

Compaq StorageWorks

Data Replication Manager HSG80 ACS Version 8.5P Operations Guide

EK-DRVMS-TE. A01

Fourth Edition (January 2001)
Compaq Computer Corporation

© 2001 Compaq Computer Corporation.

Compaq, the Compaq logo, and StorageWorks Registered in U. S. Patent and Trademark Office.

SANworks, Tru64 UNIX, and OpenVMS are trademarks of Compaq Information Technologies Group, L.P.

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

Intel, Pentium, Intel Inside, and Celeron are trademarks of Intel Corporation in the United States and other countries.

Motif, OSF/1, UNIX, the "X" device, IT DialTone, and The Open Group are trademarks of The Open Group in the United States and other countries.

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

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

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

Compaq service tool software, including associated documentation, is the property of and contains confidential technology of Compaq Computer Corporation. Service customer is hereby licensed to use the software only for activities directly relating to the delivery of, and only during the term of, the applicable services delivered by Compaq or its authorized service provider. Customer may not modify or reverse engineer, remove, or transfer the software or make the software or any resultant diagnosis or system management data available to other parties without Compaq's or its authorized service provider's consent. Upon termination of the services, customer will, at Compaq's or its service provider's option, destroy or return the software and associated documentation in its possession.

Printed in the U.S.A.

Data Replication Manager HSG80 ACS V8.5P Operations Guide
Fourth Edition January 2001
Part Number EK-DRVMS-TE. A01

Contents

About This Guide

Getting Help	xiii
Compaq Website	xiii
Compaq Technical Support	xiii
Precautions	xiv
Electrostatic Discharge Precautions	xiv
Component Precautions	xiv
Conventions	xv
Special Notices	xv
Text Conventions	xvi
Related Publications	xvii
Data Replication Manager Solution Kits for Windows NT/Intel and OpenVMS	xviii

Chapter 1

Introduction to Data Replication Manager

Data Replication Manager Overview	1-1
Peer-to-Peer Remote Copy Function	1-2
Hardware Redundancy	1-2
Hardware Components	1-2
ESA12000 Cabinet	1-3
Fibre Channel Gigabit Switch	1-6
Gigabit Interface Converters (GBIC)	1-6
Power Distribution Unit (PDU)	1-6
Fully-Redundant Power—Optional	1-7
Host Bus Adapter	1-7

Final Assembly	1-8
Software Components	1-9
Array Controller Software	1-9
Secure Path	1-9
StorageWorks Command Console (optional)	1-9
Required Hardware and Software	1-10

Chapter 2

Data Replication Manager Concepts

Remote Copy	2-2
Remote Copy Sets	2-2
Operation Modes	2-3
Synchronous Operation Mode	2-3
Asynchronous Operation Mode	2-3
Operation Mode Considerations	2-4
Outstanding_IO Settings	2-4
Synchronous	2-4
Outstanding Write Operations	2-5
High Outstanding I/O Values	2-5
Low Outstanding I/O Values	2-5
Suspend/Resume	2-5
Error Mode	2-6
Association Sets	2-6
Association Set characteristics	2-6
FAIL_ALL	2-7
Write History Logging	2-8
Mini-merge	2-8
Fast-Failback	2-8
Log Unit Restrictions	2-8
Log Unit	2-9
ORDER_ALL	2-9
Failover	2-10
Planned Failover	2-10
Unplanned Failover	2-10
Failback	2-10

Chapter 3

Getting Started

Site, Host, and Solution Preparation	3-1
Host Bus Adapter Requirements	3-2
Setting Up the Fibre Channel Switches	3-2

Setting Up the Fiber Optic Cables	3-3
Host-to-Switch Connections	3-4
Switch-to-Controller Connections.	3-4
Chapter 4	
Configuring a Data Replication Manager Solution	
Introduction	4-2
Restrictions.	4-3
Configuring Overview.	4-4
Configuration Procedures Outline	4-4
TARGET SITE	4-4
INITIATOR SITE	4-5
Configure the Controllers at the Target Site	4-6
Example Display 1	4-7
Example Display 2	4-8
Example Display 3	4-9
Example Display 4	4-10
Example Display 5	4-11
Example Display 6	4-12
Example Display 7	4-12
Example Display 8	4-13
Example Display 9	4-14
Configure Storage at the Target Site	4-15
Devices and StorageSets.	4-15
Example Display 10	4-16
Connect Fiber Optic Cables Between the Controllers and Fiber Channel Switches	4-17
Connect the Target Site to the External Fiber Link	4-18
Long Wave GBICs	4-18
Other Transport Modes	4-18
Configure the Host from the Target Site	4-19
Install Secure Path (NT only).	4-20
Install SWCC (optional).	4-20
Connect Fiber Optic Cables Between the Hosts and the Switches	4-20
Fiber Optic Cable Connection Procedure	4-21
Example Display 11	4-22
Example Display 12	4-23
Rename the Host Connections	4-23
Example Display 13	4-24
Configure the Controllers at the Initiator Site	4-25
Controller Pre-Configuration Procedure	4-25
Controller Configuration Procedure.	4-25

