Compaq SANworks™

Data Replication Manager
HSG80 ACS Version 8.5P

Operations Guide for Tru64 UNIX Version 5.0a and Version 5.1

Part Number: AA-RMPRB-TE / 211068-002
Second Edition (January 2001)
Compaq Computer Corporation

Printed in the U.S.A.

Data Replication Manager HSG80 ACS Version 8.5P
Operations Guide for Tru64 UNIX Version 5.0a and Version 5.1
Second Edition January 2001
Part Number: AA-RMPRB-TE / 211068-002

Contents

Chapter 2

Remote Copy Set Features

Chapter 3

Getting Started

Chapter 4

Configuring a Data Replication Manager Solution

Chapter 5

Managing Site Failover and Failback Procedures

Chapter 6

Troubleshooting

Chapter 7

Configuration Variations

Figures

Tables

About This Guide

This guide describes the Data Replication Manager and how to use it during a failover/failback situation running on the HSG80 Array Controller.

See the documentation that accompanied the subsystem for detailed information about subsystem enclosures and their components.

Getting Help

If you have a problem and cannot find the information you need to resolve it in this guide, you can get further information and other help in the following locations:

Compaq Website

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

http://www.compaq.com/storage

Compaq Technical Support

Use the following to connect with Compaq technical support:

- United States and Canada, call 1-800-652-6672
- Outside the United States and Canada, visit the Compaq website at:

  http://www.compaq.com/storage

Precautions

Follow the precautions listed in the sections “Electrostatic Discharge Precautions” and “Component Precautions” when carrying out the procedures outlined in this guide.

Electrostatic Discharge Precautions

Static electricity collects on all nonconducting material, such as paper, cloth, and plastic. An electrostatic discharge (ESD) can easily damage a controller or other subsystem component, even though you may not see or feel the discharge. Follow these precautions whenever you’re servicing a subsystem or one of its components:

- Always use an ESD wrist strap when servicing the controller or other components in the subsystem. Make sure that the strap contacts bare skin and fits snugly, and that its grounding lead is attached to a bus that is a verified earth ground.
- Before touching any circuit board or component, always touch a verifiable earth ground to discharge any static electricity that may be present in your clothing.
- Always keep circuit boards and components away from nonconducting material.
- Always keep clothing away from circuit boards and components.
- Always use antistatic bags and grounding mats for storing circuit boards or components during replacement procedures.
- Always keep the ESD cover over the program card when the card is in the controller. If you remove the card, put it in its original carrying case. Never touch the contacts or twist or bend the card while you’re handling it.
- Never touch the connector pins of a cable when it is attached to a component or host.

Component Precautions

System components referenced in this guide comply with regulatory standards and the use of other components in their place may violate country standards, negate regulatory compliance, or invalidate the warranty on your product.

Conventions

To help you find what you are looking for, this book uses the special notices and typographical conventions described in the following sections.

Special Notices

This book does not contain detailed descriptions of standard safety procedures; however, it does contain warnings for procedures that could cause personal injury, and cautions for procedures that could damage the controller or its related components. Look for these symbols when you are carrying out the procedures in this book:

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

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

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

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

Text Conventions

Convention	Meaning
ALLCAPS	Command syntax that must be entered exactly as shown, for example: SET FAILOVER COPY=OTHER
ALLCAPS	Command syntax that is discussed within text, for example: "Use the SHOW SPARESET command to show the contents of the spareset."
`Monospaced`	`Screen displays are in monospaced font.`
Sans serif italic *Sans serif italic*	Command variables or numeric values that you supply, for example: SHOW *RAIDset-name* or SET THIS_CONTROLLER PORT_1_SCS_NODENAME=*xxxxxx*
Serif italic	References to other books, for example: "See the *Compaq StorageWorks HSJ80 Array Controller Configuration Guide* for details."
. . .	Indicates that a portion of an example or figure has been omitted.

Related Publications

The following table lists some of the documents that you will need to reference when connecting, configuring, and operating your DRM solution.

Document Title	Part Number
Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide	EK-HSG85-CG. B01
Compaq StorageWorks HSG80 Array Controller ACS V8.5 CLI Reference Guide	EK-HSG85-RG. A01
Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide	EK-HSG84-SV. C01
Compaq StorageWorks Fibre Channel Storage Switch Service Guide	AA-RHBZA-TE
Compaq StorageWorks Fibre Channel Storage Switch User's Guide	AA-RHBYA-TE
Compaq StorageWorks Fibre Channel SAN Switch Management Guide	AA-RMMJA-TE
Compaq StorageWorks RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX Installation Reference Guide	AA-RFAUC-TE
Compaq KGPSA-BC PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide	AA-RF2JB-TE
Compaq StorageWorks Ultra SCSI RAID Enclosure (DS-BA370-Series) User's Guide	EK–BA370–UG
Compaq StorageWorks Command Console Version 2.2 (HSG80) for RA8000/ESA12000 User's Guide	AA-RFA2D-TE
Compaq StorageWorks Command Console Version 2.3 User Guide	AA-RFA2F-TE
Compaq StorageWorks Model 2100 and 2200 Ultra SCSI Controller Enclosure User Guide	EK-SE2C8-UA. B01
Compaq StorageWorks Enclosure 4200 Family LVD Disk Enclosure User Guide	EK-SW2ZS-UA. B01
Compaq SANworks SAN Switch Zoning Reference Guide	EK-P20ZG-GA. A01
Tru64 UNIX AdvFS Administration Guide	AA-RH96A-TE

Data Replication Manager Solution Kit for Tru64 UNIX Version 5.1

The Data Replication Manager (DRM) Solution kit part number for DRM operating on this operating system is:

- Tru64 UNIX Solution Kit Part Number = 128693-B21, QB-6BTAB-SA

The following components are included in the DRM Solution kit:

DRM Solution CD

- *Compaq SANworks Data Replication Manager HSG80 Array Controller ACS Version 8.5P Operations Guide for Tru64 UNIX Version 5.1*
- *Compaq SANworks Data Replication Manager HSG80 Array Controller ACS Version 8.5P Release Notes for Tru64 UNIX Version 5.1*

Additional components not included in the DRM Solution kit that are required to run the Data Replication Manager solution are available from other kits:

- HSG80 Solution for Tru64 UNIX Platform Kit:

  Part Number = 380553-001 / QB-65RAB-SA

Revision History

Revision Level	Date
Revision B, AA-RMPRB-TE / 211068-002	January 2001
Revision A, AA-RMPRA-TE / 211068-001	July 2000

Chapter 1

Introduction to Data Replication Manager

This chapter introduces Data Replication Manager and describes the required hardware and software components.

This chapter covers the following topics:

Data Replication Manager Overview

Data Replication Manager (DRM) provides a disaster-tolerant (DT) solution through the use of hardware redundancy and data replication across multiple sites kilometers apart.

A Data Replication Manager configuration consists of multiple sites. The initiator site carries out primary data processing. Target sites are used for data replication. Since data processing occurs at the initiator site, the data is replicated or mirrored to the target sites. If a significant failure happens to the initiator site, data processing can be resumed at the target sites, where the data is intact.

The DRM sites are connected over some distance via fiber optic cable or ATM. Data Replication Manager uses Fibre Channel gigabit switches to send the data between the sites. If sites are too distant to communicate via Fibre Channel, other hardware may be inserted between the sites.

Peer-to-Peer Remote Copy Function

Data Replication Manager uses the peer-to-peer remote copy function of the HSG80 controller to achieve data replication. HSG80 controller pairs at the initiator site are connected to their partner HSG80 controller pairs at the target site. Remote copy sets are created from units at the initiator and target sites. These remote copy sets are mirrors of each other. As data is written to a unit at the initiator, it is mirrored to its remote copy set partner unit at the target site.

The HSG80 controllers provide failover and failback capabilities in case of failures. Failover makes the data available at the target site after a failure. Failback is used to move data operations back to the initiator once it has been brought back on-line.

Hardware Redundancy

Data Replication Manager provides hardware redundancy. In the face of single component failures at a site, Data Replication Manager fails over to a redundant component at that site to allow continued operations. For example, if one of the dual-redundant Fibre Channel links between the sites were to fail, Data Replication Manager would switch to the other link.

Hardware Components

Data Replication Manager uses a minimum of two HSG80 array controller pairs: one at the initiator site and one at the target site. Each site must have one or more ESA12000 cabinets or EMA12000 modular storage cabinets:

- ESA12000 cabinets are equipped with one or more BA370 enclosures and disk Storage Building Blocks (SBBs). Each BA370 enclosure holds 24 disks.
- EMA12000 modular storage cabinets are equipped with one or more controllers and modular disk SBBs.

See Table 1–3 at the end of this chapter for a complete list of hardware and software required to operate a Disaster Tolerant storage subsystem with Data Replication Manager.

The hosts at the initiator and target sites are connected to a pair of dual redundant HSG80 array controllers, which are located inside of these enclosures. For complete details on this equipment, refer to the *Compaq StorageWorks RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX Installation Reference Guide*.

> **NOTE:** While this documentation addresses ESA12000 storage cabinets as a primary unit for Data Replication Manager configurations, Compaq's DT solution functions in any equivalent cabinet that houses a BA370 enclosure.
>
> For the EMA12000 modular storage cabinets, Compaq's DT solution functions in any equivalent cabinet that houses the same number of controllers and an equivalent drive configuration.
>
> Cabinet configurations may also be combined between the ESA12000 cabinets and EMA12000 modular storage cabinets. For example, if the configurations are equivalent, an ESA12000 cabinet may be used at one site (initiator or target) and an EMA12000 modular storage cabinet may be used at the associated site (target or initiator).

Connections between the controllers and hosts are made at each site with two Fibre Channel gigabit switches and two host bus adapters. Short-wave Gigabit Interface Converters (GBICs) connect the host and controllers to the switches at each site. Extended GBICs or ATM connect the initiator and target switches together if they are more than 500 meters apart.

Cabinet Components

Tables and figures throughout this chapter show hardware that is necessary or optional to complete a modular Data Replication Manager solution for each of two types of cabinet configurations, the ESA12000 cabinet and the EMA12000 modular storage cabinet.

For detailed information about these components, refer to the following documents:

- *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS V8.5 CLI Reference Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide*.

ESA12000 Storage Cabinet

The ESA12000 storage building block cabinet houses the BA370 enclosures, which contain the components listed in Table 1–1.

Table 1–1 ESA12000 Storage Cabinet Components

Two HSG80 Fibre Channel RAID array controllers
One Environmental Monitoring Unit (EMU)
One or two AC input power controllers
Up to 24 disk drive storage building blocks (SBBs) per BA370 enclosure
Five to eight 180-watt power supplies
Dual external cache batteries (ECBs)
Eight cooling fans
Six single-ended I/O Ultra SCSI modules
One Power Verification and Addressing (PVA) module
Two cache modules **(512 MB each required)**

Figure 1–1 shows the initial storage building block parts inside the ESA12000 cabinet with a 24 disk drive capacity. Throughout this chapter, figures, and text outline additional hardware that is necessary or optional to complete a Data Replication Manager solution.

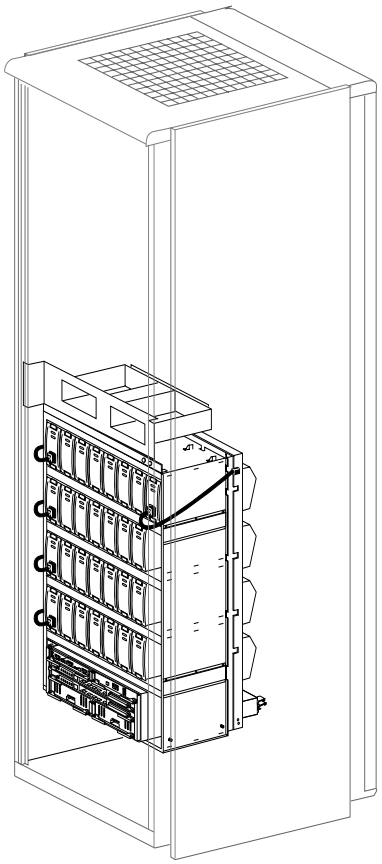

Figure 1–1. ESA12000 storage building block

Figure 1–2 shows additional components that must be added to the ESA12000 building block to support a Data Replication Manager solution, including Fibre Channel gigabit switches. The optional redundant power distribution unit also is shown.

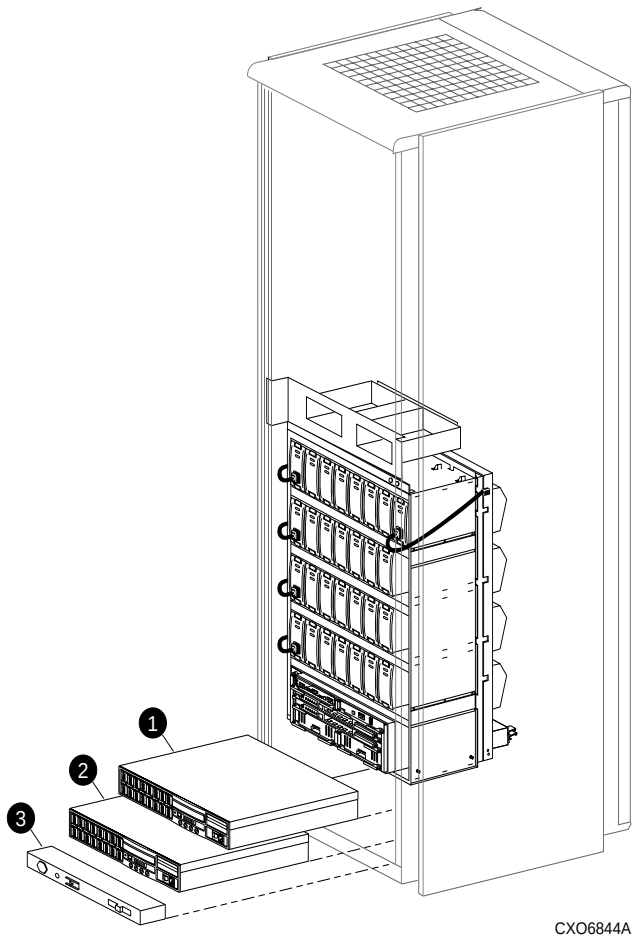

❶ Top Fibre Channel gigabit switch, 8- or 16-port

❷ Bottom Fibre Channel gigabit switch, 8- or 16-port

❸ Redundant power distribution unit (optional)

Figure 1–2. Additional components for ESA12000 Data Replication Manager

EMA12000 Modular Storage Cabinet

The EMA12000 modular storage building block cabinets include power distribution units, are pre-cabled, and contain the components listed in Table 1–2.

Table 1–2 EMA12000 Modular Storage Cabinet Components

Two HSG80 Fibre Channel RAID array controllers
Two Environmental Monitoring Units (EMUs)
Two AC input power controllers
Modular disk drive storage building blocks (SBBs):
D14 - Up to 42 drives per controller subsystem
S14 - Up to 72 drives per controller subsystem
Blue - Up to 42 drives per controller subsystem
Dual power supplies, one per enclosure
Dual external cache batteries (ECBs)
Cooling fans
Six single-ended I/O Ultra SCSI modules
Two cache modules (**512 MB each required**)

Figure 1–3 shows an EMA12000 modular building block that supports a Data Replication Manager solution. The modular storage building block consists of the controller enclosure and the disk enclosure. The redundant power distribution unit also is shown.

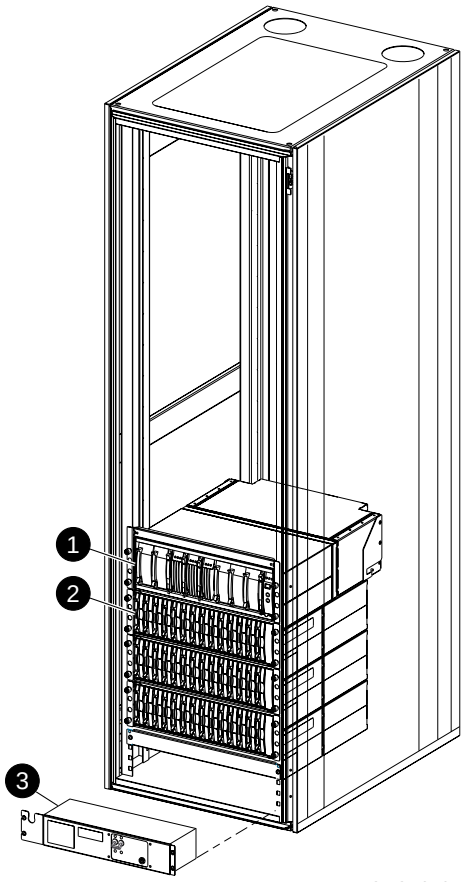

❶ Controller enclosure component of modular storage building block

❷ Disk enclosure component of modular storage building block

❸ Redundant power distribution unit

Figure 1–3. Components for EMA12000 modular storage Data Replication Manager

Switches are placed in a different cabinet in this configuration of the EMA12000 D14 high performance modular storage cabinet. A command center cabinet can be used to contain the switches.

Throughout this chapter, figures, and text outline additional hardware that is necessary or optional to complete a modular Data Replication Manager solution.

Fibre Channel Gigabit Switches

The Fibre Channel gigabit switches shown in Figure 1–4 and Figure 1–5 are two types of SAN switches used to connect the controllers to the hosts and to link the initiator and target sites together. The figures illustrate, respectively, the Fibre Channel SAN Switch 16 and the Fibre Channel SAN Switch 8-EL. The ports hold short- or long-wave Gigabit Interface Converters, which are described in the next section.

See Table 1–3 at the end of this chapter for a complete list of hardware and software required to operate a Disaster Tolerant storage subsystem with Data Replication Manager. Refer to the *Compaq StorageWorks Fibre Channel Storage Switch User's Guide* for an in-depth look at the features and functions of the Fibre Channel gigabit switches. Refer to the *Compaq StorageWorks Fibre Channel SAN Switch Management Guide* for additional detail on SAN switches.

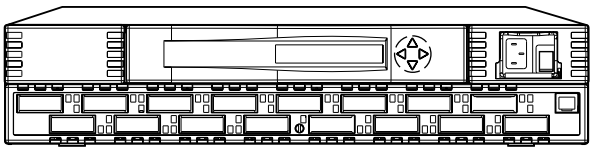

CXO7085A

Figure 1–4. Fibre Channel SAN Switch 16

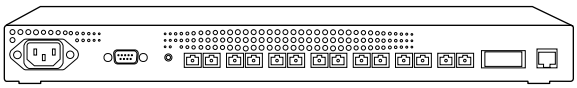

CXO7337A

Figure 1–5. Fibre Channel SAN Switch 8-EL

Gigabit Interface Converters (GBICs)

GBICs are the converters that are inserted into the ports of the Fibre Channel switch and serve as the interface between the fiber optic cables and the switch. Short-wave GBICs are used with a 50-micron multi-mode fiber optic cable (SC-terminated) to connect the components at the initiator and target sites (host-to-switch; controller-to-switch). The maximum distance that short-wave GBICs support is 500 meters.

Long-wave GBICs are used with 9-micron single-mode fiber optic cables (SC-terminated) to link the initiator and target sites. Standard long-wave GBICs connect switches that are up to 10 kilometers apart.

See Table 1–3 at the end of this chapter for a complete list of hardware and software required to operate a Disaster Tolerant storage subsystem with Data Replication Manager. Refer to the *Compaq StorageWorks Fibre Channel Storage Switch User's Guide* to learn more about GBICs.

Power Distribution Unit (PDU)

The PDU component is included with the ESA12000 and EMA12000 cabinets:

- For the ESA12000 cabinets, the PDU is used to distribute power to the BA370s and switches. You can order a second PDU to support a fully-redundant ESA12000 power configuration.
- For the EMA12000 cabinets, the PDU is used to distribute power to the modular configurations of controllers and switches. PDU redundancy is included with EMA12000 modular configurations.

For more detailed information, refer to the *Compaq StorageWorks RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX Installation Reference Guide*.

Fully-Redundant Power (Optional)

Fully-redundant power is an optional feature designed to offer a more secure source of power in case one or more units fail. If less than five power components are operational, the entire cabinet shuts down.

The fully-redundant power feature requires three additional power supplies, as well as one additional AC power controller that plugs into one additional PDU. For the ESA12000 cabinet, these additional components must be supplied for each BA370 enclosure. For the EMA12000 modular storage cabinet, the preconfiguration models (D14, S14, Blue) feature fully-redundant shelves.

For more details about power supply Storage Building Blocks (SBBs), refer to the *Compaq StorageWorks RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX Installation Reference Guide*.

Host Bus Adapters

The host bus adapters are inserted into the available slots on the host computer's PCI Bus. A Fibre Channel connection is made by inserting a multi-mode fiber optic cable between each adapter and an individual port on the Fibre Channel switch.

See Table 1–3 at the end of this chapter for a complete list of hardware and software required to operate a Disaster Tolerant storage subsystem with Data Replication Manager. For more information, refer to the *Compaq KGPSA-BC PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide* and the *Compaq KGPSA-CB PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide.*

Hardware Configurations

Figures shown previously in this chapter have reflected the build of a DT solution for each of two types of cabinet configurations. Figure 1–6 shows a completed Disaster Tolerant (DT) setup for the ESA12000 cabinet. Figure 1–7 shows a completed DT setup for the EMA12000 modular storage cabinet.

NOTE: If you prefer to join cabinets for more storage capacity, follow the instructions in the *Compaq StorageWorks RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX Installation Reference Guide*. Be sure to establish the same setup at both the initiator and target sites. Keep in mind that an additional cabinet not include switches or controllers. It does, however, include a PDU and is able to support redundant power.

.

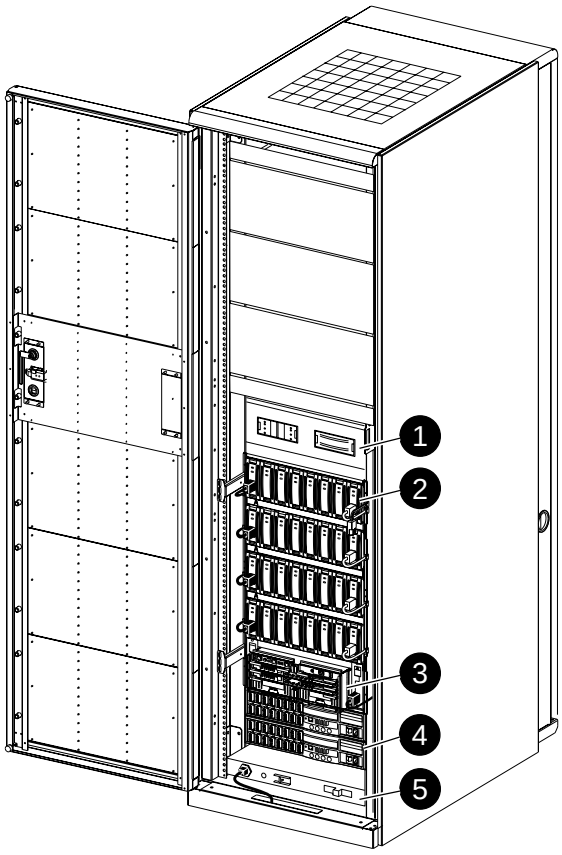

❶ Cache battery

❷ Storage building blocks

❸ Controllers

❹ Fibre Channel gigabit switches, 8- or 16 port

❺ Redundant power distribution unit

CXO7433A

Figure 1–6. Fibre Channel-based, ESA12000 DT storage subsystem (with fully-redundant power)

.

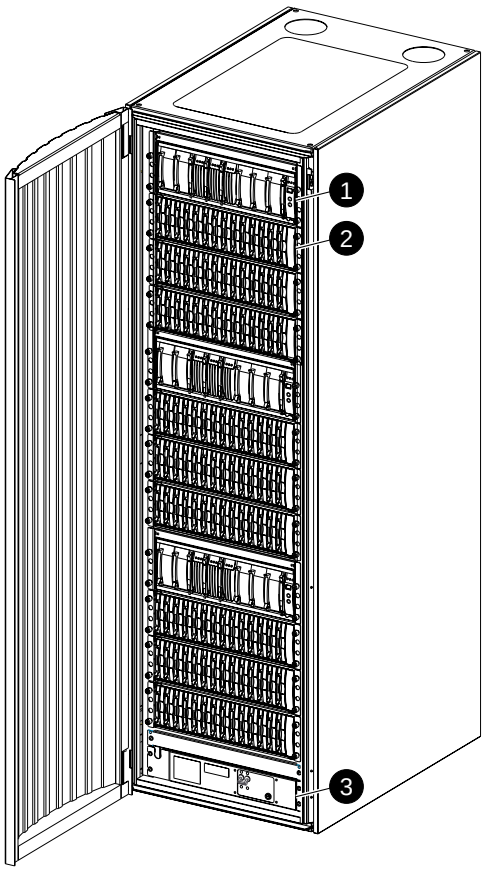

❶ Controller enclosure component of modular storage building block

❷ Disk enclosure component of modular storage building block

❸ Redundant power distribution unit

Figure 1–7. Fibre Channel-based, EMA12000 DT modular storage subsystem (with fully-redundant power)

Software Components

This section describes the software components necessary to configure and manage a DT storage subsystem. For installation instructions, see Chapter 4, "Configuring a Data Replication Manager Solution."

The following list shows the software required to enable Data Replication Manager:

- Array Controller Software (ACS) Version 8.5P
- Storage Works Command Console (SWCC) Version 2.3 (optional)

NOTE: At this time, SWCC is supported in Version 5.0a, but not in Version 5.1.

Array Controller Software

The HSG80 Array Controller Software (ACS) is the software component of the HSG80 array controller subsystem. ACS software executes on the HSG80 controller and processes I/O requests from the host, performing the device-level operations required to satisfy the requests.

StorageWorks Command Console (Optional)

StorageWorks Command Console (SWCC) provides local and remote management of StorageWorks controllers and their attached storage devices. SWCC consists of two major components: the SWCC client and the SWCC agent. SWCC can be used to configure and manage the DT storage subsystem.

The SWCC client is a graphical user interface (GUI) that runs on a local host and displays the logical and physical layout and status of a selected subsystem, in graphical form.

The SWCC agent is a companion program to the client. This host-resident program is an interface between the client and the host's storage subsystem that allows the two to communicate over a network.

For a full description of SWCC and how it operates, refer to the *Compaq StorageWorks Command Console User Guide*.

Required Hardware and Software

Table 1–3 is a checklist of equipment that is mandatory for operating a DT storage subsystem with Data Replication Manager.

Use the list to verify that you have everything, or you may not be able to configure your system to replicate data in a DT subsystem.

Table 1–3 Hardware and Software Requirements Checklist

Required Hardware		Part Number	Quantity (at each site)
ESA12000		380590–B21 (50 HZ, blue) 380590–B22 (50 HZ, opal) 380580–001 (60 HZ, blue) 380580–002 (60 HZ, opal)	Minimum of 1
EMA12000			
	(D14, 60 HZ)	175990–B21/ DS-SWXEB-AA	
	(D14, 50 HZ)	175990–B22 / DS-SWXEB-AB	
	(S14, 60 HZ)	175991–B21/ DS-SWXEB-BA	
	(S14, 50 HZ)	175991–B22 / DS-SWXEB-BB	Minimum of
	(Blue, 60 HZ)	175993–B21/ DS-SWXEB-DA	1
	(Blue, 50 HZ)	175993–B22/ DS-SWXEB-DB	
MA8000			
	(60 HZ)	175992–B21/ DS-SWXEB-CA	
	(50 HZ)	175992–B22/ DS-SWXEB-CB	
Fibre Channel gigabit switches and options:			
SAN Switch 8:			
8-port (Fabric OS v1.6d and v2.1.7)		158222-B21 / DS-DSGGB-AA	
SAN Switch 16:			Minimum of
16-port (Fabric OS v1.6d and v2.1.7)		158223-B21 / DS-DSGGB-AB	2
SAN Switch 8-EL		176219-B21 / PS-DSGGC-AA	
Multiple E-port Connectivity software option			
(SAN Switch 8-EL only)		207104-B21 / QM-GKAAA-AV	
Fibre Channel host cables			
5 meters		234457-B22 / BNGBX-05	2 per host
15 meters		234457-B23 / BNGBX-15	
30 meters		234457-B24 / BNGBX-30	
50 meters		234457-B25 / BNGBX-0	
Fibre Channel controller cables (2 meters)		234457-B21 / DS-BNGBX-02	4

Table 1–3 Hardware and Software Requirements Checklist (Continued)

Short-wave Gigabit Interface Converters (GBICs)	380561-B21 / DS-DXGGA-SA	6
Optional GBICs:		
Long-wave gigabit interface converters (GBICs)	127508-B21 / DS-DSGGA-MA	4
GBIC Very Long Distance (VLD) connector kit	169887–B21 / 3R–A1836–AA	4
Optional Power Distribution Units (ESA 12000):		
PDU, 60 Hz	380582-001 / DS-SW4IU-XA	1 per
PDU, 50 Hz	380583-B21 / DS-SW4IU-XB	cabinet
ATM Gateway Additional Hardware	**Part Number**	**Quantity**
ATM Gateway vOSG 2.2.4	166296-B21	2 per site
ATM Gateway Service Kit	166297-B21	1 per site
Software Requirements	**Part Number**	**Quantity**
ACS V8.5P	128698-B21 / QB-6CAAA-SA	

Chapter 2

Remote Copy Set Features

This chapter discusses Data Replication Manager concepts you need to know for configuring a Data Replication Manager solution. These descriptions include Remote Copy Sets and Association Sets.

This chapter discusses the following topics:

Remote Copy

Data Replication Manager uses the peer-to-peer remote copy function of the HSG80 controller to achieve data replication. The HSG80 controller pairs at the initiator site are connected to their partner HSG80 controller pairs at the target site. Remote Copy Sets are mirrors of each other and are created from units at the initiator and target sites. As data is written to a unit at the initiator site, it is mirrored to its remote copy set partner unit at the target site.

The remote copy feature is intended not only for disaster recovery but also to replicate data from one storage subsystem or physical site to another subsystem or site. It also provides methods of performing a backup at either the local or remote site.

With remote copy, user applications continue to run while data movement goes on in the background over a separate interconnect. Data warehousing, continuous computing, and enterprise applications all require remote copy capabilities. The remote copy feature is the major component in the Compaq Storageworks Data Replication Manager solution.

Remote Copy Sets

A remote copy set is a bound set of two units, one located on the initiator site and the other at the target site, for long-distance mirroring. The term “units” is defined as a single disk, storageset, mirrorset or RAIDset. The local controller is designated as the *Initiator.* The initiator acts as the director of the replication process. The corresponding remote controller is designated as the *Target*. The target receives I/O requests from the initiator to replicate the data at its location.

The ADD REMOTE_COPY_SETS *RemoteCopySetName InitiatorUnitName RemoteNodeName|TargetUnitName* command creates a remote copy set and starts a normalization copy to the target unit. During normalization, the controllers copy all data from the initiator unit to the target unit.

Remote copy sets are created only at the initiator site. There can be up to 12 remote copy sets per controller.

Non-Remote Copy Sets

Non-remote copy sets can exist on the same subsystem at the initiator and/or the target sites and are generally used for local storage at each site. This allows cloning of remote copy sets at the local site for activities like testing and backup. Since the non-remote copy sets may be different at the initiator and target sites, data is not disaster tolerant.

Operation Modes

There are two possible remote copy operation modes: *Synchronous* or *Asynchronous.* Figure 2–1 shows the timeline differences between the two.

Synchronous Operation Mode

In synchronous operation mode, data is simultaneously written to the cache of the initiator subsystem and the cache of the target subsystems. The I/O completion status is not sent to the host until all members of the remote copy set are updated. Synchronous operation ensures the highest possible level of data consistency, which makes this process especially appropriate for business applications that require a high level of currency.

Synchronous is the default setting.

Asynchronous Operation Mode

In asynchronous operation mode, the write operation is reported to the host as complete *before* the data is written to the remote unit of the remote copy set. Asynchronous mode can provide improved performance and response time, but the data on all members of the remote copy set cannot be assumed to be the same at all times.

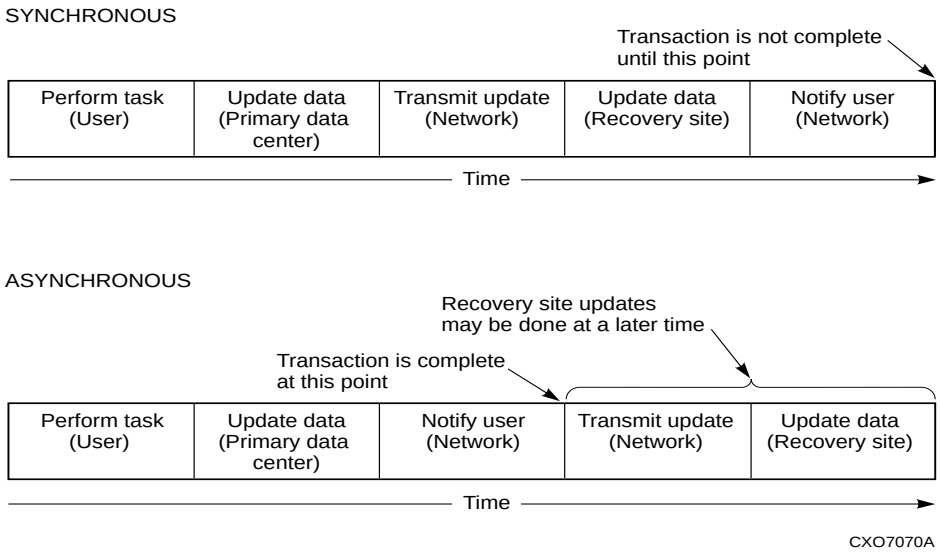

Figure 2–1. Remote copy set operation modes

Operation Mode Considerations

These characteristics should be taken into account when designing your DRM configuration:

- Synchronous replication is appropriate when exact consistency is critical to the business application. The application or application recovery depends upon data being written to both local and remote sites when completion is restored.
- Synchronous operation may deliver best response time for heavy host write operations.
- Asynchronous operation mode improves response time for some workloads.

