

Compaq StorageWorks™

HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX

Installation and Configuration Guide

First Edition (December 2000)
Part Number: AA-RNJ8A-TE
Compaq Computer Corporation

2000 Compaq Computer Corporation.

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

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

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

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

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

THE INFORMATION IN THIS PUBLICATION 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 AND WHETHER IN AN ACTION OF CONTRACT OR TORT, INCLUDING NEGLIGENCE.

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.

HGS60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX
Installation and Configuration Guide
First Edition (December 2000)
Part Number: AA-RNJ8A-TE

Contents

About This Guide

Text Conventions	xv
Symbols in Text	xvi
Symbols on Equipment	xvii
Rack Stability	xviii
Solution Software Kit Contents	xviii
Installation and Configuration Guide	xviii
Command Console Version 2.3 User Guide	xviii
CLI Reference Guide	xix
Maintenance and Service Guide	xix
CD-ROM	xix
Solution Kit Release Notes	xix
Warranty Terms and Conditions	xix
Installation and Setup of the CD-ROM	xx
Configuration Flowchart	xx
Getting Help	xxiii
Compaq Technical Support	xxiii
Compaq Website	xxiv
Compaq Authorized Reseller	xxiv

Chapter 1

Planning a Subsystem

Defining the Subsystems	1- 2
For the HSG60	1- 2
Controller Designations A and B	1- 2
Controller Designations “This Controller” and “Other Controller”	1- 3
For the HSG80	1- 4
Controller Designations A and B	1- 4
Controller Designations “This Controller” and “Other Controller”	1- 5
Failover Mode	1- 7
Transparent Failover Mode	1- 7
Selecting a Cache Mode	1- 9
Read Caching	1- 10
Read-Ahead Caching	1- 10
Write-Back Caching	1- 10
Write-Through Caching	1- 10
Enabling Mirrored Caching	1- 11
The Command Console LUN	1- 11
Determining the Address of the CCL	1- 12
In SCSI II Mode	1- 12
Disabling the CCL	1- 12
Enabling the CCL	1- 12
In SCSI III Mode	1- 12
Connections	1- 13
Naming Connections	1- 13
Numbers of Connections	1- 13
Assigning Unit Numbers	1- 16
Matching Units to Host Connections in Transparent Failover Mode	1- 16
Assigning Unit Numbers Depending on SCSI_VERSION	1- 18
The CCL in SCSI-3 mode	1- 18
The CCL in SCSI-2 Mode	1- 19
Restricting Host Access (Selective Storage Presentation)	1- 19
Restricting Host Access in Transparent Failover Mode	1- 20
Restricting Host Access by Separate Links	1- 20
Restricting Host Access by Disabling Access Paths	1- 21
Restricting Host Access by Offsets	1- 22
Worldwide Names (Node IDs and Port IDs)	1- 22
Restoring Worldwide Names (Node IDs)	1- 23
For the HSG60	1- 23
For the HSG80	1- 24
Unit World Wide Names (LUN IDs)	1- 25

Chapter 2

Planning Storage

Where to Start	2- 2
Configuration Rules	2- 3
Device PTL Addressing Convention	2- 3
For the HSG60	2- 3
For the HSG80	2- 6
Determining Storage Requirements	2- 13
Choosing a Container Type	2- 14
Creating a Storageset Profile	2- 15
Storageset Planning Considerations	2- 17
Stripeset Planning Considerations	2- 17
Mirrorset Planning Considerations	2- 19
RAIDset Planning Considerations	2- 21
Striped Mirrorset Planning Considerations	2- 23
Partition Planning Considerations	2- 24
Defining a Partition	2- 25
Guidelines for Partitioning Storagesets and Disk Drives	2- 25
Changing Characteristics through Switches	2- 26
Enabling Switches	2- 26
Changing Switches	2- 26
Storageset and Partition Switches	2- 27
RAIDset Switches	2- 27
Mirrorset Switches	2- 27
Partition Switches	2- 28
Initialization Switches	2- 28
Chunk Size	2- 28
Increasing the Request Rate	2- 29
Increasing the Data Transfer Rate	2- 30
Increasing Sequential Write Performance	2- 31
Save Configuration	2- 31
Destroy/Nodestroy	2- 31
Geometry	2- 32
Unit Switches	2- 32
Storage Maps	2- 32
Creating a Storage Map	2- 32
Example Storage Map for the HSG60	2- 33
Example Storage Maps for the HSG80	2- 35
Example Storage Map - Model 4310R Disk Enclosure	2- 35
Example Storage Map - Model 4350R Disk Enclosure	2- 37
Example Storage Map - Model 4314R Disk Enclosure	2- 38

vi *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

Example Storage Map - Model 4354R Disk Enclosure	2- 41
Using the LOCATE Command to Find Devices	2- 43

Chapter 3

Preparing the Host System

HSG60\HSG80 Array Controller Setup	3- 2
Arbitrated Loop Setup	3- 2
Fabric Setup	3- 3
Host Bus Adapter Installation	3- 5
HBA Driver Installation	3- 5
HBA Driver Removal	3- 6
Hardware Installation	3- 6
Configuration	3- 6
Connection	3- 8
Loading StorageWorks Solutions Software Version 8.5c onto Hard Disk	3- 8
Copying the File Archive from CD-ROM to Temporary Storage	3- 9
Executing the Installation and Configuration Script	3- 9
New Connection Operating System	3- 10
Using the StorageWorks Installation Utility Maintenance Options	3- 10
Upgrading the StorageWorks Solution Software	3- 11
HBA Driver Upgrade	3- 11
Updating ODM Entries	3- 13
Agent Upgrade	3- 15
Importing Volume Groups	3- 17
Host Device Initialization	3- 17
Uninstalling StorageWorks Solutions Software Version 8.5	3- 18

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
Before You Install the Agent	4- 5
Installing and Configuring the Agent	4- 5
Installing the Agent for the First Time	4- 6
Upgrading the Agent	4- 8
Running the Agent	4- 9
Reconfiguring the Agent	4- 9
Removing the Agent	4- 9
High Availability Cluster Multi-Processing (HACMP)	4- 10
Enabler Software	4- 10
General Note on Software Installation	4- 11

SWCC Clients	4- 11
Running the Agent	4- 12
HACMP Implementation for the Agent	4- 12
Prerequisites	4- 12
Implementation	4- 13
Procedure for Installing the Agent on an HACMP Cluster	4- 14
Procedure for Configuring the Agent on an HACMP cluster	4- 14
Installing and Configuring the Client on an HACMP cluster	4- 15
Operation of the Agent on an HACMP Cluster	4- 15
Normal Operation of the Agent on an HACMP Cluster	4- 16
Starting, Stopping and Checking for the Agent	4- 17
Password and Notification	4- 17
Client Functions and Storage Subsystem Options	4- 18
Failover with Agent as Part of a Resource Group	4- 18
Client IP Address	4- 18
Event Scripts	4- 18
Notification within HACMP	4- 18
Node Isolation Recovery	4- 19
Failure of the Agent	4- 19
Polling all nodes in cluster for presence of a running Agent	4- 19
Restarting the Agent on surviving node	4- 19
Notification	4- 19
Supporting Tools	4- 19
HACMP for AIX Primary Event Scripts	4- 20
HACMP for AIX Secondary Event Scripts	4- 20

Chapter 5

Configuration Procedures

Establishing a Local Connection	5- 2
Setting Up a Single Controller	5- 3
Power On and Establish Communication	5- 3
Cabling a Single Controller	5- 3
Configuring a Single Controller using CLI	5- 4
Setting Up a Controller Pair	5- 9
Power On and Establish Communication	5- 9
Cabling a Controller Pair	5- 9
Configuring a Controller Pair using CLI	5- 10
Configuring Devices	5- 15
Configuring a Stripeset	5- 15
Configuring a Mirrorset	5- 15
Configuring a RAIDset	5- 16

Configuring a Striped Mirrorset	5- 17
Configuring a Single-Disk (JBOD) Unit.	5- 18
Configuring a Partition	5- 18
Assigning Unit Numbers and Unit Qualifiers.	5- 19
Assigning a Unit Number to a Storageset	5- 19
Assigning a Unit Number to a Single (JBOD) Disk	5- 19
Assigning a Unit Number to a Partition.	5- 20
Configuration Options.	5- 20
Changing the CLI Prompt	5- 20
Mirroring cache	5- 20
Adding Disk Drives	5- 21
Adding a Disk Drive to the Spareset	5- 21
Removing a Disk Drive from the Spareset	5- 21
Enabling Autospare	5- 22
Deleting a Storageset	5- 22
Changing Switches for a Storageset or Device	5- 23
Displaying the Current Switches	5- 23
Changing RAIDset and Mirrorset Switches	5- 23
Changing Device Switches	5- 24
Changing Initialize Switches	5- 24
Changing Unit Switches	5- 24

Chapter 6

Verifying Storage Configuration from the Host

Chapter 7

Configuration Example Using CLI

HSG60 Example	7- 2
CLI Configuration Example for the HSG60.	7- 4
HSG80 Example	7- 8
CLI Configuration Example for the HSG80.	7- 11

Chapter 8

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

Backing Up the Subsystem Configuration	8- 1
Cloning Data for Backup.	8- 2
Moving Storagesets.	8- 5

Appendix A

Subsystem Profile Templates

Storageset Profile	A- 2
------------------------------	------

Templates for the HSG60	A- 3
Storage Map Template 1 for the Model 4214R Disk Enclosure	A- 3
Storage Map Template 2 for the Model 4254 Disk Enclosure.	A- 4
Storage Map Template 3 for Two Model 4214R Disk Enclosures	A- 5
Templates for the HSG80	A- 6
Storage Map Template 1 for the BA370 Enclosure	A- 6
Storage Map Template 2 for the Second BA370 Enclosure.	A- 7
Storage Map Template 3 for the Third BA370 Enclosure	A- 8
Storage Map Template 4 for the Model 4214R Disk Enclosure	A- 9
Storage Map Template 5 for the Model 4254 Disk Enclosure.	A- 10
Storage Map Template 6 for the Model 4310R Disk Enclosure	A- 12
Storage Map Template 7 for the Model 4350R Disk Enclosure	A- 14

Appendix B

Installing, Configuring, and Removing the Client

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

Figures

- Figure 1. General configuration flowchart xxi
- Figure 2. Configuring storage with the CLI xxii
- Figure 3. Configuring storage with SWCC xxiii
- Figure 1-1. Location of controllers and cache modules in a Model 2100 enclosure 1-2
- Figure 1-2. “This controller” and “other controller” for the Model 2100 enclosure 1-3
- Figure 1-3. Location of controllers and cache modules in a Model 2200 enclosure 1-4
- Figure 1-4. Location of controllers and cache modules in a BA370 enclosure 1-5
- Figure 1-5. “This controller” and “other controller” for the Model 2200 enclosure 1-6
- Figure 1-6. “This controller” and “other controller” for the BA370 enclosure 1-6
- Figure 1-7. Transparent failover - normal operation. 1-8
- Figure 1-8. Transparent failover - after failover from controller B to controller A 1-9
- Figure 1-9. Mirrored caching 1-11
- Figure 1-10. Connections in separate-link, transparent failover mode configurations. 1-14
- Figure 1-11. Connections in single-link, transparent failover mode configurations. 1-15
- Figure 1-12. LUN presentation to hosts, as determined by offset 1-17
- Figure 1-13. Limiting host access in transparent failover mode. 1-21
- Figure 1-14. Placement of the worldwide name label on the Model 2100 enclosure. 1-23
- Figure 1-15. Placement of the worldwide name label on the Model 2200 enclosure. 1-24
- Figure 1-16. Placement of the worldwide name label on the BA370 enclosure. 1-24
- 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 drive	2-6
Figure 2-5. PTL naming convention.	2-7
Figure 2-6. PTL addressing in a configuration in a BA370 enclosure	2-7
Figure 2-7. PTL addressing in a single-bus configuration, using six Model 4310R enclosures	2-9
Figure 2-8. PTL addressing in a dual-bus configuration, using three Model 4350R enclosures	2-10
Figure 2-9. PTL addressing in a single-bus configuration, using six Model 4314R enclosures	2-11
Figure 2-10. PTL addressing in a dual-bus configuration, using three Model 4354R enclosures	2-12
Figure 2-11. Mapping a unit to physical disk drives	2-13
Figure 2-12. Container types	2-14
Figure 2-13. An example storageprofile	2-16
Figure 2-14. A three-member RAID 0 stripeset (example 1).	2-17
Figure 2-15. A three-member RAID 0 stripeset (example 2).	2-18
Figure 2-16. Mirrorsets maintain two copies of the same data.	2-20
Figure 2-17. Mirrorset example 2.	2-20
Figure 2-18. A five-member RAIDset using parity	2-21
Figure 2-19. Striped mirrorset (example 1)	2-23
Figure 2-20. Striped mirrorset (example 2)	2-24
Figure 2-21. One example of a partitioned single-disk unit.	2-24
Figure 2-22. Chunk size larger than the request size	2-29
Figure 2-23. Chunk size smaller than the request size	2-30
Figure 2-24. Example storage map-single bus for the HSG60.	2-33
Figure 2-25. Example storage map-dual bus for the HSG60	2-34
Figure 2-26. Model 4310R disk enclosure - example storage map	2-36
Figure 2-27. Model 4350R disk enclosure - example storage map	2-38
Figure 2-28. Model 4314R disk enclosure - example storage map	2-40
Figure 2-29. Model 4354R disk enclosure - example storage map	2-42
Figure 4-1. An example of a network connection	4-4
Figure 4-2. Command Console Client/Agent connection in a typical HACMP cluster.	4-13
Figure 5-1. Maintenance port connection.	5-2
Figure 5-2. Single controller cabling	5-4
Figure 5-3. Controller pair failover cabling	5-10
Figure 7-1. Example storage map for HSG60	7-2
Figure 7-2. Example system for the HSG60.	7-3
Figure 7-3. Example virtual layout from the hosts' point of view	7-3
Figure 7-4. Example storage map for the BA370.	7-8
Figure 7-5. Example system for the HSG80.	7-9
Figure 7-6. Example virtual layout from the hosts' point of view	7-10
Figure 7-7. Example system for the HSG60.	7-10
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-5
Figure B-3. Navigation window showing expanded "Atlanta" host icon.	B-6

Tables

Table 1–1 Unit Assignments and SCSI_VERSION.....	1–19
Table 2–1 A Comparison of Container Types	2–15
Table 2–2 Example Chunk Sizes	2–30
Table 3–1 Correct Controller Parameter Settings for Arbitrated Loop	3–2
Table 3–2 HSG60/HSG80 Controller Parameters and Recommended Settings for Fabric	3–4
Table 4–1 SWCC Features and Components	4–2
Table 4–2 Installation and Configuration Overview	4–3
Table 4–3 Minimum IBM-AIX Host Requirements	4–5
Table 4–4 Information Required for Install Program Prompts.....	4–6
Table 4–5 HACMP Software Prerequisites	4–12
Table 4–6 HACMP Event Scripts.....	4–21
Table 5–1 Correct Controller Parameter Settings.....	5–4
Table 5–2 Correct Controller Parameter Settings.....	5–10

About This Guide

This guide is designed to be used as step-by-step instructions for installation and as a reference for operation, troubleshooting, and future upgrades.

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

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

www.compaq.com/products/registration

This guide describes:

- considerations while planning a configuration
- configuration procedures

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

Text Conventions

This document uses the following conventions to distinguish elements of text:

Keys

Keys appear in boldface. A plus sign (+) between two keys indicates that they should be pressed simultaneously.

USER INPUT	User input appears in a different typeface and in uppercase
<i>FILENAMES</i>	File names appear in uppercase italics.
Menu Options, Command Names, Dialog Box Names	These elements appear in initial capital letters.
COMMANDS, DIRECTORY NAMES, and DRIVE NAMES	These elements appear in upper case. [NOTE: UNIX commands are case sensitive and will not appear in uppercase]
Type	When you are instructed to <i>type</i> information, type the information without pressing the Enter key.
Enter	When you are instructed to enter information, type the information and then press the Enter key.

Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.



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



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

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

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

Symbols on Equipment



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

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



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

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



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

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



Power Supplies or Systems marked with these symbols indicate the equipment is supplied by multiple sources of power.

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



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal INJURY or damage to the equipment, observe local occupational health and safety requirements and guidelines for manual material handling.

Rack Stability



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

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

Solution Software Kit Contents

The solution software kit contains the software and documentation you need to install and configure your new Compaq StorageWorks subsystem. A description of the kit components follows:

Installation and Configuration Guide

Begin with this guide for your RAID system installation. It describes how to:

- Plan your subsystem
- Plan your storage configuration
- Make your controller work with the Host O/S
- Make the Host O/S compatible with the controller
- Cable and configure your controller
- Verify your storage
- Back up the subsystem, clone data and move storagesets
- Install Client

Command Console Version 2.3 User Guide

This guide explains how to use the SWCC Agent and the SWCC Client Graphical User Interface (GUI) to configure and monitor the storage system from a host server.

CLI Reference Guide

This guide contains a comprehensive description of the CLI commands available by way of a maintenance terminal.

Maintenance and Service Guide

This guide contains illustrated instructions for accessing and restoring/replacing system components should the subsystem require servicing.

CD-ROM

You can find all solutions and setup software required for RAID system installation and soft copy documentation on the CD-ROM that is included in this Solution Software Kit.

Solution Kit Release Notes

Refer to the release notes for platform-specific information on features and characteristics of the StorageWorks HSG60/HSG80 Array Controller Solution Software that may not be covered elsewhere in the documentation.

Warranty Terms and Conditions

This document contains the warranty terms and conditions for the products specified.

Installation and Setup of the CD-ROM

Contents of the CD-ROM:

- Device Drivers
- Solution Software
- SWCC Agent and Client software

Documentation included on the CD-ROM:

- Installation and Configuration Guide
- Storage Configuration Application Notes (please check our website for the latest storage configuration Application Notes)
- Heterogeneous Storage Area Network Application Notes

NOTE: If any of the components listed in the Solution Software Kit are missing, please contact your customer service representative.

Configuration Flowchart

This guide provides a three-part flowchart to follow while configuring a new storage subsystem:

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

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

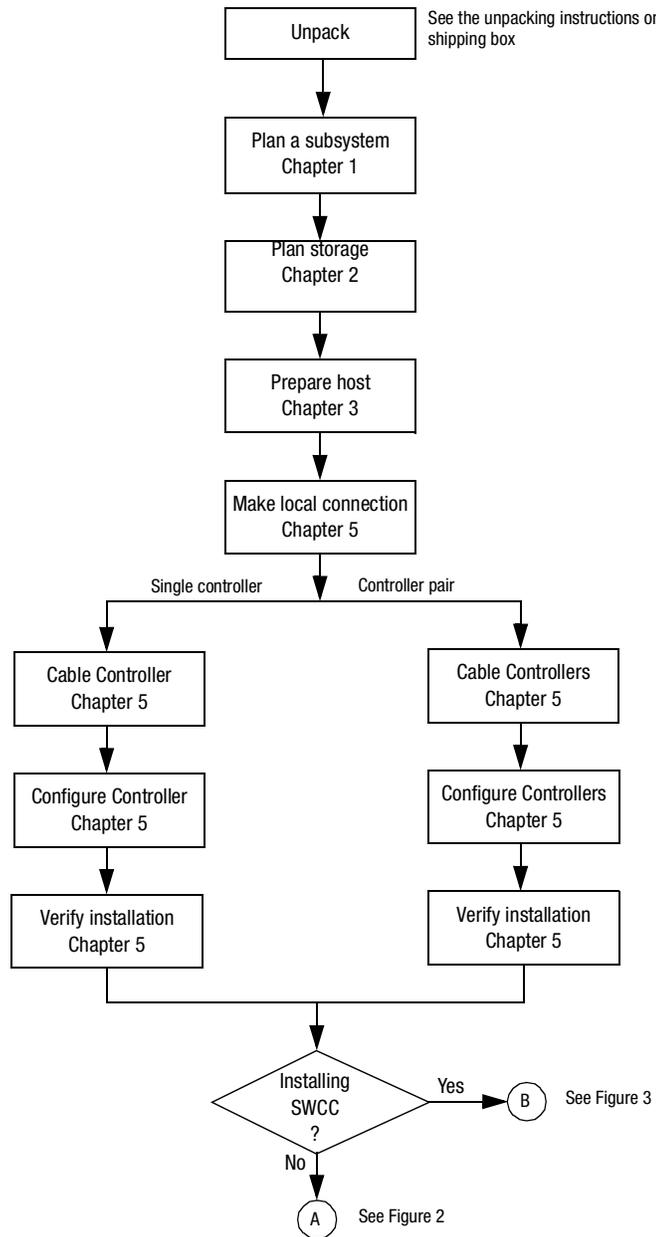


Figure 1. General configuration flowchart

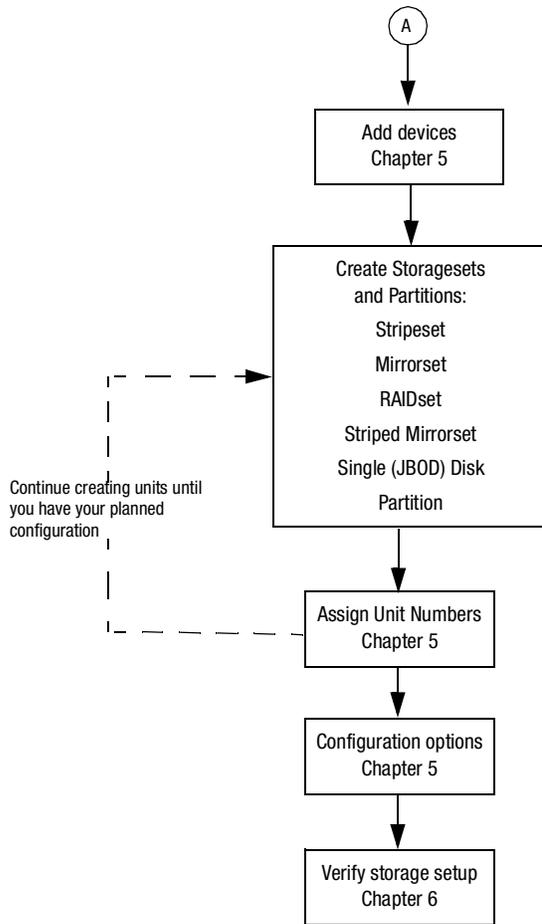


Figure 2. Configuring storage with the CLI

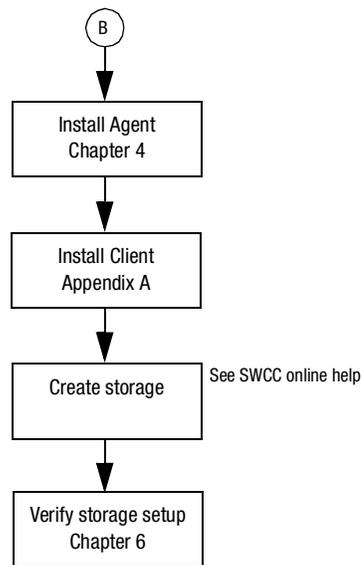


Figure 3. Configuring storage with SWCC

Getting Help

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

Compaq Technical Support

You are entitled to free hardware technical telephone support for your product for as long you own the product. A technical support specialist will help you diagnose the problem or guide you to the next step in the warranty process.

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

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

xxiv *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

Outside North America, call the nearest Compaq Technical Support Phone Center. Telephone numbers for world wide Technical Support Centers are listed on the Compaq website. Access the Compaq website by logging on to the Internet at:

<http://www.compaq.com>

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

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

Compaq Website

The Compaq website has latest information on this product as well as the latest drivers. You can access the Compaq website by logging on to the Internet at:

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

Compaq Authorized Reseller

For the name of your nearest Compaq Authorized Reseller:

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

Chapter 1

Planning a Subsystem

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:

- “Defining the Subsystems,” page 1–2
- “Failover Mode,” page 1–7
- “Selecting a Cache Mode,” page 1–9
- “Enabling Mirrored Caching,” page 1–11
- “The Command Console LUN,” page 1–11
- “Connections,” page 1–13
- “Assigning Unit Numbers,” page 1–16
- “Restricting Host Access (Selective Storage Presentation),” page 1–19
- “Worldwide Names (Node IDs and Port IDs),” page 1–22

Defining the Subsystems

For the HSG60

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.