Example Display 14	4-26
Example Display 15	4-27
Example Display 16	4-28
Example Display 17	4-29
Example Display 18	4-30
Example Display 19	4-31
Example Display 20	4-31
Example Display 21	4-32
Example Display 22	4-33
Configure the Storage at the Initiator Site.	4-33
Devices and StorageSets	4-33
Units	4-33
Example Display 23	4-34
Cabling the Initiator Site	4-35
Connect Fiber Optic Cables Between the Controllers and Switches	4-35
Connect the Initiator Site to the External Fiber Link	4-36
Long Wave GBICs	4-37
Other Transport Modes	4-37
Create Remote Copy Sets	4-38
Initiator Site Preparation	4-38
Create Connections From the Target Site	4-38
Create RCS from the Initiator Site	4-39
Example Display 24	4-39
Set Failsafe at the Initiator Site (optional)	4-40
Creating Log Units and Association Sets (optional)	4-41
Creating a Log Unit	4-41
Example Display 25	4-41
Example Display 26	4-42
Creating Association Sets and Assigning a Log Unit	4-42
Example Display 27	4-43
Configure the Host from the Initiator Site	4-44
Install the Host Bus Adapters and Drivers	4-44
Example Display 28	4-44
Install Secure Path (NT only)	4-45
Install SWCC (optional)	4-45
Connecting Fibre Optic Cables Between the Hosts and the Switches	4-45
Example Display 29	4-47
Rename the Host Connections from the Initiator Site	4-48
Example Display 30	4-49
Enable Access to the Hosts at the Initiator Site	4-50
Install Cluster Server for Windows NT (optional)	4-51
Documenting Your Configuration	4-51

Terminal Emulator Session	4-51
SHOW Commands	4-51
Example Display 31	4-52
Example Display 32	4-53
Example Display 33	4-54
Example Display 34	4-55
Example Display 353	4-55

Chapter 5

Managing Site Failover and Failback Procedures

Power Up the Data Replication Manager Systems	5-2
Target Site Powerup Procedures	5-2
Initiator Site Powerup Procedures	5-2
Power Down Data Replication Manager Systems	5-3
Initiator Site Power Down Procedures	5-3
Target Site Power Down Procedures	5-3
Site Failover Basic Description	5-4
Site Failback Basic Description	5-5
Data Replication Manager Configuration Basics	5-6
Planning Considerations	5-7
Planned Failover Procedures	5-8
Initiator Site Preparation Procedure	5-9
Target Site Failover Procedure	5-10
Target Host Setup Procedure	5-11
Simple Failback Procedure	5-12
Initiator Site Failback Preparation Procedure	5-13
Target Site Failback Procedure	5-13
Initiator Site Clean Up Procedure	5-15
Unplanned Failover	5-16
Target Site Failover Procedures	5-16
Full Failback Procedure	5-18
Initiator Site Preparation Procedure	5-18
Target Site Failback Procedure	5-20
Initiator Site Clean-Up Procedure	5-21
Failing Back New Remote Copy Sets	5-23
Target Site: Add Remote Copy Set Targets	5-23
Initiator Site: Failover Procedure	5-24
Target Site: Restart Controllers Procedure	5-24
New Hardware Failback Procedure	5-25
Initiator Site Preparation Procedures	5-25
Target Site Preparation Procedure	5-27

Initiator Site Connections Procedure	5-28
Target Site Copy Data Procedure.	5-28
Initiator Site Return Control Procedure	5-29
Target Site Clean-Up Procedure	5-30
Initiator Site Restoration of Target Connections	5-30
Initiator Site Clean-Up Procedure	5-30

Chapter 6

Troubleshooting

HSG80 Array Controller Operating Characteristics	6-2
Forced Errors Detected During Copy	6-2
Read Errors Detected During Full Copy	6-2
Dual Redundancy During Failback	6-3
Failsafe Lock Management	6-3
Link Failure Management	6-3
Remote Copy Set Member Failures	6-3
Remote Copy Set Worldwide LUN ID	6-4
Write History Logging	6-4
Failure Notification	6-5
HSG80 Array Controller Failure	6-5
SWCC Failure	6-6
Failure of One Member in a Dual Redundant Controller Pair	6-6
Failure of Both Fiber Optic Cables or Switch	6-6
Failure Modes of a DT System in Normal Operation	6-7
Failure at Target Site after Failover	6-9

Appendix A

Status Comparison

Target Site Terminal Emulator Session	A-1
SHOW Commands	A-2
Example Display 1	A-3
Example Display 2	A-4
Example Display 3	A-4
Example Display 4	A-5
Example Display 5	A-5

Glossary

Index

Figures

Figure 1–1. ESA12000 Storage building block	1–4
Figure 1–2. Additional components for Data Replication Manager.	1–5
Figure 1–3. Fibre channel gigabit switch	1–6
Figure 1–4. Fibre channel-based DT storage subsystem (with fully-redundant power)	1–8
Figure 2–1. Remote copy set operation modes.	2–3
Figure 2–2. Association sets reside on the initiator controller pair	2–7
Figure 3–1. Component locations and names.	3–3
Figure 4–1. Data Replication Manager basic configuration.	4–2
Figure 4–2. Switch Port Locations	4–17
Figure 4–3. Cabling between the controllers and switches	4–18
Figure 4–4. Cabling from the target site to the initiator site.	4–19
Figure 4–5. Cabling between the host and the switches.	4–22
Figure 4–6. Host renaming worksheet	4–24
Figure 4–7. Port Locations	4–35
Figure 4–8. Cabling between the controllers and switches	4–36
Figure 4–9. Cabling from the initiator site to the target site.	4–37
Figure 4–10. Cabling between the hosts and the switches	4–46
Figure 4–11. Data Replication Manager cabling at initiator and target sites	4–46
Figure 4–12. Host renaming worksheet	4–48
Figure 5–1. Data Replication Manager basic configuration.	5–6

Tables

Table 1-1 Hardware Requirements Checklist and Part Numbers	1-10
Table 1-2 Software Requirements Checklist and Part Numbers	1-11
Table 2-1 Data Replication Manager Switch Settings	2-11
Table 4-1 Restrictions	4-3
Table 5-1 Failover Scenarios	5-4
Table 5-2 DRM Planning Considerations	5-7
Table 6-1 Failure Modes of a DT System with Normal Operation	6-7
Table 6-2 Target Site DT Failure Modes After Failover	6-9

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/products/storageworks>

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/products/storageworks/>

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, 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 to 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

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

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: A *Warning* contains information essential to people's safety. It advises users that failure to take or avoid a specific action could result in physical harm to the user or hardware. Use a warning, not a caution, when such damage is possible.