Outstanding_IO Settings

The OUTSTANDING_IO setting allows you to control the number of initiator to target writes for a remote copy set. It does not refer to the write queue depth between the host and the controller. This setting can be applied to both synchronous and asynchronous remote copy sets. However, this setting causes different behavior, depending on the remote copy set operation mode.

The default setting is 20 for each remote copy set.

Synchronous

For the synchronous operation mode, the OUTSTANDING_IO setting refers to the number of initiator to target writes that can be outstanding at any one time. If OUTSTANDING_IO is set to 1 and the host issues four writes to a remote copy set, then only one write will be in progress between the initiator and target at a time. The other three writes are queued in the initiator controller. As each write completes at the target, another write is issued from the initiator controller write queue.

Asynchronous

For the asynchronous operation mode, the OUTSTANDING_IO setting applies to the number of non-committed host writes that can be outstanding at one time between the initiator and target. Non-committed means the write completion status has been returned to the initiator host, but the write has not been completed at the target.

Suppose, for example, that the outstanding I/O is set to 5 and that the host issues a request, waits for completion from the controller, then immediately issues another request. In asynchronous mode, each request issued by the host is completed by the controller very quickly. As a result, the host issues five requests before the remote site has completed the first request. If the host then issues another (sixth) request, it exceeds the value of the outstanding I/O.

Once you exceed the outstanding I/O value, the system switches to synchronous mode. You must then return to zero outstanding I/Os in order to return to asynchronous mode.

Outstanding Write Operations

Keep in mind that there is a controller-wide limit of 240 outstanding write operations even if the total number of writes is greater than 240. For example, you might have 12 synchronous remote copy sets, each with a value of 100. The maximum outstanding writes are 240, and not 1200. When 240 outstanding writes are reached, any new writes to the controller are queued.

High Outstanding I/O Values

Use caution when choosing an OUTSTANDING_IO setting, since writes to the targets are handled in a FIFO (First In, First Out) manner. As a result, remote copy sets with higher OUTSTANDING_IO values could potentially starve other remote copy sets if the write rates become very high at any one time.

Low Outstanding I/O Values

On the other hand, choosing a lower setting may starve a very active remote copy set. In the case of asynchronous remote copy sets, a lower OUTSTANDING_IO value may be appropriate. This lower value limits the number of outstanding non-committed writes in the event of an initiator site disaster.

Suspend/Resume

The SUSPEND switch suspends the update to the remote copy set target and starts the write history logging of write commands and data from the unit.

NOTE: This switch is valid only in normal error mode (not failsafe).

The RESUME switch initiates the mini-merge restore of the specified remote target unit. This switch enables the initiator to read the log unit and send the write commands, in order, to the target, which brings the target into congruency with the initiator. For more information on mini-merge, see the "Write History Logging" section in this chapter.

NOTE: The SET *AssociationSetName* NOLOG_UNIT command terminates any suspended targets that are currently active.

Error Mode

The following command sets the error mode condition:

SET *RemoteCopySetName* ERROR_MODE=FAILSAFE *or* NORMAL.

The FAILSAFE ERROR MODE causes a remote copy set to become failsafe locked if the target becomes inaccessible or the initiator unit fails. When failsafe locked, the remote copy set is inaccessible.

If a dual link failure occurs, the target is not removed, but is marked invalid. When the target is accessible again, a full copy operation is started. When the copy operation is completed, the failsafe locked condition is cleared.

If the error mode switch is set to NORMAL, write operations are allowed to continue even when a dual link or disk error is present. NORMAL is the default setting.

IMPORTANT: You cannot enable the failsafe switch with write history logging enabled.

Association Sets

An association set is a group of remote copy sets that share common attributes. Members of an association set can be configured to transition to the same state at the same time. For example, if one association set member assumes the failsafe locked condition, all other members of the same association set assume the failsafe locked condition, as well.

An association set may also be used to simply share a log between a group of remote copy set members that require efficient use of the log space.

Association Set Characteristics

Things to remember about association sets include:

- Association sets can have up to 12 remote copy sets as members
- Association sets can share a log unit
- Synchronous or asynchronous operation mode and members may be set differently
- If ORDER_ALL is set, *in order* execution of commands across the remote copy sets in the association set is required
- If FAIL_ALL is set, if one member assumes the failsafe locked condition, then all members of the association set assume the failsafe locked condition
- Association sets reside on the initiator controller pair, as illustrated in Figure 2–2

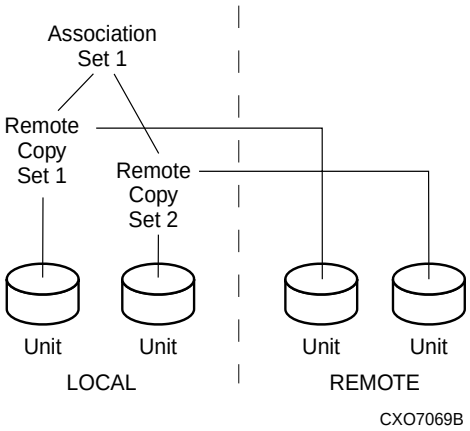

Figure 2–2. Location of association sets on the initiator controller pair

All members of an association set must be on the same controller, to enforce cache coherency. When members are added to an association set, they are moved to reside on the same controller and will failover together.

Association set members can be either synchronous or asynchronous. This allows for grouping only those members that use the write history log unit.

When you use the

ADD ASSOCIATIONS *AssociationSetName RemoteCopySetName*

command, it adds an association set with one member to the controller pair's configuration. Use this command on the node on which the initiator resides. Use the

SET *AssociationSetName ADD = RemoteCopySetName*

command to add additional members.

Upon site failover, you must re-create the association sets and log units (see "Log Unit" on page 2–10) at the target (failover) site, using the attributes that were set at the initiator site.

FAIL_ALL

Because all copy sets within an association set are moved as a unit during failover and failback, all remote copysets within an association set must be owned by the same server.

If the FAIL_ALL switch is enabled and one member of the association set assumes the failsafe locked condition, all members of the association set assume the failsafe locked condition. The failsafe locked condition prevents further host access.

IMPORTANT: This applies only to remote copy sets with Failsafe Error Mode enabled. Failsafe Error Mode is enabled through the ERROR_MODE switch of the SET *RemoteCopySets* command.

When NOFAIL_ALL is specified, the members of the association set react independently to failsafe locked conditions. One member of the association set becoming failsafe locked has no effect on the other members of the association set. NOFAIL_ALL is the default setting.

This switch has no effect if all members of the association set have failsafe locked error mode disabled (normal error mode) or if there is only one remote copy set in the association.

Write History Logging

Write history logging means using a log unit (see "Log Unit" on page 2–10) to log a history of write commands and data from the host. Write history logging is used for *mini-merge* and *fast-failback*.

Mini-merge

If the target becomes inaccessible, the writes that would have gone to the target are logged to the association set's assigned log unit. An inaccessible target in this context refers to both links or target controllers going down. When the target becomes accessible, a full copy is not necessary. Only those host writes that were performed while the links were down are re-issued. This is referred to as a *mini-merge*. If a full copy was in progress at the time of the disconnect, write history logging is not initiated and the full copy is restarted when the target is accessible again.

Fast-Failback

During a planned failover, if write history logging has been enabled at the target site, then when the failback is performed, the initiator site is synchronized through a process called *fast-failback*. The writes are logged to the target site write history log. Then, during a fast-failback, the initiator site is updated from the write history log.

Log Unit Restrictions

Things to remember about log units include:

- Up to 12 log units can be assigned (12 possible remote copy sets).
- There can be only one log unit assigned to an association set.
- The log unit must be either a mirrorset or striped mirrorset.
- Host access must be disabled to create a log unit.
- Write-back caching must be disabled to create a log unit.
- Other unit settings must be the default settings to create a log unit.
- The log unit must reside at the current initiator site.
- Upon site failover, the log unit and association set must be reconfigured.
- The log unit cannot be a partitioned unit.
- The log unit must be a fixed size.

Choose the size of the log unit carefully. When the end of the log unit is encountered, a full copy is initiated when the link is restored. The amount of time before hitting the end of a log unit depends upon how long the links are down, how long a target backup takes, host write workload, the size of the log unit, and the number of remote copy sets actively logging to the same log unit.

Display the log status by using the following CLI command:

```
SHOW REMOTE_COPY FULL
```

Log Unit

The LOG_UNIT switch assigns a single, dedicated log unit for an association set.

NOTE: This switch is valid only if all members of the association set are in normal (not failsafe) error mode. Error mode is determined by the ERROR_MODE switch of the SET *RemoteCopySet* command.

IMPORTANT: When this command is entered, a header is immediately written to the log unit, which may make it difficult or impossible to recover any user data previously written on the unit. Take great care when you specify which unit should be the log unit.

If NOLOG_UNIT is specified, the association set's log unit is deassigned. NOLOG_UNIT is the default setting.

NOTE: A full copy occurs if you disable write history logging after logging operations have begun.

ORDER_ALL

When ORDER_ALL is enabled, the order of all asynchronous write operations across all members of the association set is preserved. No log unit is required.

With the ORDER_ALL switch enabled and write history logging enabled, if one member of the association set starts write history logging, all members of the association set start write history logging. This allows the mini-merge to re-play the writes in the same order received from the host.

If NOORDER_ALL is enabled, the members of the association set can start and finish write history independently. NOORDER_ALL is the default setting.

Failover

There are two types of failover:

- Planned failover (due to a planned take-down of one of the systems; for example to perform maintenance)
- Unplanned failover (due to a failure within the DRM system)

Planned Failover

A planned failover allows for an orderly shutdown of controllers. The host applications are quiesced and all write operations are permitted to complete before shutting down the controllers, so that no data is lost or jeopardized. A planned failover requires a synchronous operation mode.

> **NOTE:** To implement a planned failover while in asynchronous operation mode, you must first switch to synchronous operation.

Unplanned Failover

An unplanned failover does not allow for an orderly shutdown of controllers. An unplanned failover is initiated when:

- The initiator site is lost, or
- There is no host access, or
- There is no access to both initiator controllers.

> **NOTE:** If both links are severed, and the initiator configuration is functional, the system administrator must determine which site to use as the primary site.

Failback

The failback method (full copy or fast-failback) is determined by the enabling of Logging or Failsafe switches, the selected operation mode, and whether the failover is planned or unplanned as detailed in Table 2–1. Table 2–1also shows the availability of the Association Set switches, ORDER_ALL and FAIL_ALL.

Table 2–1 Data Replication Manager Switch Settings

Logging Enabled					Association Sets	
Logging	**Error Mode Failsafe**	**Operation Mode**	**Failover**	**Failback**	**Order All**	**Fail All**
Enabled	Disabled	Synchronous	Planned	Fast-Failback	Settable	Not Applicable for DRM
Enabled	Disabled	Synchronous	Unplanned	Full Copy	Settable	Not Applicable for DRM
Enabled	Disabled	Asynchronous (switch to Synchronous)	Planned	Fast-Failback	Settable	Not Applicable for DRM
Enabled	Disabled	Asynchronous	Unplanned	Full Copy	Settable	Not Applicable for DRM
NOTE: Logging is recommended for operations that can tolerate temporary loss of currency at the target site.						
Failsafe Enabled					**Association Sets**	
Logging	**Error Mode Failsafe**	**Operation Mode**	**Failover**	**Failback**	**Order All**	**Fail All**
Disabled	Enabled	Synchronous	Planned	Full Copy	Not Applicable for DRM	Settable
Disabled	Enabled	Synchronous	Unplanned	Full Copy	Not Applicable for DRM	Settable
Disabled	Enabled	Asynchronous	Planned	Full Copy	Settable	Settable
Disabled	Enabled	Asynchronous	Unplanned	Full Copy	Settable	Settable
NOTE: Failsafe is recommended for operations that can tolerate application halt during temporary target inaccessibility.						
Logging and Failsafe both Disabled: Not Recommended. Not Disaster Tolerant						
Logging and Failsafe both Enabled: Not Permitted. Logging and Failsafe may not be enabled simultaneously.						

Chapter 3

Getting Started

This chapter explains how to get your Data Replication Manager solution ready for setup.

> **NOTE:** It a good idea to keep a copy of this manual at both the initiator and target sites, to ensure a successful and identical setup at both sites. Two copies also eliminate confusion if more than one person is configuring Data Replication Manager.

This chapter covers the following topics:

Site, Host, and Solution Preparation

Before you start operating your DT subsystem, you must complete the following:

- Ensure that you have enough clearance to install and store the subsystems and have adequate power resources
- If you choose to use more than one cabinet, understand the proper methods for positioning and joining them
- Have the proper devices installed
- Verify that all of the BA370 components are in place

To learn more about adding additional storage, refer to the *Compaq StorageWorks RA8000/ESA12000 and MA8000/EMA12000 Fibre Channel Solution Software V8.5b for Tru64 UNIX Installation Reference Guide*.

Host Bus Adapter Requirements

To run your Data Replication Manager solution, you must have two host bus adapters installed into your host system. Refer to the *Compaq KGPSA-BC PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide* that came with your adapter for detailed information on this hardware.

At this time, it is important to locate and record the World Wide Names of each host bus adapter. For the host bus adapter at the target site, you can record the worldwide name in the worksheet provided in Chapter 4, "Configuring a Data Replication Manager Solution". The initiator site host bus adapter World Wide Names can be recorded in the worksheet found in Chapter 4, "Configuring a Data Replication Manager Solution". You must have this number handy when you rename the host connections in the Chapter 4.

NOTE: The World Wide Name can be found on the bottom of the host bus adapter board. Look for a small bar code label with an IEEE (Institute of Electrical and Electronics Engineers) precursor. World Wide Name example: 1000-0000-C920-A5BA.

Setting Up the Fibre Channel Switches

The Fibre Channel gigabit switches must be in place before the DRM subsystems can be cabled and configured. You need the following to install your Fibre Channel switches:

- Power cord
- 10BASE-T cable with RJ45 plug (to be connected to an Ethernet hub or switch)
- Fixed IP address and subnet mask (one of each per switch)

The Ethernet cable and IP address are required to monitor and administer the Fibre Channel switch. Configure the Ethernet IP address and the Ethernet IP subnet mask with the front panel buttons of the Fibre Channel switch. Refer to the *Compaq StorageWorks Fibre Channel Storage Switch User's Guide* for more details.

Once the Ethernet IP settings are established, perform the following steps:

1. Update the *\winnt\system32\drivers\etc\hosts* file with the IP address and the name of the Fibre Channel switch.

2. *Ping* using the Ethernet IP address of the switch. If this is successful, you have access to the switch.
3. *Ping* using the name of the switch. This verifies the operation of the name resolution.
4. *Telnet* into the switch (Username = **admin** and password = **password** [default setting]). Refer to the *Compaq StorageWorks Fibre Channel SAN Switch Management Guide* for Telnet session procedures. Make the following adjustments to the switch:
   - ❏ Type **switchName** to configure the switch name. Be sure to designate a name that enables you to easily identify the switch you are trying to access.

     Example: switchName *NewSwitchName*
   - ❏ Type **switchShow** to reveal the status of the switch and some of its ports.
   - ❏ Type **version** to display the firmware levels. You must be running version 1.6B or higher.
5. Using a Java-capable browser, go to http://<FC switch DNS name> to view a visual representation of the switch. (You need to know the Domain Server Name.) You can double-click on this picture for further information.

Setting Up the Fiber Optic Cables

Before you connect the fiber optic cables to your subsystems, it is important to understand the designated names of each component. See Table 3–1 at the end of this chapter for an overview list of required connections for each site and between the sites.

> **NOTE:** For instructions on making connections, refer to Chapter 4, "Configuring a Data Replication Manager Solution."

Figure 3–1 shows how each component is referenced in this document.

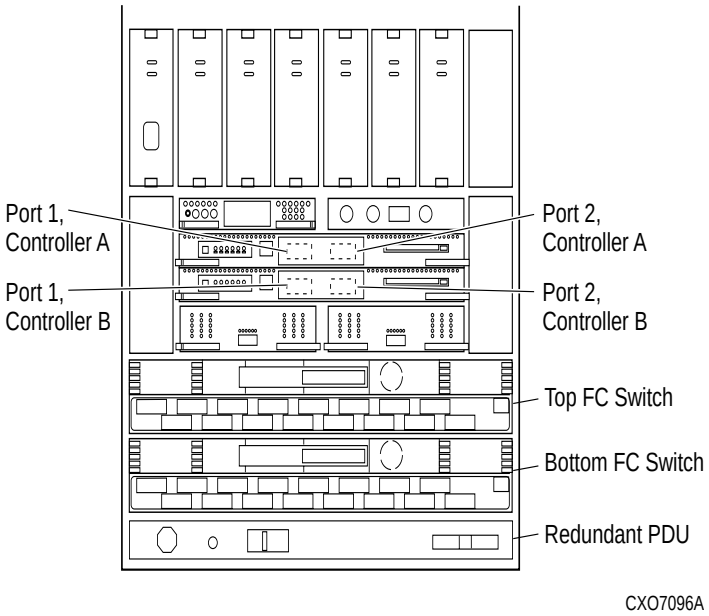

Figure 3–1. Locations and names of components for connecting fiber optic lines

Before you connect the fiber optic cables, Compaq recommends that you tag each end of the cables with the following information:

Host-to-Switch Connections

- Host name and rank number or PCI slot number of the host bus adapter
- Switch name and port number on switch

Switch-to-Controller Connections

- Fibre Channel switch name (top or bottom)
- Fibre Channel switch port number (0-15)
- Site name (initiator or target)
- Controller name
- Controller port number (1 or 2)
- Host port number

- Host Bus Adapter worldwide name

The DT solution requires two different types of fiber optic cables, depending on where the connections are made. Cabling at each individual site that involves the controller, the switch, and the host is made with 50 micron multi-mode fiber optic cables. The maximum length that these cables support is 500 meters. When cabling between initiator and target sites that are more than 500 meters apart and using GBIC, you must use a 9 micron single-mode fiber optic cable, which can run a distance of up to 10 kilometers.

> **NOTE:** The 9-micron single mode fiber optic cable may also be referenced by some manufacturers as an 8.3-micron cable. In addition, to increase the reliability of the cable or to reduce the likelihood of having to re-pull or re-install the cable over a long distance, Compaq recommends that you use multiple conductor cable.

Data Replication Manager uses other long distance transport modes. For connection information, refer to: http://www.compaq.com/storage.

CAUTION: If the Fibre Channel cable is not properly connected to the controller, failure may result. Because of the cable's frail nature, it must be regularly maintained or its performance and life span will be affected. Before continuing, make sure that you follow the precautions listed in the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide*.

Table 3–1 provides an overview of the connections that you need to make at each site and between the sites. Specific connection information is covered in more detail in Chapter 4, "Configuring a Data Replication Manager Solution." Do not make any connections at this time.

Table 3–1 Overview of Required Connections

Initiator Site			Target Site		
Host Port 1	→	Top Switch, Port 0	Host Port 1	→	Top Switch, Port 0
Host Port 2	→	Bottom Switch, Port 0	Host Port 2	→	Bottom Switch, Port 0
Controller A, Port 1	→	Top Switch, Port 2	Controller A, Port 1	→	Top Switch, Port 2
Controller A, Port 2	→	Top Switch, Port 4	Controller A, Port 2	→	Top Switch, Port 4
Controller B, Port 1	→	Bottom Switch, Port 2	Controller B, Port 1	→	Bottom Switch, Port 2
Controller B, Port 2	→	Bottom Switch, Port 4	Controller B, Port 2	→	Bottom Switch, Port 4

Between Initiator and Target Sites				
Top Switch, Port 6	→	External Fiber Link	←	Top Switch, Port 6
Bottom Switch, Port 6	→	External Fiber Link	←	Bottom switch, Port 6

Chapter 4

Configuring a Data Replication Manager Solution

This chapter provides procedures for configuring your Data Replication Manager (DRM) solution. Since a DRM system spans multiple sites, you must configure the DRM system at each site.

The procedures take you through the configuration process. Set up the target site first and then the initiator site. The setup for each site is similar. At each site, you must configure the controllers by defining controller characteristics specific to DRM. You must then define storagesets, units, remote copy sets, and association sets. After the controllers are configured, you will make fiber optic cable connections between the controllers and switches. Finally, you will install the necessary software and drivers on each host.

This chapter covers the following topics:

IMPORTANT: In this chapter, *initiator* site procedures appear in *shaded text*. This distinguishes them visually from *target* site procedures, which are *not shaded*.

Introduction

The disaster tolerant (DT) configuration that supports Data Replication Manager (DRM) involves two HSG80 Array Controller subsystems—one at an initiator site and one at a target site.

IMPORTANT: Because of the complexity of the configuration process, it is a good idea to have all Data Replication Manager documentation available at both sites to eliminate confusion and minimize the risk of error. Please follow the steps precisely in the order provided in this documentation.

Figure 4–1 depicts a basic DRM configuration that is referenced throughout this chapter.

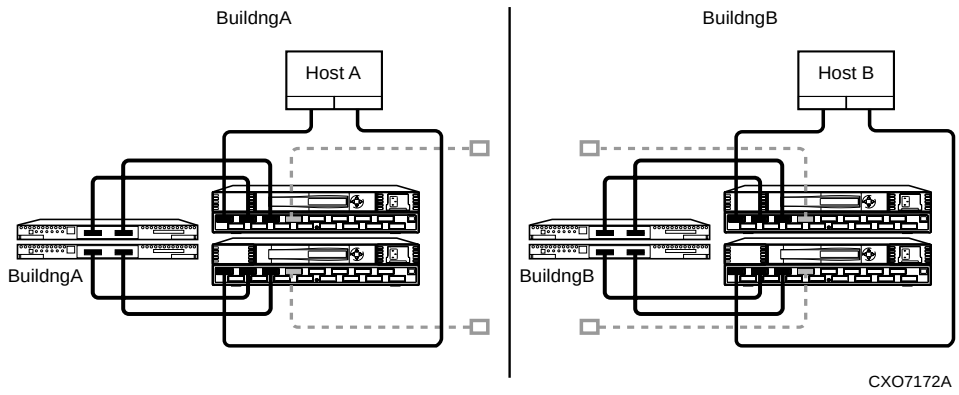

Figure 4–1. Data Replication Manager basic configuration

Restrictions

It is important to understand the operating restrictions before configuring your Data Replication Manager solution. Table 4–1 lists the points to consider when proceeding to the configuration process:

Table 4–1 Restrictions

Restriction	Comments
Two HSG80 controller pairs are required.	All controllers must run ACS Version 8.5P. HSG60 controller pairs are not supported.
Four Fibre Channel Switches are required.	These switches provide two separate fabrics connecting controllers at the initiator and target sites.
HSG80 controllers must be configured in Multiple-Bus Failover mode.	Additional software support required on the host: Must run TRU64 Version 5.0a or Version 5.1 with the appropriate patch kit.
HSG80 controllers can be configured for Fibre Channel Switched protocol or SCSI-3 protocol.	Host operating system and adapter must support Fibre Channel Switched protocol and SCSI-3 protocol.
Mirrored write-back cache must be enabled.	Requires 512 MB cache per controller (256 MB effective capacity once mirrored).
Maximum configuration: ■ 12 host equivalents per system ■ 6 host bus adapters (HBA) per host ■ 24 units per host ■ 4 subsystems per site.	
Maximum of 12 remote copy sets allowed per HSG80 controller pair.	If more than 12 remote copy sets are needed, additional subsystems are required.
Maximum of 2 members allowed per Remote Copy Set.	Composed of 1 initiator and 1 target unit.
Target unit cannot reside on the same controller pair as its initiator unit.	One controller pair required for initiator; one controller pair required for target.
Controller replication conducted through port 2 on each controller.	■ Link between initiator and target site is made through Port 2 ■ Both links must be up when DRM setup is configured

Table 4–1 Restrictions (Continued)

Restriction	Comments
Maximum 64 connections	Effective number of connections is 64 minus the 4 default remote copy connections. Over 64 connections may require use of switch zoning to restrict visible devices.
It is not possible to run DILX on units used by remote copy sets.	Run DILX prior to creating the Remote Copy Set configuration.
The LUN/unit at the initiator and target sites must be identical.	Keep the unit number, RAID level, disk geometry used, and other parameters the same to eliminate confusion and the risk of error.
Controller-based partitions are not supported within remote copy sets.	Host software may be capable of partitioning units.
Units at the initiator and target sites cannot be transportable units.	Units cannot be moved to non-controller configurations without potential data loss.
Cannot use FRUTIL on remote site while I/O is in progress to target site.	
Max_cached_transfer_size should be set to **one** on target units and to whatever is optimal for the Initiator host applications.	Write-behind caching allows for the best remote copy performance.
Log units must: ■ Reside at the initiator site ■ Not be moved to the target site ■ Not be a partitioned unit ■ Have write-back cache disabled ■ Have access disabled ■ Must be re-created at target site after failover ■ Must be a mirrorset.	
Maximum of 12 Non-Remote Copy Set LUNs at Initiator and Target sites.	
ACS 8.5F and ACS 8.5S may co-exist on the same Storage Area Network with a DRM configuration using ACS 8.5P.	
The LP7000 and LP8000 host bus adapters may co-exist on the same DRM Storage Area Network	

Configuration Overview

Both the initiator and target sites need some type of command line interface (CLI) to the controller. You can connect the serial maintenance port of both the initiator and target site controllers to a terminal from which you issue CLI commands. You can also start a terminal emulator session on Windows 2000. Use the HyperTerminal emulator. Settings are: 9600 baud, 8 bits, No parity, 1 stop bit, XON/XOFF.

The procedures for configuring a disaster tolerant (DT) system using the controller's serial maintenance port and CLI commands are listed in the Configuration Procedures section under "Target Site Procedures" and "Initiator Site Procedures."

Configuration Procedures Outline

IMPORTANT: In this chapter, *initiator* site procedures appear in *shaded text*. This distinguishes them visually from *target* site procedures, which are *not shaded*.

Target Site

- Configure the controllers at the target site
- Configure the storage at the target site
  - Configure devices and StorageSets
  - Configure LUNs
  - Disable access to units
- Connect fiber optic cables between the controllers and switches
- Connect the target site to the external fiber link
- Configure the host
  - Install the host bus adapters and drivers
  - Set up the operating system
  - Install SWCC version 2.3 (optional)
  - Connect fiber optic cables between the hosts and the switches
  - Rename the host connections on the controllers

Initiator Site

- **Configure the controllers at the initiator site**
- **Configure the storage at the initiator site**
  - **Configure devices and StorageSets**
  - **Configure LUNs**
- **Connect fiber optic cables between the controllers and switches**
- **Connect the initiator site to the external fiber link**
- **Create controller Connections**
- **Enable controller connections on target and initiator units**
- **Create remote copy sets**
- **Create log unit and association sets (optional)**
- **Set failsafe at the initiator site (optional)**
- **Configure the host**
  - **Install the host bus adapters and drivers**
  - **Install operating system**
  - **Install SWCC version 2.3 (optional)**
  - **Connect fiber optic cables between the hosts and the switches**
  - **Rename the host connections on the controllers**
  - **Enable access to the hosts on the Initiator controllers**
- **Verify system operation**

Each of these steps is detailed in the following sections.

Configure the Controllers at the Target Site

Prior to configuring the controllers at the target site, be sure to follow these preparatory steps:

- Identify the World Wide Name on the host bus adapters.
- Establish the names to assign to the target and initiator sites. Use a naming scheme that will be meaningful, such as building or city names; for example, Initiator site = *BuildngA* and Target site = *BuildngB*.

To get your DT system up and running you must set up and configure the controllers. These tasks are outlined in the following procedure:

1. Ensure that all BA370 enclosures, Fibre Channel switches, Power Distribution Units (PDUs), and the main power supply are off.
2. Plug all cabinet PDU power cords into the main power receptacles.
3. Be sure that you have a serial connection to each maintenance port of each controller.
4. Apply power to the main power source.
5. Turn on all PDUs.
6. Ensure that the Fibre Channel switches are powered on, but not cabled.

   NOTE: When the BA370 enclosures are turned on, the controllers boot if the PCMCIA cards are already installed. If there are no cards in the controller slots, insert them now and press the reset button.

7. Turn on the BA370 enclosures.
8. Establish a local connection to the controllers. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide* for instructions.
9. Verify that all controllers are on and functional by looking for the CLI prompt on the maintenance port of each controller.

   NOTE: Unless otherwise noted, all operations may be conducted from controller A.

10. Issue the following CLI command:

    SHOW THIS_CONTROLLER

    You will see a display similar to that in “Example Display 1.”

Example Display 1

```
Controller:
   HSG80 ZG8nnnnnnn Software V85P, Hardware  E03
   NODE_ID          = nnnn-nnnn-nnnn-nnnn
   ALLOCATION_CLASS = 0
   SCSI_VERSION     = SCSI-2
   Not configured for dual-redundancy
   Controller misconfigured -- other controller present
   Device Port SCSI address 7
   Time: NOT SET
   Command Console LUN is disabled
Host PORT_1:
   Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
   PORT_1_TOPOLOGY  = OFFLINE (offline)
Host PORT_2:
   Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
   PORT_2_TOPOLOGY  = OFFLINE (offline)
   NOREMOTE_COPY
Cache:
   512 megabyte write cache, version 0012
   Cache is GOOD
   No unflushed data in cache
   CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
Mirrored Cache:
   Not enabled
Battery:
   FULLY CHARGED
   Expires:   . . . . . .
   NOCACHE_UPS
Controllers misconfigured. Type SHOW THIS_CONTROLLER
```

11. Verify that the subsystem World Wide Name, also called the NODE_ID, is set. (If zeros are displayed, the name is not set.) If the name is set, go to Step 15. If the World Wide Name has not been assigned to the controller, obtain the name and set it before proceeding.

    NOTE: The subsystem's World Wide Name and checksum can be found on a sticker, which is located on top of the frame that houses the controllers, EMU, PVA, and cache modules. This sticker also includes a checksum which is required to verify that the World Wide Name is valid. If there is no label there, contact your Compaq customer service representative for assistance. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide* for more information on World Wide Names. Each subsystem's World Wide Name begins with 5000 and ends in zero, for example 5000-1FE1-FFOC-EEOO. The controller port IDs are derived from the World Wide Name.

    

    CAUTION: Never set two subsystems to the same World Wide Name, or data corruption will occur.

12. After the World Wide Name has been located, assign it to the controller using the following CLI command:

    SET THIS NODE_ID=*node_ID checksum*

    You will see a display similar to that in "Example Display 2."

Example Display 2

```
Warning 4000: A restart of this controller is required before all the
 parameters modified will take effect
%CER--HSG80> --09-FEB-1999 10:07:54-- Restart of this controller required
Restart of this controller required
```

13. Issue a SHOW THIS_CONTROLLER command to verify that the World Wide Name has been set.

    You will see a display similar to that in "Example Display 3."

Example Display 3

```
Controller:
            HSG80 ZG8nnnnnnn Software V85P, Hardware  E03
            NODE_ID          = yyyy-yyyy-yyyy-yyyy
            ALLOCATION_CLASS = 0
            SCSI_VERSION     = SCSI-2
            Not configured for dual-redundancy
                Controller misconfigured -- other controller present
            Device Port SCSI address 7
            Time: NOT SET
            Command Console LUN is lun 0 (NOIDENTIFIER)
               .
               .
               .
```

14. Configure the controllers for multiple bus failover mode by issuing the following CLI command:

    SET MULTIBUS_FAILOVER COPY = THIS_CONTROLLER

    This command automatically restarts the "other" controller.

    You will see a %LFL and a %EVL prompt. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide* for more details on these reports.

15. To ensure that the setting from step 14 has been applied, enter:

    SET THIS SCSI = SCSI-3

    NOTE: D0 can no longer be used as a device LUN.

16. To ensure that the setting from step 15 has been applied, enter:

    SHOW OTHER_CONTROLLER FULL

    A display shows that the controllers have been configured to support multiple bus failover mode. You should see a display similar to that in "Example Display 4."

Example Display 4

```
Controller:
   HSG80 ZG8nnnnnnn Software V85P, Hardware  E03
   NODE_ID          = nnnn-nnnn-nnnn-nnnn
   ALLOCATION_CLASS = 0
   SCSI_VERSION     = SCSI-3
   Configured for MULTIBUS_FAILOVER with ZG8nnnnnnn
   In dual-redundant configuration
   Device Port SCSI address 7
   Time: NOT SET
   Command Console LUN is lun 0 (NOIDENTIFIER)
   .
   .
```

NOTE: These settings are automatically applied to controller B. It is not necessary to repeat these steps on controller B.

17. Verify that the settings have been accepted on controller B by issuing the following CLI command:

    SHOW OTHER_CONTROLLER FULL

18. Change your controller prompts to help you easily identify which controller you are working on. Enter the following CLI commands:

    SET THIS_CONTROLLER PROMPT=*"TargetControllerNameTop> "*

    SET OTHER_CONTROLLER PROMPT=*"TargetControllerNameBottom> "*

    Example: SET THIS_CONTROLLER PROMPT="BuildngBTop> "

    Example: SET OTHER_CONTROLLER PROMPT="BuildngBBottom> "

    NOTE: This step takes effect immediately.

19. Check to see whether the mirrored write-back cache is enabled by issuing the following CLI command:

    SHOW THIS_CONTROLLER

    You will see a display similar to that in "Example Display 5."

