

**MA6000 HSG60 Array Controller  
ACS Version 8.5 Solution Software  
for Sun Solaris**

Installation and Configuration Guide

First Edition (May 2000)  
Part Number: AA-RMBMA-TE  
**Compaq Computer Corporation**

## Notice

© 2000 Compaq Computer Corporation.

COMPAQ, the Compaq logo, and StorageWorks are Registered in U. S. Patent and Trademark Office. NonStop is a trademark and/or service mark of Compaq Information Technologies Group, L.P.

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

UNIX is a trademark of The Open Group.

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

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

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE FOREGOING SHALL APPLY REGARDLESS OF THE NEGLIGENCE OR OTHER FAULT OF EITHER PARTY AND REGARDLESS OF WHETHER SUCH LIABILITY SOUNDS IN CONTRACT, NEGLIGENCE, TORT, OR ANY OTHER THEORY OF LEGAL LIABILITY, AND NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Printed in the U.S.A.

MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris  
Installation and Configuration Guide  
First Edition (May 2000)  
Part Number: AA-RMBMA-TE

# Table of Contents

## About This Guide

### Chapter 1

#### Planning a Subsystem

Controller Designations A and B .....	1-1
Controller Designations “This Controller” and “Other Controller” .....	1-2
Selecting a Failover Mode .....	1-3
Transparent Failover Mode .....	1-4
Multiple-Bus Failover Mode .....	1-7
Selecting a Cache Mode .....	1-9
Read Caching .....	1-9
Read-Ahead Caching .....	1-9
Write-Back Caching .....	1-9
Write-Through Caching .....	1-10
Enabling Mirrored Caching .....	1-10
The Command Console LUN .....	1-11
Determining the Address of the CCL .....	1-11
In SCSI II Mode: .....	1-11
Disabling the CCL .....	1-11
Enabling the CCL .....	1-12
In SCSI III Mode: .....	1-12
Connections .....	1-12
Naming Connections .....	1-12
Numbers of Connections .....	1-13
Assigning Unit Numbers .....	1-17

**iv**      *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

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 . . . . .	1-20
The CCL in SCSI-3 mode. . . . .	1-20
The CCL in SCSI-2 Mode . . . . .	1-20
Restricting Host Access (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
Restricting Host Access by Disabling Access Paths . . . . .	1-24
Restricting Host Access by Offsets . . . . .	1-27
Worldwide Names (Node IDs and Port IDs) . . . . .	1-28
Restoring Worldwide Names (Node IDs) . . . . .	1-29

## **Chapter 2**

### **Planning Storage**

Where to Start . . . . .	2-2
Configuration Rules . . . . .	2-3
Device PTL Addressing Convention. . . . .	2-4
Determining Storage Requirements . . . . .	2-7
Choosing a Container Type. . . . .	2-7
Creating a Storageset Profile. . . . .	2-10
Storageset Planning Considerations . . . . .	2-12
Stripeset Planning Considerations . . . . .	2-12
Mirrorset Planning Considerations . . . . .	2-15
RAIDset Planning Considerations . . . . .	2-16
Striped Mirrorset Planning Considerations . . . . .	2-18
Partition Planning Considerations. . . . .	2-19
Defining a Partition . . . . .	2-20
Guidelines for Partitioning Storagesets and Disk Drives. . . . .	2-21
Changing Characteristics through Switches . . . . .	2-21
Enabling Switches . . . . .	2-22
Changing Switches. . . . .	2-22
Storageset and Partition Switches . . . . .	2-22
RAIDset Switches. . . . .	2-23
Mirrorset Switches . . . . .	2-23
Partition Switches . . . . .	2-23

Initialization Switches . . . . .	2-24
Chunk Size . . . . .	2-24
Increasing the Request Rate . . . . .	2-24
Increasing the Data Transfer Rate . . . . .	2-26
Increasing Sequential Write Performance . . . . .	2-27
Save Configuration . . . . .	2-27
Destroy/Nodestry . . . . .	2-27
Geometry . . . . .	2-28
Unit Switches . . . . .	2-28
Storage Maps . . . . .	2-28
Creating a Storage Map . . . . .	2-28
Example Storage Map . . . . .	2-29
Using the LOCATE Command to Find Devices . . . . .	2-31

## Chapter 3

### Preparing the Host System

Making a Physical Connection . . . . .	3-1
Preparing to Install the Host Adapter . . . . .	3-1
Installing the Host Adapter . . . . .	3-2
Verifying/Installing Required Versions . . . . .	3-3
Deciding between Automatic and Manual Installation Modes . . . . .	3-4
Installing the Solution Software Packages . . . . .	3-5
Mounting the CD-ROM . . . . .	3-5
If vold is running: . . . . .	3-5
If vold is not running: . . . . .	3-6
Installing with Automatic Mode . . . . .	3-6
Installing with Manual Mode . . . . .	3-8
Preparing LUNs for Use by the FileSystem . . . . .	3-9
Labeling LUNs . . . . .	3-9
Labeling a LUN Automatically Using the Format Utility . . . . .	3-10
File System Creation and Tuning . . . . .	3-12
. . . . .	3-12

## Chapter 4

### Installing and Configuring the HS-Series Agent

Why Use StorageWorks Command Console? . . . . .	4-1
About the Network Connection for the Agent . . . . .	4-3
Configuring the Agent . . . . .	4-5
Configuring the Agent with install.sh (for First-Time Configurations) . . . . .	4-5
Configuring the Agent with config.sh . . . . .	4-6
Choosing a Password . . . . .	4-7

Adding a Subsystem Entry . . . . .	4-7
Adding a Client System Entry . . . . .	4-8
Restarting the SWCC Agent . . . . .	4-8
Configuring the Agent within FirstWatch . . . . .	4-8
Raid Manager Mail Messages . . . . .	4-9

## Chapter 5

### Configuration Procedures

Establishing a Local Connection . . . . .	5-2
Establishing Connection with a SPARC System . . . . .	5-3
Setting Up a Single Controller . . . . .	5-4
Power On and Establish Communication . . . . .	5-4
Cabling a Single controller . . . . .	5-5
Configuring a Single Controller using CLI . . . . .	5-6
Verifying Installation . . . . .	5-9
Setting Up a Controller Pair . . . . .	5-10
Power On and Establish Communication . . . . .	5-10
Cabling a Controller Pair . . . . .	5-11
Configuring a Controller Pair using CLI . . . . .	5-12
Verifying Installation . . . . .	5-16
Configuring Devices . . . . .	5-16
Configuring a Stripeset . . . . .	5-17
Configuring a Mirrorset . . . . .	5-17
Configuring a RAIDset . . . . .	5-18
Configuring a Striped Mirrorset . . . . .	5-19
Configuring a Single-Disk (JBOD) Unit . . . . .	5-20
Configuring a Partition . . . . .	5-20
Assigning Unit Numbers and Unit Qualifiers . . . . .	5-21
Assigning a Unit Number to a Storageset . . . . .	5-21
Assigning a Unit Number to a Single (JBOD) Disk . . . . .	5-22
Assigning a Unit Number to a Partition . . . . .	5-22
Configuration Options . . . . .	5-22
Changing the CLI Prompt . . . . .	5-22
Mirroring cache . . . . .	5-22
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 . . . . .	5-25
Changing Device Switches . . . . .	5-26
Changing Initialize Switches. . . . .	5-26
Changing Unit Switches . . . . .	5-26

*Chapter 6*

**Managing Storage Access for Both Loop and Fabric Environments**

General Considerations. . . . .	6-1
The Configurable Objects - Servers . . . . .	6-2
The Configurable Objects - Storage. . . . .	6-3
A Special Note on Controller Pairs (Active-Passive vs. Active-Active) . . . . .	6-3
The Configurable Objects - Hubs. . . . .	6-4
The Configurable Objects - Switches. . . . .	6-5
The Fibre Channel Access/Transport paths: Fibre Fabric vs. Fibre Loop. . . . .	6-6
Sun Solaris Implementation of Loop and Fabric. . . . .	6-6
Bindings . . . . .	6-6
Loop Bindings. . . . .	6-7
Fabric Bindings . . . . .	6-7
Loop Connections at the Solaris Level. . . . .	6-8
Loop Mappings. . . . .	6-8
Analysis. . . . .	6-9
Fabric Mappings. . . . .	6-9
Analysis. . . . .	6-10
The /opt/steam/bin/config.sh utility . . . . .	6-10
Option 1 - View Adapters . . . . .	6-12
Option 2 - Add an Adapter . . . . .	6-12
Option 3 - Remove an Adapter . . . . .	6-12
Option 4 - Modify an Adapter. . . . .	6-13
Configuration Recipes . . . . .	6-16
Loop Configuration(s) . . . . .	6-16
Fabric Configuration(s) . . . . .	6-17
Loop to Fabric Reconfiguration . . . . .	6-17
Fabric to Loop Reconfiguration . . . . .	6-18

*Chapter 7*

**Configuration Example Using CLI**

CLI Configuration Example. . . . .	7-4
------------------------------------	-----

## Chapter 8

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

Backing Up the Subsystem Configuration . . . . .	8-1
Cloning Data for Backup . . . . .	8-2
Moving Stagesets . . . . .	8-6

## Appendix A

### Subsystem Profile Templates

Stageset Profile . . . . .	A-2
Storage Map Template 1 . . . . .	A-3
Storage Map Template 2 . . . . .	A-4
Storage Map Template 3 . . . . .	A-5

## Appendix B

### Installing, Configuring, and Removing the Client

Why Install the Client? . . . . .	B-1
Before You Install the Client . . . . .	B-2
Installing the Client . . . . .	B-3
Troubleshooting the Client Installation . . . . .	B-3
Invalid Network Port Assignments	
During Installation . . . . .	B-3
“There is no disk in the drive” Message . . . . .	B-4
Adding the Storage Subsystem and its Host to the Navigation Tree . . . . .	B-5
Removing the Command Console Client . . . . .	B-7
Where to Find Additional Information . . . . .	B-8
About the User Guide . . . . .	B-8
About the Online Help . . . . .	B-8

## Index



# Figures

- Figure 1-1. Location of controllers and cache modules . . . . . 1-2
- Figure 1-2. “This controller” and “other controller” . . . . . 1-3
- Figure 1-3. Transparent failover - normal operation . . . . . 1-5
- Figure 1-4. Transparent failover - after failover from controller B to controller A. . . . . 1-6
- Figure 1-5. Typical multiple-bus configuration . . . . . 1-8
- Figure 1-6. Mirrored caching . . . . . 1-10
- Figure 1-7. Connections in separate-link, transparent failover mode configurations . . . . . 1-14
- Figure 1-8. Connections in single-link, transparent failover mode configurations . . . . . 1-15
- Figure 1-9. Connections in multiple-bus failover mode . . . . . 1-16
- Figure 1-10. LUN presentation to hosts, as determined by offset . . . . . 1-18
- Figure 1-11. Limiting host access in transparent failover mode . . . . . 1-22
- Figure 1-12. Limiting host access in multiple-bus failover mode . . . . . 1-25
- Figure 1-13. Placement of the worldwide name label . . . . . 1-29
- Figure 2-1. PTL naming convention . . . . . 2-4
- Figure 2-2. PTL addressing in a single-bus configuration, using two shelves . . . . . 2-5
- Figure 2-3. PTL addressing in a dual-bus configuration . . . . . 2-5
- Figure 2-4. Mapping a unit to physical disk drives . . . . . 2-6
- Figure 2-5. Container types . . . . . 2-8
- Figure 2-6. An example storageset profile . . . . . 2-11
- Figure 2-7. A 3-member RAID 0 stripeset (example 1) . . . . . 2-12
- Figure 2-8. A 3-member RAID 0 stripeset (example 2) . . . . . 2-13
- Figure 2-9. Mirrorsets maintain two copies of the same data . . . . . 2-15
- Figure 2-10. Mirrorset example 2 . . . . . 2-15
- Figure 2-11. A five member RAIDset using parity . . . . . 2-17
- Figure 2-12. Striped mirrorset (example 1) . . . . . 2-18

Figure 2–13. Striped mirrorset (example 2) . . . . .	2–19
Figure 2–14. One example of a partitioned single-disk unit . . . . .	2–20
Figure 2–15. Chunk size larger than the request size. . . . .	2–25
Figure 2–16. Chunk size smaller than the request size . . . . .	2–26
Figure 2–17. Example storage map-single bus . . . . .	2–29
Figure 2–18. Example storage map-dual bus. . . . .	2–30
Figure 4–1. An example of a network connection . . . . .	4–4
Figure 4–2. Configuration menu . . . . .	4–6
Figure 5–1. Maintenance port connection . . . . .	5–2
Figure 5–2. Single controller cabling. . . . .	5–5
Figure 5–3. Controller pair failover cabling. . . . .	5–11
Figure 7–1. Example storage map . . . . .	7–2
Figure 7–2. Example system . . . . .	7–3
Figure 7–3. Example virtual layout from the hosts’ point of view . . . . .	7–3
Figure 8–1. Steps the CLONE utility follows for duplicating unit members . . . . .	8–3
Figure B–1. Navigation window . . . . .	B–5
Figure B–2. Navigation window showing storage host system “Atlanta”. . . . .	B–6
Figure B–3. Navigation window showing expanded “Atlanta” host icon . . . . .	B–6

## Tables

Table 1-1 Unit Assignments and SCSI_VERSION.....	1-21
Table 2-1 A Comparison of Container Types .....	2-9
Table 2-2 Example Chunk Sizes .....	2-26
Table 3-1 StorageWorks Solution Software Packages .....	3-2
Table 4-1 SWCC Features and Components .....	4-2
Table 4-2 Installation and Configuration Overview .....	4-3
Table 6-1 Case 1: Active-Passive (Describes the controller pair) .....	6-3
Table 6-2 Case 2: Active-Active (Describes the controller pair) .....	6-4
Table 6-3 Example of Active-Passive with a Single Controller .....	6-4



# About This Guide

## Conventions

### Text Conventions

Table 1 Text Conventions	
Convention	Meaning
<b>Bold</b>	Keyboard keys appear in boldface. For example: <b>Enter/Return</b> or <b>Y</b> (es) key
SMALL CAPS	Used to indicate the status of an LED. For example: FLASHING on and off or ON steady
ALLCAPS	Command syntax that must be entered exactly as shown and for commands discussed within text, for example: SET FAILOVER COPY=OTHER_CONTROLLER “Use the SHOW SPARESET command to show the contents of the spareset.”
Monospaced	Screen display.
<i>italic</i>	User entered command variable or numeric value, for example: SHOW RAIDset-name or set this_controller id=(n,n,n,n,n) Reference to other publications, for example: “See the HSG60 Array Controller ACS Version 8.5 Release Notes for details.”

Table 1 Text Conventions (Continued)	
Convention	Meaning
.	Indicates that a portion of an example or figure has been omitted.
“this controller”	The controller serving the current CLI session through a local or remote terminal.
“other controller”	The controller in a dual-redundant pair that’s connected to the controller serving the current CLI session.