Model 2100 Enclosure

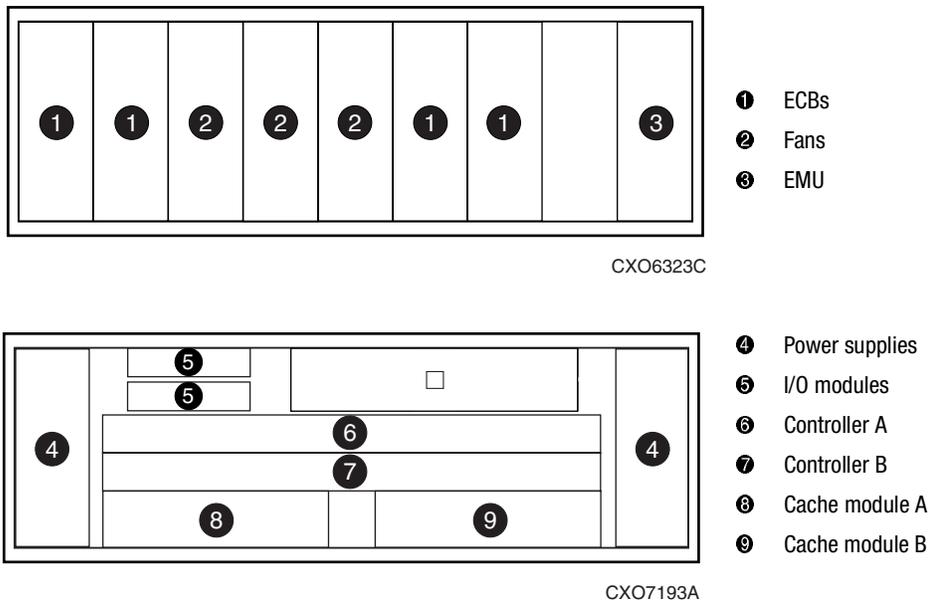


Figure 1-1. Location of controllers and cache modules in a Model 2100 enclosure

Controller Designations “This Controller” and “Other Controller”

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

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

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

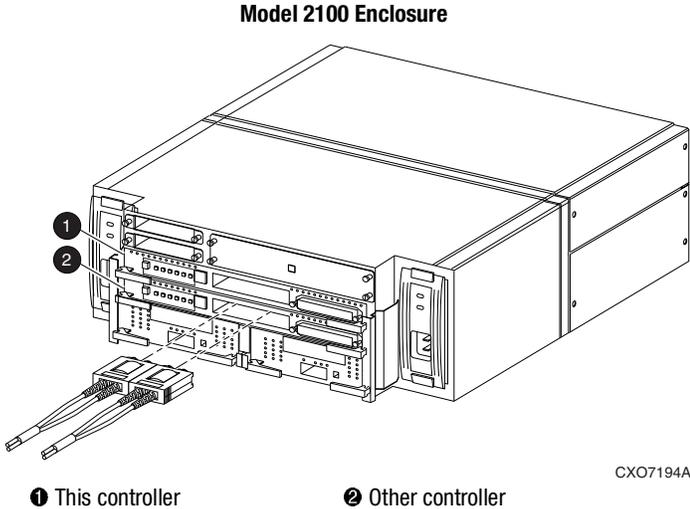


Figure 1-2. “This controller” and “other controller” for the Model 2100 enclosure

For the HSG80

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 1-3 for the Model 2200 enclosure and as shown in Figure 1-4 for the BA370 enclosure.

Model 2200 Enclosure

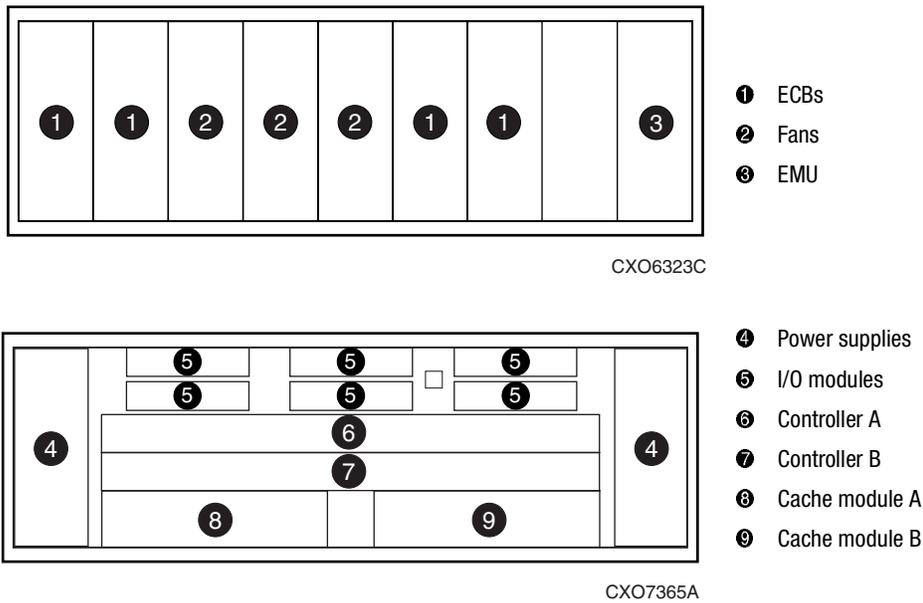


Figure 1-3. Location of controllers and cache modules in a Model 2200 enclosure

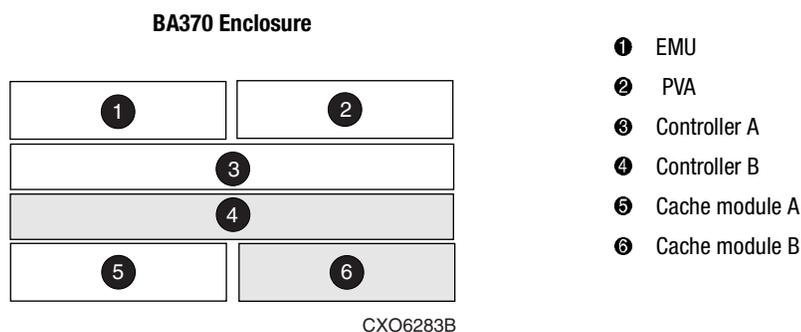


Figure 1-4. Location of controllers and cache modules in a BA370 enclosure

Controller Designations “This Controller” and “Other Controller”

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

- **“this controller”**—the controller that is the focus of the CLI session. “This controller” is the controller to which the maintenance terminal is attached and through which the CLI commands are being entered. “This controller” can be 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 1-5 shows the relationship between “this controller” and “other controller” in a Model 2200 enclosure and Figure 1-6 shows the same relationship in a BA370 enclosure.

1-6 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

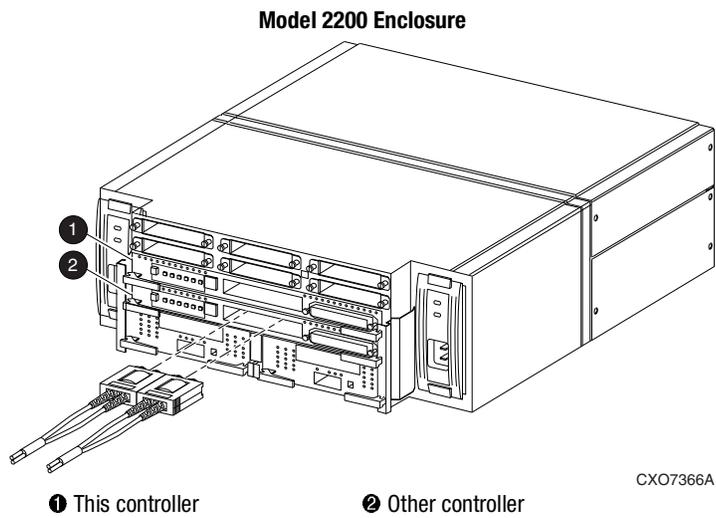


Figure 1-5. "This controller" and "other controller" for the Model 2200 enclosure

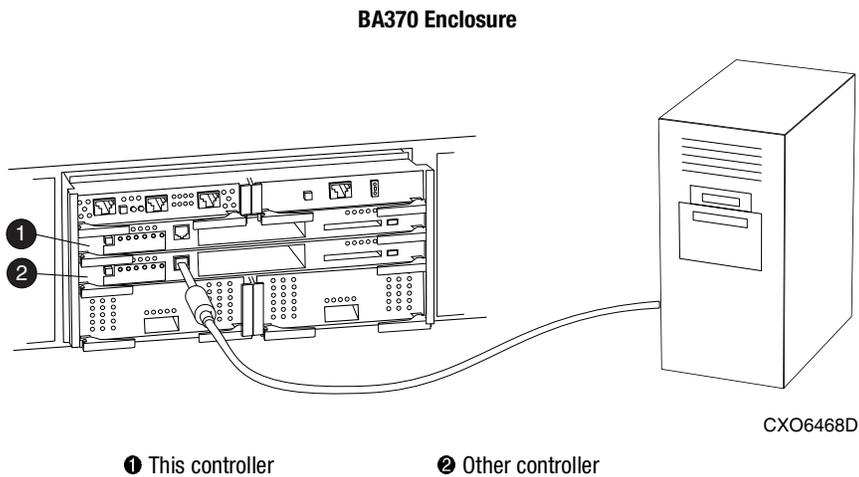


Figure 1-6. "This controller" and "other controller" for the BA370 enclosure

Failover Mode

Failover is a way to keep the storage array available to the host if one of the controllers becomes unresponsive. A controller can become unresponsive because of a hardware failure such as a host bus adapter or 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 and is only supported when Secure Path is used.

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

NOTE: Multiple-bus failover is not supported on IBM-AIX host platforms.

Transparent Failover Mode

Transparent failover mode has the following characteristics:

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

A unit or storage unit is a physical or virtual device of your subsystem. It is typically assigned a logical unit number (LUN) and is managed by the HSG60/HSG80 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 storage sets. Storage sets 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 1-7. The active and standby ports share port identity, enabling the

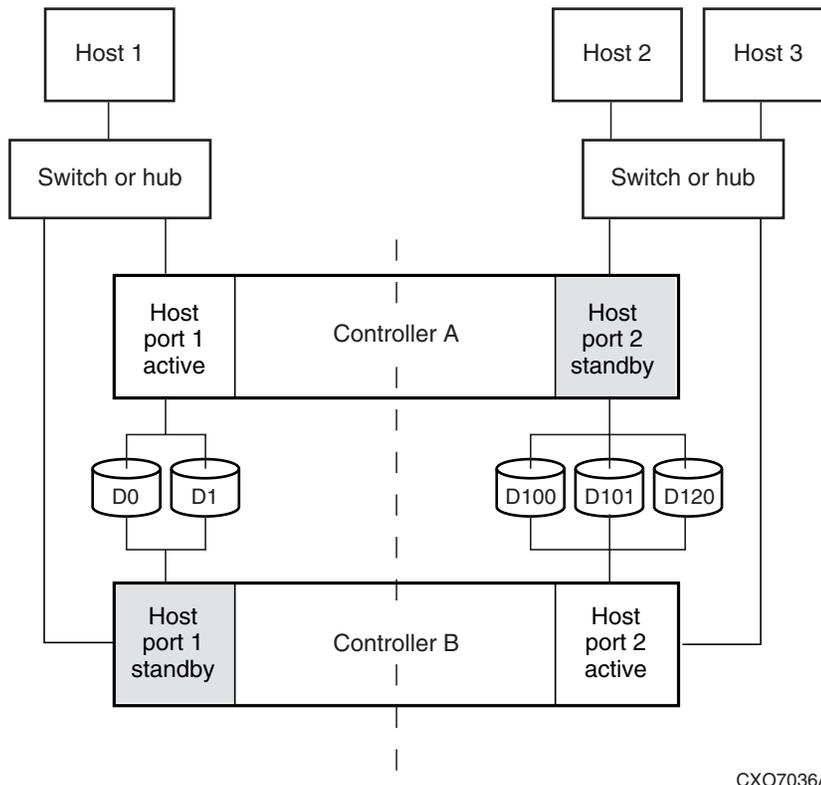
1-8 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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 1-8.

Units are divided between the host ports:

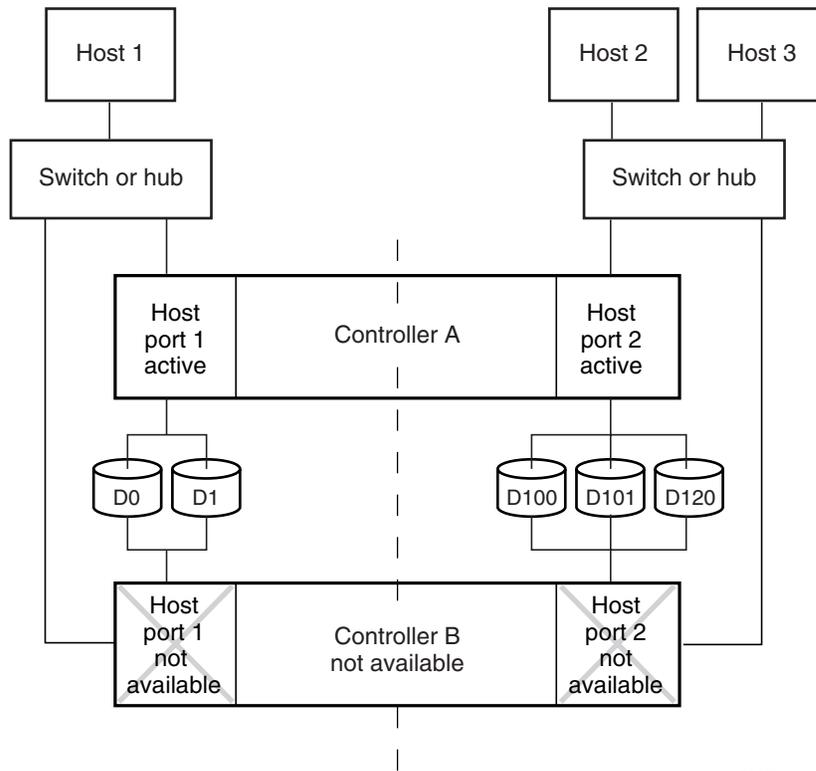
- Units 0-99 are on host port 1 of both controllers (but accessible only through the active port)
- Units 100-199 are on host port 2 of both controllers (but accessible only through the active port)

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 1-7. Transparent failover - normal operation



CXO7035A

Figure 1-8. Transparent failover - after failover from controller B to controller A

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 1-9.

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

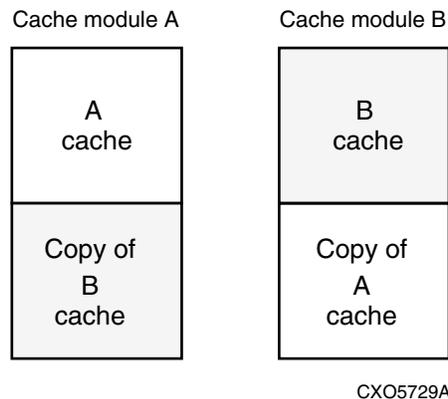


Figure 1-9. Mirrored caching

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/HSG80 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/HSG80 controller on Host Port 1. There are two modes that HSG60/HSG80 uses; SCSI-2 and SCSI-3. The Command Console LUN (CCL) does the following:

1–12 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

In dual-redundant controller configurations, the following commands alter the setting of the CCL on both controllers. The CCL is enabled only on Host Port 1.

Determining the Address of the CCL

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

```
HSG > SHOW THIS_CONTROLLER
```

NOTE: Set the prompt for the correct controller, HSG60 or HSG80. For the purpose of showing commands, this book uses the prompt HSG.

In SCSI II Mode

Disabling the CCL

To disable the CCL, enter the following CLI command:

```
HSG > SET THIS_CONTROLLER NOCOMMAND_CONSOLE_LUN
```

Enabling the CCL

To enable the CCL, enter the following CLI command:

```
HSG > 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.

NOTE: SCSI-2 mode is used for IBM-AIX platforms.

In SCSI III Mode

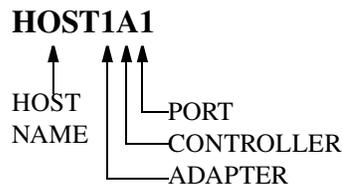
The CCL is enabled all the time. There is no option to enable/disable.

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.

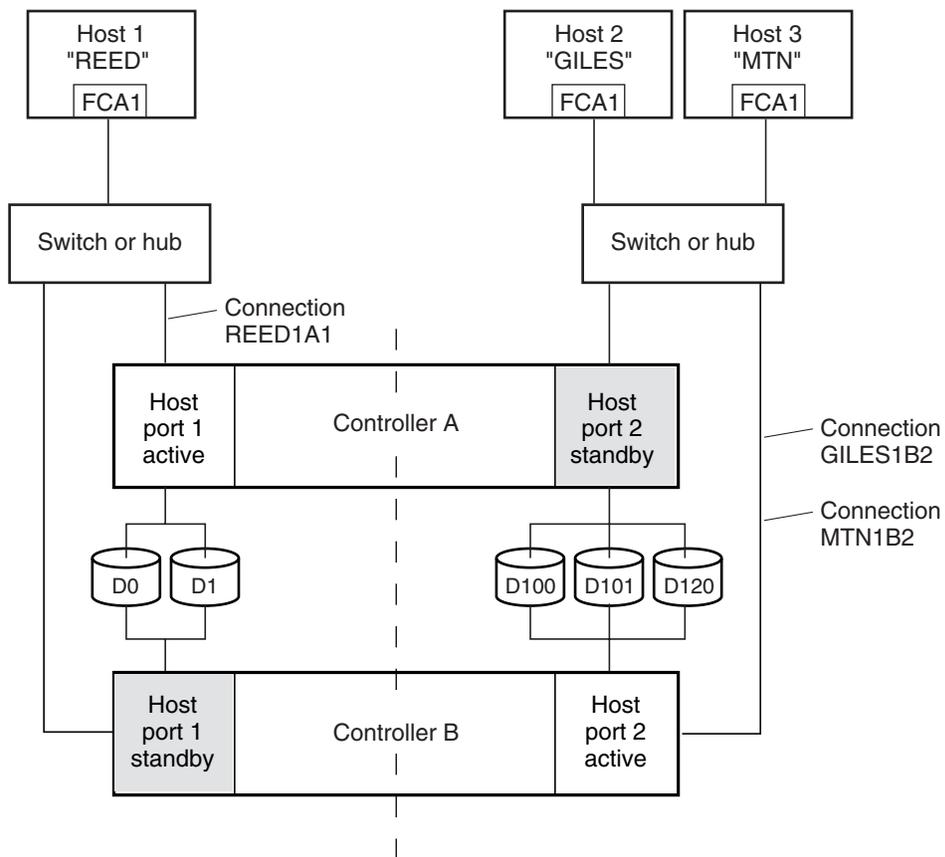
NOTE: Connection names can have a maximum of nine characters.

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 1-10.
- 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 1-11.

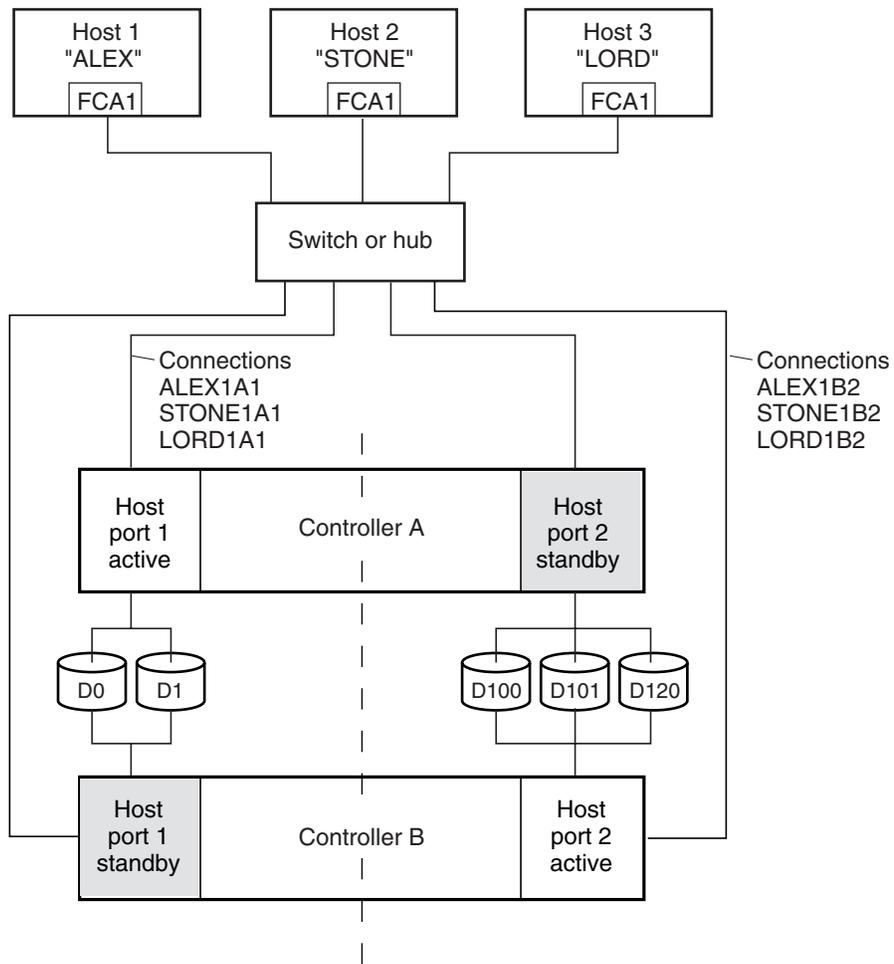
1-14 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide



NOTE: FCA = Fibre Channel Adapter

CXO7081A

Figure 1-10. Connections in separate-link, transparent failover mode configurations



NOTE: FCA = Fibre Channel Adapter

CXO7079A

Figure 1-11. Connections in single-link, transparent failover mode configurations

Assigning Unit Numbers

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

- Failover mode of the controller pair
- The *ENABLE_ACCESS_PATH* and *PREFERRED_PATH* switches in the *ADD UNIT* (or *SET unit*) commands
- The *UNIT_OFFSET* switch in the *ADD CONNECTIONS* (or *SET connections*) commands
- 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.

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

LUN number = unit number – 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)

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 1-12 shows how units are presented as different LUNs, depending on the offset of the host. In this illustration, host connection 1 and host connection 2 would need to be on host port 1; host connection 3 would need to be on host port 2.

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

CXO6455B

Figure 1-12. LUN presentation to hosts, as determined by offset

Offsets other than the default values can be specified. For example, unit D17 would be visible to a host connection on port 1 that had an offset of 10 as LUN 7 (unit number of 17 minus offset of 10). The unit would not be visible 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).

1–18 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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 1–18.

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 1–20.

Assigning Unit Numbers Depending on *SCSI_VERSION*

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

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

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

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:

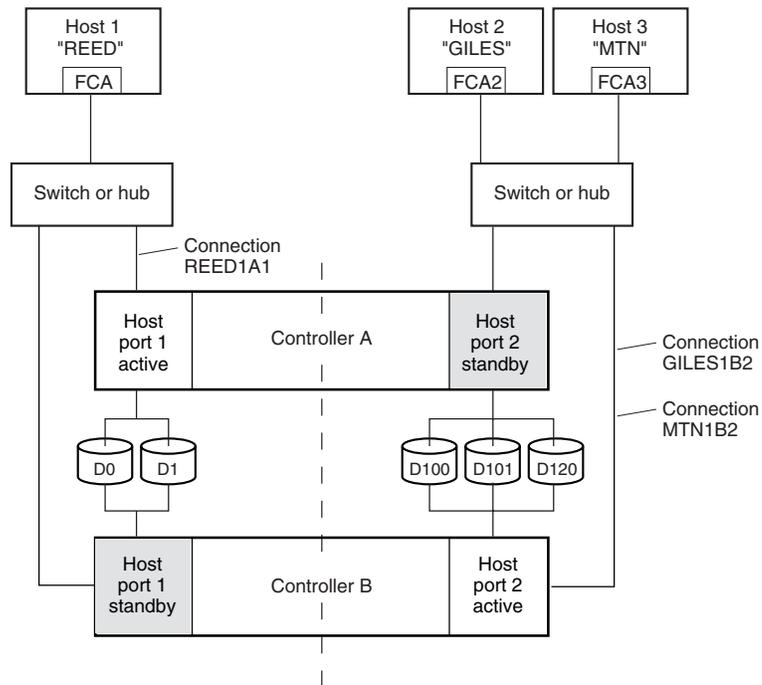
- Using separate Fibre Channel links (either loop or fabric)
- Enabling the access path of selected host connections on a shared loop or fabric
- Setting offsets

NOTE: These techniques also work for a single controller.

Restricting Host Access by Separate Links

In transparent failover mode, host port 1 of controller A and host port 1 of controller B share a common Fibre Channel link. Host port 2 of controller A and host port 2 of controller B also share a common Fibre Channel link. If the host 1 link is separate from the host 2 link, the simplest way 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 1-13. 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 1-21.



NOTE: FCA = Fibre Channel Adapter

CX07081A

Figure 1-13. Limiting host access in transparent failover mode

Restricting Host Access by Disabling Access Paths

If more than one host is on a link (that is, attached to the same port), host access can be limited by enabling the access of certain host connections and disabling the access of others. This is done through the *ENABLE_ACCESS_PATH* and *DISABLE_ACCESS_PATH* switches of the *ADD UNIT* or *SET unit* commands. 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 1-13, 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
```