Example Display 5

```
          Mirrored Cache:
              Not enabled
```

.
.
.

If mirrored cache is not enabled, issue the following CLI command:

```
SET THIS_CONTROLLER MIRRORED_CACHE
```

The controllers restart after mirrored write-back cache has been set and you see %LFL and %EVL displays.

NOTE: It may take up to five minutes after restart to check cache. The controllers reject this command until the cache check is complete. If this command is rejected, do not restart the controllers. Wait a few minutes and then try again.

20. After the controllers restart, issue the following CLI command to confirm that mirrored write-back cache is enabled:

```
SHOW THIS_CONTROLLER
```

You will see a display similar to that in "Example Display 6."

Example Display 6

```
.
.
.
Mirrored Cache:
256 megabyte write cache, version 0012
Cache is GOOD
No unflushed data in cache
.
.
.
```

NOTE: It is not necessary to repeat this step on controller B.

21. Set the fabric topology for each port on both controllers by issuing the following CLI commands:

    NOTE: You are prompted to restart the controllers after each command, but you do not need to restart the controllers until all topologies have been set.

    SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC

    SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC

    SET OTHER_CONTROLLER PORT_1_TOPOLOGY=FABRIC

    SET OTHER_CONTROLLER PORT_2_TOPOLOGY=FABRIC

22. To verify that the topology is set correctly, issue the following CLI commands:

    SHOW THIS_CONTROLLER

    SHOW OTHER_CONTROLLER

You will see a display similar to that in "Example Display 7."

Example Display 7

```
.
.
.
Host PORT_1:
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
  . . . . . . . . . . . . . . . PORT_1_TOPOLOGY  = CONNECTION DOWN
Address . . . . . . . . . . . . . . . . . . . . . . . . . . . .=nnnnnn
Host PORT_2:
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
  . . . . . . . . . . . . . . . PORT_2_TOPOLOGY  = CONNECTION DOWN
Address . . . . . . . . . . . . . . . . . . . . . . . . . . . .=nnnnnn
NOREMOTE_COPY
.
.
.
```

23. You are now ready to enable Data Replication Manager. Issue the following CLI command:

    SET THIS_CONTROLLER REMOTE_COPY=*TargetControllerName*

    Example: SET THIS_CONTROLLER REMOTE_COPY=BuildngB

    NOTE: Be sure to specify a meaningful TargetName, such as a name that reflects the target node's location. **Do not use "local" and "remote"; these are reserved keywords.** The name can be up to eight characters and must be unique to all of your controllers. Follow the naming guidelines specified in the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide.*

    After issuing this CLI command, you will see a series of %LFL and %EVL displays; the controllers automatically restart.

24. Issue the following CLI command to verify that these settings are in place:

    SHOW THIS_CONTROLLER

You will see a display similar to that in "Example Display 8."

Example Display 8

```
.
.
Host PORT_2:
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
 . . . . . .  PORT_2_TOPOLOGY  = FABRIC (up)
REMOTE_COPY = BuildngB
```

Configure Storage at the Target Site

Devices and StorageSets

Before you can configure the storage for Data Replication Manager, you need to add disks, create the storagesets, and create units. Follow the instructions in the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide*, but note the restrictions listed in Table 4–1.

NOTE: The target site must have exactly the same storageset and unit configuration as the initiator site.

Creating Storage Units

Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide* for detailed information on configuring units.

NOTE: Issue the following command to create storage units and to disable all access as the units are created. Disabling access prevents any TruCluster members from applying a Persistent Reserve to the new units.

ADD UNIT *UnitName StorageSetName* DISABLE_ACCESS_PATH=ALL

NOTE: Issue this command for all units.

After all units have been created, execute the following procedure:

1. Set the maximum cached transfer size to 1:

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = 1

   Repeat this step for each unit.

2. Verify that the access on each unit is set to *none*:

   SHOW UNITS FULL

You will see a display similar to that in “Example Display 9.”

Example Display 9

```
LUN                                     Uses             Used by
------------------------------------------------------------------------------

D110 DISK1000
   LUN ID:      nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn
   NOIDENTIFIER
   Switches:
   RUN                    NOWRITE_PROTECT        READ_CACHE
   READAHEAD_CACHE        WRITEBACK_CACHE
   MAXIMUM_CACHED_TRANSFER_SIZE = 1
   Access:
   NONE
   State:
   ONLINE to this controller
   Not reserved
   NOPREFERRED_PATH
   Size: nnnnnnnn blocks
   Geometry (C/H/S): ( 7000 / 20 / 254 )
.
.
.
```

3. Distribute the units by setting their preferred path. Use either of the following CLI commands:

   SET *Unit*PREFERRED_PATH=THIS_CONTROLLER

   or

   SET *Unit*PREFERRED_PATH=OTHER_CONTROLLER

   Keep the busiest units on different host ports.

4. Restart the controllers after configuring the units. Otherwise, the preferred path settings do not go into effect.
5. To ensure that your storage settings are in place, issue the following CLI command:

   SHOW *Storagesets* FULL

Cabling the Target Site

Connect Fiber Optic Cables Between the Controllers and Fiber Channel Switches

This procedure specifies how to make fiber optic connections. Figure 4–2 shows the port locations on the Fibre Channel switch.

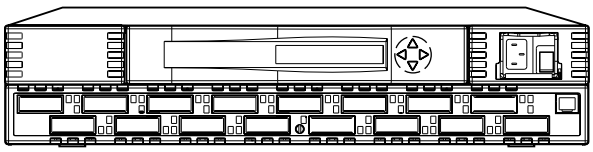

CXO7085A

Figure 4–2. Switch port locations

1. Connect a multi-mode, 50-micron fiber optic cable from port 1 of the top controller to port 2 of the top Fibre Channel switch.
2. Connect a second multi-mode, 50-micron fiber optic cable from port 2 of the top controller to port 4 of the top Fibre Channel switch.
3. Connect a third multi-mode, 50-micron fiber optic cable from port 1 of the bottom controller to port 2 of the bottom Fibre Channel switch.
4. Connect a fourth multi-mode, 50-micron fiber optic cable from port 2 of the bottom controller to port 4 of the bottom Fibre Channel switch.

NOTE: You will see an illuminated green LED on the switch as soon as the cable is inserted at both ends. This verifies that there is a good connection.

Figure 4–3 illustrates what the cabling should look like. The numbered callouts reflect the steps just completed.

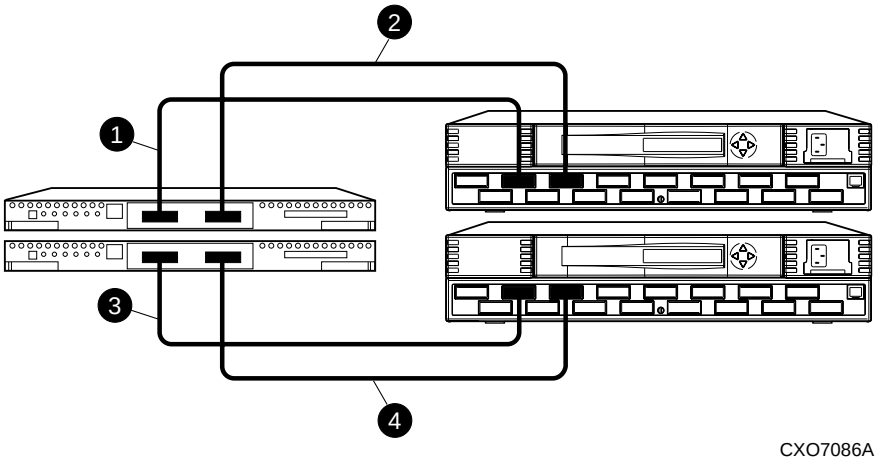

Figure 4–3. Cabling between the controllers and switches

Connect the Target Site to the External Fiber Link

Locate the connection points at the target site that link the target site to the initiator site. Look for either a fiber optic cable connector or a patch panel where you can insert a cable.

Long Wave GBICs

1. Connect a single-mode, 9-micron fiber optic cable from port 6 of the top switch to one connection point.
2. Connect another single-mode, 9-micron fiber optic cable from port 6 of the bottom switch to the other connection point.

Other Transport Modes

Data Replication Manager uses other long-distance transport modes. For connection information, refer to our website at:

http://www.compaq.com/products/storageworks/

The target site is now physically linked to the initiator site. Figure 4–4 shows this cabling.

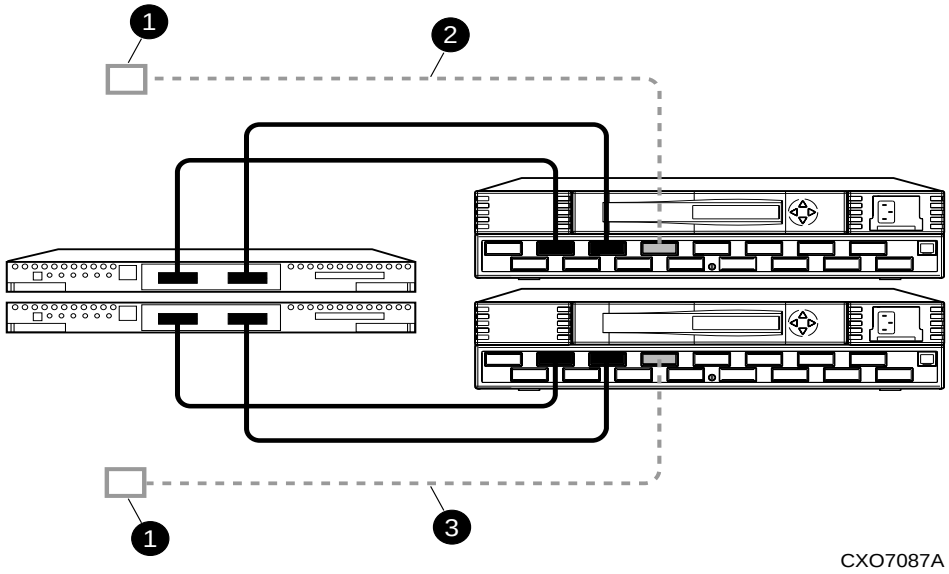

Figure 4–4. Cabling from the target site to the initiator site

Configure the Host at the Target Site

To run Data Replication Manager, you should have two host bus adapters installed in each host system. Follow the procedures outlined in the next section.

NOTE: Changes made to the topology settings do not take effect until you reboot the host.

Install SWCC (optional)

For detailed information about SWCC, refer to the *Compaq StorageWorks Command Console User Guide*.

Connect Fiber Optic Cables Between the Hosts and the Switches

Connecting the fiber optic cables requires two stages:

- Connecting the Cables
- Renaming the Host Connections

Both are described in the following sections.

Connecting the Fiber Optic Cable

1. Establish the cabling policy that you plan to follow.
2. Connect a multi-mode, 50-micron fiber optic cable from port 0 of the top switch to one adapter on a host.
3. Connect the final multi-mode, 50-micron fiber optic cable from port 0 of the bottom switch to the other adapter on the same host.

   NOTE: You may choose any available port to connect your cables to, but you must maintain that identical scheme at the initiator site. Therefore, if port 1 of controller B is connected to port 2 of the bottom switch at the target site, then port 1 of controller B must be connected to port 2 of the bottom switch at the initiator site.

4. If you have more than one host (up to twelve), connect one host bus adapter to one of the remaining ports on the top switch. Connect the other host bus adapter to the same numbered port on the bottom switch.

The host is now connected to the target site switches via the multi-mode, 50-micron fiber optic cables. Your cabling should appear as shown in Figure 4–5.

5. Verify that the connection between the host and the switch has been made by issuing the CLI command:

   SHOW CONNECTIONS

   You will see a display similar to that in “Example Display 10.”

   NOTE: You can also verify that a connection has been made by looking for the illuminated green LED that flashes on the switch ports.

.

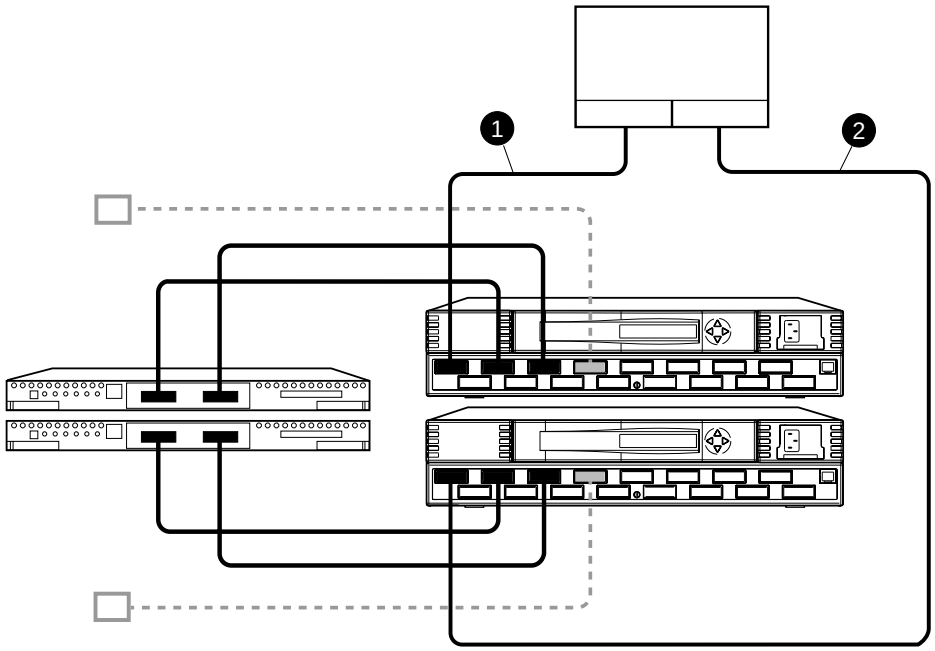

Figure 4–5. Cabling between the hosts and the switches

Example Display 10

```
Connection Unit
Name Operating system   Controller  Port    Address    Status   Offset

!NEWCON00 THIS       1      210013 online   . . . . . .0
  HOST_ID=nnnn-nnnn-nnnn-nnnn   . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

!NEWCON01 OTHER 1      200013 online  . . . . . .0
  HOST_ID=nnnn-nnnn-nnnn-nnnn   . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn
```

6. For Tru64 UNIX systems, change the operating system for each connection. Enter the following CLI command:

   SET !NEWCONnn OPERATING_SYSTEM = TRU64_UNIX

You should see a display similar to that in "Example Display 11."

Example Display 11

```
Connection Unit
Name      Operating system    Controller  Port    Address    Status Offset

!NEWCON00      TRU64_UNIX        THIS       1      210013 ol this. . 0
  HOST_ID=nnnn-nnnn-nnnn-nnnn  . . . .  ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

!NEWCON01       TRU64_UNIX        OTHER      1      200013 ol other. 0
  HOST_ID=nnnn-nnnn-nnnn-nnnn  . . . .  ADAPTER_ID=nnnn-nnnn-nnnn-nnnn
```

Rename the Host Connections

To better identify which hosts you are working with, Compaq recommends that you rename the host connections, using a meaningful connection name for each one. Each host bus adapter appears as a connection. An individual host bus adapter can be identified by its World Wide Name, which you recorded in the Chapter 3, "Getting Started," and which appears in the connection description.

Initially, each connection is named !NEWCONxx. It is much easier to track connections if the connection names are meaningful; HostA1, for example.

1. We suggest that you use the worksheet in Figure 4–6 when renaming your hosts. Fill in the fields accordingly to prepare for renaming your connections.

!NEWCON*xx*	World Wide Name	Host Name	Path Number
00	1000-XXXX-XXXX-XXXX	Tru 7 bot	F254
01	1000-XXXX-XXXX-XXXX	Tru 7 top	F156

Figure 4–6. Host renaming worksheet

2. When you have completed the worksheet, rename the connections using the following CLI commands:

   RENAME !NEWCONxx *TargetHostConnectionNamex*

   RENAME !NEWCONxx *TargetHostConnectionNamey*

   Example: RENAME !NEWCONxx hostA1

   Example: RENAME !NEWCONxx hostA2

3. When you have finished renaming your host connections, issue the following command to see your new settings:

   SHOW CONNECTIONS

You will see a display similar to that in "Example Display 12."

Example Display 12

```
Connection Unit
Name      Operating system    Controller  Port    Address    Status Offset

HostA1TRU64_UNIX        THIS       1      210013 online   . . . . . .0
  HOST_ID=nnnn-nnnn-nnnn-nnnn   . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

HostA2 TRU64_UNIX        OTHER      1      200013 online  . . . . . .0
  HOST_ID=nnnn-nnnn-nnnn-nnnn   . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn
```

Configure the Controllers at the Initiator Site

Controller Pre-Configuration Procedure

Before configuring the controllers at the initiator site, be sure to follow these preparatory steps:

- Identify the World Wide Name on the host bus adapters.
- Establish the name to be assigned to the initiator site. This name should be different from the one you assigned to the target site.

NOTE: Initiator site procedures are shown as shaded text to distinguish them from the similar target site procedures.

Controller Configuration Procedure

The first step to getting your DT system up and running involves setting up and configuring the controllers. These tasks are outlined below:

1. **Ensure that all BA370 enclosures, Fibre Channel switches, Power Distribution Units (PDUs), and the main power supply are off.**
2. **Plug all cabinet PDU power cords into the main power receptacles.**
3. **Be sure that you have a serial connection to each of the controllers.**
4. **Apply power to the main power source.**
5. **Turn on all PDUs.**
6. **Ensure that the switches are powered on, but not cabled.**

   NOTE: When the BA370 enclosures are turned on, the controllers boot if the PCMCIA cards are already installed. If there are no cards in the controller slots, insert them now, and press the reset button. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide* for complete instructions on how to properly seat the controller cards.
7. **Turn on the BA370 enclosures.**
8. **Establish a local connection to the controller. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide* for instructions.**

9. **Verify that all controllers are on and functional by looking for the CLI prompt on the maintenance port terminal.**

   NOTE: Unless otherwise noted, all operations may be conducted from controller A.

10. **Issue the following CLI command:**

    SHOW THIS_CONTROLLER

    You will see a display similar to that in "Example Display 13."

Example Display 13

```
Controller:
            HSG80 ZG8nnnnnnn Software V85P, Hardware  E03
            NODE_ID          = nnnn-nnnn-nnnn-nnnn
            ALLOCATION_CLASS = 0
            SCSI_VERSION     = SCSI-2
            Not configured for dual-redundancy
                Controller misconfigured -- other controller present
            Device Port SCSI address 7
            Time: NOT SET
            Command Console LUN is disabled
    Host PORT_1:
            Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
PORT_1_TOPOLOGY  = OFFLINE (offline)
    Host PORT_2:
            Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
PORT_2_TOPOLOGY  = OFFLINE (offline)
            NOREMOTE_COPY
Cache:
           512 megabyte write cache, version 0012
           Cache is GOOD
           No unflushed data in cache
            CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
```

```
 Mirrored Cache:
            Not enabled


Battery:
           FULLY CHARGED
            Expires:


            NOCACHE_UPS
Controller misconfigured. Type SHOW THIS_CONTROLLER
```

NOTE: The subsystem's World Wide Name and checksum can be found on a sticker located on top of the frame that houses the controllers, EMU, PVA, and cache modules. This sticker also includes a checksum that is required to verify that the World Wide Name is valid. If there is no label there, contact your Compaq customer service representative for assistance. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide* for more information on World Wide Names. Each subsystem's World Wide Name begins with 5000 and ends in zero: for example, 5000-1FE1-FF0C-EE00. The controller port IDs are derived from the World Wide Name.

11. **Verify that the subsystem World Wide Name is set. If it is, go to step 14. If the World Wide Name has not been assigned to the controller, you must obtain the name and set it before proceeding.**

    NOTE: Never set two subsystems to the same World Wide Name or data corruption will occur.

12. **After the World Wide Name has been located, assign it to the controller using the following CLI command:**

    SET THIS NODE_ID=*node_ID checksum*

    You will see a display similar to that in "Example Display 14."

Example Display 14

```
Warning 4000: A restart of this controller is required before all the
   parameters modified will take effect
%CER--HSG80> --09-FEB-1999 10:07:54-- Restart of this controller required
Restart of this controller required
```

13. **Issue a SHOW THIS CONTROLLER command to verify that the World Wide Name has been set.**

    You will see a display similar to that in "Example Display 15."

Example Display 15

```
Controller:
            HSG80 ZG8nnnnnnn Software V85P, Hardware  E03
            NODE_ID          = nnnn-nnnn-nnnn-nnnn
            ALLOCATION_CLASS = 0
            SCSI_VERSION     = SCSI-2
            Not configured for dual-redundancy
                Controller misconfigured -- other controller present
            Device Port SCSI address 7
            Time: NOT SET
            Command Console LUN is disabled
```

14. **Configure the controllers for multiple bus failover mode by issuing the following CLI command:**

    SET MULTIBUS_FAILOVER COPY=THIS_CONTROLLER

    This command automatically restarts the "other" controller.

    %LFL and %EVL prompts are displayed. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide* for more details on these reports.

15. Enter the following CLI command:

SET THIS SCSI = SCSI-3

NOTE: D0 can no longer be used as a device LUN.

16. To ensure that the settings from step 15 have been applied, enter:

SHOW THIS_CONTROLLER FULL

Check the display to verify that the controllers have been configured to support multiple bus failover mode. You will see a display similar to that in "Example Display 16."

Example Display 16

```
Controller:
            HSG80 ZG8nnnnnnn Software V85P, Hardware  E03
            NODE_ID          = nnnn-nnnn-nnnn-nnnn
            ALLOCATION_CLASS = 0
            SCSI_VERSION     = SCSI-3
            Configured for MULTIBUS_FAILOVER with ZG8nnnnnnn
                In dual-redundant configuration
            Device Port SCSI address 7
            Time: NOT SET
            Command Console LUN is lun 0 (NOIDENTIFIER)
```

NOTE: These settings are automatically applied to controller B. Therefore, it is not necessary to repeat these steps again on controller B.

17. Verify that the settings have been accepted on controller B by issuing the following CLI command:

SHOW OTHER_CONTROLLER FULL

18. Change your controller prompts to help you easily identify which component you are working on. Enter the following CLI commands:

SET THIS_CONTROLLER PROMPT=*"InitiatorControllerNameTop> "*

SET OTHER_CONTROLLER PROMPT=*"InitiatorControllerNameBottom> "*

Example: SET THIS_CONTROLLER PROMPT="BuildngATop> "

Example: SET OTHER_CONTROLLER PROMPT="BuildngABottom> "

NOTE: This step takes effect immediately

19. Check to see if mirrored write-back cache is enabled by issuing the following CLI command:

SHOW THIS_CONTROLLER

You will see a display similar to that in "Example Display 17."

Example Display 17

```
           Mirrored Cache:
               Not enabled
.
.
.
```

If it is not enabled, issue the following CLI command:

SET THIS_CONTROLLER MIRRORED_CACHE

The controllers restart after mirrored write-back cache is set, and you see %LFL and %EVL displays

NOTE: It may take up to five minutes after controller restart to check cache. The controllers reject this command until cache check is complete. If this command is rejected, do not restart the controllers. Wait a few minutes, then try again.

20. After the controllers restart, issue the following CLI command to confirm that mirrored write-back cache is enabled:

SHOW THIS_CONTROLLER

Notice that mirrored write-back cache is now set. You will see a display similar to that in "Example Display 18."

Example Display 18

```
Mirrored Cache:
            256 megabyte write cache, version 0012
            Cache is GOOD
            No unflushed data in cache
```

It is not necessary to repeat this step on controller B.

21. Set the fabric topology for each port on both controllers using the following CLI commands:

NOTE: You are prompted to restart the controllers after each command, but you do not need to restart the controllers until all topologies are set.

```
SET THIS_CONTROLLER PORT_1_TOPOLOGY=FABRIC
SET THIS_CONTROLLER PORT_2_TOPOLOGY=FABRIC
SET OTHER_CONTROLLER PORT_1_TOPOLOGY=FABRIC
SET OTHER_CONTROLLER PORT_2_TOPOLOGY=FABRIC
```

22. To ensure that fabric is up and running, issue the following CLI commands:

```
SHOW THIS_CONTROLLER
SHOW OTHER_CONTROLLER
```

You will see a display similar to that in "Example Display 19."

Example Display 19

```
Host PORT_1:
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
  . . . . . . . . . . . . . . . . . . PORT_1_TOPOLOGY  = FABRIC (up)
Address . . . . . . . . . . . . . . . . . . . . . . . . .  =nnnnnn
```

```
Host PORT_2:
   Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
   PORT_2_TOPOLOGY  = FABRIC (up)
   Address . . . . . . . . . . . . . . . . . . . . . . . . . . =nnnnnn
   NOREMOTE_COPY.
```

23. You are now ready to enable Data Replication Manager. Issue the following CLI command:

SET THIS_CONTROLLER REMOTE_COPY=*InitiatorControllerName*

Example: SET THIS_CONTROLLER REMOTE_COPY=BuildngA

NOTE: Be sure to specify a meaningful *InitiatorName*. Do not use "local" and "remote"; these are reserved keywords. The name can be up to eight characters and must be unique to all of your controllers. Follow the naming guidelines specified in the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide*.

After issuing this CLI command, you see a series of %LFL and %EVL displays; the controllers restart automatically.

24. Issue the following CLI command to verify that these settings are in place:

SHOW THIS_CONTROLLER

You will see a display similar to that in "Example Display 20."

Example Display 20

```
.
Host PORT_2:
            Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
            PORT_2_TOPOLOGY  = FABRIC (fabric up)
            REMOTE_COPY = BuildngA
```

Configure Storage at the Initiator Site

Devices and StorageSets

Before you can configure the storage for Data Replication Manager, you must add the disks, create the RAIDsets, and create units. Follow the instructions in the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide*, but note the restrictions listed at the beginning of this chapter.

NOTE: The initiator site must have exactly the same storageset and unit configuration as the target site.

Units

Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide* for detailed information.

NOTE: Use the following command to create storage units and to disable all access as the units are created. Disabling access prevents any TruCluster members from applying a Persistent Reserve to the new units.

ADD UNIT *UnitName StorageSetName* DISABLE_ACCESS_PATH=ALL

NOTE: Issue this command for all units.

1. **Verify that the access on each unit is set to none by issuing the following CLI command:**

   SHOW UNITS FULL

You will see a display similar to that in "Example Display 21."

Example Display 21

```
LUN                                      Uses             Used by
------------------------------------------------------------------------


D10  . . . . . . . . . . . . . . . . . . . DISK1000
   LUN ID:      nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn
       NOIDENTIFIER
       Switches:
          RUN                 NOWRITE_PROTECT READ_CACHE‘
          READAHEAD_CACHE     WRITEBACK_CACHE
          MAXIMUM_CACHED_TRANSFER_SIZE = 32
      Access:
           NONE
      State:
         ONLINE to this controller
         Not reserved
         NOPREFERRED_PATH
      Size: nnnnnnnn blocks
      Geometry (C/H/S): ( 7000 / 20 / 254 )
.
.
```

2. **Distribute the units by setting their preferred path. Use either of the following CLI commands:**

   SET *Unit* PREFERRED_PATH=THIS_CONTROLLER

   or

   SET *Unit* PREFERRED_PATH=OTHER_CONTROLLER

   Keep the busiest units on different host ports.

3. **Restart the controllers after configuring the units. Otherwise, the preferred path settings do not go into effect.**
4. **To ensure that your storage settings are in place, issue the following CLI command:**

   SHOW *Storagesets* FULL

Cabling the Initiator Site

This section shows you how to perform cabling at the initiator site.

Connect Fiber Optic Cables Between the Controllers and Switches

Figure 4–7 shows the switch port locations, which can assist you in connecting cables.

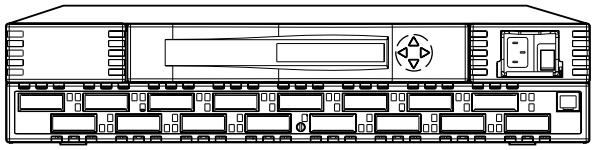

CXO7085A

Figure 4–7. Port locations

1. **Connect a multi-mode, 50-micron fiber optic cable from port 1 of the top controller to port 2 of the top Fibre Channel switch.**
2. **Connect a second multi-mode, 50-micron fiber optic cable from port 2 of the top controller to port 4 of the top Fibre Channel switch.**
3. **Connect a third multi-mode, 50-micron fiber optic cable from port 1 of the bottom controller to port 2 of the bottom Fibre Channel switch.**
4. **Connect a fourth multi-mode, 50-micron fiber optic cable from port 2 of the bottom controller to port 4 of the bottom Fibre Channel switch.**

NOTE: You should see an illuminated green LED on the switch as soon as the cable is inserted at both ends. This verifies that there is a good connection.

Figure 4–8 shows what your cabling should look like. The numbered callouts reflect the steps that you just completed.

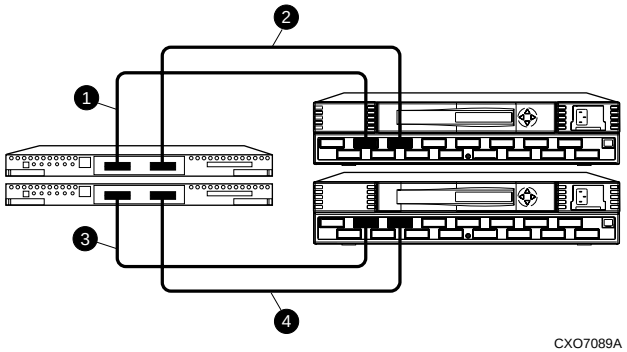

Figure 4–8. Cabling between the controllers and switches

Connect the Initiator Site to the External Fiber Link

Locate the connection points at the initiator site that link the initiator site to the target site. Look for either a fiber optic cable connector or a patch panel where you can insert a cable.

Long Wave GBICs

1. **Connect a single-mode, 9-micron fiber optic cable from port 6 of the top switch to one connection point.**
2. **Connect another single-mode, 9-micron fiber optic cable from port 6 of the bottom switch to the other connection point.**

Other Transport Modes

Data Replication Manager uses other long distance transport modes. For connection information, refer to our website at:

http://www.compaq.com/products/storageworks

The initiator site is now physically linked to the target site. See Figure 4–9 for an illustrated view of how this cabling should appear.

NOTE: You can make sure that switches and ports are connected as you have documented them by issuing the *nbrStateShow* command at the switch. Issue the *topologyShow* command at the switch to determine if you have more than one fiber optic cable between the switches on each site.

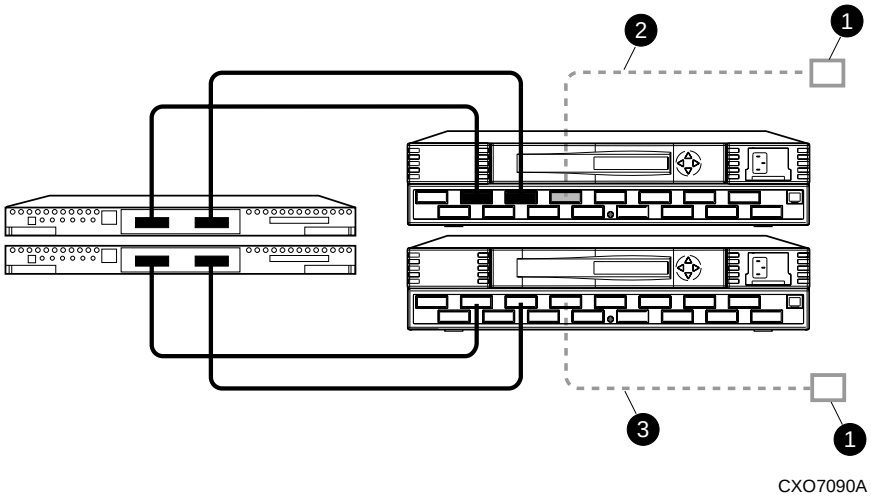

Figure 4–9. Cabling from the initiator site to the target site

Create Remote Copy Sets

Initiator Site Preparation

Before creating the remote copy set (RCS), create the connections between the initiator and target sites by issuing the following CLI command:

ADD REMOTE RCS199 D199 *TargetControllerName*\D199

Example: ADD REMOTE RCS199 D199 BuildngB\D199
Error: Initiator unit specified not found

NOTE: This command reports that it has failed, but it creates and names the connections appropriately.

Create Connections From the Target Site

1. Before creating the remote copy set, create the connections between the target and initiator sites by issuing the following CLI command:

   ADD REMOTE RCS199 D199 *InitiatorControllerName*\D199

   Example:ADD REMOTE RCS199 D199 BuildngA\D199

   NOTE: This command reports that it has failed, but it creates and names the connections appropriately.

2. Verify that the target has access to the initiator controller with this CLI command:

   SHOW CONNECTIONS

3. The target units must allow access to the controllers at the initiator site. Enable access with this CLI command:

   SET *UnitName* ENABLE_ACCESS_PATH= (*InitiatorControllerConnectionA, InitiatorControllerConnectionB, InitiatorControllerConnectionC, InitiatorControllerConnectionD)*

   Ex: SET*UnitName* ENABLE_ACCESS_PATH=(BuildngAA,BuildngAB,BuildngAC,BuildngAD)

   NOTE: Repeat this command for each *Unit Name*.

