 - Determine the initialization switches you want for your planned storage containers (“Specifying Initialization Switches,” page 2–29).
7. Create a storage map (“Creating Storage Maps,” page 2–33).
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 *StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide*.

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?

Configuration Rules for the Controller

The following list defines maximum configuration rules for the controller:

- 128 visible LUNs/200 assignable unit numbers
 - In SCSI-2 mode, if the CCL is enabled, the result is 127 visible LUNs and one CCL.
 - In SCSI-3 mode, if the CCL is enabled, the result is 126 visible LUNs and two CCLs.
- 1.024 TB storageset size
- 96 host connections
- 84 physical devices
- 20 RAID 3/5 storagesets
- 30 RAID 3/5 and RAID 1 storagesets (see note)
- 45 RAID 3/5, RAID 1, and RAID 0 storagesets (see note)

NOTE: For the previous two storageset configurations, this is a combined maximum, limited to no more than 20 RAID 3/5 storagesets in the individual combination.

- 8 partitions of a storageset or individual disk
- 6 physical devices per RAID 1 storageset (mirrorset)
- 14 physical devices per RAID 3/5 storageset (RAIDset)
- 24 physical devices per RAID 0 storageset (stripeset)
- 45 physical devices per RAID 0+1 storageset (striped mirrorset)

Addressing Conventions for Device PTL

The HSG80 controller has six SCSI device ports, each of which connects to a SCSI bus. In dual-controller subsystems, these device buses are shared between the two controllers. (The *StorageWorks* Command Console calls the device ports “channels.”) The standard BA370 enclosure provides a maximum of four SCSI target identifications (ID) for each device port. If more target IDs are needed, expansion enclosures can be added to the subsystem. For an example of how units are mapped to physical disk drives, see Figure 2–1.




Figure 2–1: Mapping a unit to physical disk drives

The HSG80 controller identifies devices based on a Port-Target-LUN (PTL) numbering scheme, shown in Figure 2–2. The physical location of a device in its enclosure determines its PTL.

- P—Designates the controller's SCSI device port number (1 through 6).
- T—Designates the target ID number of the device. Valid target ID numbers for a single-controller configuration and dual-redundant controller configuration are 0 - 3 and 8 - 15, respectively. (This applies to the BA370 cabinet only.)
- L—Designates the logical unit (LUN) of the device. For disk devices the LUN is always 0.




Figure 2–2: PTL naming convention

The controller can either operate with a BA370 enclosure or with a Model 2200 controller enclosure combined with Model 4214R, Model 4254, Model 4310R, Model 4350R, Model 4314R, or Model 4354R disk enclosures.

The controller operates with BA370 enclosures that are assigned ID numbers 0, 2, and 3. These ID numbers are set through the PVA module. Enclosure ID number 1, which assigns devices to targets 4 through 7, is not supported. Figure 2–3 shows how data is laid out on disks in an extended configuration.




Figure 2–3: How data is laid out on disks in BA370 enclosure configuration

Examples - Model 2200 Storage Maps, PTL Addressing

The Model 2200 controller enclosure can be combined with the following:

- **Model 4214R disk enclosure** — Ultra2 SCSI with 14 drive bays, single-bus I/O module.
- **Model 4254 disk enclosure** — Ultra2 SCSI with 14 drive bays, dual-bus I/O module.

NOTE: The Model 4214R uses the same storage maps as the Model 4314R, and the Model 4254 uses the same storage maps as the Model 4354R disk enclosures.

- **Model 4310R disk enclosure** — Ultra3 SCSI with 10 drive bays, single-bus I/O module. Table 2–1 shows the addresses for each device in a six-shelf, single-bus configuration. A maximum of six Model 4310R disk enclosures can be used with each Model 2200 controller enclosure.

NOTE: The storage map for the Model 4310R reflects the disk enclosure's 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.

- **Model 4350R disk enclosure** — Ultra3 SCSI with 10 drive bays, single-bus I/O module. Table 2–2 shows the addresses for each device in a three-shelf, single-bus configuration. A maximum of three Model 4350R disk enclosures can be used with each Model 2200 controller enclosure.
- **Model 4314R disk enclosure** — Ultra3 SCSI with 14 drive bays, single-bus I/O module. Table 2–3 shows the addresses for each device in a six-shelf, single-bus configuration. A maximum of six Model 4314R disk enclosures can be used with each Model 2200 controller enclosure.

NOTE: The storage map for the Model 4314R reflects the disk enclosure's 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.

- **Model 4354R disk enclosure** — Ultra3 SCSI with 14 drive bays, dual-bus I/O module. Table 2–4 shows the addresses for each device in a three-shelf, dual-bus configuration. A maximum of three Model 4354R disk enclosures can be used with each Model 2200 controller enclosure.

NOTE: Appendix A contains storageset profiles you can copy and use to create your own system profiles. It also contains an enclosure template you can use to help you keep track of the location of devices and storagesets in your shelves.

Table 2–1: PTL addressing, single-bus configuration, six Model 4310R enclosures

Model 4310R Disk Enclosure Shelf 6 (single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200

Model 4310R Disk Enclosure Shelf 5 (single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200

Model 4310R Disk Enclosure Shelf 4 (single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

Model 4310R Disk Enclosure Shelf 1 (single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk11000	Disk11100	Disk11200

Model 4310R Disk Enclosure Shelf 2 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk21000	Disk21100	Disk21200

Model 4310R Disk Enclosure Shelf 3 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200

Table 2–2: PTL addressing, dual-bus configuration, three Model 4350R enclosures

Model 4350R Disk Enclosure Shelf 1 (single-bus)

	SCSI Bus A					SCSI Bus B				
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400

Model 4350R Disk Enclosure Shelf 2 (single-bus)

	SCSI Bus A					SCSI Bus B				
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400

Model 4350R Disk Enclosure Shelf 3 (single-bus)

	SCSI Bus A					SCSI Bus B				
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400

Table 2–3: PTL addressing, single-bus configuration, six Model 4314R enclosures

Model 4314R Disk Enclosure Shelf 6 (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	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	Disk61400	Disk61500

Model 4314R Disk Enclosure Shelf 5 (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	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	Disk51400	Disk51500

Model 4314R Disk Enclosure Shelf 4 (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	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300	Disk41400	Disk41500

Model 4314R 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 4314R 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

Model 4314R Disk Enclosure Shelf 3 (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	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

Table 2–4: PTL addressing, dual-bus configuration, three Model 4354A enclosures.

Model 4354R Disk Enclosure Shelf 1 (dual-bus)

	SCSI Bus A							SCSI 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

Model 4354R Disk Enclosure Shelf 2 (dual-bus)

	SCSI Bus A							SCSI 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	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

Model 4354R Disk Enclosure Shelf 3 (dual-bus)

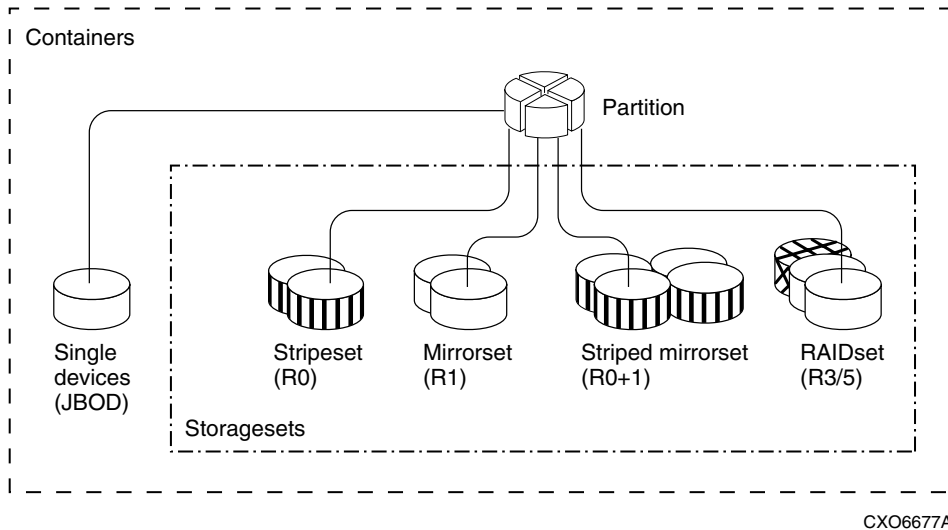
	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Choosing a Container Type

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

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in Figure 2–4. 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 2–4: Storage container types

Table 2–5 compares the different kinds of containers to help you determine which ones satisfy your requirements.

Table 2–5: Comparison of Container Types

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
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 StorageSet Profile

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

For an example of a storage set profile, see Table 2–6. This table contains blank profiles that you can copy and use to record the details for your storage sets. Use the information in this chapter to help you make decisions when creating storage set profiles.

Table 2-6: Example of 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 R1. **Disk Drives** D10300, D20300, D10400, D20400

Unit Number D101

Partitions:

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

RAIDset Switches:

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

Mirrorset Switches:

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

Initialize Switches:

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

Unit Switches:

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

Planning Considerations for StorageSet

This section contains the guidelines for choosing the storageSet type needed for your subsystem:

- “Stripeset Planning Considerations,” page 2–18
- “Mirrorset Planning Considerations,” page 2–21
- “RAIDset Planning Considerations,” page 2–22
- “Striped Mirrorset Planning Considerations,” page 2–24
- “StorageSet Expansion Considerations,” page 2–26
- “Partition Planning Considerations,” page 2–26

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 Disk 10000, Disk 20000, and Disk 10100, the first chunk of an I/O request is written to Disk 10000, the second to Disk 20000, the third to Disk 10100, the fourth to Disk 10000, until all of the data has been written to the drives (Figure 2–5).




Figure 2–5: 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 2–30, for information about setting the chunk size). Figure 2–6 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 2–6: 3-member RAID 0 stripeset (example 2)

Keep the following points in mind as you plan your 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 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. Stripsets 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 the load and provide multiple paths.
- Stripsets may contain between two and 24 members.
- If you plan to use mirror members to replace failing drives, then create the original stripeset as a stripeset of 1-member mirrorsets.
- Stripsets 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.
- Stripsets 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 six I/O device ports.


Mirrorset Planning Considerations

Mirrorsets (RAID 1) use redundancy to ensure availability, as illustrated in Figure 2–7. 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 2–8 shows a second example of a Mirrorset.



CXO7288A

Figure 2–7: Mirrorsets maintain two copies of the same data



CXO4594B

Figure 2–8: Mirrorset example 2

Keep these points in mind 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 a maximum of 20 RAID 3/5 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset may contain up to 6 members. Refer to “Configuration Rules for the Controller” on page 2–3, for detailed information on maximum numbers. 30 RAID 3/5 and RAID 1 mirrorsets are permitted, however, there is limit of no more than 20 RAID 3/5 mirrorsets in such a configuration.
- Both write-back cache modules must be the same size.
- A mirrorset should only contain disk drives of the same capacity.
- Spread mirrorset members across different device ports (drive bays).
- 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 percent 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 2–9 shows an example of a RAIDset that uses five members.




Figure 2–9: 5-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. 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 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 2–30, 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 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 2–10). 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 2–11 shows a second example of a striped mirrorset using six members.




Figure 2–10: Striped mirrorset (example 1)

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

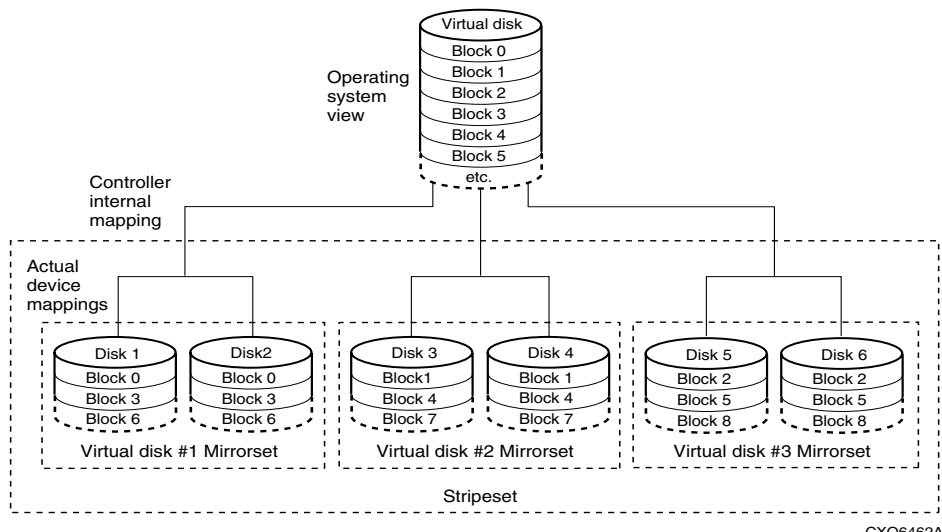


Figure 2–11: Striped mirrorset (example 2)

Plan the mirrorset members, and plan the stripeset that will contain them. Review the recommendations in “Planning Considerations for Stageset,” page 2–18, and “Mirrorset Planning Considerations,” page 2–21.

Stageset Expansion Considerations

Stageset 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 StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide describes the CLI command: ADD CONCATSETS which is used to perform concatenation.



CAUTION: Use the ADD CONCATSETS command only with host operating systems that support dynamic volume expansion. Use of this command could result in inaccessible data, if the operating system cannot handle one of its disks increasing in size.

Partition Planning Considerations

Use partitions to divide a container (stageset or individual disk drive) into smaller pieces, each of which can be presented to the host as its own storage unit. Figure 2–12 shows the conceptual effects of partitioning a single-disk container.




Figure 2–12: One example of a partitioned single-disk unit

You can create up to eight partitions per stageset (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 stageset or device. Partitions are separately addressable storage units; therefore, you can partition a single stageset 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 disks:

- Each storageset or disk drive may have 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.
- In multiple-bus failover mode, all the partitions of a particular container must be on the same controller. Partitions cannot be split across controllers.
- 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 “Creating Clones 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 also 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 SWCC during the configuration process, and SWCC automatically applies them to the storageset or device. See the SWCC online help 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 StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide contains the details of the CLI commands 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 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 data is destroyed during this procedure.

Specifying 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

Use the following types of switches to control how a RAIDset ensures 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 *StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide*.

Mirrorset Switches

Use 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 *StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide*.

Partition Switches

The following switches are available when creating a partition:

- Size
- Geometry

For details on the use of these switches refer to CREATE_PARTITION command in the *StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide*.

Specifying Initialization Switches

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

- Chunk Size (for stripesets and RAIDsets only)
- Save Configuration

- Destroy/Nodestroy
- Geometry

Each of these switches is described in the following sections.

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

Chunk Size

With ACS software, a parameter for chunk size (`chunksize=default` or `n`) on some storagesets can be set. However, unit performance may be negatively impacted if a non-default value is selected as the `chunksize`.


If a non-default chunk size has been calculated, verify that the chunk size value is divisible by 8, with no remainder. If the value is not aligned with this rule, adjust the chunk size value upward until it divisible by 8, with no remainder.

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 number of drives is less or equal to 9, then `chunk size = 256`. If the number of drives is greater than 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, separate I/O requests can be handled in parallel, which increases the request rate. This concept is shown in Figure 2–13.



CXO-5135A-MC

Figure 2–13: Large chunk size increases request rate

Large chunk sizes also tend to increase the performance of random reads and writes. StorageWorks 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 your subsystem, you must first 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), make the chunk size for RAID 0 and RAID 0+1 a small number (for example: 67 sectors). For RAID 5, make the chunk size a relatively large number (for example: 253 sectors).

Table 2–7 shows a few examples of chunk size selection.

Table 2–7: 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 Sequential Data Transfer Performance

RAID 0 and RAID 0+1 sets intended for high data transfer rates should use a relatively low chunk size (for example: 67 sectors). RAID 5 sets intended for high data rate performance should use a relatively large number (for example: 253 sectors).

Save Configuration

The *SAVE CONFIGURATION* switch is for a **single-controller configuration only**. This switch reserves an area on each of the disks for 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 areas of disks.

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

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

IMPORTANT: DO NOT use *SAVE_CONFIGURATION* in dual redundant controller installations. It is not supported and may result in unexpected controller behavior.

Destroy/Nodestroy

Specify whether to destroy or retain the user data and metadata when a disk is initialized after it has been 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 when a disk drive 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 StorageWorks HSG80 Array Controller ACS Version 8.7 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 storageset 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.

Specifying Unit Switches

Several switches control the characteristics of units. The unit switches are described under the *SET unit-number* command in the StorageWorks HSG80 Array Controller ACS Version 8.7 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 the first Chapter under the following heading:

- “Determining Connections,” page 1–13

Creating Storage Maps

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

To make a storage map, fill out the templates provided in Appendix A as you add storagesets, partitions, and JBOD disks to the configuration and assign them unit numbers. Label each disk drive in the map with the higher levels it is associated with, up to the unit level.