CAUTION: A *Caution* contains information that the user needs to know to avoid damaging the software or hardware.

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

NOTE: A *note* indicates neutral or positive information that emphasizes or supplements important points of the main text. A note supplies information that may apply only in special cases—for example, memory limitations, equipment configurations, or details that apply to specific versions of a program.

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 UPPER-CASED MONOSPACED FONT.
<i>Sans serif italic</i> <i>Sans serif italic</i>	Command variable or numeric value that you supply, for example: SHOW RAIDset-name or SET THIS_CONTROLLER PORT_1_SCS_NODENAME=xxxxxx
<i>Serif italic</i>	References to other books, for example: “See the <i>Compaq StorageWorks HSJ80 Array Controller Configuration Guide</i> for details.”
.	Indicates that a portion of an example or figure has been omitted.
.	
.	
“this controller”	The controller serving your current CLI session through a local or remote terminal.
“other controller”	The controller in a dual-redundant pair that’s connected to the controller serving your current CLI session.

Related Publications

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

Document Title	Part Number
HSG80 Array Controller ACS V8.5 Configuration Guide	165144-001 / EK-HSG85-CG
HSG80 Array Controller ACS V8.5 CLI Reference Guide	165145-001 / EK-HSG85-RG
HSG80 Array Controller ACS V8.4/8.5 Maintenance and Service Guide	118629-002 / EK-HSG84-SV
StorageWorks Fibre Channel Storage Switch Service Guide	135268-001 / AA-RHBZA-TE
StorageWorks Fibre Channel Storage Switch User's Guide	135267-001 / AA-RHBYA-TE
Compaq StorageWorks RA8000 and ESA12000 Storage Subsystems User's Guide	387404-001 / EK-SMCPR-UG
Compaq StorageWorks RA8000 and ESA12000 Fibre Channel Cluster Solutions for Windows NT Installation Guide	101471-003 / EK-NTC8K-IG
RA8000 and ESA12000 HSG80 Solution Software V8.4/V8.5 for WindowsNT Server - Intel Installation Reference Guide	387387-003 / AA-RFA9B-TE
RA8000 and ESA12000 HSG80 Solution Software V8.4/V8.5 for Open VMS - Installation Reference Guide	387401-002 / AARH4BB-TE
StorageWorks Secure Path for Windows NT, A High Availability MultiPath Solution Installation Guide	123995-002 / EK-WNTMP-MH
KGPSA PCI-to-Fibre Channel Host Adapter	EK-KGPSA-UG
Compaq StorageWorks Ultra SCSI RAID Enclosure (DS-BA370-Series) User's Guide	387403-001 / EK-BA370-UG
Command Console Version 2.2(HSG80) for RA8000/ESA12000 User's Guide	387405-003 / AA-RFA2D-TE

Data Replication Manager Solution Kits for Windows NT/Intel and OpenVMS

The part number required to order a Data Replication Manager (DRM) Solution kit depends upon which operating system the DRM will operate on:

- Windows NT Solution Kit Part Number = QB-6BTAE
- OpenVMS Solution Kit Part Number = QB-6BTAC

The following components should be included in the DRM Solution Kit:

- *Compaq StorageWorks Data Replication Manager HSG80 Array Controller ACS Version 8.5P Operations Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Configuration Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*
- *Command Console V2.2 (HSG80) for RAID Array 8000/ESA 12000 User's Guide*
- *RA8000/ESA12000 HSG80 Solution Software V8.5 for Windows NT - Intel Installation Reference Guide*
- *RA8000/ESA12000 HSG80 Solution Software V8.5 for OpenVMS Installation Reference Guide*
- *RAID Array 8000/ESA12000 Fibre Channel Cluster Solutions for WindowsNT Installation Guide*
- *RAID Array 8000/ESA12000 Fibre Channel Cluster Solutions for OpenVMS Installation Guide*
- Compaq StorageWorks Warranty Terms and Conditions
- Compaq StorageWorks Customer Letter
- Compaq License Agreement
- HSG80 Product Registration Card

Revision History

<u>Revision Level</u>	<u>Date</u>
Revision D EK-DRVMS-TE. A01	January 2001
Revision C EK-HSG84-DT.C01/128519-003	December 1999
Revision B EK-HSG84-DT.B01/128519-002	October 1999
Revision A EK-HSG84-DT.A01/128519-001	April 1999

Chapter 1

Introduction to Data Replication Manager

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

This chapter contains the following topics:

- “Data Replication Manager Overview” on page 1–1
 - “Peer-to-Peer Remote Copy Function” on page 1–2
 - “Hardware Redundancy” on page 1–2
 - “Hardware Components” on page 1–2
 - “Final Assembly” on page 1–8
 - “Software Components” on page 1–9
- “Required Hardware and Software” on page 1–10

Data Replication Manager Overview

Data Replication Manager (DRM) provides a disaster-tolerant solution through the use of hardware redundancy and data replication across multiple sites separated by some distance.

A Data Replication Manager configuration consists of multiple sites. The initiator site is where primary data processing occurs. The target site is used for data replication. Since data processing occurs at the initiator site, the data is replicated or mirrored to the target site. If a significant failure happens to the initiator site, then data processing can be resumed at the target site 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. Other hardware may be inserted between the fibre channel switches if longer distances exist between 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 will failover 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 houses one or more ESA12000 cabinets that are equipped with one or more BA370 enclosures and disk Storage Building Blocks (SBBs). Each BA370 enclosure holds 24 disks.