1-22 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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 1-13, 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
```

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.

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

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.

For the HSG60

Figure 1-14 shows the placement of the worldwide name label for the Model 2100 enclosure.

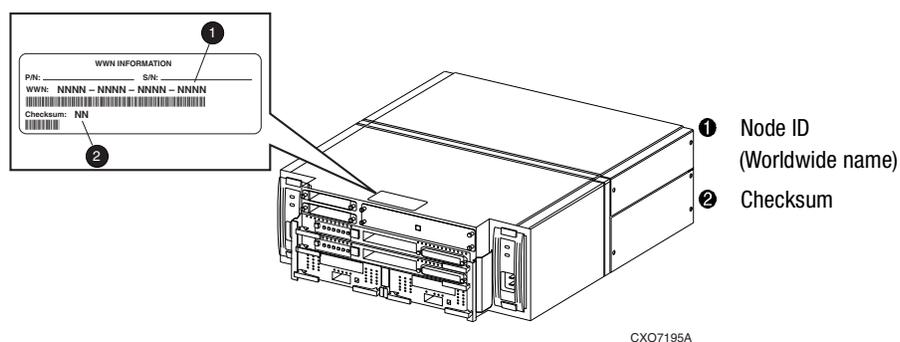


Figure 1-14. Placement of the worldwide name label on the Model 2100 enclosure

For the HSG80

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

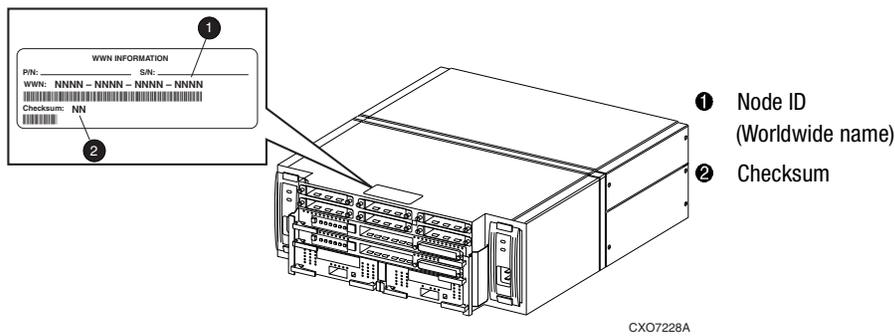


Figure 1-15. Placement of the worldwide name label on the Model 2200 enclosure

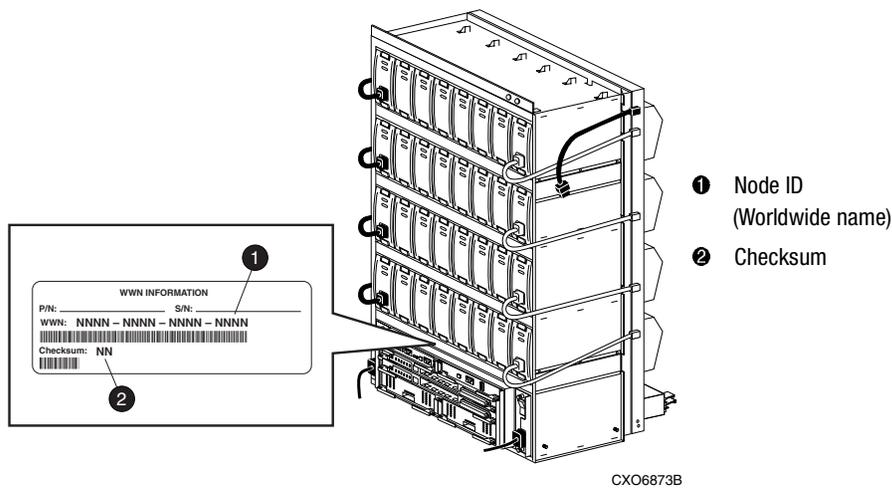


Figure 1-16. Placement of the worldwide name label on the BA370 enclosure



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

Unit 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.

Chapter 2

Planning Storage

This chapter provides information to help you plan the storage configuration of your subsystem. Use the guidelines found in this chapter 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 2–13
- “Choosing a Container Type,” page 2–14
- “Creating a Storageset Profile,” page 2–15
- “Storageset Planning Considerations,” page 2–17
- “Mirrorset Planning Considerations,” page 2–19
- “Partition Planning Considerations,” page 2–24
- “Changing Characteristics through Switches,” page 2–26
- “Storageset and Partition Switches,” page 2–27
- “Initialization Switches,” page 2–28
- “Unit Switches,” page 2–32
- “Storage Maps,” page 2–32

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 2-3.
3. Determine your storage requirements. Use the questions in “Determining Storage Requirements,” page 2-13 to help you.
4. Choose the type of storage container(s) you need to use in your subsystem. See “Choosing a Container Type,” page 2-14 for a comparison and description of each type of storageset.
5. Create a storageset profile (described in “Creating a Storageset Profile,” page 2-15). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:
 - “Storageset Planning Considerations,” page 2-17
 - “Mirrorset Planning Considerations,” page 2-19
 - “Partition Planning Considerations,” page 2-24
 - “Striped Mirrorset Planning Considerations,” page 2-23
6. Decide which switches you will need for your subsystem. General information on switches is detailed in “Storageset and Partition Switches,” page 2-27.
 - Determine what unit switches you want for your units (“Unit Switches,” page 2-32)
 - Determine what initialization switches you want for your planned storage containers (“Initialization Switches,” page 2-28)
7. Create a storage map (“Storage Maps,” page 2-32).
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* for more information.

Configuration Rules

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

- Maximum of 128 LUNs: if Command Console LUN is enabled, the result is 127 visible LUNs and one CCL
- Maximum 512-GB LUN capacity
- Maximum 24 physical devices
- Maximum 8 RAID-5 storage sets
- Maximum 8 partitions per storage set or individual disk
- Maximum 6 physical devices per RAID 1 (mirror set)
- Maximum 14 physical devices per RAID-5 storage set
- Maximum 24 physical devices per RAID 0 (stripe set)
- Maximum 24 physical devices per striped mirror set

Device PTL Addressing Convention

For the HSG60

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 Ultra SCSI 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.

2-4 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

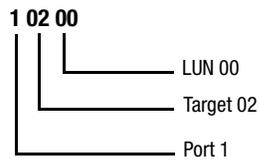


Figure 2-1. PTL naming convention

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

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	not supported	not supported

Model 4214R Disk Enclosure Shelf 2 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	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

When the controller receives an I/O request, it identifies the storage set unit number for the request, then correlates the unit number to the storage set name. From the storage set name, the controller locates the appropriate device(s) for the I/O request (For example, the RAID set “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.

2-6 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

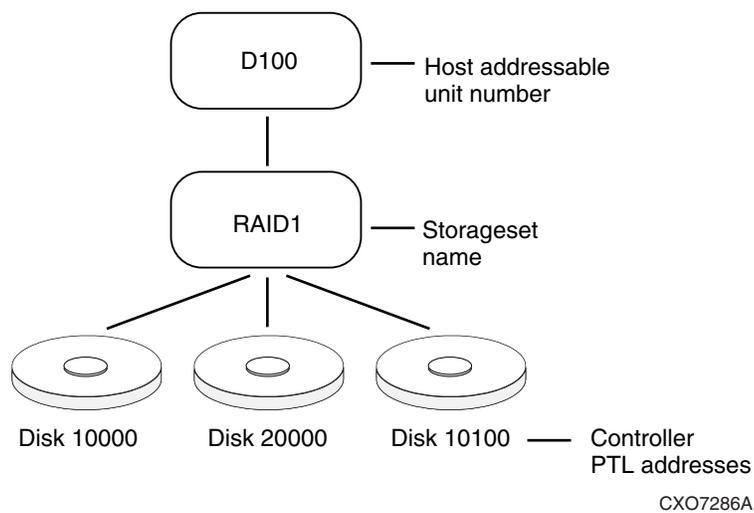


Figure 2-4. Mapping a unit to physical disk drive

For the HSG80

The controller has six SCSI device ports, each of which connects to a SCSI bus. In dual-controller subsystems, these device buses are shared between the two controllers. (The StorageWorks Command Console GUI calls the device ports “channels.”) The standard BA370 enclosure provides a maximum of four SCSI target IDs for each device port. If more target IDs are needed, expansion enclosures can be added to the subsystem.

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 through 6).
- 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 0-3 and 8-15.
- L—Designates the logical unit (LUN) of the device. For disk devices the LUN is always 0.

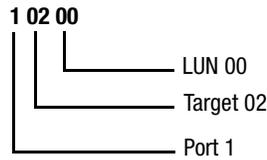
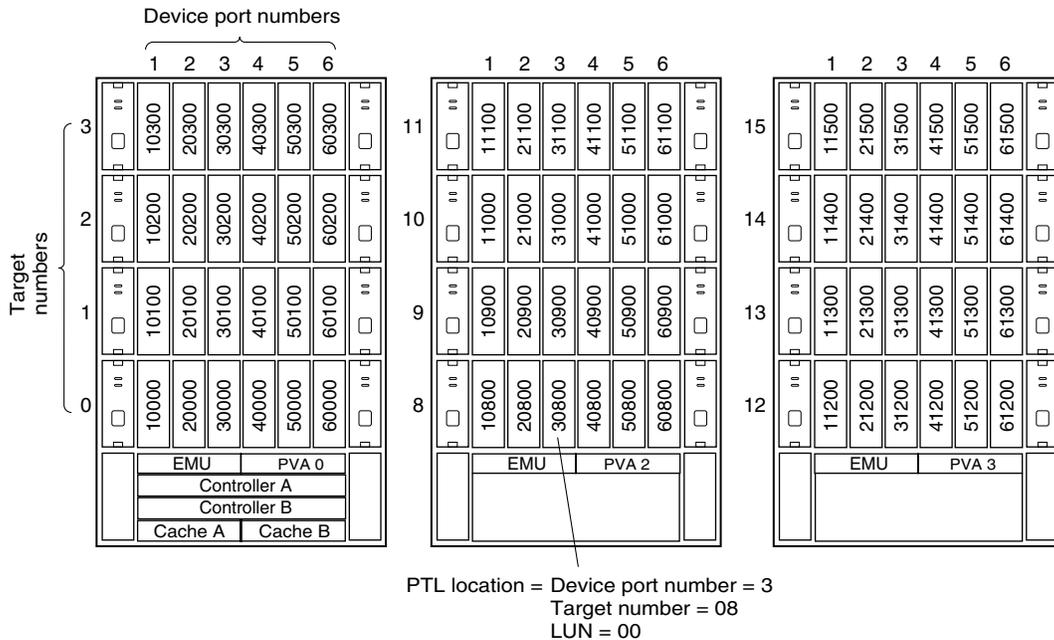


Figure 2-5. PTL naming convention

The controller can either operate with a BA370 enclosure or with a Model 2200 disk enclosure combined with Model 4214R or Model 4254 disk enclosures.

The controller operates with BA370 enclosures that are assigned ID numbers 0, 2, and 3. These ID numbers are set through the PVA module. Enclosure ID number 1, which assigns devices to targets 4 through 7, is not supported. Figure 2-6 shows the addresses for each device in an extended configuration.



CXO5851B

Figure 2-6. PTL addressing in a configuration in a BA370 enclosure

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

- **Model 4214R disk enclosure** - Ultra2 SCSI with 14 drive bays, single-bus I/O module.
- **Model 4254 disk enclosure** - Ultra2 SCSI with 14 drive bays, dual-bus I/O module.
 - NOTE:** The Model 4214R uses the same storage maps as the Model 4314R and the Model 4254 uses the same storage maps as the Model 4354R disk enclosures.
- **Model 4310R disk enclosure** - Ultra3 SCSI with 10 drive bays, single-bus I/O module. Figure 2-7 shows the addresses for each device in a six-shelf single-bus configuration. A maximum of six Model 4310R disk enclosures can be used with each Model 2200 controller enclosure.
 - NOTE:** The storage map for the Model 4310R reflects the disk enclosures physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.
- **Model 4350R disk enclosure** - Ultra3 SCSI with 10 drive bays, dual-bus I/O module. Figure 2-8 shows the addresses for each device in a three-shelf dual-bus configuration. A maximum of three Model 4350R disk enclosures can be used with each Model 2200 controller enclosure.
- **Model 4314R disk enclosure** - Ultra3 SCSI with 14 drive bays, single-bus I/O module. Figure 2-9 shows the addresses for each device in a six-shelf single-bus configuration. A maximum of six Model 4314R disk enclosures can be used with each Model 2200 controller enclosure.
 - NOTE:** The storage map for the Model 4314R reflects the disk enclosures physical location in the rack. Disk enclosures 6, 5, and 4 are stacked above the controller enclosure and disk enclosures 1, 2, and 3 are stacked below the controller enclosure.
- **Model 4354R disk enclosure** - Ultra3 SCSI with 14 drive bays, dual-bus I/O module. Figure 2-10 shows the addresses for each device in a three-shelf dual-bus configuration. A maximum of three Model 4354R disk enclosures can be used with each Model 2200 controller enclosure.

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

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

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

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

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

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

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

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

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

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

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

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

Figure 2-7. PTL addressing in a single-bus configuration, using six Model 4310R enclosures

2-10 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

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

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

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

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

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

Figure 2-8. PTL addressing in a dual-bus configuration, using three Model 4350R enclosures

Model 4314R Disk Enclosure Shelf 6 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800	Disk60900	Disk61000	Disk61100	Disk61200	Disk61300	not supported	not supported

Model 4314R Disk Enclosure Shelf 5 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk50900	Disk51000	Disk51100	Disk51200	Disk51300	not supported	not supported

Model 4314R Disk Enclosure Shelf 4 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300	not supported	not supported

Model 4314R Disk Enclosure Shelf 1 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	not supported	not supported

Model 4314R Disk Enclosure Shelf 2 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	not supported	not supported

Model 4314R Disk Enclosure Shelf 3 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	not supported	not supported

Figure 2-9. PTL addressing in a single-bus configuration, using six Model 4314R enclosures

2-12 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

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

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

	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

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

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

Figure 2-10. PTL addressing in a dual-bus configuration, using three Model 4354R enclosures

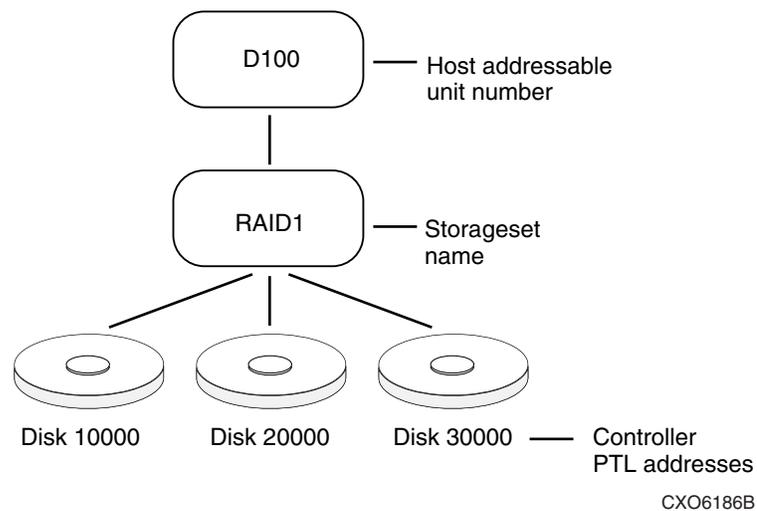


Figure 2-11. Mapping a unit to physical disk drives

Determining Storage Requirements

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

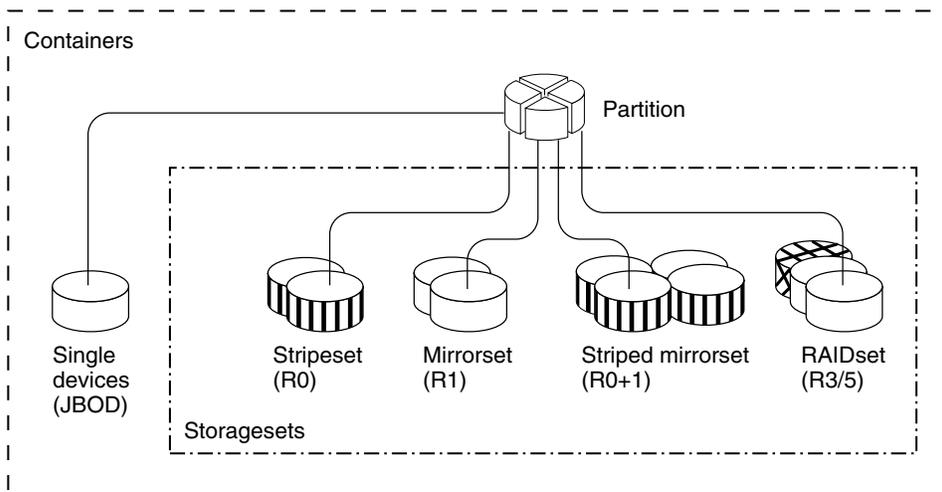
- What applications or user groups will access the subsystem? How much capacity do they need?
- What are the I/O requirements? If an application is data-transfer intensive, what is the required transfer rate? If it is I/O request intensive, what is the required response time? What is the read/write ratio 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, as shown in Figure 2-12. 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 2-12. Container types

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

Table 2-1 A Comparison of Container Types				
Container Name	Relative Availability	Request Rate (Read/Write) I/O per second	Transfer Rate (Read/Write) MB per second	Applications
Independent disk drives (JBOD)	Equal to number of JBOD disk drives	Comparable to single disk drive	Comparable to single disk drive	
Stripeset (RAID 0)	Proportionate to number of disk drives; worse than single disk drive	Excellent if used with large chunk size	Excellent if used with small chunk size	High performance for non-critical data
Mirrorset (RAID 1)	Excellent	Good/Fair	Good/Fair	System drives; critical files
RAIDset (RAID 3/5)	Excellent	Excellent/good	Read: excellent (if used with small chunk sizes) Write: good (if used with small chunk sizes)	High request rates, read-intensive, data lookup
Striped Mirrorset (RAID 0+1)	Excellent	Excellent if used with large chunk size	Excellent if used with small chunk size	Any critical response-time application

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

Creating a Storageset Profile

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

See the example storageset profile shown in Figure 2-13.

2-16 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

Type of Storage set:

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

Storage set Name *R1*

Disk Drives *D10300, D20300, D10400, D20400*

Unit Number *D101*

Partitions:

| Unit # |
|--------|--------|--------|--------|--------|--------|--------|--------|
| % | % | % | % | % | % | % | % |

RAIDset Switches:

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

Mirrorset Switches:

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

Initialize Switches:

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

Unit Switches:

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

Figure 2-13. An example storage set profile

Storageset Planning Considerations

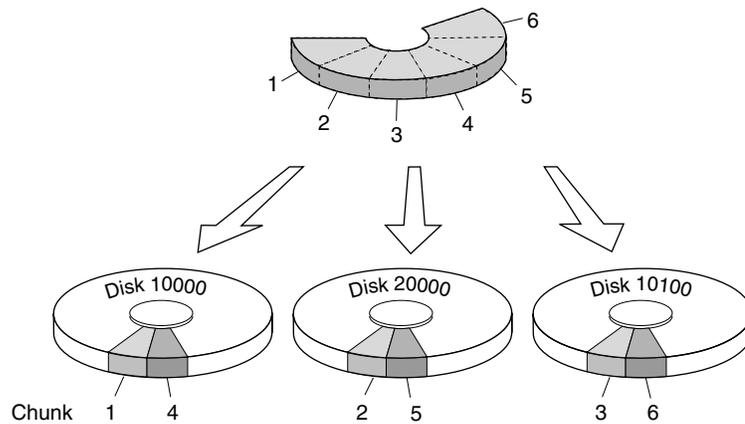
This section contains the guidelines for choosing the storageset type needed for your subsystem:

- Stripeseet Planning Considerations
- Mirrorset Planning Considerations
- Partitio Planning Considerations
- RAIDset Planning Considerations
- Striped Mirrorset Planning Considerations

Stripeseet 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 2-14).



CXO7287A

Figure 2-14. A three-member RAID 0 stripeset (example 1)

2-18 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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 2-28, for information about setting the chunk size). Figure 2-15 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.

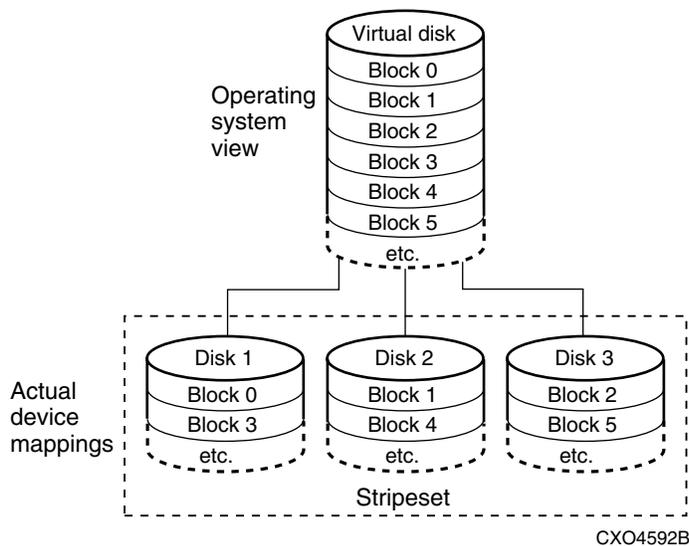


Figure 2-15. A three-member RAID 0 stripeset (example 2)

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.

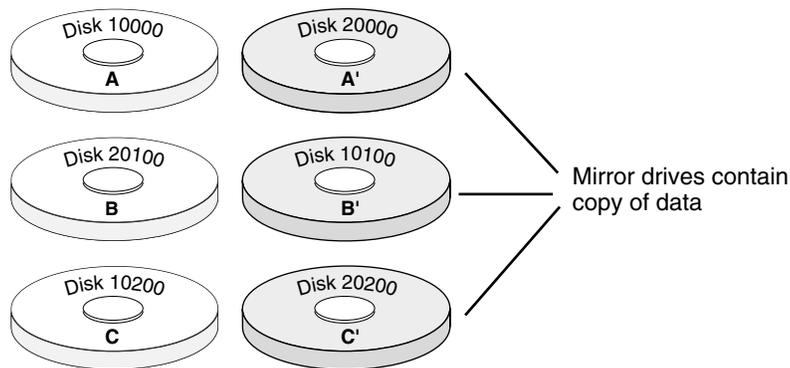
For example, if the mean time between failures (MTBF) for a single disk is 1 hour, then the MTBF for a stripeset that comprises N such disks is 1/N hours. As another example, if the MTBF of a single disk is 150,000 hours (about 17 years), a stripeset comprising four of these disks would only have an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

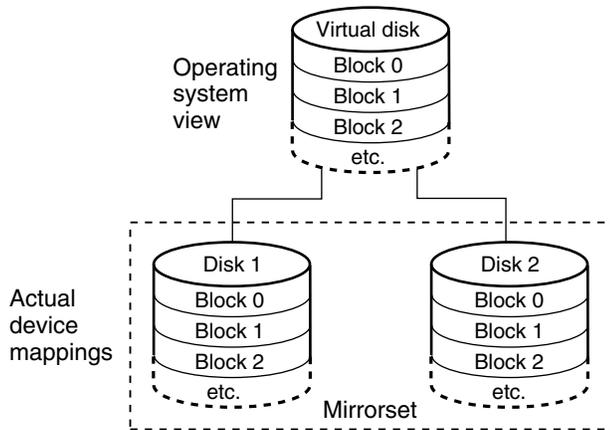
Mirrorset Planning Considerations

Mirrorsets (RAID 1) use redundancy to ensure availability, as illustrated in Figure 2-16. 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 2-17 shows a second example of a Mirrorset.



CXO7288A

Figure 2-16. Mirrorsets maintain two copies of the same data



CXO4594B

Figure 2-17. Mirrorset example 2

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 2-18 shows an example of a RAIDset that uses five members.

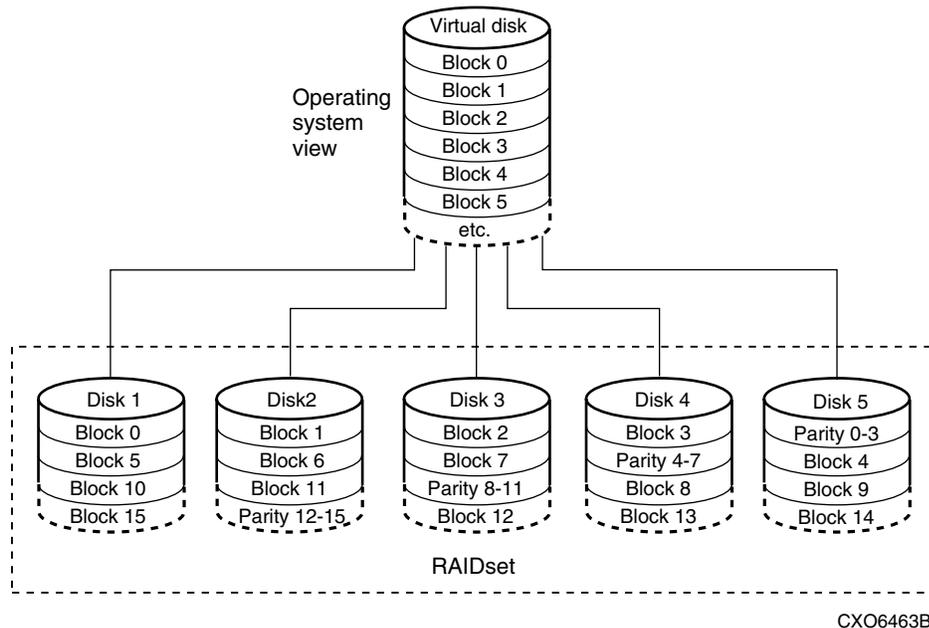


Figure 2-18. A five-member RAIDset using parity

2-22 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

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 2-28, 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 2-19). 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 2-20 shows a second example of a striped mirrorset using six members.

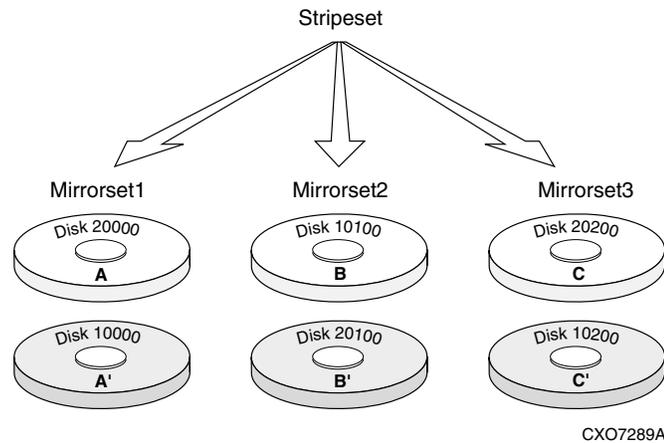


Figure 2-19. Striped mirrorset (example 1)

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.

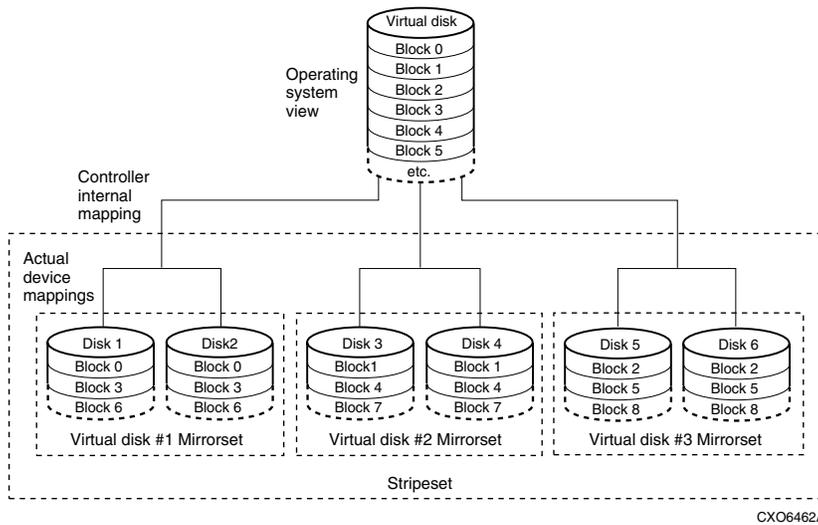


Figure 2-20. Striped mirrorset (example 2)

Plan the mirrorset members, then plan the stripeset that will contain them. Review the recommendations in “Storageset Planning Considerations,” page 2-17, and “Mirrorset Planning Considerations,” page 2-19.

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 2-21 shows the conceptual effects of partitioning a single-disk container.

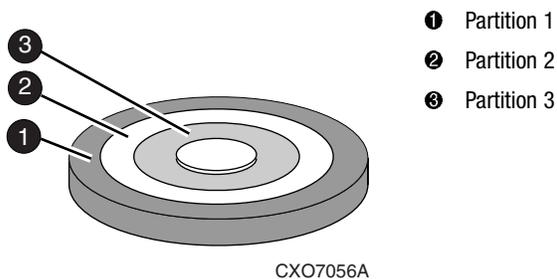


Figure 2-21. One example of a partitioned single-disk unit

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

Defining a Partition

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

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

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

Guidelines for Partitioning Storagesets and Disk Drives

Keep these points in mind as you plan your partitions:

- You can create up to eight partitions per storageset 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.
- Partitions cannot be combined into storagesets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAIDset.
- Just as with storagesets, you do not have to assign unit numbers to partitions until you are ready to use them.
- The CLONE utility cannot be used with partitioned mirrorsets or partitioned stripesets.

Changing Characteristics through Switches

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

- storageset switches
- initialization switches
- unit switches

The following sections:

- describe how to enable/modify switches
- 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/Nodestroy
- 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.

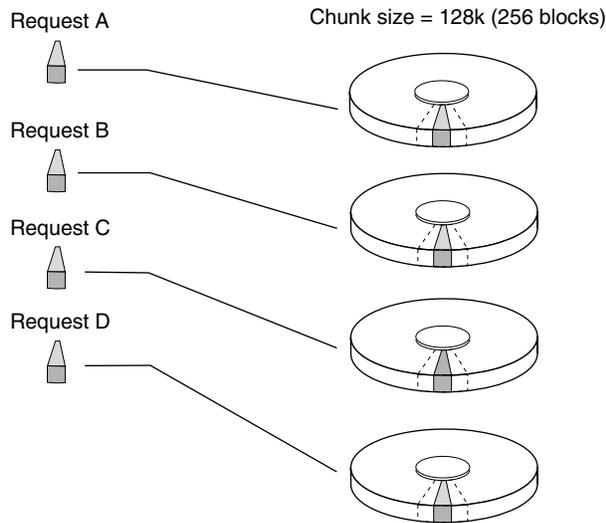
Chunk Size

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 storage set 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 2-22.



CXO-5135A-MC

Figure 2-22. 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 do not know, then you should use a chunk size of 15 times the average transfer request rate.

2-30 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

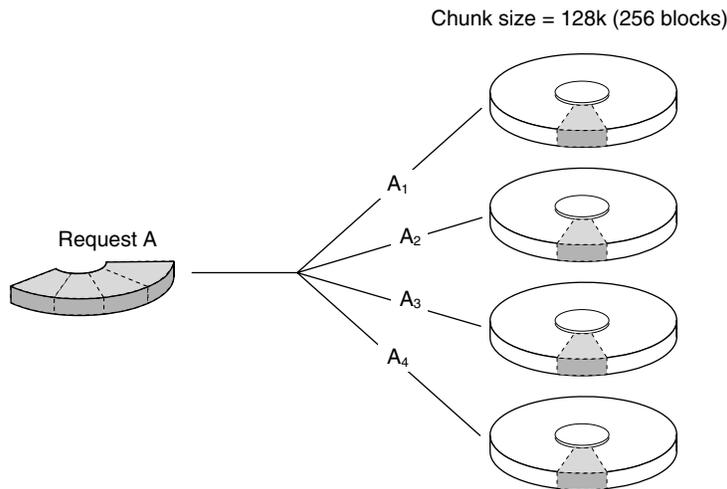
- If you have mostly sequential reads or writes (like those needed to work with large graphic files), then make the chunk size a small number (that is 17 sectors).

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

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

Increasing the Data Transfer Rate

A small chunk size relative to the average request size increases the data transfer rate by allowing multiple disk drives to participate in one I/O request. This concept is shown in Figure 2-23.



CXO-5172A-MC

Figure 2-23. 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 back up 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 1-16
- “Restricting Host Access (Selective Storage Presentation),” page 1-19

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 hard copy 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 for the HSG60

The following example (see Figure 2-24) 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 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D100 R1 DISK10000	D100 R1 DISK10100	D100 R1 DISK10200	D101 S1 M1 DISK10300	D101 S1 M1 DISK10400	D101 S1 M2 DISK10500	D101 S1 M2 DISK10800	D102 M3 DISK10900	D102 M3 DISK11000	D103 S2 DISK11100	D103 S2 DISK11200	D103 S2 DISK11300	not supported	not supported

Figure 2-24. Example storage map-single bus for the HSG60

2-34 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

The following example (see Figure 2-25) 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 2-25. Example storage map-dual bus for the HSG60

Example Storage Maps for the HSG80

Example Storage Map - Model 4310R Disk Enclosure

Figure 2-26 shows an example of four Model 4310R disk enclosures (single-bus I/O).

- Unit D100 is a 4-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, Disk30000, and Disk 40000.
- 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 Disk10100 and Disk20100.
 - M2 is a 2-member mirror set consisting of Disk30100 and Disk40100.
- Unit D102 is a 2-member mirror set named M3. M3 consists of Disk10200 and Disk20200.
- Unit D103 is a 2-member mirror set named M4. M4 consists of Disk30200 and Disk40200.
- Unit D104 is 3-member stripeset named S2. S2 consists of Disk10300, Disk20300, and Disk30300.
- Unit D105 is a single (JBOD) disk named Disk40300.
- Unit D106 is a 3-member RAID 3/5 storage set named R2. R2 consists of Disk10400, Disk20400, and Disk30400.
- Unit D107 is a single (JBOD) disk named Disk40400.
- Unit D108 is a 4-member stripeset named S3. S3 consists of Disk10500, Disk20500, Disk30500, and Disk40500.
- Unit D1 is a 2-member striped mirror set named S4. S4 consists of M4 and M5:
 - M4 is a 2-member mirror set consisting of Disk10800 and Disk20800.
 - M5 is a 2-member mirror set consisting of Disk30800 and Disk40800.
- Unit D2 is a 4-member RAID 3/5 storage set named R3. R3 consists of Disk11000, Disk21000, Disk31000, and Disk41000.
- Unit D3 is a 4-member stripeset named S5. S5 consists of Disk11100, Disk21100, Disk31100, and Disk41100.
- Unit D4 is a 2-member mirror set named M7. M7 consists of Disk11200 and Disk21200.
- Disk31200 and Disk41200 are spare set members.

2-36 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

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

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

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

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

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

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

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

Figure 2-26. Model 4310R disk enclosure - example storage map

Example Storage Map - Model 4350R Disk Enclosure

Figure 2-27 shows an example of three Model 4350R disk enclosures (dual-bus).

- Unit D100 is a 6-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, Disk30000, Disk 40000, Disk50000, and Disk60000.
- Unit D101 is a 6-member RAID 3/5 storage set named R2. R2 consists of Disk10100, Disk20100, Disk30100, Disk 40100, Disk50100, and Disk60100.
- Unit D102 is a 2-member striped mirror set named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirror set consisting of Disk10200 and Disk20200.
 - M2 is a 2-member mirror set consisting of Disk30200 and Disk40200.
- Unit D103 is a 2-member mirror set named M3. M3 consists of Disk50200 and Disk60200.
- Unit D1 is a 4-member stripe set named S2. S2 consists of Disk10300, Disk20300, Disk30300, and Disk40300.
- Unit D2 is a 2-member mirror set named M4. M4 consists of Disk50300 and Disk60300.
- Unit D3 is a 2-member striped mirror set named S3. S3 consists of M5 and M6:
 - M5 is a 2-member mirror set consisting of Disk10400 and Disk20400.
 - M6 is a 2-member mirror set consisting of Disk30400 and Disk40400.
- Unit D4 is a single (JBOD) disk named Disk50400.
- Disk60400 is a spare set member.

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

Bay	SCSI Bus A					SCSI Bus B				
	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
	D100 R1	D101 R2	D102 S1 M1	D1 S2	D3 S3 M5	D100 R1	D101 R2	D102 S1 M1	D1 S2	D3 S3 M5
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400

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

Bay	SCSI Bus A					SCSI Bus B				
	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
	D100 R1	D101 R2	D102 S1 M2	D1 S2	D3 S3 M6	D100 R1	D101 R2	D102 S1 M2	D1 S2	D3 S3 M6
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400

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

Bay	SCSI Bus A					SCSI Bus B				
	1	2	3	4	5	6	7	8	9	10
SCSI ID	00	01	02	03	04	00	01	02	03	04
	D100 R1	D101 R2	D103 M3	D2 M4	D4	D100 R1	D101 R2	D103 M3	D2 M4	spare
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400

Figure 2-27. Model 4350R disk enclosure - example storage map

Example Storage Map - Model 4314R Disk Enclosure

Figure 2-28 shows an example of four Model 4314R disk enclosures (single-bus I/O).

- Unit D100 is a 4-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, Disk30000, and Disk 40000.

- Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirrorset consisting of Disk10100 and Disk20100.
 - M2 is a 2-member mirrorset consisting of Disk30100 and Disk40100.
- Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10200 and Disk20200.
- Unit D103 is a 2-member mirrorset named M4. M4 consists of Disk30200 and Disk40200.
- Unit D104 is a 3-member stripeset named S2. S2 consists of Disk10300, Disk20300, and Disk30300.
- Unit D105 is a single (JBOD) disk named Disk40300.
- Unit D106 is a 3-member RAID 3/5 storageset named R2. R2 consists of Disk10400, Disk20400, and Disk30400.
- Unit D107 is a single (JBOD) disk named Disk40400.
- Unit D108 is a 4-member stripeset named S3. S3 consists of Disk10500, Disk20500, Disk30500, and Disk40500.
- Unit D1 is a 2-member striped mirrorset named S4. S4 consists of M4 and M5:
 - M4 is a 2-member mirrorset consisting of Disk10800 and Disk20800.
 - M5 is a 2-member mirrorset consisting of Disk30800 and Disk40800.
- Unit D2 is a 4-member RAID 3/5 storageset named R3. R3 consists of Disk10900, Disk20900, Disk30900, and Disk40900.
- Unit D3 is a 4-member stripeset named S5. S5 consists of Disk11000, Disk21000, Disk31000, and Disk41000.
- Unit D4 is a 2-member mirrorset named M7. M7 consists of Disk11100 and Disk21100.
- Unit D5 is a 2-member stripeset named S6. S6 consists of Disk31100 and Disk41100.
- Unit D6 is a 4-member RAID 3/5 storageset named R4. R4 consists of Disk11200, Disk21200, Disk31200, and Disk41200.
- Unit D7 is a 2-member mirrorset named M8. M8 consists of Disk11300 and Disk21300.
- Disk31300 and Disk41300 are spareset members.

2-40 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

Model 4314R Disk Enclosure Shelf 4 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
	D100 R1	D101 S1 M2	D103 M4	D105	D107	D108 S3	D1 S4 M6	D2 R3	D3 S5	D5 S6	D6 R4	spare	not supported	not supported
DISK ID	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800	Disk40900	Disk41000	Disk41100	Disk41200	Disk41300		

Model 4314R Disk Enclosure Shelf 1 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
	D100 R1	D101 S1 M1	D102 M3	D104 S2	D106 R2	D108 S3	D1 S4 M5	D2 R3	D3 S5	D4 M7	D6 R4	D7 M8	not supported	not supported
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300		

Model 4314R Disk Enclosure Shelf 2 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
	D100 R1	D101 S1 M1	D102 M3	D104 S2	D106 R2	D108 S3	D1 S4 M5	D2 R3	D3 S5	D4 M7	D6 R4	D7 M8	not supported	not supported
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300		

Model 4314R Disk Enclosure Shelf 3 (single-bus)

Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
	D100 R1	D101 S1 M2	D103 M4	D104 S2	D106 R2	D108 S3	D1 S4 M6	D2 R3	D3 S5	D5 S6	D6 R4	spare	not supported	not supported
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300		

Figure 2-28. Model 4314R disk enclosure - example storage map

Example Storage Map - Model 4354R Disk Enclosure

Figure 2-29 shows an example of three Model 4354R disk enclosures (dual-bus).

- Unit D100 is a 6-member RAID 3/5 storage set named R1. R1 consists of Disk10000, Disk20000, Disk30000, Disk 40000, Disk50000, and Disk60000.
- Unit D101 is a 6-member RAID 3/5 storage set named R2. R2 consists of Disk10100, Disk20100, Disk30100, Disk 40100, Disk50100, and Disk60100.
- Unit D102 is a 2-member striped mirror set named S1. S1 consists of M1 and M2:
 - M1 is a 2-member mirror set consisting of Disk10200 and Disk20200.
 - M2 is a 2-member mirror set consisting of Disk30200 and Disk40200.
- Unit D103 is a 2-member mirror set named M3. M3 consists of Disk50200 and Disk60200.
- Unit D104 is a 2-member striped mirror set named S2. S2 consists of M3 and M4:
 - M3 is a 2-member mirror set consisting of Disk10300 and Disk20300.
 - M4 is a 2-member mirror set consisting of Disk30300 and Disk40300.
- Unit D105 is a 2-member stripeset named S3. S3 consists of Disk50300 and Disk60300.
- Unit D1 is 4-member stripeset named S4. S4 consists of Disk10400, Disk20400, Disk30400, and Disk40400.
- Unit D2 is a 2-member mirror set named M5. M5 consists of Disk50400 and Disk60400.
- Unit D3 is a 2-member striped mirror set named S5. S5 consists of M6 and M7:
 - M6 is a 3-member mirror set consisting of Disk10500, Disk20500, and Disk30500.
 - M7 is a 2-member mirror set consisting of Disk40500 and Disk50500.
- Unit D4 is a single (JBOD) disk named Disk60500.
- Unit D5 is a 4-member RAID 3/5 storage set named R3. R3 consists of Disk10800, Disk20800, Disk30800, Disk and 40800.
- Disk50800 and Disk60800 are spareset members.

2-42 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
	D100 R1	D101 R2	D102 S1 M1	D104 S2 M3	D1 S4	D3 S5 M6	D5 R3	D100 R1	D101 R2	D102 S1 M1	D104 S2 M3	D1 S4	D3 S5 M6	D5 R3
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800

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

	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
	D100 R1	D101 R2	D102 S1 M2	D104 S2 M4	D1 S4	D3 S5 M6	D5 R3	D100 R1	D101 R2	D102 S1 M2	D104 S2 M4	D1 S4	D3 S5 M7	D5 R3
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

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

	SCSI Bus A							SCSI Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
	D100 R1	D101 R2	D103 M3	D105 S3	D2 M5	D3 S5 M7	spare	D100 R1	D112 R2	D103 M3	D105 S3	D2 M5	D4	spare
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Figure 2-29. Model 4354R disk enclosure - example storage map

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 the steps required to prepare an IBM host server, or a High Availability Cluster Multi-Processor (HACMP) server for a StorageWorks HSG60/HSG80 storage subsystem environment.

The following steps are required to prepare your IBM AIX host server for a StorageWorks HSG60/HSG80 controller storage subsystem:

- “HSG60\HSG80 Array Controller Setup,” page 2
- “Host Bus Adapter Installation,” page 3–5
- “Loading StorageWorks Solutions Software Version 8.5c onto Hard Disk,” page 3–8
- “Using the StorageWorks Installation Utility Maintenance Options,” page 3–10

HSG60\HSG80 Array Controller Setup

Before the StorageWorks RAID subsystem can be used by the IBM host server, its HSG60\HSG80 controllers must be properly configured. To configure the controllers, proceed as follows:

1. Establish a serial data connection to the RAID system top controller from your host (reference *Appendix A, Creating Your Storage Configuration with the CLI*).
2. From the CLI prompt, type:

```
SHOW THIS_CONTROLLER FULL
```

The current configuration of your controllers will appear.

Ensure that the HSG60/HSG80 controller(s) have been set correctly for arbitrated loop or fabric. Compare the settings displayed by the CLI with those in Table 3-1 for arbitrated loop and Table 3-2 for fabric.

Arbitrated Loop Setup

For connection to an arbitrated loop, check the configuration (parameter settings) of your controller against those listed in Table 3-1, which are the correct HSG60\HSG80 controller parameter settings for an IBM AIX environment.

Table 3-1 Correct Controller Parameter Settings for Arbitrated Loop

HSG60 Controller Parameter	Correct Parameter Setting
Node (WWN)	nnnn-nnnn-nnnn-nnnn
Command Console LUN	Enabled
Port 1 Topology	LOOP HARD
Port 2 Topology	LOOP HARD
Port 1 ALPA	nn (D5 is default)
Port 2 ALPA	nn (D6 is default)
SCSI Version	SCSI-2
New Connection Operating System Type	WINNT

If the HSG60/HSG80 controller parameter settings displayed do not match those in the recommended setting column in Table 3-1, the following CLI commands can be used to alter their values.

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

Initial setting of the Port Topology:

```
HSG >SET OTHER_CONTROLLER PORT_1_TOPOLOGY=OFFLINE
```

```
HSG >SET OTHER_CONTROLLER PORT_2_TOPOLOGY=OFFLINE
```

Setting the Port Topology:

```
HSG >SET OTHER_CONTROLLER PORT_1_TOPOLOGY=LOOP_HARD
```

```
HSG >SET OTHER_CONTROLLER PORT_2_TOPOLOGY=LOOP_HARD
```

Setting the ALPA:

```
HSG >SET THIS_CONTROLLER PORT_1_AL_PA=D5
```

```
HSG >SET THIS_CONTROLLER PORT_2_AL_PA=D6
```

Setting the SCSI Version:

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

If the StorageWorks Command Console (SWCC) will be used, the HSG60/HSG80 Command Console LUN must be enabled.

```
HSG>SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

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

```
HSG>RESTART OTHER_CONTROLLER
```

```
HSG>RESTART THIS_CONTROLLER
```

Once both controllers have been restarted, the connection type must be set. See “New Connection Operating System,” page 3-10 for more information.

Fabric Setup

Ensure that the HSG60/HSG80 controller(s) have been set correctly for fabric. Compare the settings displayed by the CLI with those in Table 3-2. The current settings can be obtained with the command:

```
HSG>SHOW THIS_CONTROLLER
```

Table 3-2 HSG60/HSG80 Controller Parameters and Recommended Settings for Fabric

Controller Parameter	Recommended Setting
Node (WWN)	nnnn-nnnn-nnnn-nnnn
Command Console LUN	Enabled
Port 1 Topology	FABRIC
Port 2 Topology	FABRIC
SCSI Version	SCSI-2
New Connection Operating System Type	WINNT

If the HSG60/HSG80 controller parameter settings displayed do not match those in the recommended setting column in Table 3-2, the following CLI commands can be used to alter their values.

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

Initial setting of the Port Topology:

```
HSG >SET THIS_CONTROLLER PORT_1_TOPOLOGY=OFFLINE
```

```
HSG >SET THIS_CONTROLLER PORT_2_TOPOLOGY=OFFLINE
```

Setting the Port Topology:

```
HSG >SET OTHER_CONTROLLER PORT_1_TOPOLOGY=FABRIC
```

```
HSG >SET OTHER_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

Setting the SCSI Version:

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

If the StorageWorks Command Console (SWCC) will be used, the HSG60/HSG80 Command Console LUN must be enabled.

