MA6000 HSG60 Array Controller
ACS Version 8.5 Solution Software
for Compaq Tru64 UNIX

Installation and Configuration Guide

Notice

Table of Contents

About This Guide

Chapter 1

Planning a Subsystem

Chapter 2

Planning Storage

Chapter 3
Preparing the Host System

Chapter 4
Installing and Configuring the HS-Series Agent

Chapter 5
Configuration Procedures

Chapter 6
Verifying Storage Configuration from the Host

Chapter 7
Configuration Example Using CLI

Chapter 8
Backing Up the Subsystem, Cloning Data for Backup, and Moving Storagesets

Appendix A
Subsystem Profile Templates

Appendix B
Installing, Configuring, and Removing the Client

Index

Figures

Tables

About This Guide

Conventions

Text Conventions

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

Convention	Meaning
.	Indicates that a portion of an example or figure has been omitted.
"this controller"	The controller serving the current CLI session through a local or remote terminal.
"other controller"	The controller in a dual-redundant pair that's connected to the controller serving the current CLI session.

Table 1 Text Conventions (Continued)

Special Notices

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

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

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

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

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

Related Publications

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

Table 2 Relevant Publications

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

Revision History

First Release: May 2000

Chapter **1**

Planning a Subsystem

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

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

The following information is included in this chapter:

■ "Controller Designations A and B," page 1–1

■ "Selecting a Failover Mode," page 1–3

■ "Selecting a Cache Mode," page 1–9

■ "Enabling Mirrored Caching," page 1–10

■ "The Command Console LUN," page 1–11

■ "Connections," page 1–12

■ "Assigning Unit Numbers," page 1–17

■ "Restricting Host Access (Selective Storage Presentation)," page 1–21

■ "Worldwide Names (Node IDs and Port IDs)," page 1–28

Controller Designations A and B

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

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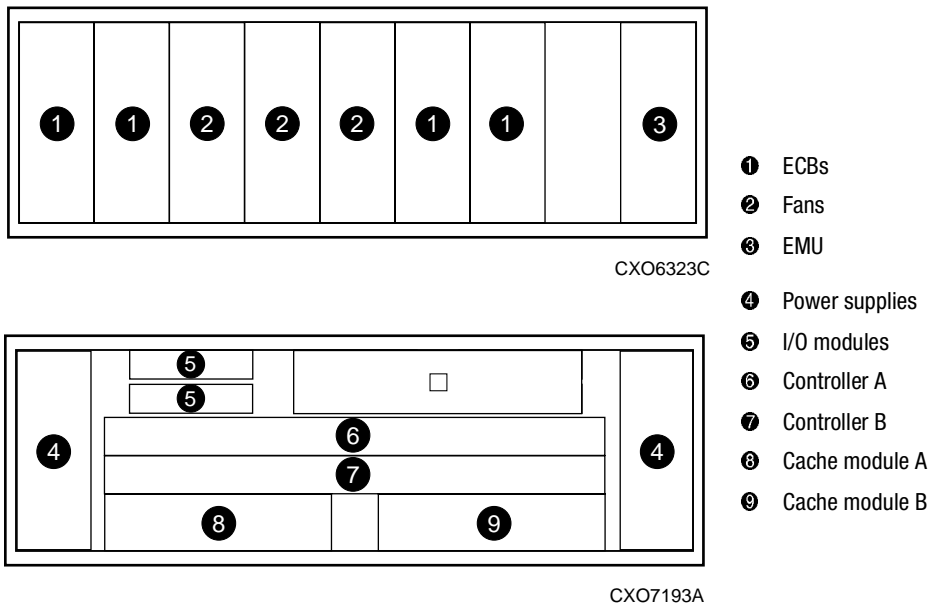
Figure 1–1. Location of controllers and cache modules

Controller Designations "This Controller" and "Other Controller"

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

■ "**this controller**"—the controller that is the focus of the CLI session. "This controller" is the controller to which the maintenance terminal is attached and through which the CLI commands are being entered. "This controller" can be shortened to "this" in CLI commands.

■ "**other controller**"—the controller that is not the focus of the CLI session and through which CLI commands are not being entered. "Other controller" can be shortened to "other" in CLI commands.

Figure 1–2 shows the relationship between "this controller" and "other controller."

CXO7194A

❶ This controller   ❷ Other controller
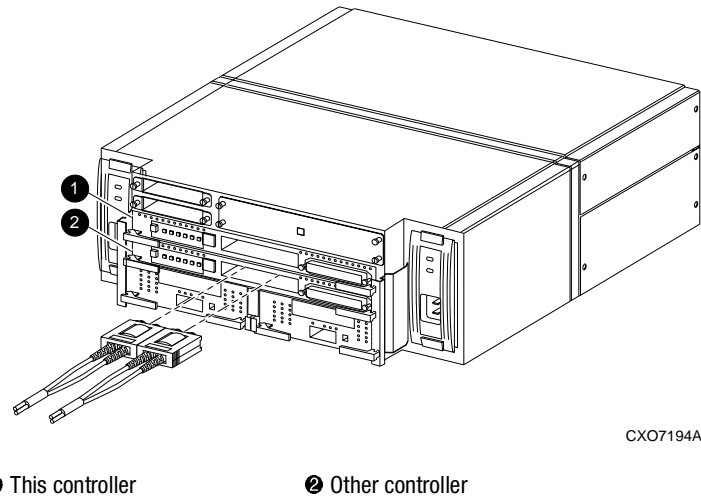
Figure 1–2. "This controller" and "other controller"

Selecting a Failover Mode

Failover is a way to keep the storage array available to the host if one controllers becomes unresponsive. A controller can become unresponsive because of a hardware failure such as a host bus adapter or controller or, in multiple-bus only, to a failure of the link between host and controller. Failover keeps the storage array available to the hosts by allowing the surviving controller to take over total control of the subsystem.

There are two failover modes:

■ **transparent**, which is handled by the surviving controller and is transparent (invisible) to the hosts

■ **multiple-bus**, which is handled by the hosts

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

Transparent Failover Mode

Transparent failover mode has the following characteristics:

■ hosts do not know failover has taken place

■ units are divided between host ports 1 and 2

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

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

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

Units are divided between the host ports:

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

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

The limitation of transparent failover is that it only compensates for a controller failure, and not for failures of either the Fibre Channel link or host Fibre Channel adapters.
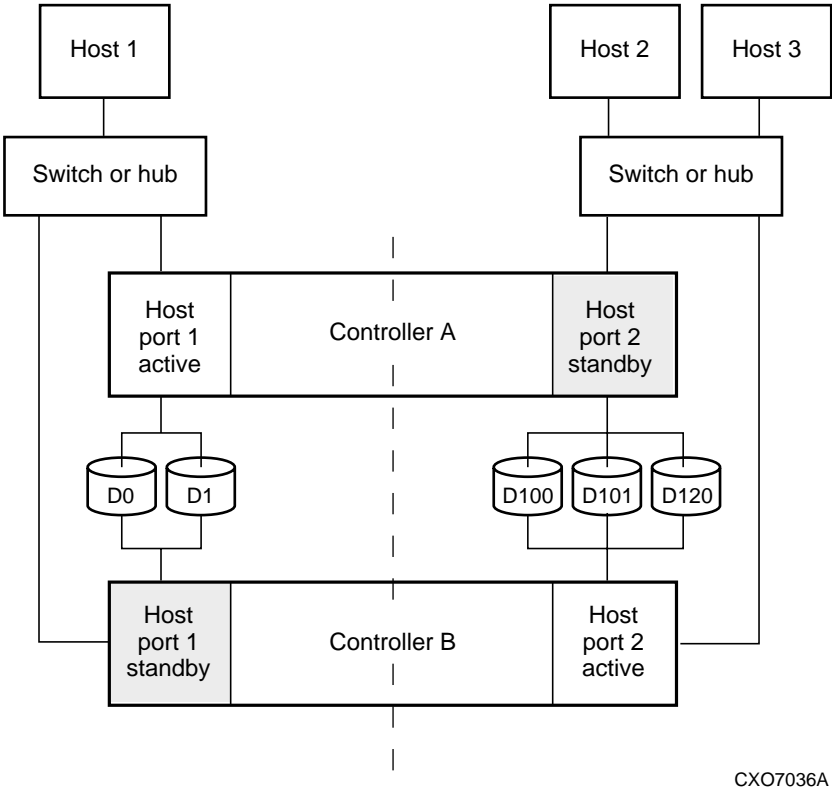
CXO7036A

Figure 1–3. Transparent failover - normal operation
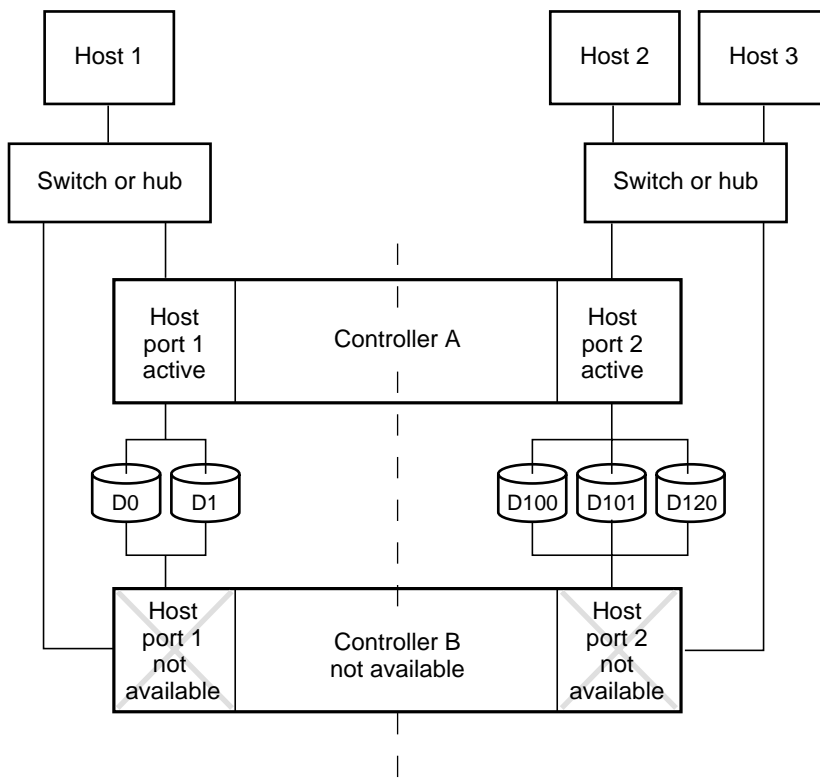
CXO7035A

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

Multiple-Bus Failover Mode

Multiple-bus failover mode has the following characteristics:

■ host controls the failover process by moving the unit(s) from one controller to another

■ all units (0 through 199) are visible at all host ports

■ each host has two or more paths to the units

Each host must have special software to control failover. With this software, the host sees the same units visible through two (or more) paths. When one path fails, the host can issue commands to move the units from one path to another. A typical multiple-bus failover configuration is shown in Figure 1–5.

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

    SET D101 PREFERRED_PATH=THIS_CONTROLLER

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

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

■ multiple-bus failover can compensate for a failure in any of the following:

  ❏ controller

  ❏ switch or hub

  ❏ Fibre Channel link

  ❏ host Fibre Channel adapter

■ a host can re-distribute the I/O load between the controllers

■ all hosts must have operating-system software that supports multiple-bus failover mode
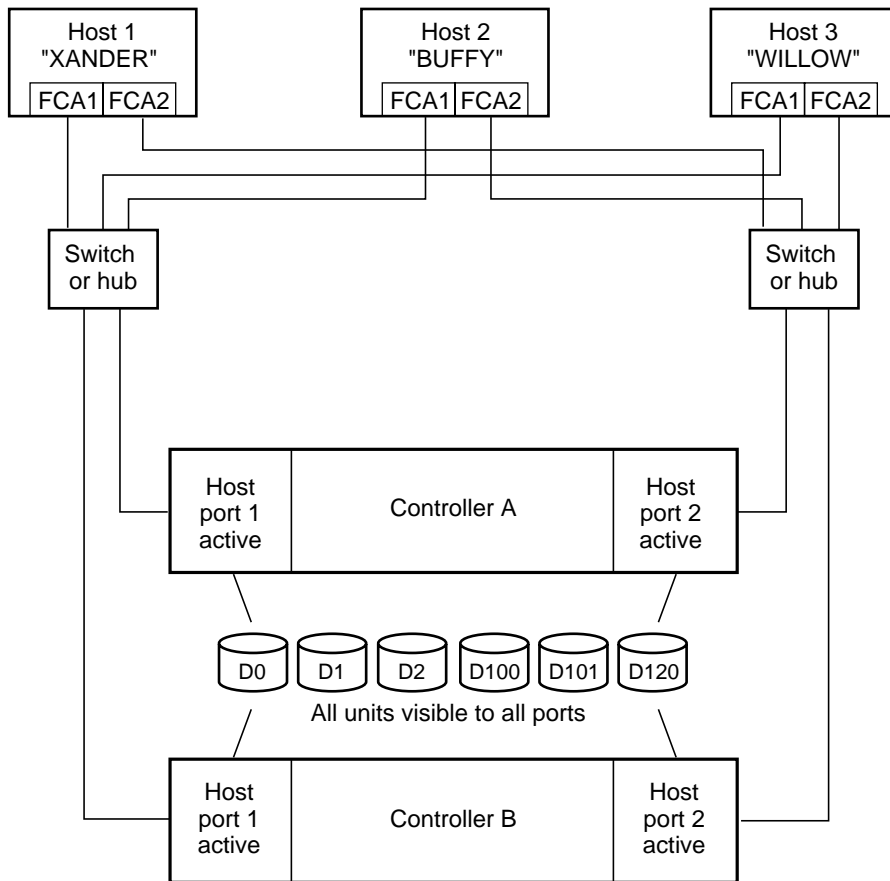
NOTE: FCA = Fibre Channel Adapter

CXO7094A

Figure 1–5. Typical multiple-bus configuration

Selecting a Cache Mode

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

Read Caching

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

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

Read-Ahead Caching

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

Write-Back Caching

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

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

Write-Through Caching

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

Enabling Mirrored Caching

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

The total memory available for cached data is reduced by half, but the level of protection is greater.
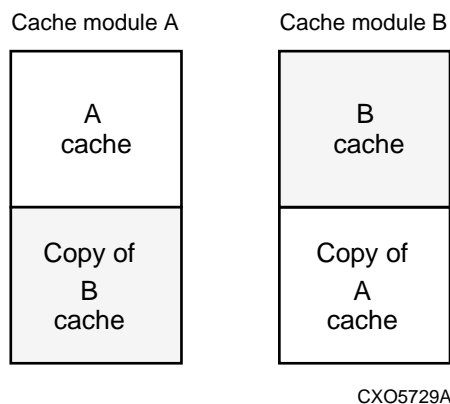
Figure 1–6. Mirrored caching

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

■ both controllers support the same size cache

■ diagnostics indicates that both caches are good

■ no unit errors are outstanding, for example, lost data or data that cannot be written to devices

■ both controllers are started and configured in failover mode

The Command Console LUN

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

■ Allows the RAID Array to be recognized by the host as soon as it is attached to the SCSI bus and configured into the operating system.

■ Serves as a communications device for the HS-Series Agent. The CCL identifies itself to the host by a unique identification string.

In general, you will need to create a storageset to be used as a CCL for Port 2. This is a requirement for the HS-Series Agent.

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

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

Determining the Address of the CCL

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

HSG60 > **SHOW THIS_CONTROLLER**

In SCSI II Mode:

Disabling the CCL

To disable the CCL, enter the following CLI command:

HSG60 > **SET THIS_CONTROLLER NOCOMMAND_CONSOLE_LUN**

Enabling the CCL

To enable the CCL, enter the following CLI command:

> HSG60 > **SET THIS_CONTROLLER COMMAND_CONSOLE_LUN**

To see the state of the CCL, use the SHOW "this controller"/ "other controller" command. Because the CCL is not an actual LUN, the SHOW UNITS command will not display the CCL location.

In SCSI III Mode:

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

Connections

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

Naming Connections

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

HOST1A1

HOST NAME

PORT
CONTROLLER
ADAPTER
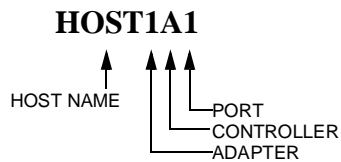
Examples:

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

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

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

Numbers of Connections

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

■ If a controller pair is in transparent failover mode and the port 1 link is separate from the port 2 link (that is, ports 1 of both controllers are on one loop or fabric, and port 2 of both controllers are on another), each adapter will have one connection, as shown in Figure 1–7.