NOTE: Data Replication Manager can run with cabinets that have 24, 48, or 72 (requires two cabinets) disk drives, but the initiator and target sites must be equipped with the identical number of disks.

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. See the *RA8000 and ESA12000 Storage Subsystems User's Guide* for complete details on this equipment.

NOTE: While this documentation addresses ESA12000 storage cabinets as the primary unit for Data Replication Manager configurations, Compaq's DT solution will function in any equivalent cabinet that houses a BA370 enclosure.

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.


ESA12000 Cabinet

The ESA12000 Storage Building Block Cabinet houses the BA370 Enclosures, which contain the following 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 (SBB) per BA370 Enclosure
- Five to Eight 180-watt Power Supplies
- External Cache Battery (ECB), dual
- Eight Cooling Fans
- Six single-ended I/O Modules
- One Power Verification and Addressing module (PVA)
- Two cache modules (**512 MB required**)

For detailed information about these components, refer to the *HSG80 Array Controller ACS Version 8.5 Configuration Guide* and *HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* and the *HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*.

Figure 1-1 shows these parts inside the ESA12000 cabinet with a 24 disk drive capacity. Subsequent sections outline additional hardware requirements needed to complete the Data Replication Manager solution.



CXO6843A

Figure 1-1. ESA12000 storage building block

Figure 1-2 highlights the specific components that must be added to the ESA 12000 building block to support a Data Replication Manager solution.




Figure 1-2. Additional components for Data Replication Manager

The following sections outline the components that are specific to the disaster tolerant solution.

Fibre Channel Gigabit Switch

The Fibre Channel gigabit switch, shown in Figure 1-3, is used to connect the controllers to the hosts and to link the initiator and target sites together. The ports hold short- or long-wave Gigabit Interface Converters, which are described in the next section. See the *StorageWorks Fibre Channel Storage Switch User's Guide* for an in-depth look at the features and functions of the Fibre Channel gigabit switch.




Figure 1-3. Fibre channel gigabit switch

Gigabit Interface Converters (GBIC)

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 together. Long-wave GBICs connect switches that are up to 10 kilometers apart. See the *StorageWorks Fibre Channel Storage Switch User's Guide* to learn more about GBICs.

Power Distribution Unit (PDU)

The PDU is another component that is included with the ESA12000 cabinet and is used to distribute power to the BA370s and switches. A second PDU can be ordered to support a fully-redundant power configuration. See the *RA8000 and ESA12000 Storage Subsystems User's Guide* for more detailed information.

Fully-Redundant Power—Optional


Fully-redundant power is an optional feature designed to offer a more secure source of power in the event that one or more units should fail. If less than five power components are operational, then the entire cabinet will shut down. This requires three additional power supplies and one additional AC power controller that plugs into one additional PDU. These additional components must be supplied for each BA370 enclosure. See the *RA8000 and ESA12000 Storage Subsystems User's Guide* for more details about power supply Storage Building Blocks (SSBs).

Host Bus Adapter

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 the *KGPSA PCI-to-Fibre Channel Host Adapter* guide for more information.

Final Assembly

Your final DT setup should reflect Figure 1-4.



CXO6842A

Figure 1-4. Fibre channel-based DT storage subsystem (with fully-redundant power)

NOTE: If you prefer to join cabinets for more storage capacity, follow the instructions in the RA8000 and ESA12000 Storage Subsystems User's Guide, and be sure to establish the same setup at both the initiator and target sites. Keep in mind that an additional cabinet will not include switches or controllers. It will, however, hold a PDU and be able to support redundant power.

Software Components

This section describes the software components necessary to configure and manage a DT storage subsystem. For installation instructions, see the Configuration chapter. The following list shows the software required to enable Data Replication Manager:

- Array Controller Software (ACS) Version 8.5P
- Secure Path Version 2.2 (Required for Windows NT)
- Storage Works Command Console (SWCC) Version 2.2 (optional)

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.

Secure Path

Secure Path is server-based software that enhances the StorageWorks RAID dual-ported storage subsystem by providing automatic error recovery from server-to-storage subsystem connection failures. Secure Path allows you to add redundant Fibre Channel paths between Windows NT hosts and a RAID storage subsystem, improving overall data availability. If any component in the path between the host and storage subsystem fails, Secure Path immediately redirects all pending and subsequent I/O requests from the failed path to an alternate path, preventing an adapter, cable, or controller failure from disrupting data access.

For more information on Secure Path, refer to the *SecurePath for Windows NT Installation Guide*.

StorageWorks Command Console (optional)

StorageWorks Command Console (SWCC) provides local and remote management of StorageWorks controllers and their attached storage. 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 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 *StorageWorks Command Console Getting Started Guide*.

Required Hardware and Software

Table 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 1-1 Hardware Requirements Checklist and Part Numbers

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
Fibre channel gigabit switch		
8-port	380591-B21 /	
16-port	DS-DSGGA-AA 380578-B21 / DS-DSGGA-AB	2
Host bus adapter	380574-001/ KGPSA-BC	2 (per system)
Fibre Channel host cables		
5 meters	234457-B22 / BNGBX-05	
15 meters	234457-B23 / BNGBX-15	
30 meters	234457-B24 / BNGBX-30	2
50 meters	234457-B25 / BNGBX-50	
Fibre Channel controller cables (2 meters)	234457-B21 / DS-BNGBX-02	4

Table 1-1 Hardware Requirements Checklist and Part Numbers (Continued)

Required Hardware	Part Number	Quantity (at each site)
Short-wave Gigabit Interface Converters (GBICs)	380561-B21 / DS-DXGGA-SA	6
ATM Gateway Additional Hardware	Part Number	Quantity
ATM Gateway	166296-B21	2 per site
ATM Gateway Service Kit	166297-B21	1 per site

Table 1-2 Software Requirements Checklist and Part Numbers

Required Software	Part Number
ACS V8.5P	128698-B21 / QB-6CAAA-SA
Secure Path 2.2.2	380594-001 / QB-6695A-AA
SWCC Version 2.2 (optional)	N/A

Chapter 2

Data Replication Manager Concepts

This chapter describes Data Replication Manager concepts for configuring a Data Replication Manager solution. These descriptions include Remote Copy Sets and Association Sets.