Using LOCATE Command to Find Devices

If you want to complete a storage map at a later time but do not remember where the disk drives and partitions are located, 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 procedure is an example of the commands 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 procedure is an example command 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
```


Example Storage Map - Model 4310R Disk Enclosure

Table 2–8 shows an example of four Model 4310R disk enclosures (single-bus I/O).

Table 2–8: Model 4310R disk enclosure, example of storage map

Model 4310R Disk Enclosure Shelf 4 (single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk40000	D101 S1 M2 Disk40100	D103 M4 Disk40200	D105 Disk40300	D107 Disk40400	D108 S3 Disk40500	D1 S4 M6 Disk40800	D2 R3 Disk41000	D3 S5 Disk41100	spare Disk41200

Model 4310R Disk Enclosure Shelf 1 (single-bus)										
Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk10000	D101 S1 M1 Disk10100	D102 M3 Disk10200	D104 S2 Disk10300	D106 R2 Disk10400	D108 S3 Disk10500	D1 S4 M5 Disk10800	D2 R3 Disk11000	D3 S5 Disk11100	D4 M7 Disk11200

Model 4310R Disk Enclosure Shelf 2 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk20000	D101 S1 M1 Disk20100	D102 M3 Disk20200	D104 S2 Disk20300	D106 R2 Disk20400	D108 S3 Disk20500	D1 S4 M5 Disk20800	D2 R3 Disk21000	D3 S5 Disk21100	D4 M7 Disk21200

Model 4310R Disk Enclosure Shelf 3 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	D100 R1 Disk30000	D101 S1 M2 Disk30100	D103 M4 Disk30200	D104 S2 Disk30300	D106 R2 Disk30400	D108 S3 Disk30500	D1 S4 M6 Disk30800	D2 R3 Disk31000	D3 S5 Disk31100	spar e Disk31200

The following explains the table in words:

- Unit D100 is a 4-member RAID 3/5 storageset named R1. R1 consists of Disk10000, Disk20000, Disk30000, and Disk40000.
- Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirrorset consisting of Disk10100 and Disk20100.
 - M2 is a 2-member mirrorset consisting of Disk30100 and Disk40100.
- Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10200 and Disk20200.
- Unit D103 is a 2-member mirrorset named M4. M4 consists of Disk30200 and Disk40200.

- Unit D104 is 3-member stripeset named S2. S2 consists of Disk10300, Disk20300, and Disk30300.
- Unit D105 is a single (JBOD) disk named Disk40300.
- Unit D106 is a 3-member RAID 3/5 storage set named R2. R2 consists of Disk10400, Disk20400, and Disk30400.
- Unit D107 is a single (JBOD) disk named Disk40400.
- Unit D108 is a 4-member stripeset named S3. S3 consists of Disk10500, Disk20500, Disk30500, and Disk40500.
- Unit D1 is a 2-member striped mirrorset named S4. S4 consists of M4 and M5:
 - M5 is a 2-member mirrorset consisting of Disk10800 and Disk20800.
 - M6 is a 2-member mirrorset consisting of Disk30800 and Disk40800.
- Unit D2 is a 4-member RAID 3/5 storage set named R3. R3 consists of Disk11000, Disk21000, Disk31000, and Disk41000.
- Unit D3 is a 4-member stripeset named S5. S5 consists of Disk11100, Disk21100, Disk31100, and Disk41100.
- Unit D4 is a 2-member mirrorset named M7. M7 consists of Disk11200 and Disk21200.
- Disk31200 and Disk41200 are spare set members.

Preparing the Host System

This chapter describes how to prepare your Sun Solaris host computer to accommodate the HSG80 controller storage subsystem.

The following information is included in this chapter:

- “Installing RAID Array Storage System,” page 3–1
- “Making a Physical Connection,” page 3–6
- “Verifying/Installing Required Versions,” page 3–8
- “Installing Solution Software Packages,” page 3–9
- “Preparing LUNs for Use by the FileSystem,” page 3–13
- “Creating and Tuning File Systems,” page 3–15
- “Solution Software Upgrade Procedures,” page 3–16

Refer to Chapter 4 for instructions on how to install and configure the HSG Agent. The Agent for HSG is operating system-specific and polls the storage.

Installing RAID Array Storage System



WARNING: A shock hazard exists at the backplane when the controller enclosure bays or cache module bays are empty.

Be sure the enclosures are empty, then mount the enclosures into the rack. **DO NOT** use the disk enclosure handles to lift the enclosure. The handles cannot support the weight of the enclosure. Only use these handles to position the enclosure in the mounting brackets.

Use two people to lift, align, and install any enclosure into a rack. Failure to use two people might cause personal injury and/or equipment damage.



CAUTION: Controller and disk enclosures have no power switches. Make sure the controller enclosures and disk enclosures are physically configured before turning the PDU on and connecting the power cords. Failure to do so can cause equipment damage.

1. Be sure the enclosures are empty before mounting them into the rack. If necessary, remove the following elements from the controller enclosure:
 - Environmental Monitoring Unit (EMU)
 - Power Supplies
 - External Cache Batteries (ECBs)
 - Fans

If necessary, remove the following elements from the disk enclosure:


- Power Supply/Blower Assemblies
- Disk Drives
- Environmental Monitoring Unit (EMU)
- I/O Modules

Refer to the *StorageWorks Model 2100 and 2200 Ultra SCSI Controller Enclosures User Guide* and *StorageWorks Model 4300 Family Ultra3 LVD Disk Enclosures User Guide* for further information.

2. Install brackets onto the controller enclosure and disk enclosures. Using two people, mount the enclosures into the rack. Refer to the mounting kit documentation for further information.
3. Install the elements. Install the disk drives. Make sure you install blank panels in any unused bays.

Fibre channel cabling information is shown to illustrate supported configurations. In a dual-bus disk enclosure configuration, disk enclosures 1, 2, and 3 are stacked below the controller enclosure—two SCSI Buses per enclosure (see Figure 3–1). In a single-bus disk enclosure configuration, 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—one SCSI Bus per enclosure (see Figure 3–2).


4. Connect the six VHDCI UltraSCSI bus cables between the controller and disk enclosures as shown in Figure 3–1 for a dual bus system and Figure 3–2 for a single bus system. Note that the supported cable lengths are 1, 2, 3, 5, and 10 meters.
5. Connect the AC power cords from the appropriate rack AC outlets to the controller and disk enclosures.



CX07383A

- ❶ SCSI Bus 1 Cable
- ❷ SCSI Bus 2 Cable
- ❸ SCSI Bus 3 Cable
- ❹ SCSI Bus 4 Cable
- ❺ SCSI Bus 5 Cable
- ❻ SCSI Bus 6 Cable
- ❼ AC Power Inputs
- ❽ Fibre Channel Ports

Figure 3–1: Dual-Bus Enterprise Storage RAID Array Storage System



CXO7382A

- | | |
|--------------------|-----------------------|
| ❶ SCSI Bus 1 Cable | ❷ AC Power Inputs |
| ❸ SCSI Bus 2 Cable | ❸ Fibre Channel Ports |
| ❹ SCSI Bus 3 Cable | |
| ❺ SCSI Bus 4 Cable | |
| ❻ SCSI Bus 5 Cable | |
| ❼ SCSI Bus 6 Cable | |

Figure 3–2: Single-Bus Enterprise Storage RAID Array Storage System

Making a Physical Connection

To attach a 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.

Preparing to Install Host Bus Adapter

Before installing the host bus adapter, perform the following steps:

1. Perform a complete backup of the entire system.
2. Shut down the computer system or perform a hot addition of the adapter based upon directions for that server.

Installing Host Bus Adapter

To make a physical connection, first install a host bus adapter.



CAUTION: Protect the host bus adapter board from electrostatic discharge by wearing an ESD wrist strap. DO NOT remove the board from the antistatic cover until you are ready to install it.

You need the following items to begin:

- Host bus adapter board
- The computer hardware manual
- Appropriate tools to service your computer

The host bus adapter board plugs into a standard PCI or Sbus slot in the host computer. Refer to the system manual for instructions on installing expansion boards.

NOTE: Take note of the worldwide name (WWN) of each adapter. Do not power on anything yet. For the FC switches to autoconfigure, power on equipment in a certain sequence. Also, the controllers in the subsystem are not yet configured for compatibility with Sun Solaris.

The *StorageWorks* Solutions Software consists of the packages as shown in Table 3–1.

Table 3–1: StorageWorks Solution Software Packages

Package	Description
CPQfcraid	Agent software and system updates required for RAID system operation. This package should always be installed.
CPQfcaw	HBA driver for the CPQ/JNI FC64-1063 Fibre Channel 64-bit Sbus adapter. This package should be loaded when a CPQ/JNI SWSA4-SC adapter is installed in the server.
CPQfca	HBA driver for the CPQ/JNI FC-1063 Fibre Channel 32-bit Sbus adapter. This package should be loaded when a CPQ/JNI SWSA4-SB adapter is installed in the server.
CPQfcaPCI	HBA driver for the CPQ/JNI FCI-1063 32-bit PCI bus adapter. This package should be loaded when a CPQ/JNI SWSA4-PC adapter is installed in the server.
CPQptisp	HBA driver for the CPQ/PTI SBS450 SCSI adapter. This should only be used when the Sun server is interfaced with existing SCSI RAID systems (RA410, RA450, RA7000), that are attached to a CPQ/PTI SBS-450 SCSI adapter. NOTE: This driver does not support Solaris in 64-bit mode. To use the SBS450 SCSI adapter under 64-bit Solaris, set the adapter to Sun emulation mode.
QLA2200-1	QLogic cPCI and SBUS drivers for Solaris 2.6.
QLA2200-2	QLogic cPCI and SBUS drivers for Solaris 7.
QLA2200-3	QLogic cPCI and SBUS drivers for Solaris 8.
QLA2300-1	QLogic PCI driver for Solaris 2.6.
QLA2300-2	QLogic PCI driver for Solaris 7.
QLA2300-3	QLogic PCI driver for Solaris 8.

To install the System Manager packages onto the hard disk, perform the following steps:

1. Mount the CD-ROM.
2. Load the Fibre Channel adapter driver packages.
3. Load the CPQptisp package (optional, only required to support legacy RAID systems attached to an SBS-450 SCSI adapter).

4. Load the CPQfcraid package.
5. Configure WWPN/target bindings (edit your sd.conf file).
6. Restart your computer.
7. Label (format) LUNs and create filesystems if required.
8. Configure the Agent to ensure subsystems are monitored. See Chapter 4 for configuration instructions.

Verifying/Installing Required Versions

To prepare your Enterprise Storage RAID Array for the RAIDManager software installation, complete the following steps:

1. Back up your entire system according to your normal procedure.
2. Select a system user with superuser privileges (for example: root) as the RAID administrator.
3. Login as the RAID administrator.
4. To find a filesystem with at least 500 KB free space, type:

```
# df -k
```
5. Choose a directory in which to install the SWCC software. The default is */opt*, however you may specify any directory. You will need to know this when running the installation script.

NOTE: The “base directory” (referred to when installing Agent) will have a “steam” subdirectory created under it. This is the directory where the Agent files will be installed.

You will need approximately 500 KB of permanent space on your hard disk to install and use Agent. If you enable logging, some additional disk space will be used by the logfiles in the */var/adm/steam* directory.

Choosing Automatic or Manual Install Modes

The installation script can be run in either *automatic* or *manual* mode. The Installation Manager examines the following packages and drivers, which are installed on your system, for compatibility with the current release of the RAID Manager software.

- Sun Solaris operating system version.
- Existence of required patches (and patch versions) on your system

- Prior versions of Solution Software packages (DECswm410, DECswm450, CPQhsz70, CPQHSG).
- Prior versions of driver packages (DECptisp, CPQfca, CPQfcaw, CPQfcaPCI).
- Incompatible versions of driver packages (PTIptisp, JNIfca, JNIfcaw, JNIfcaPCI).
- See “Solution Software Upgrade Procedures,” page 3–16 for upgrade procedures.

The key differences between Automatic and Manual modes is the way in which problems are fixed.

Automatic mode, as the name suggests, operates entirely automatically. All old or incompatible packages are removed, and the correct new packages are loaded, depending on the FC adapters installed in your system. Therefore, it is important that any FC adapters you want to use are installed prior to running the Installation Manager. You will be prompted for verification prior to any change being made to your system.

Manual mode operates manually. The Installation Manager will identify what steps to take, but the user is responsible for actually entering the correct commands. The Installation manager will also identify incompatibilities between installed packages and the new packages, and it will recommend the corrective steps. The user is responsible for manually removing old packages and installing the new packages. Be aware that it may take several iterations of running the Installation Manager to resolve all problems.

Automatic mode is the preferred installation method for most users, as it is the simplest method. Advanced users may want to use manual mode, as it allows the user ability to customize such things as supported adapters and drivers.

Installing Solution Software Packages

This section includes the following topics:

- Mounting the CD-ROM
- Installing with Automatic Mode
- Installing with Manual Mode

Installing with Automatic Mode

To install the RAID Manager software using automatic mode:

1. To start the Installation Manager, type:

```
#./install_stgwns
```

The Installation Manager asks you if you want to use Automatic or Manual mode.

2. Enter “A” to select Automatic mode.

Installation Manager verifies that you are running a supported version of Sun Solaris and that you have the correct patches loaded.

If your Sun Solaris set up is supported, Installation Manager scans your system for supported adapters, and displays a list of compatible adapter drivers.

Installation Manager prompts you to remove any old or incompatible adapter drivers.

3. Remove old and incompatible adapter drivers. Follow the on-screen prompts, always accepting the default responses.

4. Installation Manager prompts you for the package base directory. Enter the package base directory. This directory has the “steam” subdirectory, containing the RAID Manager files, created under it. The default directory is `/opt`.

- Loop mode: Each HBA is configured for LUNs 0 through 15 at targets 64 and 65 (corresponds with ALPAs 72 and 71).

- Fabric mode: You will be prompted for the SCSI targets to bind through for each HBA and the Worldwide Port Names (WWPNs) to attach to.

Installation Manager prompts you to rescan your system for supported adapter drivers.

5. If any adapter drivers were loaded in previous steps, answer “Y” to allow the Installation Manager to rescan your system for supported adapters. Since the rescan will probe every possible I/O slot for new hardware, it can take up to 5 minutes to complete on a large system. During this time, a heavily loaded system may appear to hang. This is normal.

6. Edit your `sd.conf` file. This edit can be done automatically or manually.

- Choose “M” to manually edit your `sd.conf` file. You can choose what FC adapters and what targets will be supported. By default, LUNs 0 through 15 are configured for each target that you select.

- Choose “**A**” to automatically edit the `sd.conf` file. Support is added for all supported Fibre Channel adapters, using the default SCSI targets.
- Choose “**Q**” to exit the Installation Manager.

After all the files have been loaded and edited, Installation Manager exits.

7. After RAID Manager is loaded, you are prompted to reboot your computer.
8. Change to the top level directory, by entering the following command:

```
# cd /
```

9. Enter the following command to shut down the computer:

```
# /etc/shutdown -y -i0 -g0
```

10. Wait for the OK prompt to appear, then enter the following (`boot_device` is the name of your system disk):

```
# boot boot_device -r
```

11. Wait for your computer to perform a reconfiguration boot.

NOTE: While your computer is booting, you may see the following message:

WARNING:

```
/sbus@2,0/fcaw@1,0/sd@40,0 (sd32): corrupt label - wrong magic number.
```

This message appeared because a LUN you configured earlier is not yet labeled.

Installing with Manual Mode

To install the RAID Manager software using manual mode:

1. To start the Installation Manager, enter the following:

```
# ./install_stgwks
```

Installation Manager asks you if you want to use Automatic or Manual mode. Enter “**M**” to select Manual mode.

- Installation Manager will verify that the computer is running a supported version of Solaris, and that it has the correct patches installed.
 - Installation Manager scans your system for old or incompatible drivers and asks if you would like to remove them.
2. Installation Manager scans your system for supported adapters. When the scan is complete, Installation Manager displays a list of supported adapters installed in your system. After displaying a list of installed adapters, a menu appears listing the available packages.

3. Follow the on-screen prompts to load the correct adapter driver package for the adapters installed in your system. After the adapter driver package has successfully loaded, you are returned to the Installation Manager menu.
4. If you have multiple adapters installed in your system, select additional adapter driver packages until you have successfully loaded a driver package for each adapter you will be using.
5. If you have loaded all adapter driver packages you need, install the RAID Manager (CPQfcraid) software package.
6. The default directory is `/opt`. This directory will have a “steam” subdirectory, containing the RAID Manager files, created under it.
7. If any adapter drivers were loaded in the previous steps, answer **Y** to allow the Installation Manager to rescan your system.
8. Edit your `sd.conf` file. This edit can be done automatically or manually.
 - Choose “**M**” to manually edit your `sd.conf` file. You can choose what FC adapters and what targets will be supported. By default, LUNs 0 through 15 are configured for each target that you select.
 - Choose “**A**” to automatically edit your `sd.conf` file. Support will be added for ALL supported FC adapters, using the default targets. These defaults are:
 - Loop mode: Each HBA is configured for LUNs 0 through 15 at targets 64 and 65 (corresponds with ALPAs 72 and 71).
 - Fabric mode: You will be prompted for the SCSI targets to bind through for each HBA, and the Worldwide Port Names (WWPNs) to attach to.
 - Choose “**Q**” to exit the Installation Manager.

After all files have been loaded and edited, Installation Manager exits. Installation Manager rescans your system for supported HBA drivers.

NOTE: If you chose to manually install QLogic drivers, you must type Q (Quit) in this step without configuring your storage at this time. Reboot the system. When the system comes back up, use the `/opt/steam/bin/config.sh` program, option 20 (Add/Change Adapters) to configure your storage. This will require a second reboot once your devices are configured. See the section on “The `/opt/steam/bin/config.sh` Utility” on page 5–34, for more information on the `config.sh` utility.

9. After RAID Manager is loaded, you will be prompted to reboot your computer system.

10. Change to the top level directory, by entering the following:

```
# cd /
```

11. To shut down the computer, enter the following:

```
# /etc/shutdown -y -i0 -g0
```

12. Wait for the “OK” console prompt, then enter the following (boot_device is the name of your system disk):

```
boot boot_device -r
```

13. Wait for your system to perform a reconfiguration boot.

NOTE: During boot, you may see the following message:

WARNING:

```
/sbus@2,0/fcaw@1,0/sd@40,0 (sd32): corrupt label - wrong magic number.
```

This is because a LUN you configured earlier is not yet labeled.

Preparing LUNs for Use by the FileSystem

Each logical unit number (LUN) created on the Enterprise Storage RAID Array appears as a SCSI hard disk to the host. Therefore, it must be labeled before it can be used and, in most instances, a new file system must be created.

Labeling LUNs

A LUN is labeled using the */etc/format* utility. The label contains information about the LUN such as controller-type, geometry, and partitions. More details about the use of the *fdisk* utility may be found in the *Sun 2.x Adding and Maintaining Peripherals Manual*.

Labeling a LUN Automatically Using the Format Utility

Sun Solaris releases 2.6 and later can automatically generate a label for SCSI disks. To access this capability while in the *format* utility, select the desired disk, then select **0** from the Type menu.

Example:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
0. c0t3d0 <SUN0424 cyl 1151 alt 2 hd 9 sec 80>/.../sd@3,0
1. c1t65d0 <DEC-HSG80-AX-V86L cyl 12841 alt 2 hd 14 sec57>/sd@0,0
Specify disk (enter its number): 1
```

```
selecting c1t0d0
[disk formatted]
FORMAT MENU:
  disk- select a disk
  type- select (define) a disk type
  partition - select (define) a partition table
  current   - describe the current disk
  format    - format and analyze the disk
  repair    - repair a defective sector
  label     - write label to the disk
  analyze   - surface analysis
  defect    - defect list management
  backup    - search for backup labels
  verify    - read and display labels
  save      - save new disk/partition definitions
  inquiry   - show vendor, product and revision
  volname   - set 8-character volume name
  quit

format> type
AVAILABLE DRIVE TYPES:
0. Auto configure
1. Quantum ProDrive 80S
2. Quantum ProDrive 105S
3. CDC Wren IV 94171-344
1. SUN0104
5. SUN0207
6. SUN0327
7. SUN0340
8. SUN0424
9. SUN0535
10. SUN0669
11. SUN1.0G
12. SUN1.05
13. SUN1.3G
11. SUN2.1G
15. other

Specify disk type (enter its number)[19]: 0
c1t0d0: configured with capacity of 1.89GB
<DEC-HSG80-AX-V86L cyl 12841 alt 2 hd 14 sec 57>
selecting c1t0d0
[disk formatted]
format> label
Ready to label disk, continue? yes
format> quit
```

Note that the default partition layout when using the Automatic mode is of the form,

```
partition 0: ~128MB
partition 1: ~128MB
partition 2: entire disk
partition 6: partition 2 - (partition 0 + partition 1)
partitions 3,4,5,7: empty
```

NOTE: You may modify the partition table by entering “partition” at the format menu prompt level. Remember to label the disk again after any changes to the partition table are made.

Format places an entry for the RAID Array in the */etc/format.dat* file and the label command stores the information on the RAID Array. More details about the creation of *format.dat* entries may be found in the Sun *Solaris 2.x Adding and Maintaining Peripherals Manual*.

NOTE: Although the format command will create an HSG80 disk entry in the *format.dat* file, do not try to reuse this definition. Always use type **0** to allow Solaris to retrieve the correct logical geometry for the LUN you are formatting.

Creating and Tuning File Systems

Before the new LUN can be used by the system, a new filesystem must be created on each partition that will be mounted. Use the *newfs* command to create filesystems and the *tunefs* command to modify existing filesystems. For more information, refer to the online Help for the *newfs* and *tunefs* commands.

The following commands set *maxcontig* (the maximum number of blocks belonging to one file) to 128 (*-c* to create, *-a* to modify) and rotational delay (*-d*) to 0. These recommendations may not optimize your system performance so you may want to experiment with other values.

1. To create a new filesystem, use the following command:

```
# newfs -C128 -d0 /dev/dsk/c1t0d0s2
```

This will create a filesystem on partition 2 of the device. Partition 2 usually defaults to the entire device.

2. To modify an existing filesystem use:

```
# tunefs -a128 -d0 /dev/dsk/c1t0d0s2
```

Solution Software Upgrade Procedures

Use the following procedures for upgrades to your Solution Software. It is considered best practice to follow this order of procedures:

1. Perform backups of data prior to upgrade;
2. Verify operating system versions, upgrade operating systems to supported versions and patch levels;
3. Quiesce all I/O and unmount all file systems before proceeding;
4. Upgrade switch firmware;
5. Upgrade Solution Software
6. If installing an operating system that uses Secure Path (AIX, HP-UX, Netware, Sun or Windows), upgrade Secure Path to the latest version at this time;
7. Upgrade ACS software.

NOTE: Solely for the purpose of performing upgrades to the ACS firmware, this Solution Software Kit supports previous ACS Version 8.6. It is not recommended mixing ACS versions in the same SAN.

NOTE: For upgrades in a SAN that includes HSG80 array controllers and Enterprise Virtual 2.0 Arrays controllers, the following are required:

- ACS 8.7 and VCS 2.0 require Solution Software 8.7 (SWCC 2.5) - Retrieve Solution Software drivers from EVA kit, retrieve SWCC drivers from ACS 8.7 kit
- ACS 8.6 and VCS 2.0 require Solution Software 8.6 (SWCC 2.4)
- ACS 8.7 and ACS 8.6 require Solution Software 8.7 (SWCC 2.5)

Refer to the *StorageWorks HSG60/HSG80 Array Controller ACS Version 8.7 Maintenance and Service Guide* and the *Solution Software Release Notes* for the latest information on upgrades.

The Solutions Software Version 8.7 for Sun Solaris Kit should come with the ordered items on the FCO.

This Version 8.7 Solution Software includes Version 2.5 of the CPQfcraid RAID Manager Software, and contains an upgraded HBA driver Version 2.5.9.03 (if running ACS version 8.6 prior to upgrade).

NOTE: Rolling upgrades are not supported on Sun Solaris when using Qlogic HBAs.

Use the following steps to upgrade Solution Software on Sun Solaris systems.

1. Ensure the appropriate Sun Solaris patches are installed.

2. Stop `steamd` Agent from running. To determine if the `steamd` Agent is running, type the following:

```
# /opt/steam/bin/config.sh
```

Select option **3**. The resulting display will tell you if the Agent is running, and prompt you with the option to stop it.

3. Upgrade your Solution Software on each system as follows:
 - a. Create a temporary directory and save the `sd.conf` file by typing the following:

```
# mkdir /cpqtmp  
# cp /kernel/drv/sd.conf /cpqtmp
```

4. If you are running a Secure Path multiple-bus configuration, save your driver configuration file by typing one the following.

For 64-bit Sbus:

```
# cp /kernel/drv/fcaw.conf /cpqtmp
```

For 32-bit PCI:

```
# cp /kernel/drv/fca-pci.conf /cpqtmp
```

5. Create a temporary directory and save the `steam ini` files by typing the following:

```
# mkdir /cpqtmp/etc  
# cp /opt/steam/etc/* /cpqtmp/etc
```

- b. Run the `install_stgwks` script from the Version 8.7 Solutions Software CD-ROM.
- c. Select automatic mode **A**.
- d. Select **Y** when the script asks if it should remove the older version (2.3f) of CPQfcraid.
- e. Select **Y** when the script asks if it should remove the older version (2.5.9) of CPQfca (or CPQfcaw or CPQfcaPCI). Do this only if running ACS Version 8.6 prior to upgrade.
- f. Allow the script to install the new version (2.5) of CPQfcraid.
- g. Select **Y** when prompted with **Rescan for compatible adapters now?**
- h. Select **q** to quit from configuration menu.

- i. Restore the `sd.conf` file from temporary directory by typing:

```
# cp /cpqtmp/sd.conf /kernel/drv
```

- j. If you are running a Secure Path multiple-bus configuration, restore your driver configuration file by typing one the following.

For 64-bit Sbus:

```
# cp /cpqtmp/fcaw.conf /kernel/drv
```

For 32-bit PCI:

```
# cp /cpqtmp/fca-pci.conf /kernel/drv
```

- k. If you use mail notification, delete the `notify.disabled` file found in the `/opt/steam/etc` directory before proceeding to the next step.

- l. Restore the steam ini files from temporary directory by typing:

```
# cp /cpqtmp/etc/* /opt/steam/etc
```

- m. Ensure the storage system ACS version is Version 8.7f in

```
/opt/steam/etc/storage.ini.
```

- n. If you are updating your Secure Path software in conjunction with your Solution Software, this is the appropriate point to do so. Refer to the documentation that came with your Secure Path software for more information.
- o. Reboot your server.

Installing and Configuring HSG Agent

StorageWorks Command Console (SWCC) enables real-time configuration of the storage environment and permits the user to monitor and configure the storage connected to the HSG80 controller.

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
- “Before Installing the Agent,” page 4–5
- “Configuring the Agent,” page 4–5
- “RAID Manager Mail Messages,” page 4–9

IMPORTANT: For information on how to install the solution software packages, see Chapter 3, “Preparing the Host System”

Refer to Chapter 5 for a description of how to configure a subsystem that uses Fibre Channel fabric topology.

Why Use StorageWorks Command Console (SWCC)?

StorageWorks Command Console (SWCC) enables you to monitor and configure the storage connected to the HSG80 controller. SWCC consists of Client and Agent.

- The client provides pager notification and lets you manage your virtual disks. The client runs on Windows 2000 with Service Pack 2 and 3 and Windows NT 4.0 with Service Pack 6A or above.
- The 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.

To receive information about the devices connected to your HSG80 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 HSG80 controller and its devices. Remote and local notification can be made by email and/or SNMP messages to an SNMP monitoring program.

Table 4–1: SWCC Features and Components

Features	Agent Required?	Client Required?
Creation of RAID sets: <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
Monitor multiple subsystems at once	Yes	No
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.

Installation and Configuration Overview

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 RAID Solution Software has been loaded properly. Refer to Chapter 3.
3	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.
4	Verify that there is a LUN for communications. This can be either the CCL or a LUN that was created with the CLI. See “What is the Command Console LUN?” on page 1–11 in Chapter 1.
5	Configure the SWCC Agent with the <i>install.sh</i> or <i>Config.sh</i> utilities. See “Configuring the Agent” on page 4–5
6	Install the Client software on Windows 2000 with Service Pack 2 or 3 or Windows NT 4.0 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>StorageWorks Command Console Version 2.5, 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 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

- | | |
|---|---|
| ❶ Agent system (has the Agent software) | ❺ Hub or switch |
| ❷ TCP/IP Network | ❻ HSG80 controller and its device subsystem |
| ❸ Client system (has the Client software) | ❼ Servers |
| ❹ Fibre Channel cable | |

Before Installing the Agent

The Agent requires the minimum system requirements, as defined in the release notes for your operating system. The program is designed to operate with the Client version 2.5 on Windows 2000 or Windows NT.

Configuring the Agent

There are two ways to configure your Agent:

- *install.sh*—This utility is usually used for first-time configuration or upgrades. It guides you through the configuration of your Agent. It also identifies and upgrades any old Agent configuration files from previous RAID Manager packages. See Configuring the Agent with Install.sh (for First-time Configurations) for details.
- *Config.sh*—A menu-based administration tool that is usually used to modify your Agent configuration after installation. It provides more options than *install.sh*, so it is recommended for advanced users or users with unique configuration needs. It can also be used for first-time configuration. See Configuring the Agent with RAID Array V8.7 Configuration Menu for instructions on running *Config.sh*.

Configuring the Agent with Install.sh (for First-time Configurations)

For first-time configurations, run the *install.sh* script from the \$basedir/steam/bin directory (usually /opt/steam/bin). This script will guide you through the following actions:

- Adding a subsystem entry
 - Adding a Client entry
 - Creating a password
 - Setting up email notification
 - Starting the Agent
1. To change directories, enter the following command:

```
# cd /opt/steam/bin
```

2. To run the installation, enter the following command:

```
# ./install.sh
```

Follow the on-screen prompts to set up your Agent files. If *install.sh* finds backed up Agent configuration files, you will be asked if you want to upgrade them.

Configuring the Agent with RAID Array V8.7 Configuration Menu

To view the RAID Array V8.7 Configuration Menu, run the *Config.sh* script from the \$basedir/steam/bin directory (usually /opt/steam/bin).

1. Enter the following command:

```
# cd /opt/steam/bin
```

2. Enter the following command:

```
# ./config.sh
```

The Configuration Menu shown in Figure 4–2 displays.

---- RAID Array V8.7 Configuration Menu ----	
Agent Admin Options ----- 1) Change Agent Password 2) Change SNMP Enterprise OID 3) Start/Stop the Agent 4) Toggle Agent Startup on System Boot 5) Uninstall Agent Agent Notification Options ----- 6) Toggle Error Log Notification 7) Toggle Mail Notification 8) View Mail Notification List 9) Add User to Mail Notification List 10) Remove User from Mail Notification List 11) Modify Mail Notification List Exit: ----- q) Quit Enter Selection:	Storage Subsystem Options ----- 12) View Subsystems 13) Add a Subsystem 14) Remove a Subsystem 15) Modify a Subsystem Client Options ----- 16) View Clients 17) Add a Client 18) Remove a Client 19) Modify a Client Other Options ----- 20) Add/Change Adapters 21) High Avail. Options

Figure 4–2: RAID Array V8.7 Configuration Menu

IMPORTANT: When you change the Agent configuration, you must restart the Agent for your changes to take effect.

Choosing a Password

Select option **1** from the Agent Admin Options group to set a password to protect your subsystems from unauthorized access. Any Client with configuration privileges will be asked for this password when attempting to configure the subsystem.

Adding a Subsystem Entry

Any storageset belonging to the subsystem can be used to add a subsystem entry, but be careful not to delete the LUN from the subsystem when reconfiguring, as this breaks the communication link to the Agent for the entire subsystem. Add the subsystems you want the Agent to monitor by performing the following procedure, starting at the RAID Array Configuration Menu.

1. In the Storage Subsystem Options group, select option **12**, *View Subsystems*.

An empty table appears. Before starting the Agent, add at least one subsystem with which you want to communicate.

2. Select option **13**, *Add a Subsystem*.

The subsystem name you enter is arbitrary, but use only lowercase characters to specify it. Associate the name with a subsystem by picking a LUN name. For example, *c1t65d0*, which the format command output appears as belonging to the RAID system. This is evident from the controller number and SCSI target address used.

Append s2 to the device name (*cxydz*) when entering the character special file name, as a partition must be specified. Check that this partition exists by entering format partition and format print commands.

3. Enter a monitoring interval; for example, 30 seconds, and press **Enter** twice to return to the RAID Array Configuration main menu.

Adding a Client System Entry

1. From the Client Options group, select option **16** to View Clients in the authorized client list.
2. To add a client, select option **17**, *Add Client*.
3. Enter the client's network name (for example, myhost).
4. Enter an Access level code. Specify **2** if the manager/client is allowed to configure the subsystem.
5. Add an Error Notification Level (1 for TCP sockets, 2 for SNMP, or 3 for both).

Restarting the SWCC Agent

After you make any changes to the SWCC Agent configuration, the Agent daemon must be stopped and restarted. This ensures that changes to the configuration files are read by the *steamd* program. The *steamd* program is the daemon for the Agent.

Configuring the Agent within FirstWatch

If you have VERITAS FirstWatch installed, you may configure the StorageWorks Command Console Agent to run under FirstWatch. This allows FirstWatch to start the Agent on the primary host. In the case of a failover, FirstWatch will transition the Agent to the takeover host.

To configure the Agent to run with FirstWatch, you must have FirstWatch installed and configured. You must also install the Agent on both hosts and configure the Agents identically (subsystem name, notification options, and Client options). This is required so that when the Agent is transitioned to the takeover host, all access and notification remain identical.

Once the Agent is configured on both hosts, determine which host will be the primary and which will be the takeover host. On each host, start the configuration utility (see “Configuring the Agent with RAID Array V8.7 Configuration Menu,” page 4–6) and select option **21** from the *Other Options* group.