■ If a controller pair is in transparent failover mode and port 1 and port 2 are on the same link (that is, all ports are on the same loop or fabric), each adapter will have two connections, as shown in Figure 1–8.

■ If a controller pair is in multiple-bus failover mode, each adapter has two connections, as shown in Figure 1–9.
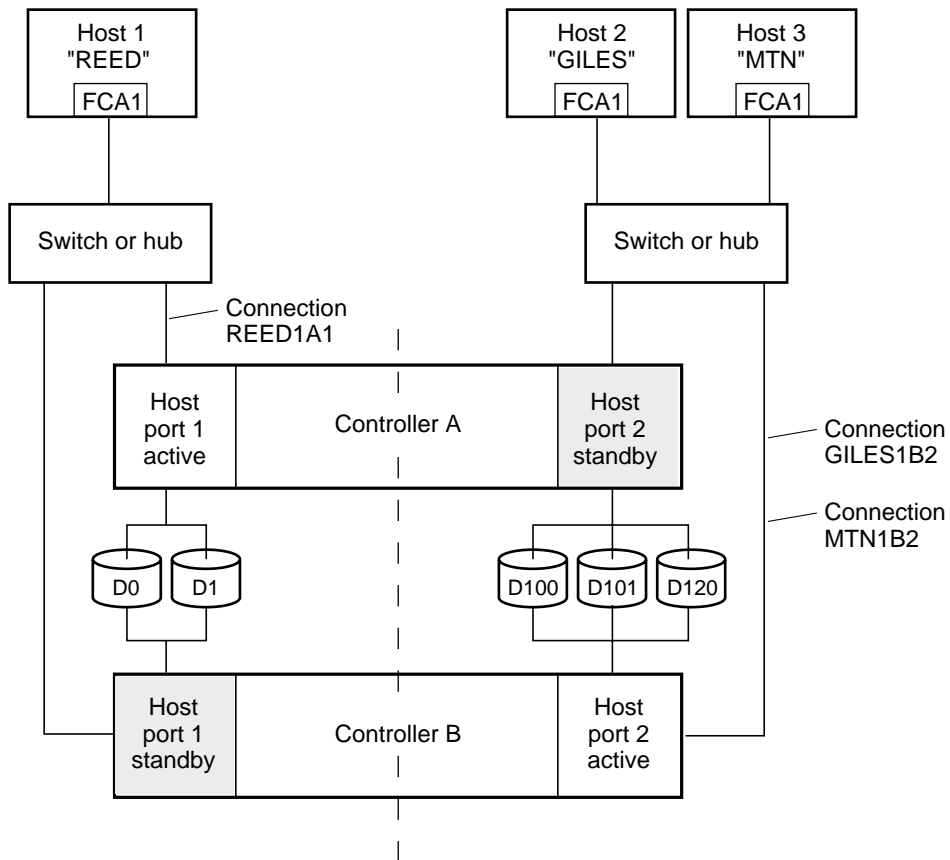
NOTE: FCA = Fibre Channel Adapter

CXO7081A

Figure 1–7. Connections in separate-link, transparent failover mode configurations
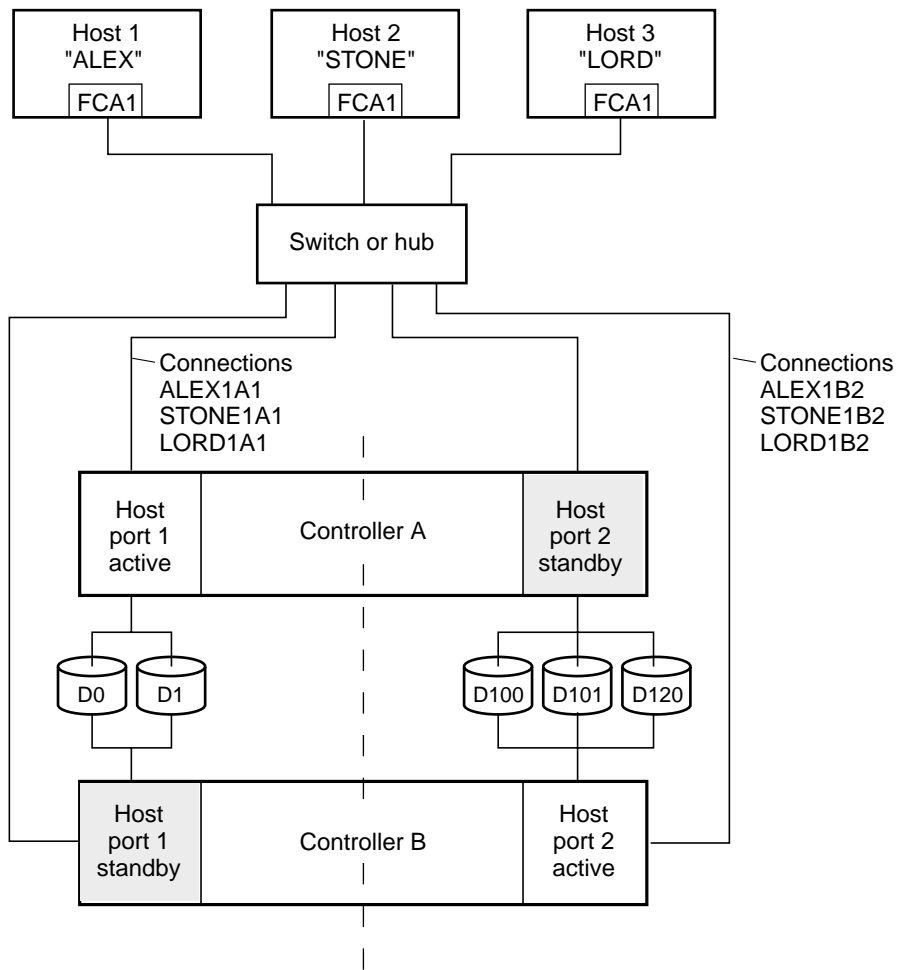
NOTE: FCA = Fibre Channel Adapter

CXO7079A

Figure 1–8. Connections in single-link, transparent failover mode configurations
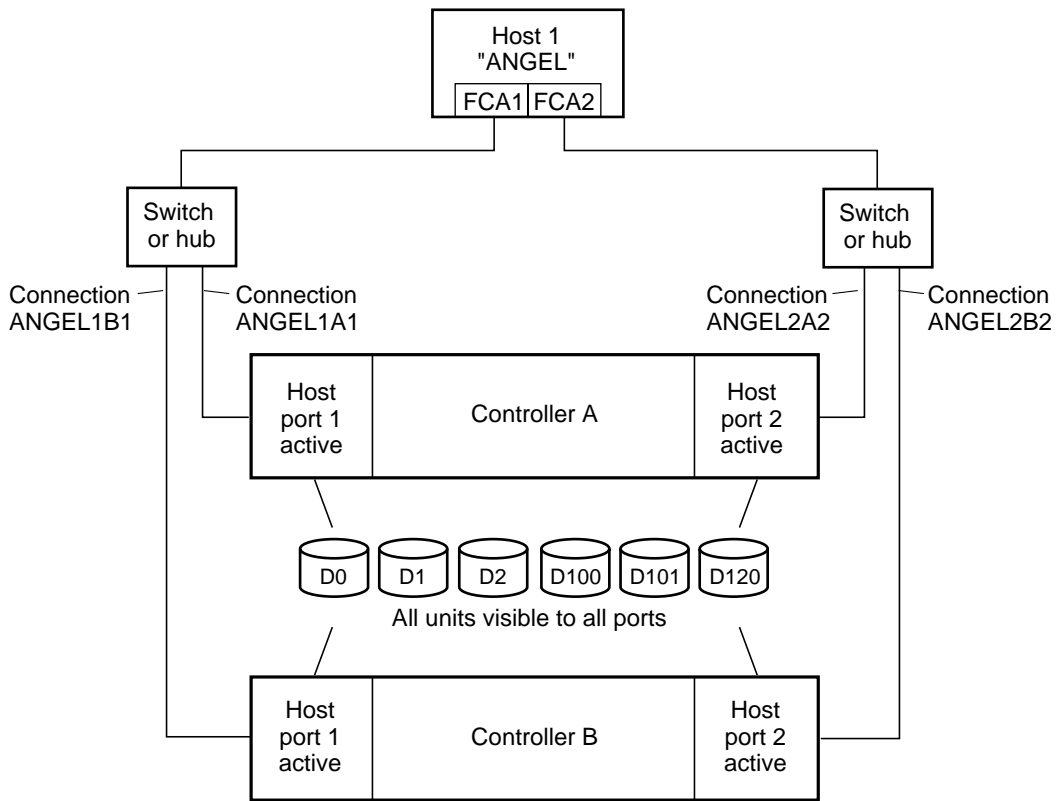
NOTE: FCA = Fibre Channel Adapter

CXO7080A

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

Assigning Unit Numbers

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

■ failover mode of the controller pair

■ the *ENABLE_ACCESS_PATH* and *PREFERRED_PATH* switches in the ADD UNIT (or SET *unit*) commands

■ the *UNIT_OFFSET* switch in the ADD CONNECTIONS (or SET *connections*) commands

■ which controller port the connection is attached to

■ the *SCSI_VERSION* switch of the SET *this_controller/other_controller* command

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

Matching Units to Host Connections in Transparent Failover Mode

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

Unit numbers are assigned to ports as follows:

■ 0 - 99 are assigned to host port 1 of both controllers

■ 100 - 199 are assigned to host port 2 of both controllers

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

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

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 1–10 shows how units are presented as different LUNs, depending on the offset of the host. In this illustration, host connection 1 and host connection 2 would need to be on host port 1; host connection 3 would need to be on host port 2.

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

CXO6455B

Figure 1–10. LUN presentation to hosts, as determined by offset

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

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

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

In addition, the access path to the host connection must be enabled for the connection to access the unit. See "Restricting Host Access in Transparent Failover Mode," page 1–21.

Matching Units to Host Connections in Multiple-bus Failover Mode

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

What LUN number a host connection assigns to a unit is a function of the *UNIT_OFFSET* switch of the ADD (or SET) CONNECTIONS command. The default offset is 0. The relationship of offset, LUN number, and unit number is as follows:

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)

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

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

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

Assigning Unit Numbers Depending on *SCSI_VERSION*

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

The CCL in SCSI-3 mode

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

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

■ Offsets should be divisible by 10 (for consistency and simplicity)

■ Unit numbers should not be assigned at connection offsets (to avoid being masked by the CCL at LUN 0)

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

LUN 0 - CCL

LUN 1 - unit 21

LUN 2 - unit 22, etc.

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

The CCL in SCSI-2 Mode

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

If *SCSI_VERSION* is set to SCSI-2 mode, the CCL floats, moving to the first available LUN location, depending on the configuration.

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

■ Offsets should be divisible by 10 (for consistency and simplicity)

■ Unit numbers should be assigned at connection offsets (so that every host connection has a unit presented at LUN 0)

Table 1-2 summarizes the recommendations for unit assignments based on the *SCSI_VERSION* switch.

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

Restricting Host Access (Selective Storage Presentation)

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

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

Restricting Host Access in Transparent Failover Mode

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

1. using separate Fibre Channel links (either loop or fabric)

2. enabling the access path of selected host connections on a shared loop or fabric

3. setting offsets

> **NOTE:** These techniques also work for a single controller.

Restricting Host Access by Separate Links

In transparent failover mode, host port 1 of controller A and host port 1 of controller B share a common Fibre Channel link. Host port 2 of controller A and host port 2 of controller B also share a common Fibre Channel link. If the host 1 link is separate from the host 2 link, the simplest way of limiting host access is to have one host or set of hosts on the port 1 link, and another host or set of hosts on the port 2 link. Each host can then see only units assigned to its respective controller port. This separation of host buses is shown in Figure 1–11. This method applies only if the host 1 link and host 2 link are separate links.

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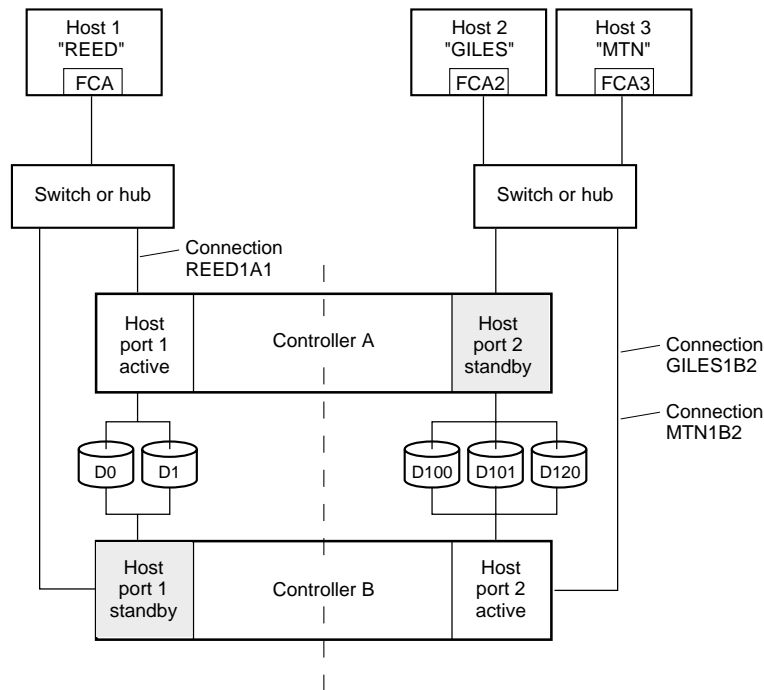
NOTE: FCA = Fibre Channel Adapter

CXO7081A

Figure 1–11. Limiting host access in transparent failover mode

Restricting Host Access by Disabling Access Paths

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

Example:

In Figure 1–11, restricting the access of unit D101 to host 3, the host named MTN, can be done by enabling only the connection to host 3. Enter the following commands:

    SET D101 DISABLE_ACCESS_PATH=ALL

    SET D101 ENABLE_ACCESS_PATH=MTN1B2

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

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

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

Restricting Host Access by Offsets

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

Example:

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

    SET MTN1B2 UNIT_OFFSET=120

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

SET D120 DISABLE_ACCESS_PATH=ALL

SET D120 ENABLE_ACCESS_PATH=MTN1B2

> **NOTE:** It is highly recommended that you still give access to only specific connections, even if there is just one connection on the link. That way, if new connections are added, they will not have automatic access to all units.

Restricting Host Access in Multiple-Bus Failover Mode

In multiple-bus mode, the units assigned to any port are visible to all ports.