The topics contained within this chapter are:

- “Remote Copy” on page 2-2
 - “Remote Copy Sets” on page 2-2
 - “Operation Modes” on page 2-3
 - “Outstanding_IO Settings” on page 2-4
 - “Suspend/Resume” on page 2-5
 - “Error Mode” on page 2-6
- “Association Sets” on page 2-6
 - “FAIL_ALL” on page 2-7
 - “Write History Logging” on page 2-8
 - “Log Unit” on page 2-9
 - “ORDER_ALL” on page 2-9
 - “Failover” on page 2-10
 - “Failback” on page 2-10

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 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, storage set, mirror set or RAID set. 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 so as 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.

Operation Modes

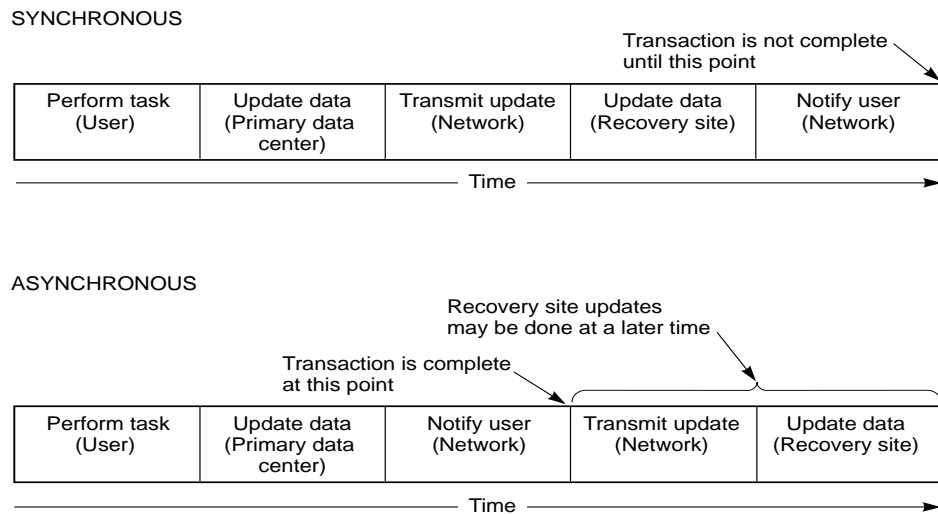
There are two possible remote copy operation modes: *Synchronous* or *Asynchronous*. Figure 2-1 shows the timeline differences between the two.

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 greater 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.



CXO7070A

Figure 2-1. Remote copy set operation modes

Operation Mode Considerations

- 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 a value of one 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.

If `OUTSTANDING_IO` is set to a value of one and the host issues four writes to a remote copy set, then the first write is handled asynchronously and the three remaining writes are handled synchronously. Once the write completes at the target, the next write operation is removed from the controller write queue and handled asynchronously.

Outstanding Write Operations

Users should be aware that there is a controller-wide limit of 240 outstanding write operations even if the total 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, then 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

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 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.

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

- Up to 12 remote copy sets as members
- Synchronous or asynchronous operation mode and members may be set differently
- If ORDER_ALL is set, requires *in order* execution of commands across the remote copy sets in the association set.
- Can share a log unit
- If FAIL_ALL is set, all members of the association set assume the failsafe locked condition if one member assumes the failsafe locked condition.
- Reside on the initiator controller pair




Figure 2-2. Association sets reside 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 will be using the write history log unit.

The `ADD ASSOCIATIONS AssociationSetName RemoteCopySetName` command 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 at the target (failover) site, using the attributes that were set at the initiator site.

FAIL_ALL

If the `FAIL_ALL` switch is enabled and one member of the association set assumes the failsafe locked condition, then 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 is a term to describe using 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*.

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 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

- 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.
- Log units must be either a mirrorset or striped mirrorset.
- Access must be disabled.
- Write-back caching must be disabled.
- Other unit settings must be the default settings.

- Log unit must reside at the current initiator site.
- Upon site failover the log unit and association set must be reconfigured.
- Log unit cannot be a partitioned unit.
- Log unit is a fixed size.

Choose the size of the log unit carefully. When the end of the log unit is encountered, a fully 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 this 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 the 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. Care should be taken in specifying 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: You may incur a full copy 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 some 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 will be 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.

Failback

The failback method, full copy or fast-failback, will be 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. The table also shows the availability of 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 is a good idea to keep a copy of this manual at both the initiator and target sites, so as to ensure a successful and identical setup at both sites. Two copies will also eliminate confusion if more than one person is configuring Data Replication Manager.

The topics contained within the Site, Host, and Solution Preparation section are:

- “Host Bus Adapter Requirements” on page 3–2
- “Setting Up the Fibre Channel Switches” on page 3–2
- “Setting Up the Fiber Optic Cables” on page 3–3
 - “Host-to-Switch Connections” on page 3–4
 - “Switch-to-Controller Connections” on page 3–4

Site, Host, and Solution Preparation

Before you start operating your DT subsystem, you will need to ensure that you have enough clearance to install and store the subsystem(s) and have adequate power resources. If you choose to use more than one cabinet, you will need to understand the proper methods for positioning and joining them. In addition, you will need to have the proper devices installed and verify that all of the BA370 components are in place.