Create RCS from the Initiator Site

1. **Verify that the initiator has access to the target controller with this CLI command:**

   SHOW CONNECTIONS

2. **The initiator units must allow access to the controllers at the target site. Enable access with this CLI command:**

   SET *UnitName* ENABLE_ACCESS_PATH= (*TargetControllerConnectionA, TargetControllerConnectionB, TargetControllerConnectionC, TargetControllerConnectionD)*

   Example: SET *UnitName* ENABLE_ACCESS_PATH=(BuildngBA,BuildngBB,BuildngBC,BuildngBD)

   NOTE: Repeat this command for each *UnitName.*

3. **The following CLI command creates remote copy sets. When this command is entered, the controllers copy all data from the initiator unit to the target unit. This process is called *normalization*.**

   ADD REMOTE *RemoteCopySetName InitiatorUnitName TargetController Name\TargetUnitName*

   Example: ADD REMOTE RMTO D1 BuildngB\D1

   NOTE: It is not necessary to repeat this step at the target site.

You will see a %EVL display that includes your remote copy set information and a display similar to that in “Example Display 22.”

Example Display 22

```
%EVL--Initra > --13-JAN-1946 05:01:56 (time not set)-- Instance Code:
0E010064
  Template: 144.(90)
  Power On Time: 0. Years, 36. Days, 6. Hours, 45. Minutes, 22. Seconds
  Controller Model: HSG80
  Serial Number: ZG8nnnnnnn Hardware Version:  Enn(2B)
  Software Version: V85P
  Informational Report
  Target Controller Board Serial Number: "      ZG8nnnnnnn"
  Initiator WWLID: nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn
  Initiator Node Name: "BuildngA"
  Initiator Unit Number: n.(nnnnnnnn)
  Target WWLID: nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn-nnnn
  Target Node Name: "BuildngB"
  Target Unit Number: n.(nnnnnnnn)
  Number of Targets: n.(nnnnnnnn)
  Remote Copy Set Name: "RMT0”
  Instance Code: 0E010064
```

Set Failsafe at the Initiator Site (optional)

When failsafe is set, the remote copy set must contain one initiator member and one target member. If the remote copy set loses the target member while failsafe is set, no further I/O is allowed to the initiator member and an error is returned to the host. This is known as a **failsafe locked** condition and prevents the system from writing data that is not protected from a disaster by a copy at the target site.

If you choose to set failsafe, enter the following command:

SET *RemoteCopySetName* ERROR_MODE=FAILSAFE

Example: SET RMT0 ERROR_MODE=FAILSAFE

NOTE: When you set failsafe, all remote copy sets must be in a normal or normalizing state. If remote copy sets are copying when you set failsafe, your command is rejected until the remote copy sets return to normal mode.

To remove the failsafe lock from a remote copy set and resume normal operation, issue the following CLI command:

SET *RemoteCopySetName* ERROR_MODE=NORMAL

Example: SET RMT0 ERROR_MODE=NORMAL

You can also use this procedure for remote copy sets where a DT-safe condition is not required.

NOTE: If the error mode is set to normal and there is no target member, the remote copy set is no longer considered DT-safe.

Creating Log Units and Association Sets (optional)

In this example, hypothetical disks 50100 and 60100 are used as the mirrorset for the log disk. The log unit is D10. The association set name is AS_D1. The association set is using remote copy set name RC_D1.

Creating a Log Unit

1. Create a mirrorset for the log disk by issuing the following CLI command:

   ADD MIRRORSET *MirrorsetName DiskName*

   Example: ADD MIRR MIR_D1LOG DISK50100 DISK60100

   NOTE: To minimize the number of devices used for logging, you can create and use one-member mirrorsets. The logged data is not protected, because the data is written only to one disk. However, all of this data is also written to the initiator unit. In the case of a log disk failure, you would incur a full normalization, rather than a mini-merge, when access to the target is reestablished. The command to create a one-member mirrorset is the same as that above, except only one disk is listed. Example: ADD MIRR MIR_D1LOG DISK 50100.

2. Initialize the mirrorset with the following CLI command:

   INITIALIZE *ContainerName*

   Example: INITIALIZE MIR_D1LOG

3. Verify that you have created a mirrorset by issuing the following CLI command:

   SHOW MIRRORSET

   You will see a display similar to that in "Example Display 23."

Example Display 23

```
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------
MIR_D1LOG     mirrorset                     DISK50100
                                            DISK60100
```

4. Present the log unit to the controller with the following CLI command:

   ADD UNIT *UnitName ContainerName*

   Example: ADD UNIT D10 MIR_D1LOG

5. Verify that the controller recognizes the log unit by issuing the following CLI command:

   SHOW UNITS

You will see a display similar to that in "Example Display 24."

Example Display 24

```
LUN                    Uses                  Used by
- — — — — — — — — — — — — — — — — — — — — — — — — — — -
D10                    MIR_D1LOG
```

Creating Association Sets and Assigning a Log Unit

1. **Create an association set with the following CLI command:**

   ADD ASSOCIATIONS *AssociationSetName RemoteCopySetName*

   Example: ADD ASSOCIATIONS AS_D1 RC_D1

   NOTE: Additional members must be added to the association set by issuing the following CLI command:

   SET *AssociationSetName* ADD=*RemoteCopySetName*

2. **Disable node access to the log unit with the following CLI command:**

   SET *UnitNumber* DISABLE_ACCESS_PATH= ALL

   Example: SET D10 DISABLE_ACCESS_PATH= ALL

3. **Disable writeback cache with the following CLI command:**

   SET *UnitNumber* NOWRITEBACK_CACHE

   Example: SET D10 NOWRITEBACK_CACHE

4. **Check to see that you have disabled access and writeback cache with the following command:**

   SHOW D10

   You will see a display similar to that in "Example Display 25."

Example Display 25

```
LUN                                     Uses             Used by
------------------------------------------------------------------
D10                                       MIR_D1LOG
LUN ID:      6000-1FE1-0001-3B10-0009-9130-8044-0066
IDENTIFIER = 10
Switches:
RUN                    NOWRITE_PROTECT        READ_CACHE
READAHEAD_CACHE        NOWRITEBACK_CACHE
MAXIMUM_CACHED_TRANSFER_SIZE = 32
Access:
    None
State:
   ONLINE to this controller
   Not reserved
   PREFERRED_PATH = THIS_CONTROLLER
Size:             35556389 blocks
Geometry (C/H/S): ( 7000 / 20 / 254 )
```

5. **To assign the log unit to the association set, issue the following CLI command:**

   SET *AssociationSetName* LOG_UNIT = D10

   Example: SET AS_D1 LOG_UNIT = D10

6. **Check to see the switch status of the association set by issuing the following CLI command:**

   SHOW *AssociationSetName*

   Example: SHOW AS_D1

   You will see a display similar to that in "Example Display 26."

Example Display 26

```
Name          Association                   Uses             Used by
------------------------------------------------------------------------
AS_D1         association                   RC_D1
Switches:
  NOFAIL_ALL
  NOORDER_ALL
LOG_UNIT = D10 (No data logged)
```

Configure the Host at the Initiator Site

This section shows how to configure the host at the initiator site.

Install the Host Bus Adapters and Drivers

You should install two host bus adapters in each host system. Follow the procedures outlined. Refer to the *Compaq KGPSA-BC PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide* for installation information.

Install SWCC (optional)

Detailed information about SWCC can be found in the *Compaq StorageWorks Command Console User Guide*.

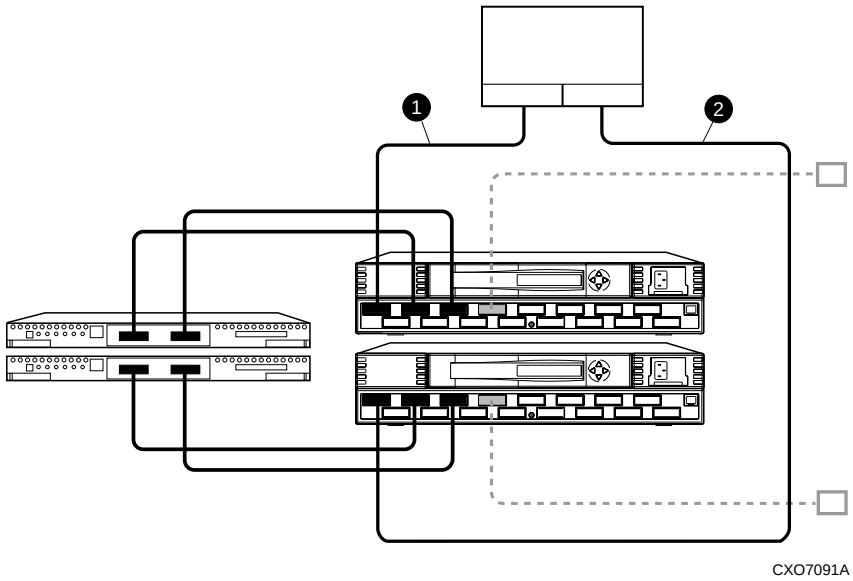

Figure 4–10. Cabling between the hosts and the switches

The cabling at each site is now complete. The initiator and target sites should be cabled according to the layout shown in Figure 4–11.

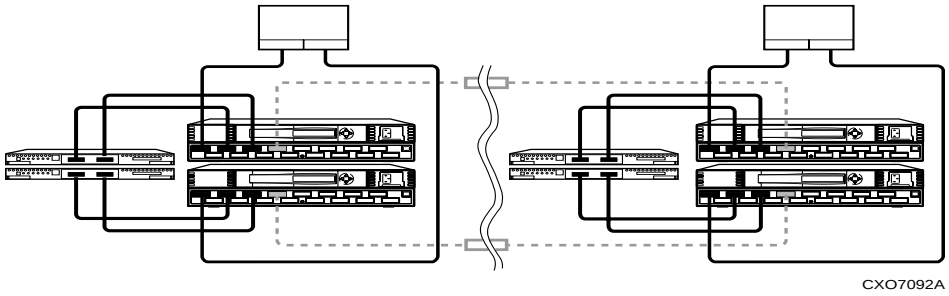

Figure 4–11. Data Replication Manager cabling at initiator and target sites

Rename the Host Connections at the Initiator Site

To identify which hosts you are working with, Compaq recommends that you rename the connection names that are reserved for the host names. You must change the !NEWCON prompt to a meaningful host name. Each host-based adapter appears as a connection. An individual host-based adapter can be identified by its World Wide Name, which you recorded in Chapter 3, "Getting Started," and which appears in the connection description.

1. Compaq suggests that you use the worksheet in Figure 4–12 when renaming your hosts. Fill in the fields accordingly to keep an accurate record of connections and host names.

!NEWCON*xx*	World Wide Name	Host Name	Path Number

Figure 4–12. Host renaming worksheet

2. **When you have completed the worksheet, rename the !NEWCON*xx* prompt using the following CLI commands:**

   RENAME !NEWCONxx *InitiatorHostConnectionNamex*

   RENAME !NEWCONxx *InitiatorHostConnectionNamey*

3. **When you have finished renaming your host connections, issue the following command to see your new settings:**

   SHOW CONNECTIONS

You will see a display similar to that in "Example Display 27."

Example Display 27

```
Connection Unit
Name      Operating system    . .Controller  Port   Address   Status
Offset

HOSTA1            TRU64_UNIX        THIS       1      210013   online 0
   HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

HOSTA2            TRU64_UNIX        OTHER      1      200113   online 0
   HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

BUILDNGBA PPRC_TARGET . . . . . .  THIS       2 online 0
   HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

BUILDNGBB PPRC_TARGET . . . . . .OTHER 2 online0
   HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

BUILDNGBCPPRC_INITIATOR . . . . . .  THIS 2 online 0
HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn
BUILDNGBDPPRC_INITIATOR . . . . . . OTHER 2 online 0
   HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn
```

At the target site, the initiator host appears as new connections, !NEWCONxx. Be sure to rename these as described above.

Enable Access to the Hosts at the Initiator Site

1. **The initiator units must have access to the hosts. Enable access with this command:**

   SET *UnitName* ENABLE_ACCESS_PATH=
   InitiatorHostConnectionNamex*,*InitiatorHostConnectionNamey

   Example: SET *UnitName* ENABLE_ACCESS_PATH=HostA1,HostA2

   NOTE: There should be two paths per host. You must repeat this sequence for each host. Reboot the host after you have enabled hosts to access the units.

2. **From a terminal window on the host, issue the following commands to recognize the new units and assign device special file numbers:**

   ```
   hwmgr-scan-comp-cat scsi_bus
   hwmgr-show scsi
   ```

Documenting Your Configuration

Keep a printed copy of your configuration for future reference. Update your records each time you modify the configuration. Follow the steps outlined below in the sections *Terminal Emulator Session* and *SHOW Commands* to obtain a status of the controllers, association sets, remote copy sets, units, and connections. After you have obtained this information for the initiator site, repeat the steps for the target.

Terminal Emulator Session

1. Use a laptop computer or another computer to connect a serial cable between the COM port on that machine and the corresponding serial port on the HSG80 controllers.
2. Start a terminal emulator session. On Windows NT-X86, use the HyperTerminal emulator. Use these settings: 9600 baud, 8 bits, No parity, 1 stop bit, XON/XOFF.
3. From the Transfer Menu, click *Capture Text*. The *Capture Text* dialog box is displayed.

   In the c:\field, type initiator.txt or target.txt.
4. Click Start.

SHOW Commands

1. To see the full information on this controller, issue the following CLI command:

   SHOW THIS_CONTROLLER FULL

   You will see a display similar to that in "Example Display **28**."

Example Display 28

```
Controller:
        HSG80 ZG91412410 Software S050P-0, Hardware  E05
        NODE_ID          = 5000-1FE1-0001-3AE0
        ALLOCATION_CLASS = 0
        SCSI_VERSION     = SCSI-3
        Configured for MULTIBUS_FAILOVER with ZG91416136
            In dual-redundant configuration
        Device Port SCSI address 6
        Time: NOT SET
        Command Console LUN is lun 0 (NOIDENTIFIER)
```

```
Host PORT_1:
        Reported PORT_ID = 5000-1FE1-0001-3AE1
        PORT_1_TOPOLOGY  = FABRIC (fabric up)
        Address          = 220113
Host PORT_2:
        Reported PORT_ID = 5000-1FE1-0001-3AE2
        PORT_2_TOPOLOGY  = FABRIC (fabric up)
        Address          = 220313
        REMOTE_COPY = BuildngA
Cache:
        256 megabyte write cache, version 0012
        Cache is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
Mirrored Cache:
        256 megabyte write cache, version 0012
        Cache is GOOD
        No unflushed data in cache
Battery:
        NOUPS
        FULLY CHARGED
        Expires:

Extended information:
        Terminal speed 9600 baud, eight bit, no parity, 1 stop bit
            Operation control: 00000000  Security state code: 75184
            Configuration backup disabled
```

2. To see the information for all association sets known to the controller pair, issue the following CLI command:

   SHOW ASSOCIATIONS FULL

   You will see a display similar to that in “Example Display **29**” for each association set.

Example Display 29

```
Name          Association                   Uses            Used by
-------------------------------------------------------------------------
AS1           association                   RC1
                                            RC2
                                            RC3
        Switches:
          NOFAIL_ALL
          NOORDER_ALL
          NOLOG_UNIT
```

3. To see information for all remote copy sets known to the controller pair, issue the following CLI command:

   SHOW REMOTE_COPY FULL

   You will see a display similar to that in "Example Display **30**" for each remote copy set.

Example Display 30

```
Name                                        Uses            Used by
-------------------------------------------------------------------------
RC1           remote copy                   D1              AS1
        Reported LUN ID: 6000-1FE1-0001-3AE0-0009-9141-6136-0038
        Switches:
          OPERATION_MODE  = SYNCHRONOUS
          ERROR_MODE      = NORMAL
          FAILOVER_MODE   = MANUAL
          OUTSTANDING_IOS = 60
        Initiator (BuildngA\D1) state:
                 ONLINE to this controller
       Target state:

                BuildngB\D1 is NORMAL
```

4. To see information for all units configured to the controller, issue the following CLI command:

   SHOW UNITS FULL

   You will see a display similar to that in "Example Display **31**" for each unit.

Example Display 31

```
D2                                          DISK10100 BuildngA\RC2
        LUN ID:      6000-1FE1-0001-3AE0-0009-9141-6136-0045
        NOIDENTIFIER
        Switches:
         RUN                 NOWRITE_PROTECT       READ_CACHE
          READAHEAD_CACHE        WRITEBACK_CACHE
          MAXIMUM_CACHED_TRANSFER_SIZE = 1
        Access:
BuildngBA,BuildngBB,BuildngBC,BuildngBD,HOSTA1,HOSTA2
        State:
          ONLINE to this controller
          Not reserved
          PREFERRED_PATH = OTHER_CONTROLLER
          Target NORMAL
        Size:             17769177 blocks

    Geometry (C/H/S): ( 5258 / 20 / 169 )
```

5. To see the connection name, operating system, controller, controller port, adapter ID address, online or offline status, and unit offset, issue the following CLI command:

   SHOW CONNECTIONS

   You will see a display similar to that in "Example Display **32**" for each connection.

Example Display 32

```
Connection                                                          Unit
   Name      Operating system    Controller  Port    Address    Status
Offset
!NEWCON28 TRU64_UNIX THIS       1      634000   OL this       0
          HOST_ID=1000-0000-C921-4B5B ADAPTER_ID=1000-0000-C921-4B5B
```

6. Save this file for future reference.
7. Repeat this procedure for the target.

Chapter 5

Managing Site Failover and Failback Procedures

This chapter describes how to manage Failover and Failback for your Data Replication Manager solution. This section contains the procedures to ensure that Failover and subsequent Failback function properly:

NOTE: All initiator site procedure text is shaded for ease of visibility and separation from target site procedures.

Power Up Data Replication Manager Systems

The procedures below outline how to power on and power off the storage subsystem after it has been configured.

CAUTION: Compaq recommends that you power up the controllers and switches at the target site before applying power to the initiator site. Powering up in the wrong sequence may cause incorrect configurations.

Power on the Data Replication Manager systems in the sequence shown in the following procedures.

Target Site Power Up Procedures

1. Ensure that all enclosures, switches, and cabinet power distribution units (PDUs) have their power switches in the OFF position.
2. Apply power to all PDUs.
3. Turn on the power switches for the cabinets from the target site.
4. Ensure that all controllers are on and functional.
5. Apply power to all Fibre Channel switches.

When completed, go to the Initiator Site Power Up Procedures.

Initiator Site Power Up Procedures

1. **Ensure that all enclosures, switches, and cabinet power distribution units (PDU) have their power switches in the OFF position.**
2. **Apply power to all PDUs.**
3. **Turn on the power switches for the cabinets from the initiator site.**
4. **Make sure that all controllers are on and functional.**
5. **Apply power to all Fibre Channel switches.**

Power Down Data Replication Manager Systems

Power down the Data Replication Manager systems in the sequence shown in the following procedures.

Initiator Site Power Down Procedures

1. **Issue the following CLI commands (in this order):**

   SHUTDOWN OTHER_CONTROLLER

   SHUTDOWN THIS_CONTROLLER
2. **Turn off the Fibre Channel switch.**
3. **Turn off the power to the enclosures.**
4. **Turn off the PDUs.**

When completed, go to the Target Site Power Down Procedures.

Target Site Power Down Procedures

1. Issue the following CLI commands (in this order):

   SHUTDOWN OTHER_CONTROLLER

   SHUTDOWN THIS_CONTROLLER
2. Turn off the Fibre Channel switch.
3. Turn off the power to the enclosures.
4. Turn off the PDUs.

Site Failover Basic Description

If the initiator site is no longer available, or if there is anticipated downtime that will prevent operation at the initiator site, you must decide whether or not to perform a site failover to the target site. Performing a failover enables the target site to assume the role of the initiator and access (write/read) data until the problem is resolved and a failback can be issued. Transferring control of system operation to the target site ensures a minimal interruption in data access after a failure.

> **NOTE:** If you decide to perform a Failover operation, keep in mind that *all* components must be failed over. Therefore, if only one component fails, fixing that single component may be preferable to performing a complete failover. Also, it is important to verify that all components at the target site are operational before you begin the site failover.

Table 5–1 outlines example scenarios that may call for a failover and those that may not.

Table 5–1 Failover Scenarios

When to Failover	When Not to Failover (recommended)
■ Both controllers fail	■ Single failed switch
■ Extended power outage at the initiator site	■ Single fiber optic cable malfunctions
■ Both host adapters fail (non-clustered hosts)	■ Single controller fails
■ Both initiator switches fail	■ Single storageset fails
■ Disaster (flooding, fire, earthquake, terrorism, etc.) that disables access to the subsystems	■ Single disk in redundant storageset fails
■ Scheduled event that prevents computing from the initiator site for an extended period	■ Target not in normal state
■ All hosts fail	

> **NOTE:** If one host in a multi-host environment fails, you must decide whether or not a failover is the best course of action.

When you decide that a site failover is necessary, identify which scenario best describes your situation: planned or unplanned failover.

The planned failover procedure should be used when failover is a scheduled event. Otherwise, Compaq suggests that you use an unplanned failover procedure.

CAUTION: Be sure to follow the steps outlined in the section *Planned Failover Procedures* accurately and completely, or you may incur data loss and extended downtime.

Failback Procedure Choices

During Failover, the remote copy sets at the target site are in a "copy ready" state, waiting for the initiator site to become available. When a new initiator site has been established or the original one has been restored, site operation can resume after a failback procedure has been performed. This involves synchronizing data on both the initiator and target subsystems so that operation can be returned to the initiator with minimal downtime.

IMPORTANT: Verify that all components at both sites are operational before performing a failback.

The failback sequence is a scheduled event. The HSG80 Array Controller requires that a viable dual-redundant subsystem be available before a failback can take place.

IMPORTANT: Failback to a single controller configuration is not supported.

The following table specifies which failback procedure to use in different circumstances:

State of the Initiator Controller Pair	Failover Procedure Used	Failback Procedure to Follow
Initiator site intact	Planned	Simple
Initiator site intact	Unplanned	Full
Initiator site not intact	Unplanned	New Hardware

Data Replication Manager Configuration Basics

The disaster-tolerant (DT) configuration that supports Data Replication Manager involves two HSG80 Array Controller subsystems—one at an initiator site and one at a target site.

IMPORTANT: Because of the complexity of the configuration process, it is a good idea to have all Data Replication Manager documentation available at both sites, to eliminate confusion and to minimize the risk of error. Please follow the steps precisely in the order provided in this documentation.

Figure 5–1 shows a basic DRM configuration; it is referenced throughout this chapter.

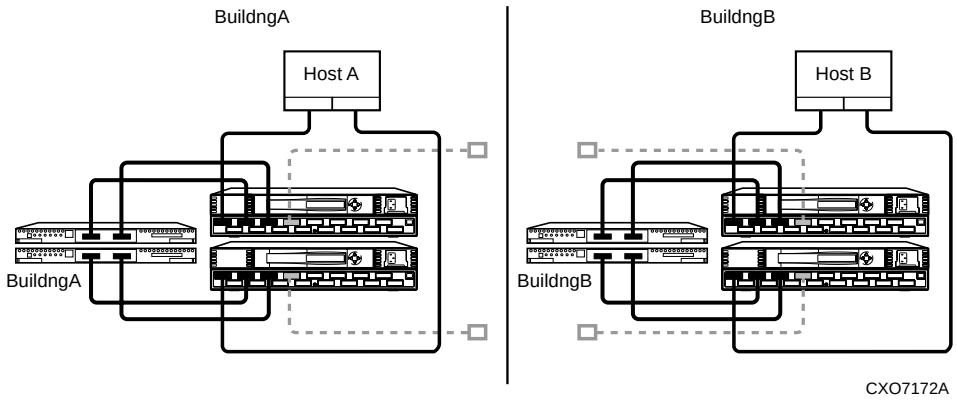

Figure 5–1. Data Replication Manager basic configuration

The example failover and failback procedures found in this chapter use fictional “Building A” as the initiator site and “Building B” as the target site. The examples failover from Building A to Building B and then failback from Building B back to Building A.

NOTE: This document consistently refers to Building A as the initiator site and Building B as the target site. This does not change even after failover has occurred to Building B (and before failback has occurred to Building A). While in failover mode, the controllers in Building B are acting as the *initiator* for all remote copy sets and are referred to as the *target* in this document.

Once the failback procedure is completed, the controllers in Building A resume their role as the initiator for remote copy sets.

Planning Considerations

The following constraints need to be considered in the initial planning of Data Replication manager (DRM) failover/failback procedures.

1. If you lose intersite connections, and both the initiator and target configurations are functional, the system administrator must determine which site to use. An intersite connection could include hardware or fiber-related equipment either at the initiator or target locations.
2. If you lose all access to the target controllers for any reason, immediately remove the remote copy set targets if either of the following two conditions applies:
   - ❏ None of the remote copy sets is running with write history logging.
   - ❏ You are running with write history logging but there is a possibility the log disk may overflow.
3. Use the following CLI command to remove remote copy set targets:

   SET *RemoteCopySetName* REMOVE=*TargetRemoteCopyName*\DiskName
4. If one of the initiator controllers fails, you can lose access to initiator units under the following conditions:
   - ❏ Access to target units in a remote copy set with no log disk assigned is lost for any reason (such as loss of both intersite links or loss of both target controllers),

     and
   - ❏ The target units are not removed from all of the remote copy sets.

   Loss of access occurs because a unit does not failover between controllers in a pair (such as from a failed top controller to a functional bottom controller) if that unit is the initiator of a remote copy set that has target units assigned, and those targets are not accessible.

   If this occurs, you are not able to access or alter the unit or its remote copy set in any way from the remaining controller. Access cannot be reestablished until the failed controller is either repaired or replaced. Furthermore, you do not see any apparent error message indicating that you no longer have access to the unit. To verify that units are inoperative you must check the status of all units by issuing the following command:

   SHOW UNITS FULL

5. An inoperative unit indicates the following state as part of its status display:

   ```
   State:
                   Unknown - Pending Remote Copy Set Validation
   ```

   This applies whether you are operating at the initiator site during normal operations or at the target site after a failover.

6. To clear this condition you must repair or replace the failed controllers, then:
   - ❑ Fix the extended link condition *or*
   - ❑ Remove the remote copy sets

7. After these conditions are met, you must restart both controllers to clear the faulted state.
8. By not removing the remote copy sets when both extended site connections are lost, you will be prohibited from moving LUNS from one HSG80 controller to the other HSG80 controller at the operating system level.

Persistent Reserve

In a TruCluster environment, cluster members apply a persistent reserve (PR) to any unit to which they have access. This prevents other devices, such as initiator or target controllers, from accessing these units. As a result, this persistent reserve must be removed at both the initiator and target sites as part of the failback procedure.

To view the persistent reserve, issue a SHOW UNIT command from the controller that the unit is online to. The following display shows how the persistent reserve is presented:

```
State:
                ONLINE To This (or Other) Controller
                Persistent reserved
```

NOTE: The persistent reserve is not displayed if you are connected to the other controller.

To prevent new units from being reserved by cluster members on the Fibre Channel fabric, use the following command format to create units:

ADD UNIT *UnitNumber StorageSetName* DISABLE=ALL

Because a cluster member reapplies the persistent reserve when it sees a unit, all cluster members must be denied access to the unit using the selective presentation feature of the HSG80. Once this is done, you must provide a standalone system access to the unit. Run the SCSI Command Utility (SCU) on the standalone system to remove the persistent reserve.

The commands to remove the persistent reserve follow.

1. From the controller CLI console issue the following commands:

   SET *UnitNumber* DISABLE=*ClusterHostConnectionNames*

   SET *UnitNumber* ENABLE=*StandaloneHostConnectionName*

2. From the standalone host, open a terminal window and issue the following commands:

   scu -f *device name* (with path to raw device) pres register key 0 skey 0x10101

   scu -f *device name* (with path) pres clear key 0x10101

An example of the host commands for a device named dsk5 with an active C partition follows:

```
scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101

scu -f /dev/rdisk/dsk5c pres clear key 0x10101
```

3. From the controller CLI console issue the following command:

   SET *UnitNumber* DISABLE=*StandaloneHostConnectionName*

Planned Failover Procedures

The Planned Failover Procedures outlined in the following sections must be used in conjunction with the Simple Failback Procedure. The planned failover consists of the following three procedures:

- Initiator Site Preparation Procedure
- Target Site Failover Procedure
- Target Host Setup Procedure

Initiator Site Preparation Procedure

1. **Before performing the failover procedure, locate your record of SHOW command output that details the current initiator configuration. (The procedure for obtaining a record of your initiator configuration is detailed in Chapter 4, "Creating a Data Replication Manager Solution.") Verify that your target controller configuration is the same as your initiator controller configuration.**
2. **If your remote copy sets are set for asynchronous operation mode, switch to synchronous mode using the following CLI command:**

   SET *RemoteCopySetName* OPERATION_MODE=SYNCHRONOUS

   Repeat this step for all remote copy sets.
3. **Turn off logging for the association sets (if enabled) and delete association sets with the following CLI command:**

   SET *AssociationSetName* NOLOG_UNIT

   DELETE *AssociationSetName*

   Repeat this step for all association sets.
4. **Unmount all file systems that are on remote copy set initiators.**

5. **Disable host access to the units by using the following CLI command:**

   SET*UnitName* **DISABLE=(***InitiatorHostConnectionNamex,InitiatorHostConnectionNamey)*

   NOTE: Do not disable access to the target connection.

   Repeat this step for all units.

6. **If in a TruCluster environment, remove the persistent reserve from all units. Provide the standalone host access to the unit using the following CLI command:**

   SET*UnitNumber* **ENABLE =** ***StandaloneHostConnection***

   At the host, run the following commands to recognize the unit and remove the reserve:

   hwmgr - scan comp - cat scsi_bus

   hwmgr - show scsi

   Look at the output to find the disk number and issue the following commands:

   scu -f /dev/rdisk/*dsk#* pres register key 0 skey 0x10101

   scu -f /dev/rdisk/*dsk#* pres clear key 0x10101

7. **Each unit that is used by a remote copy set should have four connections enabled to *TargetRemoteCopyNameA, TargetRemoteCopyNameB, TargetRemoteCopyNameC,* and *TargetRemoteCopyNameD*. To see the connections, issue the following CLI command:**

   SHOW UNITS FULL

   If access to the units is not currently enabled, issue the following command for each unit, to enable access.

   SET *UnitName* ENABLE=*TargetRemoteCopyNameA,TargetRemoteCopyNameB, TargetRemoteCopyNameC,TargetRemoteCopyNameD*

8. **Set maximum cached transfer size to 1 with the following CLI command:**

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = 1

   Repeat this step for all units.

9. **Shut down the initiator HSG80 controllers (in this order) with the following CLI commands:**

   SHUTDOWN OTHER_CONTROLLER

   SHUTDOWN THIS_CONTROLLER

10. **After the preceding command has completed and both controllers have shut down successfully, power off the controller cabinet. If the initiator site is powered down for a long time, you may need to disable cache batteries.**
11. **Continue the failover process at the target site with the "Target Site Failover Procedure."**

Target Site Failover Procedure

1. At the target site, the units must be preferred to one controller or the other. Use the following CLI command:

   SET *UnitName* PREFERRED_PATH = THIS_CONTROLLER *or* OTHER_CONTROLLER

   Repeat this step for each remote copy set unit.

2. The SITE_FAILOVER command allows you to move the initiator role to the target. This allows the target to set up write history logging for fast-failback when the connection to the initiator site is restored. To do this, type the command:

   SITE_FAILOVER *InitiatorRemoteCopyName\RemoteCopySetName*

   You will see a %EVL message on your terminal.

   Repeat this step for each remote copy set.

3. If you made the decision to **remove** remote copy set targets, continue with this step. If you made the decision to **NOT remove** remote copy set targets, go directly to Step 4.

   NOTE: See the "Planning Considerations" section of this chapter for information regarding removing remote copy set targets.

   To remove the targets, use the following CLI command:

   SET *RemoteCopySetName* REMOVE=*InitiatorRemoteCopyName\UnitNumber*

   Example: SET rcs1 REMOVE=buildngA\d1

   Repeat this step (Step 3) for all remote copy sets.

   NOTE: The ***InitiatorRemoteCopyName*** is the remote copy name of the *original* initiator.

   Go to Step 5 when completed with this step.

4. Create association sets and set up write history logging to duplicate those that are on the initiator.

   Repeat this step for each association set.

NOTE: Refer to Chapter 4, "Creating a Data Replication Manager Solution," for information on how to create association sets and configure association sets for write history logging.

5. Continue the failover procedure at the target site with the "Target Host Setup Procedure."

Target Host Setup Procedure

1. You can enhance host I/O performance by resetting the maximum cached transfer size to the value used on the initiator. Use this command:

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = *InitiatorValue*

   NOTE: The default maximum cached transfer size is 32.

   Repeat this step for each unit.

2. Give the target site hosts access to the units that are used by remote copy sets in the storage subsystems with this command:

   SET *UnitName* ENABLE=(*TargetHostConnectionNamex,TargetHostConnectionNamey*)

   If you do not recall the target host connection name, use the SHOW CONNECTION command.

   Repeat this step for each unit.

3. To verify that all of these steps have been completed successfully, issue this CLI command:

   SHOW REMOTE_COPY_SETS FULL

   The output shows you the status of remote copy sets.

   NOTE: Be sure that the units you see (listed under *Initiator State*) are at the target site.

4. To verify that the target hosts can connect to the LUNs, use this command:

   SHOW UNITS FULL

   In the Access field of the display, all units that are used by remote copy sets should show that the target hosts connections are enabled. This should also show the initiator controller connections.

5. Allow hosts to recognize new units:

   - If the target site hosts are shut down, boot them at this time. Booting the hosts enables Tru64 UNIX to recognize the drives.

- ❏ If the target site hosts are not shut down, use the following commands to recognize the drives:

  ```
  hwmgr - scan comp - cat scsi_bus
  hwmgr - scan scsi
  ```

6. Build AdvFS domains and mount filesets:

   To build the AdvFS database for the domain and fileset structure, run the *advscan* command. This rebuilds the /etc/fdmns directory to include the domains created at the initiator site.