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

■ Enabling the access path of selected host connections

■ Setting offsets

Restricting Host Access by Disabling Access Paths

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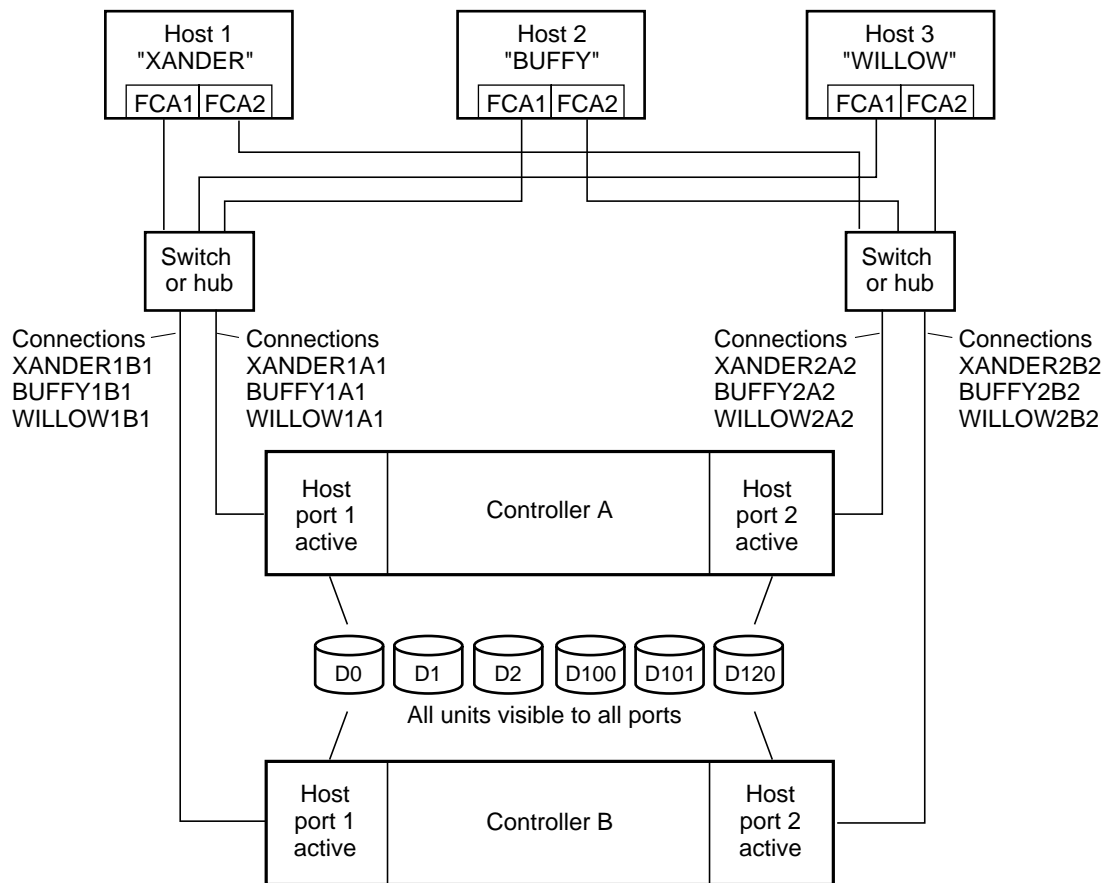
NOTE: FCA = Fibre Channel Adapter

CXO7078A

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

Example:

Figure 1–12 shows a representative multiple-bus failover configuration. Restricting the access of unit D101 to host WILLOW can be done by enabling only the connections to host WILLOW. At least two connections must be enabled for multiple-bus failover to work. For most operating systems, it is desirable to have all connections to the host enabled. To enable all connections for host WILLOW, enter the following commands:

SET D101 DISABLE_ACCESS_PATH=ALL

SET D101 ENABLE_ACCESS_PATH=(WILLOW1A1,WILLOW1B1,WILLOW2A2,WILLOW2B2)

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

SET D101 DISABLE_ACCESS_PATH=ALL

SET D101 ENABLE_ACCESS_PATH=(WILLOW1A1,WILLOW2B2)

> or

SET D101 DISABLE_ACCESS_PATH=ALL

SET D101 ENABLE_ACCESS_PATH=(WILLOW1B1,WILLOW2A2)

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

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

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

Restricting Host Access by Offsets

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

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

Example:

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

    SET WILLOW1A1 UNIT_OFFSET=120

    SET WILLOW1B1 UNIT_OFFSET=120

    SET WILLOW2A2 UNIT_OFFSET=120

    SET WILLOW2B2 UNIT_OFFSET=120

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

    SET D102 DISABLE_ACCESS_PATH=ALL

    SET D102 ENABLE_ACCESS_PATH=(WILLOW1A1,WILLOW1B1,WILLOW12A2,WILLOW2B2)

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

Worldwide Names (Node IDs and Port IDs)

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

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

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

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

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

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

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

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

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

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

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

Restoring Worldwide Names (Node IDs)

If a situation occurs that requires you to restore the worldwide name, you can restore it using the worldwide name and check sum printed on the sticker on the frame into which the controller is inserted, as shown in Figure 1–13.

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

Unit World Wide Names (LUN IDs)

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

❶ Node ID (Worldwide name)

❷ Checksum
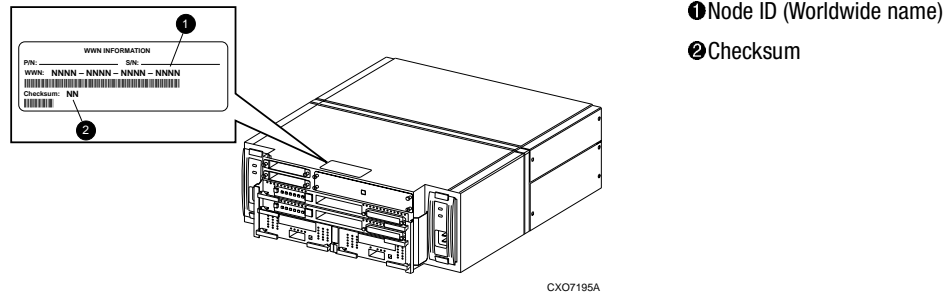
Figure 1–13. Placement of the worldwide name label

Chapter *2*

Planning Storage

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

The following information is included in this chapter:

- "Where to Start," page 2–2
- "Configuration Rules," page 2–3
- "Determining Storage Requirements," page 2–7
- "Choosing a Container Type," page 2–7
- "Creating a Storageset Profile," page 2–10
- "Storageset Planning Considerations," page 2–12
- "Mirrorset Planning Considerations," page 2–15
- "Partition Planning Considerations," page 2–20
- "Changing Characteristics through Switches," page 2–21
- "Storageset and Partition Switches," page 2–22
- "Initialization Switches," page 2–24
- "Unit Switches," page 2–29
- "Storage Maps," page 2–29

Where to Start

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

1. Review configuration rules. See "Configuration Rules," page 2–3.

2. Familiarize yourself with the current physical layout of the devices and their addressing scheme. See "Device PTL Addressing Convention," page 2–4.

3. Determine your storage requirements. Use the questions in "Determining Storage Requirements," page 2–7 to help you.

4. Choose the type of storage container(s) you need to use in your subsystem. See "Choosing a Container Type," page 2–7 for a comparison and description of each type of storageset.

5. Create a storageset profile (described in "Creating a Storageset Profile," page 2–10). Fill out the storageset profile while you read the sections that pertain to your chosen storage type:

   ■ "Storageset Planning Considerations," page 2–12

   ■ "Mirrorset Planning Considerations," page 2–15

   ■ "Partition Planning Considerations," page 2–20

   ■ "Striped Mirrorset Planning Considerations," page 2–18

6. Decide which switches you will need for your subsystem. General information on switches is detailed in "Storageset and Partition Switches," page 2–22

   ■ Determine what unit switches you want for your units ("Unit Switches," page 2–29)

   ■ Determine what initialization switches you want for your planned storage containers ("Initialization Switches," page 2–24)

7. Create a storage map ("Storage Maps," page 2–29).

8. Configure the storage you have now planned using one of the following methods:

   ■ Use SWCC. See the SWCC documentation for details.

   ■ Use the Command Line Interpreter (CLI) commands. This method allows you flexibility in defining and naming your storage containers. See the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

Configuration Rules

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

■ Maximum 128 assignable, host-accessible units

■ Maximum 512-GB unit capacity

■ Maximum 8 RAID-5 storagesets

■ Maximum 30 RAID-5 and RAID-1 storagesets

■ Maximum 45 RAID-5, RAID-1, and RAID-0 storagesets

■ Maximum 8 partitions per storageset or individual disk

■ Maximum 6 members per mirrorset

■ Maximum 14 members per RAID-5 storageset

■ Maximum 24 members per stripeset

■ Maximum 48 physical devices per striped mirrorset

Device PTL Addressing Convention

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

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

■ P—Designates the controller's SCSI device port number (1 or 2).

■ T—Designates the target identification (ID) number of the device. Valid target ID numbers for a single-controller configuration and dual-redundant controller configuration are 0 - 5 and 8 - 13.

■ L—Designates the logical unit (LUN) of the device. For disk devices the LUN is always 0.

1 02 00

LUN 00
Target 02
Port 1
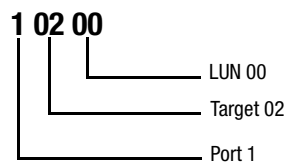
Figure 2–1. PTL naming convention

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

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

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

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

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

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

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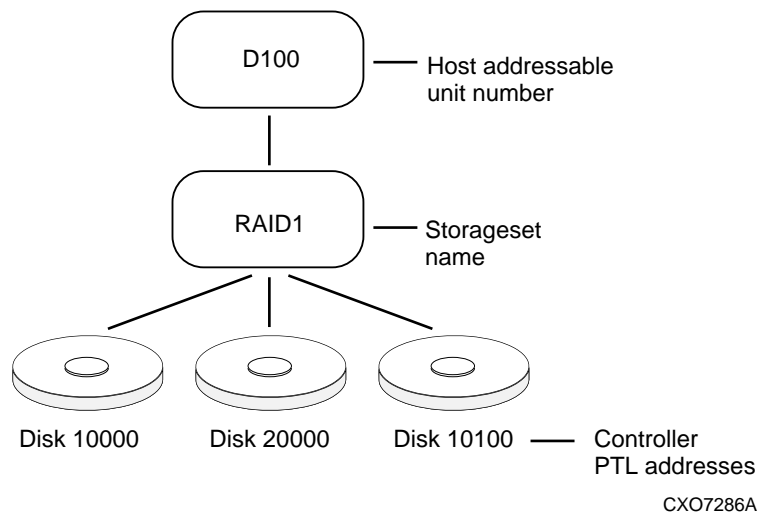
Figure 2–4. Mapping a unit to physical disk drives

Determining Storage Requirements

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

■ What applications or user groups will access the subsystem? How much capacity do they need?

■ What are the I/O requirements? If an application is data-transfer intensive, what is the required transfer rate? If it is I/O-request intensive, what is the required response time? What is the read/write ratio for a typical request?

■ Are most I/O requests directed to a small percentage of the disk drives? Do you want to keep it that way or balance the I/O load?

■ Do you store mission-critical data? Is availability the highest priority or would standard backup procedures suffice?

Choosing a Container Type

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

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

The storagesets implement RAID (Redundant Array of Independent Disks) technology. Consequently, they all share one important feature: each storageset, whether it contains two disk drives or ten, looks like one large, virtual disk drive to the host.
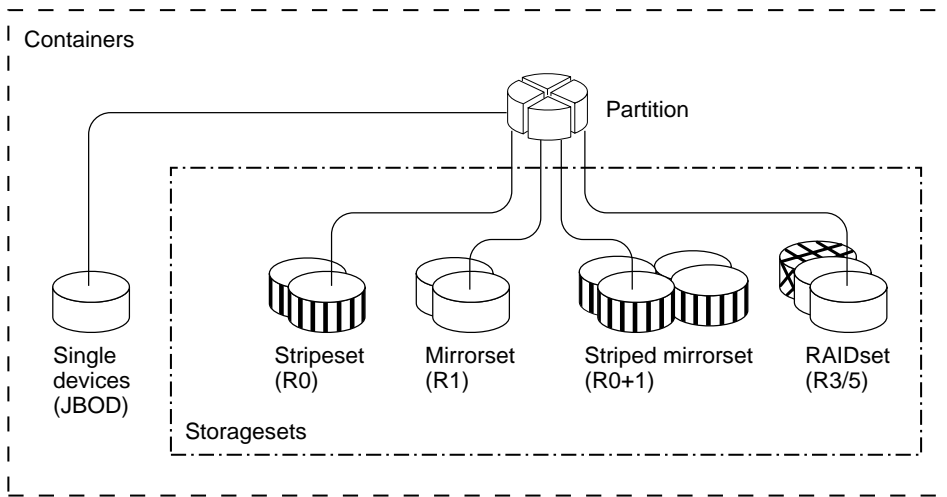
CXO6677A

Figure 2–5. Container types

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

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 (RAID1)	Excellent	Good/Fair	Good/Fair	System drives; critical files
RAIDset (RAID 3/5)	Excellent	Excellent/good	Read: excellent (if used with small chunk sizes) Write: good (if used with small chunk sizes)	High request rates, read-intensive, data lookup
Striped Mirrorset (RAID 0+1)	Excellent	Excellent if used with large chunk size	Excellent if used with small chunk size	Any critical response-time application

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

Creating a 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 2–6.

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

Type of Storageset:

_____ Mirrorset	__X_ RAIDset	_____ Stripeset	_____ Striped Mirrorset ____ JBOD

Storageset Name   *R1*

Disk Drives   *D10300, D20300, D10400, D20400*

Unit Number *D101*

Partitions:

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

RAIDset Switches:

Reconstruction Policy	Reduced Membership	Replacement Policy
_X_ Normal (default)	_X_ No (default)	_X_ Best performance (default)
___Fast	___Yes, missing:	___Best fit
		___None

Mirrorset Switches:

Replacement Policy	Copy Policy	Read Source
___Best performance (default)	___Normal (default)	___Least busy (default)
___Best fit	___Fast	___Round robin
___None		___Disk drive:

Initialize Switches:

Chunk size	Save Configuration	Metadata
_X_ Automatic (default)	___No (default)	_X_Destroy (default)
___ 64 blocks	_X_ Yes	___Retain
___ 128 blocks		
___ 256 blocks		
___ Other:		

Unit Switches:

Caching	Access by following hosts enabled
Read caching_____X__	_ALL_____
Read-ahead caching_____	_____
Write-back caching___X__	_____
Write-through caching____	_____

Figure 2–6. An example storageset profile

Storageset Planning Considerations

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

■ Stripeset Planning Considerations

■ Mirrorset Planning Considerations

■ Partition Planning Considerations

■ RAIDset Planning Considerations

■ Striped Mirrorset Planning Considerations

Stripeset Planning Considerations

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

For example, in a three-member stripeset that contains disk drives Disk10000, Disk20000, and Disk10100, the first chunk of an I/O request is written to Disk10000, the second to Disk20000, the third to Disk10100, the fourth to Disk10000, and so forth until all of the data has been written to the drives (Figure 2–7).
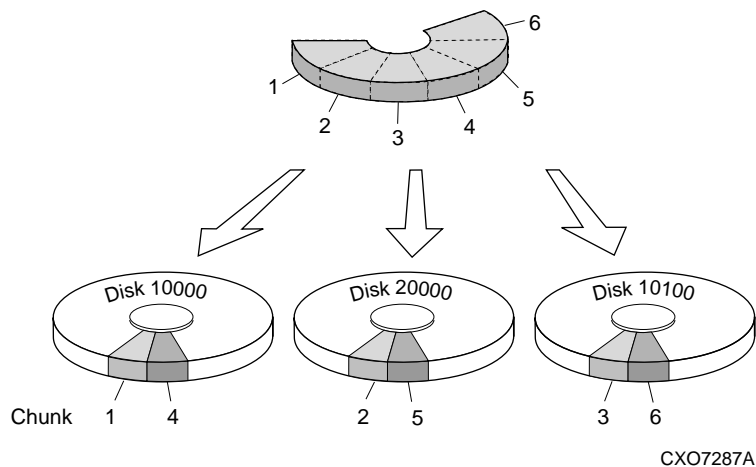
Figure 2–7. A 3-member RAID 0 stripeset (example 1)

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

A major benefit of striping is that it balances the I/O load across all of the disk drives in the storageset. This can increase the subsystem performance by eliminating the hot spots (high localities of reference), that occur when frequently accessed data becomes concentrated on a single disk drive.
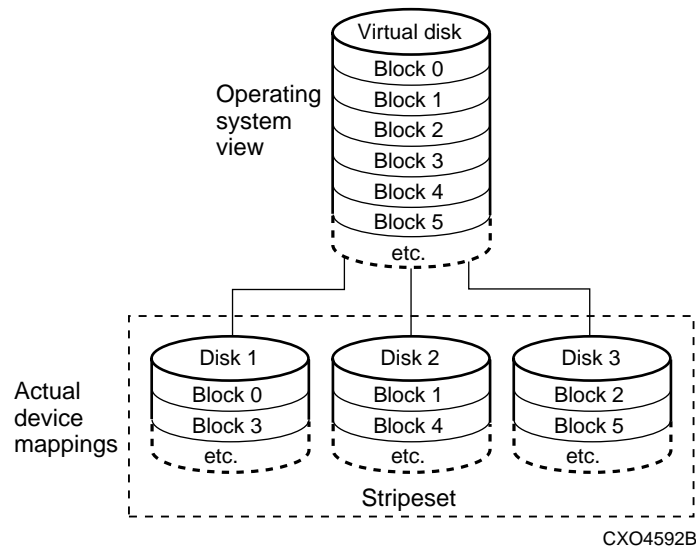
Figure 2–8. A 3-member RAID 0 stripeset (example 2)

Keep the following points in mind as you plan your stripesets:

■ Reporting methods and size limitations prevent certain operating systems from working with large stripesets.

■ A storageset should only contain disk drives of the same capacity. The controller limits the effective capacity of each member to the capacity of the smallest member in the storageset (base member size) when the storageset is initialized. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you will waste 5 GB of capacity on each 9 GB member.

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

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

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

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

■ Evenly distribute the members across the device ports to balance load and provide multiple paths.