NOTE: Any future configuration changes made on one FirstWatch server must also be made on the other host.

RAID Manager Mail Messages

Mail messages sent by the RAID Manager are useful in troubleshooting subsystem problems. To receive RAID Manager messages, you must enable “email notification” within the Agent Configuration program. For more information on how to decipher mail messages from RAID Manager, refer to StorageWorks Command Console Version 2.5, User Guide.

FC Configuration Procedures

This chapter describes procedures to configure a subsystem that uses Fibre Channel (FC) 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 5–10
- “Configuring Devices,” page 5–16
- “Configuring Storage Containers,” page 5–16
- “Assigning Unit Numbers and Unit Qualifiers,” page 5–22
- “Configuration Options,” page 5–23
- “FC Considerations for Both Loop and Fabric Environments,” page 5–26

Use the command line interpreter (CLI) or *StorageWorks* Command Console (SWCC) to configure the subsystem. This chapter uses CLI to connect to the controller. To use SWCC for configuration, see the SWCC online help for assistance.


IMPORTANT: These 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 *StorageWorks HSG60/HSG80 Array Controller ACS Version 8.7 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, 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.



- ❶ Maintenance Port
- ❷ Maintenance Port Cable

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.

Establishing Connection with a SPARC System

To set up your SPARC system for connection with the HSG80 Controller, follow these steps:

1. Use the supplied serial cable and the 9 to 25 pin RS-232 adapter (P/N=12-45238-01) to connect the serial port on the SPARC system to the serial port on the RAID array controller.
2. If you use a SUN A/B Serial Splitter Cable, and/or you attach the controller to serial port A, you may need to modify the *remote* file to specify *ttya* as follows:

- a. Open the file `/etc/remote` with an editor.
- b. Locate the lines:

```
hardwire:\  
:dv:/dev/ttyb:br#9600...
```

- c. Change `ttyb` to `ttya`
 - d. Save the file.
3. Open a terminal window from the Desktop Program menu.
 4. Start the Terminal Interface Program(TIP), at the prompt type:

```
tip hardwire
```

Press **Enter**.

5. Press **Enter** again, and the CLI prompt appears, similar to the following:

```
HSG80>
```

Setting Up a Single Controller

Power On 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 port.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Verify that the computer or terminal is configured as follows:

- 9600 baud
 - 8 data bits
 - 1 stop bit
 - no parity
 - no flow control
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 Cable from the switch to the host Fibre Channel adapter |
| ❸ Host port 1 | |
| ❹ Host port 2 | 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 Controller.
- Set Time and Verify all Commands.
- 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 1–27, for the location of the sticker.

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

```
HSG80> SHOW THIS
```

```
Controller:
```

```
  HSG80 ZG80900583 Software V8.7, Hardware E11
```

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

```
  ALLOCATION_CLASS = 0
```

If the node ID is present, go to step 5.

If the node ID is all zeroes, enter 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
```

Where:*NNNN-NNNN-NNNN-NNNN* is the node ID, and *nn* is the checksum.

2. When using a controller that 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 the Delete command for each of the listed connections. When completed, no connections will be displayed.

Configure Controller Settings

5. Set the SCSI version using the following command syntax:

```
SET THIS SCSI_VERSION=SCSI-2, or  
SET THIS SCSI_VERSION=SCSI-3
```

NOTE: If SCSI-2 is selected, you must disable CCL using the command:

```
SET THIS NOCOMMAND_CONSOLE_LUN
```

6. Assign an identifier for the communication LUN (also called the command console LUN, or CCL). The CCL must have a unique identifier that is a decimal number in the range 1 to 32767, and which is different from the identifiers of all units. Use the following syntax:

```
SET THIS IDENTIFIER=N
```

Identifier must be unique among all the controllers attached to the fabric within the specified allocation class.

7. 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

- Restart the controller, using the following command:

```
RESTART THIS
```

Set Time and Verify all Commands

- Set the time on the controller by entering the following syntax:

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

- 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 press the **Enter** key.

- Set up any additional optional controller settings, such as changing the CLI prompt. See the *SET THIS CONTROLLER/OTHER CONTROLLER* command in the StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide for the format of optional settings.
- 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 sample is a result of a SHOW THIS command, with the areas of interest in bold.

Controller:

HSG80 ZG94214134 Software V8.7, Hardware 0000

NODE_ID = 5000-1FE1-0007-9750

ALLOCATION_CLASS = 0

SCSI_VERSION = SCSI-3

Configured for dual-redundancy with ZG9421461

In dual-redundant configuration

Device Port SCSI address 7

Time: 10-Mar-2002:12:30:34

Command Console LUN is disabled

Smart Error Eject Disabled

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0007-9751

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 7D4000

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0007-9752

PORT_2_TOPOLOGY = FABRIC (standby)

Address = 210513

NOREMOTE_COPY

Cache:

512 megabyte write cache, version 0022

Cache is GOOD

No unflushed data in cache

CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

Mirrored Cache:

Not enabled

Battery:

NOUPS

FULLY CHARGED

Expires: 25-JUN-2003

.....

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

If you want to communicate with the Fibre Channel 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

6. Plug the Fibre Channel cable from the first host bus adapter into the switch. Enter the SHOW CONNECTIONS command to view the connection table:

```
SHOW CONNECTIONS
```

7. Rename the connections 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 1–13.

8. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=SUN
```

9. 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

10. Repeat step 7, 8, and 9 for each of that adapter’s host connections, or delete the unused connections from the table.
11. For each host adapter, repeat steps 6 through 10.

Verify Installation

To verify installation for your Sun Solaris host, restart your system with the -r option (restart). After the system is booted, use the format command to verify that your LUNS are accessible.

Sun Solaris restart your system with the -r option (restart). After the system is booted, use the format command to verify that your LUNS are accessible.

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
 - no flow control
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 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 5–11).

Figure 5–3 shows a controller pair with failover cabling showing one HBA per server with HSG80 controller in transparent failover mode.



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

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

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

```
SHOW THIS
```

See “Worldwide Names (Node IDs and Port IDs),” page 1–27, for the location of the sticker.

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

```
HSG80> show this
Controller:
  HSG80 ZG80900583 Software V8.7, Hardware E11
  NODE_ID           = 5000-1FE1-0001-3F00
  ALLOCATION_CLASS = 0
```

If the node ID is present, go to step 5.

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
```

Where: *NNNN-NNNN-NNNN-NNNN* is the node ID and *nn* is the checksum.

2. If the controller is not new from the factory, 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 the 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
```

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, first take both ports offline, 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

- Restart the controller, using the following command:

```
RESTART THIS
```

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

Set Time and Verify All Commands

- Set the time on the controller by entering the following syntax:

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

- 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 press **Enter**.

- Set up any additional optional controller settings, such as changing the CLI prompt. See the *SET THIS CONTROLLER/OTHER CONTROLLER* command in the StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide for the format of optional settings. Perform this step on both controllers.
- Verify that all commands have taken effect by entering the following command:

```
SHOW THIS
```

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

Controller:

HSG80 ZG94214134 Software V8.7, Hardware 0000

NODE_ID = 5000-1FE1-0007-9750

ALLOCATION_CLASS = 0

SCSI_VERSION = SCSI-3

Configured for dual-redundancy with ZG9421461

In dual-redundant configuration

Device Port SCSI address 7

Time: 10-Mar-2002:12:30:34

Command Console LUN is disabled

Smart Error Eject Disabled

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0007-9751

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 7D4000

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0007-9752

PORT_2_TOPOLOGY = FABRIC (standby)

Address = 210513

NOREMOTE_COPY

Cache:

512 megabyte write cache, version 0022

Cache is GOOD

No unflushed data in cache

CACHE_FLUSH_TIMER = DEFAULT (10 seconds)

Mirrored Cache:

Not enabled

Battery:

NOUPS

FULLY CHARGED

Expires: 25-JUN-2003

13. 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

14. Plug the FC cable from the first host adapter 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 1–13.

15. Rename the connections to something meaningful to the system and easy to remember. For example, to assign the name ANGEL1A1 to connection !NEWCON01, enter:

```
RENAME !NEWCON01 ANGEL1A1
```

StorageWorks recommends using a naming convention, see “Naming Connections,” page 1–13.

16. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=SUNAIX_CAMBEX
```

17. 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 Connection

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

Verify Installation

To verify installation for your Sun Solaris host, restart your system with the `-r` option (reconfigure boot). After the system is booted, use the `format` command to verify that your LUNs are accessible.

Configuring Devices

The disks on the device bus of the HSG80 can be configured manually or with the CONFIG utility. The CONFIG utility is easier. Invoke CONFIG with the following command:

```
RUN CONFIG
```



WARNING: It is highly recommended to use the CONFIG utility only at reduced I/O loads.

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


Configuring Storage Containers

For a technology refresher on this subject, refer to “Choosing a Container Type,” page 2–14.

In choosing a container, you choose between independent disks (JBODs) or one of several storageset types, as shown in Figure 5–4. The independent disks and the selected storageset may also be partitioned.

The following configurations are detailed in the following section:

- “Configuring a Stripeset” on page 5–17
- “Configuring a Mirrorset” on page 5–18
- “Configuring a RAIDset” on page 5–19
- “Configuring a Striped Mirrorset” on page 5–19
- “Configuring a Single-Disk Unit (JBOD)” on page 5–20
- “Configuring a Partition” on page 5–20



CXO6677A

Figure 5-4: Storage container types

Configuring a Stripese

1. Create the stripeset by adding its name to the controller's list of storagesets and by specifying the disk drives it contains. Use the following syntax:

```
ADD STRIPESET STRIPESET-NAME DISKNNNNN DISKNNNNN.....
```

2. Initialize the stripeset, specifying any desired switches:

```
INITIALIZE STRIPESET-NAME SWITCHES
```

See “Specifying Initialization Switches” on page 2-29 for a description of the initialization switches.

3. Verify the stripeset configuration:

```
SHOW STRIPESET-NAME
```

4. Assign the stripeset a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5-22.

For example:

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

```
ADD STRIPESET STRIPE1 DISK10000 DISK20000 DISK30000
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 StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide for a description of the mirrorset switches.

2. Initialize the mirrorset, specifying any desired switches:

```
INITIALIZE MIRRORSET-NAME SWITCHES
```

See “Specifying Initialization Switches” on page 2–29 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 5–22.

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 by 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 StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide for a description of the RAIDset switches.

2. Initialize the RAIDset, specifying any desired switches:

```
INITIALIZE RAIDSET-NAME SWITCH
```

NOTE: StorageWorks 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 “Specifying Initialization Switches” on page 2–29 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 5–22.

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 DISK30000
INITIALIZE RAID1
SHOW RAID1
```

Configuring a Striped Mirrorset

1. Create, but do not initialize, at least two mirrorsets.

See “Configuring a Mirrorset” on page 5–18.

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 “Specifying Initialization Switches” on page 2–29 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 5–22.

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 Unit (JBOD)

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

```
INITIALIZE DISK-NAME SWITCHES
```

See “Specifying Initialization Switches” on page 2–29 for a description of the initialization switches.

2. Verify the configuration by entering 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 5–22.

Configuring a Partition

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

```
INITIALIZE STORAGESET-NAME SWITCHES
```

or

```
INITIALIZE DISK-NAME SWITCHES
```

See “Specifying Initialization Switches” on page 2–29 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`, on the last partition only, to let the controller assign the largest free space available to the partition.

NOTE: See the `CREATE_PARTITION` command in the StorageWorks HSG80 Array Controller ACS Version 8.7 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 5–22.

For example:

The commands to create RAID1, a three-member RAIDset, then partition it into two storage units are shown below.

```
ADD RAIDSET RAID1 DISK10000 DISK20000 DISK30000
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 StorageWorks HSG80 Array Controller ACS Version 8.7 CLI Reference Guide.

Because of different SCSI versions, refer to the section “Assigning Unit Numbers Depending on SCSI_VERSION,” page 1–19. The choice for SCSI_VERSION effects how certain unit numbers and host connection offsets interact.

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 1–21

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, according to the following syntax:

```
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 syntax:

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

NOTE: It is suggested that the prompt name reflect some information about the controllers. For example, if the subsystem is the third one in a lab, name the top controller prompt, LAB3A and the bottom controller, LAB3B.

Mirroring cache

To specify mirrored cache, use the following syntax:

```
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, use the following syntax:

```
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, use the following syntax:

```
ADD SPARESET DISKNNNNN
```

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

For example:

The following example shows the syntax 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.

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

```
SHOW SPARESET
```

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

```
DELETE SPARESET DISKNNNNN
```

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, you must use the command: `DELETE DISKNNNNN`.

3. Verify the contents of the spareset by entering 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, use the following command:

```
SET FAILEDSET AUTOSPARE
```

To disable autospare, use 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 the disk drive is configured into a storageset. Therefore, the presence of metadata indicates that 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 nontransportable and initialize it.)

Deleting a Storageset

NOTE: If the storageset you are deleting is partitioned, you must 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 profile when changing 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
```