   Domains are named based on the volumes that they contain. For instance, if a domain includes dsk5c and dsk6g (partition c of dsk5 and partition g of dsk6) the domain is named domain_dsk5c_dsk6g. If you wish to rename these domains, please refer to your AdvFS documentation for instructions.

   Before running the following command, determine the *dsk* numbers and active partitions for all of your remote copy sets. From a terminal window type:

   ```
   /sbin/advfs/advscan -r /dev/rdisk/dsk#1 /dev/rdisk/dsk#2
   ```

   NOTE: Include all active partitions for all units.

   The domains are listed in the /etc/fdmns directory. You can determine which filesets are active by running the following command for all domains:

   ```
   showfsets domain_name
   ```

   Once you have identified all filesets and have created corresponding directories, you can mount them.

   The following example shows how to create a directory for, and then mount, a fileset named fileset5 on domain_dsk5c.

   NOTE: Any legal directory would work as well.

   ```
   mkdir -p /usr/fileset5
   mount -t advfs domain_dsk5c#fileset5 /usr/fileset5
   ```

This completes the target host setup procedure. The following section describes the Simple Failback Procedure from a planned failover.

Simple Failback Procedure

The Simple Failback Procedure is used in conjunction with the Planned Failover Procedure. Before performing the failback procedure, locate your record of SHOW command output that details the initiator configuration. Verify that your target controller configuration is the same as your initiator controller configuration. Compare the status of the controllers, association sets, remote copy sets, units, and connections. Appendix A contains the full procedure.

The simple failback consists of the following three procedures:

- Initiator Site Failback Preparation
- Target Site Failback
- Initiator Site Cleanup

Target Site Unit Persistent Reserve Removal

Before powering up the controllers at the initiator site you must remove any persistent reserves from the units at the target site (if you are not in a cluster environment this does not apply and you can skip to the Initiator Site Failback Preparation Procedure). This allows the initiator controllers to see the target site units, which is necessary for the initiator units to come up in a state that allows the failback procedure to proceed.

Because a cluster member reapplies the persistent reserve when it sees a unit, all cluster members must be denied access to the unit using the selective presentation feature of the HSG80. Once this is done, you must provide a standalone system access to the unit. Run the SCSI Command Utility (SCU) on the standalone system to remove the persistent reserve. The commands to remove the persistent reserve follow.

Unmount file systems that are on remote copy set units.

From the controller CLI console, issue the following commands:

SET *UnitNumber* DISABLE=*ClusterHostConnectionNames*

SET *UnitNumber* ENABLE=*StandaloneHostConnectionName*

From the standalone host, open a terminal window and issue the following commands:

scu -f *device name* (with path to raw device) pres register key 0 skey 0x10101

scu -f *device name* (with path) pres clear key 0x10101

An example of the host commands for a device named dsk5 with an active C partition follows:

```
scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101
scu -f /dev/rdisk/dsk5c pres clear key 0x10101
```

From the controller CLI console, issue the following command:

SET *UnitNumber* DISABLE=*StandaloneHostConnectionName*

Initiator Site Failback Preparation Procedure

1. **The controllers should have been powered down since the failover procedure. If this is not the case and it is possible that writes have occurred to the units since the failover procedure was executed, use the "Full Failback Procedure" on page 5–24 of this chapter.**

   Warning: Any data that had been written to the initiator unit is destroyed during the copy back of data.

2. **If you made the decision to remove the remote copy set targets, use the "Full Failback Procedure" on page 5-16 of this chapter.**
3. **Power up the controller cabinets. Once the connection between the initiator and the target has been re-established, the remote copy sets begin to merge.**
4. **Continue the simple failback process at the target site with the "Target Site Simple Failback Procedure."**

Target Site Simple Failback Procedure

1. Now that you have powered up the controller cabinets, the remote copy sets should be in the process of merging. Ensure the targets have not been dropped by checking the status of the merge periodically with the following CLI command:

   SHOW REMOTE_COPY_SETS FULL

   If targets have been dropped, use the following command to re-add them. This results in a full normalization:

   SET REMOTECOPYSETNAME **ADD**=INITIATORCONTROLLERNAME\UNITNAME

   When the remote copy sets have completed merging, the Target field of the display is NORMAL.

IMPORTANT: You must wait for merge on all remote copy sets to complete before you can proceed. However, you can continue performing I/O to the units at the target site during normalization.

2. When all remote copy sets are done normalizing and you decide to move operations back to the initiator site, shut down the target site hosts.
3. Unmount file systems that are on remote copy set units.
4. Disable host access to the target units for all remote copy sets by using the following CLI command:

   SET *UnitName* DISABLE=(*TargetHostConnectionNamex,TargetHostConnectionNamey*)

5. If you are not in a clustered environment, skip to the next step. Now that you have removed cluster member access to the units at the target site you can remove the persistent reserve from these units. This allows the initiator units to use them as targets. The commands to remove the persistent reserve follow.

   From the controller CLI console, issue the following command:

   SET *UnitNumber* ENABLE=*StandaloneHostConnectionName*

   From the standalone host, open a terminal window and issue the following commands:

   scu -f *device name* (with path to raw device) pres register key 0 skey 0x10101

   scu -f *device name (*with path) pres clear key 0x10101

   An example of the host commands for a device named dsk5 with an active C partition follows:

   scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101

   scu -f /dev/rdisk/dsk5c pres clear key 0x10101

   From the controller CLI console, issue the following command:

   SET *UnitNumber* DISABLE=*StandaloneHostConnectionName*

6. You may now boot hosts for non-remote copy set units.
7. Turn off write history logging if enabled:

   SET *AssociationSetName* NOLOG_UNIT

   Repeat this procedure for each association set.

8. Delete the association set by using the following CLI command:

   DELETE *AssociationSetName*

   Repeat this procedure for each association set.

9. Move control of the remote copy sets to the original initiator using the following CLI command:

   SET *RemoteCopySetName* INITIATOR=*InitiatorRemoteCopyName\UnitName*

   NOTE: If after issuing this command for one of the Remote Copy Sets, you get the error message: `Error: Rem Cp Set specified is currently in a transient state,` wait a few seconds and try again. The command eventually succeeds.

   NOTE: Repeat this step for all remote copy sets.

10. If maximum cached transfer size was changed for the target units as part of the failover procedure, set it back to 1 with the following CLI command:

    SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE=1

    Repeat this procedure for each unit.

11. Continue with the simple failback procedure at the initiator site with "Initiator Site Cleanup Procedure."

Initiator Site Cleanup Procedure

1. **You can enhance host I/O performance by resetting the maximum cached transfer size to the original value used on the initiator. Use this command:**

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = *InitiatorValue*

2. **Create the association sets and then add the log units.**

   NOTE: Refer to Chapter 4, "Creating a Data Replication Manager Solution," for information on how to create association sets and configure association sets for write history logging.

3. **Enable access to the initiator site host by using the following CLI command:**

   SET *UnitName* ENABLE=(*InitiatorHostConnectionNamex,InitiatorHostConnectionNamey*)

4. **Optional: Set failsafe by using the following CLI command:**

   SET *RemoteCopySetName* ERROR MODE=FAILSAFE

NOTE: Failsafe cannot be set if the remote copy set is within an association set that is to be used for write history logging.

5. **Enable write history logging, if appropriate.**

   NOTE: Refer to the Chapter 4, "Creating a Data Replication Manager Solution," for information on how to configure association sets for write history logging.

6. **If you changed an asynchronous remote copy set to synchronous during failover, change back to asynchronous mode by issuing the following CLI command:**

   SET *RemoteCopySetName* OPERATION_MODE=ASYNCHRONOUS

   Repeat this step for all applicable remote copy sets.

7. **If you have shut down the host, boot the host at this time. Booting the host enables the host to recognize the drives. If you did not shut down the host, from a terminal window on the host, run the following command to recognize the new units and assign device special file numbers:**

   # hwmgr - scan comp - cat scsi_bus

   # hwmgr - show scsi

8. Build AdvFS domains and mount filesets:

To build the AdvFS database for the domain and fileset structure, run the *advscan* command. This rebuilds the /etc/fdmns directory to include the domains created at the initiator site.

Domains are named based on the volumes in the domain. For instance, if a domain includes dsk5c and dsk6g (partition c of dsk5 and partition g of dsk6) the domain is named domain_dsk5c_dsk6g. If you wish to rename these domains, please refer to your AdvFS documentation for instructions.

Before running the following command, determine the *dsk* numbers and active partitions for all of your remote copy sets. From a terminal window type:

/sbin/advfs/advscan -r /dev/rdisk/dsk#1 /dev/rdisk/dsk#2 **(include all active partitions for all units)**

The domains are listed in the /etc/fdmns directory. You can determine which filesets are active by running the following command for all domains:

showfsets *domain_name*

Once you have identified all of the filesets and created corresponding directories, you can mount them.

The following example shows how to create a directory for and mount a fileset named fileset5 on domain_dsk5c. Any legal directory would work as well.

```
mkdir -p /usr/fileset5

mount -t advfs domain_dsk5c#fileset5 /usr/fileset5
```

This completes the Initiator Site Cleanup Procedure.

Unplanned Failover

Use the Target Site Failover Procedure outlined in this section in conjunction with the Full Failback or New Hardware Failback procedures whenever a situation occurs at the initiator site to bring it down (unable to perform its functions as an initiator).

Target Site Failover Procedures

IMPORTANT: Since the initiator may be running and perhaps write history logging, care must be taken to ensure that the connection between the sites be severed and not be restored until directed to do so in the proper failback procedure.

1. Ensure that the connection between sites is not restored by typing the following CLI commands:

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = OFFLINE

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = OFFLINE

2. At the target site, prefer the units to one controller or the other using the following CLI command:

   SET *UnitName* PREFERRED_PATH = THIS_CONTROLLER *or* OTHER_CONTROLLER

   Repeat this step for all units.

3. Use the following CLI command to failover each remote copy set (maximum of twelve per subsystem):

   SITE_FAILOVER *InitiatorRemoteCopyName*\RemoteCopySetName

   You will see a %EVL message on your terminal.

   Repeat this step for all remote copy sets.

4. To remove the targets use the following CLI command:

   SET *RemoteCopySetName* REMOVE=*InitiatorRemoteCopyName*\UnitNumber

   Example: SET rcs1 REMOVE=buildngA\d1

   Repeat this step for all remote copy sets.

5. (Optional) Re-create association sets and log disks that are on the initiator.

   NOTE: Refer to Chapter 4, "Creating a Data Replication Manager Solution," for information on how to create association sets and log disks.

   Repeat this step for all association sets.

6. Give the target site hosts access to the units in its storage subsystems with this command:

   SET *UnitName* ENABLE = *TargetHostConnectionNamex,TargetHostConnectionNamey*

   If you do not recall the target host name, use the SHOW CONNECTION command.

   Repeat this step for all units.

7. To verify that all of the steps have been completed successfully, issue this CLI command:

   SHOW REMOTE_COPY FULL

8. To verify that the target host can connect to the LUNs, use this command:

   SHOW UNITS FULL

   In the Access field of the display, all units should show that the target hosts are enabled. You should also see the connections to the initiator controller.

9. You can enhance host I/O performance by resetting the maximum cached transfer size to the original value used on the initiator. Use this command:

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = *InitiatorValue*

10. If you have shut down the host, boot the host at this time. Booting the host enables the host to recognize the drives. If you did not shut down the host, from a terminal window on the host run the following command to recognize the new units and assign device special file numbers:

    # hwmgr - scan comp - cat scsi_bus

    # hwmgr - show scsi

11. Build AdvFS domains and mount filesets:

    To build the AdvFS database for the domain and fileset structure, run the *advscan* command. This rebuilds the /etc/fdmns directory to include the domains created at the initiator site.

    Domains are named based on the volumes in the domain. For instance, if a domain includes dsk5c and dsk6g (partition c of dsk5 and partition g of dsk6) the domain is named domain_dsk5c_dsk6g. If you wish to rename these domains, please refer to your AdvFS documentation for instructions.

    Before running the following command, determine the *dsk* numbers and active partitions for all of your remote copy sets. From a terminal window type:

    /sbin/advfs/advscan -r /dev/rdisk/dsk#1 /dev/rdisk/dsk#2

    NOTE: Include all active partitions for all units.

The domains are listed in the /etc/fdmns directory. You can determine which filesets are active by running the following command for all domains:

showfsets *domain_name*

Once you have identified all of the filesets and created corresponding directories, you can mount them.

The following example shows how to create a directory for and mount a fileset named fileset5 on domain_dsk5c. Any legal directory would work as well.

```
mkdir -p /usr/fileset5

mount -t advfs domain_dsk5c#fileset5 /usr/fileset5
```

12. (Optional) If, while performing failover, you decided to create a new unit protected by a new remote copy set, use the following CLI command:

    ADD REMOTE_COPY_SETS *RemoteCopySetName\UnitName*

This completes the failover procedure. When the problem that disabled the initiator site is remedied, refer to “Failback Procedure Choices” on page 5–5 for proper failback procedure.

Full Failback Procedure

After powering up initiator controllers and before continuing the Full Failback Procedure, verify that your initiator controller configuration is the same as your target controller configuration.

Compare the status of the controllers, association sets, remote copy sets, units, and connections. A full procedure is detailed in Appendix A. Make sure any status change is reflected on the target. A status comparison is accomplished by bringing up a terminal emulator session and entering a SHOW THIS command.

Initiator Site Preparation Procedure

1. **Shut down any initiator hosts that are still up and running.**
2. **If the subsystems are powered off, power on at this time.**
3. **If any of the remote copy set units are connected to a cluster and the unit is online, check those units for a persistent reserve. If a persistent reserve is present, remove the reserve with the following procedure:**

   From the controller CLI console, issue the following commands:

   SET *UnitNumber* DISABLE=*ClusterHostConnectionNames*

   SET *UnitNumber* ENABLE=*StandaloneHostConnectionName*

   From the standalone host, open a terminal window and issue the following commands:

   scu -f *device name* (with path to raw device) pres register key 0 skey 0x10101

   scu -f *device name* (with path) pres clear key 0x10101

   An example of the host commands for a device named dsk5 with an active C partition follows:

   scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101

   scu -f /dev/rdisk/dsk5c pres clear key 0x10101

If the unit's status is "unknown pending remote copy set verification," you will not be able to remove the persistent reserve while the unit is part of a remote copy set. To remove the persistent reserve on cluster-connected units, delete the remote copy set using the following command:

DELETE *RCSName*

Now use the following procedure to remove the persistent reserve.

NOTE: From this point, these units are treated as parts of new remote copy sets created at the target site.

From the controller CLI console, issue the following commands:

SET *UnitNumber* DISABLE=*ClusterHostConnectionNames*

SET *UnitNumber* ENABLE=*StandaloneHostConnectionName*

From the standalone host, open a terminal window and issue the following commands:

scu -f *device name* (with path to raw device) pres register key 0 skey 0x10101

scu -f *device name* (with path) pres clear key 0x10101

An example of the host commands for a device named dsk5 with an active C partition follows:

scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101

scu -f /dev/rdisk/dsk5c pres clear key 0x10101

4. **Disable access to all units by issuing the following CLI command:**

   SET *UnitName* DISABLE=ALL

   Repeat this step for all units.

5. **Check all units to make sure there is no lost data. If there is lost data, clear it with the following CLI command:**

   CLEAR_ERRORS *UnitName* LOST_DATA

   NOTE: Use the SHOW UNITS FULL command to check for lost data.

   Repeat this step for all units.

6. **Both controllers on the initiator site must be restarted (even if you just powered on). Do this with the following CLI commands:**

   RESTART OTHER_CONTROLLER

   RESTART THIS_CONTROLLER

   NOTE: Wait 5 minutes for controller memory diagnostics to complete before proceeding.

7. **Remove initiator site host access by using the following CLI command:**

   SET *UnitName* DISABLE = *InitiatorHostConnectionNamex,InitiatorHostConnectionNamey*

8. **For each remote copy set, use the following CLI command to set the error mode to normal:**

   SET *RemoteCopySetName* ERROR_MODE = NORMAL

9. **Use the following CLI commands to turn off logging for all of the association sets and to remove association sets:**

   SET *AssociationSetName* NOLOG_UNIT

   DELETE *AssociationSetName*

   NOTE: Repeat this step for all association sets.

10. **Delete all remote copy sets using the following CLI command:**

    DELETE *RemoteCopySetName*

11. **Set up *new units* for any additional remote copy sets that were added at the target site while failed over, by using the following CLI command:**

    ADD UNIT *UnitName ContainerName* DISABLE=ALL

12. **At the initiator site, prefer the units to one controller or the other using the following CLI command:**

    SET *UnitName* PREFERRED_PATH = THIS_CONTROLLER *or* OTHER_CONTROLLER

13. **Set maximum cached transfer size back to 1 with the following CLI command:**

    SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = 1

14. **Disable Port 1 connections to fabric by using the following CLI command:**

    SET THIS (and OTHER) PORT_1_TOPOLOGY=OFFLINE

15. **Continue with the full failback procedures at the target site with "Target Site Preparation Procedure."**

Target Site Preparation Procedure

This section describes the preparation of the target site and the creation of connections from the initiator site to the target.

1. Disable initiator controller access to all remote copy set units by issuing the following command:

   SET *UnitName* DISABLE=(*InitiatorRemoteCopyNameA,InitiatorRemoteCopyNameB, InitiatorRemoteCopyNameC,InitiatorRemoteCopyNameD*)

   Repeat this step for each unit.

2. Verify that you have disabled access with the following CLI command:

   SHOW *UnitName* FULL

3. Delete the connections to the original controllers at the initiator site using the following CLI command:

   DELETE *InitiatorRemoteCopyNameA*

   DELETE *InitiatorRemoteCopyNameB*

   DELETE *InitiatorRemoteCopyNameC*

   DELETE *InitiatorRemoteCopyNameD*

   The only access to the target units will now be from the hosts.

4. To restore the connections to the initiator site, type the following CLI commands:

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = FABRIC

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = FABRIC

5. Issue the following CLI command:

   ADD REMOTE_COPY_SETS RCS199 D199 InitiatorRemoteCopyName\D199

   NOTE: This command reports as failed, but it creates and names the connections appropriately.

6. Continue with the full failback procedure at the initiator site with “Initiator Site Connections Procedure.”

Initiator Site Connections Procedure

This section describes the creation of initiator site connections to the target.

1. **Set target access to all remote copy units by issuing the following CLI command:**

   SET *UnitName* ENABLE = (*TargetRemoteCopyNameA*,*TargetRemoteCopyName*B, *TargetRemoteCopyNameC*,*TargetRemoteCopyName*D)

   Repeat this procedure for all units.

2. **Continue with the full failback procedure at the target site with "Target Site Copy Data procedure."**

Target Site Copy Data Procedure

The section describes the copying of the data from the target site to the initiator.

1. Set initiator access to all remote copy units.

   SET *UnitName* ENABLE = (*InitiatorRemoteCopyNameA*,*InitiatorRemoteCopyNameB*, *InitiatorRemoteCopyNameC*,*InitiatorRemoteCopyNameD*)

   Repeat this step for all units.

2. Add back the initiator unit to the remote copy sets with the following CLI command:

   SET *RemoteCopySetName* ADD = *InitiatorRemoteCopyName**UnitName*

   Repeat this step for all remote copy sets.

IMPORTANT: You must wait for normalization on all remote copy sets to complete before you can proceed.

---

3. Enter the following command to see the percentage of completion.

   SHOW REMOTE_COPY_SETS FULL

   When the units are all normalized, the Target field of the display is NORMAL.

4. Stop I/O from the target hosts to the remote copy set units.
5. Unmount all units that are part of a remote copy set.
6. In a TruCluster environment, remove persistent reserves from all units using the following procedure:

   Removing persistent reserves allows the initiator units to use them as targets. If you are not in a clustered environment, you can skip to the next step. The commands to remove the persistent reserve follow.

   From the controller CLI console, issue the following commands:

   SET *UnitNumber* DISABLE=ClusterHostConnectionNames

   SET *UnitNumber* ENABLE=StandaloneHostConnectionName

   From the standalone host, open a terminal window and issue the following commands:

   scu -f device name *(with path to raw device)* pres register key 0 skey 0x10101

   scu -f device name *(with path)* pres clear key 0x10101

   An example of the host commands for a device named dsk5 with an active C partition follows:

   scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101

   scu -f /dev/rdisk/dsk5c pres clear key 0x10101

   From the controller CLI console, issue the following command:

   SET *UnitNumber* DISABLE=*StandaloneHostConnectionName*

7. Disable host access to the target units by using the following CLI command:

   SET*UnitName* DISABLE=(*TargetHostConnectionNamex,TargetHostConnectionNamey)*

   Repeat this step for all units.

8. Delete all association sets and log units by using the following CLI commands:

   SET *AssociationSetName* NOLOG_UNIT

   DELETE *AssociationSetName*

   Repeat this procedure for all association sets.

9. Shut down the target HSG80 controllers (in this order) with the following CLI commands:

   SHUTDOWN OTHER_CONTROLLER

   SHUTDOWN THIS_CONTROLLER

10. Continue with the full failback procedure at the initiator site with "Initiator Site Return Control Procedure."

Initiator Site Return Control Procedure

This section describes how to return Data Replication Manager control to the initiator site.

1. **Disconnect controller access by using the following CLI commands:**

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = OFFLINE

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = OFFLINE

2. **Issue the site_failover command to properly set up the initiator site:**

   SITE_FAILOVER *TargetRemoteCopyName*\RemoteCopySetName

   Observe a %EVL message on your terminal.

   Repeat this step for each remote copy set.

3. **Continue with the full failback procedure at the target site with the *Target Site Restore Procedure*.**

Target Site Restore Procedure

1. Both controllers on the target site must be restarted after the site failover has taken place. Press the *Reset* button or turn on the power.
2. Delete all remote copy sets using the following CLI command:

   DELETE *RemoteCopySetName*

3. Set the maximum cached transfer size, if it was changed for all the remote copy units, with the following CLI command:

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = 1

4. Continue with the full failback procedure at the initiator site with "Initiator Site Restoration of Target Connections."

Initiator Site Restoration of Target Connections

This section describes the how to restore all target connections from the initiator site.

1. **To restore the connections to the target site, type the following CLI commands:**

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = FABRIC

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = FABRIC

2. **Re-enable logging or failsafe, if desired. To set failsafe, type the following CLI command:**

   SET *RemoteCopySetName* ERRORMODE=FAILSAFE

   NOTE: Failsafe cannot be set if the remote copy set is within an association set that is to be used for write history logging.

3. **Create the association set and then add the log unit.**

   NOTE: Refer to Chapter 4, "Creating a Data Replication Manager Solution," for information on how to create association sets and configure association sets for write history logging.

4. **Enable host access by issuing the following CLI commands:**

   SET THIS_CONTROLLER PORT_1_TOPOLOGY = FABRIC

   SET OTHER_CONTROLLER PORT_1_TOPOLOGY = FABRIC

5. **Enable host access to the units by using the following CLI command:**

   SET*UnitName* ENABLE=(*InitiatorHostConnectionNamex,InitiatorHostConnectionNamey)*

6. **To verify that all of these steps have been completed successfully, issue this CLI command:**

   SHOW REMOTE_COPY FULL

   The output shows you a list of remote copy sets. Be sure the *Initiator State* points to the initiator and the *Target State* points to the target.

7. **Set maximum cached transfer size to the original value using the following CLI command:**

   SET *UnitName* MAXIMUM_CACHE_TRANSFER_SIZE = *initiator value*

8. To verify that the initiator host can connect to the LUNs, use this command:

```
SHOW UNITS FULL
```

In the Access field of the display, all units should show that the initiator hosts are enabled.

This completes the full Failback Procedure.

New Hardware Failback Procedure

Use the New Hardware Failback Procedure when the initiator site is not intact and you are working with all new hardware that is not configured.

Initiator Site Preparation Procedure

1. **Shut down any initiator hosts that are still up and running.**
2. **Manually reconfigure the controllers, but do not re-create the original remote copy sets. This procedure includes the following steps:**

   NOTE: Steps b, e, and f cause the controller pair to restart.

   a. **Set node ID and checksum (this information can be found on the original initiator BA370 cabinet). See Chapter 4, "Creating a Data Replication Manager Solution," for information on the node ID and worldwide name.**

   b. **Enter the following command:**

      SET MULTIBUS_FAILOVER COPY = THIS

   c. **Set the controller to SCSI-3 using the following CLI command:**

      SET THIS_CONTROLLER SCSI_VERSION = SCSI-3

   NOTE: Do not restart the controller.

   d. **Designate a controller prompt name using the following CLI commands:**

      SET THIS_CONTROLLER PROMPT= ***"InitiatorControllerNameTop > "***

      SET OTHER_CONTROLLER PROMPT="*InitiatorControllerNameBottom > "*

   e. **Set mirrored cache using the following CLI command:**

      SET THIS_CONTROLLER MIRRORED_CACHE

   f. **SET THIS REMOTE_COPY =** ***InitiatorRemoteCopyName***

g. **Run the Configuration utility to assign a disk name to physical disks, using the following CLI command:**

   RUN CONFIGURATION

h. **Create and initialize storage sets and units. The units that are to be part of remote copy sets must be identical to the corresponding units at the target site.**

3. **Disable access to all units by issuing the following CLI command:**

   SET *UnitName* DISABLE=ALL

   Repeat this step for all units.

4. **Set up *new units* for any additional remote copy set that were added at the target site while failed over, by using the following CLI command:**

   ADD UNIT *UnitName ContainerName* DISABLE=ALL

5. **At the initiator site, prefer the units to one controller or the other using the following CLI command:**

   SET *UnitName* PREFERRED_PATH = THIS_CONTROLLER *or* OTHER_CONTROLLER

6. **Set maximum cached transfer size back to 1 with the following CLI command:**

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = 1

7. **Disable Port 1 and enable Port 2 connections to fabric by using the following CLI commands:**

   SET THIS (and OTHER) PORT_1_TOPOLOGY=OFFLINE

   SET THIS (and OTHER) PORT_2_TOPOLOGY=FABRIC

8. **Compare the status of the controllers, association sets, remote copy sets, units, and connections. Detailed description of the procedure is shown in Appendix A. Make sure any status change is reflected on the target. To make a status comparison, bring up a terminal emulator session and enter a SHOW THIS command.**

9. **Continue with the new hardware failback procedures at the target site with "Target Site Preparation Procedure."**

Target Site Preparation Procedure

This section describes how to prepare a target site and how to create connections from the initiator site to the target.

1. Remove the targets from the remote copy sets with the following CLI command:

   SET *RemoteCopySetName* REMOVE = *InitiatorName*\UnitName

2. Disable initiator controller access to all remote copy set units by issuing the following command:

   SET *UnitName* DISABLE=(*InitiatorRemoteCopyNameA,InitiatorRemoteCopyNameB, InitiatorRemoteCopyNameC,InitiatorRemoteCopyNameD*)

   Repeat this step for each unit.

3. Verify that you have disabled access with the following CLI command:

   SHOW *UnitName* FULL

4. Delete the connections to the original controllers at the initiator site using the following CLI command:

   DELETE *InitiatorRemoteCopyNameA*

   DELETE *InitiatorRemoteCopyNameB*

   DELETE *InitiatorRemoteCopyNameC*

   DELETE *InitiatorRemoteCopyNameD*

   The only access to the target units will now be from the hosts.

5. To restore the connections to the initiator site, type the following CLI commands:

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = FABRIC

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = FABRIC

6. Issue the following CLI command:

   ADD REMOTE_COPY_SETS RCS199 D199 InitiatorRemoteCopyName\D199

   NOTE: This command reports as failed, but it creates and names the connections appropriately.

7. Continue with the new hardware failback procedure at the initiator site with “Initiator Site Connection Connections Procedure.”

Initiator Site Connections Procedure

This section describes the how to create initiator site connections to the target.

1. **Issue the following CLI command:**

   ADD REMOTE_COPY_SETS RCS199 D199 *TargetRemoteCopyName*\D199

   NOTE: This command reports as failed, but it creates and names the connections appropriately.

2. **Set target access to all remote copy units by issuing the following CLI command:**

   SET *UnitName* ENABLE = (*TargetRemoteCopyNameA,TargetRemoteCopyName*B, *TargetRemoteCopyNameC,TargetRemoteCopyName*D)

   Repeat this procedure for all units.

3. **Continue with the new hardware failback procedure at the target site with "Target Site Copy Data procedure."**

Target Site Copy Data Procedure

The section describes how to copy data from the target site to the initiator.

1. Set initiator access to all remote copy units with the CLI command:

   SET *UnitName* ENABLE = (*InitiatorRemoteCopyNameA,InitiatorRemoteCopyNameB, InitiatorRemoteCopyNameC,InitiatorRemoteCopyNameD*)

   Repeat this step for all units.

2. Add back the initiator unit to the remote copy sets with the following CLI command:

   SET *RemoteCopySetName* ADD = *InitiatorRemoteCopyName\UnitName*

   Repeat this step for all remote copy sets.

IMPORTANT: You must wait for normalization on all remote copy sets to complete before you can proceed.

3. Enter the following command to see the percentage of completion.

   SHOW REMOTE_COPY_SETS FULL

   When the units are all normalized, the Target field of the display is NORMAL.

4. Stop I/O from the target hosts to the remote copy set units.
5. Unmount all units that are part of a remote copy set.
6. In a TruCluster environment, remove persistent reserves from all units using the following procedure:

   Removing persistent reserves allows the initiator units to use them as targets. If you are not in a clustered environment, you can skip to the next step. The commands to remove the persistent reserve follow.

   From the controller CLI console:

   SET *UnitNumber* DISABLE=ClusterHostConnectionNames

   SET *UnitNumber* ENABLE=StandaloneHostConnectionName

   From the standalone host, open a terminal window and issue the following commands:

   scu -f device name *(with path to raw device)* pres register key 0 skey 0x10101

   scu -f device name *(with path)* pres clear key 0x10101

   An example of the host commands for a device named dsk5 with an active C partition follows:

   scu -f /dev/rdisk/dsk5c pres register key 0 skey 0x10101

   scu -f /dev/rdisk/dsk5c pres clear key 0x10101

   From the controller CLI console:

   SET *UnitNumber* DISABLE=*StandaloneHostConnectionName*

7. Disable host access to the target units by using the following CLI command:

   SET*UnitName* DISABLE=(*TargetHostConnectionNamex,TargetHostConnectionNamey)*

   Repeat this step for all units.

8. Delete all association sets and log units by using the following CLI commands:

   SET *AssociationSetName* NOLOG_UNIT

   DELETE *AssociationSetName*

   Repeat this procedure for all association sets.

9. Shut down the target HSG80 controllers (in this order) with the following CLI commands:

   SHUTDOWN OTHER_CONTROLLER

   SHUTDOWN THIS_CONTROLLER

10. Continue with the new hardware failback procedure at the initiator site with "Initiator Site Return Control Procedure."

Initiator Site Return Control Procedure

This section describes how to return Data Replication Manager control to the initiator site.

1. **Disconnect controller access by using the following CLI command:**

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = OFFLINE

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = OFFLINE

2. **Issue the site_failover command to properly set up the initiator site:**

   SITE_FAILOVER *TargetRemoteCopyName*\RemoteCopySetName

   You will see a %EVL message on your terminal.

   Repeat this step for each remote copy set.

3. **Continue with the new hardware failback procedure at the target site with the *Target Site Restore Procedure*.**

Target Site Restore Procedure

1. Both controllers on the target site must be restarted after the site failover has taken place. Press the *Reset* button or turn on the power.
2. Delete all remote copy sets using the following CLI command:

   DELETE *RemoteCopySetName*

3. Set the maximum cached transfer size, if it was changed for all the remote copy units, with the following CLI command:

   SET *UnitName* MAXIMUM_CACHED_TRANSFER_SIZE = 1

4. Continue with the new hardware failback procedure at the initiator site with the "Initiator Site Restoration of Target Connections" section.

Initiator Site Restoration of Target Connections

This section describes how to restore all target connections from the initiator site.

1. **To restore the connections to the target site, type the following CLI commands:**

   SET THIS_CONTROLLER PORT_2_TOPOLOGY = FABRIC

   SET OTHER_CONTROLLER PORT_2_TOPOLOGY = FABRIC

2. **Re-enable logging or failsafe, if desired. To set failsafe, type the following CLI command:**

   SET *RemoteCopySetName* ERRORMODE=FAILSAFE

   NOTE: Failsafe cannot be set if the remote copy set is within an association set that is to be used for write history logging.

3. **Create the association set and then add the log unit.**

   NOTE: Refer to Chapter 4, "Creating a Data Replication Manager Solution," for information on how to create association sets and configure association sets for write history logging.

4. **Enable host access by issuing the following CLI commands:**

   SET THIS_CONTROLLER PORT_1_TOPOLOGY = FABRIC

   SET OTHER_CONTROLLER PORT_1_TOPOLOGY = FABRIC

5. **If the initiator hosts were shut down, reboot them at this time. Rename all !NEWCONxx connections to their previous names.**

   RENAME !NEWCONxx *InitiatorHostConnectionName*

   Example: RENAME !NEWCONxx hostA1

6. **Change the operating system to Tru64 UNIX for each new connection with the following CLI command:**

   SET !NEWCONXX OPERATING_SYSTEM = TRU64_UNIX

7. **Enable host access to the units by using the following CLI command:**

   SET*UnitName* ENABLE=(*InitiatorHostConnectionNamex,InitiatorHostConnectionNamey)*

8. **To verify that all of these steps have been completed successfully, issue this CLI command:**

   SHOW REMOTE_COPY FULL

   The output shows you a list of remote copy sets. Be sure the *Initiator State* points to the initiator and the *Target State* points to the target.