## Special Notices

This guide does not contain detailed descriptions of standard safety procedures. However, it does contain warnings for procedures that might cause personal injury and cautions for procedures that might damage the controller or its related components. Look for these symbols when performing the procedures in this guide:



**WARNING:** A warning indicates the presence of a hazard that can cause personal injury if precautions in the text are not observed.



**CAUTION:** A caution indicates the presence of a hazard that might damage hardware, corrupt software, or cause a loss of data.

**IMPORTANT:** An *important* note is a type of note that provides information essential to the completion of a task. Users can disregard information in a note and still complete a task, but they should not disregard an important note.

**NOTE:** A note provides additional information that is related to the completion of an instruction or procedure.

## Related Publications

The following table lists some of the documents related to the use of the controller, cache module, and external cache battery.

<b>Table 2 Relevant Publications</b>	
<b>Document Title</b>	<b>Part Number</b>
<i>Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide</i>	EK-HSG84-SV. D01
<i>Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide</i>	EK-HSG85-RG. B01
<i>Compaq StorageWorks Command Console Version 2.3 for the MA6000 User Guide</i>	AA-RMBDA-TE
<i>Compaq StorageWorks Fibre Channel Storage Switch Service Guide</i>	AA-RHBZA-TE / 135268-001
<i>Compaq StorageWorks Fibre Channel Storage Switch User Guide</i>	AA-RHBYA-TE / 135267-001
<i>Compaq StorageWorks Fibre Channel Switch Quick Setup Guide</i>	AA-RHC0A-TE / 135269-001
<i>Compaq StorageWorks KGPSA PCI-to-Fibre Channel Host Adapter</i>	EK-KGPSA-UG
<i>The RAIDBOOK—A Source Book for Disk Array Technology</i>	RAID Advisory Board: I.S.B.N 1-879936-90-9
<i>Enclosure 4200 Family LVD Disk Enclosures User Guide</i>	EK-SW2ZS-UA.B01 / 148451-002
<i>Model 2100 and 2200 Ultra SCSI Controller Enclosure</i>	EK-SE2C8-UA.B01 / 127422-002

## Revision History

First Release: May 2000





# Chapter 1

## Planning a Subsystem

This chapter provides information that helps you plan the configuration of the subsystem. Read Chapter 2 to plan the types of storage containers you need.

This chapter frequently references the command line interface (CLI). For the complete syntax and descriptions of the CLI commands, see the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

The following information is included in this chapter:

- “Controller Designations A and B,” page 1–1
- “Selecting a Failover Mode,” page 1–3
- “Selecting a Cache Mode,” page 1–9
- “Enabling Mirrored Caching,” page 1–10
- “The Command Console LUN,” page 1–11
- “Connections,” page 1–12
- “Assigning Unit Numbers,” page 1–17
- “Restricting Host Access (Selective Storage Presentation),” page 1–21
- “Worldwide Names (Node IDs and Port IDs),” page 1–28

### Controller Designations A and B

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

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




Figure 1-1. Location of controllers and cache modules

## Controller Designations “This Controller” and “Other Controller”

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

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

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




Figure 1-2. “This controller” and “other controller”

## Selecting a Failover Mode

Failover is a way to keep the storage array available to the host if one controller becomes unresponsive. A controller can become unresponsive because of a hardware failure such as a host bus adapter or controller or, in multiple-bus only, to a failure of the link between host and controller. 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 transparent (invisible) to the hosts
- **multiple-bus**, which is handled by the hosts

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 unit is a physical or virtual device in the MA6000 RAID system. It is typically assigned a logical unit number (LUN) and is managed by the HSG60 controller and presented to a server through the Fibre Channel bus and servers host bus adapter. Disks which are set up as JBODs or RAIDsets are referred to as storagesets. Storagesets are units.

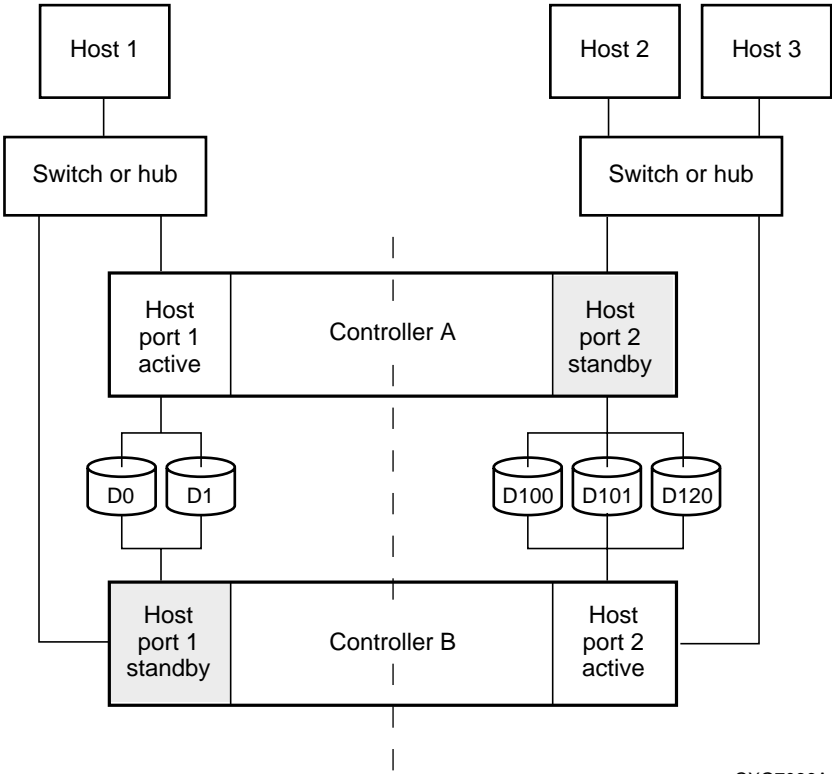
In transparent failover mode, host port 1 of controller A and host port 1 of controller B must be on the same Fibre Channel link. Host port 2 of controller A and host port 2 of controller B must also be on the same Fibre Channel link. Depending on operating system restrictions and requirements, the port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

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

Units are divided between the host ports:

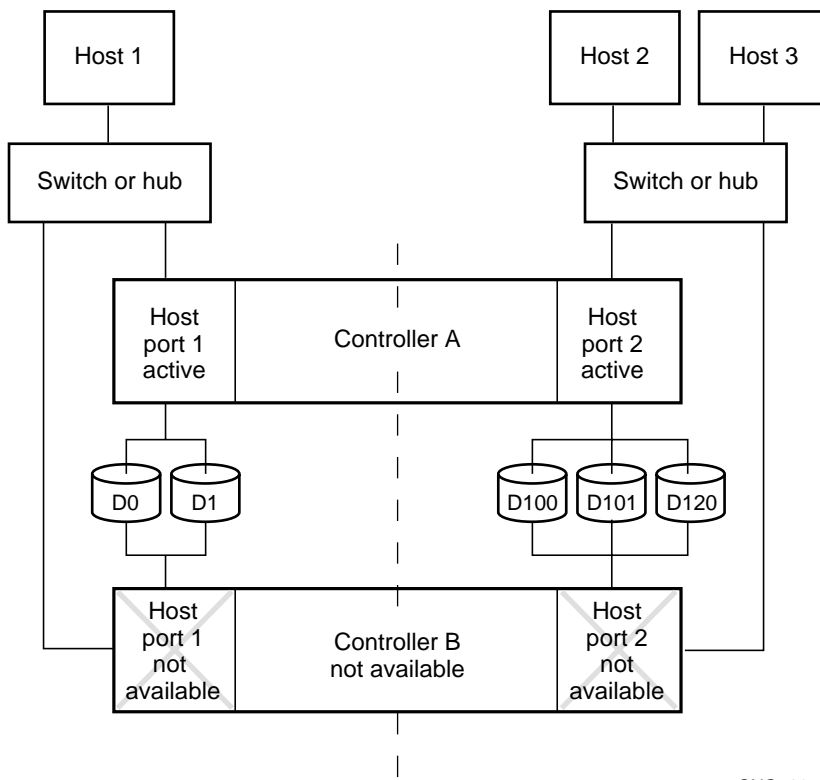
- Units 0-99 are on host port 1 of both controllers (but accessible only through the active port)
- Units 100-199 are on host port 2 of both controllers (but accessible only through the active port)

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



CXO7036A

Figure 1-3. Transparent failover - normal operation



CXO7035A

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

## Multiple-Bus Failover Mode

Multiple-bus failover mode has the following characteristics:

- host controls the failover process by moving the unit(s) 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

Each host must have special software to control failover. 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 1-5.


In multiple-bus failover mode, you can specify which units are normally serviced by a specific controller of a controller pair. This process is called preferring or preferment. Units can be preferred to one controller or the other by the *PREFERRED\_PATH* switch of the ADD (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 re-distribute 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

CXO7094A

Figure 1-5. Typical multiple-bus configuration



## Selecting a Cache Mode

The cache module supports read, read-ahead, write-through, and write-back caching techniques. 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 give improved 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

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. Write-through caching is enabled when write-back caching is disabled.

## Enabling Mirrored Caching

In mirrored caching, half of each controller's cache mirrors the companion controller's cache, as shown in Figure 1-6.

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




Figure 1-6. Mirrored caching

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

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

## The Command Console LUN

StorageWorks Command Console (SWCC) Software communicates with HSG60 controllers through an existing storage unit, or LUN (logical unit number). The dedicated LUN that SWCC uses is called the CCL (Command Console LUN). CCL serves as the communication device for the HS-Series Agent and identifies itself to the host by a unique identification string. By default, a CCL device is enabled within the HSG60 controller on Host Port 1. There are two modes that HSG60 uses, SCSI-2 and SCSI-3. The Command Console LUN (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 general, you will need to create a storageset to be used as a CCL for Port 2. This is a requirement for the HS-Series Agent.

In dual-redundant controller configurations, the following commands alter the setting of the CCL on both controllers. The CCL is enabled only on Host Port 1. At least one storage device must be configured on Host Port 2 before installing the Agent on a host connected to Host Port 2. This can be any type of storage device.

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

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

```
HSG60 > SHOW THIS_CONTROLLER
```

### In SCSI II Mode:

#### Disabling the CCL

To disable the CCL, enter the following CLI command:

```
HSG60 > SET THIS_CONTROLLER NOCOMMAND_CONSOLE_LUN
```