```
HSG>SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
```

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

```
HSG>RESTART OTHER_CONTROLLER
```

```
HSG>RESTART THIS_CONTROLLER
```

Once both controllers have been restarted, the connection type must be set. See “New Connection Operating System,” page 3-10 for more information.

Host Bus Adapter Installation

The Host Bus Adapter (HBA) installation is accomplished by first installing the HBA driver.

NOTE: If you are installing the RAID subsystem on an HACMP cluster, you will need to install the HBA and its driver into each IBM host server in the cluster.

HBA Driver Installation

The HBA driver is contained on the CD-ROM in your StorageWorks Solutions Software kit. The HBA driver was designed to be installed from writable media and will need to be copied to a temporary hard disk location.

If this is an existing configuration, see “HBA Driver Removal,” page 3-6 or “Upgrading the StorageWorks Solution Software,” page 3-11, depending on your requirements.

To install the HBA Driver in a new configuration, perform the following steps:

1. Insert the StorageWorks Solutions Software distribution CD-ROM into the CD-ROM drive.

NOTE: Be sure you have a mount point, usually a directory in root, such as /mnt, and you know the device number of your CD-ROM drive, such as /dev/cd0.

2. Mount the CD-ROM, type:

```
# mount -vcdafs -r /dev/cd0 /mnt
```

NOTE: Your CD-ROM device may not be the same as specified in the example command.

3. Change directories to the location of the driver, type:

```
# cd /mnt/fcdriver/pci
```

4. Create a temporary directory to load the driver, type:

```
# mkdir /tmp/fcdriver
```

5. Copy the driver from the CD-ROM to the temporary directory, type:

```
# cp PC1000.image /tmp/fcdriver
```

6. Change directories, type:

```
# cd /tmp/fcdriver
```

7. Install the driver, type:

```
# installp -acd PC1000.image all
```

3-6 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

HBA Driver Removal

If this HBA driver (PC1000), or a previous version of the HBA driver (q2100) needs to be removed from the system, use the following procedure. The q2100 driver is used as an example.

1. Back up the data
2. Delete hdisks (HSG60/HSG80)
3. Delete the Cambex Adapter (scsix)
4. Uninstall the previous driver, using the command:
`# installp -u q2100.driver.obj`
5. Remove the device configuration data from the kernel, type:
`# bosboot -ad /dev/ipldevice`

Hardware Installation

Power down the host server and install the adapter in a free PCI slot, but preferably not PCI0 (system cards). If a 64-bit slot is not available, a 32-bit slot may be used. Make sure that airflow is not blocked by cables. Observe proper static precautions. Detailed procedures for opening the case may be found in the Operator Guide for the specific RS/6000 model.

Do not attach the cables from the host adapter to the fibre channel devices at this time.

Power up the host server.

Configuration

During power-up initialization, AIX will find and configure the new adapter(s) using the driver's configuration method. To repeat this process and capture detailed output to a text file, use the command:

```
# cfgmgr -v > /tmp/cfgmgr.log 2>&1
```

Find out the names of the installed adapters with the command:

```
# lsdev -Cc adapter
```

An example of this command is as follows:

```
# lsdev -Cc adapter
```

fda0	Available 01-A0 Standard I/O Diskette Adapter
sa0	Available 01-C0 Standard I/O Serial Port 1
sa1	Available 01-D0 Standard I/O Serial Port 2
sioka0	Available 01-E0 Keyboard Adapter
sioma0	Available 01-F0 Mouse Adapter
scsi0	Available 04-C0 Wide SCSI I/O Controller
scsi1	Available 04-05 Wide SCSI I/O Controller
scsi2	Available 04-06 Wide/Fast-20 SCSI I/O Controller
scsi3	Available 04-07 Wide/Fast-20 SCSI I/O Controller
iga0	Available 04-01 GXT110P Graphics Adapter
ppa0	Available 01-B0 Standard I/O Parallel Port Adapter
supp0	Available 01-G0 Mini-Service Processor
ent0	Available 04-04 IBM PCI Ethernet Adapter (22100020)
scsi4	Available 04-02 Cambex Fibre Channel I/O Controller

For arbitrated loop, the ID of each device on the loop must be unique. The adapter ID is ϕ (zero) by default, and may be examined after it is installed with the `lsattr` command. If the adapter name found by the `lsdev` command is `scsi4`, the command would be:

```
lsattr -E -l scsi4
```

An example of the `lsattr` command is as follows:

```
# lsattr -El scsi4
```

bus_io_addr	0x01000200 Mapped address of board I/O registers	False
bus_intr_lvl	4 Bus interrupt level	False
intr_priority	3 Interrupt priority	False
id	0 Adapter card loop ID	True
exec_throttle	1 Max commands per port	True
max_iocb	1 Max buffers per port	True
retry_delay	3 Delay amount after bus reset (s)	True
login_on_lip	yes Login all devices after LIP	True
full_duplex	no Use open full duplex (OPNyx)	True

3-8 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

topology_list		Prioritized topology list	True
P2P_LOOP			
frame_size	1024	Max payload frame length	True
nfinity	300	Command timeout infinity (s)	True
dma_window	256M	DMA window size for 64 bit machine	True

The ID may be changed in the database to any non-conflicting value in the range 0 to 126 with a command such as:

```
# chdev -l scsi4 -a id=6 -P
```

Connection

When the HBA has been configured, connect the fibre connection from the HBA to the Hub that the RAID Array is connected to.

Loading StorageWorks Solutions Software Version 8.5c onto Hard Disk

The StorageWorks Solutions Software consists of the RAID Array software installation script, the AIX support files, and the Command Console (CC) Agent. Installation and operation of the software will require temporary and permanent space on your hard disk. You will need approximately 4 Mbytes of permanent space on your hard disk to use the StorageWorks Solutions Software and an additional 2 Mbytes of temporary space on your hard disk for installation purposes.

The StorageWorks Solutions Software consists of:

- An installation and configuration script
- ODM database files
- ODM catalog files

Loading the StorageWorks Solutions Software requires the following steps:

- Mounting the CD-ROM that contains the software
- Copying the files from the file archive to a temporary storage area
- Executing the installation and configuration script

Copying the File Archive from CD-ROM to Temporary Storage

Before the installation script or upgrade script can be run, the file archive that contains the StorageWorks Solutions Software will need to be copied from the CD-ROM to a temporary directory. The scripts can then be executed. You will need approximately 2 Mbytes of temporary space to load the file archive from CD-ROM.

NOTE: If you are using HACMP cluster software, you will need to perform this procedure on each node in the cluster.

1. Insert the StorageWorks Solutions Software distribution CD-ROM into the drive.
2. Login as root (superuser).
3. Change the directory to a temporary storage area (/tmp). For example, type:

```
# cd /tmp
```

4. Create a directory into which the file archive can be loaded:

```
# mkdir HSG
```

5. Change to that directory:

```
# cd HSG
```

6. Mount the CD-ROM:

```
# mount -vcdvfs -r /dev/cd0 /mnt
```

NOTE: The name of /dev/cd0 may be different on your system

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

```
# tar xvf /mnt/swcc/agent/stm_aix.tar
```

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

8. If you are upgrading, proceed to “Upgrading the StorageWorks Solution Software,” page 3-11

Executing the Installation and Configuration Script

Perform the following procedure to install the StorageWorks Solutions Software version 8.5c onto your system from path /tmp/HSG:

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

```
# ./stgws_aix.sh
```

3-10 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

The installation options menu appears.

2. Type 1 to install the StorageWorks Software. The Copyright Notice will be displayed:
Copyright Compaq Computer Corp. 2000. All rights reserved.
Restricted Rights: Use, duplication by the U.S. Government is subject to restriction as set forth in subparagraph (c) (1) (ii) of DFARS252.227-7013, or in FAR 52.227-19, or in FAR 52.227-14 Alt. III, as applicable
Portions of this software are proprietary to and embody the confidential technology of Compaq Computer Corporation. Possession, use, or copying of this software and media is authorized only pursuant to a valid written license from Compaq or an authorized sublicensor.

The installation of StorageWorks software will begin automatically.

New Connection Operating System

Determine if the operating system set for the online controller connection is WINNT. This is done with the CLI using (either HSG60 or HSG80 can be used, depending on which controller you have):

```
HSG60> SHOW CONNECTIONS
```

Although an IBM operating system entry exists in the CLI and SWCC, for the Fibre Channel connection, use the default setting WINNT for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover. To use the default setting, type:

```
HSG60> SET !NEWCONNxx OPERATING_SYSTEM=WINNT
```

where: xx is the online connection id produced in the SHOW CONNECTIONS command above.

NOTE: Further information on command syntax can be found in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

Using the StorageWorks Installation Utility Maintenance Options

The StorageWorks Solutions Software Installation Utility offers the following maintenance options:

- Upgrading StorageWorks Solutions Software
- Uninstalling StorageWorks Solutions Software

Upgrading the StorageWorks Solution Software

The StorageWorks Solution Software Kit contains a method to upgrade to the latest StorageWorks software and update AIX ODM entries. This upgrade is intended for users who already have HSG60 or HSG80 software installed.

Before you begin the upgrade, it is highly recommended that you quiesce all I/O to your StorageWorks storage system. The upgrade is subdivided into four operations; HBA driver upgrade, updating ODM entries, Agent upgrade and importing volume groups.

HBA Driver Upgrade

To upgrade the HBA driver, use the following procedure:

1. Write down your configuration using one of the configuration tables provided in Appendix A.
2. Unmount your mounted filesystems:
 - a. Type `smi t`
 - b. Go to System Storage Management, File Systems, Unmount a file system
 - c. Select the file system name that you choose to unmount.
 - d. Press **Enter**.
3. Deactivate volume groups.
 - a. Type `smi t`
 - b. Go to System Storage Management, Logical Volume Manager, Volume Groups, Deactivate a volume group.
 - c. Select the volume groups you want to deactivate.
 - d. Press **Enter**.
4. Export volume groups.
 - a. Type `smi t`
 - b. Go to System Storage Management, Logical Volume Manager, Volume Groups, Export a volume group.
 - c. Select the volume group that you want to export
 - d. Press **Enter**.
5. Stop the Agent if it is running.

3-12 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

- a. Run the script.
./stgwks_aix.sh
 - b. Select option 4
 - c. Select option 3
 - d. Type **y** for yes.
6. Remove the hdisks that you have configured.
rmdev -dl hdiskx
where x is the hdisk number.
 7. Remove the adapter device file.
rmdev -dl scsix
where x is the scsi number that is associated with the cambex fibre channel adapter.
 8. Uninstall the q2100 driver.
installp -u q2100.driver.obj
 9. Bosboot (regen kernel)
bosboot -ad /dev/ipldevice
 10. Install new driver.
 - a. Login as root.
 - b. cd /(driver location)
 - c. Install the PC1000 driver
installp -acd PC1000.image all
 11. Reboot system.
reboot
 12. Disconnect fibre cable.
Pull cable from hub or switch.
 13. Run cfmgr
cfmgr -v
 14. See if adapter is there.
lsdev -Cc adapter
Your output should appear similar to the following:
scsi2 Available 04-02 Cambex Fibre Channel I/O Controller

15. Connect fibre cable.

Put fibre cable back in hub.

16. Run `cfgmgr`

```
# cfgmgr -v
```

17. Check to see if disks are there.

```
# lsdev -Cc disk
```

Your output should appear similar to the following:

```
hdisk6 Available 04-02-00-0,0 DEC HSG80 RAID Array
```

If you see "OTHER SCSI" in place of "DEC HSG80 RAID Array", see "Importing Volume Groups," page 3-17, for more information.

Updating ODM Entries

The StorageWorks upgrade script will replace the current HSG60/HSG80 ODM files installed on your system. The ODM upgrade should not affect your configuration.

NOTE: If you are using HACMP cluster software, you will have to perform this procedure on each node in the cluster.

To update ODM entries, use the following procedure:

1. Type the following command:

```
# cd /tmp/HSG (or to the directory created in Section "Copying the File Archive from CD-ROM  
to Temporary Storage")  
# ./stgws_aix.sh
```

The script executes and displays the Installation and Configuration Menu.

2. Type 2 and the StorageWorks Maintenance Menu is displayed.

3. Type 1 to upgrade the software. The copyright notice and the following information will be displayed (HSG60 or HSG80 will appear, depending on you controller):

3-14 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

```
Copyright Compaq Computer Corp. 2000. All rights reserved.

Restricted Rights: Use, duplication by the U.S. Government is subject to
restriction as set forth in subparagraph (c) (1) (ii) of DFARS 252.227-7013,
or in FAR 52.227-19, or in FAR 52.227-14 Alt. III, as applicable

Portions of this software are proprietary to and embody the confidential tech-
nology of Compaq Computer Corporation. Possession, use, or copying of
this software and media is authorized only pursuant to a valid written license
from Compaq or an authorized sublicensor.

The purpose of this script is to upgrade existing HSG60 Storage Works soft-
ware to the current version.

Checking User Privileges...OK
Checking AIX Revision Level...OK
Checking that Cambex FC Adapter is installed and correctly configured...OK
Checking for ODM add files...OK
Checking for ODM catalog files...OK
Adding HSG60 to ODM Database...OK
Moving New ODM Catalogs into place...OK
Copying the new installation files to /usr/stgwks2 directory...OK
Configuring system...OK
Saving base configuration...OK
Issuing bosboot command, please standby...OK

Please reboot the system at this time to allow configuration changes to take
effect.

RAID Array Upgrade Complete

-- Hit RETURN to continue --
```

```
StorageWorks Maintenance Menu
-----
1) Upgrade StorageWorks software
2) Uninstall StorageWorks software
Q) Quit
Enter Selection: q
```

4. Type Q to return to the main menu.
5. Reboot your RS/6000 System.

NOTE: If you are using HACMP cluster software, you will need to perform this procedure on each node in the cluster.

Agent Upgrade

To upgrade the Agent (v2.x Agents only), use the following procedure:

1. Copy the /usr/stgwks2/steam/etc directory contents into a temporary area. These files are what the Agent uses to poll the Client as well as send out messages about the status of your array.
 - a. Go to the area where your Agent is installed, it should be /usr/stgwks2, the command is:

```
# cd /usr/stgwks2
```
 - b. Make a temporary storage directory to store your backed up configuration files, the command is:

```
# mkdir temp
```
 - c. Go to that directory:

```
# cd temp
```
 - d. Copy the contents of the /usr/stgwks2/steam/etc directory to the temp area:

```
# cp ../steam/etc/* .
```
2. Once you have these files copied, you need to uninstall the Agent.
 - a. Go to the area where you installed your Agent (/usr/stgwks2):

```
# cd /usr/stgwks2
```
 - b. Run the StorageWorks script:

```
# ./stgwks_aix.sh
```
 - c. Select option 4: Agent Maintenance
 - d. Select option 5 from the next menu: Uninstall Agent
 - e. Type **yes** to the question "Are you sure you want to continue? [yes/no]"
3. Next install your new Agent like you would during a normal Agent install
 - a. Go to where you unpacked your .tar file:

```
# cd /tmp/HSG80 (or location you picked)
```
 - b. Run the StorageWorks script:

```
# ./stgwks_aix.sh
```
 - c. Select option 3: Install Agent

3-16 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

- d. Follow the steps as you would a normal install. You will be required to go and reset your Agent password in this step. Proceed until you reach the place where you add a Client. At this point you would hit ctrl-c to exit from the script.
4. Copy your configuration files back into their proper locations.
 - a. Go to the /usr/stgwks2/steam/etc area:

```
# cd /usr/stgwks2/steam/etc
```
 - b. Copy those files into place:

```
# cp ../../temp/*.ini .
```
5. Verify the installation by looking at the storage configuration.
 - a. Go to your Agent installation directory:

```
# cd /usr/stgwks2/
```
 - b. Run the StorageWorks script:

```
# ./stgwks_aix.sh
```
 - c. Pick option 4: Agent Maintenance
 - d. Look at options 9, 13, and 18 to make sure they are the same as before the upgrade.
6. Restart the Agent.
 - a. Go to your Agent installation directory:

```
# cd /usr/stgwks2/
```
 - b. Run the StorageWorks script:

```
# ./stgwks_aix.sh
```
 - c. Select option 4: Agent Maintenance
 - d. Select option 3: Start/Stop Agent
 - e. Press **Enter** to accept default yes to starting the Agent
7. Delete the temp directory created in the beginning of the upgrade.
 - a. Go to /usr/stgwks2

```
# cd /usr/stgwks2
```
 - b. Remove contents of temp:

```
# rm temp/*
```
 - c. Remove temp:

```
# rmdir temp
```

Importing Volume Groups

To import pre-existing volume groups, use the following procedure:

1. Import volume group.
 - a. Type `smi t`
 - b. Go to System Storage Management, Logical Volume Manager, Volume Groups, Import a volume group.
 - c. Type volume group name and hdisk that it was previously on.
 - d. Press **Enter**.
2. Activate volume group.
 - a. Type `smi t`
 - b. Go to System Storage Management, Logical Volume Manager, Volume Groups, Activate a volume group.
 - c. Type in the volume group name that you previously imported.
 - d. Press **Enter**.
3. Mount file system.
 - a. Type `smi t`
 - b. Go to System Storage Management, Filesystems, Mount a file system.
 - c. Type the name, where you want to mount it, and the type of fs.
 - d. Press **Enter**.

Host Device Initialization

Once the StorageWorks Solutions Software has been installed, logical units may be created on the RAID Array and configured into the system. These logical units can be created using the StorageWorks Command Console (SWCC) Client or the HSG60/HSG80 Command Line Interpreter (CLI). In order to utilize SWCC Client to create logical units, the SWCC Agent and Client must first be installed on the AIX and Windows Systems respectively. Refer to *Chapter 4, Installing and Configuring the HS-Series Agent* for details.

New logical units created on the RAID Array must be recognized by the operating system before they can be utilized. The following steps outline the process:

1. Using the CLI or the StorageWorks Command Console Client, create the logical units on the RAID Array.

3-18 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

2. As the superuser (root), run the command:

```
# cfgmgr -v
```

This command will identify and configure all of the new RAID Array devices in the operating system.

3. Verify that the host has located and configured the new RAID Array devices in the operating system using the command:

```
# lsdev -Cc disk
```

The output of the command looks like (HSG60 or HSG80 will appear on disks 2, 3 and 4, depending on you controller):

```
hdisk0 Available 10-58-00-3, 0 Other SCSI Disk Drive
hdisk1 Available 10-58-00-4, 0 Other SCSI Disk Drive
hdisk2 Available 20-70-00-9, 0 DEC HSG60 RAID Array
hdisk3 Available 20-70-00-10, 0 DEC HSG60 Command Console LUN
hdisk4 Available 20-70-00-10, 1 DEC HSG60 RAID Array
```

4. If the devices were not recognized, re-check the status of the RAID Controller, logical units, and cable connections.
5. You can now use the System Management Interface Tool (SMIT) to make volume groups, logical volumes and file systems using these logical units.

NOTE: If you are using HACMP cluster software, you need only perform this procedure once and it will apply for the entire cluster. However, shared storage must have the same "hdisk" number on each node in the cluster and shared volume groups, logical volumes and the file systems must be organized as elements of shared resource groups for the defined cluster(s). See *High Availability or Multi-Processing for AIX Installation Guide, Version 4.3*.

Uninstalling StorageWorks Solutions Software Version 8.5

To uninstall all of the StorageWorks RAID Array software, follow these steps:

NOTE: If you are using HACMP cluster software, you will have to perform this procedure on each node in the cluster.

1. Type the following commands:

```
# cd /usr/stgwks2
# ./stgwks_aix.sh
```

The script executes and displays the Installation and Configuration menu.

2. Select 2 to bring up the StorageWorks Maintenance menu.
3. Select 2 (Uninstall StorageWorks software).

If you respond (Y) YES, the system starts the deletion process.

4. Type **Q** to return to main menu.

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
- “Before You Install the Agent,” page 4-5
- “Installing and Configuring the Agent,” page 4-5
- “Running the Agent,” page 4-9
- “Reconfiguring the Agent,” page 4-9
- “Removing the Agent,” page 4-9
- “High Availability Cluster Multi-Processing (HACMP),” page 4-10

Why Use StorageWorks Command Console?

StorageWorks Command Console (SWCC) lets you monitor and configure the storage connected to the HSG60/HSG80 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 5.

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.

4-2 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

To receive information about the devices connected to your HSG60/HSG80 controller over a TCP/IP network, you must install the Agent on a computer that is connected to a controller.

The Agent can also be used as a standalone application without Client. In this mode, which is referred to as Agent only, Agent monitors the status of the subsystem and provides local and remote notification in the event of a failure. A subsystem is comprised of the HSG60/HSG80 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

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

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

Table 4-2 provides an overview of the installation.

Table 4-2 Installation and Configuration Overview	
Step	Procedure
1	Verify that your hardware has been set up correctly. See the previous chapters in this Guide.
2	Verify that you have a network connection for the Client and Agent systems. See “About the Network Connection for the Agent.”
3	Verify that there is a LUN to communicate through. This can be either the CCL or a LUN that was created with the CLI. See “The Command Console LUN” in Chapter 1.
4	Install the Agent (TCP/IP network connections) on a system connected to the HSG60/HSG80 controller. See “Installing and Configuring the Agent.”
5	Add the name of the Client system to the Agent’s list of Client system entries (TCP/IP network connections). This can be done during installation or when reconfiguring the Agent.
6	Install the Client software on Windows 2000 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. User Guide</i> .

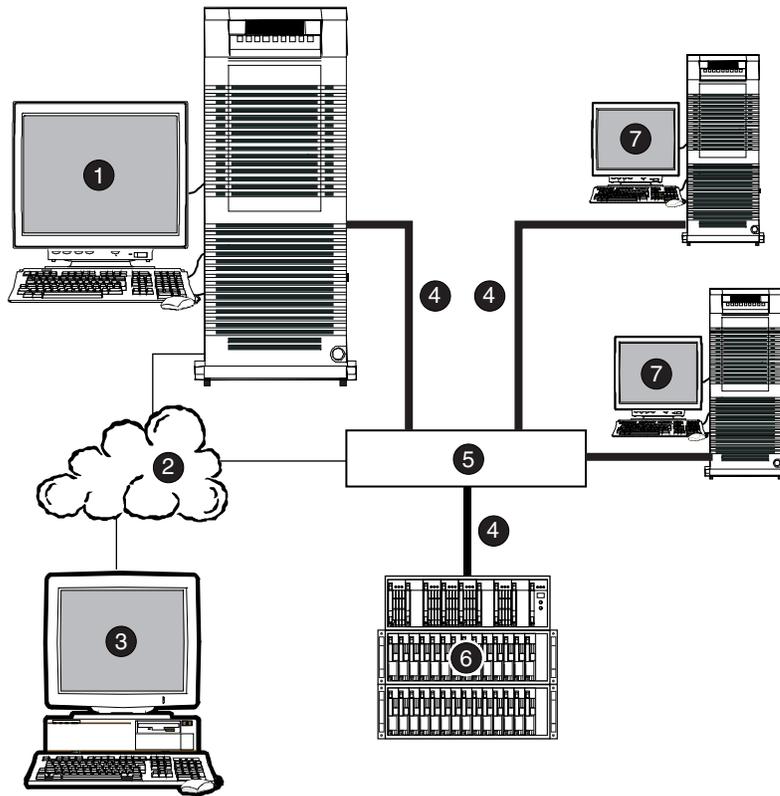
About the Network Connection for the Agent

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

By using a network connection, you can configure and monitor your subsystem from anywhere on your LAN. If you 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).

4-4 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide



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/HSG80 controller and its device subsystem |
| 3 Client system (has the Client software) | 7 Servers |
| 4 Fibre Channel cable | |

Before You Install the Agent

The Agent requires the minimum Host System resources shown in Table 4-3. The program is designed to operate with the Client Version 2.3 for Windows 2000 and Windows NT.

Table 4-3 Minimum IBM-AIX Host Requirements

Host Feature	Requirements
Architecture	IBM RS/6000 2MB free space in /tmp directory, and 4 MB free space for Agent installation directory
Operating System	IBM AIX V4.2 or V4.3, Version 4.3.3 for HACMP
Controller Compatibility	StorageWorks HSG60/HSG80 Controller running operating firmware V8.5

Installing and Configuring the Agent

This installation assumes that you copied the installation and configuration script to a temporary directory.

You must gather some information prior to starting the installation. Entries listed in Table 4-4 appearing with an asterisk specify the information you need prior to starting the installation.

The Agent, once installed and configured properly will automatically start when you install it and will restart automatically if it stops for any reason.

This section provides information about two types of installations:

- Installing the Agent for the First Time
- Upgrading the Agent

Installing the Agent for the First Time

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

```
# cd /tmp/HSG
```

3. Enter the following to start the StorageWorks installation and configuration.