9. **Set maximum cached transfer size to the original value using the following CLI command:**

   SET *UnitName* MAXIMUM_CACHE_TRANSFER_SIZE = *initiator value*

10. **To verify that the initiator host can connect to the LUNs, use this command:**

    SHOW UNITS FULL

    In the Access field of the display, all units should show that the initiator hosts are enabled.

This completes the Full Failback Procedure.

Chapter 6

Troubleshooting

This chapter describes possible failure modes of a Data Replication Manager solution. Isolation of errors and detailed error analysis require a complete understanding of how a Data Replication Manager subsystem operates. While it is not possible to document every error and failure condition, key failures of the Data Replication Manager subsystem and its components are discussed.

Troubleshooting information on specific Data Replication Manager components can also be found in their respective user manuals.

This section contains the following topics:

HSG80 Array Controller Operating Characteristics

The HSG80 array controller has certain characteristics that may become evident when used in a Data Replication Manager solution. The following sections will help you understand these characteristics and educate you on how to respond to them.

Forced Errors Detected During Copy

A forced error is a data bit indicating that a corresponding logical data block contains unrecoverable data. If a read request from the initiator to the target encounters a forced error during a full copy, then the data in that block is copied to the target and marked with a forced error. These forced errors are then reported to the host and reappear each time the block is read. The file containing the forced error should be restored from a known good backup.

Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide* for complete details on how to recover from a forced error situation.

Read Errors Detected During Full Copy

During normal operation, when a read error is detected an unrecoverable error is reported to the host, the offending block is re-vectored, and the new block is marked with a forced error. During a full copy, however, the handling is slightly different because the block that is unrecoverable may not be within normal file system space. In that case, the controller terminates the copy and reports the event.

Unrecoverable read errors on the source member terminate the copy and send a fault management report to the host. Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Configuration Guide* and *Compaq StorageWorks HSG80 Array Controller ACS V8.5 CLI Reference Guide* for more information on how to interpret these logs.

Dual Redundancy During Failback

The failback sequence is a scheduled event based upon the configuration at the failback site. The HSG80 array controller requires that a viable dual-redundant subsystem be available before a failback can take place. Failback to a single-controller configuration is not supported.

Failsafe Lock Management

If failsafe mode is set for a remote copy set, that set can become failsafe locked if a unit fails or the target becomes inaccessible.

If a unit fails, then the target is removed from the remote copy set. Once the unit failure has been eliminated, the target can be re-added to the remote copy set that initiates a full copy.

If a dual-link failure occurs, the remote copy set is placed in a failsafe locked condition. The target remains a member of the remote copy set but is marked invalid. Once the link has been restored to the target, a full copy is initiated. Once completed, the failsafe locked condition is cleared.

If the initiator unit fails, the remote copy set goes into failsafe locked condition.

Link Failure Management

When an initiator controller detects that the link to its target controller is unavailable, the initiating controller restarts. This causes all remote copy sets on the initiating controller to failover to its dual-redundant partner controller. The restart of the initiator controller is an intended action and is not indicative of a defective controller.

Remote Copy Set Member Failures

While most remote copy set members are based on protected storage, in the unlikely event of a remote copy set member failure, the following operating characteristics should be understood:

- If a remote copy set target member fails, a write issued to that remote copy set causes a write failure at the target. The target member is removed, and the remote copy set is put in failsafe lock condition. If you wish to continue operation at the initiator site, be sure to change the remote copy set error mode to normal before proceeding.

- If a remote copy set member at the initiator fails, the unit becomes unavailable to the host. The target member of the remote copy set is not read/write accessible through the initiator controller. Recovery from this condition requires a failover to the target site.

Remote Copy Set Worldwide LUN ID

Remote copy sets are assigned a unique worldwide LUN ID (WWLID) that represents their specific LUN. The controller identifies a remote copy set by its WWLID and presents it to the target when a failover is executed for that unit. If the remote copy set is failed over to a target site, its WWLID is transferred with that unit, even though it may not be consistent with the controller's worldwide ID or the IDs of the other units presented on the new controller. The remote copy set does not assume a new WWLID, regardless of those that appear at the target site.

Write History Logging

Once write history logging commences to a log unit, care must be taken when choosing to disable logging. Issuing the SET *AssociationSetName* NOLOG_UNIT command may incur a full copy operation on the remote copy set. For example, this may happen if the controller is logging updates for a remote copy set because the links to the target are down. If the log unit is disabled during this time, the controller cannot use the write history log to update the target when the links are restored, as some operations were not written to the log. Therefore, a full copy is initiated. Also, the log disk is no longer known to the controller.

Component Failures

The service and maintenance of a Data Replication Manager solution is based on failure of subsystem components. When a component fails, you must determine the cause of the failure, the most appropriate workaround to eliminate down time, and the best course of action to resolve the problem.

Failure Notification

It is important to understand the operation of the DT subsystem and the individual component error logging methods that are used to analyze failures on a DT subsystem. Each component within the DT subsystem provides error and failure information specific to the function being performed. The array controllers maintain and log specific information relevant to the operation and to the devices connected to both the host ports and device ports of the controllers. Events, errors, and failures related to a DT subsystem are provided to the host. Information is available from the HSG80 controller via the serial maintenance port.

With Data Replication Manager, fault management events that occur on the target controllers are "passed through" and reported on the initiator controllers. The initiator then reports these events to the host via Template 90 (Data Replication Manager Services Event Sense Data Response). Refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide* for more information.

HSG80 Array Controller Failure

The HSG80 array controller provides event and error reporting via the controller's serial maintenance port. To help capture random disk errors associated with the controllers, a terminal can be connected to this maintenance port. For a complete explanation and examples of these codes, refer to the *Compaq StorageWorks HSG80 Array Controller ACS V8.5 Maintenance and Service Guide*.

SWCC Failure

SWCC notifies the user of any component loss in the system via an active SWCC Client Graphical User Interface (GUI). This GUI window on the command console monitor is a graphical representation of the controllers and their physical and logical storage elements. SWCC periodically queries the controllers for status. Clients connected to the GUI *.ini* file are notified at the GUI screen of any changes in status. The user can manipulate controllers and storage through the GUI and can intervene in the DT process when there is a problem.

Refer to the *Compaq StorageWorks Command Console User Guide* and the on-line user help for more information.

Failure of One Member in a Dual Redundant Controller Pair

In a dual redundant setup, each of the controller pairs can lose a single member to failure. When this happens, a normal controller failover occurs automatically, and the preferred devices are automatically moved to the remaining controller. A decrease in I/O speed may occur. The faulty controller must be replaced using conventional controller troubleshooting techniques.

NOTE: It is not possible to set up a DT configuration unless both controllers are operational.

Failure of Both Fiber Optic Cables or Switch

If you are operating in failsafe mode and both links between the initiator and target sites are lost, the remote copy set is put into failsafe locked mode. If you are operating in normal mode, I/O continues through the initiator host, and the target is still removed.

If you lose the fiber optic cable connection of a switch at either site, see Table 6–1 for information on how to resolve the problem.

Failure Modes of a DT System In Normal Operation

Table 6–1 details the failure modes of a DT system operating in normal mode. While this table concentrates on the major failure possibilities, keep in mind that there are several other combinations that may occur. In most cases, when there is a loss of a major component, a failover is necessary to continue operation.

Table 6–1 Failure Modes of a DT System with Normal Operation

Initiator Host	Target Host	Initiator Switch A	Initiator Switch B	Target Switch A	Target Switch B	Initiator Controller A	Initiator Controller B	Target Controller A	Target Controller B	**Failure Mode** *Loss of:*	**Action**
X										Applications	Failover; Repair Host
	X									Remote host	Repair Host
X	X									Both sites	Failover not possible; Repair Hosts
		X								Data path	Repair Switch
			X							Data path	Repair Switch
				X						Data path	Repair Switch
					X					Data path	Repair Switch
		X	X							Data access	Failover; Repair Switches
				X	X					Remote copy set targets	Repair Switches; Target member must incur mini-merge or full copy
		X		X						Data path	Repair Switches
						X				Data path	Repair Controller
							X			Data path	Repair Controller
								X		Data path	Repair Controller
									X	Data path	Repair Controller
						X	X			Data access	Failover
								X	X	Remote Copy Set Targets	Repair controllers; Normalize remote copy sets
						X		X		Data path	Repair controllers

Failure at Target Site after Failover

After a failover has occurred, failures at the target site are detected in the same way as in a non-disaster tolerant state. Table 6–2 shows the possible failure modes at the target site, assuming that the initiator site is not available to failback to.

Table 6–2 Target Site DT Failure Modes After Failover

Target Host	Target Top Switch	Target Bottom Switch	Target Controller A	Target Controller B	Failure Mode *Loss Of*	Action
X					Remote site	Repair host
	X				Data path	Repair switch
		X			Data path	Repair switch
	X	X			Data access	Repair switches
			X		Data path	Repair controller
				X	Data path	Repair controller
			X	X	Data access	Replace controllers

Chapter 7

Configuration Variations

This chapter describes Data Replication Manager concepts and variations for alternative Data Replication Manager configurations. These descriptions include Cascaded Switches, Multiple Intersite Links, DataSafe Solutions, and Switch Zoning.

The topics contained within this chapter are:

Cascaded Switches

Using cascaded switches provides a DRM configuration variation that lets you:

- Increase the distance between sites (expand the fabric)
- Increase host or controller port connections

The term *cascaded switch* means that the output of switch A is connected to the input of switch B. Switch B may then be connected to another switch, host, or controller.

Hopping

The cascading of switches employs hopping. A *hop* is defined as one or more connections between two Fibre Channel switches. For example, two switches cascaded are equal to one hop. Server to Fibre Channel switch segments and storage to Fibre Channel switch segments are not counted as hops.

Table 7–1 lists rules that apply to hops in a Data Replication Manager environment.

Table 7–1 DRM Hop Rules

1.	Only three hops are allowed from the initiator controller to the target controller, as shown in Figure 7-1.
2.	Only three hops are allowed from the host to the initiator controller, as shown in Figure 7-2.
3.	Within a single fabric where switches are interconnected, each Fibre Channel switch must have a unique domain number (Domain_ID).
4.	The maximum distance allowed using short wavelength laser GBICs and 50-micron multi-mode fiber optic cable is 500 meters per cable segment.
5.	The maximum distance allowed using long wavelength laser GBICs and 9-micron single-mode fiber optic cable is 10 kilometers per intersite cable segment.
6.	Very Long Distance GBICs can extend intersite links up to 100 kilometers.
7.	Dense Wave-Length Division Multiplexing (DWDM) can extend up to 120 kilometers.
8.	Only one extended long wavelength intersite is allowed per fabric.
9.	Cascaded switches are not supported in ATM configurations.

Cascaded Switches Configurations

Figure 7–1 shows a Data Replication Manager configuration that increases the distance between sites by using cascaded switches and hopping to expand the distance limitations of the fabric.

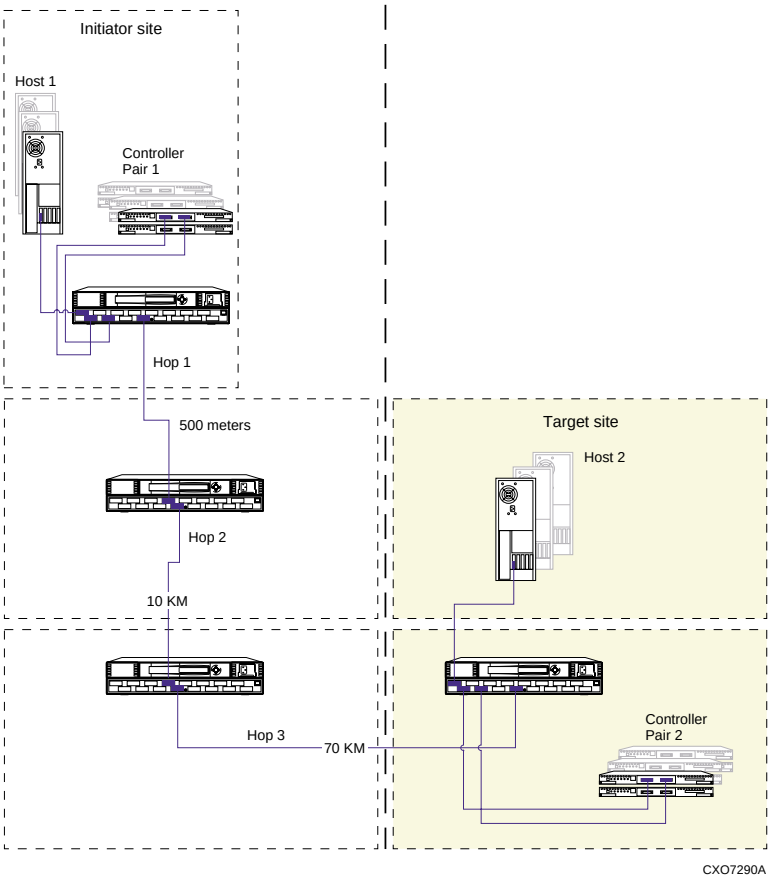

Figure 7–1. Cascaded switches in DRM environment, with 0 hops from host and 3 hops to controller

Figure 7–1 features switches cascaded over varying distances for different configurations. This example configuration shows four Fibre Channel cascaded switches, no hops from the host, and three hops to controller pair 2. Hop 1 spans the shortest distance (up to 500 meters), hop 2 spans up to 10 kilometers, and hop 3 spans the longest distance, at 70 kilometers.

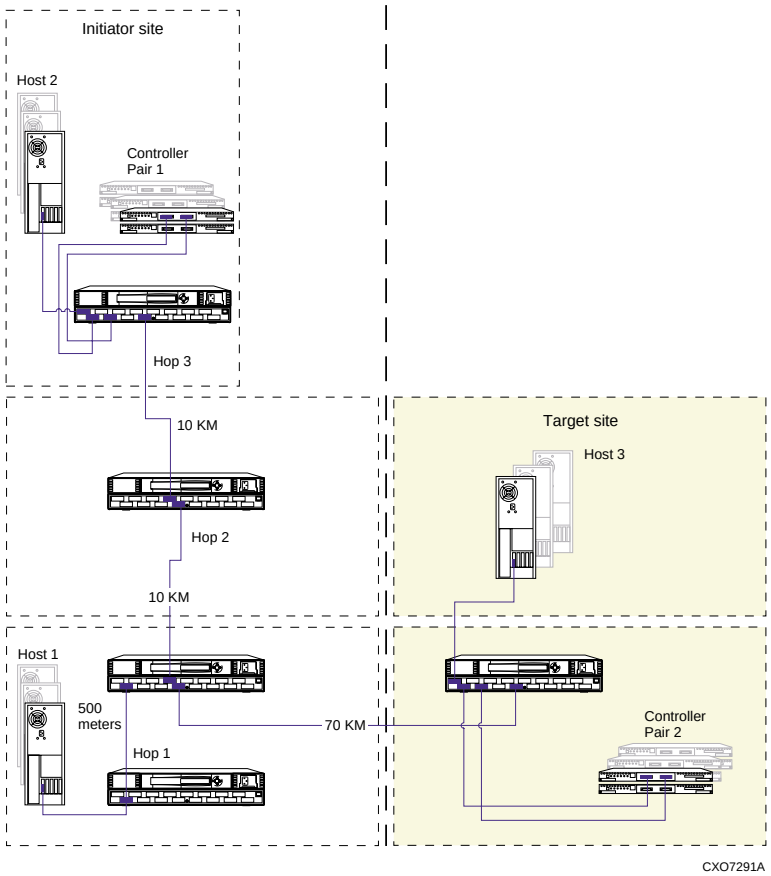

Figure 7–2. Cascaded switches in DRM environment, with 3 hops from host and 3 hops to controller

Figure 7–2 shows a Data Replication Manager configuration that increases the number of host to controller port connections using cascaded switches and hopping. The figure features switches cascaded from Host 1 to the initiator controller. This example configuration shows four Fibre Channel cascaded switches, three hops from Host 1, and three hops to controller pair 1.

Multiple Intersite Links

Multiple intersite links (ISLs) may be used to provide additional bandwidth between local and remote sites. Each ISL is a fiber link between two switches.

Table 7–2 lists restrictions that apply when using multiple intersite links in a Data Replication Manager environment.

Table 7–2 Multiple Intersite Link Restrictions

DRM supports a maximum of two intersite link connections.
To access more than one e-port when using multiple intersite links with the SAN Switch 8-EL, the Multiple E-port Connectivity software option is required.
For a Cascaded Switches configuration, the SAN switch 8-EL must be placed at the end of the cascade to provide access to the e-port, unless the Multiple E-port Connectivity software option is used.

Increasing bandwidth between switches is handled by adding additional connections between the switches, as shown in Figure 7–3. DRM supports a maximum of two connections. The switches are shown as physically connected, although the connections are transparent to the fabric. Functionally, the devices appear to be connected directly together.

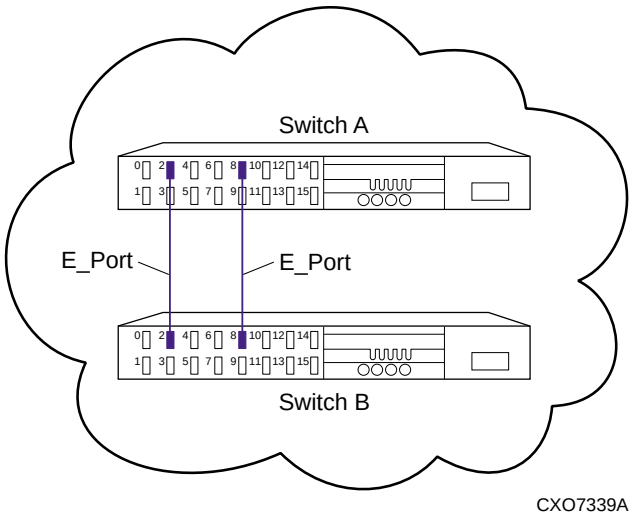

Figure 7–3. Multiple intersite links

DataSafe Solutions

The Compaq Departmental DataSafe solution (also known as “firewall”) is a pre-tested configuration variation that uses specific hardware, Data Replication Manager software, and installation practices to protect operations from hardware or software outages. The DataSafe solution provides a cost-effective configuration variation in the Data Replication Manager environment, that can be implemented at a single site.

DataSafe is based on a standard Data Replication Manager implementation that maintains two separate fabrics, fault tolerance, and no single point of failure. Bidirectional and stretch cluster configurations can be implemented. The DataSafe configuration varies in two important ways, since it:

- Uses fewer switches (2 instead of the standard 4)
- Eliminates intersite links for improved bandwidth performance

DataSafe uses 50-micron multi-mode cable and short wave GBICs, which limits to 500 meters the maximum distance between the server and switch, and between the switch and storage array.

DataSafe Configuration

Figure 7–4 shows the pre-tested configuration for the DataSafe solution.

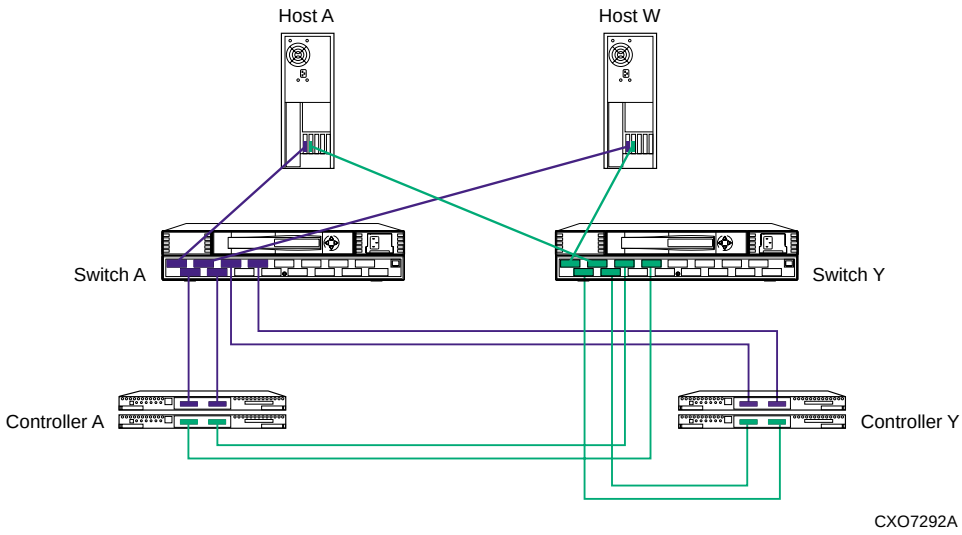

Figure 7–4. DataSafe (firewall) configuration in DRM environment

DataSafe Configuration Procedures

Carefully follow the procedures outlined in Chapter 4 of this guide to configure a Data Replication Manager solution, with these exceptions:

- The two Fibre Channel switches should be wired in a variation of the procedure described in the "*Connect Fiber Optic Cables Between the Controllers and Fiber Channel Switches"* section of Chapter 4, for both the target site and the initiator site. For the DataSafe solution, follow these instructions:

    1) Connect a multi-mode, 50-micron fiber optic cable from port 1 of the top controller of Controller Pair A to port 1 of the Fibre Channel switch A.
    2) Connect a second multi-mode, 50-micron fiber optic cable from port 2 of the top controller of Controller Pair A to port 3 of the Fibre Channel switch A.
    3) Connect a third multi-mode, 50-micron fiber optic cable from port 1 of the bottom controller of Controller Pair A to port 6 of the Fibre Channel switch Y.
    4) Connect a fourth multi-mode, 50-micron fiber optic cable from port 2 of the bottom controller of Controller Pair A to port 4 of the Fibre Channel switch Y.

5) Connect a multi-mode, 50-micron fiber optic cable from port 1 of the top controller of Controller Pair Y to port 4 of the Fibre Channel switch A.
6) Connect a second multi-mode, 50-micron fiber optic cable from port 2 of the top controller of Controller Pair Y to port 6 of the Fibre Channel switch A.
7) Connect a third multi-mode, 50-micron fiber optic cable from port 1 of the bottom controller of Controller Pair Y to port 3 of the Fibre Channel switch Y.
8) Connect a fourth multi-mode, 50-micron fiber optic cable from port 2 of bottom controller of Controller Pair Y to port 1 of the target Fibre Channel switch Y.

NOTE: You should see an illuminated green LED on the switch as soon as the cable is inserted at both ends. This verifies that there is a good connection.

- To connect fiber optic cables between the hosts and the switches, follow these instructions:

1) Connect a multi-mode, 50-micron fiber optic cable from port 0 of switch A to one adapter in the initiator host (Switch A, Host A).
2) Connect a multi-mode, 50-micron fiber optic cable from port 2 of switch Y to the other adapter in the initiator host (Switch Y, Host A).
3) Connect a multi-mode, 50-micron fiber optic cable from port 2 of switch A to one adapter in the target host (Switch A, Host W).
4) Connect a multi-mode, 50-micron fiber optic cable from port 0 of switch Y to the other adapter in the target host (Switch Y, Host W).

NOTE: You may choose any available port to connect your cables, but you must maintain that identical scheme at the initiator and target sites. For additional hosts, connect to the remaining switch ports in the manner described above.

5) Verify that the connection between the host and the switch has been made by entering this CLI command:

```
SHOW CONNECTIONS
```

NOTE: You can also verify that a connection has been made by looking for the illuminated green LED that flashes on the switch ports.

- Eliminate the entire procedure for connecting the External Fiber Links for both the target site and the initiator site. In each instance this procedure occurs in Chapter 4 between “Connect Fiber Optic Cables Between the Controllers and Switches” and “Create Controller Connections.” This procedure is not needed since the DataSafe solution does not accommodate long wave GBICs or other transport modes.

Switch Zoning

The Fibre Channel switch zoning feature provides a means to control Storage Area Network (SAN) access at the node port level. Zoning can be used to separate one physical fabric into many virtual fabrics consisting of selected server and storage ports. This capability allows you to:

- Set up barriers between different operating environments
- Deploy logical fabric subsets by creating defined user groups
- Create separate test and maintenance areas within the fabric
- Flexibly manage a SAN while meeting the different objectives of closed user groups

SAN Management

The intelligent infrastructure of the Fibre Channel fabric is created using one or more switches. This fabric is the backbone for deploying and managing IT resources as a network and is enhanced when zoning is used. Zoning provides management for the SAN, both automatically and transparently, while arranging fabric-connected devices into logical groups over the physical fabric configuration.

Although the fabric provides fast, reliable, seamless information access within the SAN, zoning increases SAN control by creating segmentation or zones within a fabric. Selected storage devices, servers, and/or workstations comprise these zones, and zoning enforces limiting access to information only to those devices in the defined zone.

Zone Membership

A zone configuration is a set of one or more zones that are enabled together. Compaq Fibre Channel switches allow multiple zone configurations to be defined, of which only one is enabled at a time.

Zone configuration is dynamic. Nodes can be in multiple zones to allow for overlapping, depending on the desired access control. The number of zones and zone memberships are unlimited and, depending on the number of fabric-connected devices and device locations, zones may vary in size and shape.

Devices may be members of more than one zone; however, since zone members see only members in their zones, they can access only one another. This is applicable for enterprise backup, since temporary zones can be created.

Zone members can be any of the following:

- Physical port number

  The physical port number is identified by the switch ID and the port number (such as **1,5).** In this example, **1** is the switch ID and domain ID, and **5** is the port number on that switch.

- Node World Wide Name

  The node World Wide Name is the 16-digit hexadecimal identifier of a Fibre Channel device, such as **50:00:00:20:3c:76:df:00.**

- Port World Wide Name

  The port World Wide Name is the 16-digit hexadecimal identifier of a single port on a multi-port Fibre Channel device, such as **50:00:00:20:3c:76:df:01.**

Zone Alias identifies one or more zone members by a single name, such as **group1**, which is the name of a list of zone members.

NOTE: A zone alias cannot have another zone alias as one of its zone members.

Resource Partitioning

Zoning allows a finer segmentation of SANs by creating logical device subsets and enabling resource partitioning for management and access control. This is accomplished by allowing ports or World Wide Name (WWN) addresses to be used to confine access to devices in a common zone.

Data Security

Zoning adds an additional level of information management and data security when used in addition to the Selective Storage Presentation (SSP) feature of the HSG80 controllers. Switch zoning and SSP complement each other; they do not replace each other. Switch zoning controls communication access. SSP controls data access.

While SSP controls which servers have access to each storage unit, switch zoning controls which servers can communicate with each other and which servers can communicate with each storage controller host port. Zoning controls access at the storage system level. SSP controls access at the storage unit level. Switch zoning is a higher-level access control than SSP.

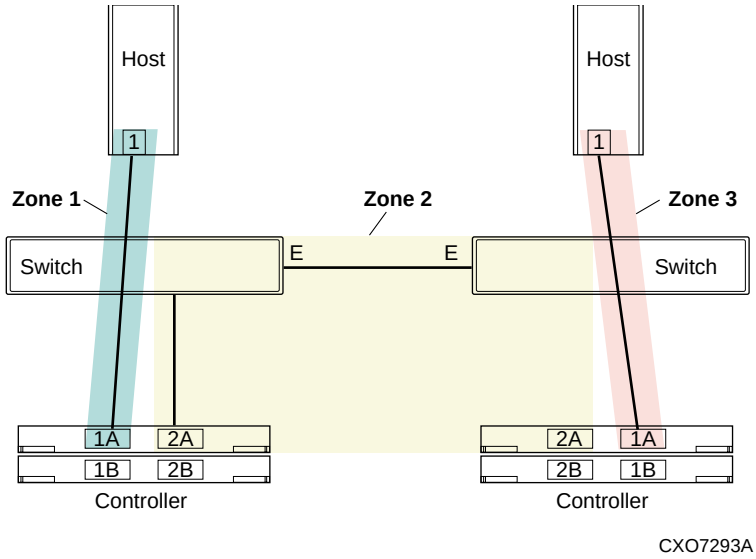

Figure 7–5. Zoning in homogeneous environment

Homogeneous Environment

Figure 7–5 illustrates a simple configuration for three zones in a homogeneous environment. While the figure has been simplified to show a vertical, non-stretch cluster environment for only fabric A, fabric B could be shown as well.

Zoning can be used in heterogeneous mixed platform SAN configurations. However, only homogeneous zoning is supported at this time for Data Replication Manager.

In a zoned environment, access to each device is limited to the other devices that are in the same zone. Any device outside of the zone is unaware of the existence of the devices inside the zone. Zones can be distinct or they can overlap. If a device is defined to be in the overlap of multiple zones, it is aware of the devices in all of those zones and, likewise, all of the devices in all of those zones are aware of it.

A zone is specified by a zone name. A zone member is specified by a physical fabric port number, a node World Wide Name, or a port World Wide Name. A set of zones is configured during zone specification. Zoning specifications are provided in Table 7–3.

.

Table 7–3 Zoning Specifications

Zone Administration	■ Addition and deletion of zone members ■ Display of zone members for all zones, or by individual zone or configuration ■ Creation of aliases ■ Configuration of a set of zones ■ Creation of temporary zones
Zoning Enforcement	■ By Simple Name Server (via software) ■ By specifying the zone by physical fabric port number (via hardware)
Zoning Management	Via telnet
Zoning Backup	Recorded in switch flash memory
Fabric OS Support	Fabric OS, Version 2.0 and above

Zoning Commands

Switch zoning is enabled on the SAN through commands entered in a telnet session logged into one switch using a switch account with administrative privileges. Zone configurations and zones are managed with the commands shown in Table 7–4. These zoning commands are used and described in the zoning examples in “Zoning A DRM Configuration” on page 7–14.

Table 7–4 Zoning Commands

Zoning Component	Available Commands
zones	zoneCreate, zoneDelete, zoneAdd, zoneRemove, zoneShow
zone aliases	aliCreate, aliDelete, aliAdd, aliRemove, aliShow
zone configurations	cfgCreate, cfgDelete, cfgAdd, cfgRemove, cfgShow, cfgEnable, cfgDisable, cfgSave, cfgClear

Planning Considerations for Homogeneous Configurations that Require Zoning

Planning is an essential part of the zoning process. Four planning considerations are applicable to homogeneous Data Replication Manager configurations, as detailed below.

Windows 2000 Using Secure Path

In no single point of failure (NSPOF) configurations, Windows 2000 servers must use Compaq Secure Path V3.1. Secure Path uses a pair of physically separate fabrics, or a pair of zones, to provide the data path redundancy.

For each Windows 2000 server, Secure Path supports connecting to only two active host ports on a storage system. Only Host Port 1 of both controllers is connected to the pair of zones or pair of fabrics; each is connected to a separate zone or fabric.

Host Port 2 of both controllers cannot be connected to the same pair of zones or pair of fabrics. Host Port 2 of both controllers can be used if they are in a second pair of zones, or are connected to a second pair of redundant fabrics, and they are used by a different set of Windows 2000 servers running Secure Path.

More than 64 Host Connections

A host connection is a data path from one host bus adapter to one active controller host port, regardless of whether the host connection uses storage units on that storage system.

The ESA12000 / RA8000 has a limit of 64 host connections. When a 65th connection is attempted, the connection name capacity of HSG80 controllers is exceeded. This could result in a controller fatal error.

In transparent failover mode two host ports are active. One host server with one host bus adapter has two data paths, one to each active controller host port. This allows up to 32 servers to have data paths to the HSG80 controllers.

In multiple bus failover mode, four host ports are active. One server with two host bus adapters, as in an NSPOF configuration, has eight data paths, four for each host bus adapter. This allows up to eight servers to have data paths to the HSG80 controllers.

Switch zoning is used to prevent more than 64 host connections from being made to a single storage system. This allows more servers to be attached to the SAN to use other storage systems.

Prevent Host Bus Adaptor (HBA) from Seeing All Active Host Ports

Every storage controller host port on the SAN is a separate target for the HBA. If there are more host ports than the HBA has targets for, it is possible the storage it needs to access will not be given a target by the HBA.

For example, a server's HBA can have a maximum of 16 targets. On a specific SAN there are 3 storage systems configured for transparent failover mode and 2 storage systems configured for multiple bus failover mode. Each of the storage systems configured for transparent failover mode has 2 active host ports.

In this example, each of the storage systems configured for multiple bus failover mode has 4 active host ports. There are 20 active host ports in the SAN. This means that 4 of the host ports cannot be given a target by the HBA.

Switch zoning prevents the HBA from seeing active host ports that it does not use. This allows the host ports it does use to be given targets.

Multiple Tru64 UNIX Clusters

Tru64 UNIX communicates on the SAN in both initiator and target mode. When a Tru64 UNIX server configured for clustering boots, it queries the SAN for all targets with which it can communicate. When another Tru64 UNIX server configured for clustering responds, both servers consider each other to be members of the same cluster, when actually they may be in different clusters. Switch zoning precludes this confusion by placing each Tru64 UNIX cluster in separate zones.

Zoning A DRM Configuration

The Fibre Channel fabric can be customized for zoning in numerous ways. The DRM uses for zoning include:

- Creating fabric functional areas by separating test or maintenance areas from production areas.
- Designating closed user groups by including certain zone devices for exclusive use by zone members.
- Simplifying resource utilization by logically consolidating equipment for convenience.
- Facilitating time-sensitive functions by creating a temporary zone used to back up a set of devices that are members of other zones.
- Securing fabric areas by providing another level of software security to control port level access.

Zoning a Data Replication Manager configuration is described on the following pages. For additional information, refer to the *SAN Switch Zoning Reference Guide.*

Figure 7–6 shows a simplified Data Replication Manager configuration comprised of three zones. These zones have been designated for this example as “green” for the initiator site, “blue” for the target site, and “red,” which contains the remote copy sets or paths.