### **Enabling the CCL**

To enable the CCL, enter the following CLI command:

```
HSG60 > SET THIS_CONTROLLER_COMMAND_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.

### **In SCSI III Mode:**


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

## **Connections**

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

### **Naming Connections**

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



**Examples:**

A connection from the first adapter in host SPGS that goes to port 1 of controller A would be called SPGS1A1.

A connection from the third adapter in host LONDON that goes to port 2 of controller B would be called LONDON3B2.

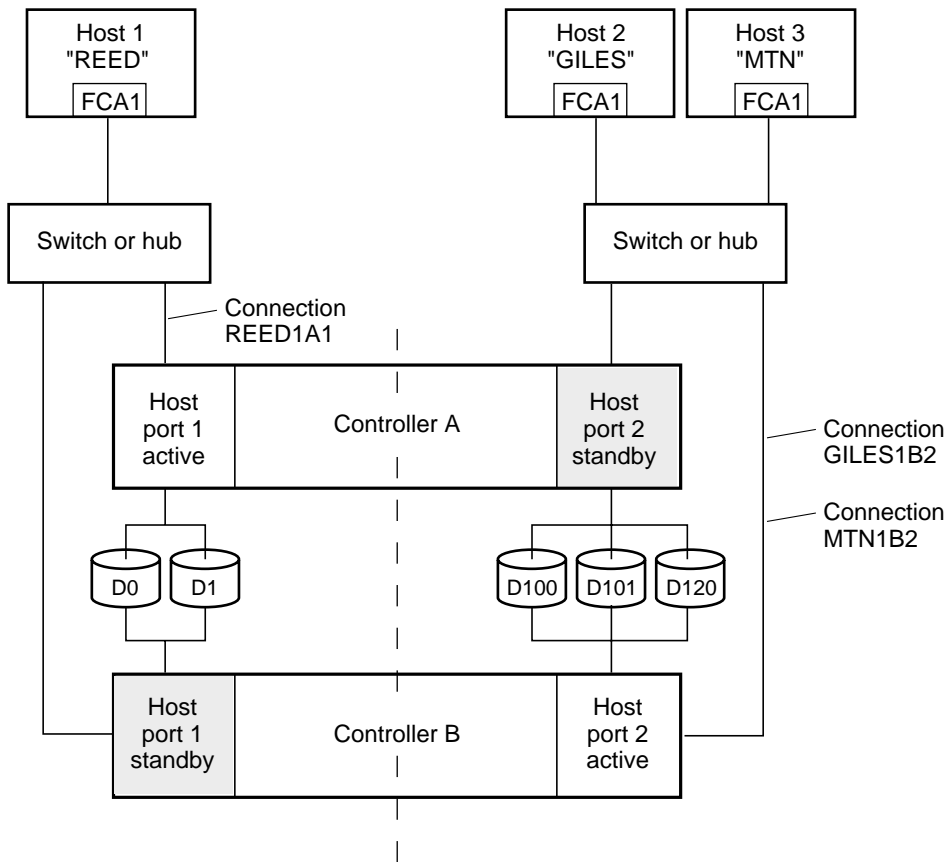
**NOTE:** Connection names can have a maximum of 9 characters.

## **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 1-7.
- 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 1-8.
- If a controller pair is in multiple-bus failover mode, each adapter has two connections, as shown in Figure 1-9.

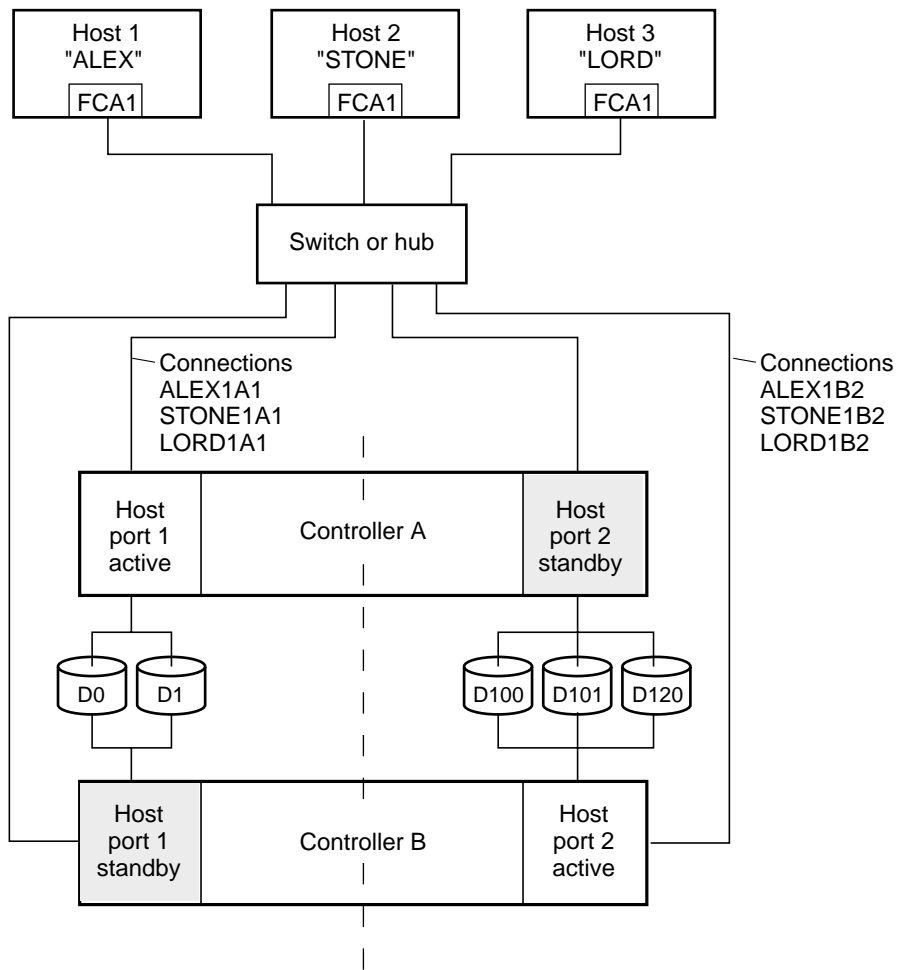
**1-14** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide



NOTE: FCA = Fibre Channel Adapter

CXO7081A


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



NOTE: FCA = Fibre Channel Adapter

CXO7079A

Figure 1-8. Connections in single-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CXO7080A

Figure 1-9. Connections in multiple-bus failover mode



## Assigning Unit Numbers

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

- failover mode of the controller pair
- the *ENABLE\_ACCESS\_PATH* and *PREFERRED\_PATH* switches in the ADD UNIT (or SET *unit*) commands
- the *UNIT\_OFFSET* switch in the ADD CONNECTIONS (or SET *connections*) commands
- which controller port the connection is attached to
- 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(s) to access and assigns it to either port 1 of both controllers or to port 2 of both controllers.

Unit numbers are assigned to ports as follows:

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

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

What 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 relationship of offset, LUN number, and unit number is as follows:

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

- 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)

**1-18** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

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 102 will be presented to a port 2 host connection as LUN 2 (unit number of 102 minus offset of 100).

Figure 1-10 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

CX06455B

Figure 1-10. 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 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).

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 at all to a host connection with a unit offset of 128 or greater, because that offset is not within the units range (unit number of 127 minus offset of 128 is a negative number).

An additional factor to consider when assigning unit numbers and offsets is SCSI version. If the *SCSI\_VERSION* switch of the SET *this\_controller/other\_controller* command is set to SCSI-3, the Command Console LUN (CCL) is presented as LUN 0 to every connection, superseding any unit assignments. The interaction between SCSI version and unit numbers is explained further in “Assigning Unit Numbers Depending on SCSI\_VERSION” on page 1-20.

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.

What 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 as follows:

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

- 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)

For example, unit D17 would be visible to a host connection with an offset of 0 as LUN 17 (unit number of 17 minus offset of 0). The same unit 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.

Which controller of a dual-redundant pair initially accesses the unit is determined by the *PREFERRED\_PATH* switch of the ADD UNIT (or SET *unit*) command. 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/other controller* command determines how the Command Console LUN (CCL) is presented. There are two choices: SCSI-2 and SCSI-3. The choice for *SCSI\_VERSION* effects how certain unit numbers and certain host connection offsets interact.

---

**IMPORTANT:** SCSI-2 mode should be used with Sun Solaris hosts. SCSI-3 mode will not work and is not supported.

---

### The CCL 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 are recommendations 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.

### The CCL 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.

Recommendations for assigning host connection offsets and unit numbers in SCSI-2 mode are as follows:

- 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 1-2 summarizes the recommendations for unit assignments based on the *SCSI\_VERSION* switch.

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

## Restricting Host Access (Selective Storage Presentation)

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

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

### Restricting Host Access in Transparent Failover Mode

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


1. using separate Fibre Channel links (either loop or fabric)
2. enabling the access path of selected host connections on a shared loop or fabric
3. 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 of limiting host access is to have one host or set of hosts on the port 1 link, and another host or set of hosts on the port 2 link. Each host can then see only units assigned to its respective controller port. This separation of host buses is shown in Figure 1-11. This method applies only if the host 1 link and host 2 link are separate links.

**NOTE:** Compaq highly recommends that you still give access to only specific connections. That way, if new connections are added, they will not have automatic access to all units. See "Restricting Host Access by Disabling Access Paths," page 1-23.



NOTE: FCA = Fibre Channel Adapter

CXO7081A

Figure 1-11. 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. Access path is a unit switch, meaning it must be specified for each unit. Default access is for the unit to be accessible to all hosts.

#### Example:

In Figure 1-11, restricting the access of unit D101 to host 3, the host named MTN, 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=MTN1B2
```

If the storage subsystem has more than one host connection, carefully specify the access path 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* 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.

The procedure of restricting access by enabling all access paths then disabling selected paths is particularly 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.

#### Example:

In Figure 1-11, assume both host connections on port 2 (connections GILES1B2 and MTN1B2) initially have the default port 2 offset of 100. Setting the offset of connection MTN1B2 to 120 will present unit D120 to host MTN as LUN 0.

```
SET MTN1B2 UNIT_OFFSET=120
```

## 1-24 MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

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

```
SET D120 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=MTN1B2
```

**NOTE:** It is highly recommended that you still give access to only specific connections, even if there is just one connection on the link. That 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.

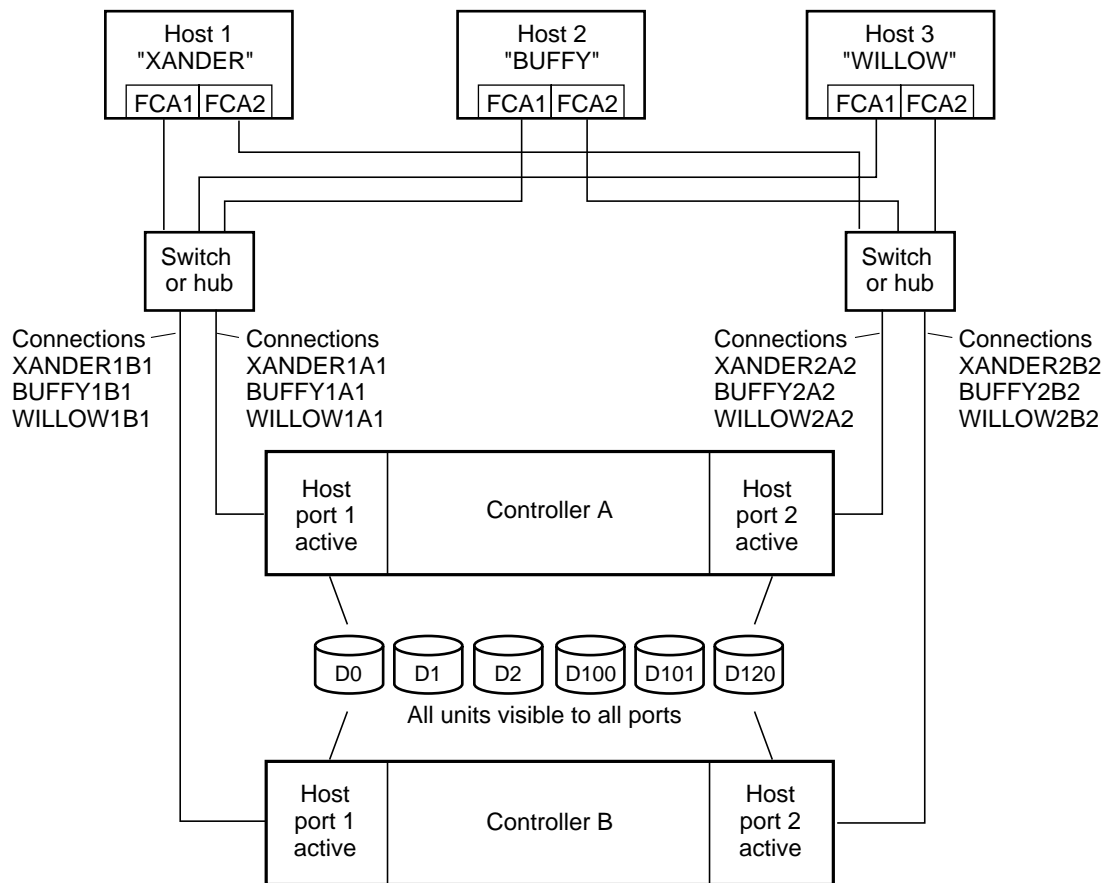
There are two ways to limit host access in multiple-bus failover mode:

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

### Restricting Host Access by Disabling Access Paths

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 is for the unit to be accessible to all hosts. Also, 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

CXO7078A

Figure 1-12. Limiting host access in multiple-bus failover mode

**Example:**

Figure 1-12 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host WILLOW can be done by enabling only the connections to host WILLOW. 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 WILLOW, enter the following commands:

```
SET D101 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=(WILLOW1A1,WILLOW1B1,WILLOW2A2,WILLOW2B2)
```

To enable only two connections for host WILLOW (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=(WILLOW1A1,WILLOW2B2)
```

or

```
SET D101 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=(WILLOW1B1,WILLOW2A2)
```

If the storage subsystem has more than one host connection, 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.

The procedure of restricting access by enabling all access paths then disabling selected paths is particularly 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.

### Example:

In Figure 1-12, assume all host connections initially have the default offset of 0. Giving all connections to host WILLOW an offset of 120 will present unit D120 to host WILLOW as LUN 0. Enter the following commands:

```
SET WILLOW1A1 UNIT_OFFSET=120
SET WILLOW1B1 UNIT_OFFSET=120
SET WILLOW2A2 UNIT_OFFSET=120
SET WILLOW2B2 UNIT_OFFSET=120
```

Host WILLOW 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 WILLOW, enable only host WILLOW's access, as follows:

```
SET D102 DISABLE_ACCESS_PATH=ALL
SET D102 ENABLE_ACCESS_PATH=(WILLOW1A1,WILLOW1B1,WILLOW12A2,WILLOW2B2)
```

**NOTE:** It is highly recommended that you always give access to only specific connections. That way, if new connections are added, they will not have automatic access to all units. See "Restricting Host Access by Disabling Access Paths," page 1-24.

## 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 increment 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 check sum printed on the sticker on the frame into which the controller is inserted, as shown in Figure 1-13.



**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 World Wide Names (LUN IDs)

In addition, each unit has its own world wide 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.




Figure 1-13. Placement of the worldwide name label



# *Chapter 2*

## **Planning Storage**

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

The following information is included in this chapter:

- “Where to Start,” page 2–2
- “Configuration Rules,” page 2–3
- “Determining Storage Requirements,” page 2–7
- “Choosing a Container Type,” page 2–7
- “Creating a Storageset Profile,” page 2–10
- “Storageset Planning Considerations,” page 2–12
- “Mirrorset Planning Considerations,” page 2–15
- “Partition Planning Considerations,” page 2–19
- “Changing Characteristics through Switches,” page 2–21
- “Storageset and Partition Switches,” page 2–22
- “Initialization Switches,” page 2–24
- “Unit Switches,” page 2–28
- “Storage Maps,” page 2–28

## Where to Start

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

1. Review configuration rules. See “Configuration Rules,” page 2-3.
2. Familiarize yourself with the current physical layout of the devices and their addressing scheme. See “Device PTL Addressing Convention,” page 2-4.
3. Determine your storage requirements. Use the questions in “Determining Storage Requirements,” page 2-7 to help you.
4. Choose the type of storage container(s) you need to use in your subsystem. See “Choosing a Container Type,” page 2-7 for a comparison and description of each type of storageset.
5. Create a storageset profile (described in “Creating a Storageset Profile,” page 2-10). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:
  - “Storageset Planning Considerations,” page 2-12
  - “Mirrorset Planning Considerations,” page 2-15
  - “Partition Planning Considerations,” page 2-19
  - “Striped Mirrorset Planning Considerations,” page 2-18
6. Decide which switches you will need for your subsystem. General information on switches is detailed in “Storageset and Partition Switches,” page 2-22
  - Determine what unit switches you want for your units (“Unit Switches,” page 2-28)
  - Determine what initialization switches you want for your planned storage containers (“Initialization Switches,” page 2-24)
7. Create a storage map (“Storage Maps,” page 2-28).
8. Configure the storage you have now planned using one of the following methods:
  - Use SWCC. See the SWCC documentation for details.
  - Use the Command Line Interpreter (CLI) commands. This method allows you flexibility in defining and naming your storage containers. See the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.



## **Configuration Rules**

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

- Maximum 128 assignable, host-accessible units
- Maximum 512-GB unit capacity
- Maximum 8 RAID-5 storage sets
- Maximum 30 RAID-5 and RAID-1 storage sets
- Maximum 45 RAID-5, RAID-1, and RAID-0 storage sets
- Maximum 8 partitions per storage set or individual disk
- Maximum 6 members per mirror set
- Maximum 14 members per RAID-5 storage set
- Maximum 24 members per stripe set
- Maximum 48 physical devices per striped mirror set

## Device PTL Addressing Convention

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

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

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




Figure 2-1. PTL naming convention

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

Model 4214 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	not supported	not supported

Model 4214 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	not supported	not supported

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

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

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

**2-6** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

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




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

## Determining Storage Requirements

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


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

## Choosing a Container Type

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

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

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



CXO6677A

Figure 2-5. Container types

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

<b>Table 2-1 A Comparison of Container Types</b>				
<b>Container Name</b>	<b>Relative Availability</b>	<b>Request Rate (Read/Write) I/O per second</b>	<b>Transfer Rate (Read/Write) MB per second</b>	<b>Applications</b>
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 (RAID1)	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 Stageset Profile**

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

See the example stageset profile shown in Figure 2-6.

Appendix A contains blank profiles that you can copy and use to record the details for your stagesets. Use the information in this chapter to help you make decisions when creating stageset profiles.



**Type of StorageSet:**

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

**StorageSet Name** *R1*

**Disk Drives** *D10300, D20300, D10400, D20400*

**Unit Number** *D101*

**Partitions:**

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

**RAIDset Switches:**

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

**Mirrorset Switches:**

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

**Initialize Switches:**

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

**Unit Switches:**

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

Figure 2-6. An example storageSet profile

## Storageset Planning Considerations


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

- Stripeset Planning Considerations
- Mirrorset Planning Considerations
- Partition Planning Considerations
- RAIDset Planning Considerations
- Striped Mirrorset Planning Considerations

### 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 allowing several disk drives to participate in one I/O request.

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



CXO7287A

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

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rate. You can set the chunk size or use the default setting (see “Chunk Size,” page 2-24, for information about setting the chunk size). Figure 2-8 shows another example of a three-member RAID 0 Stripeset.

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




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

Keep the following points in mind 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 will waste 5 GB of capacity on each 9 GB member.

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

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

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


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

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

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


## Mirrorset Planning Considerations

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



CXO7288A

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



CXO4594B

Figure 2-10. Mirrorset example 2

**2-16** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

Keep these points in mind as you plan your mirrorsets:

- Data availability with a mirrorset is excellent but comes with a higher cost—you need twice as many disk drives to satisfy a given capacity requirement. If availability is your top priority, consider using dual-redundant controllers and redundant power supplies.
- You can configure up to 20 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset may contain up to 6 members.
- Both write-back cache modules must be the same size.
- A mirrorset should only contain disk drives of the same capacity.
- Mirrorsets are well-suited for the following:
  - Any data for which reliability requirements are extremely high
  - Data to which high-performance access is required
  - Applications for which cost is a secondary issue
- Mirrorsets are not well-suited for the following applications:
  - Write-intensive applications (a performance hit of 10% will occur)
  - Applications for which cost is a primary issue

## **RAIDset Planning Considerations**

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




Figure 2-11. A five member RAIDset using parity

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

The relationship between the chunk size and the average request size determines if striping maximizes the request rate or the data-transfer rates. You can set the chunk size or use the default setting. See “Chunk Size,” page 2-24, for information about setting the chunk size.

Keep these points in mind as you plan your RAIDsets:

- Reporting methods and size limitations prevent certain operating systems from working with large RAIDsets.
- Both cache modules must be the same size.
- A RAIDset must include at least 3 disk drives, but no more than 14.

- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you will waste 5 GB of capacity on each 9 GB member.
- RAIDsets are particularly well-suited for the following:
  - ❑ Small to medium I/O requests
  - ❑ Applications requiring high availability
  - ❑ High read request rates
  - ❑ Inquiry-type transaction processing
- RAIDsets are not particularly well-suited for the following:
  - ❑ Write-intensive applications
  - ❑ Database applications in which fields are continually updated
  - ❑ Transaction processing

## Striped Mirrorset Planning Considerations

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




Figure 2-12. Striped mirrorset (example 1)



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

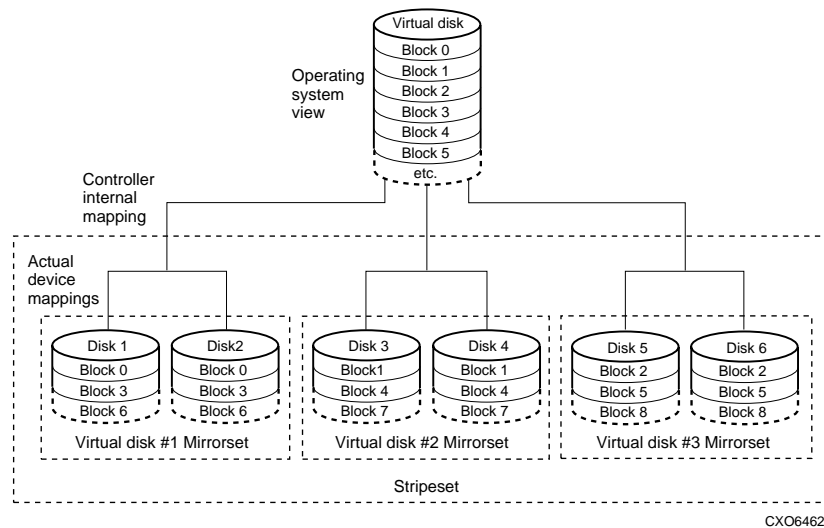


Figure 2-13. Striped mirrorset (example 2)

Plan the mirrorset members, then plan the stripeset that will contain them. Review the recommendations in “Storageset Planning Considerations,” page 2-12, and “Mirrorset Planning Considerations,” page 2-15.

## Partition Planning Considerations

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




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

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

## Defining a Partition

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

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

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

## Guidelines for Partitioning Storage Sets and Disk Drives

Keep these points in mind as you plan your partitions:

- You can create up to eight partitions per storage set or disk drive.
- In transparent failover mode, all the 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 storage sets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAID set.
- Just as with storage sets, 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 mirror sets or partitioned stripe sets.

## Changing Characteristics through Switches

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

- storage set switches
- initialization switches
- unit switches

The following sections:

- describe how to enable/modify switches
- contain a description of the major CLI command switches

## Enabling Switches

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

If you use CLI commands to configure the storageset or device manually, the configuration procedure found in Chapter 4 of this guide indicates when and how to enable each switch. The *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* contains the details of the CLI command and their switches.

## Changing Switches

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



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

---

## Storageset and Partition Switches

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

- RAIDset
- Mirrorset

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

### **RAIDset Switches**

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

- Replacement policy
- Reconstruction policy
- Remove/replace policy

For details on the use of these switches refer to SET RAIDSET and SET *RAIDset-name* commands in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

### **Mirrorset Switches**

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

- Replacement policy
- Copy speed
- Read source
- Membership

For details on the use of these switches refer to ADD MIRRORSET and SET *mirrorset-name* commands in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

### **Partition Switches**

When you create a partition, you can specify the following switches:

- Size
- Geometry

For details on the use of these switches refer to CREATE\_PARTITION command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

## Initialization Switches

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

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

Each of these is described in the following sections.

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


### Chunk Size

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

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

### Increasing the Request Rate

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



CXO-5135A-MC

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

Large chunk sizes also tend to increase the performance of random reads and writes. It is recommended 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 first must analyze the types of requests that are being made to the subsystem:


- Many parallel I/Os that use a small area of disk should use a chunk size of 10 times the average transfer request rate.
- Random I/Os that are scattered over all the areas of the disks should use a chunk size of 20 times the average transfer request rate.
- If you don't know, then you should use a chunk size of 15 times the average transfer request rate.
- If you have mostly sequential reads or writes (like those needed to work with large graphic files), then make the chunk size a small number (that is 17 sectors).

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

<b>Transfer Size (KB)</b>	<b>Small Area of I/O Transfers</b>	<b>Unknown</b>	<b>Random Areas of I/O Transfers</b>
2	41	59	79
4	79	113	163
8	157	239	317

### Increasing the Data Transfer Rate

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



CXO-5172A-MC

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



## Increasing Sequential Write Performance

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

## Save Configuration

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

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

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

## Destroy/Nodestroy

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

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

- *DESTROY* (default) overwrites the user data and forced-error metadata on a disk drive when it is initialized.
- *NODESTROY* preserves the user data and forced-error metadata when a disk drive is initialized. Use *NODESTROY* to create a single-disk unit from any disk drive that has been used as a member of a mirrorset. See the *REDUCED* command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 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.

## Unit Switches

There are several switches that control the characteristics of units. The unit switches are described under the SET *unit-number* command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

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

- “Assigning Unit Numbers,” page 1-17
- “Restricting Host Access (Selective Storage Presentation),” page 1-21

## Storage Maps

Configuring your 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 (a storage map).

### Creating a Storage Map

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

### Example Storage Map

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

- Unit D100 is a 3-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk10100, and Disk10200.
- Unit D101 is a 2-member striped mirror set named S1. S1 consists of M1 and M2:
  - M1 is a 2-member mirror set consisting of Disk10300 and Disk10400.
  - M2 is a 2-member mirror set consisting of Disk10500 and Disk10800.
- Unit D102 is a 2-member mirror set named M3. M3 consists of Disk10900 and Disk11000.
- Unit D103 is a 3-member stripeset named S2. S2 consists of Disk 11100, Disk11200, and Disk11300.

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

Model 4214 Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D100 R1 DISK10000	D100 R1 DISK10100	D100 R1 DISK10200	D101 S1 M1 DISK10300	D101 S1 M1 DISK10400	D101 S1 M2 DISK10500	D101 S1 M2 DISK10800	D102 M3 DISK10900	D102 M3 DISK11000	D103 S2 DISK11100	D103 S2 DISK11200	D103 S2 DISK11300	not supported	not supported

Figure 2-17. Example storage map-single bus

**2-30** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

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

- Unit D100 is a 3-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, and Disk10100.
- Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:
  - M1 is a 2-member mirrorset consisting of Disk10200 and Disk20200.
  - M2 is a 2-member mirrorset consisting of Disk20300 and Disk10300.
- Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10400 and Disk20400.
- Unit D103 is a 3-member stripeset named S2. S2 consists of Disk 10500, Disk20500, and Disk10800.
- Unit D104 is a single (JBOD) disk named Disk20100.
- Disk20800 is a spareset member.

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

Figure 2-18. Example storage map-dual bus

## Using the LOCATE Command to Find Devices

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

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

```
LOCATE D104
```

The LEDs on the disk drives that make up unit D104 will flash. After you have noted the position of all the drives contained within D104, enter the following command to turn off the flashing LEDs:

```
LOCATE CANCEL
```

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

```
LOCATE R1
```

After you have noted the position of the drives contained within R1, enter the following command to turn off the flashing LEDs:

```
LOCATE CANCEL
```



# *Chapter 3*

## **Preparing the Host System**

This chapter describes how to prepare your Sun Solaris host computer to accommodate the HSG60 controller storage subsystem.

The following information is included in this chapter:

- “Making a Physical Connection,” page 3–1
- “Verifying/Installing Required Versions,” page 3–3
- “Installing the Solution Software Packages,” page 3–5
- “Preparing LUNs for Use by the FileSystem,” page 3–9
- “File System Creation and Tuning,” page 3–12

### **Making a Physical Connection**

In order to attach your host computer to the storage subsystem, you must install one or more host bus adapters into the computer. A Fibre Channel cable goes from the host bus adapter to a Fibre Channel (FC) switch.

### **Preparing to Install the Host Adapter**

Perform the following steps:

1. Perform a complete backup of your entire system.
2. Shut down your computer system.

## Installing the Host Adapter

The first step in physical connection is installation of a host bus adapter.

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

You need the following to begin:

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

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

**NOTE:** Don't power on anything yet. In order for the FC switches to autoconfigure, you must 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	The 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	The 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	The 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.
CPQoptisp	The 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 SB5450 SCSI adapter under 64-bit Solaris, set the adapter to Sun emulation mode.



Installing the MA6000 Manager packages onto the hard disk requires the following steps:

1. Mount the CD-ROM.
2. Load the Fibre Channel adapter driver package(s).
3. Load the CPQoptisp 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. Reboot your computer.
7. Label (format) LUNs and create filesystems if required.
8. Configure the Agent to ensure subsystems are monitored (see Chapter 4).

## Verifying/Installing Required Versions

To prepare your system 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. Log in 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.

## Deciding between Automatic and Manual Installation Modes

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

- Solaris operating system version.
- Existence of required patches (and patch versions) on your system
- Prior versions of Solution Software packages (DECswm410, DECswm450, CPQhsz70, CPQhsg60).
- Prior versions of driver packages (DECptisp, CPQfca, CPQfcaw, CPQfcaPCI).
- Incompatible versions of driver packages (PTIptisp, JNIfca, JNIfcaw, JNIfcaPCI).
- It will also identify and backup any old Agent configuration files for later upgrading.

The key differences between Automatic and Manual modes is how they fix any problems found.

**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 the Solution Software Packages

This section includes the following topics:

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

### Mounting the CD-ROM

1. To check whether the volume management daemon (`vold`) is currently running, enter the following:  

```
# ps -ea | grep vold
```
2. Follow the steps below for “If `vold` is running;,” page 3-5 or “If `vold` is not running;,” page 3-6 as applicable:

#### If *vold* is running:

- a. Insert the CD-ROM into the computer that is connected to the controller.
- b. Use the `mount` command to verify that the volume manager has automatically mounted the CD-ROM, by entering:

```
# mount
```

**NOTE:** The system may take a few seconds to mount the CD-ROM. If the `mount` command does not indicate that the CD-ROM has been mounted, wait a short interval and then repeat the command. You may need to use the `volcheck` command to force `vold` to check for mounted media.

- c. Change to the Solaris directory, enter:

```
# cd /cdrom/hsg60_V85_sun/agents/solaris
```

- d. Do one of the following:

- If you want to use the automatic installation (recommended), See “Installing with Automatic Mode,” page 3-6.
- If you want to use the manual installation, See “Installing with Manual Mode,” page 3-8. The manual installation allows you to choose what is installed, such as supported adapters and drivers.

**If *bold* is not running:**

- a. Insert the CD-ROM into the drive of the computer that is connected to the controller.
- b. Mount the CD-ROM. For example, enter the following (your device name may be different):  

```
# mount -F hsfs -r /dev/dsk/c0t6d0s2 /cdrom
```
- c. Change to the Solaris directory by entering the following:  

```
# cd /cdrom/agents/solaris
```
- d. Do one of the following:
  - If you want to use the automatic installation (recommended), See “Installing with Automatic Mode,” page 3–6.
  - If you want to use the manual installation, See “Installing with Manual Mode,” page 3–8. The manual installation allows you to choose what is installed, such as supported adapters and drivers.

## **Installing with Automatic Mode**

To install the RAID Manager software using automatic mode:

1. To start the Installation Manager, type:  

```
# ./install_stgwks
```

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 this is true, the Installation Manager scans your system for supported adapters, and it displays a list of compatible adapter drivers.

Installation Manager prompts you to remove any old or incompatible packages.
3. Remove old and incompatible software. Follow the on-screen prompts, always accepting the default responses. Installation Manager prompts you for the package base directory.
4. Enter the package base directory. This directory has the “steam” subdirectory, containing the Raid Manager files, created under it. The default directory is /opt.  

Installation Manager rescans your system for supported HBA 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 a while (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.
  - If you 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.
  - If you choose “**A**” to automatically edit the `sd.conf` file, support is added for all supported Fibre Channel adapters, using the default SCSI 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 World Wide Port Names (WWPNs) to attach to.

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

7. After Raid Manager is loaded, select “**Q**” to exit the Installation Manager. You are prompted to reboot your computer.
8. Change to the top level directory, by entering the following:

```
# cd /
```
9. Enter the following command to shut down the computer:

```
# /etc/shutdown -y -i0 -g0
```
10. Wait for the “ok” *console* prompt, 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 RAIDManager 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 you are running a supported version of Solaris, and that you have the correct patches loaded.
  - Installation Manager scans your system for old or incompatible drivers and it will display a message describing which commands to use to correct the problem.
2. Follow the on-screen prompts to quit the Installation Manager.
  3. Remove the offending packages.
  4. Restart Installation Manager.

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.

5. 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.
6. 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.
7. If you have loaded all adapter driver packages you need, install the RAIDManager (CPQhsg60) software package. This is described in the following steps:
8. Enter the package base directory. This directory will have a “steam” subdirectory, containing the RAIDManager files, created under it. The default directory is /opt.  
Installation Manager re-scans your system for supported HBA drivers.
9. If any adapter drivers were loaded in the previous steps, you must answer “Y” to allow the Installation Manager to re-scan your system.
10. Edit your sd.conf file. This edit can be done automatically or manually.
  - If you 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.

- If you 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 World Wide Port Names (WWPNs) to attach to.

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

11. After Raid Manager is loaded, select “Q” to exit Installation Manager. You will be prompted to reboot your computer system.
12. Change to the top level directory, by entering the following:
 

```
# cd /
```
13. To shut down the computer, enter the following:
 

```
# /etc/shutdown -y -i0 -g0
```
14. Wait for the “ok” console prompt, then enter the following (boot\_device is the name of your system disk):
 

```
boot boot_device -r
```
15. 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 MA6000 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 format 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 2.4 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-HSG60-AX-V85Z 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
4. 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-HSG60-AX-V85Z 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 HSG60 disk entry in the format.dat file, do not try to re-use this definition. Always use type 0 to allow Solaris to retrieve the correct logical geometry for the LUN you are formatting.

## File System Creation and Tuning

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

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
```

These commands set *maxcontig* to 128 and *rotational delay* to 0. These recommendations may not optimize your system performance so you may want to experiment with other values.