■ Stripesets may contain between two and 24 members.

■ Stripesets are well-suited for the following applications:

❏ Storing program image libraries or run-time libraries for rapid loading.

❏ Storing large tables or other structures of read-only data for rapid application access.

❏ Collecting data from external sources at very high data transfer rates.

■ Stripesets are not well-suited for the following applications:

❏ A storage solution for data that cannot be easily reproduced or for data that must be available for system operation.

❏ Applications that make requests for small amounts of sequentially located data.

❏ Applications that make synchronous random requests for small amounts of data.

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

Mirrorset Planning Considerations

Mirrorsets (RAID 1) use redundancy to ensure availability, as illustrated in Figure 2–9. For each primary disk drive, there is at least one mirror disk drive. Thus, if a primary disk drive fails, it's mirror drive immediately provides an exact copy of the data. Figure 2–10 shows a second example of a Mirrorset.
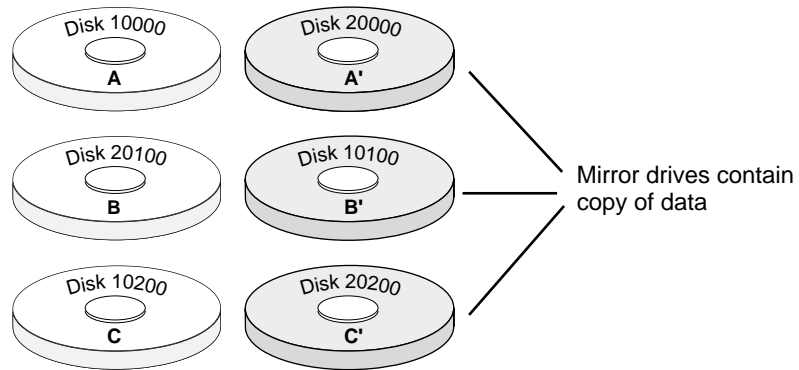
CXO7288A

Figure 2–9. Mirrorsets maintain two copies of the same data
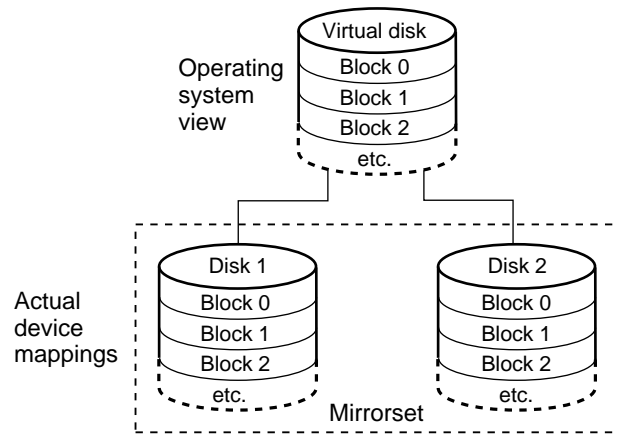
CXO4594B

Figure 2–10. Mirrorset example 2

Keep these points in mind as you plan your mirrorsets:

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

■ You can configure up to 20 mirrorsets per controller or pair of dual-redundant controllers. Each mirrorset may contain up to 6 members.

■ Both write-back cache modules must be the same size.

■ A mirrorset should only contain disk drives of the same capacity.

■ Mirrorsets are well-suited for the following:

❏ Any data for which reliability requirements are extremely high

❏ Data to which high-performance access is required

❏ Applications for which cost is a secondary issue

■ Mirrorsets are not well-suited for the following applications:

❏ Write-intensive applications (a performance hit of 10% will occur)

❏ Applications for which cost is a primary issue

RAIDset Planning Considerations

RAIDsets (RAID 3/5) are enhanced stripesets—they use striping to increase I/O performance and distributed-parity data to ensure data availability. Figure 2–11 shows an example of a RAIDset that uses five members.
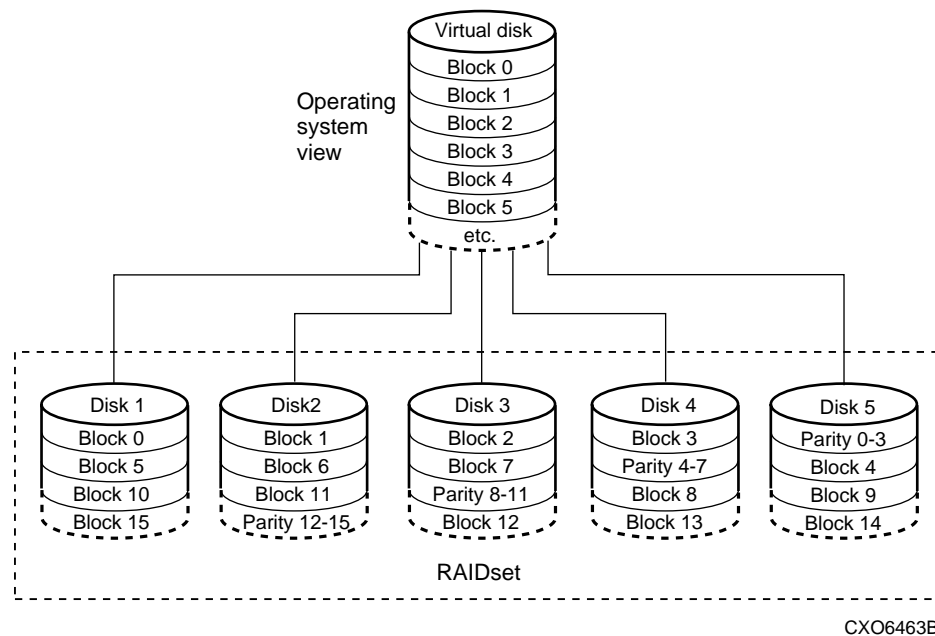
CXO6463B

Figure 2–11. A five member RAIDset using parity

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

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

Keep these points in mind as you plan your RAIDsets:

■ Reporting methods and size limitations prevent certain operating systems from working with large RAIDsets.

■ Both cache modules must be the same size.

■ A RAIDset must include at least 3 disk drives, but no more than 14.

- A storageset should only contain disk drives of the same capacity. The controller limits the capacity of each member to the capacity of the smallest member in the storageset. Thus, if you combine 9 GB disk drives with 4 GB disk drives in the same storageset, you will waste 5 GB of capacity on each 9 GB member.

- RAIDsets are particularly well-suited for the following:

  - ❏ Small to medium I/O requests

  - ❏ Applications requiring high availability

  - ❏ High read request rates

  - ❏ Inquiry-type transaction processing

- RAIDsets are not particularly well-suited for the following:

  - ❏ Write-intensive applications

  - ❏ Database applications in which fields are continually updated

  - ❏ Transaction processing

Striped Mirrorset Planning Considerations

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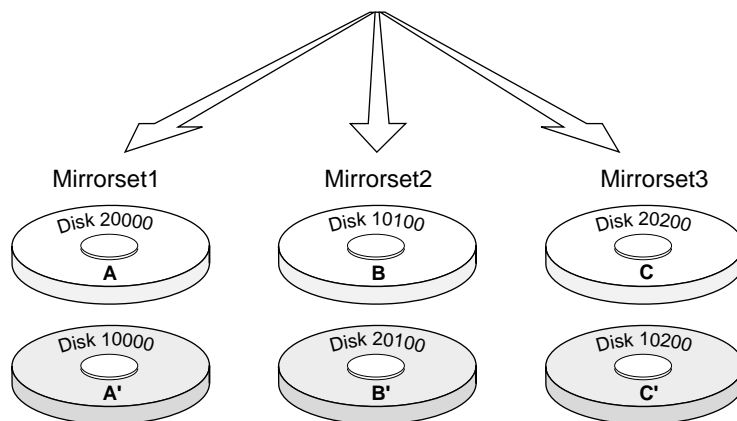
Figure 2–12. Striped mirrorset (example 1)

The failure of a single disk drive has no effect on the ability of the storageset to deliver data to the host. Under normal circumstances, it also has very little effect on performance. Because striped mirrorsets do not require any more disk drives than mirrorsets, this storageset is an excellent choice for data that warrants mirroring.
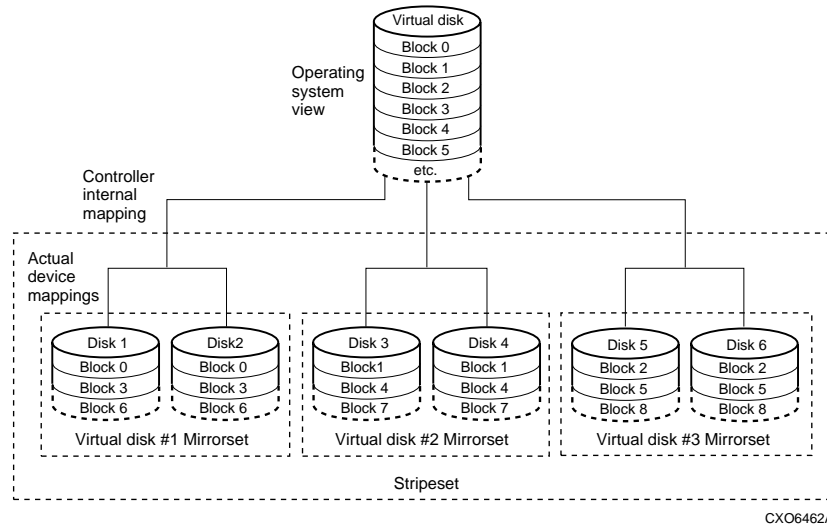
Figure 2–13. Striped mirrorset (example 2)

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

Partition Planning Considerations

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

CXO7056A
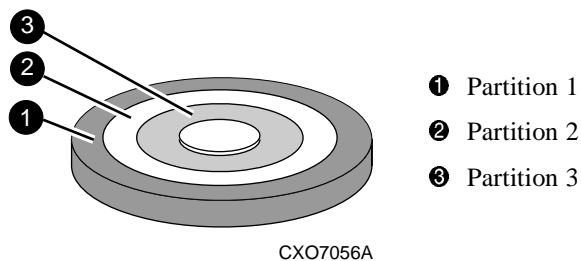
Figure 2–14. One example of a partitioned single-disk unit

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

Defining a Partition

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

■ Mirrorsets and single disk units—the controller allocates the largest whole number of blocks that are equal to or less than the percentage you specify.

■ RAIDsets and stripesets—the controller allocates the largest whole number of stripes that are less than or equal to the percentage you specify.

   ❑ Stripesets—the stripe size = chunk size × number of members.

   ❑ RAIDsets—the stripe size = chunk size × (number of members minus 1)

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

Guidelines for Partitioning Storagesets and Disk Drives

Keep these points in mind 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.

■ In multiple-bus failover mode, all the partitions of a particular container must be on the same controller. Partitions cannot be split across controllers.

■ Partitions cannot be combined into storagesets. For example, you cannot divide a disk drive into three partitions, then combine those partitions into a RAIDset.

■ Just as with storagesets, you do not have to assign unit numbers to partitions until you are ready to use them.

■ The CLONE utility cannot be used with partitioned mirrorsets or partitioned stripesets.

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 ≤ 9 then chunk size = 256. If d > 9 then chunk size = 128.

■ *CHUNKSIZE=n* lets you specify a chunk size in blocks. The relationship between chunk size and request size determines whether striping increases the request rate or the data-transfer rate.

Increasing the Request Rate

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

Request A   Chunk size = 128k (256 blocks)

Request B
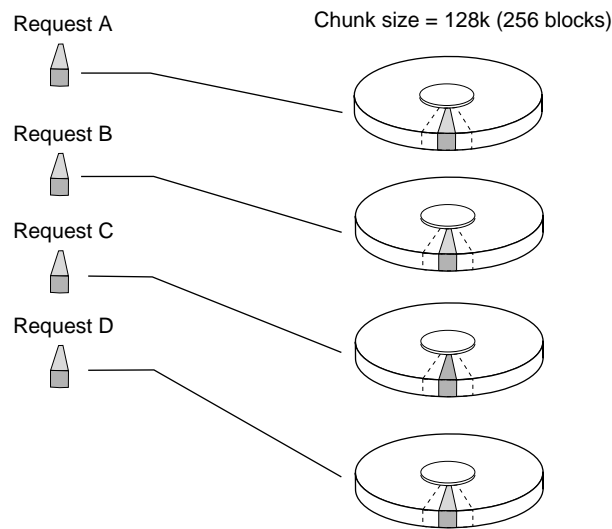
Request C

Request D

CXO-5135A-MC

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

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

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

- Many parallel I/Os that use a small area of disk should use a chunk size of 10 times the average transfer request rate.

- Random I/Os that are scattered over all the areas of the disks should use a chunk size of 20 times the average transfer request rate.

- If you don't know, then you should use a chunk size of 15 times the average transfer request rate.

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

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

Table 2–2 Example Chunk Sizes

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

Increasing the Data Transfer Rate

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

Chunk size = 128k (256 blocks)

Request A

A_1

A_2

A_3

A_4
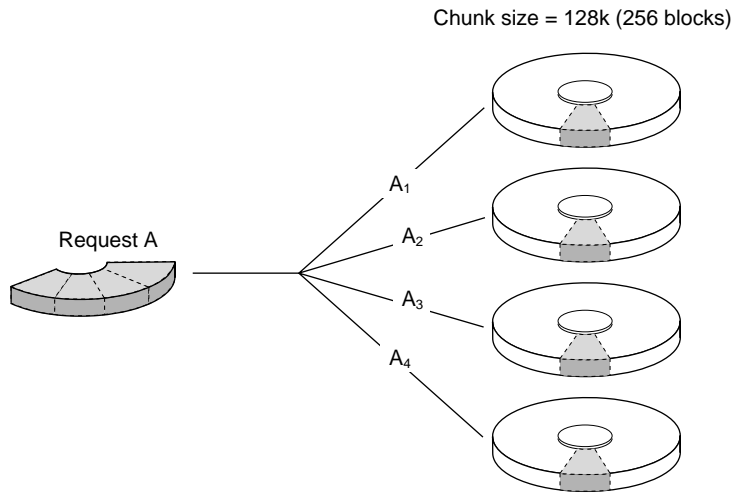
CXO-5172A-MC

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

Increasing Sequential Write Performance

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

Save Configuration

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

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

> **CAUTION:** If user data already exists on a storageset, do NOT reinitialize it with the *save configuration* option, as this will change the site and position of the user data on the storageset. Compaq recommends backing up user data prior to reinitializing any storageset.

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

Destroy/Nodestroy

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

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

■ *DESTROY* (default) overwrites the user data and forced-error metadata on a disk drive when it is initialized.

■ *NODESTROY* preserves the user data and forced-error metadata when a disk drive is initialized. Use *NODESTROY* to create a single-disk unit from any disk drive that has been used as a member of a mirrorset. See the REDUCED command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* for information on removing disk drives from a mirrorset.

  NODESTROY is ignored for members of a RAIDset.

Geometry

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

■ *CAPACITY*—the number of logical blocks. The range is from 1 to the maximum container size.

■ *CYLINDERS*—the number of cylinders used. The range is from 1 to 16777215.

■ *HEADS*—the number of disk heads used. The range is from 1 to 255.

■ *SECTORS_PER_TRACK*—the number of sectors per track used. The range is from 1 to 255.

Unit Switches

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

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

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

Storage Maps

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

Creating a Storage Map

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

Example Storage Map

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

■ Unit D100 is a 3-member RAID 3/5 storageset named R1. R1 consists of Disk10000, Disk10100, and Disk10200.

■ Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:

❏ M1 is a 2-member mirrorset consisting of Disk10300 and Disk10400.

❏ M2 is a 2-member mirrorset consisting of Disk10500 and Disk10800.

■ Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10900 and Disk11000.

■ Unit D103 is a 3-member stripeset named S2. S2 consists of Disk 11100, Disk11200, and Disk11300.

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

Model 4214 Disk Enclosure Shelf 1 (single-bus)													
Bay	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
SCSI ID	00	01	02	03	04	05	08	09	10	11	12	13	14	15
DISK ID	D100 R1 DISK10000	D100 R1 DISK10100	D100 R1 DISK10200	D101 S1 M1 DISK10300	D101 S1 M1 DISK10400	D101 S1 M2 DISK10500	D101 S1 M2 DISK10800	D102 M3 DISK10900	D102 M3 DISK1000	D103 S2 DISK11100	D103 S2 DISK11200	D103 S2 DISK11300	not supported	not supported

Figure 2–17. Example storage map-single bus

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

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

■ Unit D101 is a 2-member striped mirrorset named S1. S1 consists of M1 and M2:

❏ M1 is a 2-member mirrorset consisting of Disk10200 and Disk20200.