```
# ./stgwks_aix.sh
```

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

4. Select option **3** to install the Agent.
5. Follow the text-based prompts and enter the appropriate information. Table 4-4 guides you in responding to many of the prompts.
6. When you are asked to start the Agent, select option **Y**. The Agent's default is to start automatically after a system reboot or if it stops for any reason.
7. Remove the temporary directory if desired. For example, to remove the /tmp/HSG directory, enter the following:

```
# rm -r /tmp/HSG
```

Table 4-4 Information Required for Install Program Prompts

Information Required to Complete Installation	How to Respond
Target directory for installation	/usr/stgwks2 is the Compaq recommended directory for StorageWorks. It contains the software you copied in Chapter 3. You need a directory with a minimum of 4 MB of space available. Enter the name of a directory and the install program automatically creates it if it doesn't already exist.
*Port that the Command Console Client uses to communicate with the Command Console Agent for RAID subsystem configuration and maintenance.	*Default is 4998. If the default port is already used, press N and specify a new port.
*Port number that the Command Console Agent uses to send notifications to the Command Console Client.	*Default is 4999. If the default port is already used, press N and specify a new port.

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

Information Required to Complete Installation	How to Respond
The directory which contains the device special files.	AIX device special files are located in the /dev directory.
A password to gain access to the Agent to configure the RAID subsystem with the Client.	Enter a 4-16 character password. The install program does not echo what you enter.
*The name of CC Agent manufacturer and the Enterprise object identifier for use with its MIB, so that it can operate with other SNMP applications.	*Enter Compaq Computer Corporation as the CC Agent manufacturer.
*The Enterprise Object Identifier (OID) is a vendor specific value used to identify SNMP messages.	*Enter: 36.2.15.21.
*Names of the Host Systems on which you will run the Client.	*Enter the network names for the hosts from which you will run the Client.
<p>Level of access a Client has for communicating with the storage systems.</p> <p>The possible options are:</p> <p>0 = Navigation Tree monitoring only</p> <p>1 = Navigation Tree and Storage Window monitoring only</p> <p>2 = Navigation Tree monitoring, Storage Window monitoring and configuration</p>	You can give a Client configuration level access only if you specified an access password.
<p>How do you want a CC Agent server to notify a Client when an error condition occurs?</p> <p>Notification schemes available are:</p> <p>0 = No Error Notification</p> <p>1 = Notification via a TCP/IP Socket</p> <p>2 = Notification via the SNMP protocol</p> <p>3 = Notification via both TCP/IP and SNMP</p>	<p>Selecting this option results in:</p> <p>0 = No notification over network, but local notification through e-mail and an entry in the CC Agent error log file.</p> <p>1 = Notification messages sent to the GUI and local notification sent through e-mail and an entry in the CC Agent error log file.</p> <p>2 = Notification messages sent to the SNMP monitor and local notification sent through e-mail and an entry in the CC Agent error log file.</p> <p>3 = Notification messages sent to the GUI, the SNMP monitor, and local notification sent through e-mail and an entry in the CC Agent error log file.</p>
Names of your subsystems.	Specify unique names for each subsystem in your configuration.

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

Information Required to Complete Installation	How to Respond
*The character special file name for accessing the storage system (for example, hdisk4).	Enter the device name that your operating system assigned to the communications logical unit that you enabled in Appendix B. To determine the device name enter: lsdev -Cc disk Then locate the hdisk that corresponds to the communications LUN ("HSGxxCCL").
Do you want e-mail notification of errors or failures on the RAID subsystem?	Enter Y to receive e-mail and syslog notification. Enter N for syslog notification only.
*If yes to e-mail notification, specify e-mail address and notification level.	*Enter e-mail address and notification level for each user.
1 = Fatal error only	1 = Notification for serious errors, which require immediate attention.
2 = Warning and fatal errors	2 = Notification for serious errors and warning messages.
3 = Info, warning and fatal errors	3 = Notification for serious errors, warnings and informational messages.

*Information you need prior to starting the installation.

Upgrading the Agent

The RAID subsystem, Client, and notification information is preserved when you upgrade the Agent. To upgrade the Agent from earlier versions, refer to *Upgrading or Reinstalling the StorageWorks Solution Software* in Chapter 3.

Running the Agent

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

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

NOTE: The HACMP cluster software installation script does not place an entry into the “/etc/inittab” file. During operation of the HACMP as a cluster, Agent failure is handled by means of the HACMP event scripts.

Reconfiguring the Agent

Reconfigure the Agent by using the configuration and installation program.

1. Change to the */usr/stgwks2* directory and enter the following:

```
# ./stgwks_aix.sh
```

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

2. Select option 4, Agent Maintenance. The configuration program menu enables you to perform numerous modifications to the Agent and the way it performs. You can also use this script to manually start and stop the Agent daemon.

NOTE: If you are using HACMP cluster software you will need to perform this procedure on each node in the cluster.

Removing the Agent

Remove the Agent using the Agent Maintenance option from the *stgwks_aix.sh* Bourne-shell script.

1. Run the *./stgwks_aix.sh* script from the */usr/stgwks2* directory.
2. From the Installation and Configuration menu, select option 4, Agent Maintenance.
3. Select option 5, Uninstall Agent, from the Agent Maintenance Menu.

NOTE: If you are using HACMP cluster software you will need to perform this procedure on each node in the cluster.

High Availability Cluster Multi-Processing (HACMP)

The remainder of this chapter contains information about HACMP support. If you are not using HACMP, the Agent software instructions and menus are as outlined in the earlier sections of this chapter.

Enabler Software

If you are running HACMP software, you need to install the Enabler software found at:

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

Check the index for the HACMP Enabler bits for your subsystem type. Install Enabler on each node in the cluster using SMIT or the *installp* command line process. The enabler software provides HSG60/HSG80 concurrency on HACMP clusters of S70 and SP machines using the StorageWorks Fibre Channel Adapter, DS-SWIA1-PC and DS-SWIA1-PD.

General Note on Software Installation

In the HACMP environment, all Solutions Software must be installed on each node in the cluster. Follow the instructions in Chapter 3 for each node using the CD-ROM kit to install the StorageWorks Software and Agent on all nodes.

To clarify, from the *stgwks_aix.sh* menu, any installation, upgrade, configuration or removal of the StorageWorks Software including ODM files, Agent software and Agent configuration must be done on each node in the cluster. HACMP users have a slightly different configuration menu that allows for local Agent control.

Agent Maintenance Menu

HACMP Options

- | | |
|-------------------------------|--|
| 1) Change Agent Password | 21) Check ALL nodes for running Agent |
| 2) Change SNMP Enterprise OID | 22) Stop running Agent on ALL nodes |
| 3) (RESERVED) | 23) Check LOCAL node for running Agent |
| 4) (RESERVED) | 24) Stop running Agent on LOCAL node |
| 5) Uninstall Agent | 25) Start Agent on LOCAL node |

Client Options:

- 6) Add a Client
- 7) Remove a Client
- 8) Modify a Client
- 9) View Clients
- 10) Remove User from Mail Notification List
- 11) Modify Mail Notification List

Storage Subsystem Options:

- 10) Add a subsystem
- 11) Remove a subsystem
- 12) Modify a Subsystem
- 13) View subsystems

Agent Notification Options:

- 14) Enable/Disable system notification
- 15) Add user to mail notification list
- 16) Delete user from mail notification list
- 17) Modify mail notification list
- 18) View mail notification list
- 19) Enable/Disable mail notification
- 20) (RESERVED)

Q)

Enter Selection:

SWCC Clients

Client specifications must be the same for each node of the cluster to avoid unpredictable system behavior. Check that the \$STEAMDIR/steam/etc/client.ini file is identical on each node.

Running the Agent

There is an important distinction between running the Agent in a single host environment or in a cluster. In a single host environment, the Agent is started when you install it and it restarts automatically on reboot of the machine or if it stops for any reason.

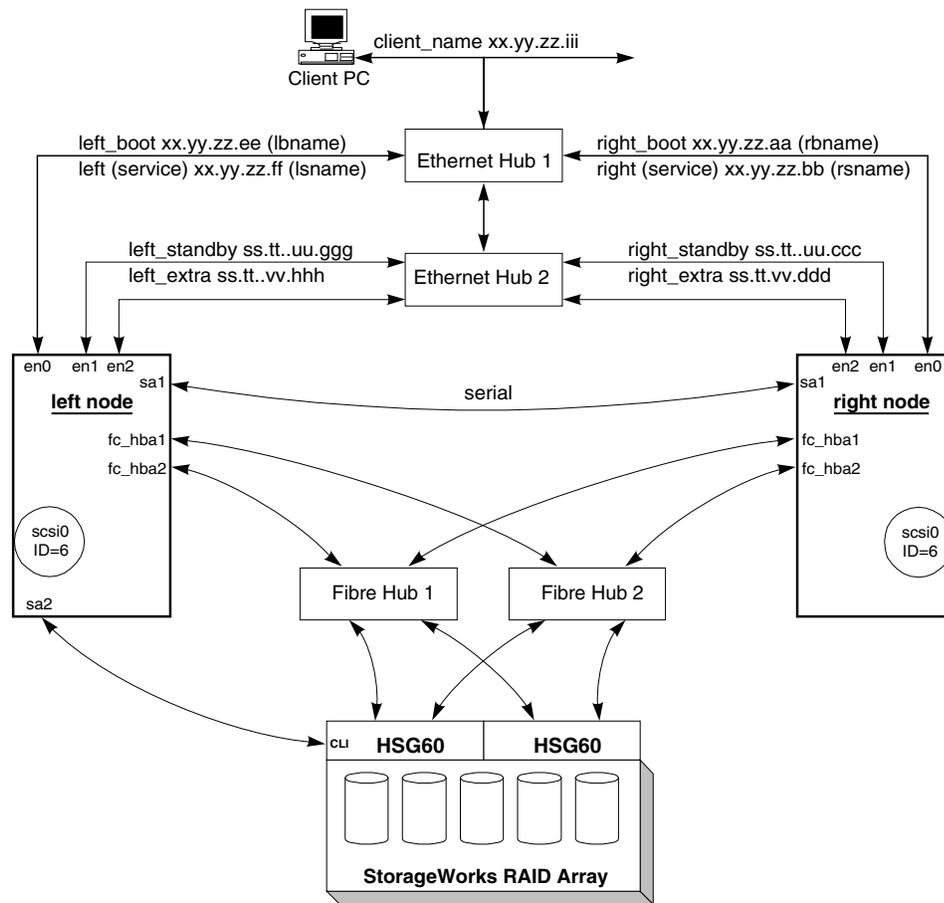
In an HACMP cluster, the installation script does not put an entry into the `/etc/inittab` file, therefore the Agent is not restarted if stopped. During operation of the HACMP environment as a cluster, Agent failure is handled by means of the HACMP event scripts provided in the Installation Kit.

HACMP Implementation for the Agent

Prerequisites

Table 4-5 lists the software programs that must be correctly installed and configured on each node in the cluster for the Agent to operate properly in an HACMP environment.

Prerequisites	Comments
AIX 4.3.2 or 4.3.3	This product is supported on AIX 4.3.2 and 4.3.3 platforms with StorageWorks (DS-SWIA1-PC, DS-SWIA1-PD).
HACMP 4.3 for AIX	The cluster single point of control (C-SPQC) utility for HACMP is available for versions of HACMP from 4.2 and later.
Simple HACMP configuration using SMIT	Supported cluster configurations are illustrated the Agent Maintenance Menu. Refer also to IBM Redbook: SG24-4498-00 HACMP/6000 customization examples or http://www.redbooks.ibm.com/abstracts/sg244498.html



CXO7325A

Figure 4-2. Command Console Client/Agent connection in a typical HACMP cluster

Implementation

The Agent installation and operation in the HACMP environment includes the procedures and functions involved in typical single node Agent operation. It also involves customizations to the HACMP environment itself. In addition there are several areas where the interaction of these two software packages can produce unique exceptions to ordinary performance of either HACMP or the Agent separately. Customization of the installation parameters and operational limitations will always be necessary to ensure

4-14 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

proper operation of the Agent in any High-Availability Cluster installation. The implementation of a fully functional Agent capable of high availability failover operation in an HACMP environment may be divided into four major components as follows:

- Installation
- Agent Configuration
- Client Installation & Configuration
- Operation

Procedure for Installing the Agent on an HACMP Cluster

There are four steps required to install the Agent on an HACMP cluster:

1. Copy the `stm_aix.tar` file from the `\swcc\agent` directory on the CD-ROM to a temporary directory on each node in the cluster as described in Chapter 3.
2. Expand the `stm_aix.tar` file into a suitable directory, preferably the default: `/usr/stgwks2`.
Perform this step on each node in the cluster.
3. Install the Agent on each node in the cluster as described earlier in this chapter.
4. Install Client Software on Remote Machine(s) as described in Appendix B.

Procedure for Configuring the Agent on an HACMP Cluster

The `HACMP_config.sh` utility, which is called automatically by the `config.sh` script, contains five new functions for use within the HACMP environment. The functions that act on the ALL nodes on the cluster are different from the equivalent functions used in previous versions of the Agent. These appear in the HACMP menu as Options 21 through 25 and are implemented as follows:

- Option 21. CHECK ALL nodes in the cluster for a running Agent.
This function checks for a running Agent by executing the `ps -eaf | grep steam` command automatically on each node in the cluster & reports the results back to the interactive operator.
- Option 22. STOP running Agent on ALL nodes in the cluster.
This function checks each node in the cluster for a running Agent, extracts the process ID for any Agent found and then executes the `kill -9` command to that process on that node.

- Option 23. CHECK LOCAL node for running Agent.

This function checks only the local node for a running Agent and reports the results back to the interactive operator.

- Option 24. STOP running Agent on the LOCAL node.

This function checks only the local node for a running Agent, extracts the process ID for any Agent found and then executes the `kill -9` command to that process on the local node.

- Option 25. START the Agent on the LOCAL node

This function checks for a running Agent by executing the `ps -eaf | grep steam` command automatically on each node in the cluster. If no Agents are running on any node in the cluster, the Agent is started on the local node.

Installing and Configuring the Client on an HACMP Cluster

Installing and configuring the Client is essentially the same whether using HACMP or not. The Client installation and configuration procedures are described in Appendix B.

Operation of the Agent on an HACMP Cluster

The three areas of operation of the Agent on an HACMP cluster that are supported in this release are:

- Normal operation of the Agent on an HACMP cluster.
- Failover with Agent as part of a Resource Group.
- Failure of the Agent itself.

Normal Operation of the Agent on an HACMP Cluster

Once the Agent has been installed onto each node of the cluster and the Client(s) has been installed and configured, the normal operating mode of the Agent in the HACMP environment may be activated by the following procedure.

1. Bring up the HACMP Cluster

Installation, configuration and operation of an HACMP cluster is described in the IBM publication: *High Availability Cluster Multi-Processing for AIX Installation Guide Version 4.3 (SC23-4278-00)*. Examples of typical HACMP configurations are described in IBM Redbook: *SG24-4498 -00 "HACMP/6000 Customization Examples* which is available directly from IBM or on the internet at:

<http://www.redbooks.ibm.com/abstracts/sg244498.html>

2. Configure the Agent Within HACMP

After the Agent has been installed and the HACMP cluster software is running, the Agent should be configured on each node and started on one node of the cluster. This is achieved using the unique features of the HACMP utility script:

```
/usr/stgwnks2/steam/bin/HACMP_config.sh.
```

3. Start the Agent Within HACMP

Use the **config.sh** utility to start the Agent on the local node. The **config.sh** utility automatically calls the **HACMP_config.sh** script containing the necessary functions. These functions are used to control the operation of the Agent.

There are five new menu options designed to support the operation of the Agent in the HACMP environment. The implementation details of these functions are described in the section on Agent configuration. Their application within the HACMP environment is described as follows:

a. Option **21**, CHECK ALL Nodes for Running Agent

This function determines if there are any instances of the Agent process running on any node in the cluster. It automatically polls each node through each known IP address that is currently active and reports the results back to the local operator. It is important to note that a node that is temporarily isolated, or is waiting to rejoin the cluster following a **node_down** event, will not be reported by this utility.

b. Option **22**, STOP Running Agent on ALL Nodes

This function is used to stop any instances of the Agent process running on any node in the cluster. Since it uses the same polling mechanism as Option 21, it can only act on nodes that are currently connected to the local node by at least one known IP address.

c. Option **23**, CHECK LOCAL Node for Running Agent

This function is used to determine if there are any instances of the Agent process running on the local node. It uses the `ps -eaf | grep steam` command and reports the results to the operator.

d. Option **24**, STOP Running Agent on LOCAL Node

This function is used to stop any instances of the Agent process running on the local node. It may be invoked manually on each node in a cluster whose nodes may have become isolated or have not yet rejoined the cluster.

e. Option **25**, START the Agent on LOCAL Node

This function is used to start the Agent process running on the local node. It automatically polls each node first as in Option 21, and it will not start the Agent if one is running already on any node in the cluster.

Starting, Stopping and Checking for the Agent

The difference between a single-node configuration and an HACMP cluster is critical when starting or stopping the Agent. To STOP a running Agent, the script utility must search the entire cluster and stop any Agent(s) found. This involves both stopping the process with the `kill -9` command and the removing the entry from `/etc/inittab`.

To START the Agent, the utility must first verify that there is no Agent running already, and then it must start the Agent on one node only. The `/etc/inittab` service is not used in the HACMP environment, because it would cause the Agent process to be re-spawned on each reboot of each node. This could easily lead to multiple Agents running on the same HSG60/HSG80 array, which is not allowed.

Password and Notification

The menu Options 1 and 2, involving Agent passwords and SNMP Enterprise OID's respectively are the same in appearance as in the single node configuration. However, they behave differently within the `HACMP_config.sh` script because this information should be kept consistent throughout the entire cluster. The `HACMP_config.sh` script ensures this consistency. Agent Notification Options 14 – 19 may be configured differently from node-to-node within a cluster without interfering with the operation of HACMP or the Agent. There may be valid reasons to customize an installation this way.

Client Functions and Storage Subsystem Options

The Client functions, Options 6 – 9, and storage subsystem Options, items 10 – 13, must be consistent on each node within the cluster. Since the `HACMP_config.sh` script does not synchronize these functions automatically, any changes made to Client Options or storage subsystem options must be repeated manually on each node in the cluster.

Failover with Agent as Part of a Resource Group

When HACMP is started, each participating node activates its service IP address and deactivates its boot IP address. Unless they have been configured to operate using the service IP addresses, the Agent will be inaccessible to Clients. In order for the Agent to maintain communication with its Client(s), the service IP address must be part of the HA Resource Group. These procedures are documented in the section “Normal Operation of the Agent on an HACMP Cluster,” page 4-16

Client IP Address

The Client IP Address must be placed into a Resource Group during the HACMP cluster initialization. The characteristics of this Resource Group must match those of the other shared resources in the cluster. This means that if the shared Volume Group(s) are identified as part of a Cascading Resource Group, then the Client IP Address must be placed into a Cascading Resource Group. If the shared Volume Group(s) is identified as part of a Rotating Resource Group, then the Client IP Address must be placed into a Rotating Resource Group.

Event Scripts

The operation of HACMP is directed by several event scripts, which will be automatically customized during the Agent installation as described in Chapter 4 of this document. These Event Scripts are usually located in `/usr/sbin/cluster/events` by default.

Notification within HACMP

Refer to the previous section on Password/Notification configuration.

Node Isolation Recovery

When any single node experiences a loss of services in a multi-node cluster such that it loses all contact with the rest of the cluster yet still remains up with a running Agent, it is considered an “isolated” node. It is important that the Agent process be killed prior to starting another instance of the Agent on another node. This scenario is handled by appending code to shut down the Agent unconditionally to the **“node_down_local”** event script. During Agent installation, this code is automatically appended to the **node_down_local** event script.

Failure of the Agent

Failure of the Agent during an HACMP session is handled by the HACMP event scripts. This is achieved by placing the Client IP address into a Resource Group and issuing the **server_down** event upon failure of the Agent process.

Polling all Nodes in Cluster for Presence of a Running Agent

When the Agent fails, HACMP invokes the **start_server** event script. This script will restart the Agent automatically using the **HACMP_config.sh** script, Option 25.

Restarting the Agent on Surviving Node

When the HACMP software completes the process of failover after a **node_down** event. The **start_server** event script will restart the Agent automatically using the **HACMP_config.sh** script, Option 25.

Notification

Notification type and level are determined by the settings of these parameters in the configuration of the surviving node at the time the Agent process is restarted. These parameters may or may not be the same for every node in the cluster.

Supporting Tools

The **HACMP_install.sh** installation script is called directly by the **install.sh** script and performs all necessary event modifications to the HACMP event scripts on the local node. This process must be repeated for each node in the cluster.

HACMP for AIX Primary Event Scripts

The HACMP Primary Event Scripts are part of the HACMP package and none require modification for Agent support. The purpose of the Primary Event Scripts is for the HACMP Event Manager to call one or more Secondary Event Scripts in response to the occurrence of certain events within the HACMP cluster. Some of the Secondary Event Scripts must be modified to support the Agent in the HACMP environment.

HACMP for AIX Secondary Event Scripts

Once it is installed and configured in an HACMP cluster environment, the operation and behavior of the Agent are completely determined by the action of “Secondary Event Scripts”. Once each of these scripts is examined and/or modified to support the Agent, the Agent will be “HACMP aware” for all normal topologies and configurations of HACMP for AIX. After installation of the Agent in the HACMP cluster, the only actions that will ever be necessary on the part of the Agent are:

1. Check for Agent running:
(SWCC_ag_check_local.sh SWCC_ag_check_all.sh)
2. Stop Agent:
(SWCC_ag_stop_local.sh SWCC_ag_stop_all.sh)
3. Start Agent:
(SWCC_ag_start_local.sh)

These actions are implemented using five additional scripts that are called from the modified event scripts. The five additional HACMP Agent scripts are:

- a. SWCC_ag_check_all.sh
- b. SWCC_ag_stop_all.sh
- c. SWCC_ag_check_local.sh
- d. SWCC_ag_stop_local.sh
- e. SWCC_ag_start_local.sh

Table 4-6 lists the scripts located in the “`$steamdir`”/bin directory which, by default, is the `/usr/stgwns2/steam/bin` directory created as part of the Agent installation. At the time of installation of the Agent, the following modifications are made automatically by the `HACMP_install.sh` script:

Table 4-6 HACMP Event Scripts

HACMP Event Script	Action Required	Appended to Event Script by HACMP_install.sh
Node_down_local	Always stop Agent on local node	"\$steamdir"/bin/SWCC_ag_stop_local.sh
Node_down_remote	(none)	(nothing)
Node_up_local	Check for Agent running on any node and start Agent on local node if not already running	"\$steamdir"/bin/SWCC_ag_start_local.sh
Node_up_remote	(none)	(nothing)
Node_down_local_complete	Check for SWCC Agent running on local node and only issue this event after Agent is stopped on local node	"\$steamdir"/bin/SWCC_ag_check_local.sh
Node_down_remote_complete	Check for Agent running on any node	"\$steamdir"/bin/SWCC_ag_check_all.sh
Node_up_local_complete	Check for Agent running on any node and start Agent on local node if not already running	"\$steamdir"/bin/SWCC_ag_start_local.sh
Node_up_remote_complete	(none)	(nothing)
Acquire_aconn_service	(none)	(nothing)
Acquire_service_addr	(none)	(nothing)
Acquire_takeover_addr	Check for Agent running and start it if not running	"\$steamdir"/bin/SWCC_ag_start_local.sh
Start_server	Check for Agent running on any node and start Agent on local node if not already running	"\$steamdir"/bin/SWCC_ag_start_local.sh

Table 4-6 HACMP Event Scripts (Continued)

HACMP Event Script	Action Required	Appended to Event Script by HACMP_install.sh
Stop_server	Check for Agent running on any node and stop any instance of the Agent	"\$steamdir"/bin/SWCC_ag_stop_all.sh
Get_aconn_rs	(none)	(nothing)
Get_disk_vg_fs	(none)	(nothing)
Release_aconn_rs	(none)	(nothing)
Release_service_addr	(none)	(nothing)
Release_takeover_addr	Check for Agent running and notify if takeover address is Client address	"\$steamdir"/bin/SWCC_ag_check_local.sh
Release_vg_fs	(none)	(nothing)
Swap_aconn_protocols	(none)	(nothing)

Chapter 5

Configuration Procedures

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

The following information is included in this chapter:

- “Establishing a Local Connection,” page 5–2
- “Setting Up a Single Controller,” page 5–3
- “Setting Up a Controller Pair,” page 5–9
- “Configuring Devices,” page 5–15
- “Configuring a Stripeset,” page 5–15
- “Configuring a Mirrorset,” page 5–15
- “Configuring a RAIDset,” page 5–16
- “Configuring a Striped Mirrorset,” page 5–17
- “Configuring a Single-Disk (JBOD) Unit,” page 5–18
- “Configuring a Partition,” page 5–18
- “Assigning Unit Numbers and Unit Qualifiers,” page 5–19
- “Configuration Options,” page 5–20

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 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

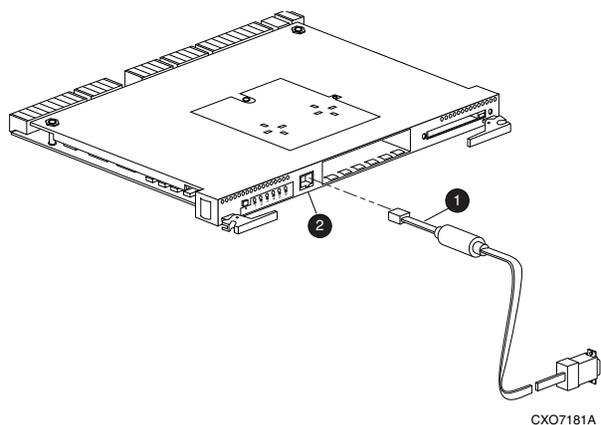
The configuration procedures assume that controllers and cache modules are installed in a fully functional and populated enclosure and that the PCMCIA cards are installed.

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.

Setting Up a Single Controller

Power On and Establish Communication

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

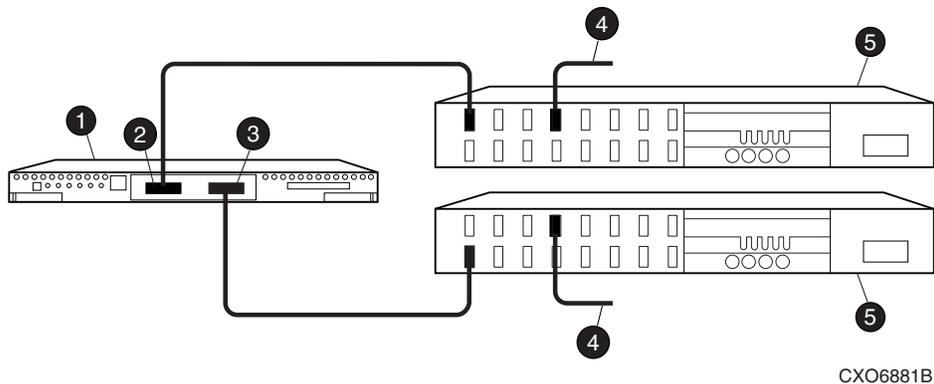
A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

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 (See “Configuring a Single Controller using CLI,” for more information).

5-4 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide



- 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. Establish a serial data connection to the RAID system Controller A from your host.
2. From the CLI prompt, enter:

```
SHOW THIS_CONTROLLER FULL
```

The current configuration of your RAID HSG60/HSG80 controllers will appear.
3. Check the configuration (parameter settings) of your controller against those listed in Table 5-1, which are the correct HSG60/HSG80 controller parameter settings for an IBM-AIX environment.

Table 5-1 Correct Controller Parameter Settings		
HSG60/HSG80 Parameter Settings for:	Loop	Fabric
Prompt	HSG60> HSG80>	HSG60> HSG80>
Node (WWN)	nnnn-nnnn-nnnn-nnnn	nnnn-nnnn-nnnn-nnnn
Command Console LUN	Enabled	Enabled
Port 1 Topology	LOOP HARD	FABRIC

Table 5-1 Correct Controller Parameter Settings

HSG60/HSG80 Parameter Settings for:	Loop	Fabric
Port 2 Topology	LOOP HARD	FABRIC
Port 1 ALPA	D5	N/A
Port 2 ALPA	D6	N/A
SCSI Version	SCSI-2	SCSI-2
New Connection OS Type	WINNT	WINNT

Setting the SCSI version to SCSI-3 does not make the controller fully compliant with all SCSI-3 standards. IBM AIX does fully support SCSI-3 mode at this time.

- Change any controller parameter settings inconsistent with those in Table 5-1 by entering the appropriate command for that parameter, as described in the following steps.

NOTE: After completing all of the required parameter changes, you must restart the controller for the new settings to take effect.

Ignore the Restart Controller messages until you have finished making all the required configuration parameter changes.

- To set the prompt for the HSG60/HSG80 controller, enter:
SET THIS_CONTROLLER PROMPT="HSG>"
- Clear the persistent error messages using:
HSG> CLEAR CLI
- Set the node by acquiring the World Wide Name (WWN) node number and checksum from the enclosed label (located on the top surface of the controller enclosure), and using:
HSG> SET THIS_CONTROLLER NODE=NNNN-NNNN-NNNN-NNNN CHECKSUM
- Enable the command console LUN using:
HSG> SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
- Set the port topology offline using:
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=OFFLINE
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=OFFLINE

5-6 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

10. Set the port topology using:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC
```

```
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

See Chapter 3 for an example of LOOP_HARD topology setup.

11. Set the SCSI version using:

```
HSG> SET THIS_CONTROLLER SCSI VERSION=SCSI-2
```

12. Determine if the operating system set for the online controller connection is WINNT, using:

```
HSG> SHOW CONNECTIONS
```

Although an IBM operating system entry exists in the CLI and SWCC for the fibre channel connection, use default setting WINNT for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover. If so, use:

```
HSG> SET !NEWCONxx OPERATING_SYSTEM=WINNT
```

where: xx is the online connection ID produced in the SHOW CONNECTIONS command above.

NOTE: Further information on command syntax can be found in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*

13. Restart the controllers for the new settings to take effect, using:

```
HSG> RESTART OTHER_CONTROLLER
```

```
HSG> RESTART THIS_CONTROLLER
```

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

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

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

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

```
HSG> 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.