# Chapter 4

## Installing and Configuring the HS-Series Agent

The following information is included in this chapter:

- “Why Use StorageWorks Command Console?,” page 4-1
- “About the Network Connection for the Agent,” page 4-3
- “Configuring the Agent,” page 4-5
- “Raid Manager Mail Messages,” page 4-9

---

**IMPORTANT:** To install the Solution software packages, see Chapter 3, Preparing the Host System

---

### Why Use StorageWorks Command Console?

StorageWorks Command Console (SWCC) lets you monitor and configure the storage connected to the HSG60 controller. SWCC is comprised of a Client and Agent. The Client provides pager notification and lets you manage your virtual disks. Client runs on Windows 2000 and Windows NT 4.0 (Intel) with Service Pack 4 or later.

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 e-mail notification and error logging

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

**4-2** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

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 is comprised of the HSG60 controller and its devices. Remote and local notification can be made by e-mail and/or SNMP messages to an SNMP monitoring program.

**Table 4-1 SWCC Features and Components**

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

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

The following table provides an overview of the installation.

**Table 4-2 Installation and Configuration Overview**

Step	Procedure
1	Verify that your hardware has been setup correctly. See the previous chapters.
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," page 4-3.
4	Verify that there is a LUN to communicate through. This can be either the CCL or a LUN that was created with the CLI.
5	Configure the SWCC Agent with the install.sh or config.sh utilities. See "Configuring the Agent," page 4-5.
6	Install the Client software on Windows 2000 or Windows NT 4.0 (Intel) with Service Pack 4 or later. See Appendix B.
7	Add the name of the agent system to the Navigation Tree of each client system that is on the Agent's list of client system entries (TCP/IP network connections). See Appendix B.
8	Set up pager notification (TCP/IP network connections). Refer to "Setting Up Pager Notification" in the <i>Compaq StorageWorks Command Console Version 2.3. for the MA6000 User Guide</i> .

## About the Network Connection for the Agent


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

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

---

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

---



CXO7240A

Figure 4-1. An example of a network connection

- |   |   |
|---|---|
| ① Agent system (has the Agent software)   | ⑤ Hub or switch                             |
| ② TCP/IP Network                          | ⑥ HSG60 controller and its devices (MA6000) |
| ③ Client system (has the Client software) | ⑦ Servers                                   |
| ④ Fibre Channel cable                     |   |

## Configuring the Agent

There are two ways to configure your Agent:

- **install.sh** – This utility is usually used for the first-time configuration or for 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)”
- **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 `config.sh`” 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:

- Adding a subsystem entry
  - Adding a client entry
  - Creating a password
  - Setting up e-mail notification
  - Starting the Agent
1. To change directories, enter the following:  
# `cd /opt/steam/bin`
  2. To run the installation, enter the following:  
# `./install.sh`

Follow the on-screen prompts to setup 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 config.sh

To view the RAID Array V8.5 Configuration Menu, run the *config.sh* script from the \$BASEDIR/steam/bin directory (usually /opt/steam/bin).

1. Enter the following:

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

2. Enter the following:

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

This displays the Configuration Menu shown in Figure 4-2.

---

---- RAID Array v8.5 Configuration Menu ----

---

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

---

Figure 4-2. Configuration menu



---

**IMPORTANT:** When you change the Agents 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 for this procedure, 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 subsystems you want the Agent to monitor.

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

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

2. Select option **13**, to add a subsystem.

The subsystem name is arbitrary, but use only lower-case 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's partition* and *print* commands.

3. Enter a monitoring interval; for example, 30 seconds, and press **Enter** twice to return to the *Main Menu*.

### **Adding a Client System Entry**

1. From the *Client Options* group, select option **16**, *View Clients* to see 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 remains 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 *config.sh*," 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 *Compaq StorageWorks Command Console Version 2.3. for the MA6000 User Guide*.



# Chapter 5

## Configuration Procedures

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

The following information is included in this chapter:

- “Establishing a Local Connection,” page 5–2
- “Setting Up a Single Controller,” page 5–4
- “Setting Up a Controller Pair,” page 5–10
- “Configuring Devices,” page 5–16
- “Configuring a Stripeset,” page 5–17
- “Configuring a Mirrorset,” page 5–17
- “Configuring a RAIDset,” page 5–18
- “Configuring a Striped Mirrorset,” page 5–19
- “Configuring a Single-Disk (JBOD) Unit,” page 5–20
- “Configuring a Partition,” page 5–20
- “Assigning Unit Numbers and Unit Qualifiers,” page 5–21
- “Configuration Options,” page 5–22

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

**5-2** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*


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

To install a controller or cache module and the PCMCIA card, see the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*.

## Establishing a Local Connection

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

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



- ① Maintenance Port Cable
- ② Maintenance Port

Figure 5-1. Maintenance port connection



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

---

## Establishing Connection with a SPARC System

To setup your SPARC system for connection with the HSG60 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 #1, 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 in the window similar to the following:

```
HSG60>
```

## 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 COMM1 or COMM2 ports.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Configure the computer or terminal as follows:
  - 9600 baud
  - 8 data bits
  - 1 stop bit
  - no parity
5. Press **Enter**. A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.



## Cabling a Single controller

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

**NOTE:** It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure (“Configuring a Single Controller using CLI,” page 5-6).



- ❶ controller
- ❷ host port 1
- ❸ host port 2
- ❹ cable from the switch to the host Fibre Channel adapter
- ❺ FC switch

Figure 5-2. Single controller cabling

## Configuring a Single Controller using CLI

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

```
SHOW THIS
```

See “Worldwide Names (Node IDs and Port IDs),” page 1–28, for the location of the sticker.

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

```
HSG> show this
Controller:
  HSG60 ZG80900583 Software V8.5L-0, Hardware E01
  NODE_ID           = 5000-1FE1-0001-3F00
  ALLOCATION_CLASS = 0
```

If the node ID is present, go to step 4.

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

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

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

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

```
SET NOFAILOVER
```

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

```
CLEAR_ERRORS CLI
```

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

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

4. 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, using the following command:

```
RESTART THIS
```

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

- Set the time on the controller using 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 won't give you a prompt. Ignore the procedure and just 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 *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 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. Following is a sample result of a SHOW THIS command, with the areas of interest in bold.

```
Controller:
  HSG60 (C) DEC ZG09030200 Software V85L, Hardware 0000
  NODE_ID = 5000-1FE1-0000-0000
  ALLOCATION_CLASS = 0
  SCSI_VERSION = SCSI-2
  Not configured for dual-redundancy
  Device Port SCSI address 7
  Time: 10-Mar-1999:12:30:34
  Command Console LUN is lun 0

Host PORT_1:
  Reported PORT_ID = 5000-1FE1-0000-0001
  PORT_1_TOPOLOGY = FABRIC (fabric up)
  Address = 210313

Host PORT_2:
  Reported PORT_ID = 5000-1FE1-0000-0002
  PORT_2_TOPOLOGY = FABRIC (fabric up)
  Address = 210513
  NOREMOTE_COPY
  .....
```

10. 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.

11. Plug in the Fibre Channel cable from the first adapter in the first host into the switch. Enter a *SHOW connections* command to view the connection table:

```
SHOW CONNECTIONS
```

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

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

For an explanation about why there may be multiple connections shown for one adapter, see page 1-13.

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

```
RENAME !NEWCON01 ANGEL1A1
```

For a recommended naming convention, see “Naming Connections,” page 1-12.

13. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=SUN
```

14. Verify the changes:

```
SHOW CONNECTIONS
```

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

15. Repeat step 12, 13, and 14 for each of that adapter’s host connections. (Or delete the unused connections from the table.
16. For each host adapter, repeat steps 11 through 15.

## Verifying Installation

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

## Setting Up a Controller Pair


### 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 COMM1 or COMM2 ports.
2. Turn on the computer or terminal.
3. Apply power to the storage subsystem.
4. Configure the computer or terminal as follows:
  - 9600 baud
  - 8 data bits
  - 1 stop bit
  - no parity
5. Press **Enter**. A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

## Cabling a Controller Pair

The cabling for a controller pair 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 5-12).



- ❶ controller A
- ❷ controller B
- ❸ host port 1
- ❹ host port 2
- ❺ cable from the switch to the host FC adapter
- ❻ FC switch

Figure 5-3. Controller pair failover cabling

CXO6887B

## Configuring a Controller Pair using CLI

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

```
SHOW THIS
```

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

```
HSG1402> show this
```

```
Controller:
```

```
  HSG60 ZG80900583 Software V8.5L-0, Hardware E01
```

```
  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
```

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

See “Worldwide Names (Node IDs and Port IDs),” page 1-28 for the location of the sticker.

2. If working with anything but factory-new controllers, enter the following command to remove any failover mode that may have been previously configured:

```
SET NOFAILOVER
```

3. Enter the following command to stop the CLI from reporting a misconfiguration error resulting from having no failover mode specified:

```
CLEAR CLI
```

4. Put the controller pair into transparent failover mode, using the following command:

```
SET FAILOVER COPY=THIS
```

The copy qualifier specifies where there is a good copy of the array configuration.

Initially, the copy qualifier is meaningless because there is no array configuration yet, but it is part of the command syntax and must be specified.

**NOTE:** If there is configuration information that you want to keep, which will probably be the case if the controllers are not factory-new, set COPY to the controller that contains the good configuration information

When the command is entered, the “other” controller (the one to which the serial line is not attached) will restart. The restart may set off the audible alarm of the EMU. To silence the alarm, press once and release the button on the EMU. The alarm will stop, but the LED will remain on until the controllers bind into transparent failover mode.



The binding process takes about 15 seconds. The CLI will print out a last event report from the “other” controller. This just indicates that the “other” controller restarted. The CLI will continue reporting this condition until cleared with the following command:

```
CLEAR CLI
```

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

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

6. Set the topology for both ports of both controllers:

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

```
SET OTHER PORT_1_TOPOLOGY=FABRIC
```

```
SET OTHER PORT_2_TOPOLOGY=FABRIC
```

**NOTE:** If the controllers are not factory-new, they may have another topology set, in which case these commands will result in an error message. If this happens, take all ports offline first, then reset the topology:

```
SET THIS PORT_1_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_2_TOPOLOGY=OFFLINE
```

```
SET OTHER PORT_1_TOPOLOGY=OFFLINE
```

```
SET OTHER PORT_2_TOPOLOGY=OFFLINE
```

```
SET THIS PORT_1_TOPOLOGY=FABRIC
```

```
SET THIS PORT_2_TOPOLOGY=FABRIC
```

```
SET OTHER PORT_1_TOPOLOGY=FABRIC
```

```
SET OTHER PORT_2_TOPOLOGY=FABRIC
```

7. Restart the controllers:

```
RESTART OTHER
```

```
RESTART THIS
```

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

When the “other” controller restarts the EMU alarm will sound, but the restart of “this” controller turns the alarm off.

8. Set the time on “this” controller using the following syntax:

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

**5-14** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

9. Use the FRUTIL utility to set up the battery discharge timer in “this” controller. Enter the following command to start FRUTIL:

```
RUN FRUTIL
```

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

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

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

10. Move the serial cable to the lower controller, and repeat step 8.
11. Set up any additional optional controller settings, such as changing the CLI prompt. See the SET *this\_controller/other\_controller* command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* for the format of optional settings.
12. Verify that all commands have taken effect. Use the following commands:

```
SHOW THIS
SHOW OTHER
```

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

```

Controller:
  HSG60   (C) DEC ZG09030200 Software V85L, Hardware 0000
  NODE_ID           = 5000-1FE1-0000-0000
  ALLOCATION_CLASS = 0
  SCSI_VERSION     = SCSI-2
  Configured for TRANSPARENT_FAILOVER with ZG80200284
  In dual-redundant configuration
  Device Port SCSI address 6
  Time: 10-Mar-1999:12:30:34
  Command Console LUN is lun 0

Host PORT_1:
  Reported PORT_ID = 5000-1FE1-0000-0001
  PORT_1_TOPOLOGY = FABRIC (fabric up)
  Address          = 210313

Host PORT_2:
  Reported PORT_ID = 5000-1FE1-0000-0002
  PORT_2_TOPOLOGY = FABRIC (standby)
  Address          = 210513
  NOREMOTE_COPY
  .....
  
```

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

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

14. Plug in the Fibre Channel cable from the first adapter on the first host into the switch. Enter a **SHOW connections** command to view the connection table:

```
SHOW CONNECTIONS
```

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

For a description of why plugging in one adapter can result in multiple connections, see “Numbers of Connections,” page 1-13.

**5-16** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

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

```
RENAME !NEWCON01 ANGEL1A1
```

Compaq recommends using a naming convention, see “Naming Connections,” page 1-12.

16. Specify the operating system for the connection:

```
SET ANGEL1A1 OPERATING_SYSTEM=SUN
```

17. Verify the changes:

```
SHOW CONNECTIONS
```

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

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.

## Verifying Installation

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

## Configuring Devices

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

```
RUN CONFIG
```

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

## Configuring a Stripese

1. Create the stripese by adding its name to the controller's list of storagesets and specifying the disk drives it contains. Use the following syntax:  
ADD STRIPESET STRIPESET-NAME DISKNNNNN DISKNNNNN .....
2. Initialize the stripese, specifying any desired switches:  
INITIALIZE STRIPESET-NAME SWITCHES  
See "Initialization Switches" on page 2-24 for a description of the initialization switches.
3. Verify the stripese configuration:  
SHOW STRIPESET-NAME
4. Assign the stripese a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5-21.

### Example:

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

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

## Configuring a Mirrorset

1. Create the mirrorset by adding its name to the controller's list of storagesets and specifying the disk drives it contains. Optionally, you can append mirrorset switch values:  
ADD MIRRORSET MIRRORSET-NAME DISKNNNNN DISKNNNNN SWITCHES  
  
**NOTE:** See the ADD MIRRORSET command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* for a description of the mirrorset switches.
2. Initialize the mirrorset, specifying any desired switches:  
INITIALIZE MIRRORSET-NAME SWITCHES  
See "Initialization Switches" on page 2-24 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 5-21.

**Example:**

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

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

## Configuring a RAIDset

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

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

**NOTE:** See the ADD RAIDSET command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* for a description of the RAIDset switches.

2. Initialize the RAIDset, specifying any desired switches:

```
INITIALIZE RAIDSET-NAME SWITCH
```

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

See “Initialization Switches” on page 2-24 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 5-21.

**Example:**

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

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

## Configuring a Striped Mirrorset

1. Create, but do not initialize, at least two mirrorsets.  
See “Configuring a Mirrorset” on page 5-17.
2. Create a stripeset and specify the mirrorsets it contains:  
ADD STRIPESET STRIPESET-NAME MIRRORSET-1 MIRRORSET-2...MIRRORSET-N
3. Initialize the striped mirrorset, specifying any desired switches:  
INITIALIZE STRIPESET-NAME SWITCH  
See “Initialization Switches” on page 2-24 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 5-21.

### **Example:**

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

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

## Configuring a Single-Disk (JBOD) Unit

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

```
INITIALIZE DISK-NAME SWITCHES
```

See “Initialization Switches” on page 2–24 for a description of the initialization switches.

2. Verify the configuration using the following command:

```
SHOW DISK-NAME
```

3. Assign the disk a unit number to make it accessible by the hosts. See “Assigning Unit Numbers and Unit Qualifiers” on page 5–21.

## Configuring a Partition

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

```
INITIALIZE STORAGESET-NAME SWITCHES
```

or

```
INITIALIZE DISK-NAME SWITCHES
```

See “Initialization Switches” on page 2–24 for a description of the initialization switches.

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

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

or

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

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

**NOTE:** See the CREATE\_PARTITION command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 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 5-21.

**Example:**

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

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

## Assigning Unit Numbers and Unit Qualifiers

Each storageset, partition, or single (JBOD) disk must be assigned a unit number for the host to access. As the units are added, their properties can be specified through the use of command qualifiers, which are discussed in detail under the ADD UNIT command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

Each unit can be reserved for the exclusive use of a host or group of hosts. See “Restricting Host Access in Transparent Failover Mode,” page 1-21

### 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
```

**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
```

**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
```

**Example:**

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

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

# Configuration Options

## Changing the CLI Prompt

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

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

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

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

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

## Mirroring cache

To specify mirrored cache:

```
SET THIS_MIRRORED_CACHE
```

## Adding Disk Drives

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

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

```
ADD DISK DISKNNN PTL-LOCATION
```

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

```
RUN CONFIG
```

## Adding a Disk Drive to the Spareset

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

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

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

```
ADD SPARESET DISKNNNNN
```

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

### **Example:**

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

```
ADD SPARESET DISK11300
```

```
ADD SPARESET DISK21300
```

## Removing a Disk Drive from the Spareset

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

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

```
SHOW SPARESET
```

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

```
DELETE SPARESET DISKNNNNN
```

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

```
SHOW SPARESET
```

## Enabling Autospare

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

To enable autospare 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 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 re-used disk drive might. To erase metadata from a disk drive, add it to the controller's list of devices, then set it to be transportable 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's profile when you change its switches.

### Displaying the Current Switches

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

```
SHOW STORAGESET-NAME
```

or

```
SHOW DEVICE-NAME FULL
```

### Changing RAIDset and Mirrorset Switches

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

```
SET RAID1 POLICY=BEST_FIT
```

### **Changing Device Switches**

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

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

### **Changing Initialize Switches**

The initialization switches cannot be changed without destroying the data on the storageset or device. 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
```

# Chapter 6

## Managing Storage Access for Both Loop and Fabric Environments

This chapter is divided into the following three sections:

- **General Considerations**, is a general discussion of the elements of a Fibre Channel environment.
- **StorageWorks' Sun Solaris Implementation of Loop and Fabric**, is a discussion of the StorageWorks' implementation of the Sun Solaris' approach for both Fibre Channel loop (FC\_AL) and Fibre Channel switch (FC\_SW) support at the system level.
- **Configuration Recipes**, is a recipe section for Arbitrated Loop and switch configurations.

In such 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 will help in the description of the differences and assist in the understanding of the configuration of Solaris servers and adapters for loop and switch support.

### General Considerations

Fibre Channel Configurations are assemblages of servers, storage (RAID Arrays), hubs, switches, and repeaters. Together they present a multitude of options for configuring a storage area network and require attention to their configuration in design, installation, testing, and documentation.

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 a greater discipline and control.

This chapter does not attempt to answer the larger question of how to connect the assemblage in an installation but *rather* how to configure the Solaris server and the storage ( RAID array ) components so that given a path with a loop or switch configuration, the Solaris server and storage will be able to communicate successfully.

## **The Configurable Objects - Servers**

Servers are the computing power and the communication access to the Storage that the computing applications require. In this chapter we are focusing the discussion on Sun servers running the Solaris operating system.

Servers can connect to local storage or RAID Arrays and 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: 1) number of adapters on a server and 2) the number of RAID systems that may be connected to an adapter. The current rule of thumb is that a server may have up to 4 adapters and may connect up to 4 RAID arrays. These numbers represent a relatively safe area of loading and are acceptable for most general usage server-storage situations. For some servers (e.g., in the 4m class) these values are may be too high and for some servers (e.g., 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. We have tested up to 24 adapters 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 value(s) are the actual I/O loading from server to storage. The major components of I/O values are: a) reading vs. writing data, b) the size of the records being read/written, and c) the frequency of data to/from a storage device.

Additionally, another set of parameters has a large impact on the viability of a configuration from the RAID array point of view. The major components such as: a) the nature of the controller (SCSI vs. Fibre), b) the amount of cache on the controller, and c) the configuration of the disks enclosed by the RAID Array.



Where does all this lead? The answer is quite simple but somewhat unacceptable from a "perfect world" point of view. The answer is that a specific configuration may or may not replicate the tested laboratory configuration(s). When treating a specific configuration the best guideline is to empirically determine the viability of a configuration that is outside of the 'safe' guidelines.

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.

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

## The Configurable Objects - Storage

Storage as used in this chapter is a RAID Array. Specifically, RAID arrays that are supported by the Compaq HSG80 family of controllers. This controller supports all of the standard RAID set configurations in a loop or fabric environment. This chapter *does not* attempt to answer the questions of best strategies for RAID set configurations with Fibre Channel.

However, the controller must be configured to either communicate in loop mode or fabric mode. These configuration elements *are* discussed herein.

## A Special Note on Controller Pairs (Active-Passive vs. Active-Active)

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 should one controller fail, the second controller will automatically service the I/O stream, protecting the integrity of the data.

Typically, this redundancy is created for failover protection. Failover can be implemented as either Active-Passive or Active-Active.

**Table 6-1 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

**6-4** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

In Case 1, all I/O would be loaded on a single controller, controller A, and 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 not a benefit. It is a drawback to the processing power of the controller pair.

**Table 6-2 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

In Case 2, we are able to protect both I/O channels as well as distribute 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 describing a single controller and not the controller pair.

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 below.

**Table 6-3 Example of Active-Passive with a Single Controller**

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

However, when applied to the **controller pair**, both controllers are active, the correct designation is ACTIVE-ACTIVE.

## The Configurable Objects - Hubs

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

**NOTE:** We do not support cascading hubs, so the largest supported loop configuration using hubs consists of 12 ports maximum.

## The Configurable Objects - Switches

Switches serve to allow multiple connections to each other. With switches, Fibre Channel is removed from some of the limitations of arbitrated loop: throughput, single point of failure, sizing of configuration, etc.

Earlier we noted that the complexity of a configuration can increase dramatically with components added in 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. So, this is a *first caution*: **take the time to plan your configuration, test the portions of it and record the various connections and paths.**

From the simplest case of a single server connected to a storage subsystem, we can extend the fabric by adding more subsystems, more switches, or more servers.

Adding switches is probably the easiest. To add a switch, all that must be done is to 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 then 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 (Silkorm 1000 Series) and new (Silkorm 2000 Series) switches. It is important to keep in mind that 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.

Connecting a server or a switch to another switch is a simple process. The Fibre Channel cables are inserted into the gigabit interface converters (GBICs) that reside in each port of the switch.

So, connections between switches are as easy as connections between server and switch with the added ability of adding more switching capability and/or more storage.

## **The Fibre Channel Access/Transport paths: Fibre Fabric vs. Fibre Loop.**

In a Fibre Loop (simply, loop), a server's Target is specified by the controller's port(s). These ports have physical addresses (PA) and in a loop that is arbitrated, the ports are designated ALPAs or 'arbitrated loop, physical addresses'. Since all servers and storage systems connected to a loop share it, it is important that every port on the loop have a unique ALPA.

In the Fibre Fabric (simply, fabric) the basic Target becomes the controller's World Wide 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.

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

In summary, we will manage two entities with Fibre Channel: ALPAs for loop and WWPNs for fabric.

## **Sun Solaris Implementation of Loop and Fabric**

### **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 and the *fca-pci.conf* for the JNI PCI adapter.

## Fabric Bindings

Fabric bindings rely on the World Wide Names and the StorageWorks' implementation uses the more specific relationship, the World Wide 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, we have 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 we extend the idea across servers, then we have many possibilities of configuring and sharing a storage in a fabric. Thus, in a shorthand version we can connect as follows:

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

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

In every case we can access storage subsystems 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 examples below have been reformatted for clarity and discussion.

## 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 we are enabling. 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";
```

**6-10** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

**Analysis**

- Device Class:     name="sd"  
The device class that we are enabling. 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 and they are documented in the Solaris Getting Started documentation. 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.



```

      ---- RAID array v8.5 Configuration Menu ----
Agent Admin Options                Storage Subsystem Options
-----
1) Change Agent Password           12) View Subsystems
2) Change SNMP Enterprise OID      13) Add a Subsystem
3) Start/Stop the Agent            14) Remove a Subsystem
4) Toggle Agent Startup on System  15) Modify a Subsystem
   Boot
5) Uninstall Agent
Agent Notification Options          Client Options
-----
6) Toggle Error Log Notification    16) View Clients
7) Toggle Mail Notification         17) Add a Client
8) View Mail Notification List      18) Remove a Client
9) Add User to Mail Notification    19) Modify a Client
   List
10) Remove User from Mail Notifica  Other Options
   tion List
11) Modify Mail Notification List   -----
Exit:                               20) Add/Change Adapters
-----
q) Quit                             21) High Avail. Options
Enter Selection:
```

All the options in this menu apply to any subsystem and any server, regardless of the Fibre Channel access that we have chosen.

When the V8.5 kits are being installed, the user has the opportunity to declare the specific type of Fibre Channel configuration for each adapter, either loop or fabric. Additionally, there will be many times that a change must be made to the adapters and we direct our attention to the most powerful option, Option 20, Add/Change Adapters.

When selected , we have 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
q) Exit
Enter choice:
```

**6-12** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

**NOTE:** In the header of this sub-menu we are reminded 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 file, fca.conf or fcaw.conf or fca-pci.conf.

### Option 1 - View Adapters

The sample display below shows one SBUS 64-bit Adapter and two PCI Adapters. The SBUS adapter is configured for Arbitrated Loop and the PCI adapters are configured for Fabric mode. Note also that both PCI adapters are communicating with the same World Wide Port Name entities on the Fabric. We do not know anything about the RAID configuration that is employed.

```
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 F 16 64   5000-1fe1-0000-0d41 /pci@a,2000/fibre-channel@2
                  65 5000-1fe1-0000-0d42
3 fibre-channel 1 fca-pci F 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

When selecting Option 2, the user will 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

When selecting Option 3, the list of adapters will be displayed and the user can 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

When selecting Option 5, Modify an Adapter we will see a listing of the adapters and may elect 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).

The most important detail to remember is that changing values or the driver mode for one adapter instance will not affect the mode for the other adapter instances. In this way we can have unique instances of each adapter for a specific type of a fibre transport medium, loop or fabric.

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	F 16	64 65	5000-1fe1-0000-0d41 5000-1fe1-0000-0d42	/pci@a,2000/fibre-channel@2
3 fibre-channel	1 fca-pci	F 16	64 65	5000-1fe1-0000-0d41 5000-1fe1-0000-0d42	/pci@b,2000/fibre-channel@2

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.

**NOTE:** After all prompts have been completed, the user has the option of accepting or rejecting the changes as shown in the change summary.  
(Shown below)

- Driver fca-pci is configured for Fabric mode.

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. Obviously, given the starting mode of 'fabric', answering 'y' will set the driver for loop mode. Recall however, that to implement loop mode, the ALPAs are used for access.

Similarly, if changing from loop to fabric mode, we would need to supply the WWPNs.

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

**6-14** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

By default, the Solaris Platform kit software supplies LUNs 0-15 for a total of 16 LUNs per Target. This number has been found to be more than adequate for many purposes. However, if a change of LUNs per Target is required and the user selects 'y', then the following prompt is displayed:

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

There are now two options available. Either press RETURN to accept the default value of 16 LUNs or enter the desired value.

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

The default Targets are 64 and 65. For a server that is communicating to more than one subsystem, we would need to have more than the two Targets. If the user selects 'y', then the following prompt appears:

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

The user has three options at this point:

1. Press RETURN and accept the default Target Ids of 64, 65.
2. 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.
3. Enter the letter, 'n'. This value specifies 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 the user has selected 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.
- Please 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:

1. Press RETURN to restore the old value for the WWPN.
2. Enter the WWPN value for each Target.
3. 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.

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]

From the display and the discussion, it is clear that we changed mode from fabric to loop, changed the number of LUNs per Target to 8 and set the Target IDs to 67 and 68.

At this point, we have two options.

1. Press RETURN and choose the default option of 'Y' and we will have modified the adapter's driver accordingly, making all the necessary changes to the sd.conf entries and the specific driver file, fca-pci.conf.
2. Enter 'n' and discard the proposed changes. This means that 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. Thus,

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

## Configuration Recipes

This section is included as a reminder of the steps that must be taken to configure a loop environment or a fabric environment or change from one to the other ( loop Ø fabric ) from the Solaris and storage points of view. This is not intended as a complete configuration guide, but rather a set of macro-level change recipes. Some assumptions have been made in these recipes. For example, it is assumed that the server in question has an adapter to support the Fibre Channel loop or fabric; that a storage system is available that will support a fibre connection, that hubs and switches are available and so on.

**NOTE:** The next two recipes assume that the Solaris platform kit is being loaded onto the system and a specific mode of operation is being selected.

### Loop Configuration(s)

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:
  - verify that the controller is configured for loop mode and verify that the specific, desired ALPAs are assigned to each destination controller port using the CLI.
  - verify that the server has the specific ALPAs assigned as Targets for the adapter by using `/opt/steam/bin/config.sh` - Option 20 Add/Change Adapters, Option 1, View adapters.
5. In cases that require a change of ALPAs at the server a reboot of the system will be required. The command `reboot -r` should be used to force a reconfiguration boot.
6. After reboot, the intended Targets should be visible under the *format* listing. This assumes that at least one unit has been created on the subsystem(s).

## Fabric Configuration(s)

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 subsystem(s) controller(s) 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.)

**NOTE:** The next two recipes assume that the Solaris platform kit is loaded and everything is running as desired but a change of configuration is required.