To learn more about adding additional storage, refer to the *RA8000 and ESA12000 Storage Subsystems User's 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 *KGPSA PCI-to-Fibre Channel Host Adapter* guide that came with your adapter for detailed information on this hardware.

At this time, it is important to locate and record the worldwide 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 the Configuration chapter. The initiator site host bus adapter worldwide names can be recorded in the worksheet found in the Configuration chapter. You will need to have this number handy when you rename the host connections in the Configuration chapter.

NOTE: The worldwide 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. Worldwide name example: 1000-0000-C920-A5BA

Setting Up the Fibre Channel Switches

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

- Power cord
- 10BASE-T cable with RJ45 plug (to be connected to a low-cost Ethernet hub or switch)
- Fixed IP address and subnetmask (one of each per switch)

The Ethernet cable and IP address are required to monitor and administer the Fibre Channel switch. You will have to configure the Ethernet IP address and the Ethernet IP subnet mask with the front panel buttons of the Fibre Channel switch. See the *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 will verify the operation of the name resolution.

4. *Telnet* into the switch. Username = **admin** and password = **password** (default setting). Make the following adjustments to the switch:
 - ❑ Type **switchName** to configure the switch name. Be sure to designate a name that will enable you to easily identify the switch that you are trying to access.
 - ❑ 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 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. Figure 3-1 shows what each component will be referred to in this document:



CX07096A

Figure 3-1. Component locations and names

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

- Rank number or PCI slot number of the host bus adapter
- 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 and 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 will 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, it is recommended that multiple conductor cable be used.

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

<http://spdinline.shr.dec.com/pg/ecg/spd/prodcenter.asp?ProdCode=56&RefreshCode=1&QLID=3#here>



CAUTION: If the Fibre Channel optical cable is not properly connected to the controller, failure may result. Because of the cable's frail nature, it must also be regularly maintained, or its performance and life span will be affected. Before proceeding, it is important to administer the precautionary measures detailed in the *HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*.

The tables (below) provide an overview of the connections that you will need to make at and between each site. Specific connection information is covered in more detail in the Configuration chapter. Do not make any connections at this time.

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	← Top switch, Port 6

Chapter 4

Configuring a Data Replication Manager Solution

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

The procedures take you through the configuration process. You will first set up the target site, then the initiator site. Setup for each site is very similar. At each site, you will configure the controllers by defining controller characteristics specific to Data Replication Manager. Then storage sets, units, remote copy sets, and association sets are defined. Once the controllers are configured, you will make fiber optic cable connections between the controllers and switches. Finally, the necessary software and drivers are installed on each host.

This chapter contains the following topics:

- “Introduction” on page 4–2
- “Configuring Overview” on page 4–4
- “Configure the Controllers at the Target Site” on page 4–6
- “Configure Storage at the Target Site” on page 4–15
- “Configure the Host from the Target Site” on page 4–19
- “Configure the Controllers at the Initiator Site” on page 4–25
- “Configure the Storage at the Initiator Site” on page 4–33
- “Cabling the Initiator Site” on page 4–35
- “Create Remote Copy Sets” on page 4–38
- “Creating Log Units and Association Sets (optional)” on page 4–41

- “Configure the Host from the Initiator Site” on page 4-44
- “Rename the Host Connections from the Initiator Site” on page 4-49
- “Enable Access to the Hosts at the Initiator Site” on page 4-51
- “Install Cluster Server for Windows NT (optional)” on page 4-52
- “Documenting Your Configuration” on page 4-52
- “Documenting Your Configuration” on page 4-52

Introduction

The 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 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 and will be 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	Implication
Two HSG80 controller pairs are required.	All controllers must run ACS V8.5P.
Four Fibre Channel Switches are required.	These switches provide two separate fabrics connecting controllers at the initiator and target sites.
HSG80 controller(s) must be configured in Multiple-Bus Failover mode.	Additional software support required on the host: For Windows NT, Secure Path V2.2.2 For OpenVMS V7.2-1, Fibre Channel TIMA kit
HSG80 controller(s) must be configured for Fibre Channel Switched protocol.	Host operating system and adapter must support Fibre Channel Switched protocol as well.
Mirrored write-back cache must be enabled.	512 MB cache per controller (256 MB effective capacity once mirrored)
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 (1 on initiator; 1 on target)	Composed of 1 initiator and one 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.	<ul style="list-style-type: none"> ■ Link between initiator and target site is made through Port 2 ■ Both links must be up when configured
Maximum 64 connections	Effective number of connections is 64 minus the 4 default remote copy connections.
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, etc. the same to eliminate confusion and error risk.
Controller-based partitions are not supported within Remote Copy Sets.	Host software may be capable of partitioning units.
Unit at the initiator and target sites cannot be transportable units.	Units cannot be moved to non-controller configurations without potential data loss.

Table 4-1 Restrictions (Continued)

Restriction	Implication
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:	
<ul style="list-style-type: none"> ■ 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. 	

Configuring Overview

Both the initiator and target sites need some type of CLI interface 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 NT. Use the HyperTerminal emulator. Settings are: 9600 baud, 8 bits, No parity, 1 stop bit, XON/XOFF.

The procedures for configuring a 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." You may visually identify Initiator Site procedures by the text shading throughout this chapter and the next.

Configuration Procedures Outline

TARGET SITE

- Configure the Controllers at the Target Site.
- Configure the Storage at the Target Site.