NOTE: FULL is not required when showing a particular device. It is used when showing all devices, for example, SHOW DEVICES 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. These switches are integral to the formatting and can only be changed by reinitializing the storageset. Initializing a storageset 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
```

FC Considerations for Both Loop and Fabric Environments

In a mixed environment of Fibre Channel loop and Fibre Channel switched access, there are specific differences in transport and access to storage subsystems—either on the loop or in the fabric. This chapter describes the differences and configuration processes for Sun servers and adapters for loop and switch support. It is divided into the following sections:

- “General FC Environmental Considerations,” page 5–27
- “Configurable Objects,” page 5–27
- “Fibre Channel Access/Transport Paths,” page 5–31

- “Implementation of Loop and Fabric Bindings,” page 5–31
- “Configuration Procedures,” page 5–40

General FC Environmental Considerations

Fibre channel (FC) configurations include servers, storage (RAID Arrays), hubs, switches, and repeaters. There are many options for configuring a storage area network.

In the simple cases, a configuration can easily be understood and investigated. In more complex configurations, with Fibre Channel’s ability to support transmission lengths in terms of miles/kilometers, the problems of design, installation, testing and documentation require greater discipline and control.

Configurable Objects

This section describes how to configure the Sun server and the storage (RAID Array) components so that given a path with a loop or switch configuration, the Sun server and storage will be able to communicate successfully. This section describes how to configure the following objects:

- Servers
- Storage
- Hubs
- Switches

Servers

Servers are the computing power and the communication access for the storage that the computing applications require. Servers can connect to local storage or RAID Arrays. Many have tried to characterize the loading of a server in relation to a “number of” quantities involved in a server-storage relationship.

Commonly described in the loading relationships are considerable factors that include the:

- Number of adapters on a server, and
- Number of RAID systems that may be connected to an adapter.

A server may have up to four adapters and may connect up to four RAID arrays. These numbers represent a relatively safe area of loading and are acceptable for most general usage server-storage situations. For some servers (for example, in the 4m class), these

values may be too high; and for other servers (for example, in the 4u class of Enterprise servers), these numbers may be too low. The maximum number of adapters is only limited by the number of available expansion slots. Up to 24 adapters have been tested simultaneously in a mainframe-class machine (E10000).

These numbers, however, do not stand alone when describing a load on a server-storage system. The key values are the actual I/O loading from server to storage. The major components of I/O values are:

- Reading versus writing data
- Size of the records being read/written, and
- Frequency of data to/from a storage device.

The following set of major components also has a large impact on the viability of a configuration from the RAID array point of view:

- Nature of the controller (SCSI versus Fibre Channel)
- Amount of cache on the controller, and
- Configuration of the disks enclosed by the RAID Array.

A specific configuration may or may not replicate the tested laboratory configurations. When treating a specific configuration, the best guideline is to determine the viability of a configuration that is applicable to your environment.

Additionally, while a given configuration *may* support the desired I/O quantity, it *may not* support a desired performance characteristic of the I/O. Thus, both configuration and performance can be best described with some test trials of a specific configuration.

NOTE: This chapter does not attempt to answer questions of loading or performance.

Storage

The term Storage, as used in this chapter, is a RAID Array that is supported by the HSG80 controller. This controller supports all of the standard RAID set configurations in a loop or fabric environment. These configuration elements are discussed in the following sections.

Active-Passive vs. Active-Active Controller Pairs

In dual controller systems, most controller pairs are configured for redundancy, which means that the controllers are bound together in a such a way that if one controller fails, the second controller will automatically service the I/O stream, protecting data integrity. Typically, this redundancy is created for failover protection. Failover can be implemented as either Active-Passive or Active-Active.

In Case 1, (shown in Table 5–1) all I/O is loaded on a single controller, Controller A. In the event of a failure, the passive channel would handle the I/O from the failed controller. Having the load on a single controller is a drawback to the processing power of the controller pair.

Table 5–1: Case 1: Active-Passive (Describes the controller pair)

Channel 1 - Active	Channel 2 - Active	Controller A	Primary Ch 1, Primary Ch 2
Channel 1 - Passive	Channel 2 - Passive	Controller B	Standby Ch 1, Standby Ch 2

In Case 2, both I/O channels are protected as well as distributing the processing load over both controllers.

The fact that one controller is active while the other controller is passive leads to a misnomer of Active-Passive to describe the redundant pair. The reason that this is a misnomer is the fact that the description is referring to a single controller and not the controller pair.

Table 5–2: Case 2: Active-Active (Describes the controller pair)

Channel 1 - Active	Channel 2 - Passive	Controller A	Primary Ch 1, Standby Ch 2
Channel 1 - Passive	Channel 2 - Active	Controller B	Primary Ch 1, Standby Ch 2

When speaking of a single controller, it is correct to say Active-Passive as that describes the state of the channels. This is shown with a single controller in Table 5–3.

Table 5–3: Example of Active-Passive with a Single Controller

Channel 1 - Active	Channel 2 - Passive	Controller A	Primary Ch 1, Standby Ch 2
--------------------	---------------------	--------------	-------------------------------

When applied to the **controller pair**, both controllers are active, the correct designation is Active-Active.

Hubs

Hubs serve to create the connections within a loop. They also contain port-bypass circuitry to allow malfunctioning nodes to be bypassed, thus enabling the loop to continue to function.

NOTE: Because cascading hubs are not supported, the largest supported loop configuration using hubs consists of 12 ports maximum.

Switches

Switches enable multiple connections to each other. With switches, the Fibre Channel is removed from some of the limitations of an arbitrated loop, such as throughput, single point of failure, and sizing of configuration.

As indicated earlier, the complexity of a configuration can increase in direct proportion to the number of components added to the Fibre Channel storage network. While switches provide a great deal of flexibility and scalability of storage access and sharing, they also tend to make transparent the path from server to storage.

IMPORTANT: Take the time to plan your configuration, test the portions of it and record the various connections and paths.

You can extend the fabric's capability by adding more subsystems, more switches, or more servers. Adding switches is probably the easiest way to extend the fabric's capability. To add a switch, connect it to the existing fabric.

In a fabric configuration of cascaded switches, there is only one *principal* switch. If there are two separate switch configurations there will be two principal switches and care must be taken when integrating the two switch configurations.

The most common problem encountered when adding a switch to the fabric is domain overlap, especially when interconnecting old (StorageWorks Fibre Channel) switches and new (StorageWorks SAN) switches. Every switch in a domain must have a unique domain ID. Refer to the documentation that came with your switch for more information on changing domain IDs and cascading switches.

To connect a server or a switch to another switch, insert the Fibre Channel cables into the gigabit interface converters (GBICs) that reside in each port of the switch.

In summary, connections between switches are similar to connections between server and switch, with the added ability of adding more switching capability and/or more storage.

Fibre Channel Access/Transport Paths

In a fibre loop, a server's target is specified by the controller's ports. These ports have Physical Addresses (PA) and in a loop that is arbitrated, the ports are called Arbitrated Loop, Physical Addresses (ALPAs). Since all servers and storage systems connected to a loop share the PA, it is important that every port on the loop have a unique ALPA.

In the fibre fabric, the basic target is the controller's Worldwide Node Name (WWNN). The WWNN is the fabric network address for the controller of the RAID array. WWNNs have a format of AAAA-BBBB-CCCC-DDDD, where A, B, C, D are alphanumeric characters. (See "Worldwide Names (Node IDs and Port IDs)," page 1-27 for details.)

Each controller has two ports and each port has a designated Worldwide Port Name (WWPN). In the StorageWorks' Sun implementation, the specific WWPNs are used to define what nodes in the fabric that Sun will attempt to bind to.

In summary, StorageWorks manages two entities with Fibre Channel technology: ALPAs for loop and WWPNs for fabric.

Implementation of Loop and Fabric Bindings

Bindings

In general, bindings for drivers may be either dynamic or persistent. The drawbacks to dynamic bindings are threefold: 1) indiscriminate selection of Targets, 2) usage of system resources, and 3) long scanning delays while booting a server.

To remedy these three shortcomings, the StorageWorks' Solaris implementation uses a persistent set of bindings for both the loop and switch implementations.

Loop Bindings

Loop bindings rely on the ALPAs of the controller ports to define the SCSI Targets that Solaris binds to. The StorageWorks Solution software installation process (or the *Config.sh* utility) will create entries in the */kernel/drv/sd.conf* file that allow the driver to have a direct entry/path for that Target. Additionally, modifications are made to the driver configuration files per adapter. The driver files are: *fca.conf* for the 32-bit, JNI SBUS adapter, *fcaw.conf* for the 64-bit, JNI SBUS adapter, *fca-pci.conf* for the JNI PCI adapter and *qlc.conf* for Qlogic cards.

Fabric Bindings

Fabric bindings rely on the Worldwide Names and the StorageWorks' implementation uses the more specific relationship, the Worldwide Port Names(WWPN). Similar to the loop implementation, these WWPNs are the basis for the `/kernel/drv/sd.conf` entries that identify a specific fabric access. WWPNs are derived from the WWNN as follows:

If a WWNN is FFFF-SSSS-TTTT-LLL0 the corresponding Port Names are:

FFFF-SSSS-TTTT-LLL1 for Port 1

FFFF-SSSS-TTTT-LLL2 for Port 2

By utilizing the WWPN, the user has complete control over the mappings (and subsequently the system binding) from a server to a port in the fabric. By enabling the Port Names, there is the flexibility of mapping a server to a specific port of a fabric array or to multiple ports of a fabric array.

Additionally, the user is able to map a server with multiple adapters to multiple arrays in the fabric. And, if the idea extends across servers, then there are many possibilities of configuring and sharing a storage in a fabric. Thus, in a shorthand version, connections are as follows:

Server 1 => Switch => array1, Port 1 & array1, Port 2

Or

Server 1 => Switch => array1, Port 1 & array2, Port 1

In every case, storage subsystems can be accessed that are not directly connected to a single switch by resolving the access path through the fabric to a specific WWPN.

Obviously, this ability brings with it complexity and the need to both design and consider the growth possibilities for the fabric and the storage capability. So, this is a *second caution*: **take the time to plan a configuration, test the portions of it and record the various connections and paths.** This activity will provide a valuable tool for support of Fibre Channel fabric installations.

Loop Connections at the Solaris Level

SCSI Target bindings are defined in the `/kernel/drv/sd.conf` file for both loop and fabric definitions. These entries look similar to the following:

```
name="sd" parent="/iommu@f,e0000000/sbus@f,e0001000/fcaw@0,0"  
Target=65 lun=1
```


sd.conf entries reflect the choices that the user has made regarding the Targets, and the number of LUNs they wish to map to.

The following examples have been reformatted for clarity.

Loop Mappings

Three entries from a *sd.conf* file are shown below.

```
name="sd" parent="/sbus@a,0/fcaw@2,0" Target=64 lun=0;
name="sd" parent="/sbus@a,0/fcaw@2,0" Target=64 lun=1;
...
name="sd" parent="/sbus@a,0/fcaw@2,0" Target=64 lun=15;
```

Analysis

Device Class: name="sd"

The device class that is enabled. The sd driver is the SCSI disk, or Target driver for Solaris.

Adapter Path: parent="/iommu@f,e0000000/sbus@f,e0001000/fcaw@0,0"

Identifies the adapter, the instance and the adapter/driver name. Note that the specific adapter path will vary with the specific architecture of the hardware, and the specific slot the adapter is installed in.

Target: Target=64

The default bindings for Solaris use Targets of 64 and 65 for Arbitrated Loop. These Targets correspond with ALPAs 72 and 71. If multiple HSG80 controllers are configured on the same loop, you should have Targets 62 and 63 defined also (corresponding with ALPAs 74 and 73).

LUN: lun=0

The default logical units of 0..15 have been supplied and may be extended to a total of 64.

Fabric Mappings

Two entries from a *sd.conf* file are shown below.

```
name="sd" parent="/iommu@f,e0000000/sbus@f,e0001000/fcaw@1,0"
Target=64
lun=0
wwpn="5000-1fe1-0000-13f1";
...
name="sd" parent="/iommu@f,e0000000/sbus@f,e0001000/fcaw@1,0"
Target=64
lun=15
ww pn="5000-1fe1-0000-13f1";
```

Analysis

Device Class: name="sd"

The device class that is enabled. As with the loop settings, the SCSI disk driver is used.

Adapter Path: parent="/iommu@f,e0000000/sbus@f,e0001000/fcaw@1,0"
Identifies the adapter, the instance and the adapter/driver name.

Target: Target=64

The default bindings for Solaris use Targets 64 and 65 for Fabric. These values are used in conjunction with the WWPN (see below)

Unlike loop mode, these Targets are arbitrary; that is, they do not correspond to an ALPA. Targets 64 and 65 are simply used for compatibility with loop mode.

LUN: lun=0

The default logical units of 0..15 have been supplied and may be extended to a total of 64. (See Adapter Configuration Menu below.)

Fabric Port: wwpn="5000-1fe1-0000-13f1";

The WWPN that Solaris will attempt to bind to. This WWPN identifies the first port on the HSG80 controller.

The /opt/steam/bin/config.sh Utility

When installing the StorageWorks RAID Manager Software, a utility script file is loaded onto the system. While installing the kit, the user may choose an installation directory or take the default, /opt directory. These discussions assume that the default directory was used (/opt/steam).

This utility script has many options. Refer to the Solaris documentation for more information. Our attention here is the ability to configure either loop ALPAs or fabric WWPNs. Both are accomplished with the /opt/steam/bin/*Config.sh* script.

The menu is shown below. All the options in this menu apply to any subsystem and any server, regardless of the Fibre Channel access chosen.

---- RAID Array V8.7 Configuration Menu ----

<p>Agent Admin Options -----</p> <ol style="list-style-type: none"> 1) Change Agent Password 2) Change SNMP Enterprise OID 3) Start/Stop the Agent 4) Toggle Agent Startup on System Boot 5) Uninstall Agent <p>Agent Notification Options -----</p> <ol style="list-style-type: none"> 6) Toggle Error Log Notification 7) Toggle Mail Notification 8) View Mail Notification List 9) Add User to Mail Notification List 10) Remove User from Mail Notification List 11) Modify Mail Notification List <p>Exit: -----</p> <p>q) Quit</p> <p>Enter Selection:</p>	<p>Storage Subsystem Options -----</p> <ol style="list-style-type: none"> 12) View Subsystems 13) Add a Subsystem 14) Remove a Subsystem 15) Modify a Subsystem <p>Client Options -----</p> <ol style="list-style-type: none"> 16) View Clients 17) Add a Client 18) Remove a Client 19) Modify a Client <p>Other Options -----</p> <ol style="list-style-type: none"> 20) Add/Change Adapters 21) High Avail. Options
--	--

Adding/Changing Adapters

When the V8.7 kits are being installed, the user has the opportunity to define the specific type of Fibre Channel configuration for each adapter, either loop or fabric (QLogic HBAs are supported in fabric-mode only).

IMPORTANT: Changing values or the driver mode for one adapter instance will not affect the mode for the other adapter instances. In this way, unique instances of each adapter can exist for a specific type of a fibre transport medium, loop or fabric.

To make a change to the adapters, select Option 20, *Add/Change Adapters*. When selected, you access the following sub-menu:

```

--- Adapter Configuration Menu ---
(sd.conf & fc*.conf)
1) View Adapters
2) Add an Adapter
3) Remove an Adapter
4) Modify an Adapter
5) View available WWPNS
q) Exit
Enter choice:
    
```

NOTE: In the header of this sub-menu, remember that actions that are performed here will affect the specific instance mappings in the /kernel/drv/sd.conf file as well as the specific driver configuration file, fca.conf, fcaw.conf, fca-pci.conf, qla2200.conf, or qla2300.conf.

Option 1 - View Adapters

The following example shows one JNI SBUS 64-bit Adapter and two JNI PCI Adapters. The JNI SBUS adapter is configured for Arbitrated Loop and the JNI PCI adapters are configured for Fabric mode. Note also that both PCI adapters are communicating with the same Worldwide Port Name entities on the Fabric. The RAID configuration employed is unknown.

```

NN Adapter # Control. M LN Targets
Port IDs ( WWPNS ) Driver Path
1 fcaw0 fcaw L 16 64,65
/sbus@2,0/fcaw@2,0
2 fibre-channel 0 fca-pci F 16 64
5000-1fe1-0000-0d41
/pci@a,2000/fibre-channel@2
65 5000-1fe1-0000-0d42
3 fibre-channel 1 fca-pciF 16 64
5000-1fe1-0000-0d41
/pci@b,2000/fibre-channel@2
65 5000-1fe1-0000-0d42
-- Hit RETURN to continue --
    
```

Option 2 - Add an Adapter

Select Option 2 to be informed of and prompted for permission to scan the bus(es) on the system and report any adapters not already listed. The steps that follow when adding an adapter are self-explanatory.

Option 3 - Remove an Adapter

Select Option 3 to view a list of the adapters. You can then either choose to quit the 'remove' operation by pressing **Enter** or select an adapter by inputting the adapter **NN** (numeric identification at the left in the list). The user must then confirm that this adapter is to be removed from the list.

Option 4 - Modify an Adapter

Select Option 5, *Modify an Adapter*, to see a list of the adapters. To quit the 'modify' action by pressing **Enter** or select an adapter to modify by entering the **NN** (numeric identification at the left in the list).

Option 5 - View Available WWPNS

Select this option to run the Scan Adapters Utility. This utility detects available adapters and WWPNS for this host. Perform the following procedure to run the Scan Adapters Utility.

1. `cd opt/steam/bin`
2. `./config.sh`

In the following example you will:

- Change the mode from fabric to loop.
- Change the number of LUNs per Target.
- Change the number of LUNs per Target.
- Set the Target IDs.

```

NN Adapter # Control. M LN Targets Port IDs ( WWPN ) Driver Path
-----
1 fcaw 0 fcaw L 16 64,65 /sbus@2,0/fcaw@2,0
2 fibre-channel 0 fca-pci F16 64 5000-1fe1-0000-0d41
  /pci@a,2000/fibre-channel@2
  65 5000-1fe1-0000-0d42
3 fibre-channel1 fca-pci F 16 64 5000-1fe1-0000-0d41
  /pci@b,2000/fibre-channel@2
  65 5000-1fe1-0000-0d42

Select record number NN (Hit RETURN to escape.):
    
```

Choose Record 2, and the following dialogue takes place. (Each information message or prompt is prefaced with a bullet).

- Driver fca-pci is configured for Fabric mode.

Change the Mode from Fabric to Loop (For JNI HBAs only)

This is always the first statement during modify and serves as a reminder of the mode of the adapter's driver.

- Do you want to change this mode? [y,N]

This question prompts for a change of mode. Since, in this example, the mode is currently set to 'fabric', answering y will set the driver for loop mode. When you implement loop mode, the ALPAs are used for access.

Change the Number of LUNs per Target

When changing from loop to fabric mode, you need to supply the WWPNs.

```
Do you want to change the number of LUNs per Target? [y,N]
```

By default, the Solaris Platform kit software supplies LUNs 0-15 for a total of 16 LUNs per Target. This number is more than adequate for many purposes. However, if you want to change the number of LUNs per Target, select **y**, and the following prompt is displayed:

```
Enter the number of LUNs per Target in a range of 1..64 (default:
16):
```

Set the Target IDs

Either press **Enter** to accept the default value of 16 LUNs or enter the desired value.