❏ M2 is a 2-member mirrorset consisting of Disk20300 and Disk10300.

■ Unit D102 is a 2-member mirrorset named M3. M3 consists of Disk10400 and Disk20400.

■ Unit D103 is a 3-member stripeset named S2. S2 consists of Disk 10500, Disk20500, and Disk10800.

■ Unit D104 is a single (JBOD) disk named Disk20100.

■ Disk20800 is a spareset member.

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

Figure 2–18. Example storage map-dual bus

Using the LOCATE Command to Find Devices

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

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

LOCATE D104

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

LOCATE CANCEL

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

LOCATE R1

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

LOCATE CANCEL

Chapter **3**

Preparing the Host System

This chapter describes how to prepare your Tru64 UNIX host computer to accommodate the HSG60 controller storage subsystem.

The following information is included in this chapter:

- "Making a Physical Connection," page 3–1
- "Preparing LUNs for Access by Tru64 UNIX FileSystem," page 3–3
- "DECsafe Available Server Environment (ASE)," page 3–5
- "HSG60 Units and Tru64 UNIX Utilities," page 3–6

Making a Physical Connection

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

Preparing to Install the Host Adapter

Perform the following steps:

1. Perform a complete backup of your entire system.

2. Shut down your computer system.

Installing the Host Adapter

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

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

You need the following to begin:

- Host bus adapter board

- Your computer hardware manual

- Appropriate tools to service your computer

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

> **NOTE:** Don't power on anything yet. In order for the FC switches to autoconfigure, you must power on equipment in a certain sequence. Also, the controllers in the subsystem are not yet configured for compatibility with Tru64 UNIX.

Preparing LUNs for Access by Tru64 UNIX FileSystem

Tru64 UNIX treats a HSG60 LUN much like a SCSI disk; therefore, to prepare your LUN for access by the Tru64 UNIX filesystem, you must do the following:

■ Create the partitions on the LUN using *disklabel*

■ Create a filesystem on the LUN

■ Mount the filesystem to be able to access it

Creating Partitions on a LUN Using disklabel

Create the partitions on a LUN by issuing a *disklabel* command. The *disklabel* command partitions the LUN for access by the Tru64 UNIX Operating System. Tru64 UNIX defines only partitions a, b, c, and g for the HSG60 controller. In addition, refer to the *rz* and *disktab man* pages for more information about the *disklabel* utility.

To create the read/write partitions on a LUN using the default partition sizes, enter the following:

For V4.0F

```
# disklabel -rw device HSG60
```

Where:

■ HSG60 defines this LUN as attached to the HSG60 array controller. Always use the device label HSG60 regardless of what drive types make up the LUN.

■ *device* is the character device name.

For example, to create partitions on block device, *rza8,* enter:

```
# disklabel -rw rza8 HSG60
```

For V5.0x

```
# disklabel -rw disk HSG60
```

Where:

■ HSG60 specifies the type of disk as found in /etc/disktab or the driver for that device.

■ *disk* is the character device name found in /dev/rdisk.

For example, to create partitions on block device, *rza8,* enter:

```
# disklabel -rw dsk8 HSG60
```

To view a LUN partition, enter:

```
# disklabel -r device
```

Creating a Filesystem on a LUN

> **NOTE:** The *newfs* command is given here as an example. For Advanced File System (ADVFS) and for making devices available for Logical Storage Manager (LSM), similar types of commands exist. For additional information, please consult the related documentation.

Use the *newfs* command to create a UFS filesystem on a LUN the same way that you would create a filesystem on a disk device by entering the following:

```
# newfs special-device disk-type
```

Where:

■ *disk-type* is the disk as listed in the etc/disktab or HSG60.

■ *special device* can be either of the following examples:

For V4.0F:

```
/dev/rrza8c
```

To create a UFS filesystem using the above examples, enter the following:

```
# newfs /dev/rrza8c HSG60
```

For V5.0x:

```
/dev/rdisk/dsk8c
```

To create a UFS filesystem using the above examples, enter the following:

```
# newfs /dev/rdisk/dsk8c HSG60
```

This creates a UFS filesystem on the c partition (whole disk-device).

Mounting the Filesystem

To access the LUN, mount it as a device filesystem to a mount point. For example:

For V4.0F

```
# mount /dev/rza8c/mnt
```

For V5.0x

```
# mount /dev/disk/dsk8c/mnt
```

To view the mounted filesystem, enter:

```
df
```

The LUN is now accessible to the filesystem just as a disk device would be. The filesystem can not see the RAID functionality and number of physical devices attached to the HSG60 controller. This device appears as a single LUN or "disk" to the user as viewed by the filesystem.

DECsafe Available Server Environment (ASE)

HSG60 disk devices can be used with the DECsafe Available Server Environment (ASE) for Tru64 UNIX provided a valid Host configuration (including Host adapters) is used to support them. Refer to the *Tru64 UNIX ASE Installation and User's Guide, Software Product Description, SPD: 44.17.xx,* for further information. Refer to the *Release Notes* for supported host adapters and Tru64 UNIX version levels for ASE.

Using genvmunix

> **NOTE:** This section on genvmunix does not apply to Tru64 UNIX V5.0x, since the configuration file has a different format for devices

If you use *genvmunix* to initialize the system and *doconfig* to build a new configuration file, the new configuration file will only list the HSG60 LUN 0 units; nonzero LUNs will not appear in the new configuration file.

Before rebuilding a configuration file using *genvmunix*, save any existing customized configuration file that has entries for HSG60 units. After rebuilding the configuration file, add the entries from the saved configuration file to the new configuration file.

HSG60 Units and Tru64 UNIX Utilities

This section contains notes on the interactions of some Tru64 UNIX utilities with storage units in your MA6000 storage system.

File

You can use the Tru64 UNIX *file* utility to determine if a Controller Unit can be accessed from the host.

The unit that you want to test must already have a character mode device special file and the correct disk label.

The following example uses the HSG60 unit D101 on SCSI Bus-2. Run the file command and specify the character mode device special file, such as:

For V4.0F

```
# /usr/bin/file /dev/rrzb17a
```

The device activity indicator (green light) will illuminate on the device if the information is not in cache. If the unit is a multi-device container, only one of the devices from that container will illuminate. The Tru64 UNIX Operating System should display something like the following output after the command is entered:

```
/dev/rrzb17a character special (8/33856) SCSI #2 HSG60 disk #146 (SCSI ID #1)
```

■ 8 is the major number

■ 33856 represents the minor number

■ 2 is the SCSI host-side bus number

■ 146 is the drive number as listed in the Configuration File

■ 1 is the Controller Target ID

For V5.0x

/usr/bin/file /dev/rdisk/dsk7c

The device activity indicator (green light) will illuminate on the device if the information is not in cache. If the unit is a multi-device container, only one of the devices from that container will illuminate. The Tru64 UNIX Operating System should display something like the following output after the command is entered:

/dev/rdisk/dsk7c character special (19/150) SCSI #2 HSG60 disk #0 (SCSI ID #1) (SCSI LUN #1)

■ 19 is the major number

■ 150 represents the minor number

■ 2 is the SCSI host-side bus number

■ 0 is the drive number as listed in the Configuration File

■ 1 is the Controller Target ID

■ 1 is the LUN number

If the only output that is returned from the file command is the major and minor number, then either the device is not answering or the device special file does not have the correct minor number. Check the minor number to be sure that it matches the host SCSI bus number, the Controller target ID, and the LUN of Controller Unit.

If an error occurs regarding the disk label, there is good probability that the device can be accessed. This error can usually be fixed by creating the disk label with the Tru64 UNIX *disklabel* utility.

> **NOTE:** For major and minor number calculations see the manual pages, "man SCSI".

Reading from the device, dd

Check the created device using *dd* on the 'raw' device to see if there is a full communication path between devices and Tru64 UNIX; for example:

For V4.0F

dd if=/dev/rrzb17a of=/dev/null

For V5.0x

dd if=/dev/rdisk/dsk7a of=/dev/null

This will read the full disk-device until you press *Ctrl-c* or the device has been read. If the test is successful, the device activity LED (green) on the device lights. If the device consists of multiple disks, all these will be lit.

scu

You can use the *SCSI CAM Utility (scu)* program to see which HSG60 units are available to the Tru64 UNIX Operating System. It is located in the /sbin directory and documented in the REF Pages.

The *scu* command, *scan ed*t, polls all devices on the host-side SCSI buses. This allows you to show what devices are available from all host-side SCSI buses. The device special files do not have to exist for *scu* to see the devices. Example: scan SCSI bus 2, where your MA6000 is connected:

For V4.0F and V5.0x

/sbin/scu scan edt bus 2

/sbin/scu show edt bus 2

V4.0F output:

CAM Equipment Device Table (EDT) Information:

Bus: 2, Target: 1, Lun: 0, Device Type: Direct Access
Bus: 2, Target: 1, Lun: 1, Device Type: Direct Access
Bus: 2, Target: 1, Lun: 2, Device Type: Direct Access
Bus: 2, Target: 1, Lun: 3, Device Type: Direct Access

V5.0x output:

CAM Equipment Device Table (EDT) Information:

Bus/Target/Lun Device Type ANSI Vendor ID Product ID Revision N/W

---------------- -------------- ---- ---------- ------------ -------- ----

2	1	0	Direct	SCSI-2 DEC	HSG60	V85F W
2	1	1	Direct	SCSI-2 DEC	HSG60	V85F W
2	1	2	Direct	SCSI-2 DEC	HSG60	V85F W
2	1	3	Direct	SCSI-2 DEC	HSG60	V85F W

For detailed information about one of the units, you may use:

    # /sbin/scu show device bus 2 target 1 lun 0

The preceding command line gives you the SCSI inquiry of the selected device.

HSG60 units appear like any other SCSI device. All four entries for Bus 2, Target 1 in the example display are HSG60 units. The last entry would be for unit D103 on host SCSI bus 2.

iostat

You can use the *iostat* utility to view performance statistics on MA6000 storage units. (Set your terminal screen to 132 columns before running *iostat*.)

The output from *iostat* shows the number of devices (LUNs) that have been defined in the configuration file. It is much easier to interpret the output if the configuration file contains entries for all eight devices. If the configuration file does not contain entries for all devices, the *iostat* output has fewer columns and it is difficult to correlate each column with a specific device.

For V4.0F

iostat rznn s t

The output from *iostat* shows all devices that have device name rznn. The information for LUN 0 is in the first column, the information for LUN 1 is in the second column, and so forth.

iostat rz16 5 4

```
rza16 rzb16 rzc16 rzd16 rze16 rzf16 rzg16 rzh16
bps tps bps tps bps tps bps tps bps tps bps tps bps tps bps tps
0 0 0 0 0 0 0 0 0 0 0 0 0 0 126 3
0 0 0 0 0 0 0 0 0 0 0 0 0 0 1618 34
0 0 0 0 0 0 0 0 0 0 0 0 0 0 1639 34
0 0 0 0 0 0 0 0 0 0 0 0 0 0 1610 34
```

The above display shows activity on all 8 LUNs attached to one SCSI-bus (2), one Target 0. The above list represents information on rza16, rzb16... rzh16. The device with activity is device rzh16.

The *iostat* utility only recognizes device names in the format rznn where nn is calculated as:

(8 * Host SCSI Bus #) + (HSG60 Target ID)

(This is the same formula and format that is used for the configuration file.)

Invoke the *iostat* utility using the following format:

For V5.0x

iostat dsk3 s t

Where:

■ *rznn* or *dsk3* is the device name

■ The *s* is optional and denotes the amount of time, in seconds, between screen updates

■ The *t* is optional and denotes the total number of screen updates

The output from *iostat* shows all devices that have device name rznn. The information for LUN 0 is in the first column, the information for LUN 1 is in the second column, and so forth.

```
    # iostat dsk3 1 5
```

tty		floppy0		dsk3		cpu			
tin	tout	bps	tps	bps	tps	us	ni	sy	id
0	5	0	0	10	2	0	0	1	99
0	52	0	0	16	2	0	0	0	100
0	52	0	0	16	2	0	0	0	100
0	52	0	0	16	2	0	0	0	100
0	52	0	0	8	1	0	0	1	99

Chapter **4**

Installing and Configuring the HS-Series Agent

The following information is included in this chapter:

■ "Why Use StorageWorks Command Console?" on page 4–1

■ "Installation and Configuration Overview," page 4–3

■ "Before You Install the Agent," page 4–5

■ "Installing and Configuring the Agent," page 4–5

■ "Reconfiguring the Agent," page 4–9

■ "Removing the Agent," page 4–11

Why Use StorageWorks Command Console?

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

The Agent obtains the status of the storage connected to the controller. It also passes the status of the devices connected to the controller to other computers and provides e-mail notification and error logging

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

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

Table 4–1 SWCC Features and Components

Features	Agent Required?	Client Required?
Able to create the following:	Yes	Yes
■ 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)		
Able to monitor many subsystems at once	Yes	Yes
Event logging	Yes	No
E-mail notification	Yes	No
Pager notification	Yes	Yes

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

The following table provides an overview of the installation.

Table 4–2 Installation and Configuration Overview

Step	Procedure
1	Verify that your hardware has been setup correctly. See the previous chapters.
2	Verify that you have a network connection for the client and agent systems.
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 controller. See "Installing and Configuring the Agent," page 4–5.
5	Add the name of the client system to the Agent's list of client system entries (TCP/IP network connections). This can be done during installation or when reconfiguring the Agent.
6	Install the Client software on Windows 2000 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 *Compaq StorageWorks Command Console Version 2.3. for the MA6000 User Guide.*
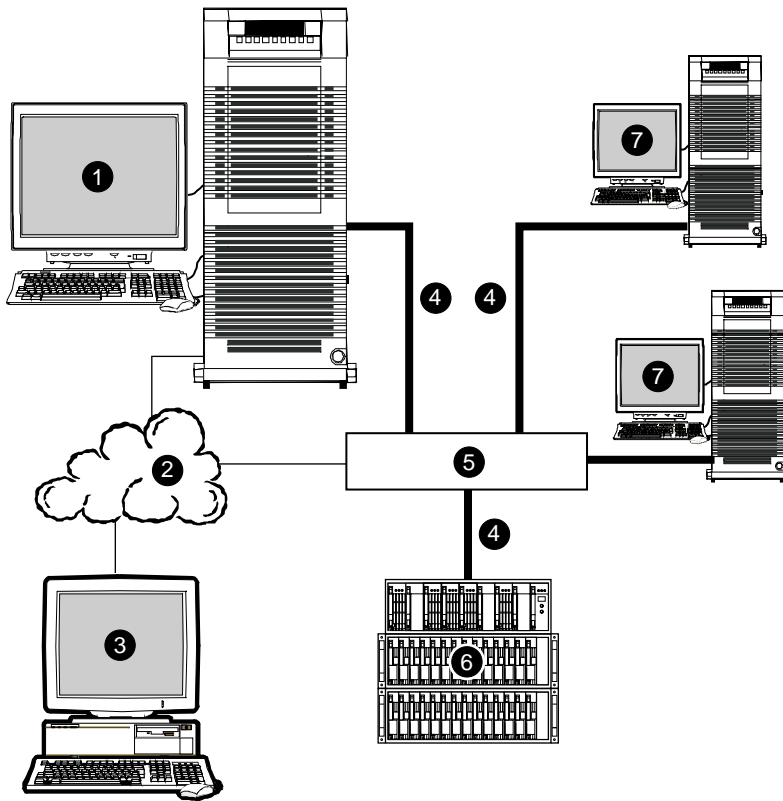
CXO7240A

Figure 4–1. An example of a network connection

❶ Agent system (has the Agent software)

❷ TCP/IP Network

❸ Client system (has the Client software)

❹ Fibre Channel cable

❺ Hub or switch

❻ HSG60 controller and its devices (MA6000)

❼ Servers

Before You Install the Agent

1. Log in as root (superuser). Agent installations on Tru64 UNIX must be done locally. Do not attempt to install the Agent over the network.

2. Remove previous versions of the Agent from your computer.

3. Read the release notes.

> ⚠ **CAUTION:** Do not install the Agent by using Dataless Management Services (DMS) or Remote Installation Services (RIS) sever installations.

Installing and Configuring the Agent

The following instructions assume that you have a directory /mnt to which you can mount the CD-ROM. If you do not, you must create a mount point and replace /mnt in the following sequence with the mount point you created. It also assumes your CD-ROM device is /dev/rz4c. If not, replace /dev/rz4c with the actual CD-ROM device.