- Devices and StorageSets.
- 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.
 - Operating System Setup.
 - Install Secure Path for Windows NT.
 - Install SWCC (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.
 - Devices and StorageSets.
 - 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 Secure Path for Windows NT.

- Install SWCC (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 worldwide name on the host bus adapters.
- Establish the name that you will assign to the target and initiator sites. Use a naming convention 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 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 Fibre Channel switches are powered on but not cabled.
7. Turn on the cabinets.

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

8. Establish a local connection to the controller. Refer to the *HSG80 Array Controller ACS Version 8.5 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 should see a display similar to that contained 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 worldwide name, also called the `NODE_ID`, is set. (If zeros are displayed, the name is not set.) If the name is set, proceed to step 15. If the worldwide name has not been assigned to the controller, you will need to obtain the name and set it before proceeding.

NOTE: The subsystem's worldwide 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 worldwide name is valid. If there is no label there, contact your Compaq customer service representative for assistance. Refer to the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, for more information on worldwide names. Each subsystem's worldwide name begins with 5000 and ends in zero, for example 5000-1FE1-FFOC-EE00. The controller port IDs are derived from the worldwide name.

 **CAUTION:** Never set two subsystems to the same worldwide name, or data corruption will occur.

12. Once the worldwide name has been located, assign it to the controller using the following CLI command:

```
SET THIS NODE_ID=node_ID checksum
```

You should see a display similar to that contained 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. Restart the controller using the following CLI command:

```
RESTART THIS_CONTROLLER
```


NOTE: Once you have restarted the controller, you will see a series of %LFL, %CER, and %EVL prompts. These indicate a Last Failure Log, a CLI Event Report, and an Event Log, respectively. For a complete explanation of these event reports, refer to the *HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*.

14. Issue a **SHOW THIS** command to verify that the worldwide name has been set. You should see a display similar to that contained 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 disabled.
.
.
.
```

15. Configure 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 *HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide* for more details on these reports.

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

```
SHOW OTHER_CONTROLLER FULL
```

A display will show that the controllers have been configured to support multiple bus failover mode. You should see a display similar to that contained 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-2
  Configured for MULTIBUS_FAILOVER with ZG8nnnnnnn
  In dual-redundant configuration
  Device Port SCSI address 7
  Time: NOT SET
  Command Console LUN is disabled
  .
  .
  .
```

NOTE: These settings will be 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 using 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> "

19. If you are working in an OpenVMS environment, proceed to step 22. If you are using Windows NT, disable the Command Console LUN (CCL) and set the SCSI version to SCSI-2 by using the following CLI commands:

```
SET THIS NOCOMMAND_CONSOLE_LUN
```

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

20. Verify that the CCL has been disabled and that the SCSI version is set to SCSI-2 by using the following CLI command:

```
SHOW THIS_CONTROLLER
```

The display will indicate that the CCL has been disabled. You should see a display similar to that contained in “Example Display 5”.

Example Display 5

Controller:

```
HSG80 ZG8nnnnnnn Software V85P, Hardware E03
NODE_ID          = nnnn-nnnn-nnnn-nnnn
ALLOCATION_CLASS  = 0
SCSI_VERSION     = SCSI-2
Configured for MULTIBUS_FAILOVER with ZG8nnnnnnn
In dual-redundant configuration
Device Port SCSI address 7
Time: NOT SET
Command Console LUN is disabled
.
.
.
```

NOTE: For OpenVMS, use SCSI-3.

21. Verify that the settings you have established from controller A have been applied to controller B by using the following CLI command:

```
SHOW OTHER_CONTROLLER
```

Proceed to step 23 if you are working in a Windows NT environment.

22. **For OpenVMS:** Set the SCSI version, alloclass, and identifier with the following commands:

```
SET THIS SCSI_VERSION = SCSI-3
SET THIS_CONTROLLER ALLOCATION_CLASS = 1
SET THIS IDENTIFIER = 99
```

Setting this switch causes the host to load the SYSS\$GDRIVER, which provides the GG devices. The value range is 1 - 99.

23. **For OpenVMS and Windows NT:** Check to see if mirrored write-back cache is enabled by using the following CLI command:

```
SHOW THIS_CONTROLLER
```

You should see a display similar to that contained in “Example Display 6”.

Example Display 6

```
Mirrored Cache:  
    Not enabled
```

```
.  
. .  
. .
```

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

```
SET THIS_CONTROLLER MIRRORED_CACHE
```

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

NOTE: It may take up to five minutes after restart to check cache. The controllers will reject this command until cache check is complete. If this command is rejected, don't restart the controllers. Wait a few minutes, then retry.

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

```
SHOW THIS_CONTROLLER
```

You should see a display similar to that contained in “Example Display 7”.

Example Display 7

```
.  
. .  
. .
```

```
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.

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

NOTE: You will be 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
```

26. Restart the controllers (in this order) with the following CLI commands:

```
RESTART OTHER_CONTROLLER
RESTART THIS_CONTROLLER
```

27. To verify that the topology is set correctly, issue the following CLI command:

```
SHOW THIS_CONTROLLER
```

You should see a display similar to that contained in “Example Display 8”.

Example Display 8

```
.
.
.
Host PORT_1:
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
. . . . .PORT_1_TOPOLOGY = FABRIC (offline)
Address . . . . .=nnnnnn
Host PORT_2:
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
. . . . .PORT_2_TOPOLOGY = FABRIC (offline)
Address . . . . .=nnnnnn
NOREMOTE_COPY
.
```

.
.

NOTE: If the port_1_topology = fabric (**point-to-point**) is displayed, this indicates a switch configuration error. Please refer to the *StorageWorks Fibre Channel Storage Switch User's Guide*.