```
Do you want to change the Target Ids? [y,N]
```

The Loop Target defaults are 64 and 65. For a server that is communicating to more than one subsystem, you need to have more than the two Targets. If you select **y**, then the following prompt appears:

```
Enter Targets (n= none, default: 64, 65) :
```

The Fabric Target defaults are 0,1 (zero, one)

You now have three options:

- Press **Enter** and accept the default Loop Target Ids of 64, 65.
- Enter the new Target Ids. These Ids should be entered as decimal values, separated by spaces. Up to 16 Targets can be entered, and each Target should be in the range of 17—65.
- Enter the letter, **n** to indicate that there are no Targets specified. The impact of this action is very important as it will preserve the adapter in the mode requested but will remove all the `sd.conf` entries for that driver. The result is that the driver will no longer communicate with any Targets.

If you select fabric mode, then the next prompt will appear in an attempt to map the Target Ids to the WWPNs.

- You configured driver `fcaw` for Fabric Mode.
- Provide Port ID (WWPN) for each default Target, otherwise the Target will be suppressed.
- Driver `fcaw` for adapter `fcaw`.
 - Target: 64, WWPN: \$\$\$\$-\$\$\$\$-\$\$\$\$-\$\$\$\$
 - Enter WWPN (n=none, RETURN=old value, if any): 5000-1fe1-0000-03f1

These prompts are provided for each Target that has been defined. The user has three choices at this prompt:

- Press **Enter** to restore the old value for the WWPN.
- Enter the WWPN value for each Target.
- Enter the letter, **n**. This value specifies that there are no WWPNs specified. The impact of this action is very important as it will preserve the adapter in the mode requested but will remove all the sd.conf entries for that driver. The result is that the driver will no longer communicate with any Targets.

Accept/Reject Changes

At this point, all the modification prompts have been provided and the final actions are to show a summary of the modifications and allow the user to accept or skip the changes.

- Adapter: fcaw
 Path: /pci@a,2000/fibre-channel@2
 Mode: Loop
 LUNs/Target: 8
 Targets: 67,68
- Is this correct? [Y,n]

At this point, you have two options.

- Press **Enter** and choose the default option of **Y** and to accept the adapter driver's modification, including changes to the sd.conf entries and the specific driver file, fca-pci.conf.
- Enter **n** and discard the proposed changes. No changes will be applied to the adapter and driver and everything remains as previously.

Since these changes impact the system level driver mode, LUNs and Targets, the server must be rebooted with the (-r) reconfigure switch. Enter the following command:

```
# reboot -- -r
```

Configuration Procedures

This section presents a set of macro-level procedures to perform:

- “Loop Configurations,” page 5-41
- “Fabric Configurations,” page 5-41

- “Reconfiguring from Loop to Fabric,” page 5–42
- “Reconfiguring from Fabric to Loop,” page 5–43

The following assumptions have been made in these procedures:

- The server in question has an adapter to support the Fibre Channel loop or fabric.
- A storage system is available that will support a fibre connection.
- Hubs and switches are available.

Loop Configurations

NOTE: This procedure assumes that the Solaris platform kit is being loaded onto the system and a specific mode of operation is being selected.

1. Load the Solaris Platform kit and select the Loop mode for the driver when prompted. By default, the Targets that are associated with the adapter are 64, 65.
2. Reboot the system. This is the normal reboot after driver/kit installation.
3. Complete the installation as documented with `/opt/steam/bin/Install.sh`.
4. If the server is in a hub configuration with multiple subsystems:
 - a. Verify that the controller is configured for loop mode and that the specific, desired ALPAs are assigned to each destination controller port using the CLI.
 - b. Verify that the server has the specific ALPAs assigned as Targets for the adapter by using `/opt/steam/bin/Install.sh - Option 20 Add/Change Adapters, Option 1, View adapters`.
5. Reboot the system for cases that require a change of ALPAs at the server. Enter the command `reboot -- -r` to force a reconfiguration boot.
6. After reboot, enter the `format` command to see the intended Targets list. This assumes that at least one unit has been created on the subsystems.

Fabric Configurations

NOTE: This procedure assumes that the Solaris platform kit is being loaded onto the system and a specific mode of operation is being selected.

1. Load the Solaris Platform kit and select the Fabric mode for the driver when prompted. There are no default WWPNs so the software will prompt for the specific WWPNs. When entering the values, case is not important, but the hyphenation is.

2. Reboot the system. This is the normal reboot after driver/kit installation.
3. Complete the installation as documented with `/opt/steam/bin/Install.sh`
4. Verify that the destination subsystems controllers are configured for fabric mode using the CLI.
5. If the Targets cannot be seen under the format command then a reboot `-r` of the server will be required. (This assumes that some change was made to the server or controller.)

Reconfiguring from Loop to Fabric

NOTE: QLogic HBAs are supported in fabric-mode only.

NOTE: This procedure assumes that the Solaris platform kit is loaded and everything is running as desired but a change of configuration is required.

1. Stop I/O to the subsystem. Perform backups, as required.
2. Remove the connections from the hub.
3. Reconfigure the controllers from loop to fabric mode using the CLI.
4. Create the connections from server to switch and from switch to subsystem. For configurations with multiple adapters and multiple subsystems, verify all the Fibre Channel paths from server to subsystems.
5. At the servers, reconfigure the adapters to support the fabric WWPNs and remove the unused loop Targets. This is accomplished with `opt/steam/bin/Install.sh - Option 20 Add/Change Adapters`, Option 4, *Modify Adapters*. This option allows the mode of the driver to be changed and the WWPNs to be added for the configuration. For each adapter that is used in the changed configuration, this modify step must be repeated. For each server in the changed configuration, this modify step must be repeated.
6. The changes from Step 5 require that the servers be rebooted (reboot `-r`) to reinitialize the drivers in the correct mode and acquire the new entries/Targets from the `sd.conf` files as discussed above. Shutdown and power-off the Sun server before rebooting. The adapters must be power-cycled for reliable operation when changing from loop to fabric mode.
7. Using the format command, verify that the units of the subsystems are now available.

Reconfiguring from Fabric to Loop

NOTE: QLogic HBAs are supported in fabric-mode only.

NOTE: This procedure assumes that the Solaris platform kit is loaded and everything is running as desired, but a change of configuration is required.

1. Stop I/O to the subsystem. Perform backups as required.
2. Remove connections from the switch and any other intermediate switches to the subsystem.
3. Reconfigure the controllers from fabric support to loop using SWCC or the CLI. At this time verify that the intended ALPAs are assigned to the controller ports as desired. The defaults may suffice.
4. Create the connections from server to hub and from hub to subsystem. For configurations with multiple adapters and multiple subsystems, verify all the Fibre Channel paths from server to subsystems through the hubs.
5. At the servers, reconfigure the adapters to support loop. This is accomplished with `opt/steam/bin/Install.sh` - Option 20, *Add/Change Adapters*, Option 4, *Modify Adapters*. This option allows the mode of the driver to be changed and the Targets to be verified for the configuration. For each adapter that is used in the changed configuration, this modify step must be repeated. For each server in the changed configuration, this modify step must be repeated.
6. The changes from Step 5 require that the servers be rebooted (`reboot -r`) to reinitialize the drivers in the correct mode and acquire the new entries/Targets from the `sd.conf` files as discussed above. Shutdown and power-off the Sun server before rebooting. The adapters must be power-cycled for reliable operation when changing from fabric to loop mode.
7. Using the `format` command, verify that the units of the subsystems are now available.

Using CLI for Configuration

This chapter presents an example of how to configure a storage subsystem using the Command Line Interpreter (CLI). The CLI configuration example shown assumes:

- A normal, new controller pair, which includes:
 - 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 Sun Solaris hosts, as shown in Figure 6–2. The resulting virtual system, from the host's point of view, is shown in Figure 6–3.

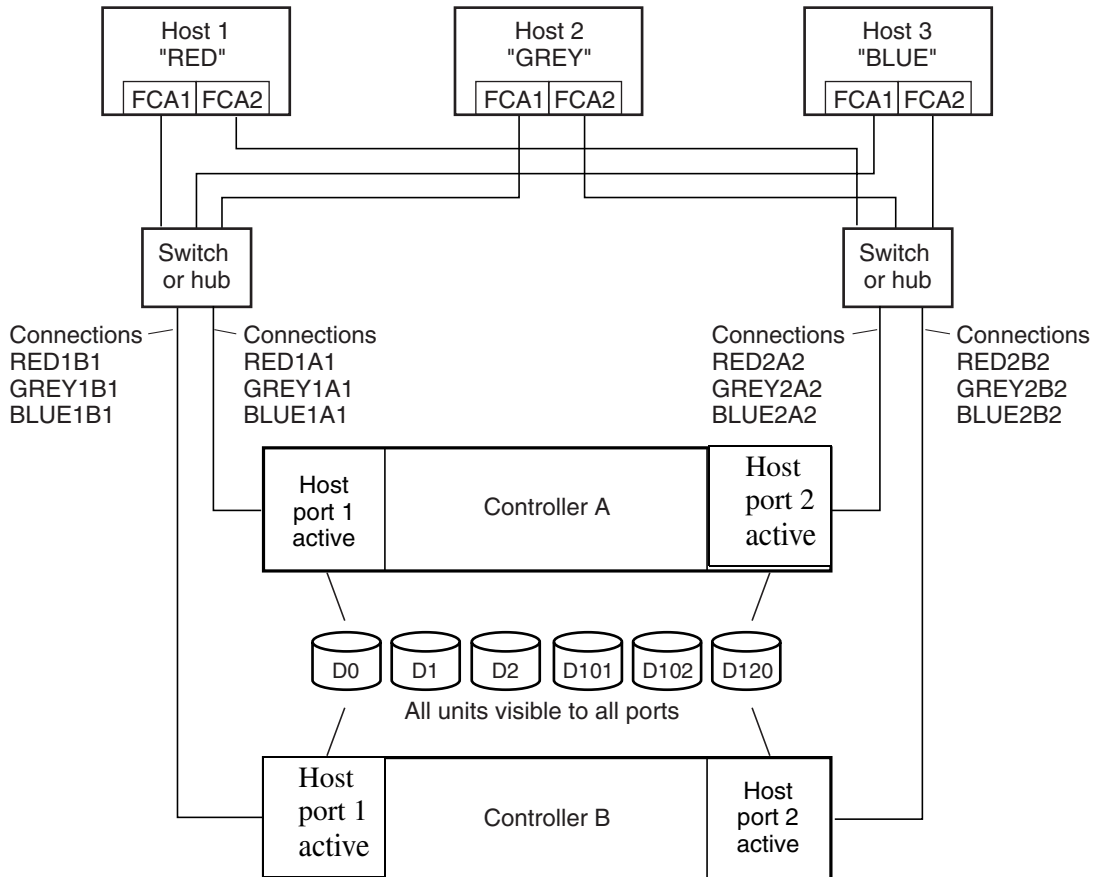
Figure 6–1 shows an example storage system map for the BA370 enclosure. Details on building your own map are described in Chapter 2. Templates to help you build your storage map are supplied in Appendix A.

	Port							
	1	2	3	4	5	6		
Power Supply	D2 S2 DISK103 00	D2 S2 DISK203 00	D2 S2 DISK303 00	D2 S2 DISK403 00	D2 D101 DISK503 00	spareset member DISK603 00	Power Supply	3
Power Supply	D0 S1 MI DISK102 00	D0 S1 M1 DISK202 00	D0 S1 M2 DISK302 00	D0 S1 M2 DISK402 00	D1 M3 DISK502 00	D1 M3 DISK602 00	Power Supply	2
Power Supply	D120 R2 DISK101 00	D120 R2 DISK201 00	D120 R2 DISK301 00	D120 R2 DISK401 00	D120 R2 DISK501 00	D120 R2 DISK601 00	Power Supply	1
Power Supply	D102 R1 DISK100 00	D102 R1 DISK200 00	D102 R1 DISK300 00	D102 R1 DISK400 00	D102 R1 DISK500 00	D102 R1 DISK600 00	Power Supply	0

Figure 6–1: Example storage map for the BA370 Enclosure

Figure 6–2 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host BLUE can be done by enabling only the connections to host BLUE. At least two connections must be enabled for multiple-bus failover to work. For most operating systems, it is desirable to have all connections to the host enabled. The example system, shown in Figure 6–2, contains three non-clustered Sun

Solaris hosts. Port 1 link is separate from 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

CX07547A

Figure 6–2: Example, three non-clustered host systems

Figure 6–3 represents units that are logical or virtual disks comprised of storagesets configured from physical disks.

Figure 6–4 represents an example of virtual system layout from the hosts' point of view.




Figure 6–3: Example, logical or virtual disks comprised of storagesets




Figure 6–4: Example, virtual system layout from hosts' point of view

CLI Configuration Example

Text conventions used in this example 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: “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 ALLOCATION_CLASS=0
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 [Y]
```

Y

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 [Y]
```

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	WINNT	THIS	1	XXXXX X	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXX		ADAPTER_ID=XXXX-XXXX-XXXX-XX XX				

```

RENAME !NEWCON00 PURPLE1A1
SET PURPLE1A1 OPERATING_SYSTEM=SUN
SHOW CONNECTIONS
    
```

NOTE: Connection table sorts alphabetically.

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

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	WINNT	THIS	2	XXXXXX X	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
PURPLE1A1	SUN	OTHER	1	XXXXXX X	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

```

RENAME !NEWCON01 WHITE1B2
SET WHITE1B2 OPERATING_SYSTEM=SUN
SHOW CONNECTIONS

```

Mark or tag both end of Fibre Channel cables.

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

```

RENAME !NEWCON02 TAN1B2
SET TAN1B2 OPERATING_SYSTEM=SUN
SET TAN1B2 UNIT_OFFSET=120
SHOW CONNECTIONS
    
```

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
PURPLE1A1	SUN	OTHER	1	XXXXXX X	OL other	0
			HOST_ID=XXXX-XXXX-XXXX-XXXX ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
WHITE1B2	SUN	THIS	2	XXXXXX X	OL this	100
			HOST_ID=XXXX-XXXX-XXXX-XXXX ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
TAN1B2	SUN	THIS	2	XXXXXX X	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 DISK30000 DISK40000 DISK50000
DISK60000
INITIALIZE R1
ADD UNIT D102 R1 DISABLE_ACCESS_PATH=ALL
SET D102 ENABLE_ACCESS_PATH=PURPLE1A1
ADD RAIDSET R2 DISK10100 DISK20100 DISK30100 DISK40100 DISK50100
DISK60100
INITIALIZE R2
ADD UNIT D120 R2 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(TAN1B2)
ADD MIRRORSET MI DISK10200 DISK20200
ADD MIRRORSET M2 DISK30200 DISK40200
ADD STRIPESSET S1 M1 M2
INITIALIZE S1
ADD UNIT D0 S1 DISABLE_ACCESS_PATH=ALL
SET D0 ENABLE_ACCESS_PATH=PURPLE1A1
ADD MIRRORSET M3 DISK50200 DISK60200
INITIALIZE M3
ADD UNIT D1 M3 DISABLE_ACCESS_PATH=ALL
SET D1 ENABLE_ACCESS_PATH=(WHITE1B2)
ADD STRIPESSET S2 DISK10300 DISK20300 DISK30300 DISK40300
INITIALIZE S2
ADD UNIT D2 S2 DISABLE_ACCESS_PATH=ALL
SET D2 ENABLE_ACCESS_PATH=PURPLE1A1
INITIALIZE DISK50300
ADD UNIT D101 DISK50300 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=(WHITE1B2)
ADD SPARESET DISK60300
SHOW UNITS FULL
```

Backing Up, Cloning, and Moving Data

This chapter includes the following topics:

- “Backing Up Subsystem Configurations,” page 7–1
- “Creating Clones for Backup,” page 7–2
- “Moving Stagesets,” page 7–7

Backing Up Subsystem Configurations

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.

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.

IMPORTANT: DO NOT use `SAVE_CONFIGURATION` in dual redundant controller installations. It is not supported and may result in unexpected controller behavior.

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

```
SHOW DEVICES FULL
```

Creating Clones 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 back up. The CLONE utility uses steps shown in Figure 7-1 to duplicate each member of a unit.

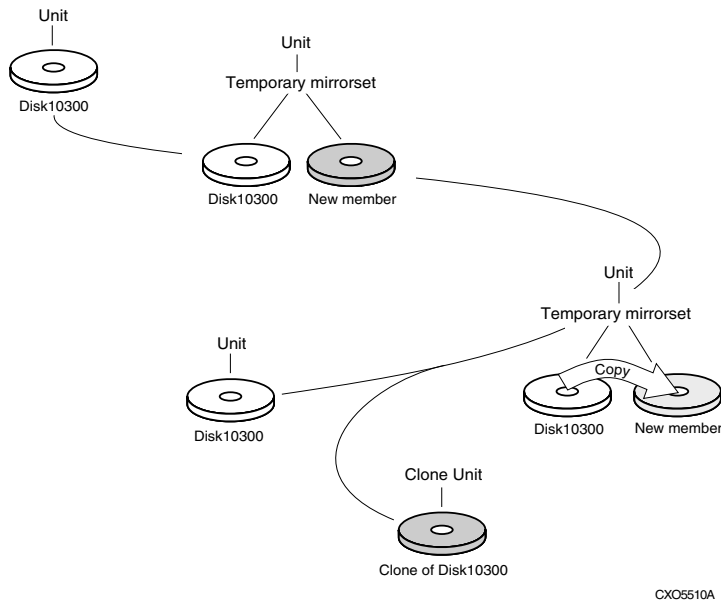


Figure 7-1: CLONE utility steps for duplicating unit members

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 your 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 D6. The clone command terminates after it creates storage unit D33, a clone or copy of D6.

```
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

```

Linking WWIDs for Snap and Clone Units

LUN WWIDs (World Wide Identifiers) for snap and clone units are different each time they are created. This causes more system data records to keep track of the WWIDs as well as script changes at the customer sites. To eliminate this issue, a linked WWID scheme has been created, which keeps the WWIDs of these units constant each time they are created.

The WWID of a LUN is 128bits long, and is composed as follows:

LW3	LW2	LW1	LW0
Controller Node ID		Controller Serial	VSN_Seed

Controller Node ID is the 64-bit fibre channel node id of the controller pair.

Controller Serial is the low-order 48 bits of the controller that “initialized” the storage set. The Controller_Serial is composed from several fields, but the high order 12 bits are a Reserved Field.

VSN_Seed is a counter that is incremented every time we initialize a storage set.

In the case that the linked WWID is already in use, a unique WWID will be allocated, and a message to this effect will be displayed. This is not a syntax error and does not cause the command to fail.

CLI format

```
CLI> add snapshot_units <SnapUnit> <UnitName> <SourceName>
use_parent_wwid
```

Snapshot Unit - unit number to be assigned to the snap unit.

Unit Name - name of storage set or disk that will become the snap unit.

Source Name - unit number of source storage set or disk.