## Loop to Fabric Reconfiguration

1. Stop I/O to the subsystem. Perform backups, as required.
2. Remove the connections from the hub.
3. Reconfigure the controller(s) 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 server(s), reconfigure the adapters to support the fabric WWPNs and remove the unused loop Targets. This is accomplished with `/opt/steam/bin/config.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 server(s) 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. We have found that 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 subsystem(s) are now available.

## **Fabric to Loop Reconfiguration**

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 controller(s) 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 hub(s).
5. At the server(s), reconfigure the adapters to support loop. This is accomplished with `/opt/steam/bin/config.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 server(s) 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. We have found that 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 subsystem(s) are now available.



# Chapter 7

## Configuration Example Using CLI

This chapter presents an example of configuring a particular storage subsystem using the CLI.

The CLI configuration example shown assumes:

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

An example storage subsystem is shown in Figure 7-1. The example system contains three non-clustered Sun Solaris hosts, as shown in Figure 7-2. The resulting virtual system from the hosts' point of view is shown in Figure 7-3.

The name of the example storage subsystem is "HLNDR." The top controller will be named "HLNDR TOP" and the bottom controller will be named "HLNDR BOTTOM." In the example, the CLI prompts will be changed to the controllers' names. The only other option chosen for the controller pair is mirrored cache.

The example uses these conventions:

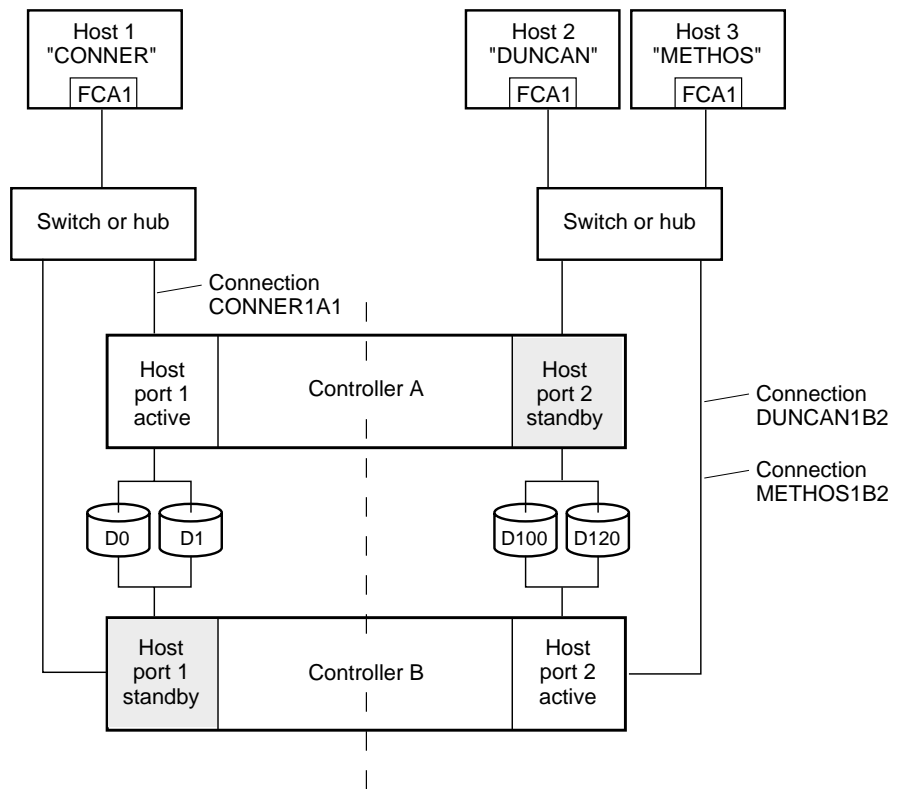
- Screen messages from CLI or utilities are boxed.
- Instructions to users are in italics.

7-2 MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
<b>SCSI ID</b>	00	01	02	03	04	05	08	09	10	11	12	13	14	15
<b>DISK ID</b>	D0 R1  DISK10000	D0 R1  DISK10100	D0 R1  DISK10200	D1 S1 M1 DISK10300	D1 S1 M1 DISK10400	D1 S1 M2 DISK10500	D1 S1 M2 DISK10800	D100 M3 DISK10900	D100 M3 DISK11000	D120 S2 DISK11100	D120 S2 DIS112000	Spare DISK11300	not supported	not supported

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


Figure 7-1. Example storage map



NOTE: FCA = Fibre Channel Adapter

CXO7296A

Figure 7-2. Example system



CXO7297A

Figure 7-3. Example virtual layout from the hosts' point of view

## CLI Configuration Example

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

*Plug serial cable from maintenance terminal into top controller.*

```
CLEAR CLI
SET FAILOVER COPY=THIS
CLEAR CLI
SET THIS SCSI_VERSION=SCSI-2
SET THIS PORT_1_TOPOLOGY=FABRIC
SET THIS PORT_2_TOPOLOGY=FABRIC
SET THIS ALLOCATION_CLASS=0
RESTART OTHER
RESTART THIS
SET THIS TIME=11-JAN-2000:10:20:00
RUN FRUTIL
```

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

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 [N]

Y

```
SET THIS MIRRORED_CACHE.
```

**NOTE:** This command causes the controllers to restart.

```
SET THIS PROMPT="HLNDR BOTTOM"
SET OTHER PROMPT="HLNDR TOP"
SHOW THIS
SHOW OTHER
RUN CONFIG
```

```

ADD RAIDSET R1 DISK10000 DISK20000 DISK10100 DISK20100 DISK10200 DISK20200
INITIALIZE R1
ADD UNIT D0 R1 DISABLE_ACCESS_PATH=ALL
SET D0 ENABLE_ACCESS_PATH=(CONNER1A1)
ADD MIRRORSET MI DISK10300 DISK20300 DISK10400 DISK20400
ADD MIRRORSET M2 DISK10500 DISK20500 DISK10800 DISK20800
ADD STRIPESET S1 M1 M2
INITIALIZE S1
ADD UNIT D1 S1 DISABLE_ACCESS_PATH=ALL
SET D1 ENABLE_ACCESS_PATH=(CONNER1A1)
ADD MIRRORSET M3 DISK10900 DISK20900 DISK11000 DISK21000
INITIALIZE M3
ADD UNIT D100 M3 DISABLE_ACCESS_PATH=ALL
SET D100 ENABLE_ACCESS_PATH=(DUNCAN1B2)
ADD STRIPESET S2 DISK11100 DISK11200 DISK21100 DISK21200
INITIALIZE S2
ADD UNIT D120 S2 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(METHOS1B2)
ADD SPARESET DISK11300
ADD SPARESET DISK21300
SHOW UNITS ALL

```

*Plug in the Fibre Channel cable from the adapter in host "CONNER."*

```
SHOW CONNECTIONS
```

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

```

RENAME !NEWCON00 CONNER1A1
SET CONNER1A1 OPERATING_SYSTEM=SUN

```

**7-6** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

SHOW CONNECTIONS

**NOTE:** Connection table sorts alphabetically.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
XANDER1A1		OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

Mark or tag both ends of Fibre Channel cables.

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

SHOW CONNECTIONS

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

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON01	WNNT	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
CONNER1A1	SUN	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

RENAME !NEWCON01 DUNCAN1B2

SET DUNCAN1B2 OPERATING\_SYSTEM=SUN

SHOW CONNECTIONS

Mark or tag both end of Fibre Channel cables.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
CONNER1A1	SUN	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
DUNCAN1B2	SUN	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

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

SHOW CONNECTIONS RENAME !NEWCON02 METHOS1B2

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON02	WNNT	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				
CONNER1A1	SUN	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				
DUNCAN1B2	SUN	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				

SETMETHOS1B2 OPERATING\_SYSTEM=SUN

SET METHOS1B2 UNIT\_OFFSET=120

SHOW CONNECTIONS *Mark or tag both end of Fibre Channel cables.*

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
CONNER1A1	SUN	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				
DUNCAN1B2	SUN	THIS	2	XXXXXX	OL this	100
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				
METHOS1B2	SUN	THIS	2	XXXXXX	OL this	120
HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX				









## Chapter 8

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

This chapter describes some common procedures that are not part of configuration.

The following information is included in this chapter:

- “Backing Up the Subsystem Configuration,” page 8-1
- “Cloning Data for Backup,” page 8-2
- “Moving Stagesets,” page 8-6

## Backing Up the Subsystem Configuration

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

You can use the `SHOW this_controller FULL` command to find out if the save configuration feature is active and which devices are being used to store the configuration. The display includes a line that indicates status and how many devices have copies of the configuration, as shown in the following example.

```
SHOW THIS_CONTROLLER FULL
```

The last line of the status the CLI returns shows on how many devices the configuration is backed up.

Also, the `SHOW devices FULL` command shows which disk drive are set up to back up the configuration. Enter this command:


```
SHOW DEVICES FULL
```

## Cloning Data for Backup

Use the `CLONE` utility to duplicate the data on any unpartitioned single-disk unit, stripeset, mirrorset, or striped mirrorset in preparation for backup. When the cloning operation is done, 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 the clones contain the most up-to-date data.

After the `CLONE` utility copies the data from the members to the clones, it restores the unit to its original configuration and creates a clone unit you can backup. The `CLONE` utility uses steps shown in Figure 8-1 to duplicate each member of a unit.



CXO5510A

Figure 8-1. Steps the CLONE utility follows for duplicating unit members

**8-4** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

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.

**Example:**

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

```
RUN CLONE
CLONE LOCAL PROGRAM INVOKED
UNITS AVAILABLE FOR CLONING:
    98
ENTER UNIT TO CLONE? 98
CLONE WILL CREATE A NEW UNIT WHICH IS A COPY OF UNIT 98.
ENTER THE UNIT NUMBER WHICH YOU WANT ASSIGNED TO THE NEW UNIT? 99
THE NEW UNIT MAY BE ADDED USING ONE OF THE FOLLOWING METHODS:
1. CLONE WILL PAUSE AFTER ALL MEMBERS HAVE BEEN COPIED. THE USER MUST THEN PRESS
    RETURN TO CAUSE THE NEW UNIT TO BE ADDED.
2. AFTER ALL MEMBERS HAVE BEEN COPIED, THE UNIT WILL BE ADDED AUTOMATICALLY.
UNDER WHICH ABOVE METHOD SHOULD THE NEW UNIT BE ADDED[]?1
DEVICES AVAILABLE FOR CLONE TARGETS:
DISK20200 (SIZE=832317)
DISK20300 (SIZE=832317)
USE AVAILABLE DEVICE DISK20200(SIZE=832317) FOR MEMBER DISK10300(SIZE=832317) (Y,N) [Y]? Y
MIRROR DISK10300 C_MA
SET C_MA NOPOLICY
SET C_MA MEMBERS=2
SET C_MA REPLACE=DISK20200
DEVICES AVAILABLE FOR CLONE TARGETS:
DISK20300 (SIZE=832317)
USE AVAILABLE DEVICE DISK10400(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=DISK10400
```

**8-6** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

COPY IN PROGRESS FOR EACH NEW MEMBER. PLEASE BE PATIENT...

.  
.

COPY FROM DISK10300 TO DISK20200 IS 100% COMPLETE

COPY FROM DISK10000 TO DISK10400 IS 100% COMPLETE

PRESS RETURN WHEN YOU WANT THE NEW UNIT TO BE CREATED

REDUCE DISK20200 DISK10400

UNMIRROR DISK10300

UNMIRROR DISK10000

ADD MIRRORSET C\_MA DISK20200

ADD MIRRORSET C\_MB DISK10400

ADD STRIPESET C\_ST1 C\_MA C\_MB

INIT C\_ST1 NODESTROY

ADD UNIT D99 C\_ST1

D99 HAS BEEN CREATED. IT IS A CLONE OF D98.

CLONE - NORMAL TERMINATION

## Moving 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's software for information on which drives can be supported.



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

---



Use the following procedure to move a stageset while maintaining the data it contains:

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

```
SHOW STAGESET-NAME
```

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

If you do not have a stageset 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 stageset-name command. Use the following syntax:

```
DELETE UNIT-NUMBER
```

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

```
DELETE STAGESET-NAME
```

5. Delete each disk drive, one at a time, that the stageset 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. Add again 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 stageset by adding its name to the controller's list of valid stagesets and specifying the disk drives it contains. (Although you have to recreate the stageset from its original disks, you do not have to add them in their original order.) Use the following syntax:

```
ADD STAGESET-NAME DISK-NAME DISK-NAME
```