16. 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.

17. Verify that all commands have taken effect. Use the following command:

```
HSG> 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:

HSG ZG09030200 Software V85L, Hardware E01

NODE_ID = 5000-1FE1-0000-00A0

ALLOCATION_CLASS = 0

SCSI_VERSION = 2

Device Port SCSI address 6

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

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0000-0001

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 011121

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0000-0002

PORT_2_TOPOLOGY = FABRIC (fabric up)

Address = 011221

Cache: 128 megabyte write cache, version 0012 Cache is GOOD

.....

18. Turn on the switches, if not done previously.
19. 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:

```
HSG> SHOW CONNECTIONS
```

IMPORTANT: Although an IBM operating system entry exists in the CLI and SWCC for the Fibre Channel connection, use the default setting WINNT for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

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	WINNT	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 “Connections,” page 1–13.

20. 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:

```
HSG> RENAME !NEWCON01 ANGEL1A1
```

For a recommended naming convention, see “Naming Connections,” page 1–13.

21. Specify the operating system for the connection:

```
HSG> SET ANGEL1A1 OPERATING_SYSTEM=WINNT
```

IMPORTANT: Although an IBM operating system entry exists in the CLI and SWCC, for the Fibre Channel connection, use the default setting WINNT for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

22. Verify the changes:

```
SHOW CONNECTIONS
```

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

23. Repeat step 21, 22, and 23 for each of that adapter’s host connections or delete the unused connections from the table.
24. For each host adapter, repeat steps 20 through 24.

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 COM1 or COM2 (serial) 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**.

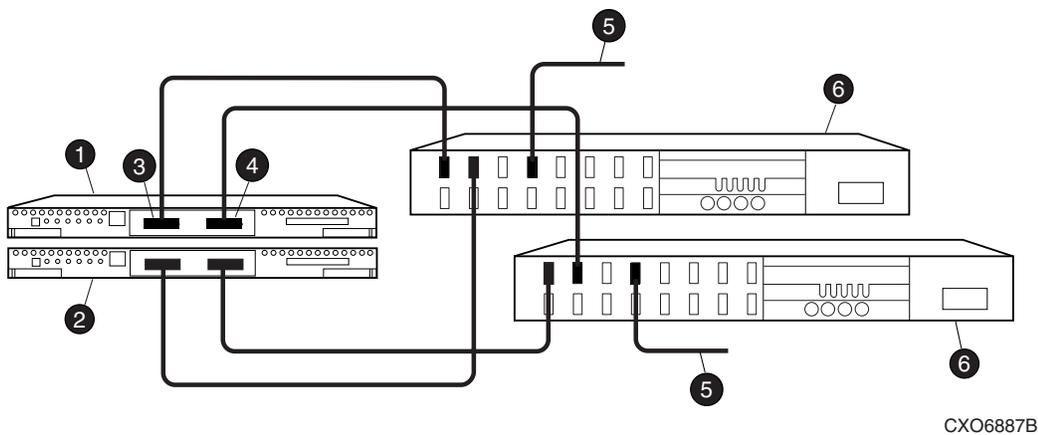
A copyright notice and the CLI prompt appear, indicating that you established a local connection with the controller.

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 hub or switch. The host cables are plugged into the hub or switch as part of the configuration procedure (See “Configuring a Controller Pair using CLI,” for more information).

5-10 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide



- 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

Configuring a Controller Pair using CLI

1. Establish a serial data connection to the RAID system Controller A from your host.
2. From the CLI prompt, enter:

```
SHOW THIS_CONTROLLER FULL
```

The current configuration of your RAID HSG60/HSG80 controllers will appear.
3. Check the configuration (parameter settings) of your controller against those listed in Table 5-1, which are the correct HSG60/HSG80 controller parameter settings for an IBM-AIX environment.

Table 5-2 Correct Controller Parameter Settings

HSG60/HSG80 Parameter Settings for:	Loop	Fabric
Prompt	HSG60> HSG80>	HSG60> HSG80>
Node (WWN)	nnnn-nnnn-nnnn-nnnn	nnnn-nnnn-nnnn-nnnn

Table 5-2 Correct Controller Parameter Settings

HSG60/HSG80 Parameter Settings for:	Loop	Fabric
Command Console LUN	Enabled	Enabled
Port 1 Topology	LOOP HARD	FABRIC
Port 2 Topology	LOOP HARD	FABRIC
Port 1 ALPA	D5	N/A
Port 2 ALPA	D6	N/A
SCSI Version	SCSI-2	SCSI-2
New Connection OS Type	WINNT	WINNT

Setting the SCSI version to SCSI-3 does not make the controller fully compliant with the SCSI-3 standards. IBM AIX does not fully support SCSI-3 mode at this time.

- Change any controller parameter settings inconsistent with those in Table 5-2 by entering the appropriate command for that parameter.

NOTE: After completing all of the required parameter changes, you must restart the controller for the new settings to take effect.

Ignore the Restart Controller messages until you have finished making all the required configuration parameter changes.

- To set the prompt for the HSG60/HSG80 controller, enter:
SET THIS_CONTROLLER PROMPT="HSG>"
- Clear the persistent error messages using:
HSG> CLEAR CLI
- Set the node by acquiring the World Wide Name (WWN) node number and checksum from the enclosed label (located on the top surface of the controller enclosure), and using:
HSG> SET THIS_CONTROLLER NODE=NNNN-NNNN-NNNN-NNNN CHECKSUM
- Enable the command console LUN using:
HSG> SET THIS_CONTROLLER COMMAND_CONSOLE_LUN
- Set the port topology offline using:
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=OFFLINE
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=OFFLINE

5-12 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

10. Set the port topology using:

```
HSG> SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC
```

```
HSG> SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

See Chapter 3 for an example of LOOP_HARD topology setup.

11. Set the SCSI version using:

```
HSG> SET THIS_CONTROLLER SCSI VERSION=SCSI-2
```

12. Determine if the operating system set for the online controller connection is WINNT, using:

```
HSG> SHOW CONNECTIONS
```

Although an IBM operating system entry exists in the CLI and SWCC, for the fibre channel connection, use default setting WINNT for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover. If so, use:

```
HSG> SET !NEWCONxx OPERATING_SYSTEM=WINNT
```

where: xx is the online connection ID produced in the SHOW CONNECTIONS command above.

NOTE: Further information on command syntax can be found in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*

13. Restart the controllers for the new settings to take effect, using:

```
HSG> RESTART OTHER_CONTROLLER
```

```
HSG> RESTART THIS_CONTROLLER
```

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

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

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

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

```
HSG> RUN FRUTIL
```

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

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

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

16. 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.

17. Verify that all commands have taken effect. Use the following command:

```
HSG> 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:

HSG ZG09030200 Software V85L, Hardware E01

NODE_ID = 5000-1FE1-0000-00A0

ALLOCATION_CLASS = 0

SCSI_VERSION = 2

Device Port SCSI address 6

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

Host PORT_1:

Reported PORT_ID = 5000-1FE1-0000-0001

PORT_1_TOPOLOGY = FABRIC (fabric up)

Address = 011121

Host PORT_2:

Reported PORT_ID = 5000-1FE1-0000-0002

PORT_2_TOPOLOGY = FABRIC (standby)

Address = 011221

Cache: 128 megabyte write cache, version 0012 Cache is GOOD

.....

18. Turn on the switches, if not done previously.
19. Plug in the Fibre Channel cable from the first adapter on the first host into the hub. Enter a SHOW *connections* command to view the connection table:

```
HSG> 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.

5-14 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

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

20. 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:

```
HSG> RENAME !NEWCON01 ANGEL1A1
```

Compaq recommends using a naming convention. See “Naming Connections,” page 1-13.

21. Specify the operating system for the connection:

```
HSG> SET ANGEL1A1 OPERATING_SYSTEM=WINNT
```

IMPORTANT: Although an IBM operating system entry exists in the CLI and SWCC, for the Fibre Channel connection, use the default setting **WINNT** for operation with your IBM host. The IBM setting may cause problems when LUNs are not available to AIX or during controller failover.

22. Verify the changes:

```
HSG> SHOW CONNECTIONS
```

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

23. Repeat steps 21, 22, and 23 for each of that adapter’s host connections or delete the unwanted connections from the table.
24. For each host adapter, repeat steps 20 through 24.

Configuring Devices

The disks on the device bus of the HSG60/HSG80 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 two 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 STRIPESSET STRIPESSET-NAME DISKNNNNN DISKNNNNN .....
```
2. Initialize the stripese, specifying any desired switches:

```
INITIALIZE STRIPESSET-NAME SWITCHES
```

See "Initialization Switches," page 2-28 for a description of the initialization switches.
3. Verify the stripese configuration:

```
SHOW STRIPESSET-NAME
```
4. Assign the stripese a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers," page 5-19.

Example:

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

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

Configuring a Mirrorset

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

```
ADD MIRRORSET MIRRORSET-NAME DISKNNNNN DISKNNNNN SWITCHES
```

5-16 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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,” page 2–28 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,” page 5–19.

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,” page 2–28 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,” page 5–19.

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,” page 5–15.
2. Create a stripeset and specify the mirrorsets it contains:
ADD STRIPESET STRIPESET-NAME MIRRORSET-1 MIRRORSET-2...MIRRORSET-N
3. Initialize the striped mirrorset, specifying any desired switches:
INITIALIZE STRIPESET-NAME SWITCH
See “Initialization Switches,” page 2–28 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,” page 5–19.

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,” page 2–28 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,” page 5–19.

Configuring a Partition

1. Initialize the storage set or disk drive, specifying any desired switches:
INITIALIZE STORAGESET-NAME SWITCHES
or
INITIALIZE DISK-NAME SWITCHES
See “Initialization Switches,” page 2–28 for a description of the initialization switches.
2. Create each partition in the storage set 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 storage set 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,” page 5-19.

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 1-20 and “Worldwide Names (Node IDs and Port IDs),” page 1-22.

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
```

5-20 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

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

Verifying Storage Configuration from the Host

Once the StorageWorks Solution Software has been installed, logical units may be created on the RAID Array and configured into the system. These logical units can be created using the Client or the CLI. In order to utilize Client to create logical units, the Agent and Client must first be installed on the AIX and Windows systems, respectively.

New logical units created on the RAID Array must be recognized by the operating system before they can be utilized. The following steps outline the process:

1. Using the CLI or the Client, create the logical units on the RAID Array.
2. As the superuser (*root*), run the command:

```
# cfgmgr -v
```

3. Verify that the host has located and configured the new RAID Array devices in the operating system using the command:

```
# lsdev -Cc disk
```

The output of the command looks like:

```
hdisk0 Available 00-07-A1-3,0 4.5 GB 16 Bit SCSI Disk Drive
hdisk1 Available X0-06-01-0,0 DEC HSG Command Console LUN
hdisk2 Available X0-06-01-0,1 DEC HSG RAID Array
hdisk3 Available X0-06-01-0,2 DEC HSG RAID Array
hdisk4 Available X0-06-01-0,3 DEC HSG RAID Array
hdisk5 Available X0-06-01-0,4 DEC HSG RAID Array
hdisk6 Available X0-06-01-0,5 DEC HSG RAID Array
hdisk7 Available X0-06-01-0,6 DEC HSG RAID Array
hdisk8 Available X0-06-01-0,7 DEC HSG RAID Array
hdisk9 Available X0-06-01-2,2 DEC HSG RAID Array
hdisk10 Available X0-06-01-2,4 DEC HSG RAID Array
hdisk11 Available X0-06-01-4,4 DEC HSG RAID Array
hdisk12 Available X0-06-01-5,1 DEC HSG RAID Array
```

6-2 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

You should have an entry of type:

hdisknn Available X0-06-01-2,4 DEC HSG RAID Array

for each logical unit that you have created on the RAID Array.

4. If the devices were not recognized, re-check the status of the RAID controller, logical units, and cable connections.

You can now use the *System Management Interface Tool (SMIT)*, to make volume groups, logical volumes and file systems using these logical units.

NOTE: If you are using HACMP cluster software, you need only perform this procedure once and it will apply for the entire cluster. However, shared storage must have the same “hdisk” number on each node in the cluster and shared volume groups, logical volumes and file systems must be organized as elements of shared resource groups for the defined clusters. For more information, see *High Availability for Multi-Processing for AIX Installation Guide Version 4.3*.

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 4214R 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 AIX 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:

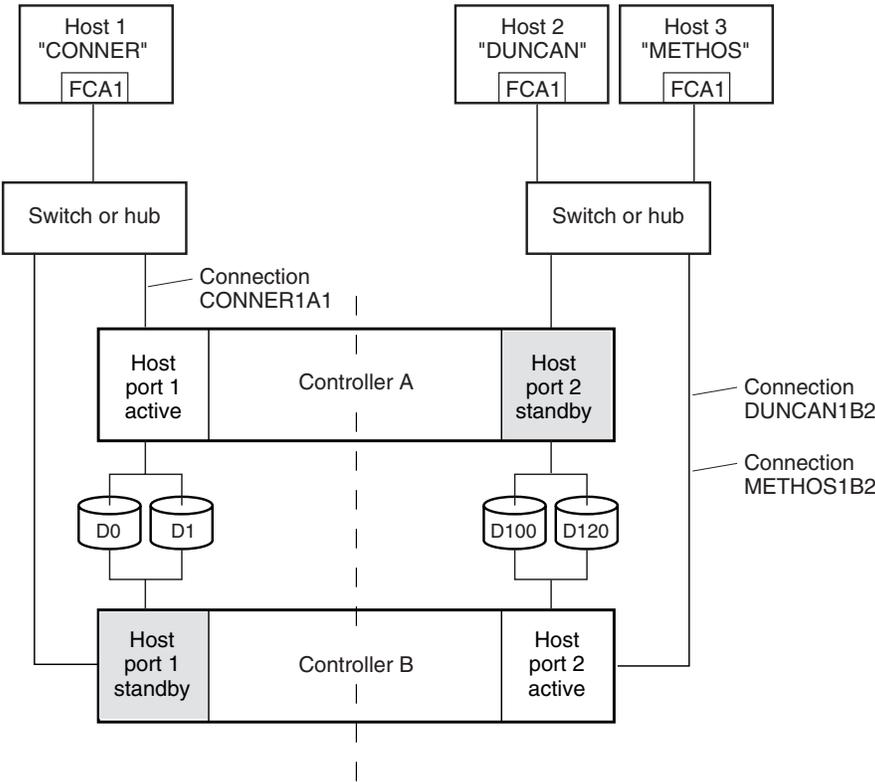
- Screen messages from CLI or utilities are boxed.
- Instructions to users are in italics.

HSG60 Example

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D0 R1 DISK10000	D0 R1 DISK10100	D0 R1 DISK10200	D1 S1 M1 DISK10300	D1 S1 M1 DISK10400	D1 S1 M2 DISK10500	D1 S1 M2 DISK10800	D100 M3 DISK10900	D100 M3 DISK11000	D120 S2 DISK11100	D120 S2 DISK11200	Spare DISK11300	not supported	not supported

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

Figure 7-1. Example storage map for HSG60

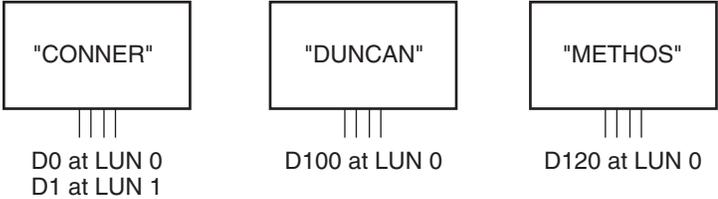


NOTE: FCA = Fibre Channel Adapter

CXO7296A

Figure 7-2. Example system for the HSG60

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



CXO7297A

CLI Configuration Example for the HSG60

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_PROFILE=PLDA
SET THIS PORT_2_PROFILE=PLDA
SET OTHER PORT_1_PROFILE=PLDA
SET OTHER PORT_2_PROFILE=PLDA
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

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

Plug serial cable from maintenance terminal into bottom controller.

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

```
RUN FRUTIL
```

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

Y

```
SET THIS MIRRORED_CACHE
```

NOTE: This command causes the controllers to restart.

```
SET THIS PROMPT="HLNDR BOTTOM"
```

```
SET OTHER PROMPT="HLNDR TOP"
```

```
SHOW THIS
```

SHOW OTHER

Plug in the Fibre Channel cable from the adapter in host "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=WINNT

SHOW CONNECTIONS

NOTE: Connection table sorts alphabetically.

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

Mark or tag both ends of Fibre Channel cables.

Plug in the Fibre Channel cable from the second adapter in host "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	WINNT	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=WINNT

SHOW CONNECTIONS

7-6 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

Mark or tag both end of Fibre Channel cables.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
CONNER1A1	WINNT	OTHER	1	XXXXXX	OL other	0
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
DUNCAN1B2	WINNT	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

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

RENAME !NEWCON02 METHOS1B2

SET METHOS1B2 OPERATING_SYSTEM=WINNT

SET METHOS1B2 UNIT_OFFSET=120

SHOW CONNECTIONS

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

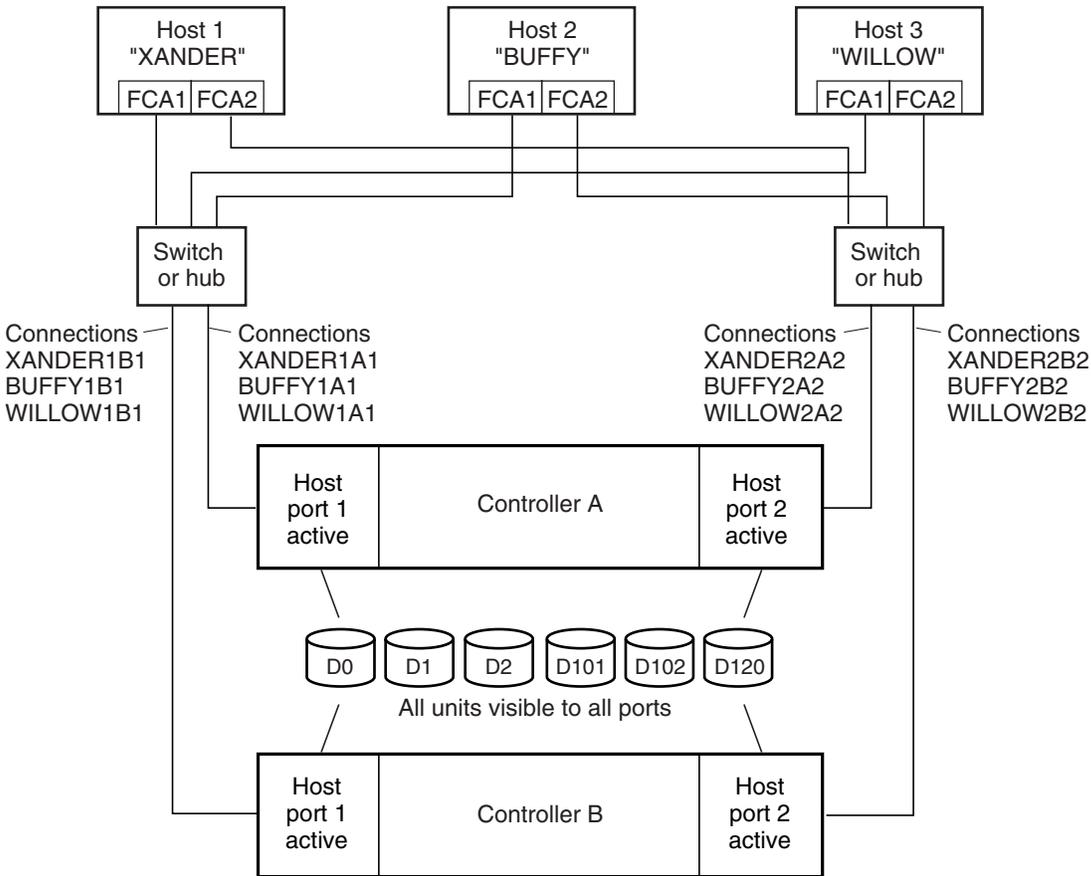
Mark or tag both end of Fibre Channel cables.

```
RUN CONFIG
ADD RAIDSET R1 DISK10000 DISK20000 DISK10100 DISK20100 DISK10200 DISK20200
INITIALIZE R1
ADD UNIT D0R1 DISABLE_ACCESS_PATH=ALL
SET D0 ENABLE_ACCESS_PATH=(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
```

HSG80 Example

	Port						
	1	2	3	4	5	6	
Power Supply	D2 S2 DISK10300	D2 S2 DISK20300	D2 S2 DISK30300	D2 S2 DISK40300	D2 D101 DISK50300	spareset member DISK60300	Power Supply 3
Power Supply	D0 S1 MI DISK10200	D0 S1 M1 DISK20200	D0 S1 M2 DISK30200	D0 S1 M2 DISK40200	D1 M3 DISK50200	D1 M3 DISK60200	Power Supply 2
Power Supply	D120 R2 DISK10100	D120 R2 DISK20100	D120 R2 DISK30100	D120 R2 DISK40100	D120 R2 DISK50100	D120 R2 DISK60100	Power Supply 1
Power Supply	D102 R1 DISK10000	D102 R1 DISK20000	D102 R1 DISK30000	D102 R1 DISK40000	D102 R1 DISK50000	D102 R1 DISK60000	Power Supply 0

Figure 7-4. Example storage map for the BA370



NOTE: FCA = Fibre Channel Adapter

CXO7109A

Figure 7-5. Example system for the HSG80

7-10 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

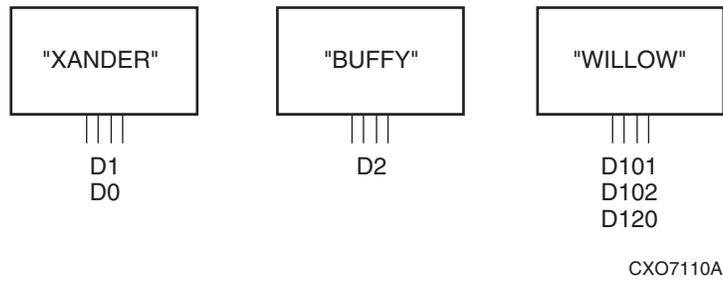
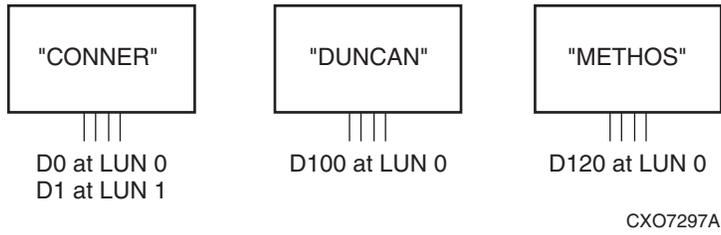


Figure 7-6. Example virtual layout from the hosts' point of view

Figure 7-7. Example system for the HSG60



CLI Configuration Example for the HSG80

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_PROFILE=PLDA
SET THIS PORT_2_PROFILE=PLDA
SET OTHER PORT_1_PROFILE=PLDA
SET OTHER PORT_2_PROFILE=PLDA
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
```

7-12 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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=WINNT

SHOW CONNECTIONS

NOTE: Connection table sorts alphabetically.

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

Mark or tag both ends of Fibre Channel cables.

Plug in the Fibre Channel cable from the second adapter in host "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	WINNT	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=WINNT

SHOW CONNECTIONS

Mark or tag both end of Fibre Channel cables.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
CONNER1A1	WINNT	OTHER	1	XXXXXX	OL other	0
		HOST_ID=XXXX-XXXX-XXXX-XXXX		ADAPTER_ID=XXXX-XXXX-XXXX-XXXX		
DUNCAN1B2	WINNT	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

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

RENAME !NEWCON02 METHOS1B2

SET METHOS1B2 OPERATING_SYSTEM=WINNT

SET METHOS1B2 UNIT_OFFSET=120

SHOW CONNECTIONS

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

Mark or tag both end of Fibre Channel cables.

7-14 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