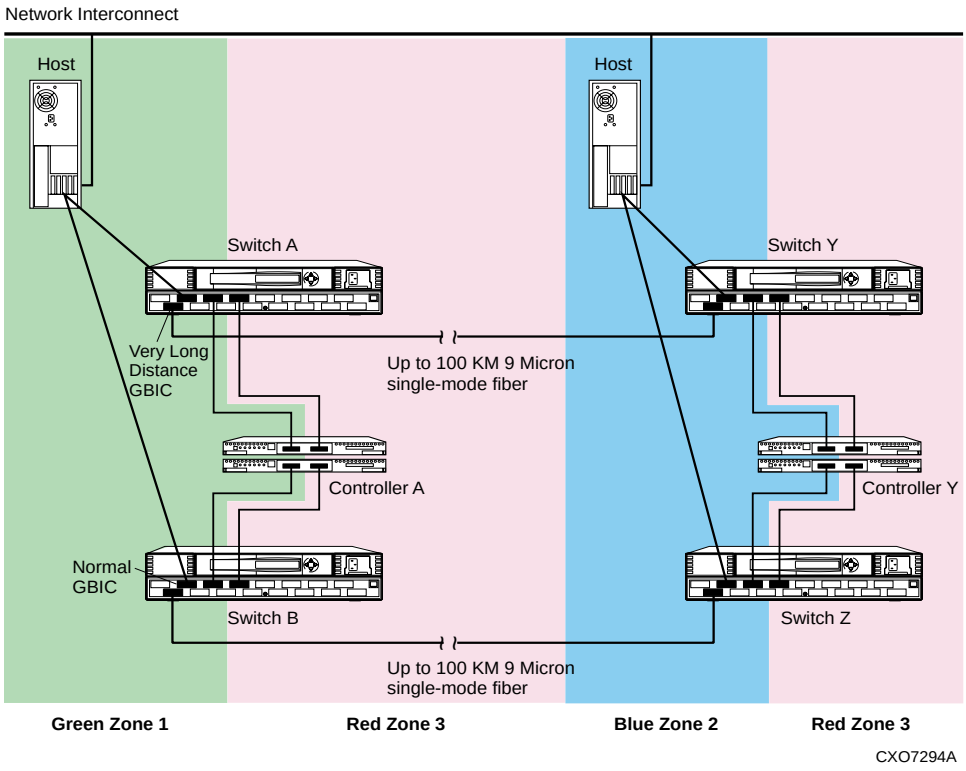

Figure 7–6. Zoning in a DRM homogeneous environment

Compaq suggests that you use Figure 7–7, “Blank Zoning Input Form (Template)”, to capture and track the required device and command information. The form has been designed to support 2 paths, 2 switches, and a maximum of 16 entries per switch. Copy and use a separate form to track information for each zone (such as Green Zone, Blue Zone, Red Zone). Organize alias name, function, and site data by either World Wide ID (WWID) number or Port ID number.

Zoning Configuration Name=

Zone Name= **Switch Name=** **Path=**

WWID #	Domain ID #	Port #	Alias Name	Function	Site

Zoning Configuration Name=

Zone Name= **Switch Name=** **Path=**

WWID #	Domain ID #	Port #	Alias Name	Function	Site

Sheet #

Figure 7–7. Blank Zoning Input Form (Template)

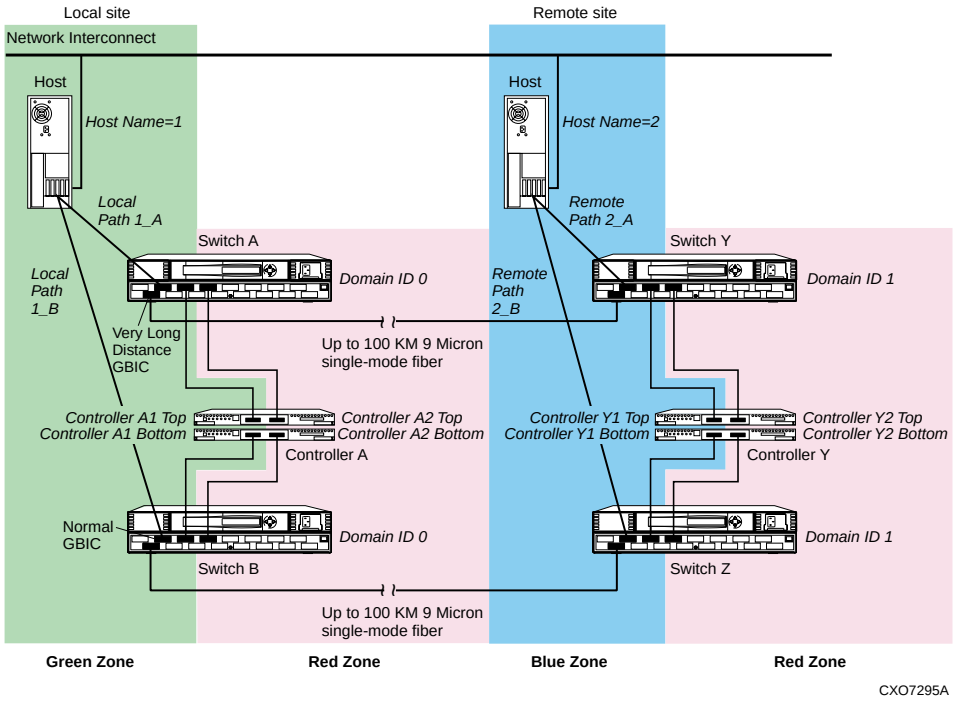

Figure 7–8. Zoning a DRM example

Figure 7–8 illustrates a zoned example for a Data Replication Manager configuration. The steps required to produce this zoning configuration are shown on the following pages. Since multiple configurations can be created and saved, the one currently in use is known as the *effective configuration*. The process to create and save an effective configuration that ensures the switches are enabled after reboot or shutdown includes:

1. Begin by zoning the initiator site (designated in example as Green Zone).
2. Then zone the target site (designated as Blue Zone).
3. Finally, zone the remote copy sets or paths (designated as Red Zone).

Three examples follow that illustrate these zoning concepts.

Zoning the Green Zone (Example)

Figure 7–9, the Green Zone Input Form, was created from a blank template (Figure 7–7) and was added to throughout this example.

Zoning Configuration Name=Zoning_DRM

Zone Name=Green Zone **Switch Name=Switch A** **Path=A**

WWID #	Domain ID #	Port #	Alias Name	Function	Site
	0	1	-	E port	Local
	0	2	Host1_A	Host	Local
	0	4	Controller A1_top	Controller	Local

Zoning Configuration Name=Zoning_DRM

Zone Name=Green Zone **Switch Name=Switch B** **Path=B**

WWID #	Domain ID #	Port #	Alias Name	Function	Site
	0	1	-	E port	Local
	0	2	Host1_B	Host	Local
	0	4	Controller A1_bottom	Controller	Local

Sheet # 1 of 3

Notes;

Figure 7–9. Green Zone Input Form

Figure 7–8 shows zoning using the **Domain ID #** and the **Port #**, rather than the **WWID #**. The World Wide ID could also have been used. A general rule of thumb is that if you are changing connections within the Data Replication Manager, use the World Wide ID for zoning; however, if you are changing out hardware, use the Domain ID and port number.

1. Identify and write down the Domain ID of the switch. To get this information, use the switchShow command from the switch in a Telnet session or from the front console of the switch. Refer to the *"SHOW Command Examples"* section in this chapter for more information on the switchShow command. The example in Figure 7–8 shows that switches A and B both have a Domain ID of 0.

   NOTE: Refer to the *Compaq StorageWorks Fibre Channel SAN Switch Management Guide* for telnet session procedures.

After logging this information on a blank template form, give it the zone name “Green Zone.” On this input form, list the Green Zone in two blocks, one for switch A and one for Switch B. Refer to the Green Zone Input Form for entries.

2. Log the ports that are being used to connect to the hosts and E ports.

   The example shows that Host 1 is the host name and that there are two connections from Host 1 to the switches. Host 1 path 1_A is connected to port 2 of switch A, and Host 1 path 1_B is connected to port 2 of switch B. The E port is located in port 1 on both switches A and B.

3. List the controller connections.

   From the example in Figure 7-8, the controller pair is listed as Controller A1_top (meaning top controller, port 1), and Controller A1_bottom (bottom controller port 1). Switch A port 4 connects to Controller A1_top. Switch B port 4 connects to Controller A1_bottom.

4. Open a telnet session to switch A.

5. Create the alias names in the zone. The naming convention used in the example refers to Host 1_A as the host name and path used for the connection to port 2 of switch A. The command used for this alias is:

   ```
   aliCreate "Host 1_A", "0,2"
   ```

   This means that an alias is generated named “Host 1_A” with a switch domain of 0 and port number of 2.

IMPORTANT: Be sure to issue all aliCreate, ZoneAdd, and CfgAdd commands from a switch within the fabric for which the alias is being created.

   The next alias is for controller A1_top. The command for this alias is:

   ```
   aliCreate "Controller A1_top", "0,4"
   ```

   This generates an alias named “Controller A1_top” with a switch domain of 0 and a port number of 4.

   Since E ports cannot be zoned, an alias is not needed for domain 0 port 1 in this example.

6. The next alias to create is from the Switch B Green name template with the following command:

   ```
   aliCreate "Host 1_B", "0,2"
   ```

7. Next, create the controller A-1 bottom alias with the following command:

   ```
   aliCreate "Controller A1_bottom", "0,4"
   ```

8. Save the configuration. The command is:

   cfgSave

9. Open a telnet session to switch B.
10. Perform a switchShow command to confirm host, controller, and E port connections.
11. Repeat steps 4 through 7 above for Switch B.

This concludes the alias naming for the Green Zone.

Next, we'll configure the Blue Zone in this exercise.

Zoning the Blue Zone (Example)

Zoning the Blue portion of this example is similar to the steps for the Green Zone, with a few name and number changes. See Figure 7–10 for the sample Blue Zone Input Form.

Zoning Configuration Name= Zoning_DRM

Zone Name=Blue Zone **Switch Name=Switch Y** **Path=A**

WWID #	Domain ID #	Port #	Alias Name	Function	Site
	1	1	--	E port	Remote
	1	2	Host 2_A	Host	Remote
	1	4	Controller Y1_top	Controller	Remote

Zoning Configuration Name= Zoning_DRM

Zone Name=Blue Zone **Switch Name=Switch Z** **Path=B**

WWID #	Domain ID #	Port #	Alias Name	Function	Site
	1	1	--	E port	Remote
	1	2	Host 2_B	Host	Remote
	1	4	Controller Y1_bottom	Controller	Remote

Sheet # 2 of 3

Notes;

Figure 7–10. Blue Zone Input Form

As shown in Figure 7–8, the Host in the Blue Zone is named “Host 2.“

1. Log the domain IDs of switches Y and Z. In this example, they are both Domain ID 1.

   After logging this information on a blank template form, give it the zone name “Blue Zone.” On this input form, list the Blue Zone in two blocks, one for switch Y and one for switch Z. See Figure 7–10 for entries.

2. Log the ports that are being used to connect to the hosts and E ports.

   Figure 7–8 shows that Host 2 is the host name and that there are two connections from Host 2 to the switches. Host 2 path 2_A is connected to port 2 of switch Y, and Host 2 path 2_B is connected to port 2 of switch Z. The E port is located in port 1 on both switches Y
   and Z.

3. List the controller connections.

   Figure 7–8 shows that the controller pair is listed as Controller Y1_top (meaning top controller, port 1), and Controller Y1_bottom (bottom controller port 1). Switch Y port 4 connects to Controller Y1_top. Switch Z port 4 connects to Controller Y1_bottom.

4. Select the telnet session from switch A.

5. Create the alias names in the zone.

   The naming convention used in the example refers to Host 2_A as the host name and path used for the connection to port 2 of switch A. The command used for this alias is:

   ```
   aliCreate "Host 2_A", "1,2"
   ```

   An alias is generated, named “Host 2_A”, with a switch domain of 1 and port number 2.

   The next alias is for Controller Y1_top. The command for this alias is:

   ```
   aliCreate "Controller Y1_top", "1,4"
   ```

   An alias is generated, named “Controller Y1_top”, with a switch domain of 1 and a port number 4.

   Since E ports cannot be zoned, an alias is not needed for domain 1 port 1.

6. The next alias to create is from the Switch Z Blue name template. The command is:

   ```
   aliCreate "Host 2_B", "1,2"
   ```

7. Next, create the controller Y-1 bottom alias with the following command:

   ```
   aliCreate "Controller Y1_bottom", "1,4"
   ```

8. Save the configuration. The command is:

   ```
   cfgSave
   ```

9. Select the telnet session from switch B.
10. Repeat above steps 4 through 7 for Switch Z.

This concludes the alias naming for the Blue Zone.

Next, we'll configure the Red Zone.

Zoning the Red Zone (Example)

Zoning the Red Zone portion of this example is similar to the steps for the Blue and Green zones, with a few exceptions. See Figure 7–11 for the sample Blue Zone Input Form.

Zoning Configuration Name=Zoning_DRM

Zone Name=Red Zone **Switch Name=Switch A & B** **Path=A & B**

WWID #	Domain ID #	Port #	Alias Name	Function	Site
	0	6	Controller A2_Top	Controller	Local
	0	6	Controller A2_Bottom	Controller	Local

Zoning Configuration Name=Zoning_DRM

Zone Name=Red Zone **Switch Name=Switch Y & Z** **Path=A & B**

WWID #	Domain ID #	Port #	Alias Name	Function	Site
	1	6	Controller Y2_Top	Controller	Remote
	1	6	Controller Y2_Bottom	Controller	Remote

Sheet # 3 of 3

Notes;

Figure 7–11. Red Zone Input Form

As shown in Figure 7–8, the Red Zone contains only switch ports and controller ports; no hosts.

These are the DRM remote copy paths.

1. Log the Domain Ids switches A, B,Y and Z. As shown earlier, switches A and B are Domain ID 0, and switches Y and Z are Domain ID 1.

   After logging this information on a blank template form, give it the zone name “Red Zone.” On this input form, list the Red Zone in two blocks; one for switches A and B, and one for switches Y and Z. Refer to Figure 7–11 for entries.

2. List the controller connections.

   Figure 7–8 shows the controller pair listed as: Controller A2_top (meaning top A controller, port 2), Controller A2_bottom (bottom A controller port 2), Controller Y2_top (top Y controller, port 2), and Controller Y2_bottom (bottom Y controller, port2).

   Switch A port 6 connects to Controller A2_top. Switch B port 6 connects to Controller A2_bottom. Switch Y port 6 connects to Controller Y2_top. Switch Z port 6 connects to Controller Y2_bottom.

3. Select the telnet session from switch A.

4. Create the alias names in the zone. For the connections from switches A and Y to the controllers A2_top and Y2_top, the command is:

   ```
   aliCreate "Controller A2_top", "0,6"
   aliCreate "Controller Y2_top", "1,6"
   ```

   These commands create the alias “Controller A2_top” with a Domain ID of 0 and switch port 6, and the alias “Controller Y2_top” with a Domain ID of 1 and switch port 6.

5. For the connections from switches B and Z to the controllers A2_bottom and Y2_bottom, enter the following commands:

   ```
   aliCreate "Controller A2_bottom", "0,6"
   aliCreate "Controller Y2_bottom", "1,6"
   ```

   These commands create the alias “Controller A2_bottom” with a Domain ID of 0 and switch port 6, and the alias “Controller Y2_bottom” with a Domain ID of 1 and switch port 6.

6. Save the configuration. The command is:

   ```
   cfgSave
   ```

7. Select a telnet session from Switch B.

8. Repeat above steps 3 through 5 using the Red Zone template alias names for Switches A, B, Y, and Z.
9. After entering the alias names, perform another configuration save. The command is:

```
cfgSave
```

Create the Zone Names

1. Select the telnet session from switch A.
2. Create the Green zone name and add the zone members. The command is:

```
zoneCreate "Green Zone", "Host 1_A; Controller A1_top; Host 1_B; Controller
A1_bottom"
```

3. Create the Blue zone name and add the zone members. The command is:

```
zoneCreate "Blue Zone", "Host 2_A; Controller Y1_top; Host 2_B; Controller
Y1_bottom"
```

4. Create the Red zone name and add the zone members. The command is:

```
zoneCreate "Red Zone", "Controller A2_top; Controller A2_bottom; Controller
Y2_top; Controller Y2_bottom"
```

   These three steps created the zone names that are stored in flash memory in both switches A and Y. The next step is to repeat the above commands for switches B and Z.

5. Select the telnet session from switch B.
6. Create the Green zone name and add the zone members. The command is:

```
zoneCreate "Green Zone", "Host 1_A; Controller A1_top; Host 1_B; Controller
A1_bottom"
```

7. Create the Blue zone name and add the zone members. The command is:

```
zoneCreate "Blue Zone", "Host 2_A; Controller Y1_top; Host 2_B; Controller
Y1_bottom"
```

8. Create the Red zone name and add the zone members. The command is:

```
zoneCreate "Red Zone", "Controller A2_top; Controller A2_bottom; Controller
Y2_top; Controller Y2_bottom"
```

   These three steps created the zone names that are stored in flash memory in both switches B and Z.

9. Save the configuration. The command is:

```
cfgSave
```

Create the Configuration Name

1. Select the telnet session from switch A.
2. Create the configuration using “Zoning_DRM” as the example name and add all of the zone members.
   The command is:

   ```
   cfgCreate "Zoning_DRM", "Green Zone; Blue Zone; Red Zone"
   ```

   This creates a configuration file titled “*Zoning_DRM”*, containing the Green, Blue, and Red Zones and their alias members, which is stored in flash memory for switches A and Y.
3. Save the configuration. The command is:

   ```
   cfgSave
   ```
4. Enable the new zone configuration with the following command:

   ```
   cfgEnable "Zoning_DRM"
   ```

   This now becomes the effective (in use) configuration for both switches A and Y.
5. To make this the active configuration after a restart or power down, issue one final cfgSave command:

   ```
   cfgSave
   ```

   This ensures the effective configuration of the switches after a restart or power down.
6. Select the telnet session from switch B.
7. Create the configuration using “*Zoning_DRM*” as the filename and add all of the zone members.
   The command is:

   ```
   cfgCreate "Zoning_DRM", "Green Zone; Blue Zone, Red Zone"
   ```

   This creates a configuration file titled Zoning_DRM, containing the Green, Blue, and Red Zones and their alias members, which is stored in flash memory for switches B and Z.
8. Save the configuration. The command is:

   ```
   cfgSave
   ```
9. Enable the new zone configuration with the following command:

   ```
   cfgEnable "Zoning_DRM"
   ```

This now becomes the effective (in use) configuration for both switches B and Z.

10. To make this the active configuration after a restart or power down, issue one final cfgSave command:

    ```
    cfgSave
    ```

    This ensures the effective configuration of the switches after a restart or power down.

Zoning is now complete.

SHOW Command Examples

The following SHOW commands, created from a telnet session, are examples of switch commands that can provide valuable configuration information. Refer to the *Compaq StorageWorks Fibre Channel SAN Switch Management Guide* for Telnet session procedures.

The examples show actual command output, followed by a description of the output.

switchShow command

```
----------------------------------------------------------------------------------------------
drm_switch_1: admin>switchShow
switchName:          drm_switch_1
switchState:         Online
switchRole:          Principal
switchDomain:        0
switchId:            fffc40
switchWwn:           10:00:00:60:69:00:51:9d
port  0:    sw    Online   F-Port  10:00:00:00:c9:20:ca:a0
port  1:    sw    Online   F-Port  50:00:1f:e1:00:00:41:61
port  2:    sw    Online   F -Port  50:00:1f:e1:00:00:41:62
port  3:    lw    Online   E-Port  10:00:00:60:69:00:52:be “drm_switch_2” (downstream)
port  4:    --    No_Module
port  5:    --    No_Module
port  6:    --    No_Module
port  7:    --    No_Module
```

```
port   8:      --      No_Card
port   9:      --      No_Card
port  10:      --      No_Card
port  11:      --      No_Card
port  12:      --      No_Card
port  13:      --      No_Card
port  14:      --      No_Card
port  15:      --      No_Card
value = 16 = 0x10
drm_switch_1:admin>
value = 16 = 0x10
drm_switch_1:admin>
-----------------------------------------------------------------------------------------------
```

The **switchShow** command provides the following direct or implied information:

1. The name of the switch is *drm_switch_1*.
2. The status of the switch is *Online,* which means it is not disabled.
3. The switch role is the *Principal* switch in the fabric. In other words, it generates the N-port identifiers (the 24-bit address).
4. The WWN of the switch itself (*drm_switch_1*) is 10:00:00:60:69:00:51:9d.
5. The Domain ID of the switch is *0*.
6. There are four GBICs installed (ports 0-3). Port 3 is longwave. This is the link to the other switch, because it is an E-port.
7. On port 0, a host with WWN *10:00:00:00:c9:20:ca:a0* is connected. The operating system this host is running is not identifiable from this information.
8. On ports 1 and 2, an HSG80 controller is connected. Host port 2 is connected to switch port 2.
9. Looking at the addresses of the HSG80 on ports 1 and 2, note that these are the top and bottom controllers of the HSG80 controller pair. (The addresses are the same except for the last numbers, 61 and 62, and are one number apart.)

fabricShow command

```
---------------------------------------------------------------------------------------------
drm_switch_1:admin> fabricShow
Switch    ID     Worldwide Name           Enet IP Address     FC IP Address    Name
---------------------------------------------------------------------------------------------
   0:    fffc40  10:00:00:60:69:00:51:9d   16.82.208.186        0.0.0.0        >"drm_switch_1"
   1:    fffc41  10:00:00:60:69:00:52:be   16.82.208.189        0.0.0.0        >"drm_switch_2"
value = 1 = 0x1
drm_switch_1:admin>
---------------------------------------------------------------------------------------------
```

The **fabricShow** command provides information about the configuration of the fabric:

1. There are two switches in the fabric.
2. The other switch in the fabric has Domain ID 1 and is called *drm_switch_2*. The WWN is *10:00:00:60:69:00:52:be*.
3. The IP addresses of both switches are identified as Domain ID 0 (*10:00:00:60:69:00:51:9d)* and Domain ID 1 (*10:00:00:60:69:00:52:be*)

uRouteShow command

```
---------------------------------------------------------------------------------------------
drm_switch_1:admin> uRouteShow
Local Domain ID:    0
In Port      Domain      Out Port    Metric      Hops      Flags          Next (Dom, Port)
---------------------------------------------------------------------------------------------
   0          1            3          1000         1         D                1, 3
Type <CR> to continue, Q<CR> to stop:     <CR>
   1          1            3          1000         1         D                1, 3
Type <CR> to continue, Q<CR> to stop:     <CR>
   2          1            3          1000         1         D                1, 3
value = 1 = 0x1
drm_switch_1:admin>
---------------------------------------------------------------------------------------------
```

The **uRouteShow** command gives the route of an F-port to another domain. Since there are only two domains in the fabric, it specifies the path to the other switch.

If F-port 0,1,2 needs to access Domain 1, the E-port 3 is used, which is connected in Domain 1 to E-port 3. (F-port and E-port information is taken from the ***switchShow*** command output.)

topologyShow command

```
---------------------------------------------------------------------------------------------
drm_switch_1:admin> topologyShow
Local Domain ID:      0
Domain    Metric      Hops    Out Port        In Ports          Flags          Name
---------------------------------------------------------------------------------------------
    1       1000        1          3          0x00000007        D           “drm_switch_2”
value = 1 = 0x1
drm_switch_1:admin>
---------------------------------------------------------------------------------------------
```

The **topologyShow** command identifies usage of E-ports by corresponding F-ports. In a configuration with multiple host bus adapters and/or multiple storage systems per site, this information indicates whether or not performance problems can be expected.

The value listed under *In Ports* is a bitmap. Only the last two bytes are used. 0x00000007 breaks down to bit 2,1 and 0 set to a “1”. It says:

To go to domain “1” I will use E-port 3. The F-ports using E-port 3 are 2, 1 and 0.

With multiple links between the sites and more than one storage system, it is definitely desirable to have host port 2 of the controllers operate on different E-ports.

There is currently no way to enforce this mapping. It is dynamic and depends on the timing during FLOGI and PLOGI of the different controllers.

nsShow command

```
---------------------------------------------------------------------------------------------
drm_switch_1:admin> nsShow
The Local Name Server has 3 entries:
Type  Pid          COS    PortName                  NodeName                  TTL(sec)
 N    200013;   1,2,3    10:00:00:00:c9:20:ca:a0   10:00:00:00:c9:20:ca:a0;     na
 N    200213;      3      50:00:1f:e1:00:00:41:61   50:00:1f:e1:00:00:41:60:     na
      FC4s:  FCP  [DEC     HSG80                  0000]
 N    200213;      3      50:00:1f:e1:00:00:41:62   50:00:1f:e1:00:00:41:60:     na
      FC4s:  FCP
      PortSymb:  [23]  0x4853475f505052435f434e54525f494e49545242 02020
value = 0 = 0x0
drm_switch_1:admin>
---------------------------------------------------------------------------------------------
```

The **nsShow** command provides the following direct or implied information:

1. There are three ports on this switch logged into the fabric.
2. All three 3 ports are N-ports.
3. Port 0 (derived from the Pid) offers Class Of Service 1, 2 and 3.
4. Port 1 (derived from the Pid) offers Class 3 service only.

   The port name is: 50:00:1f:e1:00:00:41:61 (which is known from the *switchShow* command output).

   The node name is: 50:00:1f:e1:00:00:41:60 (which is the WWN of the BA370).
5. Port 1 is an FCP-capable FC4 device. FCP means *Fibre Channel protocol*. The identification of that part of the standard is FC-FCP.
6. Here it is clearly seen that port 1 is host port 1 of the HSG80 controller, since it holds the Inquiry string.
7. Port 2 is very similar to port 1. However, there is a port symbol,

   0x4853475f505052435f434e54525f494e49545242

   (Or for those who prefer ASCII),

   H S G_P P R C_C N T R_I N I T R B

SET THIS REMOTE = INTR has been done on this controller pair. The REMOTE_COPY_NAME of this controller pair is therefore now known.

The *B* after the REMOTE_COPY_NAME indicates the bottom controller of the HSG80 controller pair.

The **nsShow** command must be given on both switches in the fabric to get a complete view, including addresses.

nsAllShow command

```
---------------------------------------------------------------------------------------------
drm_switch_1:admin> nsAllShow
6 Nx_Ports in the Fabric
  200013      200113      200213      210013      210113      210213
value = 0 = 0x0
drm_switch_1:admin>
---------------------------------------------------------------------------------------------
```

The **nsAllShow** command shows that the other domain (with Domain ID “0”) has 3 F-ports: ports 0, 1 and 2 on the other switch. The WWNs of the N-Ports connected to the other switch are not identifiable.

errShow command

```
-------------------------------------------------------------------------------------------------
drm_switch_1:admin> errShow
Error 01
----------
0x103e9500  (tSwitch): Mar 10 02:49:56
        Error SYS-BOOT, 3, Restart reason: Power-on
value = 1 = 0x1
drm_switch_1:admin>
-------------------------------------------------------------------------------------------------
```

The **errShow** command indicates that the switch was booted on March 10, due to a power- on situation.

NOTE: The error log is in volatile RAM, so it is lost after restart and power down.

Appendix A

Status Comparison

This appendix describes the procedure for comparing the status of:

- Controllers
- Association sets
- Remote copy sets
- Units
- Connections

Performing a status comparison consists of the following three procedures:

- Target Site Terminal Emulator Session
- Issuing SHOW Commands

Target Site Terminal Emulator Session

1. Using a serial cable, connect the COM port of a laptop computer or another computer to the corresponding serial port on the HSG80 controllers.
2. Start a terminal emulator session that is capable of capturing text to a file (which will later be saved as step 6 of SHOW Commands procedure). Use the following settings: 9600 baud, 8 bits, No parity, 1 stop bit, XON/XOFF.

Issuing SHOW Commands

1. To see the full information on this controller, issue the following CLI command:

   ```
   SHOW THIS_CONTROLLER FULL
   ```

   You will see a display similar to that shown in Example Display 1.

2. To see the information for all association sets known to the controller pair, issue the following CLI command:

   ```
   SHOW ASSOCIATIONS FULL
   ```

   You will see a display similar to that of Example Display 2 for each association set.

3. To see information for all remote copy sets known to the controller pair, issue the following CLI command:

   ```
   SHOW REMOTE_COPY FULL
   ```

   You will see a display similar to that in Example Display 3 for each remote copy set.

4. To see information for all units configured to the controller, issue the following CLI command:

   ```
   SHOW UNITS FULL
   ```

   You will see a display similar to that of Example Display 4 for each unit.

5. To see the connection name, operating system, controller, controller port, adapter ID address, online or offline status, and unit offset, issue the following CLI command:

   ```
   SHOW CONNECTIONS
   ```

   You will see a display similar to that of Example Display 5 for each connection.

6. Save for future reference the file started during the terminal emulator session procedure (step 2). This file contains the text captured throughout steps 1-5 of the SHOW Commands.

"Example Display 1" corresponds to step 1 of the "SHOW Commands" section:

Example Display 1

```
Controller:
        HSG80 ZG91412410 Software V85P, Hardware  E05
        NODE_ID          = nnnnnnnnnnn
        ALLOCATION_CLASS = 0
        SCSI_VERSION     = SCSI-2
        Configured for MULTIBUS_FAILOVER with ZG91416136
            In dual-redundant configuration
        Device Port SCSI address 6
        Time: NOT SET
        Command Console LUN is lun 0 (NOIDENTIFIER)
Host PORT_1:
        Reported PORT_ID = 5000-1FE1-0001-3AE1
        PORT_1_TOPOLOGY  = FABRIC (fabric up)
        Address          = 220113
Host PORT_2:
        Reported PORT_ID = 5000-1FE1-0001-3AE2
        PORT_2_TOPOLOGY  = FABRIC (fabric up)
        Address          = 220313
        REMOTE_COPY = BuildingB
Cache:
        256 megabyte write cache, version 0012
        Cache is GOOD
        No unflushed data in cache
        CACHE_FLUSH_TIMER = DEFAULT (10 seconds)
Mirrored Cache:
        256 megabyte write cache, version 0012
        Cache is GOOD
        No unflushed data in cache
Battery:
        NOUPS
        FULLY CHARGED
        Expires:
```

```
Extended information:
        Terminal speed 9600 baud, eight bit, no parity, 1 stop bit
        Operation control: 00000000  Security state code: 75184
        Configuration backup disabled
```

"Example Display 2" corresponds to step 2 of the "SHOW Commands" section:

Example Display 2

```
Name          Association                   Uses             Used by
-------------------------------------------------------------------------
AS1           association                   RC1
                                            RC2
                                            RC3
        Switches:
          NOFAIL_ALL
          NOORDER_ALL
          NOLOG_UNIT
```

"Example Display 3" corresponds to step 3 of the "SHOW Commands" section:

Example Display 3

```
Name                                        Uses             Used by
-------------------------------------------------------------------------
RC1           remote copy                   D1               AS1
        Reported LUN ID: nnnnnnnnnnnnnnn
        Switches:
          OPERATION_MODE  = SYNCHRONOUS
          ERROR_MODE      = NORMAL
          FAILOVER_MODE   = MANUAL
          OUTSTANDING_IOS = 60
                            .
                            .
                            .
```

"Example Display 4" corresponds to step 4 of the "SHOW Commands" section:

Example Display 4

```
D2                                          DISK10100      BuildingB\RC2
        LUN ID: nnnnnnnnnnnnnnnnnnnnn
        NOIDENTIFIER
        Switches:
         RUN                 NOWRITE_PROTECT       READ_CACHE
          READAHEAD_CACHE        WRITEBACK_CACHE
          MAXIMUM_CACHED_TRANSFER_SIZE = 1
        Access:
    BuildngAA, BuildngAB, BuildngAC, BuildngAD, HostCon_1, HostCon_2
        State:
          ONLINE to this controller
          Not reserved
          PREFERRED_PATH = OTHER_CONTROLLER
          Target NORMAL
        Size:             17769177 blocks

    Geometry (C/H/S): ( 5258 / 20 / 169 )
```

"Example Display 5" corresponds to step 5 of the "SHOW Commands" section:

Example Display 5

```
Connection                                              Unit
Name Operating system Controller  Port    Address    Status
Offset !NEWCON28 WINNT        THIS       1      634000   OL this       0
          HOST_ID=1000-0000-C921-4B5B ADAPTER_ID=1000-0000-C921-4B5B.
```

Appendix B

Replicating Storage Units

This chapter describes Data Replication Manager concepts and procedures for making point-in-time copies of a storage unit.

The topics contained within this chapter are:

Cloning and *Snapshot* are methods of making a point-in-time copy of a storage unit. Table B–1 provides an overview comparison.

Table B–1 Cloning and Snapshot Comparison

Cloning	Snapshot
Can be done at Initiator site or Target site	Can be done at Target site only
Resides on disk Source and Cloned units	Resides in cache Requires 512 MB cache for both controllers
Read/Write capability	Read/Write capability - Source Unit Read Only capability - Snapshot Unit
Data captured in hours (for moderate I/O loads, possibly at a rate of 60 GB/hour)	Data captured in seconds
Unlimited clones	4 Snapshot Units allowed per Source Unit
The Source Unit must have the following characteristic: Write-back cache disabled	The Source Unit must have the following characteristics: ■ Less than 512 GB ■ Write-back cache enabled ■ Non-transportable
Can clone an unpartitioned single-disk unit, stripeset, or mirrorset	The Snapshot Unit must have the following characteristics: ■ Write-back cache enabled ■ Capacity equal to or greater than the Source Unit ■ Made of any storage container, including partitions, with the exception of log containers
Source Unit and Clone both reside on and failover on the same controller.	Source Unit and Snapshot Unit both reside on and failover on the same controller.
Operates in both multi-bus and controller failover modes.	Operates in both multi-bus and controller failover modes.