**8-8** *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

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
```

**Example:**

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

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

# *Appendix* **A**

## **Subsystem Profile Templates**

This appendix contains storageset profiles you can copy and use to create your system profiles. It also contains an enclosure template you can use to help keep track of the location of devices and storagesets in your shelves. For HSG60, the maximum number of shelves is two. Four (4) templates will be needed for the MA6000.

- “Storageset Profile,” page A-2
- “Storage Map Template 1,” page A-3 for a single Model 4214R enclosure
- “Storage Map Template 2,” page A-4 for a single Model 4254 enclosure
- “Storage Map Template 3,” page A-5 for two Model 4214R enclosures

## StorageSet Profile

### Type of StorageSet:

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

### 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

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

Model 4214 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	not supported	not supported

## Storage Map Template 2

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

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

## Storage Map Template 3

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

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

Model 4214 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	not supported	not supported

Model 4214 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	not supported	not supported





## *Appendix* **B**

### **Installing, Configuring, and Removing the Client**

The following information is included in this appendix:

- “Why Install the Client?,” page B-1
- “Before You Install the Client,” page B-2
- “Installing the Client,” page B-3
- “Troubleshooting the Client Installation,” page B-3
- “Adding the Storage Subsystem and its Host to the Navigation Tree,” page B-5
- “Removing the Command Console Client,” page B-7
- “Where to Find Additional Information,” page B-8

#### **Why Install the Client?**

Client allows you to monitor and manage your storage subsystem. When you install the Client, you are able to do the following:

- 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 (RAID 3/5)
- Create an individual device (JBOD)
- Monitor many subsystems at once
- Set up pager notification

## **Before You Install the Client**

1. Verify that you are logged into an account that is a member of the administrator group.
2. Check the software product description that came with your software for a list of supported hardware.
3. Verify that you have the SNMP service installed on your computer. You must have SNMP installed on your computer for this software to work properly. The Client software uses SNMP to receive traps from the Agent. The SNMP service is available on your Windows NT or Windows 2000 installation CD-ROM. To verify that you have the SNMP service:
  - For Windows NT, double-click Services in *Start > Settings > Control Panel*. The entry for SNMP is shown in this window. If you install the SNMP service and you already have a Windows NT service pack on your computer, reinstall the service pack after installing the SNMP service.
  - For Windows 2000, click *Start > Settings > Control Panel > Administrative Tools > Component Services*. The entry for SNMP is shown in the Component Services window.
4. Read the release notes.
5. Read “Troubleshooting the Client Installation,” in this appendix.
6. If you have the Command Console Client open, exit the Command Console Client.
7. If you have Command Console Client version 1.1b or earlier, remove the program by going into Add/Remove Programs.
8. If you have a previous version of Command Console, you can save your Navigation Tree configuration by copying the swcc2.mdb file to another directory. After you have installed the product, move swcc2.mdb to the directory to which you installed SWCC.
9. Install the HS-Series Agent. For more information, see Chapter 4.

## Installing the Client

1. Insert the CD-ROM into a computer running Windows 2000 or Windows NT 4.0 (Intel) with Service Pack 4 or later.
2. Using Microsoft Windows Explorer, go to the top-level directory on the CD-ROM, and double-click setup.exe. The SWCC Setup window appears.
3. Select HSG60 Controller and click Next.

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

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

**NOTE:** For more information on AES, see *Compaq StorageWorks Command Console Version 2.3 for the MA6000 User Guide*.

## Troubleshooting the Client Installation

This section provides the following information on how to resolve some of the problems that may appear when you install 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.

**B-4** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide

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 the Storage Subsystem and its Host to the Navigation Tree

The Navigation Tree allows you to manage your storage over the network by using the Storage Window. If you are planning 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.

**B-6** MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide




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

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



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

**NOTE:** You can create virtual disks by using the Storage Window. For more information on the Storage Window, refer to *Compaq StorageWorks Command Console Version 2.3 for the MA6000 User Guide*.

## Removing the Command Console Client

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

**NOTE:** When you remove the Command Console Client, the `swcc2.mdb` file is deleted. This file contains the Navigation Tree configuration. If you want to save this information, move the file to another directory.

1. Click *Start > Programs > Command Prompt* and change to the directory to which you installed the Command Console Client.
2. Enter the following command:  
`C:\Program Files\Compaq\SWCC> AsyncEventService -remove`
3. Do one of the following:
  - On Windows NT 4.0, click *Start > Settings > Control Panel*, and then double-click the Add/Remove Programs icon in the Control Panel. The Add/Remove Program Properties window appears.
  - On Windows 2000, click *Start > Settings > Control Panel > Add/Remove Programs*. The Add/Remove Program window appears.
4. Select Command Console in the window.
5. Do one of the following:
  - On Windows NT 4.0, click *Add/Remove*.
  - On Windows 2000, click *Change/Remove*.
6. Follow the instructions on the screen.

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

## **Where to Find Additional Information**

You can find additional information about SWCC by referring to the online Help and to *Compaq StorageWorks Command Console Version 2.3 for the MA6000 User Guide*.

### **About the User Guide**

*Compaq StorageWorks Command Console Version 2.3 for the MA6000 User Guide* contains additional information on how to use SWCC. Some of the topics in the user guide are the following:

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

### **About the Online Help**

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

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



# Index

## 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 4–7
- virtual disks B–8

### 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 4–8
- functions 4–1
- installing 4–5
- restarting 4–8

### array of disk drives 2–9

### assigning unit numbers 1–17

### assignment

- unit numbers
  - fabric topology 5–21
- unit qualifiers
  - fabric topology 5–21

### assignment of unit numbers

- fabric topology
  - partition 5–22
  - single disk 5–22

### asynchronous event service B–8

### automatic mode

- installation 3–6

### autospare

- enabling
  - fabric topology 5–24

### availability 2–16

## B

### backup

- cloning data 8–2
- subsystem configuration 8–1

## C

### cabling

- controller pair 5–11
- multiple-bus failover
  - fabric topology configuration 5–11
- single controller 5–5

### cache modules

- location 1–2
- read caching 1–9
- write-back caching 1–9
- write-through caching 1–10

### caching techniques

- mirrored 1–10
- read caching 1–9
- read-ahead caching 1–9

**I-2**     **MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide**

- write-back caching 1-9
- write-through caching 1-10
- CD-ROM
  - mounting 3-5
- changing switches
  - configuration options 5-25, 5-26
- choosing passwords
  - Agent 4-7
- chunk size
  - choosing for RAIDsets and stripesets 2-24
  - controlling stripesize 2-24
  - using to increase data transfer rate 2-26
  - using to increase request rate 2-24
  - using to increase write performance 2-27
- CHUNKSIZE 2-24
- CLI commands
  - installation verification 5-9, 5-16
- CLI prompt
  - changing
    - fabric topology 5-22
- Client
  - adding 4-8
  - removing B-7
- CLONE utility
  - backup 8-2
- cloning
  - backup 8-2
- command console LUN 1-11
  - SCSI-2 mode 1-20
  - SCSI-3 mode 1-20
- Compaq Insight Manager B-8
- comparison of container types 2-9
- configuration
  - backup 8-1
  - fabric topology
    - devices 5-16
    - multiple-bus failover
      - cabling 5-11
    - multiple-bus failover using CLI 7-4
    - single controller cabling 5-5
  - fibre channel fabric topology
    - using CLI 5-1
    - using SWCC 5-1
  - restoring 2-27
  - rules 2-3
- 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 5-22
  - 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-8
- configuring devices
  - fabric topology 5-16
- connection
  - hardware 3-5
- connections 1-12
  - naming 1-12
- containers
  - attributes 2-7
  - comparison 2-9
  - illustrated 2-8
  - mirrorsets 2-15
  - planning storage 2-7
  - stripesets 2-13
- controller verification
  - installation 5-9, 5-16

- controllers
  - cabling 5-5, 5-11
  - location 1-2
  - node IDs 1-28
  - other defined xiv
  - this defined xiv
  - verification of installation 5-9, 5-16
  - worldwide names 1-28
- creating
  - storage set and device profiles 2-10
- creating and tuning
  - filesystem 3-12
- creation
  - file system 3-12

## D

- data transfer rate 2-26
- DESTROY 2-27
- device switches
  - changing
    - fabric topology 5-26
- devices
  - changing switches
    - fabric topology 5-25
  - configuration
    - fabric topology 5-16
  - creating a profile 2-10
- disk drives
  - adding
    - fabric topology 5-23
  - adding to the spareset
    - fabric topology 5-23
  - array 2-9
  - corresponding storage sets 2-28
  - dividing 2-19
  - removing from the spareset
    - fabric topology 5-24
- displaying the current switches
  - fabric topology 5-25
- dividing storage sets 2-19

## E

- enabling
  - switches 2-22
- erasing metadata 2-27
- establishing a local connection 5-2

## F

- fabric topology
  - configuration
    - single controller cabling 5-5
- failover 1-3
  - multiple-bus 1-7
  - transparent 1-4
- file system
  - creation 3-12
  - tuning 3-12
- filesystem
  - creating and tuning 3-12
  - preparing LUNs 3-9
- format utility
  - labeling LUNs 3-10
- functions
  - Agent 4-1

## G

- geometry
  - initialize switches 2-28

## H

- hardware
  - connections 3-5
  - installation 3-5
- host access
  - restricting by offsets
    - multiple-bus failover 1-27
    - transparent failover 1-23
  - restricting in multiple-bus failover mode
    - disabling access paths 1-24
  - restricting in multiple-bus failover mode 1-24
  - restricting in transparent failover mode 1-21
    - disabling access paths 1-23
    - separate links 1-22

**I-4**     *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Sun Solaris Installation and Configuration Guide*

host adapter  
  installation 3-1  
  preparation 3-1  
host connections 1-12  
  naming 1-12

**I**

initialize switches  
  changing  
    fabric topology 5-26  
  CHUNKSIZE 2-24  
  DESTROY 2-27  
  geometry 2-28  
  NODESTROY 2-27  
  NOSAVE\_CONFIGURATION 2-27  
  SAVE\_CONFIGURATION 2-27  
installation  
  automatic mode 3-4, 3-6  
  controller verification 5-9, 5-16  
  hardware 3-5  
  host adapter 3-2  
  invalid network port assignments B-3  
  manual mode 3-4, 3-8  
  solution software 3-5  
  there is no disk in the drive message B-4  
installation verification  
  CLI commands 5-9, 5-16  
installing  
  Agent 4-5  
integrating SWCC B-8  
invalid network port assignments B-3

**J**

JBOD 2-9

**L**

labeling LUNs  
  format utility 3-10  
LOCATE  
  find devices 2-31  
location  
  cache module 1-2  
  controller 1-2

LUN IDs  
  general description 1-29  
LUN presentation 1-18  
LUNs  
  labeling 3-9  
  using format utility 3-10  
  preparing for use by filesystem 3-9

**M**

mail messages  
  RAIDManager 4-9  
maintenance port connection  
  establishing a local connection 5-2  
  illustrated 5-2  
manual mode  
  installation 3-8  
mapping storagesets 2-28  
messages  
  there is no disk in the drive B-4  
mirrored caching  
  enabling 1-10  
  illustrated 1-10  
mirrorset switches  
  changing  
    fabric topology 5-25  
mirrorsets  
  planning considerations 2-15  
  important points 2-16  
  switches 2-23  
mounting the CD-ROM 3-5  
moving storagesets 8-6  
multiple-bus failover 1-7  
  ADD CONNECTIONS command 1-19  
  ADD UNIT command 1-19  
  ADD UNITcommand 1-19  
  CLI configuration procedure  
    fabric topology 7-4  
  fabric topology configuration  
    cabling 5-11  
  host connections 1-19  
  restricting host access 1-24  
    disabling access paths 1-24

- restricting host access by offsets 1-27
  - SET CONNECTIONS command 1-19
  - SET UNITcommand 1-19
- N**
- network port assignments B-3
  - node ID
    - restoring 1-29
  - node IDs 1-28
  - NODE\_ID
    - worldwide name 1-28
  - NODESTROY 2-27
  - NOSAVE\_CONFIGURATION 2-27
- O**
- offset
    - LUN presentation 1-18
    - restricting host access
      - multiple-bus fafileover 1-27
      - transparent fafileover 1-23
    - SCSI version factor 1-19
  - online help
    - SWCC B-8
  - options
    - for mirrorsets 2-23
    - for RAIDsets 2-23
    - initialize 2-24
  - other controller 1-2
    - defined xiv
  - overwriting data 2-27
- P**
- pager notification B-8
    - configuring B-8
  - partitions
    - assigning a unit number
      - fabric topology 5-22
    - defining 2-20
    - planning considerations 2-19
      - guidelines 2-21
  - passwords
    - choosing 4-7
  - performance 2-16
  - planning
    - overview 2-10
    - striped mirrorsets 2-19
    - stripesets 2-13
  - planning considerations 2-16
  - planning storage
    - containers 2-7
    - where to start 2-2
  - planning storagesets
    - characteristics
      - changing switches 2-22
      - enabling switches 2-22
      - initialization switch 2-21
      - storage switch 2-21
      - unit switch 2-21
    - switches
      - initialization 2-24
      - storageset 2-22
      - unit 2-28
  - profiles
    - creating 2-10
    - description 2-10
    - storageset A-1
      - example A-2
- R**
- RAIDManager
    - mail messages 4-9
  - RAIDset switches
    - changing
      - fabric topology 5-25
  - RAIDsets
    - choosing chunk size 2-24
    - maximum membership 2-17
    - planning considerations 2-16
      - important points 2-17
    - switches 2-23
  - read caching
    - enabled for all storage units 1-9
    - general description 1-9

- read requests
  - decreasing the subsystem response time with
    - read caching 1-9
- read-ahead caching 1-9
  - enabled for all disk units 1-9
- related publications xv
- removing
  - Client B-7
- request rate 2-24
- requirements
  - host adapter installation 3-2
- 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 2-27
- saving configuration 2-27
- SCSI version
  - offset 1-19
- SCSI-2
  - assigning unit numbers 1-20
  - command console lun 1-20
- SCSI-3
  - assigning unit numbers 1-20
  - command console lun 1-20
- selective unit presentation 1-21
- SET CONNECTIONS
  - multiple-bus failover 1-19
  - transparent failover 1-17
- SET UNIT
  - multiple-bus failover 1-19
- settings
  - controller configuration handling 2-27
  - single disk (JBOD)
    - assigning a unit number
      - fabric topology 5-22
  - single Model 4214R disk enclosure
    - storage map
      - template 1 A-3
  - single Model 4254 disk enclosure
    - storage map
      - template 2 A-4
  - software
    - verification 3-3
  - solution software 3-2
    - installation 3-5
  - storage
    - creating map 2-28
    - profile
      - example A-2
  - storage map 2-28
    - example 2-30
    - template 1 A-3
      - single Model 4214R disk enclosure A-3
    - template 2 A-4
      - single Model 4254 disk enclosure A-4
    - template 3 A-5
      - two Model 4254 disk enclosures A-5
  - storage set profile 2-10
  - storage set switches
    - SET command 2-22
  - storage sets
    - creating a profile 2-10
    - deleting
      - fabric topology 5-25
    - fabric topology
      - changing switches 5-25
    - moving 8-6
    - planning considerations 2-12
      - mirrorsets 2-15
      - partitions 2-19
      - RAIDsets 2-16
      - striped mirrorsets 2-18
      - stripesets 2-12

- profile 2-10
  - example 2-11
- profiles A-1
- striped mirrorsets
  - planning 2-19
  - planning considerations 2-18
- stripesets
  - distributing members across buses 2-14
  - planning 2-13
  - planning considerations 2-12
    - important points 2-13
- subsystem
  - adding 4-7
  - saving configuration 2-27
- subsystem configuration
  - backup 8-1
- SWCC 4-1
  - additional information B-8
  - integrating B-8
  - online help B-8
- switches
  - changing 2-22
  - changing characteristics 2-21
  - CHUNKSIZE 2-24
  - DESTROY 2-27
  - enabling 2-22
  - mirrorsets 2-23
  - NODESTROY 2-27
  - NOSAVE\_CONFIGURATION 2-27
  - RAIDset 2-23
  - SAVE\_CONFIGURATION 2-27
- switches for storagesets
  - overview 2-21
- system
  - preparing LUNs 3-9

**T**

- templates
  - subsystem profile A-1
- terminology
  - cache module A and B 1-1
  - controller A and B 1-1

- other controller 1-2
- this controller 1-2
- this controller 1-2
  - defined xiv
- transparent failover 1-4
  - 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-3
  - there is no disk in the drive message B-4
- tuning
  - file system 3-12
- two Model 4254 disk enclosures
  - storage map
    - template 3 A-5

**U**

- uninstalling
  - Client Client
    - uninstalling B-7
- unit numbers
  - assigning 1-17
    - fabric topology 5-21
  - assigning depending on SCSI version 1-20
  - assigning in fabric topology
    - partition 5-22
    - single disk 5-22
- unit qualifiers
  - assigning
    - fabric topology 5-21
- unit switches
  - changing
    - fabric topology 5-26
- units
  - LUN IDs 1-29

**V**

- verification
  - controller installation 5-9, 5-16
  - software 3-3
- verification of installation
  - controller 5-9, 5-16
- virtual disks
  - adding B-8
- vold
  - not running 3-6
  - running 3-5

**W**

- worldwide name
  - NODE\_ID 1-28
  - REPORTED PORT\_ID 1-28
  - restoring 1-29
- worldwide names 1-28
- write performance 2-27
- write requests
  - improving the subsystem response time with
    - write-back caching 1-9
  - placing data with with write-through caching 1-10
- write-back caching
  - general description 1-9
- write-through caching
  - general description 1-10