> **IMPORTANT:** Before you install the Agent, complete the steps in "Before You Install the Agent."

1. Obtain the following information:

   ❏ Names of systems you want the Agent to contact

   ❏ Names of subsystems connected to the host

   ❏ E-mail addresses for the Agent to send status updates.

2. Insert the CD-ROM into a computer.

3. Enter one of the following commands at the command prompt:

   ■ For V4.0x enter:

   # mount -r-t cdfs -o rrip /dev/rz6c /mnt

   (Substituting rz6c if necessary for your CD-ROM)

   ■ For V5.0x enter:

   # mount -r-t cdfs -o rrip /dev/cdrom0c /mnt

   (Substituting cdrom0c if necessary for your CD-ROM)

4. Change directories on the CD-ROM by entering:

   # cd /mnt/agent

> **CAUTION:** The version of *Compaq Tru64 UNIX* that you are using, either V4.0x or V5.0x will determine which *Device Special File Name* format you will need to enter.

5. To run the installation program, enter the following at the command prompt:

   # setld -l .

   NOTE: The -l is a lowercase L.

   You are asked if you want to install the listed subsets.

6. Enter the corresponding number of the software and then press *RETURN*.

7. Select option **y** and then press *RETURN*.

   As the installation continues, you see several system messages, telling you that the subset has been installed and is being loaded. The computer asks you to add a client system entry.

8. Enter the name of the client system that you want to receive updates from this Agent. Press *RETURN*.

   NOTE: Enter your most important client system first. Enter client systems that are connected infrequently last. The Agent puts the client systems entered first at the top of its list. It first contacts the client system located at the top of the list.

9. From the access options menu, enter an access level for the client system.

   The access privilege level controls the client system's level of access to the subsystems. The following explains each access option:

Table 4–3 Client System Access Options

Options	SWCC Function
0 = Overall Status	■ Can use the Client to add a system to a Navigation Tree, set up a pager, and view properties of the controller and the system
	■ Cannot use Client to open a Storage Window
1 = Detailed Status	■ Can use the Client to open a Storage Window, but you cannot make modifications in that window
2 = Configuration and Status	■ Can use the Client to make changes in a Storage Window to modify a subsystem configuration

10. Press *RETURN*.

    A menu for selecting a client system notification scheme appears:

11. From the menu, enter a notification scheme for the client system, **0**, **1**, **2**, or **3**.

The notification scheme defines the network protocol to be used by the Agent when notifying the selected client system of a change in the state of a subsystem. The following describes how the Transmission Control Protocol/Internet Protocol (TCP/IP) and the Simple Network Management Protocol (SNMP) work with SWCC:

Table 4–4 Client System Notification Options

Options	SWCC Function
Transmission control Protocol/Internet Protocol (TCP/IP)	■ Automatically updates the Storage Window of subsystem changes provided AES is running
	■ Required for Windows NT event logging and pager notification
	■ If you do not select TCP/IP, you will need to refresh the Storage Window (depending on the subsystem) to obtain the latest status of a subsystem.
Simple Network Management Protocol (SNMP)	■ Requires you to use an SNMP-monitoring program to view SNMP traps

12. Press *RETURN*.

    The computer asks you if the entered information is correct.

13. Select option **y** and then press *RETURN*.

    A message, asking if you would like to add another Client, appears.

14. To stop adding client systems, select option **n** and then press *RETURN*.

    You are asked for a password, which is required to do configurations within the Client software. If an old password is found, you are asked if you want to use it.

15. Enter your case-sensitive password that has 4 to 16 characters, and then press *RETURN*. You are asked to retype the password.

16. Retype the password and then press *RETURN*.

    Once the password has been entered, the system scans for subsystems. You will be asked for the name and monitoring interval of each RAID subsystem found.

17. Enter the subsystem name in lowercase letters, and then press *RETURN*. You are asked for the monitoring interval. The monitoring interval is the rate at which the Agent queries the specified subsystem for status.

18. Enter the monitoring interval in seconds, and then press *RETURN*. The monitoring interval determines how often the Agent polls the subsystems. You are asked if the displayed information is correct.

19. If the displayed information is correct, select option **y** and then press *RETURN*.

    You are asked if you want E-mail enabled.

20. To enable the E-mail, select option **y** and then press *RETURN*. The software asks for the address of the person to notify.

21. Enter the address of the person to notify, and then press *RETURN*.

22. From the displayed menu, enter the notification level, **1**, **2**, or **3**.

    Table 2-5 provides the definitions of the E-mail notification options:

<p align="center">**Table 4–5 Definitions of e-mail Notification Options**</p>

Term	Definition
Critical Errors	Notifies you of errors that would prevent you from accessing data; for example, a disk failing or a client system not available on the network.
Warnings	Notifies you when something is broken, but its breakage does not prevent you from accessing data; for example, the RAID degrading or losing a fan.
Information	Provides messages, but they do not indicate that something is broken. Examples of informational messages are the following: an Agent startup message or a message saying that an error has been resolved.

23. Press *RETURN*.

    The software asks if the displayed information is correct.

24. If the displayed information is correct, select option **y** and then press *RETURN*.

    A message, asking if you would like to add another person to notify, appears.

25. Select option **n** to end the dialog or **y** to add another person to notify. Press *RETURN*.

    After you select option **n**, the system will start the new Agent. If the installation script detects a problem at Agent startup, it will alert you.

    You are asked if you are running Agent in a TruCluster environment.

26. Do one of the following:

    ■ If you are running the Agent in cluster, select option **y**. The software is configured to scan the SCSI-bus for subsystems at startup.

    ■ If you are not running the Agent in a cluster, select option **n**. The software is configured to look at the storage.ini file at startup time. The storage.ini file contains the listing of your subsystems.

> **NOTE:** It is recommend that if you are running Tru64 UNIX 5.0a, especially within a fabric, configure the software to scan the SCSI-bus at startup. When asked if running the Agent in a cluster, answer **y**.

The system starts the new Agent.

27. Enter the following and then press *RETURN* to unmount the CD-ROM:

cd /

umount /mnt

Reconfiguring the Agent

You can change your configuration by accessing the Agent Configuration menu. To access the Configuration menu, enter the following:

/usr/opt/SWCC500/scripts/swcc_config

The following is an example of the menu.

```
    SWCC Agent Configuration Utility
-------------------------------

Options Available Are:

1) Add/Delete Client PC Information.
2) Modify Storage Subsystem Information.
3) Change SWCC Agent Password.
4) Turn Email Notification ON/OFF.
5) Add/Delete Email Notification Users.
6) Restart Agent With Changes.
7) Enable/Disable Agent.
8) Cluster/No Cluster

Enter a number or 'q' to QUIT:
```

> △ **CAUTION:** After you make a change to the configuration, such as adding a client system, you must stop and then start the Agent for your changes take effect. When you stop and then start the Agent, the Storage Windows for the subsystems connected to the agent system lose their connection. To regain that connection, close and then reopen the Storage Windows connected to the agent system after you restart the Agent.

Table 4–6 Information Needed to Configure Agent

Term/Procedure	Description
Adding a Client system entry	For a client system to receive updates from the Agent, you must add it to the Agent's list of client system entries. The Agent will only send information to client system entries that are on this list. In addition, adding a client system entry allows you to access the Agent system from the Navigation Tree on that Client system.
Adding a subsystem entry	You need to tell the Agent the subsystem that it needs to monitor.
Client system	Network names for the computers on which the Client software runs.
Client system access options	The access privilege level controls the Client system's level of access to the subsystems.
	0=No Access - Can use the Client software to add a system to a Navigation Tree, set up a pager, and view properties of the controller and the system. You cannot use Client to open a Storage Window.
	1=Show Level Access - Can use the Client software to open a Storage Window, but you cannot make modifications in that window.
	2=Storage Subsystem Configuration Capability - Can use the Client software to make changes in a Storage Window to modify a subsystem configuration.
Client system notification options Note: For all of the client system notification options, local notification is available through an entry in the system error log file and e-mail (provided that e-mail notification in PAGEMAIL.COM has not been disabled).	0 = No Error Notification – No error notification is provided over network.
	1 = Notification via a TCP/IP Socket (Transmission Control Protocol/Internet Protocol) - It updates the Storage Window of subsystem changes provided AES is running. Required for Windows NT event logging and pager notification. If you do not select TCP/IP, you will need to refresh the Storage Window to obtain the latest status of a subsystem.
	2 = Notification via the SNMP protocol (Simple Network Management Protocol) – Requires you to use an SNMP-monitoring program to view SNMP traps.
	3 = Notification via both TCP/IP and SNMP – Combination of options 1 and 2.

<div style="text-align:center">**Table 4–6 Information Needed to Configure Agent (Continued)**</div>

Term/Procedure	Description
Deleting a client system entry	When you remove a client system from the Agent's list, you are instructing the Agent to stop sending updates to that client system. In addition, you will be unable to access this agent system from the Navigation Tree.
E-mail notification	Modify file pagemail.com in directory sys$manager. System is the default. When an error is logged, the Agent executes the PAGEMAIL.COM command. You can modify this file for Agent to log errors in a log file and/or change the account to which Agent sends messages. You can also modify for which level of errors you will be notified. Client does not need to be running to perform these actions.
Monitoring interval in seconds	How often the subsystem is monitored
Password	It must be a text string that has 4 to 16 characters. It can be entered from the client system to gain configuration access. Accessing the Agent Configuration menu can change it.

Removing the Agent

1. Enter the following and then press *RETURN*:

   # setld -d SWCC500

   The procedure to remove the Agent asks the following:

   Are you sure you want to delete all Agent data?

2. Do one of the following:

   ❏ Select option **y** and then press *RETURN*. All configuration files for the Agent is deleted.

   ❏ Select option **n** and then press *RETURN*. The computer saves the configuration data, which includes the following:

      ▲ Subsystem list of entries

      ▲ Client system access options

      ▲ Password

      ▲ E-mail notification

You are able to use the data if you reinstall the Agent or install a newer version of the Agent.

3. Enter the following and then press *RETURN*:

   # setld -d SWCC500

   The procedure to remove the Agent asks the following:

   Are you sure you want to delete all Agent data?

4. Do one of the following:

   ❑ Select option **y** and then press *RETURN*. All configuration files for the Agent is deleted.

   ❑ Select option **n** and then press *RETURN*. The computer saves the configuration data, which includes the following:

      ▲ Subsystem list of entries

      ▲ Client system access options

      ▲ Password

      ▲ E-mail notification

   You are able to use the data if you reinstall the Agent or install a newer version of the Agent.

Chapter **5**

Configuration Procedures

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

The following information is included in this chapter:

■ "Establishing a Local Connection," page 5–2

■ "Setting Up a Single Controller," page 5–3

■ "Setting Up a Controller Pair," page 5–9

■ "Configuring Devices," page 5–16

■ "Configuring a Stripeset," page 5–16

■ "Configuring a Mirrorset," page 5–17

■ "Configuring a RAIDset," page 5–18

■ "Configuring a Striped Mirrorset," page 5–19

■ "Configuring a Single-Disk (JBOD) Unit," page 5–20

■ "Configuring a Partition," page 5–20

■ "Assigning Unit Numbers and Unit Qualifiers," page 5–21

■ "Configuration Options," page 5–23

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.

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.

CXO7181A
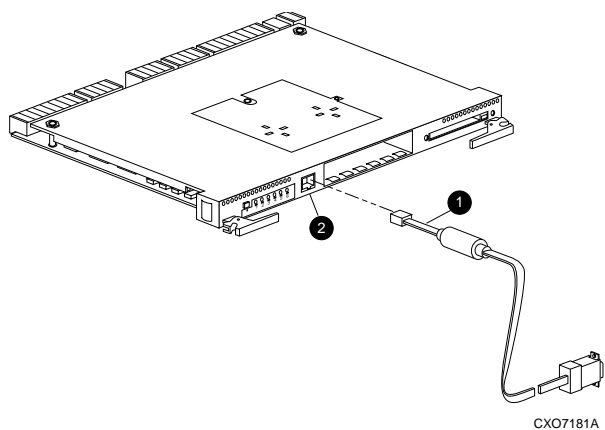
❶ Maintenance Port Cable   ❷ Maintenance Port

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 COMM1 or COMM2 ports.

2. Turn on the computer or terminal.

3. Apply power to the storage subsystem.

4. Configure the computer or terminal as follows:

   ■ 9600 baud

   ■ 8 data bits

   ■ 1 stop bit

   ■ no parity

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

Cabling a Single controller

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

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

CXO6881B

❶ controller
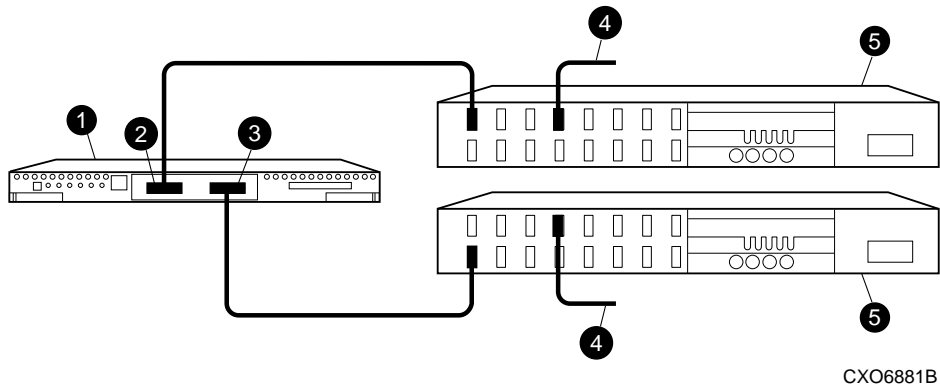
❷ host port 1

❸ host port 2

❹ cable from the switch to the host Fibre Channel adapter

❺ FC switch

Figure 5–2. Single controller cabling

Configuring a Single Controller using CLI

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

   SHOW THIS

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

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

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

   If the node ID is present, go to step 4.
   If the node ID is all zeroes, enter the node ID and checksum, which are located on a sticker on the controller enclosure. Use the following syntax to enter the node ID:

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

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

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

   SET NOFAILOVER

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

   CLEAR_ERRORS CLI

3. Set the version to SCSI-3 for Tru64 UNIX V5.0A only, using the following command:

   SET THIS SCSI_VERSION=SCSI-3

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

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

   SET THIS PORT_1_TOPOLOGY=FABRIC

   SET THIS PORT_2_TOPOLOGY=FABRIC

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

SET THIS PORT_1_TOPOLOGY=OFFLINE

SET THIS PORT_2_TOPOLOGY=OFFLINE

SET THIS PORT_1_TOPOLOGY=FABRIC

SET THIS PORT_2_TOPOLOGY=FABRIC

5. Restart the controller, using the following command:

RESTART THIS

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

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

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

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

RUN FRUTIL

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

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

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

8. 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.

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

SHOW THIS

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

```
Controller:
    HSG60  (C) DEC ZG09030200 Software V85L, Hardware 0000
    NODE_ID         = 5000-1FE1-0000-0000
    ALLOCATION_CLASS = 0
    SCSI_VERSION     = SCSI-3
    Not configured for dual-redundancy
    Device Port SCSI address 7
    Time: 10-Mar-1999:12:30:34
    Command Console LUN is lun 0
Host PORT_1:
    Reported PORT_ID = 5000-1FE1-0000-0001
     PORT_1_TOPOLOGY = FABRIC (fabric up)
    Address         = 210313
Host PORT_2:
    Reported PORT_ID = 5000-1FE1-0000-0002
     PORT_2_TOPOLOGY  = FABRIC (fabric up)
    Address         = 210513
    NOREMOTE_COPY
.......
```

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

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

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

    SHOW CONNECTIONS

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

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

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

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

    RENAME !NEWCON01 ANGEL1A1

    For a recommended naming convention, see "Naming Connections," page 1–12.

13. Specify the operating system for the connection:

    SET ANGEL1A1 OPERATING_SYSTEM=TRU64_UNIX

14. Verify the changes:

    SHOW CONNECTIONS

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

15. Repeat step 12, 13, and 14 for each of that adapter's host connections. (Or delete the unused connections from the table.

16. For each host adapter, repeat steps 11 through 15.

Verifying Installation

To verify installation for your Tru64_UNIX host, enter the following command:

For 4.0F

file /dev/rrz*c | grep HSG60

For 5.0A

file /dev/cport/ *

Your host computer should report one CCL device special file for each HSG60 configured.