Example:

```
CLI> add snap d2 disk10100 d1 use_parent_wwid
```

Example of error message text:

```
CLI> add snap d2 disk10100 d1 use_parent_wwid
```

A new WWID has been allocated for this unit because the linked WWID for d2 is already in use.

```
CLI> run clonew
```

Implementation Notes

Add Snap with Linked WWID - The user has a script that runs every night to create a snapshot, run a backup to tape from the snapshot, then delete the snapshot. Each time this is done, a new WWID is allocated. When the operating system runs out of room for all of these “orphaned” WWIDs, the host system must be rebooted. Therefore the user decides to update the script such that the “add snap” command reads as follows:

```
CLI> add snap d2 disk10100 d1 use_parent_wwid.
```

This results in same WWID being used for the snapshot each night.

Run Clonew - Operates the same as “CLI> run clone” with the exception that clonew will use the linked WWID associated with the source unit instead of allocating a new one.

```
CLI> run clonew
```

Run Clone - Works the same as v86. In other words a unique WWID is always allocated to the clone unit.

Clonew of a Snap - The user wants to clone a snap unit without using any more WWIDs. The clone created from the snap unit will be created using the linked WWID associated with the snap unit.

Exception: A new WWID will be allocated if the snapshot was created using the `use_parent_wwid` switch. Each WWID only has one linked WWID variation. If the linked WWID is already in use, then a new unique WWID is allocated.

Snap of a Clone - The user want to snapshot a clone without using any more WWIDs, and issues "`CLI> add snap d3 disk10300 d2 use_parent_wwid`". D3, the snapshot unit, will be created with the linked WWID associated with d2.

Exception: A new WWID will be allocated if the clone was created using the "`run clonew`" command. Each WWID only has one linked WWID variation. If the linked WWID is already in use, then a new unique WWID is allocated.

Manual Clone Creation - The user has his own set of scripts that create clones, and wants to update them to use linked WWIDs. At some point in the script there will be an "`add unit`" command. The switch "`parent_wwid=<unit>`" must be provided.

For example,

```
CLI> add unit d2 disk10100 parent_wwid=d1
```

would create a unit d2 from device disk10100 whose WWID would be the linked WWID associated with unit d1.

Exception: A new WWID will be allocated if the linked WWID associated with d1 is already in use (for example, a clone or snapshot is already using the linked WWID)

Moving StorageSets

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



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

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



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

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 your 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
```

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

Old cabinet

```
DELETE D100
DELETE RAID99
DELETE DISK10000
DELETE DISK10100
DELETE DISK20000
DELETE DISK20100
```

New cabinet

```
ADD DISK DISK10000
ADD DISK DISK10100
ADD DISK DISK20000
ADD DISK DISK20100
ADD RAIDSET RAID99 DISK10000 DISK10100 DISK20000 DISK20100
ADD UNIT D100 RAID99
```


Subsystem Profile Templates

This appendix contains storageset profiles to copy and use to create your 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 physical location of the disk enclosures 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 BA370 Enclosure,” page A-4
- “Storage Map Template 2 for the second BA370 Enclosure,” page A-5
- “Storage Map Template 3 for the third BA370 Enclosure,” page A-6
- “Storage Map Template 4 for the Model 4214R Disk Enclosure,” page A-7
- “Storage Map Template 5 for the Model 4254 Disk Enclosure,” page A-9
- “Storage Map Template 6 for the Model 4310R Disk Enclosure,” page A-11
- “Storage Map Template 7 for the Model 4350R Disk Enclosure,” page A-14
- “Storage Map Template 8 for the Model 4314R Disk Enclosure,” page A-16
- “Storage Map Template 9 for the Model 4354R Disk Enclosure,” page A-19

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 BA370 Enclosure

Use this template for:

- BA370 single-enclosure subsystems
- first enclosure of multiple BA370 enclosure subsystems

	Port							
	1	2	3	4	5	6		
Power Supply	D10300	D20300	D30300	D40300	D50300	D60300	Power Supply	3
Power Supply	D10200	D20200	D30200	D40200	D50200	D60200	Power Supply	2
Power Supply	D10100	D20100	D30100	D40100	D50100	D60100	Power Supply	1
Power Supply	D10000	D20000	D30000	D40000	D50000	D60000	Power Supply	0

Targets

Storage Map Template 2 for the second BA370 Enclosure

Use this template for the second enclosure of multiple BA370 enclosure subsystems.

	Port							
	1	2	3	4	5	6		
Power Supply	D11100	D21100	D31100	D41100	D51100	D61100	Power Supply	11
Power Supply	D11000	D21000	D31000	D41000	D51000	D61000	Power Supply	10
Power Supply	D10900	D20900	D30900	D40900	D50900	D60900	Power Supply	9
Power Supply	D10800	D20800	D30800	D40800	D50800	D60800	Power Supply	8

Targets

Storage Map Template 3 for the third BA370 Enclosure

Use this template for the third enclosure of multiple BA370 enclosure subsystems.

	Port							
	1	2	3	4	5	6		
Power Supply	D11500	D21500	D31500	D41500	D51500	D61500	Power Supply	15
Power Supply	D11400	D21400	D31400	D41400	D51400	D61400	Power Supply	14
Power Supply	D11300	D21300	D31300	D41300	D51300	D61300	Power Supply	13
Power Supply	D11200	D21200	D31200	D41200	D51200	D61200	Power Supply	12

Targets

Storage Map Template 4 for the Model 4214R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4214R disk enclosure (single-bus). You can have up to six Model 4214R disk enclosures per controller shelf.

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

Model 4214R Disk Enclosure Shelf 3 (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	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

Storage Map Template 5 for the Model 4254 Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4254 disk enclosure (dual-bus). You can have up to three Model 4254 disk enclosures per controller shelf.

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

Model 4254 Disk Enclosure Shelf 2 (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	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

continued on the following page

continued from previous page

Model 4254 Disk Enclosure Shelf 3 (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	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Storage Map Template 6 for the Model 4310R Disk Enclosure

Use this template for a subsystem with a six-shelf Model 4310R disk enclosure (single-bus). You can have up to six Model 4310R disk enclosures per controller shelf.

Model 4310R Disk Enclosure Shelf 6 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200

Model 4310R Disk Enclosure Shelf 5 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200

Model 4310R Disk Enclosure Shelf 4 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

Model 4310R Disk Enclosure Shelf 1 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk11000	Disk11100	Disk11200

Model 4310R Disk Enclosure Shelf 2 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk21000	Disk21100	Disk21200

Model 4310R Disk Enclosure Shelf 3 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk31000	Disk31100	Disk31200

Storage Map Template 7 for the Model 4350R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4350R disk enclosure (single-bus). You can have up to three Model 4350R disk enclosures per controller shelf.

Model 4350R Disk Enclosure Shelf 6 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk61000	Disk61100	Disk61200

Model 4350R Disk Enclosure Shelf 5 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk51000	Disk51100	Disk51200

Model 4350R Disk Enclosure Shelf 4 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	05	08	10	11	12
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk41000	Disk41100	Disk41200

Storage Map Template 8 for the Model 4314R Disk Enclosure

Use this template for a subsystem with a six-shelf Model 4314R disk enclosure. You can have a maximum of six Model 4314R disk enclosures with each Model 2200 controller enclosure.

Model 4314R Disk Enclosure Shelf 6 (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	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	Disk61400	Disk61500

Model 4314R Disk Enclosure Shelf 5 (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	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	Disk51500	Disk51500

Model 4314R Disk Enclosure Shelf 4 (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	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300	Disk41400	Disk41500

continued from previous page

Model 4314R 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 4314R 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

Model 4314R Disk Enclosure Shelf 3 (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	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	Disk31400	Disk31500

Storage Map Template 9 for the Model 4354R Disk Enclosure

Use this template for a subsystem with a three-shelf Model 4354R disk enclosure (dual-bus). You can have up to three Model 4354R disk enclosures per controller shelf.

Model 4354R Disk Enclosure Shelf 1 (dual-bus)

	SCSI Bus A							SCSI 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

Model 4354R Disk Enclosure Shelf 2 (dual-bus)

	SCSI Bus A							SCSI 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	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

Model 4354R Disk Enclosure Shelf 3 (dual-bus)

	SCSI Bus A							SCSI 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	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Installing, Configuring, and Removing the Client

The following information is included in this appendix:

- “Why Install the Client?,” page B-2
- “Before You Install the Client,” page B-2
- “Installing the Client,” page B-4
- “Installing the Integration Patch,” page B-5
- “Troubleshooting Client Installation,” page B-8
- “Adding Storage Subsystem and its Host to Navigation Tree,” page B-10
- “Removing Command Console Client,” page B-12
- “Where to Find Additional Information,” page B-13

Why Install the Client?

The Client monitors and manages a 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 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. SNMP must be 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 Windows NT Service Pack 6A 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 “Installing the Integration Patch,” page B-5 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 with 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

The following restriction should be observed when installing SWCC on Windows NT 4.0 Workstations.

If you select all of the applets during installation, the installation will fail on the HSG60 applet and again on one of the HSG80 applets. The workaround is to install all of the applets you want except for the HSG60 applet and the HSG80 ACS 8.5 applet. You can then return to the setup program and install the one that you need.

1. In a SAN environment where you would need both HSG60 and HSG80 subsystems, StorageWorks recommends you install both, but one at a time. This problem is not seen under Windows NT 4.0 Server.
2. Insert the CD-ROM into a computer running Windows 2000 with Service Pack 2 or Windows NT 4.0 (Intel) with Service Pack 6.0A.
3. A dialog box should automatically appear.
4. One of the items in the dialog box should say “SWCC Client Software” and has a button that says “INSTALL” next to it. Click on the button to start the SWCC client installation procedure.
5. Select “HSG80 Controller for ACS87 or newer” menu option to properly install SWCC client, and click **Next**.

If this method does not work, go to the \windows directory on the CD-ROM and run the *setup.exe* program.

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

6. 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 *StorageWorks Command Console Version 2.5, User Guide*.

Installing the Integration Patch

The integration patch determines which version of firmware the controller is using and launches the appropriate StorageWorks Command Console (SWCC) Storage Window within Insight Manager (CIM) version 4.23.

Should I Install the Integration Patch?

Install this patch if your HSG80 controller uses ACS 8.6 or later. This patch enables you to use the controller's SWCC Storage Window within CIM to monitor and manage the controller.

How to Install the Integration Patch

Perform the following steps to install the integration patch.

1. Verify that you have installed the HSG80 Storage Window for ACS 8.6 or later in the Add/Remove Programs in the Windows Control Panel. The HSG80 Storage Window for ACS 8.6 or later is needed to display the correct Storage Window for your version of the firmware.
2. Verify that you have installed HSG80 Storage Window version 2.1 in the Add/Remove Programs (StorageWorks HSG80 V2.1) in the Windows Control Panel. The HSG80 Storage Window version 2.1 is required to run the integration patch.
3. Verify that you have installed CIM version 4.23.
4. Install the integration patch from the Solution Software CD-ROM by double-clicking on setup.exe in the following directory:

```
\SWCC\Client\HSG80shim
```

The patch is installed in the same location as the original SWCC installation.

IMPORTANT: IMPORTANT: Do not remove the HSG80 Client from your computer. If you remove the HSG80 Client, you will no longer be able to access its Storage Window.

Integrating Controller's SWCC Storage Window with CIM

You can open the controller's Storage Window from within the Windows-based CIM version 4.23 by doing the following:

1. Verify that you have installed the following by looking in Add/Remove Programs in Control Panel:
 - The HSG80 Storage Window for ACS 8.6 or later (Required to open the correct Storage Window for your firmware).
 - The HSG80 Storage Window version 2.1 (StorageWorks HSG80 V2.1) The CIM integration patch uses files in this program.
 - CIM version 4.23.
 - CIM integration patch (HSG80 Insight Manager Shim).
2. Verify that you have installed the CIM Agent and the StorageWorks Command Console HS-Series Agent on the same computer.
3. Add the name of the client system that has CIM to the Agent's list of client system entries and choose SNMP as a notification scheme.
4. Open Insight Manager.
5. To open the Server window, click on the device you want to observe in the CIM Navigation window.
6. Click on the Mass Storage button in the Server window. The CIM Navigation Tree is displayed.
7. Click on the + symbol next to RAID Storage System. The Navigation Tree expands to display a listing called Storage System Information.
8. Double-click Storage System Information. You are given the status of the system.
9. Click Launch. The controller's Storage Window is displayed.

Insight Manager Unable to Find Controller's Storage Window

If you installed Insight Manager before SWCC, Insight Manager will be unable to find the controller's Storage Window. To find the controller's Storage Window, perform the following procedure:

1. Double-click the Insight Agents icon (Start > Settings > Control Panel). A window appears showing you the active and inactive Agents under the Services tab.
2. Highlight the entry for Fibre Array Information and click Add. The Fibre Array Information entry is moved from Inactive Agents to Active Agents.

Removing the Integration Patch Will Corrupt Storage Window

If you remove the integration patch, HSG80 Storage Windows version 2.1 will no longer work and you will need to reinstall HSG80 Storage Windows version 2.1. The integration patch uses some of the same files as the HSG80 Storage Window version 2.1.

Troubleshooting Client Installation

This section provides information on how to resolve some of the problems that may appear when installing 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
kzpccdiscoveryport	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 Storage Subsystem and its Host to Navigation Tree

The Navigation Tree enables you to manage storage over the network by using the Storage Window. If you plan to use pager notification, you must 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 StorageWorks Command Console Version 2.5, User Guide.

Removing Command Console Client

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

NOTE: When you remove the CCL, 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 CCL.

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 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 HSG80 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 StorageWorks Command Console Version 2.5, User Guide.

About the User Guide

StorageWorks Command Console Version 2.5, 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 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.

8B/10B	A type of byte definition encoding and decoding to reduce errors in data transmission patented by the IBM Corporation. This process of encoding and decoding data for transmission has been adopted by ANSI.
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.
alias address	An AL_PA value recognized by an arbitrated loop port in addition to the assigned AL_PA.
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.
arbitrate	A process of selecting one L_Port from a collection of several ports that request use of the arbitrated loop concurrently.
arbitrated loop	A loop type of topology where two or more ports can be interconnected, but only two ports at a time can communicate.
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.
association set	<p>A group of remote copy sets that share selectable attributes for logging and failover. Members of an association set transition to the same state simultaneously. For example, if one association set member assumes the failsafe locked condition, then other members of the association set also assume the failsafe locked condition.</p> <p>An association set can also be used to share a log between a group of remote copy set members that require efficient use of the log space.</p>
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.
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.
battery hysteresis	The ability of the software to allow write-block caching during the time a battery is charging, but only when a previous down time has not drained more than 50 percent of rated battery capacity.
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.
cache module	A fast storage buffer
CCL	CCL-Command Console LUN, a “SCSI Logical Unit Number” virtual-device used for communicating with Command Console Graphical User Interface (GUI) software.
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.
CLI	See Command Line Interface.
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 interface	CLI. A command line entry utility used to interface with the HS-series controllers. CLI enables the configuration and monitoring of a storage subsystem through textual commands.
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 HSG80 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 subsystem cabinets, such as the cabinets in which <i>StorageWorks</i> 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.
DDL	Dual data link. The ability to operate on the CI bus using both paths simultaneously to the same remote node.
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 (See DWZZA).
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.

DWZZA	A StorageWorks SCSI bus signal converter used to connect 8-bit single-ended devices to hosts with 16-bit differential SCSI adapters. This converter extends the range of a single-ended SCSI cable to the limit of a differential SCSI cable.
DWZZB	A StorageWorks SCSI bus signal converter used to connect a variety of 16-bit single-ended devices to hosts with 16-bit differential SCSI adapters.
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.
failback	The process of restoring data access to the newly-restored controller in a dual-redundant controller configuration. <i>See</i> also failover.

failedset	A group of failed mirrorset or RAIDset devices automatically created by the controller.
failover	<p>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.</p> <p>The ability for HSG80 controllers to transfer control from one controller to another in the event of a controller failure. This ensures uninterrupted operation. Use Transparent Failover mode for single HBA configurations. Use multiple-bus failover mode for Secure Path based configurations.</p>
FCA	Fibre Channel Adapter
FC-AL	The Fibre Channel Arbitrated Loop standard. See <i>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 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.
HBA	Host bus adapter
host compatibility mode	A setting used by the controller to provide optimal controller performance with specific operating systems. This improves the controller performance and compatibility with the specified operating 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 or hot-pluggable	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 (see SBB).
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	LUN. 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.
LUN	<i>See</i> logical unit number.
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.
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 <i>StorageWorks</i> 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.
multibus failover	Allows the host to control the failover process by moving the units from one controller to another.
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.

offset	A relative address referenced from the base element address. Event Sense Data Response Templates use offsets to identify various information contained within one byte of memory (bits 0 through 7).
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.
parity RAID	<i>See</i> RAIDset.
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.
partition	A logical division of a container, represented to the host as a logical unit.
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.
PDU	Power distribution unit. The power entry device for <i>StorageWorks</i> cabinets. The CDU provides the connections necessary to distribute power to the cabinet shelves and fans.
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.

pluggable	A replacement method that allows the complete system to remain online during device removal or insertion. The system bus must be halted, or quiesced, for a brief period of time during the replacement procedure. See also hot-pluggable.
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 login and port discovery process.
preferred address	The AL_PA which an NL_Port attempts to acquire first during initialization.
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 mirror set.
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 mirror sets.
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 storage set that, unlike RAID level 3, stores the parity information across all of the disk drives within the storage set. <i>See also</i> RAID level 3.
RAID level 3/5	A RAID storage set that stripes data and parity across three or more members in a disk array. A RAID set combines the best characteristics of RAID level 3 and RAID level 5. A RAID set is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAID set is sometimes called parity RAID.
RAID set	<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.
remote copy	A feature intended for disaster tolerance and replication of data from one storage subsystem or physical site to another subsystem or site. Remote copy also provides methods of performing a backup at either the local or remote site. With remote copy, user applications continue to run while data movement goes on in the background. Data warehousing, continuous computing, and enterprise applications all require remote copy capabilities.
remote copy set	A bound set of two units, one located locally and one located remotely, for long-distance mirroring. The units can be a single disk, or a storageset, mirrorset, or RAIDset. A unit on the local controller is designated as the “initiator” and a corresponding unit on the remote controller is designated as the “target”.
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	<i>StorageWorks</i> building block. (1) A modular carrier plus the interface required to mount the carrier into a standard <i>StorageWorks</i> 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> DOC (DWZZA-On-a-chip) 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 storage set.
storage array	An integrated set of storage devices.
storage array subsystem	<i>See</i> storage subsystem.
storage set	(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.
storage set expansion	The dynamic expansion of the storage capacity (size) of a unit. A storage container is created in the form of a concatenation set which is added to the existing storage set defined as a unit.
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 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 <i>StorageWorks</i> enclosures to form storage subsystems.</p> <p><i>StorageWorks</i> 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>
stripe	The data divided into blocks and written across two or more member disks in an array.
striped mirrorset	See RAID level 0+1.
stripeset	See 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.
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.
transfer data rate	The speed at which data may be exchanged with the central processor, expressed in thousands of bytes per second (kbytes).
transparent failover	Keeps the storage array available to the hosts by allowing the surviving controller of a dual redundant pair to take over total control of the subsystem and is transparent (invisible) to the hosts.
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 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	Write-through caching always writes directly to disk, ensuring that the application is never tricked into believing that the data is on the disk when it may not be. This results in highest data integrity, through with slightly reduced performance.

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.

A

ADD CONNECTIONS

- multiple-bus failover 1–19
- transparent failover 1–17

ADD UNIT

- multiple-bus failover 1–19
- transparent failover 1–17

adding

- Client 4–8
- subsystem 4–8
- virtual disks B–13

adding a disk drive to the spareset

- configuration options 5–23

adding disk drives

- configuration options 5–23

Agent

- choosing passwords 4–7
- configuring
 - using config.sh 4–6
- configuring within FirstWatch 4–9
- functions 4–1
- installing 4–5
- restarting 4–8

array of disk drives 2–15

assigning unit numbers 1–16

assignment

- unit numbers
 - fabric topology 5–22
- unit qualifiers
 - fabric topology 5–22

assignment of unit numbers

- fabric topology
 - partition 5–22
 - single disk 5–22

asynchronous event service B–13

automatic mode

- installation 3–10

autospare

- enabling
 - fabric topology 5–24

availability 2–22

B

Back up, Clone, Move Data 7–1

backup

- cloning data 7–2
- subsystem configuration 7–1

C

cabling

- controller pair 5–11
- multiple-bus failover
 - fabric topology configuration 5–10
- single controller 5–4

cache modules

- location 1–2, 1–3
- read caching 1–10
- write-back caching 1–10
- write-through caching 1–10

caching techniques

- mirrored 1–10

- read caching 1–10
- read-ahead caching 1–10
- write-back caching 1–10
- write-through caching 1–10
- changing switches
 - configuration options 5–25, 5–26
- choosing passwords
 - Agent 4–7
- chunk size
 - choosing for RAIDsets and stripesets 2–30
 - controlling stripesize 2–30
 - using to increase request rate 2–30
 - using to increase write performance 2–32
- CHUNKSIZE 2–30
- CLI commands
 - installation verification 5–9, 5–16
- CLI configuration
 - example 6–4
- CLI configurations 6–1
- CLI prompt
 - changing
 - fabric topology 5–23
- Client
 - adding 4–8
 - removing B–12
 - uninstalling B–12
- CLONE utility
 - backup 7–2
- cloning
 - backup 7–2
- command console LUN 1–11
 - SCSI-2 mode 1–20
 - SCSI-3 mode 1–20
- comparison of container types 2–15
- configuration
 - backup 7–1
 - fabric topology
 - devices 5–16
 - multiple-bus failover
 - cablings 5–10
 - multiple-bus failover using CLI 6–4
 - single controller cabling 5–4
 - restoring 2–32
 - rules 2–3
- Configuration Flowchart 1–xviii
- configuration options
 - fabric topology
 - adding a disk drive to the spareset 5–23
 - adding disk drives 5–23
 - changing switches
 - device 5–26
 - displaying the current switches 5–25
 - initialize 5–26
 - RAIDset and mirrorset 5–25
 - unit 5–26
 - changing switches for a storageset or device 5–25
 - changing the CLI prompt 5–23
 - deleting a storageset 5–25
 - enabling autospare 5–24
 - removing a disk drive from the spareset 5–24
- configuring
 - Agent
 - using config.sh 4–6
 - pager notification B–13
- configuring devices
 - fabric topology 5–16
- connections 1–13
 - naming 1–13
- containers
 - attributes 2–14
 - comparison 2–15
 - illustrated 2–14, 5–17
 - mirrorsets 2–21
 - planning storage 2–14
 - stripesets 2–19
- controller
 - verification of installation 5–16
- controller verification
 - installation 5–9, 5–16
- controllers
 - cabling 5–4, 5–11
 - location 1–2, 1–3

- node IDs 1–27
- verification of installation 5–9, 5–16
- worldwide names 1–27
- creating
 - storageget and device profiles 2–16
- creating and tuning
 - filesystem 3–15
- Creating Clones for Backup 7–2
- creation
 - file system 3–15

D

- Destroy/Nodestroy parameters 2–32
- device switches
 - changing
 - fabric topology 5–26
- devices
 - changing switches
 - fabric topology 5–25
 - configuration
 - fabric topology 5–16
 - creating a profile 2–16
- disk drives
 - adding
 - fabric topology 5–23
 - adding to the spareset
 - fabric topology 5–23
 - array 2–15
 - corresponding storagegets 2–33
 - dividing 2–26
 - removing from the spareset
 - fabric topology 5–24
- displaying the current switches
 - fabric topology 5–25
- dividing storagegets 2–26

E

- enabling
 - switches 2–28
- erasing metadata 2–33
- establishing a local connection 5–2

F

- fabric topology
 - configuration
 - single controller cabling 5–4
- failover 1–5
 - multiple-bus 1–7
 - transparent 1–5
- FC considerations
 - loop and fabric 5–26
- file system
 - creation 3–15
 - tuning 3–15
- file systems
 - creating and tuning 3–15
- filesystem
 - creating and tuning 3–15
 - preparing LUNs 3–13
- First enclosure of multiple-enclosure subsystem
 - storage map template 1 A–4, A–7, A–9, A–11, A–14, A–16, A–19
- format utility
 - labeling LUNs 3–13
- functions
 - Agent 4–1

G

- geometry
 - initialize switches 2–33
- Geometry parameters 2–33

H

- Host access
 - restricting in multiple-bus failover mode
 - disabling access paths 1–24
- host access
 - restricting by offsets
 - multiple-bus failover 1–26
 - transparent failover 1–23
 - restricting in multiple-bus failover mode
 - 1–24
 - restricting in transparent failover mode
 - 1–21
 - disabling access paths 1–23

- separate links 1–22
- host adapter
 - installation 3–6
 - preparation 3–6
- host connections 1–13
 - naming 1–13
- HSG Agent
 - configuration menu 4–6
 - install and configure 4–1
 - install.sh 4–5
 - network connection 4–3
 - overview 4–2
- I**
- initialize switches
 - changing
 - fabric topology 5–26
 - CHUNKSIZE 2–30
 - geometry 2–33
 - NOSAVE_CONFIGURATION 2–32
 - SAVE_CONFIGURATION 2–32
- Insight Manager B–13
- Install
 - automatic mode 3–10
 - manual mode 3–11
- Install Modes
 - automatic or manual 3–8
- installation
 - automatic mode 3–8, 3–10
 - controller verification 5–9, 5–16
 - invalid network port assignments B–8
 - manual mode 3–8, 3–11
 - solution software 3–9
 - there is no disk in the drive message B–9
- installation verification
 - CLI commands 5–9, 5–16
- installing
 - Agent 4–5
- integrating SWCC B–13
- invalid network port assignments B–8
- J**
- JBOD 2–15

- L**
- labeling LUNs
 - format utility 3–13
- LOCATE
 - find devices 2–34
- location
 - cache module 1–2, 1–3
 - controller 1–2, 1–3
- LUN IDs
 - general description 1–29
- LUN presentation 1–18
- LUNs
 - labeling 3–13
 - using format utility 3–13
 - preparing for use by filesystem 3–13
- M**
- mail messages
 - RAIDManager 4–9
- maintenance port connection
 - establishing a local connection 5–2
 - illustrated 5–2
- manual mode
 - installation 3–11
- mapping storagesets 2–33
- messages
 - there is no disk in the drive B–9
- mirrored caching
 - enabling 1–10
 - illustrated 1–11
- mirrorset switches
 - changing
 - fabric topology 5–25
- mirrorsets
 - planning considerations 2–21
 - important points 2–22
 - switches 2–29
- Model 2200 Storage Maps
 - examples 2–7
- moving storagesets 7–7
- multiple-bus failover
 - restricting host access by offsets 1–26

- multiple-bus failover 1–7
 - ADD CONNECTIONS command 1–19
 - ADD UNIT command 1–19
 - ADD UNITcommand 1–19
 - CLI configuration procedure
 - fabric topology 6–4
 - fabric topology configuration
 - cabling 5–10
 - host connections 1–19
 - restricting host access 1–24
 - disabling access paths 1–24
 - SET CONNECTIONS command 1–19
 - SET UNITcommand 1–19

N

- network port assignments B–8
- node IDs 1–27
 - restoring 1–28
- NODE_ID
 - worldwide name 1–27
- NOSAVE_CONFIGURATION 2–32

O

- offset
 - LUN presentation 1–18
 - restricting host access
 - multiple-bus failover 1–26
 - transparent failover 1–23
 - SCSI version factor 1–18
- online help
 - SWCC B–13
- options
 - for mirrorsets 2–29
 - for RAIDsets 2–28
 - initialize 2–30
- other controller 1–3

P

- pager notification B–13
 - configuring B–13
- partitions
 - assigning a unit number
 - fabric topology 5–22

- defining 2–27
- planning considerations 2–26
 - guidelines 2–27
- passwords
 - choosing 4–7
- performance 2–22
- Physical connection, making 3–6
- planning 2–1
 - overview 2–16
 - striped mirrorsets 2–25
 - stripesets 2–19
- Planning a subsystem 1–1
- planning configurations
 - where to start 2–2
- planning considerations 2–22
- planning storage
 - containers 2–14
- planning storage sets
 - characteristics
 - changing switches 2–28
 - enabling switches 2–28
 - initialization switch 2–27
 - storage switch 2–27
 - unit switch 2–27
 - switches
 - initialization 2–29
 - storage set 2–28
- profiles
 - creating 2–16
 - description 2–16
 - storage set A–1
 - example A–2

R

- RAID Manager mail messages 4–9
- RAIDManager
 - mail messages 4–9
- RAIDset switches
 - changing
 - fabric topology 5–25
- RAIDsets
 - choosing chunk size 2–30

- maximum membership 2–24
- planning considerations 2–22
 - important points 2–23
- switches 2–28
- read caching
 - enabled for all storage units 1–10
 - general description 1–10
- read requests
 - decreasing the subsystem response time with read caching 1–10
- read-ahead caching 1–10
 - enabled for all disk units 1–10
- removing
 - Client B–12
- request rate 2–30
- requirements
 - host adapter installation 3–6
- restarting
 - Agent 4–8
- restricting host access
 - disabling access paths
 - multiple-bus failover 1–24
 - transparent failover 1–23
 - multiple-bus failover 1–24
 - separate links
 - transparent failover 1–22
 - transparent failover 1–21
- S**
- SAVE_CONFIGURATION 2–32
- saving configuration 2–32
- SCSI version
 - offset 1–18
- SCSI-2
 - assigning unit numbers 1–19
 - command console lun 1–20
- SCSI-3
 - assigning unit numbers 1–19
 - command console lun 1–20
- Second enclosure of multiple-enclosure subsystem
 - storage map template 2 A–5
- selective storage presentation 1–21
- SET CONNECTIONS
 - multiple-bus failover 1–19
 - transparent failover 1–17
- SET UNIT
 - multiple-bus failover 1–19
- setting
 - controller configuration handling 2–32
- single disk (JBOD)
 - assigning a unit number
 - fabric topology 5–22
- Single-enclosure subsystem
 - storage map template 1 A–4
- solution software 3–7
 - installation 3–9
- storage
 - creating map 2–33
 - profile
 - example A–2
- storage configurations 2–1
- storage map 2–33
- Storage map template 1 A–4
 - first enclosure of multiple-enclosure subsystem A–4
 - single enclosure subsystem A–4
- Storage map template 2 A–5
 - second enclosure of multiple-enclosure subsystem A–5
- Storage map template 3 A–6
 - third enclosure of multiple-enclosure subsystem A–6
- Storage map template 4
 - first enclosure of multiple-enclosure subsystem A–7
- Storage map template 5
 - first enclosure of multiple-enclosure subsystem A–9
- Storage map template 6
 - first enclosure of multiple-enclosure subsystem A–11
- Storage map template 7

- first enclosure of multiple-enclosure subsystem A-14
 - Storage map template 8
 - first enclosure of multiple-enclosure subsystem A-16
 - Storage map template 9
 - first enclosure of multiple-enclosure subsystem A-19
 - storageset
 - deleting
 - fabric topology 5-25
 - fabric topology
 - changing switches 5-25
 - planning considerations 2-18
 - mirrorsets 2-21
 - partitions 2-26
 - RAIDsets 2-22
 - striped mirrorsets 2-24
 - stripesets 2-18
 - profile 2-16
 - profiles A-1
 - storageset profile 2-16
 - storageset switches
 - SET command 2-28
 - storagesets
 - creating a profile 2-16
 - moving 7-7
 - striped mirrorsets
 - planning 2-25
 - planning considerations 2-24
 - stripesets
 - distributing members across buses 2-20
 - planning 2-19
 - planning considerations 2-18
 - important points 2-19
 - subsystem
 - adding 4-8
 - saving configuration 2-32
 - subsystem configuration
 - backup 7-1
 - SWCC 4-1
 - additional information B-13
 - integrating B-13
 - online help B-13
 - switches
 - changing 2-28
 - changing characteristics 2-27
 - CHUNKSIZE 2-30
 - enabling 2-28
 - mirrorsets 2-29
 - NOSAVE_CONFIGURATION 2-32
 - RAIDset 2-28
 - SAVE_CONFIGURATION 2-32
 - switches for storagesets
 - overview 2-28
 - system
 - preparing LUNs 3-13
- T**
- templates
 - subsystem profile A-1
 - terminology
 - other controller 1-3
 - this controller 1-3
 - Third enclosure of multiple-enclosure subsystem
 - storage map template 3 A-6
 - this controller 1-3
 - transparent failover 1-5
 - ADD CONNECTIONS 1-17
 - ADD UNIT 1-17
 - matching units to host connections 1-17
 - restricting host access 1-21
 - disabling access paths 1-23
 - separate links 1-22
 - restricting host access by offsets 1-23
 - SET CONNECTIONS 1-17
 - troubleshooting
 - invalid network port assignments B-8
 - there is no disk in the drive message B-9
 - tuning
 - file system 3-15
- U**
- uninstalling
 - Client B-12

unit numbers
 assigning 1–16
 fabric topology 5–22
 assigning depending on SCSI version 1–19
 assigning in fabric topology
 partition 5–22
 single disk 5–22

unit qualifiers
 assigning
 fabric topology 5–22

unit switches
 changing
 fabric topology 5–26

units
 LUN IDs 1–29
Upgrade procedures
 solution software 3–16

V

verification
 controller installation 5–9, 5–16
 verification of installation

 controller 5–9, 5–16
 Verifying/Installing Required Versions 3–8
 virtual disks
 adding B–13

W

where to start 1–1
worldwide names 1–27
 NODE_ID 1–27
 REPORTED PORT_ID 1–27
 restoring 1–28
write performance 2–32
write requests
 improving the subsystem response time with
 write-back caching 1–10
 placing data with with write-through caching
 1–10
write-back caching
 general description 1–10
write-through caching
 general description 1–10