```
RUN CONFIG
ADD RAIDSET R1 DISK10000 DISK20000 DISK30000 DISK40000 DISK50000 DISK60000
INITIALIZE R1
ADD UNIT D102R1 DISABLE_ACCESS_PATH=ALL
SET D102 ENABLE_ACCESS_PATH=(CONNER1A1)
ADD RAIDSET R2 DISK10100 DISK20100 DISK30100 DISK40100 DISK50100 DISK60100
INITIALIZE R2
ADD UNIT D120 R2 DISABLE_ACCESS_PATH=ALL
SET D120 ENABLE_ACCESS_PATH=(METHOS1B2)
ADD MIRRORSET MI DISK10200 DISK20200
ADD MIRRORSET M2 DISK30200 DISK40200
ADD STRIPESET S1 M1 M2
INITIALIZE S1
ADD UNIT D0 S1 DISABLE_ACCESS_PATH=ALL
SET D0 ENABLE_ACCESS_PATH=(CONNER1A1)
ADD MIRRORSET M3 DISK50200 DISK60200
INITIALIZE M3
ADD UNIT D1 M3 DISABLE_ACCESS_PATH=ALL
SET D1 ENABLE_ACCESS_PATH=(DUNCAN1B2)
ADD STRIPESET S2 DISK10300 DISK20300 DISK30300 DISK40300
INITIALIZE S2
ADD UNIT D2 S2 DISABLE_ACCESS_PATH=ALL
SET D2 ENABLE_ACCESS_PATH=(CONNER1A1)
INITIALIZE DISK50300
ADD UNIT D101 DISK50300 DISABLE_ACCESS_PATH=ALL
SET D101 ENABLE_ACCESS_PATH=(DUNCAN1B2)
ADD SPARESET DISK60300
SHOW UNITS ALL
```

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 8-5

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.

8-2 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

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

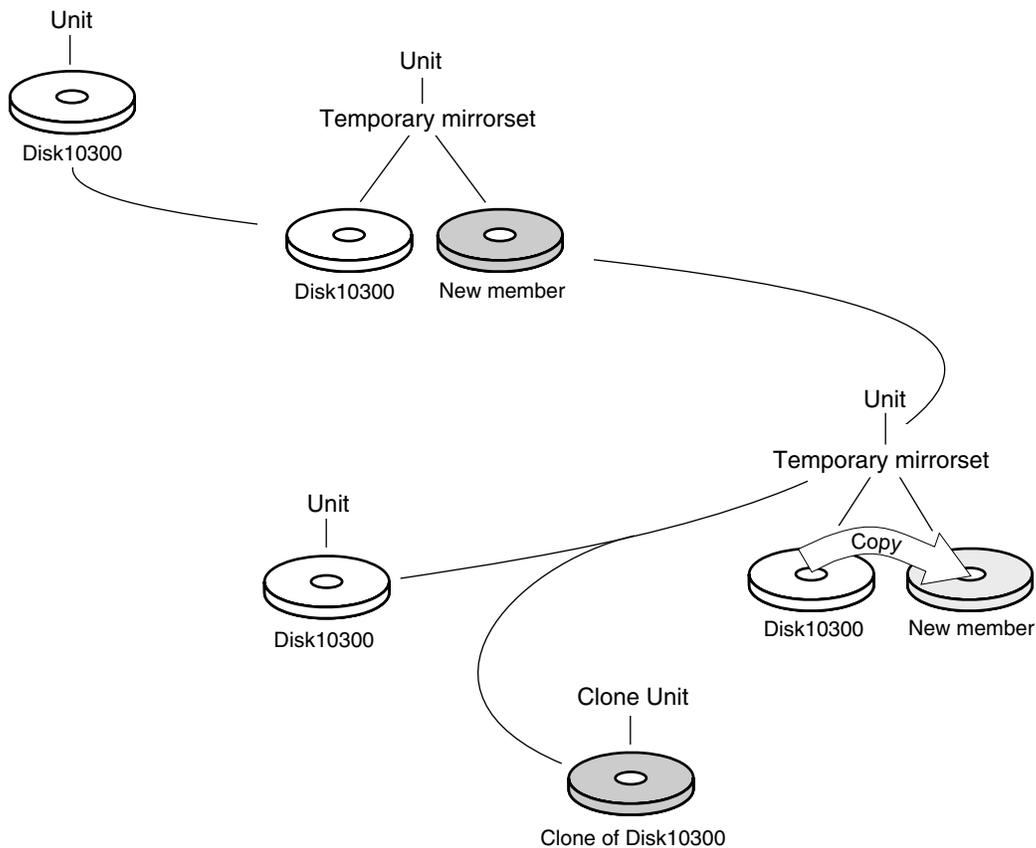
```
SHOW DEVICES FULL
```

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 back up. 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

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.

8-4 *HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide*

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
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 STRIPESSET C_ST1 C_MA C_MB
INIT C_ST1 NODESTROY
ADD UNIT D99 C_ST1
D99 HAS BEEN CREATED. IT IS A CLONE OF D98.
CLONE - NORMAL TERMINATION
```

Moving Stagesets

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



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

See the release notes for the version of 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 storage set while maintaining the data it contains:

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

```
SHOW STORAGESET-NAME
```

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

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

```
DELETE UNIT NUMBER
```

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

```
DELETE STORAGESET-NAME
```

5. Delete each disk drive, one at a time, that the storage set 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 storageSet by adding its name to the controller's list of valid storageSets and specifying the disk drives it contains. (Although you have to recreate the storageSet from its original disks, you do not have to add them in their original order.)

Use the following syntax:

```
ADD STORAGESET-NAME DISK-NAME DISK-NAME
```

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

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

Example:

The following example moves unit D100 to another cabinet. D100 is the RAIDset RAID99 that consists of members DISK10000, DISK20000, DISK10100 and DISK20100.

Delete the disks from the source cabinet through the source cabinet controller. Use the following syntax:

```
DELETE D100
```

```
DELETE RAID99
```

```
DELETE DISK10000
```

```
DELETE DISK10100
```

```
DELETE DISK20000
```

```
DELETE DISK20100
```

Move the disks to the destination cabinet, and add them through the destination cabinet controller. Use the following syntax:

```
ADD DISK DISK10000 100
```

```
ADD DISK DISK10100 110
```

```
ADD DISK DISK20000 200
```

```
ADD DISK DISK20100 210
```

```
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. Four (4) templates will be needed for the subsystem.

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

- “Storageset Profile,” page A-2
- “Templates for the HSG60,” page A-3
 - “Storage Map Template 1 for the Model 4214R Disk Enclosure,” page A-3
 - “Storage Map Template 2 for the Model 4254 Disk Enclosure,” page A-4
 - “Storage Map Template 3 for Two Model 4214R Disk Enclosures,” page A-5
- “Templates for the HSG80,” page A-6
 - “Storage Map Template 1 for the BA370 Enclosure,” page A-6
 - “Storage Map Template 2 for the Second BA370 Enclosure,” page A-7
 - “Storage Map Template 3 for the Third BA370 Enclosure,” page A-8
 - “Storage Map Template 4 for the Model 4214R Disk Enclosure,” page A-9
 - “Storage Map Template 5 for the Model 4254 Disk Enclosure,” page A-10
 - “Storage Map Template 6 for the Model 4310R Disk Enclosure,” page A-12
 - “Storage Map Template 7 for the Model 4350R Disk Enclosure,” page A-14

StorageSet Profile

Type of StorageSet:

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

StorageSet Name

Disk Drives

Unit Number

Partitions:

| Unit # |
|--------|--------|--------|--------|--------|--------|--------|--------|
| | | | | | | | |

RAIDset Switches:

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

Mirrorset Switches:

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

Initialize Switches:

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

Unit Switches:

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

Templates for the HSG60

Storage Map Template 1 for the Model 4214R Disk Enclosure

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

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	not supported	not supported

Storage Map Template 2 for the Model 4254 Disk Enclosure

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

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

Storage Map Template 3 for Two Model 4214R Disk Enclosures

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

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

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	not supported	not supported

Model 4214R Disk Enclosure Shelf 2 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	not supported	not supported

Templates for the HSG80

Storage Map Template 1 for the BA370 Enclosure

Use this template for:

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

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

Targets

Storage Map Template 2 for the Second BA370 Enclosure

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

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

Targets

Storage Map Template 3 for the Third BA370 Enclosure

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

	Port						
	1	2	3	4	5	6	
Power Supply	D11500	D21500	D31500	D41500	D51500	D61500	Power Supply
Power Supply	D11400	D21400	D31400	D41400	D51400	D61400	Power Supply
Power Supply	D11300	D21300	D31300	D41300	D51300	D611300	Power Supply
Power Supply	D11200	D21200	D31200	D41200	D51200	D61200	Power Supply

15

14

13

12

Targets

Storage Map Template 4 for the Model 4214R Disk Enclosure

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

Model 4214R Disk Enclosure Shelf 1 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk10000	Disk10100	Disk10200	Disk10300	Disk10400	Disk10500	Disk10800	Disk10900	Disk11000	Disk11100	Disk11200	Disk11300	not supported	not supported

Model 4214R Disk Enclosure Shelf 2 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk20000	Disk20100	Disk20200	Disk20300	Disk20400	Disk20500	Disk20800	Disk20900	Disk21000	Disk21100	Disk21200	Disk21300	not supported	not supported

Model 4214R Disk Enclosure Shelf 3 (single-bus)														
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk30900	Disk31000	Disk31100	Disk31200	Disk31300	not supported	not supported

Storage Map Template 5 for the Model 4254 Disk Enclosure

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

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

Model 4254 Disk Enclosure Shelf 2 (dual-bus)														
Bus A							Bus B							
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk30000	Disk30100	Disk30200	Disk30300	Disk30400	Disk30500	Disk30800	Disk40000	Disk40100	Disk40200	Disk40300	Disk40400	Disk40500	Disk40800

continued on the following page

continued from previous page

Model 4254 Disk Enclosure Shelf 3 (dual-bus)														
Bus A								Bus B						
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
SCSI ID	00	01	02	03	04	05	08	00	01	02	03	04	05	08
DISK ID	Disk50000	Disk50100	Disk50200	Disk50300	Disk50400	Disk50500	Disk50800	Disk60000	Disk60100	Disk60200	Disk60300	Disk60400	Disk60500	Disk60800

Storage Map Template 6 for the Model 4310R Disk Enclosure

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

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

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

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

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

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

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

continued on the following page

continued from previous page

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

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

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

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

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

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

Storage Map Template 7 for the Model 4350R Disk Enclosure

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

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

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

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

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

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

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

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-2
- “Troubleshooting the Client Installation,” page B-3
- “Adding the Storage Subsystem and its Host to the Navigation Tree,” page B-4
- “Removing the Command Console Client,” page B-6
- “Where to Find Additional Information,” page B-7

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, navigate to:
`\swcc\client\intel`
Double-click setup.exe
3. Select HSG60/HSG80 Controller and click *N*ext.

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 information on how to resolve some of the following 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.

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). 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.

B-4 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

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

spgui	4998/tcp	#Command Console
spagent	4999/tcp	#HS-Series 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.

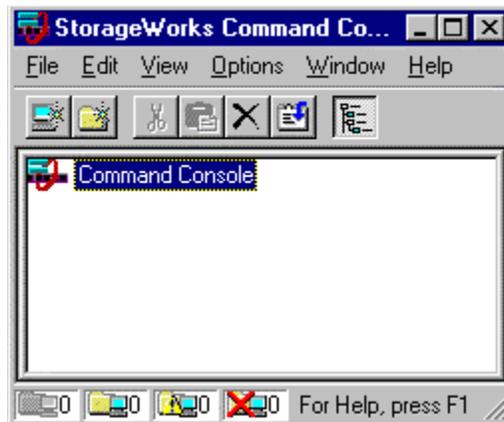


Figure B-1. Navigation window

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.

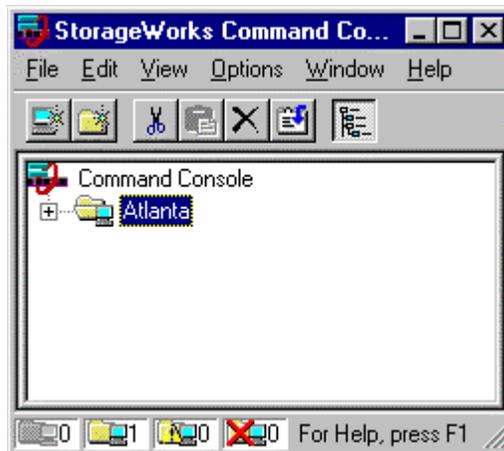


Figure B-2. Navigation window showing storage host system "Atlanta"

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/HSG80 Client by using the Add/Remove program.

Where to Find Additional Information

You can find additional information about SWCC by referring to the online Help and to *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 include:

- 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, Online Help provides 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

- transparent failover 1-16

ADD UNIT

- transparent failover 1-16

adding

- virtual disks B-7

adding a disk drive to the spareset

- configuration options 5-21

adding disk drives

- configuration options 5-21

Agent

- functions 4-1

- installing 4-5

- reconfiguring 4-9

- running 4-9

- uninstalling 4-9

array of disk drives 2-15

assigning unit numbers 1-16

assignment

- unit numbers

 - fabric topology 5-19

- unit qualifiers

 - fabric topology 5-19

assignment of unit numbers

- fabric topology

 - partition 5-20

 - single disk 5-19

asynchronous event service B-7

autospare

enabling

- fabric topology 5-22

availability 2-21

B

backup

- cloning data 8-2

- subsystem configuration 8-1

C

cabling

- controller pair 5-10

- multiple-bus failover

 - fabric topology configuration 5-9

- single controller 5-4

cache modules

- location 1-2, 1-4, 1-5

- read caching 1-10

- write-back caching 1-10

- write-through caching 1-10

caching techniques

- mirrored 1-11

- read caching 1-10

- read-ahead caching 1-10

- write-back caching 1-10

- write-through caching 1-10

changing switches

- configuration options 5-22, 5-23, 5-24

chunk size

- choosing for RAIDsets and stripesets 2-28

- controlling stripesize 2-28

Index-2 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

- using to increase data transfer rate 2-30
- using to increase request rate 2-29
- using to increase write performance 2-31
- CHUNKSIZE 2-28
- CLI prompt
 - changing
 - fabric topology 5-20
- Client
 - removing B-6
 - uninstalling B-6
- CLONE utility
 - backup 8-2
- cloning
 - backup 8-2
- command console LUN 1-11
 - SCSI-2 mode 1-19
 - SCSI-3 mode 1-18
- Compaq Insight Manager B-7
- comparison of container types 2-15
- Configuration
 - fabric topology
 - procedure flowchart
 - illustrated -xxi
- configuration
 - backup 8-1
 - fabric topology
 - devices 5-15
 - multiple-bus failover
 - cabling 5-9
 - multiple-bus failover using CLI 7-4, 7-11
 - single controller cabling 5-3
 - fibre channel fabric topology
 - using CLI 5-1
 - using SWCC 5-1
 - restoring 2-31
 - rules 2-3
 - configuration options
 - fabric topology
 - adding a disk drive to the spareset 5-21
 - adding disk drives 5-21
 - changing switches
 - device 5-24
 - displaying the current switches 5-22, 5-23
 - initialize 5-24
 - RAIDset and mirrorset 5-23
 - unit 5-24
 - changing switches for a storageset or device 5-23
 - changing the CLI prompt 5-20
 - deleting a storageset 5-22
 - enabling autospare 5-22
 - removing a disk drive from the spareset 5-21
 - configuring
 - pager notification B-7
 - configuring devices
 - fabric topology 5-15
 - connections 1-13
 - hardware 3-17
 - naming 1-13
 - containers
 - attributes 2-14
 - comparison 2-15
 - illustrated 2-14
 - mirrorsets 2-19
 - planning storage 2-14
 - stripesets 2-18
 - controllers
 - cabling 5-4, 5-10
 - location 1-2, 1-4, 1-5
 - node IDs 1-22
 - worldwide names 1-22
 - creating
 - storageset and device profiles 2-15
 - D**
 - data transfer rate 2-30
 - DESTROY 2-31
 - device switches
 - changing
 - fabric topology 5-24
 - devices
 - changing switches

- fabric topology 5-23
- configuration
 - fabric topology 5-15
- creating a profile 2-15
- disk drives
 - adding
 - fabric topology 5-21
 - adding to the spareset
 - fabric topology 5-21
 - array 2-15
 - corresponding storagesets 2-32
 - dividing 2-24
 - removing from the spareset
 - fabric topology 5-21
- displaying the current switches
 - fabric topology 5-22, 5-23
- dividing storagesets 2-24

E

- enabling
 - switches 2-26
- erasing metadata 2-31
- establishing a local connection 5-2

F

- Fabric topology
 - procedure flowchart
 - illustrated -xxi
- fabric topology
 - configuration
 - single controller cabling 5-3
- failover 1-7
 - transparent 1-7
- First enclosure of multiple-enclosure subsystem
 - storage map template 1 A-6, A-9, A-10, A-12, A-14
- first enclosure of multiple-enclosure subsystem
 - storage map
 - template 1 A-3
- functions
 - Agent 4-1

G

- geometry
 - initialize switches 2-32

H

- hardware
 - connections 3-17
 - installation 3-17
- host access
 - restricting by offsets
 - transparent failover 1-22
 - restricting in transparent failover mode 1-20
 - disabling access paths 1-21
 - separate links 1-20
- host adapter
 - installation 3-2
- host connections 1-13
 - naming 1-13

I

- initialize switches
 - changing
 - fabric topology 5-24
 - CHUNKSIZE 2-28
 - DESTROY 2-31
 - geometry 2-32
 - NODESTROY 2-31
 - NOSAVE_CONFIGURATION 2-31
 - SAVE_CONFIGURATION 2-31
- installation
 - hardware 3-17
 - invalid network port assignments B-3
 - there is no disk in the drive message B-4
- installation and configuration script
 - executing 3-9
- installation utility 3-10
- installing
 - Agent 4-5
- integrating SWCC B-7
- invalid network port assignments B-3

J

- JBOD 2-15

Index-4 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

L

LOCATE

find devices 2-43

location

cache module 1-2, 1-4, 1-5

controller 1-2, 1-4, 1-5

LUN IDs

general description 1-25

LUN presentation 1-17

M

maintenance options 3-10

maintenance port connection

establishing a local connection 5-2

illustrated 5-2

mapping storage sets 2-32

messages

there is no disk in the drive B-4

mirrored caching

enabling 1-11

illustrated 1-11

mirrorset switches

changing

fabric topology 5-23

mirrorsets

planning considerations 2-19

important points 2-20

switches 2-27

moving storage sets 8-5

multiple-bus failover

CLI configuration procedure

fabric topology 7-4, 7-11

fabric topology configuration

cabling 5-9

N

network port assignments B-3

node IDs 1-22

restoring 1-23

NODE_ID

worldwide name 1-22

NODESTROY 2-31

NOSAVE_CONFIGURATION 2-31

O

offset

LUN presentation 1-17

restricting host access

transparent failover 1-22

SCSI version factor 1-18

online help

SWCC B-7

options

for mirrorsets 2-27

for RAIDsets 2-27

initialize 2-28

other controller 1-3, 1-5

overwriting data 2-31

P

pager notification B-7

configuring B-7

partitions

assigning a unit number

fabric topology 5-20

defining 2-25

planning considerations 2-24

guidelines 2-25

performance 2-21

planning

overview 2-15

striped mirrorsets 2-24

stripesets 2-18

planning considerations 2-21

planning storage

containers 2-14

where to start 2-2

planning storage sets

characteristics

changing switches 2-26

enabling switches 2-26

initialization switch 2-26

storage switch 2-26

unit switch 2-26

- switches
 - initialization 2-28
 - storage set 2-27
 - unit 2-32
- profiles
 - creating 2-15
 - description 2-15
 - storage set A-1
 - example A-2
- R**
- RAID set switches
 - changing
 - fabric topology 5-23
- RAID sets
 - choosing chunk size 2-28
 - maximum membership 2-22
 - planning considerations 2-21
 - important points 2-22
 - switches 2-27
- read caching
 - enabled for all storage units 1-10
 - general description 1-10
- read requests
 - decreasing the subsystem response time with
 - read caching 1-10
- read-ahead caching 1-10
 - enabled for all disk units 1-10
- reconfiguring
 - Agent 4-9
- removing
 - Agent 4-9
 - Client B-6
- request rate 2-29
- restricting host access
 - disabling access paths
 - transparent failover 1-21
 - separate links
 - transparent failover 1-20
 - transparent failover 1-20
- running
 - Agent 4-9

- S**
- SAVE_CONFIGURATION 2-31
- saving configuration 2-31
- SCSI version
 - offset 1-18
- SCSI-2
 - assigning unit numbers 1-18
 - command console lun 1-19
- SCSI-3
 - assigning unit numbers 1-18
 - command console lun 1-18
- Second enclosure of multiple-enclosure subsystem
 - storage map template 2 A-7
- selective unit presentation 1-19
- SET CONNECTIONS
 - transparent failover 1-16
- setting
 - controller configuration handling 2-31
- single disk (JBOD)
 - assigning a unit number
 - fabric topology 5-19
- Single-enclosure subsystem
 - storage map template 1
 - A-6
- single-enclosure subsystem
 - storage map
 - template 1
 - A-3
- solution software
 - reinstalling 3-11
 - upgrading 3-11
- specifying LUN ID alias
 - SWCC 5-20
- storage
 - creating map 2-32
 - profile
 - example A-2
- storage map 2-32
 - example 2-34
 - template 1 A-3

Index-6 HSG60/HSG80 Array Controller ACS Version 8.5c Solution Software for IBM AIX Installation and Configuration Guide

- first enclosure of multiple-enclosure subsystem A-3
- single enclosure subsystem A-3
- Storage map template 1 A-6
 - first enclosure of multiple-enclosure subsystem A-6, A-9, A-10, A-12, A-14
 - single enclosure subsystem A-6
- Storage map template 2 A-7
 - second enclosure of multiple-enclosure subsystem A-7
- Storage map template 3 A-8
 - third enclosure of multiple-enclosure subsystem A-8
- storageset
 - deleting
 - fabric topology 5-22
 - fabric topology
 - changing switches 5-23
 - planning considerations 2-17
 - mirrorsets 2-19
 - partitions 2-24
 - RAIDsets 2-21
 - striped mirrorsets 2-23
 - stripesets 2-17
 - profile 2-15
 - example 2-16
 - profiles A-1
- storageset profile 2-15
- storageset switches
 - SET command 2-27
- storagesets
 - creating a profile 2-15
 - moving 8-5
- striped mirrorsets
 - planning 2-24
 - planning considerations 2-23
- stripesets
 - distributing members across buses 2-19
 - planning 2-18
 - planning considerations 2-17
 - important points 2-18
- subsystem

- saving configuration 2-31
- subsystem configuration
 - backup 8-1
- SWCC 4-1
 - additional information B-7
 - integrating B-7
 - online help B-7
 - specifying LUN ID alias 5-20
- switches
 - changing 2-26
 - changing characteristics 2-26
 - CHUNKSIZE 2-28
 - DESTROY 2-31
 - enabling 2-26
 - mirrorsets 2-27
 - NODESTROY 2-31
 - NOSAVE_CONFIGURATION 2-31
 - RAIDset 2-27
 - SAVE_CONFIGURATION 2-31
- switches for storagesets
 - overview 2-26

T

- templates
 - subsystem profile A-1
- terminology
 - other controller 1-3, 1-5
 - this controller 1-3, 1-5
- Third enclosure of multiple-enclosure subsystem
 - storage map template 3 A-8
- this controller 1-3, 1-5
- transparent failover 1-7
 - ADD CONNECTIONS 1-16
 - ADD UNIT 1-16
 - matching units to host connections 1-16
 - restricting host access 1-20
 - disabling access paths 1-21
 - separate links 1-20
 - restricting host access by offsets 1-22
 - SET CONNECTIONS 1-16
- troubleshooting
 - invalid network port assignments B-3

there is no disk in the drive message B-4

U

uninstalling

Agent 4-9

Client B-6

unit numbers

assigning 1-16

 fabric topology 5-19

assigning depending on SCSI version 1-18

assigning in fabric topology

 partition 5-20

 single disk 5-19

unit qualifiers

assigning

 fabric topology 5-19

unit switches

changing

 fabric topology 5-24

units

 LUN IDs 1-25

V

virtual disks

 adding B-7

W

worldwide names 1-22

 NODE_ID 1-22

 REPORTED PORT_ID 1-22

 restoring 1-23

write performance 2-31

write requests

 improving the subsystem response time with

 write-back caching 1-10

 placing data with with write-through caching

 1-10

write-back caching

 general description 1-10

write-through caching

 general description 1-10