Setting Up a Controller Pair

Power On and Establish Communication

1. Connect the computer or terminal to the controller as shown in Figure 5–1. The connection to the computer is through the COMM1 or COMM2 ports.

2. Turn on the computer or terminal.

3. Apply power to the storage subsystem.

4. Configure the computer or terminal as follows:

   ■ 9600 baud

   ■ 8 data bits

   ■ 1 stop bit

   ■ no parity

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

Cabling a Controller Pair

The cabling for a controller pair is shown in Figure 5–3.

> **NOTE:** It is a good idea to plug only the controller cables into the switch. The host cables are plugged into the switch as part of the configuration procedure ("Configuring a Controller Pair using CLI," page 5–11).

CXO6887B
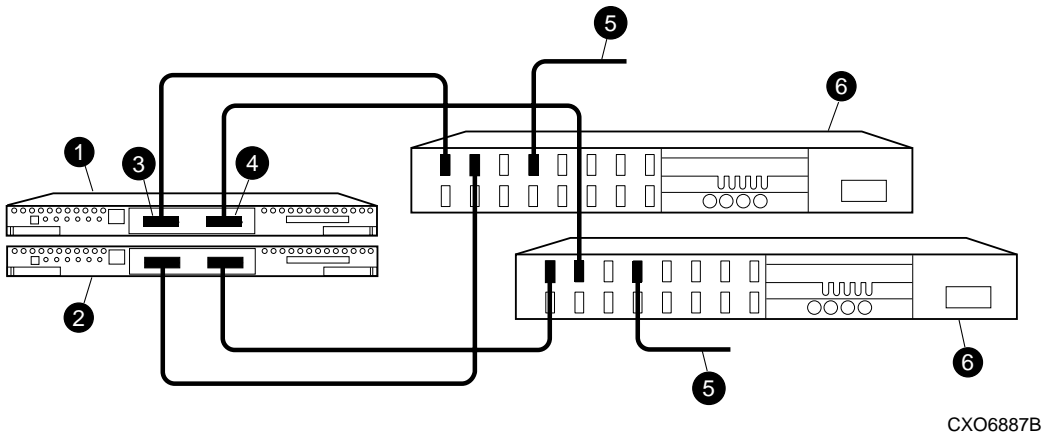
❶ controller A   ❹ host port 2

❷ controller B   ❺ cable from the switch to the host FC adapter

❸ host port 1   ❻ FC switch

Figure 5–3. Controller pair failover cabling

Configuring a Controller Pair using CLI

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

   SHOW THIS

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

   ```
   HSG1402> show this
   Controller:
      HSG60 ZG80900583 Software V8.5L-0, Hardware E01
      NODE_ID         = 5000-1FE1-0001-3F00
      ALLOCATION_CLASS = 0
   ```

   If the node ID is present, go to step 5.
   If the node ID is all zeroes, enter the node ID and checksum, which are located on a sticker on the controller enclosure. Use the following syntax to enter the node ID:

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

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

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

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

   SET NOFAILOVER

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

   CLEAR CLI

4. Put the controller pair into multiple-bus failover mode, using the following command:

   SET MULTIBUS COPY=THIS

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

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

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

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

CLEAR CLI

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

SET THIS SCSI_VERSION=SCSI-3

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

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

SET THIS PORT_1_TOPOLOGY=FABRIC

SET THIS PORT_2_TOPOLOGY=FABRIC

SET OTHER PORT_1_TOPOLOGY=FABRIC

SET OTHER PORT_2_TOPOLOGY=FABRIC

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

SET THIS PORT_1_TOPOLOGY=OFFLINE

SET THIS PORT_2_TOPOLOGY=OFFLINE

SET OTHER PORT_1_TOPOLOGY=OFFLINE

SET OTHER PORT_2_TOPOLOGY=OFFLINE

SET THIS PORT_1_TOPOLOGY=FABRIC

SET THIS PORT_2_TOPOLOGY=FABRIC

SET OTHER PORT_1_TOPOLOGY=FABRIC

SET OTHER PORT_2_TOPOLOGY=FABRIC

7. Restart the controllers:

RESTART OTHER

RESTART THIS

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

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

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

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

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

   RUN FRUTIL

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

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

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

10. Move the serial cable to the lower controller, and repeat step 8.

11. Set up any additional optional controller settings, such as changing the CLI prompt. See the SET *this_controller/other_controller* command in the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* for the format of optional settings.

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

    SHOW THIS

    SHOW OTHER

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

```
Controller:
    HSG60   (C) DEC ZG09030200 Software V85L, Hardware 0000
    NODE_ID         = 5000-1FE1-0000-0000
    ALLOCATION_CLASS = 0
    SCSI_VERSION    = SCSI-3
    Configured for MULTIBUS_FAILOVER with ZG80200284
    In dual-redundant configuration
    Device Port SCSI address 6
    Time: 10-Mar-1999:12:30:34
    Command Console LUN is lun 0
Host PORT_1:
    Reported PORT_ID = 5000-1FE1-0000-0001
     PORT_1_TOPOLOGY = FABRIC (fabric up)
    Address         = 210313
Host PORT_2:
    Reported PORT_ID = 5000-1FE1-0000-0002
    PORT_2_TOPOLOGY  = FABRIC (fabric up)
    Address         = 210513
    NOREMOTE_COPY
..........
```

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

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

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

    SHOW CONNECTIONS

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

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

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

    RENAME !NEWCON01 ANGEL1A1

    Compaq recommends using a naming convention, see "Naming Connections," page 1–12.

16. Specify the operating system for the connection:

    SET ANGEL1A1 OPERATING_SYSTEM=TRU64_UNIX

17. Verify the changes:

    SHOW CONNECTIONS

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

18. Repeat steps 15, 16, and 17 for each of that adapter's host connections. (Or delete the unwanted connections from the table.

19. For each host adapter, repeat steps 14 through 18.

Verifying Installation

To verify installation for your Tru64_UNIX host, enter the following command:

For 4.0F

file /dev/rrx*c | grep HSG60

For 5.0A

file /dev/cport/ *

Your host computer should report one CCL device special file for each HSG60 configured.

Configuring Devices

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

    RUN CONFIG

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

Configuring a Stripeset

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

   ADD STRIPESET STRIPESET-NAME DISKNNNNN DISKNNNNN ....

2. Initialize the stripeset, specifying any desired switches:

   INITIALIZE STRIPESET-NAME SWITCHES

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

3. Verify the stripeset configuration:

   SHOW STRIPESET-NAME

4. Assign the stripeset a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5–21.

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

    ADD STRIPESET STRIPE1 DISK10000 DISK20000 DISK10100

    INITIALIZE STRIPE1 CHUNKSIZE=128

    SHOW STRIPE1

Configuring a Mirrorset

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

   ADD MIRRORSET MIRRORSET-NAME DISKNNNNN DISKNNNNN SWITCHES

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

2. Initialize the mirrorset, specifying any desired switches:

   INITIALIZE MIRRORSET-NAME SWITCHES

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

3. Verify the mirrorset configuration:

   SHOW MIRRORSET-NAME

4. Assign the mirrorset a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5–21.

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

ADD MIRRORSET MIRR1 DISK10000 DISK20000

INITIALIZE MIRR1

SHOW MIRR1

Configuring a RAIDset

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

   ADD RAIDSET RAIDSET-NAME DISKNNNNN DISKNNNNN DISKNNNNN SWITCHES

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

2. Initialize the RAIDset, specifying any desired switches:

   INITIALIZE RAIDSET-NAME SWITCH

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

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

3. Verify the RAIDset configuration:

   SHOW RAIDSET-NAME

4. Assign the RAIDset a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5–21.

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

   ADD RAIDSET RAID1 DISK10000 DISK20000 DISK10100

   INITIALIZE RAID1

   SHOW RAID1

Configuring a Striped Mirrorset

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

   See "Configuring a Mirrorset" on page 5–17.

2. Create a stripeset and specify the mirrorsets it contains:

   ADD STRIPESET STRIPESET-NAME MIRRORSET-1 MIRRORSET-2....MIRRORSET-N

3. Initialize the striped mirrorset, specifying any desired switches:

   INITIALIZE STRIPESET-NAME SWITCH

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

4. Verify the striped mirrorset configuration:

   SHOW STRIPESET-NAME

5. Assign the stripeset mirrorset a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5–21.

Example:

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

   ADD MIRRORSET MIRR1 DISK10000 DISK20000

   ADD MIRRORSET MIRR2 DISK20100 DISK10100

   ADD MIRRORSET MIRR3 DISK10200 DISK20200

   ADD STRIPESET STRIPE1 MIRR1 MIRR2 MIRR3

   INITIALIZE STRIPE1

   SHOW STRIPE1

Configuring a Single-Disk (JBOD) Unit

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

   INITIALIZE DISK-NAME SWITCHES

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

2. Verify the configuration using the following command:

   SHOW DISK-NAME

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

Configuring a Partition

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

   INITIALIZE STORAGESET-NAME SWITCHES

   or

   INITIALIZE DISK-NAME SWITCHES

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

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

   CREATE_PARTITION STORAGESET-NAME SIZE=N SWITCHES

   or

   CREATE_PARTITION DISK-NAME SIZE=N SWITCHES

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

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

3. Verify the partitions:

   SHOW STORAGESET-NAME

   or

   SHOW DISK-NAME

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

4. Assign the partition a unit number to make it accessible by the hosts. See "Assigning Unit Numbers and Unit Qualifiers" on page 5–21.

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

    ADD RAIDSET RAID1 DISK10000 DISK20000 DISK10100

    INITIALIZE RAID1

    CREATE_PARTITION RAID1 SIZE=25

    CREATE_PARTITION RAID1 SIZE=LARGEST

    SHOW RAID1

Assigning Unit Numbers and Unit Qualifiers

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

Each unit can be reserved for the exclusive use of a host or group of hosts. See "Restricting Host Access in Transparent Failover Mode," page 1–21 and "Restricting Host Access in Multiple-Bus Failover Mode," page 1–24.

Assigning a Unit Number to a Storageset

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

    ADD UNIT UNIT-NUMBER STORAGESET-NAME

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

    ADD UNIT D102 R1

Assigning a Unit Number to a Single (JBOD) Disk

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

ADD UNIT UNIT-NUMBER DISK-NAME

Example:
To assign unit D4 to disk20300, use the following command:

ADD UNIT D4 DISK20300

Assigning a Unit Number to a Partition

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

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

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

ADD UNIT D100 MIRR1 PARTITION=3

Preferring Units

In multiple-bus failover mode, individual units can be preferred to a specific controller. To prefer, for example, unit D102 to "this controller," use the following command:

SET D102 PREFERRED_PATH=THIS

RESTART commands must be issued to both controllers for this command to take effect:

RESTART THIS_CONTROLLER

RESTART OTHER_CONTROLLER

> **NOTE:** The controllers need to restart together for the preferred settings to take effect. The RESTART *other_controller* command must be entered immediately after the RESTART *this_controller* command.

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

This chapter briefly describes how to verify that multiple paths exist to virtual disk units.

After you configure units (virtual disks) through either the CLI or SWCC, reboot the host to enable it to access the new storage.

After the host reboots, verify that the disk is correctly presented to the host. The command to use consists of the following syntax:

    hwmgr -view dev - category disk

Use the following command to verify the configuration:

    # hwmgr -view dev - category disk

The disk information returned is as follows:

```
HWID: Device Name       Mfg     Model       Location

 --------------------------------------------------------------------------------

 123: /dev/disk/dsk56c   DEC     HSG60CCL    bus-5-targ-0-lun-0

 124: /dev/disk/dsk57c   DEC     HSG60       bus-6-targ-0-lun-1

 125: /dev/disk/dsk58c   DEC     HSG60       bus-5-targ-0-lun-2

 126: /dev/disk/dsk59c   DEC     HSG60       bus-5-targ-0-lun-60

 127: /dev/disk/dsk60c   DEC     HSG60CCL    bus-5-targ-10-lun-0

 128: /dev/disk/dsk61c   DEC     HSG60       bus-5-targ-10-lun-60

 129: /dev/disk/dsk62c   DEC     HSG60       bus-5-targ-10-lun-61

 130: /dev/disk/dsk63c   DEC     HSG60       bus-5-targ-10-lun-62
```

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

■ Full array with no expansion cabinet

■ PCMCIA cards installed in both controllers

An example storage subsystem is shown in Figure 7–1. The example system contains three non-clustered TRU64_UNIX hosts, as shown in Figure 7–2. From the hosts' point of view, each host will have four paths to its own virtual disks. The resulting virtual system from the hosts' point of view is shown in Figure 7–3.

The name of the example storage subsystem is "BTVS." The top controller will be named "BTVS TOP" and the bottom controller will be named "BTVS 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.

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 DIS112000	Spare DISK11300	not supported	not supported

Model 4214R Disk Enclosure Shelf 2(single-bus)													
Bay	1	2	3	4	5	6	7	8	9	10	11	12	13	14
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

Host 1
"XANDER"
FCA1 FCA2

Host 2
"BUFFY"
FCA1 FCA2

Host 3
"WILLOW"
FCA1 FCA2

Switch
or hub

Switch
or hub

Connections
XANDER1B1
BUFFY1B1
WILLOW1B1

Connections
XANDER1A1
BUFFY1A1
WILLOW1A1

Connections
XANDER2A2
BUFFY2A2
WILLOW2A2

Connections
XANDER2B2
BUFFY2B2
WILLOW2B2

Host
port 1
active

Controller A

Host
port 2
active

D0   D1   D100   D120

All units visible to all ports

Host
port 1
active

Controller B

Host
port 2
active

NOTE: FCA = Fibre Channel Adapter

CXO7298A

Figure 7–2. Example system
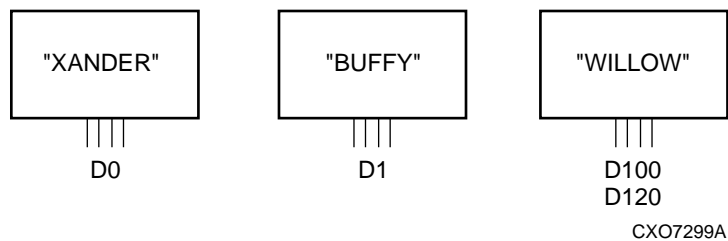
Figure 7–3. Example virtual layout from the hosts' point of view

CLI Configuration Example

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

Plug serial cable from maintenance terminal into top controller.

    CLEAR CLI
    SET MULTIBUS COPY=THIS
    CLEAR CLI
    SET THIS SCSI_VERSION=SCSI-3
    SET THIS PORT_1_TOPOLOGY=FABRIC
    SET THIS PORT_2_TOPOLOGY=FABRIC
    SET OTHER PORT_1_TOPOLOGY=FABRIC
    SET OTHER PORT_2_TOPOLOGY=FABRIC
    SET THIS ALLOCATION_CLASS=0
    RESTART OTHER
    RESTART THIS
    SET THIS TIME=11-JAN-2000:10:20:00
    RUN FRUTIL

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

    Y

Plug serial cable from maintenance terminal into bottom controller.

> **NOTE:** Bottom controller (B) becomes "this" controller.

RUN FRUTIL

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

Y

SET THIS MIRRORED_CACHE.

NOTE: This command causes the controllers to restart.

SET THIS PROMPT="BTVS BOTTOM"

SET OTHER PROMPT="BTVS TOP"

SHOW THIS

SHOW OTHER

RUN CONFIG

ADD RAIDSET R1 DISK10000 DISK20000 DISK10100 DISK20100 DISK10200 DISK20200

INITIALIZE R1

ADD UNIT D0R1 DISABLE_ACCESS_PATH=ALL

SET DO ENABLE_ACCESS_PATH=(XANDER1A1, XANDER1B1, XANDER2A2, XANDER2B2)

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=(BUFFY1A1, BUFFY1B1, BUFFY2A2, BUFFY2B2)

ADD MIRRORSET M3 DISK10900 DISK20900 DISK11000 DISK21000

INITIALIZE M3

ADD UNIT D100 M3 DISABLE_ACCESS_PATH=ALL

SET D100 ENABLE_ACCESS_PATH=(WILLOW1A1, WILLOW1B1, WILLOW2A2, WILLOW2B2)

ADD STRIPESET S2 DISK11100 DISK11200 DISK21100 DISK21200

INITIALIZE S2

ADD UNIT D120 S2 DISABLE_ACCESS_PATH=ALL

SET D120 ENABLE_ACCESS_PATH=(WILLOW1A1, WILLOW1B1, WILLOW2A2, WILLOW2B2)

ADD SPARESET DISK11300

ADD SPARESET DISK21300

SHOW UNITS ALL

Plug in the Fibre Channel cable from the first adapter in host "XANDER."