Cloning Data for Backup

Use the CLONE utility to duplicate 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 the single-disk unit, which can continue to service its I/O load. When you are cloning a mirrorset, CLONE does not need to create a temporary mirrorset. Instead, it adds a temporary member to the mirrorset and copies the data onto this new member.

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

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

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.

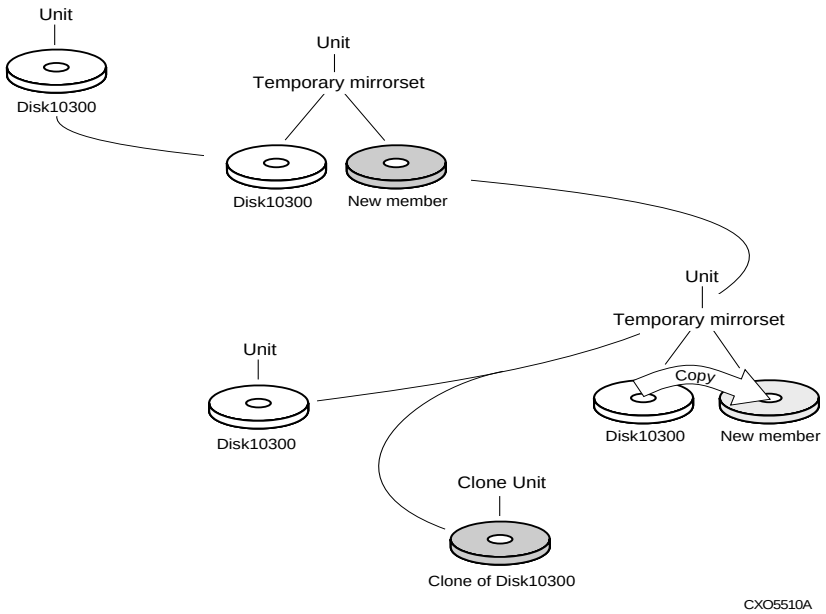

Figure B–1. Steps the CLONE utility follows for duplicating unit members

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)
DISK30100 (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)
DISK30100 (SIZE=832317)
```

```
USE AVAILABLE DEVICE DISK10400(SIZE=832317) FOR MEMBER DISK(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
```

Snapshot

With snapshot, the contents of a Source Unit are frozen in time and presented to the host as a second unit, the *snapshot*. The Snapshot Unit (Figure B–2) preserves the original data (from the time of the snapshot) while allowing writes to the Source Unit to continue. A temporary volume (the Snapshot Unit) is created and used to store the original data that has been overwritten on the Source Unit since the time of the snapshot.

Using a cache bitmap, reads are directed to the Source Unit or the Snapshot Unit.

If no data has been written to the Source Unit since the time of the snapshot, data is then read from the Source Unit.

If data has been written to the Source Unit, then data is read from the Snapshot Unit.

Figure B–2. Snapshot unit

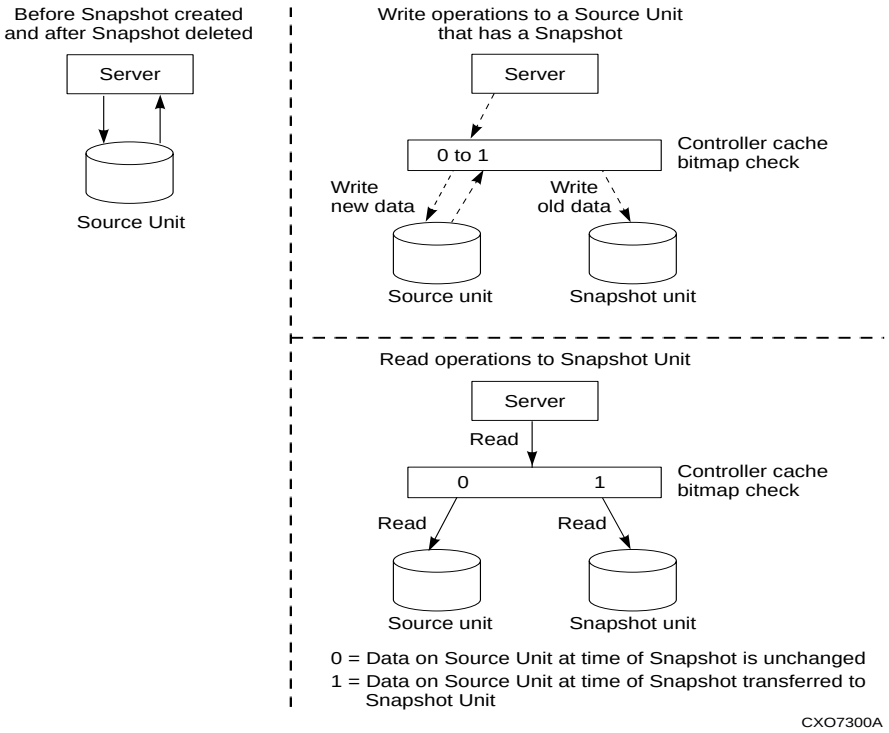

Snapshot Command

NOTE: This command is operational only in controller software versions V8.5S and V8.5P and is operational only if both controllers have 512MB mirrored cache.

This command creates and names a Snapshot Unit. A Snapshot Unit is one that reflects the contents of another unit at a particular point in time (the instant the ADD SNAPSHOT_UNITS command is entered). The Snapshot Unit can then be presented to the host. The Snapshot Unit remains until it is deleted (DELETE command).

Syntax

ADD SNAPSHOT_UNITS *snapshot-unit storageset source-unit*

Parameters

The following parameters are required for the ADD SNAPSHOT_UNITS command:

- Snapshot Unit
- Storageset
- Source Unit

The relationship of the parameters can be summarized as follows:

When the ADD SNAPSHOT_UNITS command is entered, *storageset* becomes *snapshot-unit* and archives the current contents of *source-unit* at that instant.

These parameters are described in the paragraphs that follow.

snapshot-unit

The unit number assigned to the Snapshot Unit. The unit number must start with a letter (A through Z) and may consist of a maximum of nine characters, including letters A through Z, numbers 0 through 9, periods (.), dashes (-), or underscores (_).

The Snapshot Unit is created with all host access disabled by default. Do a SET command to set up host access.

The Snapshot Unit is created on the same controller as the Source Unit, and must remain there.

storageset

Identifies the storageset that becomes the Snapshot Unit. The storageset must:

- Have a capacity equal to or greater than the Source Unit
- Be initialized
- Not be a partition or a concatset

Source Unit

The unit whose contents is frozen in time and preserved on the Snapshot Unit. The Source Unit must:

- Be less than 512 GB
- Have write-back cache enabled
- Be non-transportable

Switches

There are no switches associated with this command.

Example

To create unit D4 (Snapshot Unit), which consists of storageset RAID2, and which becomes a point-in-time snapshot of unit D1 (Source Unit), enter:

```
ADD SNAPSHOT_UNITS D4 raid2 D1
```

Glossary

This glossary defines terms pertaining to the Data Replication Manager for the HSG80 running ACS Version 8.5P. It is not a comprehensive glossary of computer terms.

ACS

An acronym for Array Controller software.

See array controller software.

adapter

A hardware device that converts the protocol and hardware interface of one bus type to another without changing the function of the bus.

AL_PA or ALPA

An acronym for Arbitrated Loop Physical Address. A two-digit hexadecimal number that expresses a port's physical position on the loop. ALPA numbers are normally not assigned in sequence (i.e., position 1 is not ALPA 1, and so on). A table in the Fibre Channel Standard equates the loop position to the default ALPA.

arbitrated loop

A Fibre Channel topology. The basic definition is a ring of ports where the transmit output of one port is attached to the receive input of the next. Each port has a unique loop address and it talks to other ports on the loop by arbitrating for loop access. Loop addresses are assigned via cooperative port intercommunication during loop initialization, which occurs any time the device configuration on the loop is physically changed. PLDA (private loop direct attach), the specific profile implemented by the controller, is a subset of arbitrated loop.

See also PL_DA or PLDA.

array controller

See controller.

array controller software Also known by the acronym ACS. ACS is software that is contained on a removable PCMCIA program card that provides the operating system for the array controller.

association sets An association set is a group of remote copy sets that share common attributes. Members of an association set can be configured to transition to the same state at the same time. An association set:

- Shares the same log unit
- Has its host access removed from all members when one member fails
- Keeps I/O order across all members

CLI commands available are ADD ASSOCIATIONS and SET *associations.*

asynchronous mode A mode of operation of the remote copy set whereby the write operation provides command completion to the host after the data is safe on the initiating controller, and prior to the completion of the target command.

Asynchronous mode can provide greater performance and faster response time, but the data on all members at any one point in time cannot be assumed to be identical.

See also synchronous mode.

ATM An acronym for Asynchronous Transfer Mode. This abbreviation refers to a technology used in LANs and WANs to enable disparate traffic (i.e., data, voice, and video) to be carried over the same Local or Wide Area Network. ATM is the transfer mode of choice for broadband integrated services digital networks (BISDNs). ATM traffic carries information in fixed-size cells.

autospare A controller feature that automatically replaces a failed disk drive. The operator can enable the *AUTOSPARE* switch for the failedset, causing physically replaced disk drives to be automatically placed into the spareset. Also called “autonewspare.”

bad block A disk drive data block that contains a physical defect.

bad block replacement A replacement routine that substitutes defect-free disk blocks for those found to have defects. This process takes place in the controller, transparent to the host.

BBR *See* bad block replacement.

block A stream of data stored on disk or tape media and transferred and error-checked as a unit. In a disk drive, a block is also called a sector (the smallest collection of consecutive bytes addressable on a disk drive). In integrated storage elements, a block contains 512 bytes of data, error codes, flags, and the block address header.

cache A fast, temporary storage buffer in a controller or computer.

cache memory A portion of high-speed memory used as an intermediary between a data user and a larger amount of storage. The objective of designing cache into a system is to improve performance by placing the most frequently used data in the highest performance memory.

CBR An acronym for Constant Bit Rate, a category of ATM service. This category supports a constant (guaranteed) data rate. CBR supports applications that require a highly predictable transmission rate.

chunk A block of data written by the host.

See also block, chunk size.

cascaded switch As applied to the Data Replication Manager: the term cascaded switch identifies that the output of a switch is connected to the input of another switch, which then may in turn be connected to another switch or host or controller.

chunk size The number of data blocks, assigned by a system administrator, written to the primary RAIDset or stripeset member before the remaining data blocks are written to the next RAIDset or stripeset member.

CLI An acronym for the Command Line Interpreter. Also known as Command Line Interface. The CLI is the configuration interface to operate the controller software.

clone Utility that duplicates data on any unpartitioned single-disk unit, stripeset, mirrorset, or striped mirrorset in preparation for backup.

command line interface or command line interpreter *See* CLI.

connection

As applied to the Data Replication Manager: this refers to a connection between two end Fibre Channel ports. An example would be the connection between a Host Bus Adapter (by way of the Fibre Channel Switches) and the HSG80 controller.

CLI commands available are ADD CONNECTIONS, SET *connection-name.*

See also link.

container

1. Any entity that is capable of storing data, whether it is a physical device or a group of physical devices.
2. A virtual internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Examples of storageset containers that the controller uses to create units include stripesets and mirrorsets.

controller

A hardware device that uses software to facilitate communications between a host and one or more storage devices organized in an array. The HS-series StorageWorks family of controllers are all array controllers.

copying member

In a mirrorset, a copying member is a container introduced to the mirrorset after it has been in use for some amount of time. None of the blocks can be guaranteed to be the same as other members of the mirrorset. Therefore the COPYING member is made the same by copying all the data from a NORMAL member. This is in contrast to NORMALIZING, where all blocks written since creation are known to be the same.

When all of the blocks on the copying member are the same as those on the normal member, the copying member becomes a normal member. Until it becomes a normal member, the copying member contains undefined data and is not useful for any purpose.

DataSafe

Also known as *firewall*. This pre-tested configuration uses specific hardware, Data Replication Manager software, and installation practices to protect operations from hardware or software outages. The solution includes No Single Point of Failure (NSPOF) functionality.

default gateway

The default path that a computer or router uses to forward and route data between two or more networks having different protocols.

device

See node, peripheral device.

disaster tolerance Disaster tolerance provides the ability for rapid recovery of user data from a remote location when a significant event (or disaster) occurs at the primary computing site.

See also remote copy sets, DT.

DT An acronym for Disaster Tolerance. The Data Replication Manager is an example of a database that is made disaster tolerant.

*See als*o disaster tolerance.

dual-redundant configuration A storage subsystem configuration consisting of two active controllers operating as a single controller. If one controller fails, the other controller assumes control of the failing controller's devices.

See also failover, failback.

ECB An acronym for External Cache Battery. The unit that supplies backup power to the cache module in the event the primary power source fails or is interrupted.

EMU An acronym for Environmental Monitoring Unit. A piece of hardware that provides increased protection against catastrophic failures. Some subsystem enclosures include an EMU, which works with the controller to detect conditions such as failed power supplies, failed blowers, elevated temperatures, and external air sense faults. The EMU also controls certain cabinet hardware, including DOC chips, alarms, and fan speeds.

external cache battery *See* ECB.

F_Port A port in a fabric where an N_Port or NL_Port may attach.

See N_port, NL_port, FL_Port.

fabric A network of switches containing a Fibre Channel arbitrated loop.

failback The process of restoring data access to the newly-restored controller in a dual-redundant controller configuration. The failback method (full copy or fast-failback) is determined by the enabling of the Logging or Failsafe switches, the selected mode of operation (synchronous or asynchronous), and whether the failover is planned or unplanned.

See also failover, dual-redundant configuration.

failedset — A group of disk drives that have been removed from RAIDsets due to a failure or a manual removal. Disk drives in the failedset should be considered defective and should be tested and repaired before being placed back into the spareset or back in their original locations.

failover — The process that takes place when one controller in a dual-redundant configuration assumes the workload of a failed companion controller. Failover continues until the failed controller is repaired or replaced.

The CLI command is: SITE_FAILOVER

See also failback, dual-redundant configuration, planned failover.

failsafe locked — The failsafe error mode can be enabled by the user to fail any I/O whenever the target is inaccessible or the initiator unit fails. When either of these conditions occurs, the remote copy set goes into the inoperative (offline) state and the failsafe error mode is "failsafe locked."

The CLI command SET *remote-copy-set-name ERROR_MODE=FAILSAFE* enables this error mode.

fast-failback — A term representing the synchronization of the initiator site with the target during a planned failover of the initiator subsystem.

The write operations are logged to the target site write history log, and during the fast-failback, the initiator site is updated from the write history log.

See also mini-merge, unplanned failover, planned failover, write history logging.

FC–AL or FCAL — An acronym for Fibre Channel Arbitrated Loop. FC-AL is the overall Fibre Channel topology whose basic definition is a ring of ports where the transmit outputs of one port are attached to the receive input of the next.

FC–ATM — An acronym for Fibre Channel Asynchronous Transfer Mode (ATM AAL5 over Fibre Channel).

FC–FG — An acronym for Fibre Channel Fabric Generic Requirements.

FG–FP — An acronym for Fibre Channel Framing Protocol.

See HIPPI–FC.

FC-GS-1 — An acronym for Fibre Channel Generic Services-1.

FC–GS-2 — An acronym for Fibre Channel Generic Services-2.

FC–IG An acronym for Fibre Channel Implementation Guide.

FC–LE An acronym for Fibre Channel Link Encapsulation (ISO 8802.2).

FC–PH An acronym for the Fibre Channel Physical and Signaling Standard.

FC–SB An acronym for the Fibre Channel Single Byte Command Code Set.

FC–SW An acronym for the Fibre Channel Switched Topology and Switch Controls. This topology involves a structure whose fabric is unknown to the end nodes. The fabric may contain multiple paths between source and destination.

FCC An acronym for the Federal Communications Commission. The federal agency responsible for establishing standards and approving electronic devices within the United States.

FCC Class A This certification label appears on electronic devices that can only be used in a commercial environment within the United States.

FCC Class B This certification label appears on electronic devices that can be used in either a home or a commercial environment within the United States.

FCP An acronym for Fibre Channel Protocol. The mapping of SCSI-3 operations to Fibre Channel.

FDDI An acronym for Fiber Distributed Data Interface. An ANSI standard for 100 megabaud transmission over fiber optic cable.

FD SCSI The fast, narrow, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s.

See FWD SCSI and SCSI.

fiber An optical strand used in fiber optic cable. Spelled ***fibre*** when used in "Fibre Channel" protocol.

See also fiber optic cable, Fibre Channel.

fiber optic cable A transmission medium designed to transmit digital signals in the form of pulses of light. Fiber optic cable is noted for its properties of electrical isolation and resistance to electrostatic contamination.

Fibre Channel An ANSI standard name given to a low-level protocol for a type of serial transmission. The Fibre Channel specifications define the physical link, the low level protocol, and all other pertinent characteristics.

FL_Port A port in a fabric where N_port or an NL_port may be connected.

See N_port, NL_port, F_Port.

See also fabric.

firewall *See* DataSafe.

frame A frame is the basic unit of communication using the Fibre Channel protocol. Each frame consists of a payload encapsulated in control information. The initiator breaks up the exchange into one or more sequences, which in turn are broken into one or more frames. The responder recombines the frames into sequences and exchanges.

See also initiator.

FWD SCSI Acronym for fast, wide, differential (FWD) Small Computer System Interface (SCSI) bus with a 16-bit data transfer rate of up to 20 MB/sec.

See also FD SCSI and SCSI.

GBIC An acronym for Gigabit Interface Converter. The hardware devices inserted into the ports of the Fibre Channel switch that hold the Fibre Channel cables. A GBIC converts fiber optic cable connections to Fibre Channel switch connections.

GLM An acronym for the Gigabit Link Modules used in Fibre Channel long distance applications. As applied to the Data Replication Manager: the GLMs provide the ability to increase the fiber optic cable transmission distances from 10 to 70 Km.

hard address The AL_PA or ALPA which an NL_port attempts to acquire during loop initialization.

heterogeneous host support Also called *noncooperating host support*. The ability to share storage between two similar (or dissimilar) hosts by way of storage partitioning.

HIPPI–FC An acronym for the high-performance parallel interface (HIPPI) over the Fibre Channel. HIPPI is a media-level, point-to-point, 12 channel, full-duplex, electrical/optical interface.

hop One or more connections between two Fibre Channel switches. For example, two switches cascaded are equal to one hop.

ILS Abbreviation for intersite link.

See also multiple intersite links.

initiator

1. A SCSI device that requests an I/O process to be performed by another SCSI device, namely, the SCSI target. The controller is the initiator on the device bus.
2. For subsystems using the disaster tolerance Data Replication Manager solution, the initiator is the site that is the primary source of information. In the event of a system outage, the database would be recovered from the target system.

See also target.

IP address An abbreviation for Internet Protocol Address. The IP address is a number that is used as the address specifying a particular computer connected to the internet.

latency The amount of time required for a transmission to reach its destination.

LBN An acronym for logical block number. A volume-relative address of a block on a mass storage device.The blocks that form the volume are labeled sequentially starting with LBN 0.

L_port A node or fabric port capable of performing arbitrated loop functions and protocols. NL_port and FL_Port are loop-capable ports.

link A connection between two adjacent Fibre Channel ports, consisting of a transmit fiber and a receive fiber. An example would be the connection between the Fibre Channel switch port and the HSG80 controller.

See also connection.

local terminal A terminal plugged into the EIA-423 maintenance port on the front bezel of the HS array controller. Also called a maintenance terminal.

Logical Block Number *See* LBN.

logical unit A physical or virtual device addressable through a target ID number. The logical unit numbers (LUNs) use their target's bus connection to communicate on the SCSI bus.

See LUN.

Logical Unit Number *See* LUN.

LOG_UNIT A CLI command switch that (when enabled) assigns a single, dedicated log unit for a particular association set. The association set members must all be in the NORMAL error mode (not failsafe).

See also write history logging.

long distance mirroring Also known as peer-to-peer remote copy. *See also* remote copy sets.

loop *See also* arbitrated loop.

loop_ID A seven-bit value numbered contiguously from zero to 126-decimal, which represents the 127 legal AL_PA or ALPA values on a loop (not all of the 256 hex values are allowed as AL_PA values per FC-AL).

loop tenancy The period of time between the following two events: when a port wins loop arbitration and when the port returns to a monitoring state.

L_Port A node or fabric port capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.

LUN An acronym for logical unit number. A value that identifies a specific logical unit belonging to a SCSI target ID number. A number associated with a physical device unit during a task's I/O operations. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.

mini-merge As applied to the Data Replication Manager: a term representing the data transfers to be made whenever a target becomes inaccessible. This happens when both links or both target controllers have gone down. The transfers that would have been made are instead logged into the association set's assigned log unit to wait until the remote copy set subsystem comes back online.

See fast-failback, write history logging.

mirroring The act of creating an exact copy or image of data.

mirrorset

1. A group of storage devices organized as duplicate copies of each other. Mirrorsets provide the highest level of data availability at the highest cost. Another name for *RAID 1*. Also called *mirrored units* or *mirrored virtual disks.*
2. Two or more physical disks configured to present one highly reliable virtual unit to the host.
3. A virtual disk drive consisting of multiple physical disk drives, each of which contains a complete and independent copy of the entire virtual disk's data.

multiple intersite links

Each intersite link (ILS) is a fiber link between two switches. As applied to Data Replication Manager: increasing bandwidth between switches is handled by adding additional connections between the switches, to a maximum of two connections.

N_port

A port attached to a node for use with point-to-point topology or fabric topology.

See point-to-point connection.

NL_port

A port attached to a node for use in all three Fibre Channel topologies: point-to-point, arbitrated loop, and switched fabric.

network

In data communication, a configuration in which two or more terminals or devices are connected to enable information transfer.

Non-L_Port

A node or fabric port that is not capable of performing the arbitrated loop functions and protocols. N_Ports and F_Ports are loop-capable ports.

non-participating mode

A mode within an L_Port that inhibits the port from participating in loop activities. L_Ports in this mode continue to retransmit received transmission words but are not permitted to arbitrate or originate frames. An L_Port in non-participating mode may or may not have an AL_PA.

See also participating mode.

non-RCS LUN

As applied to Data Replication Manager: a logical unit number (LUN) value that identifies a physical device unit which exists at the local site and does not have a mirror copy at a remote site.

See also remote copy sets, LUN.

node

1. In data communications, the point at which one or more functional units connect transmission lines.
2. In Fibre Channel, a device that has at least oneN_port or NL_port.

normal member

A mirrorset member that, block-for-block, contains exactly the same data as that on the other members within the mirrorset. Read requests from the host are always satisfied by normal members.

normalizing

A state in which, block-for-block, data written by the host to a mirrorset member is consistent with the data on other normal and normalizing members. The normalizing state exists only after a mirrorset is initialized. Therefore, no customer data is on the mirrorset.

normalizing member

A mirrorset member whose contents are the same as all other normal and normalizing members for data that has been written since the mirrorset was created or since lost cache data was cleared. A normalizing member is created by a normal member when either all of the normal members fail or all of the normal members are removed from the mirrorset. *See also* copying member

OC-3

An acronym for the optical carrier that provides high-speed bandwidth at 155.3 megabits per second.

other controller

The controller in a dual-redundant pair that is not connected to the controller serving your current CLI session with a local terminal.

See also this controller, local terminal.

participating mode

A mode within an L_port that allows the port to participate in loop activities. A port must have a valid AL_PA or ALPA to be in participating mode.

PCM

An acronym for Polycenter Console Manager

PCMCIA

An acronym for Personal Computer Memory Card Industry Association. An international association formed to promote a common standard for PC card-based peripherals to be plugged into notebook computers. A PCMCIA card, sometimes called a PC Card, is about the size of a credit card. It is used in the HSJ80 to load the controller software.

See also program card, ACS.

PCR

An acronym for peak cell rate, the maximum transmission speed of a virtual connection. PCR is a required parameter for the CBR service category.

peer-to-peer remote copy *See* remote copy sets.

peripheral device Any unit, distinct from the CPU and physical memory, that can provide the system with input or accept any output from it. Terminals, printers, tape drives, and disks are peripheral devices.

persistent reserve

planned failover As applied to the Data Replication Manager: an orderly shutdown of the controllers for installation of new hardware, updating the software, and so on. The host applications are quiesced and all write operations permitted to complete before the shutdown. The controllers must be in synchronous operation mode before starting a planned failover.

See also synchronous mode, unplanned failover.

PL_DA or PLDA An acronym for Private Loop Direct Attach. PLDA is a Fibre Channel profile, a proper subset of arbitrated loop. The PLDA profile (part of the Fibre Channel Standard), defines a specific way to implement arbitrated loop topology.

See arbitrated loop.

point-to-point connection A network configuration in which a connection is established between two, and only two, terminal installations. The connection may include switching facilities.

See N_port.

port

- In general terms, a port is:
  1) A logical channel in a communications system.
  2) The hardware and software used to connect a host controller to a communications bus, such as a SCSI bus or serial bus.
- Regarding the controller, the port is:
  1) The logical route for data in and out of a controller that can contain one or more channel, all of which contain the same type of data.
  2) The hardware and software that connects a controller to a SCSI device.

port_name A 64-bit unique identifier assigned to each Fibre Channel port. The Port_Name is communicated during the logon and port discovery process.

preferred address The AL_PA which an NL_Port attempts to acquire first during initialization.

private NL_Port An NL_Port which does not attempt login with the fabric and only communicates with NL_Ports on the same loop.

public NL_Port An NL_Port that attempts login with the fabric and can observe the rules of either public or private loop behavior. A public NL_Port may communicate with both private and public NL_Ports.

program card The PCMCIA card containing the controller's operating software.

See also PCMCIA.

PTL An acronym for Port-Target-LUN. The controller's method of locating a device on the controller device bus:

- P designates the port (1—6)
- T designates the target ID of the device (1—6 in a non-redundant configuration, or 0—5 in a dual-redundant configuration)
- L designates the LUN of the devices (0—7).

PVA module An abbreviation for Power Verification and Addressing module. The Ultra SCSI RAID enclosure assembly whose primary functions are to: (1) allow the user to select the enclosure Ultra SCSI bus ID; (2) enable the user to place the subsystem in a standby condition and return it to an operational status; and (3), in conjunction with the associated EMU, ensures that the major Ultra SCSI elements are functioning properly and notifies the user and the controller of error or fault conditions.

PVC An acronym for Permanent Virtual Circuit. PVC is a logical connection manually defined by the network administrator. The PVC is created by specifying the VPI and VCI.

quiesce To make a bus inactive or dormant. During a device warm swap, the SCSI bus must quiesce.

See also planned failover.

QoS

An acronym for Quality of Service in an ATM network. Each virtual connection in an ATM network is set to a service category. The performance of the connection is measured by the established QoS parameters (outlined by the ATM Forum).

Performance issues include data rate, cell loss rate, cell delay, and delay variation (jitter).

Categories of ATM service are:

- Constant Bit Rate (CBR)
- Variable Bit Rate-Real Time (VBR-RT)
- Variable Bit rate- Non-Real Time (VBR-NRT)
- Available Bit Rate (ABR)
- Unspecified Bit Rate (UBR)

See ATM.

RCS

See remote copy sets.

redundancy

The provision of multiple interchangeable components to perform a single function in order to cope with failures and errors. A RAIDset is considered to be redundant when user data is recorded directly to one member, and all of the other members and associated parity also are recorded. If a member is missing from the RAIDset, its data can be regenerated as needed, but the RAIDset is no longer redundant until the missing member is replaced and reconstructed.

remote copy sets

A feature that allows data to be copied (mirrored) from the originating site (initiator) to a remote site (target). The result is a mirror copy of the data (remote copy set) at two disparate sites. Used in disaster tolerance (DT) applications such as the Data Replication Manager.

CLI commands available are: ADD REMOTE_COPY_SETS, SET *remote-copy-set-name,* SET *controller REMOTE_COPY.*

See also disaster tolerance, non-RCS LUN.

remote copy set metadata Remote copy set metadata describes the remote copy set membership and state. To assist with site failover, this metadata is located in the mirrored write-back cache on the controller where each member resides. Backup copies of the metadata reside in the controller NVRAM at each site. Only the initiator modifies the metadata and ensures all copies are subsequently updated.

replacement policy The policy specified by a CLI command switch (SET FAILEDSET command) indicating whether a failed disk from a mirrorset or RAIDset is to be automatically replaced with a disk from the spareset. The two switch choices are *AUTOSPARE* and *NOAUTOSPARE*.

SCSI An acronym for Small Computer System Interface:

1. An American National Standards Institute (ANSI) interface standard defining the physical and electrical parameters of a parallel I/O bus used to connect initiators to devices.
2. A processor-independent standard protocol for system-level interfacing between a computer and intelligent devices, including hard drives, floppy disks, CD-ROMs, printers, scanners, and others.

SCSI device

1. A host computer adapter, a peripheral controller, or an intelligent peripheral that can be attached to the SCSI bus.
2. Any physical unit that can communicate on a SCSI bus.

SCSI device ID number A bit-significant representation of the SCSI address referring to one of the signal lines, numbered 0 through 7 for an 8-bit bus, or 0 through 15 for a 16-bit bus.

SCSI ID number The representation of the SCSI address that refers to one of the signal lines numbered 0 through 15.

snapshot A snapshot unit is one that reflects the contents of another unit at a particular point in time.

See also unit.

storage array An integrated set of storage devices. Storage arrays can be manipulated as one unit, with a single command.

storage unit The general term that refers to storagesets, single-disk units, and all other storage devices that are installed in a subsystem and accessed by the host. A storage unit can be any entity that is capable of storing data, whether it is a physical device or a group of physical devices.

storageset

1. A group of devices configured with RAID techniques to operate as a single container.
2. Any collection of containers, such as stripesets, mirrorsets, striped mirrorsets, JBODs, and RAIDsets.

subnet mask Also known as *address mask.* A subnet is an IP network that can be reached through a single IP address. All the members of the subnet share the mask value. Members of the subnet can then be referenced more easily. A subnetwork is a network that is part of another network, connected through a gateway, bridge, or router.

surviving controller The controller in a dual-redundant configuration pair that serves its companion's devices when the companion controller fails.

SWCC An acronym for Storage Works Command Console.

synchronous mode A mode of operation of the remote copy set whereby the data is written simultaneously to the cache of the initiator subsystem and the cache of the target subsystem. The I/O completion status is not sent until all members of the remote copy set are updated.

See also asynchronous mode.

target

1. A SCSI device that performs an operation requested by another SCSI device, namely the SCSI initiator. The target number is determined by the device's address on its SCSI bus.

   For subsystems using the disaster-tolerant Data Replication Manager solution, data processing occurs at the initiator site and the data is replicated or mirrored to the target site. In the event of a system outage, the database would be recovered from the target system.

See also initiator.

this controller The controller that is serving the current CLI session through a local or remote terminal.

See also other controller.

UBR An acronym for unspecified bit rate. The UBR is a category of ATM service that supports connections that have no specified performance requirements.

ULP An acronym for Upper Layer Protocol.

ULP process A function executing within a Fibre Channel node which conforms to the Upper Layer Protocol (ULP) requirements when interacting with other ULP processes.

UltraNet Wizard Another term for the Fibre Channel-to-ATM Configuration Wizard. This wizard is an UltraNet application that allows the designation of the default configuration settings for Fibre-Channel-ATM on the Open Systems Gateway.

unit A container made accessible to a host. A unit may be created from a single disk drive or tape drive. A unit may also be created from a more complex container, such as a RAIDset. The controller supports a maximum of eight units on each target.

unplanned failover As applied to the Data Replication Manager: An unplanned outage of the controllers. This may occur when the site communication is lost, or due to some other failure whereby remote copy sets cannot be implemented. The controllers do not perform an orderly shutdown.

See also planned failover.

VCI An acronym for virtual channel identifier. The VCI is the field of the cell header that stores the virtual channel address.

VPI An acronym for virtual path identifier. The VCI is the field of the cell header that stores the virtual path address.

World Wide Name or World wide ID Also known by the acronym WWN. A unique 64-bit number assigned to a subsystem by the Institute of Electrical and Electronics Engineers (IEEE) and set by manufacturing prior to shipping. This name is referred to as the node ID within the CLI.

write history logging As applied to the Data Replication Manager: The use of a log unit to log a history of write commands and data from the host. Write history logging is used for mini-merge and fast- failback.

See mini-merge, fast-failback.

WTI Switch An abbreviation for the Western Telematics Switch that must be installed to set up and service the ATM gateway. The WTI switch is a 16-port serial switch that is used to configure or service the OSG unit locally or remotely.

zone A set of devices that access one another. All devices connected to a fabric may be configured into one or more zones. Devices that are in the same zone can see each other; devices that are in different zones cannot.

zone alias Zone aliases simplify the entry of repetitive port numbers or World Wide Names. A zone alias is a C-style name for one or more port numbers or World Wide Names (e.g., the named host could be used as an alias for 10:00:00:60:69:00:00:8a).

zone configuration A set of zones. At any one time zoning may be disabled or one zone configuration may be in effect. When a zone configuration is in effect, all zones that are members of that configuration are in effect. You select which zone configuration is currently in effect.

zoning As applied to the Data Replication Manager: an optionally licensed feature of the SilkWorm switch that allows a finer segmentation of Storage Area Networks (SANs) by allowing ports or WWN addresses to be used to confine access to devices that are in a common zone.

Index

D

H

I

L

M

N

O

P

R

S