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			
!NEWCON01	WNNT	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

RENAME !NEWCON00 XANDER1B1

SET XANDER1B1 OPERATING_SYSTEM=TRU64_UNIX

RENAME !NEWCON01 XANDER1A1

SET XANDER1A1 OPERATING_SYSTEM=TRU64_UNIX

SHOW CONNECTIONS

NOTE: Connection table sorts alphabetically.

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
XANDER1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	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 "XANDER."

SHOW CONNECTIONS

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
!NEWCON02	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
!NEWCON03	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
...						

RENAME !NEWCON02 XANDER2B2

SET XANDER2B2 OPERATING_SYSTEM=TRU64_UNIX

RENAME !NEWCON02 XANDER2A2

SET XANDER2A2 OPERATING_SYSTEM=TRU64_UNIX

SHOW CONNECTIONS

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
XANDER1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

Mark or tag both ends of Fibre Channel cables

Repeat this process to add connections from the other two hosts. The resulting connection table should appear similar to the following:

Connection Name	Operating System	Controller	Port	Address	Status	Unit Offset
BUFFY1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BUFFY1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BUFFY2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
BUFFY2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
WILLOW1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
WILLOW1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
WILLOW2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
WILLOW2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER1A1	TRU64_UNIX	OTHER	1	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER1B1	TRU64_UNIX	THIS	1	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER2A2	TRU64_UNIX	OTHER	2	XXXXXX	OL other	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			
XANDER2B2	TRU64_UNIX	THIS	2	XXXXXX	OL this	0
HOST_ID=XXXX-XXXX-XXXX-XXXX			ADAPTER_ID=XXXX-XXXX-XXXX-XXXX			

For more information on using CLI, please see the *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*

Chapter **8**

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

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 Storagesets," page 8–6

Backing Up the Subsystem Configuration

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

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

    SHOW THIS_CONTROLLER FULL

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

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

SHOW DEVICES FULL

Cloning Data for Backup

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

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

After the CLONE utility copies the data from the members to the clones, it restores the unit to its original configuration and creates a clone unit you can backup. The CLONE utility uses steps shown in Figure 8–1 to duplicate each member of a unit.
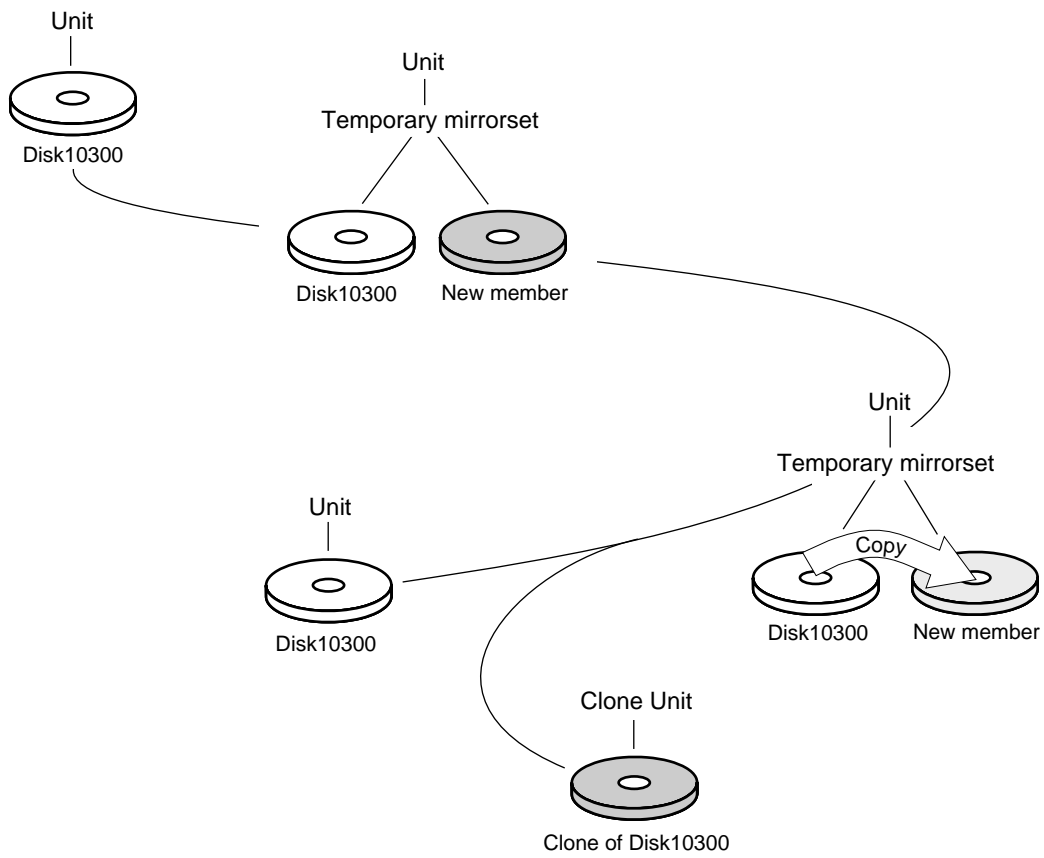
CXO5510A

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

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

1. Establish a connection to the controller that accesses the unit you want to clone.

2. Start CLONE using the following command:

   RUN CLONE

3. When prompted, enter the unit number of the unit you want to clone.

4. When prompted, enter a unit number for the clone unit that CLONE will create.

5. When prompted, indicate how you would like the clone unit to be brought online: either automatically or only after your approval.

6. When prompted, enter the disk drives you want to use for the clone units.

7. Back up the clone unit.

Example:

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

RUN CLONE

CLONE LOCAL PROGRAM INVOKED

UNITS AVAILABLE FOR CLONING:

      98

ENTER UNIT TO CLONE? **98**

CLONE WILL CREATE A NEW UNIT WHICH IS A COPY OF UNIT 98.

ENTER THE UNIT NUMBER WHICH YOU WANT ASSIGNED TO THE NEW UNIT? **99**

THE NEW UNIT MAY BE ADDED USING ONE OF THE FOLLOWING METHODS:

1. CLONE WILL PAUSE AFTER ALL MEMBERS HAVE BEEN COPIED. THE USER MUST THEN PRESS

   RETURN TO CAUSE THE NEW UNIT TO BE ADDED.

2. AFTER ALL MEMBERS HAVE BEEN COPIED, THE UNIT WILL BE ADDED AUTOMATICALLY.

UNDER WHICH ABOVE METHOD SHOULD THE NEW UNIT BE ADDED[]?**1**

DEVICES AVAILABLE FOR CLONE TARGETS:

DISK20200 (SIZE=832317)

DISK20300 (SIZE=832317)

USE AVAILABLE DEVICE DISK20200(SIZE=832317) FOR MEMBER DISK10300(SIZE=832317) (Y,N) [Y]? **Y**

MIRROR DISK10300 C_MA

SET C_MA NOPOLICY

SET C_MA MEMBERS=2

SET C_MA REPLACE=DISK20200

DEVICES AVAILABLE FOR CLONE TARGETS:

DISK20300 (SIZE=832317)

USE AVAILABLE DEVICE DISK10400(SIZE=832317) FOR MEMBER DISK10000(SIZE=832317) (Y,N) [Y]? **Y**

MIRROR DISK10000 C_MB

SET C_MB NOPOLICY

SET C_MB MEMBERS=2

SET C_MB REPLACE=DISK10400

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

.

.

COPY FROM DISK10300 TO DISK20200 IS 100% COMPLETE

COPY FROM DISK10000 TO DISK10400 IS 100% COMPLETE

PRESS RETURN WHEN YOU WANT THE NEW UNIT TO BE CREATED

REDUCE DISK20200 DISK10400

UNMIRROR DISK10300

UNMIRROR DISK10000

ADD MIRRORSET C_MA   DISK20200

ADD MIRRORSET C_MB   DISK10400

ADD STRIPESET C_ST1 C_MA C_MB

INIT C_ST1   NODESTROY

ADD UNIT D99 C_ST1

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

CLONE - NORMAL TERMINATION

Moving Storagesets

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

> ⚠ **CAUTION:** Move only normal storagesets. Do not move storagesets that are reconstructing or reduced, or data corruption will result.

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

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

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

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

   SHOW STORAGESET-NAME

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

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

   LOCATE DISK-NAME

   To cancel the locate command, enter the following:

   LOCATE CANCEL

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

   DELETE UNIT-NUMBER

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

   DELETE STORAGESET-NAME

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

   DELETE DISK-NAME

   DELETE DISK-NAME

   DELETE DISK-NAME

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

7. 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, and DISK10100.

DELETE D100

DELETE RAID99

DELETE DISK10000

DELETE DISK10100

DELETE DISK20000

DELETE DISK20100

ADD DISK DISK10000

ADD DISK DISK10100

ADD DISK DISK20000

ADD DISK DISK20100

ADD RAIDSET RAID99 DISK10000 DISK10100 DISK20000 DISK20100

ADD UNIT D100 RAID99

Appendix **A**

Subsystem Profile Templates

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

■ "Storageset Profile," page A-2

■ "Storage Map Template 1," page A-3 for a single Model 4214R enclosure

■ "Storage Map Template 2," page A-4 for a single Model 4254 enclosure

■ "Storage Map Template 3," page A-5 for two Model 4214R enclosures

Storageset Profile

Type of Storageset:

_____ Mirrorset	__X_ RAIDset	_____ Stripeset	_____ Striped Mirrorset _____ JBOD

Storageset Name

Disk Drives

Unit Number

Partitions:

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

RAIDset Switches:

Reconstruction Policy	Reduced Membership	Replacement Policy
___Normal (default)	___No (default)	___Best performance (default)
___Fast	___Yes, missing:	___Best fit
		___None

Mirrorset Switches:

Replacement Policy	Copy Policy	Read Source
___Best performance (default)	___Normal (default)	___Least busy (default)
___Best fit	___Fast	___Round robin
___None		___Disk drive:

Initialize Switches:

Chunk size	Save Configuration	Metadata
___ Automatic (default)	___No (default)	___Destroy (default)
___ 64 blocks	___Yes	___Retain
___ 128 blocks		
___ 256 blocks		
___ Other:		

Unit Switches:

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

Storage Map Template 1

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

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

Storage Map Template 2

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

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

Storage Map Template 3

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

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

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

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

Appendix **B**

Installing, Configuring, and Removing the Client

The following information is included in this appendix:

■ "Why Install the Client?," page B–1

■ "Before You Install the Client," page B–2

■ "Installing the Client," page B–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–8

Why Install the Client?

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

■ Create mirrored device group (RAID 1)

■ Create striped device group (RAID 0)

■ Create striped mirrored device group (RAID 0+1)

■ Create striped parity device group (RAID 3/5)

■ Create an individual device (JBOD)

■ Monitor many subsystems at once

■ Set up pager notification

Before You Install the Client

1. Verify that you are logged into an account that is a member of the administrator group.

2. Check the software product description that came with your software for a list of supported hardware.

3. Verify that you have the SNMP service installed on your computer. You must have SNMP installed on your computer for this software to work properly. The Client software uses SNMP to receive traps from the Agent. The SNMP service is available on your Windows NT or Windows 2000 installation CD-ROM. To verify that you have the SNMP service:

   ❑ For Windows NT, double-click Services in *Start > Settings > Control Panel*. The entry for SNMP is shown in this window. If you install the SNMP service and you already have a Windows NT service pack on your computer, reinstall the service pack after installing the SNMP service.

   ❑ For Windows 2000, click *Start > Settings > Control Panel > Administrative Tools > Component Services*. The entry for SNMP is shown in the Component Services window.

4. Read the release notes.

5. Read "Troubleshooting the Client Installation," in this appendix.

6. If you have the Command Console Client open, exit the Command Console Client.

7. If you have Command Console Client version 1.1b or earlier, remove the program by going into Add/Remove Programs.

8. If you have a previous version of Command Console, you can save your Navigation Tree configuration by copying the swcc2.mdb file to another directory. After you have installed the product, move swcc2.mdb to the directory to which you installed SWCC.

9. Install the HS-Series Agent. For more information, see Chapter 4.

Installing the Client

1. Insert the CD-ROM into a computer running Windows 2000 or Windows NT 4.0 (Intel) with Service Pack 4 or later.

2. Using Microsoft Windows Explorer, go to the top-level directory on the CD-ROM, and double-click setup.exe. The SWCC Setup window appears.

3. Select HSG60 Controller and click *Next*.

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

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

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

Troubleshooting the Client Installation

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

■ Invalid Network Port Assignments During Installation

■ "There is no disk in the drive" Message

Invalid Network Port Assignments During Installation

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

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

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

■ The KZPCC Agent and Client, (4991 and 4985)

■ The RA200 Agent and Client, (4997 and 4995)

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

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

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

"There is no disk in the drive" Message

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

■ Ignore the error message by clicking *Ignore*.

■ Replace the removable disks, and place a disk in the floppy drive and a CD-ROM in the CD-ROM drive. Then, click *Retry*.

Adding the Storage Subsystem and its Host to the Navigation Tree

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

1. Verify that you have properly installed and configured the HS-Series Agent on the storage subsystem host.

2. Click *Start > Programs > Command Console > StorageWorks Command Console*.
    Client displays the Navigation Window. The Navigation Window lets you monitor and
    manage many storage subsystems over the network.
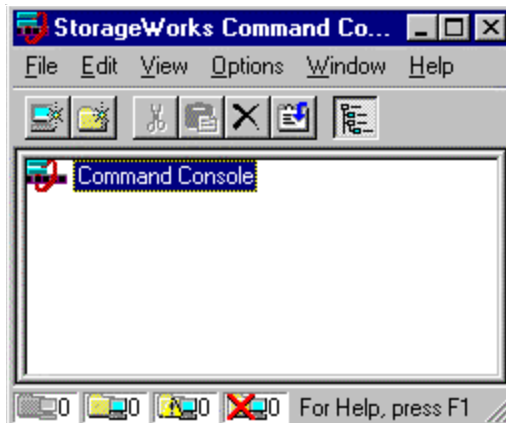
Figure B–1. Navigation window

3. Click *File > Add System*. The Add System window appears.

4. Type the host name or its TCP/IP address and click *Apply*.

5. Click Close.
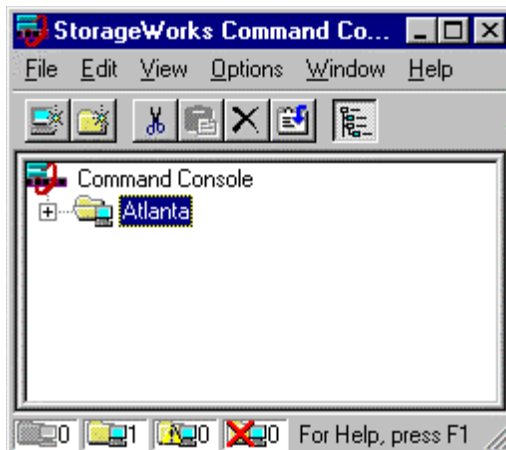
Figure B–2. Navigation window showing storage host system "Atlanta"

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

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

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

Removing the Command Console Client

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

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

1. Click *Start > Programs > Command Prompt* and change to the directory to which you installed the Command Console Client.

2. Enter the following command:

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

3. Do one of the following:

   ❏ On Windows NT 4.0, click *Start > Settings > Control Panel*, and then double-click
     the Add/Remove Programs icon in the Control Panel. The Add/Remove Program
     Properties window appears.

   ❏ On Windows 2000, click *Start > Settings > Control Panel > Add/Remove
     Programs*. The Add/Remove Program window appears.

4. Select Command Console in the window.

5. Do one of the following:

   ❏ On Windows NT 4.0, click *Add/Remove*.

   ❏ On Windows 2000, click *Change/Remove*.

6. Follow the instructions on the screen.

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

Where to Find Additional Information

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

About the User Guide

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

■ About AES

■ Adding Devices

■ Adding Virtual Disks

■ Setting Up Pager Notification

■ How to Integrate SWCC with Compaq Insight Manager

■ Troubleshooting Information

About the Online Help

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

■ Navigation Window – Online Help provides information on pager notification and a tour of the Command Console Client, in addition to information on how to add a system to the Navigation Tree.

■ Storage Window – Online Help provides detailed information about the Storage Window, such as how to create virtual disks.

Index