28. You are now ready to enable data replication. Use the following CLI command:

```
SET THIS_CONTROLLER REMOTE_COPY=TargetNodeName
```

Example: SET THIS_CONTROLLER REMOTE_COPY=BuildngB

NOTE: Be sure to specify a meaningful TargetName such as a name that reflects the Target'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 as specified in the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*.

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

29. Use the following CLI command to verify that these settings are in place:

```
SHOW THIS_CONTROLLER
```

You should see a display similar to that contained in "Example Display 9".

Example Display 9

.
.

```
Host PORT_2:
```

```
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
```

```
. . . . . PORT_2_TOPOLOGY = FABRIC (offline)
```

```
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 *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, but note the restrictions listed at the beginning of this chapter.

NOTE: Keep in mind that the target site must have the same exact storageset and unit configuration that the initiator site will have.

Refer to the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, for detailed information on configuring units.

Once all of the units have been created, you can proceed to the steps below.

1. Disable access on all units with this command:

```
SET UnitName DISABLE_ACCESS_PATH=ALL
```

NOTE: Be sure to issue this command for all units.

2. Set the maximum cached transfer size to one with the following CLI command:

```
SET UnitName MAXIMUM_CACHED_TRANSFER_SIZE = 1
```

Repeat this step for each unit

3. Verify that the access on each unit is set to none by using the following CLI command:

```
SHOW UNITS FULL
```

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

Example Display 10

```

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 = 32
  Access:
  None
  State:
  ONLINE to this controller
  Not reserved
  NOPREFERRED_PATH
  Size: nnnnnnnn blocks
  Geometry (C/H/S): ( 7000 / 20 / 254 )
.
.
.

```

NOTE: For OpenVMS, all units must have an identifier set.

4. **For OpenVMS**, set device ID on all units with the following command:

```
SET Unit Name IDENTIFIER = value
```

Example: SET D1 IDENTIFIER = 1

This becomes the VMS device ID for DGA1.

5. To ensure that your storage settings are in place, use the following CLI command:

```
SHOW StoragesetName FULL
```


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 switch.





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 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 4-3 illustrates what your cabling should look like. The numbered callouts reflect the steps that you just completed.



CXO7086A

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. See Figure 4-4 for an illustrated view of how this cabling should appear.



CXO7087A

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

Configure the Host from the Target Site

To run Data Replication Manager, you should have two host bus adapters installed in each host system. If you are using Windows NT, you will need to change the default Topology value from *loop* to *switch*. Follow the procedures outlined in the next section. Refer to the *KGPSA-BC PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide* for installation information.

1. Access REGEDT32. Select START>>RUN, and type REGEDT32.
2. Find DriverSetting with the following sequence:
 - HKEY_LOCAL_MACHINE
 - System
 - CurrentControlSet
 - Services
 - Lp6nds35
 - Parameters
 - Device
 - DriverParameters

3. Using the String Editor, make sure the value of *Topology* in the string is equal to 1. This will tell the driver to operate in a Fibre Channel switched environment.
4. Exit the Registry Editor.

Install Secure Path (NT only)

For installation information, refer to the *StorageWorks Secure Path for Windows NT, A High Availability MultiPath Solution, Installation Guide*.

1. Verify that the Secure Path Agent is installed by going to the control panel and selecting *services*. The Secure Path Agent should be set for automatic start up. Verify that HSZDisk and RAIDisk are also installed by going to the control panel and selecting *devices*. HSZDisk and RAIDisk should be set for startup at boot.
2. Use the *Secure Path Agent Configuration* to grant access to the client at both the initiator and target sites. This can be found by following these menus:
 - Program Files
 - StorageWorks
 - Secure Path Agent
 - Secure Path Cfg
3. You can set the password and allow client access via the *Secure Path Agent Configuration*.

NOTE: Compaq recommends that you set both the fully qualified and unqualified DNS names as valid, authorized clients.
4. Reboot the host.

Install SWCC (optional)

Detailed information about SWCC can be found in the *StorageWorks Command Console Getting Started Guide*.

Connect Fiber Optic Cables Between the Hosts and the Switches

Connecting the fiber optic cables involves two stages:

- Cable Connection Procedure
- Renaming the Host Connections

Both are described in the following sections.

Fiber Optic Cable Connection Procedure

1. Connect a multi-mode, 50-micron fiber optic cable from port 0 of the top switch to one adapter on a host.
2. 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.


3. If you have more than one host (up to four), 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 it does in Figure 4-5.

4. 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.



CXO7088A

Figure 4-5. Cabling between the host and the switches

You should see a display similar to that contained in “Example Display 11”.

Example Display 11

```

Connection Unit
Name      Operating system  Controller  Port  Address  Status  Offset
!NEWCON00  WINNT              THIS       1     210013  online  0
  HOST_ID=nnnn-nnnn-nnnn-nnnn . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

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

5. **For OpenVMS** systems, change the operating system for each connection. Enter the following CLI command:

```
SET !NEWCONnn OPERATING_SYSTEM = VMS
```

6. 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.

You should see a display similar to that contained in “Example Display 12”.

Example Display 12

```
Connection Unit
Name      Operating system  Controller  Port   Address  Status Offset
-----
!NEWCON00      VMS                THIS        1      210013  offline 0
HOST_ID=nnnn-nnnn-nnnn-nnnn . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

!NEWCON01      VMS                OTHER        1      200013  offline 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. Each host bus adapter will appear as a connection. An individual host bus adapter can be identified by its worldwide name that you recorded in the Getting Started chapter and appears in the connection description.

Initially, each connection is named !NEWCONxx. It is much easier to track connections if the connection names are meaningful, for example: HostnameA1.

Figure 4-6 is a helpful worksheet to use when renaming your hosts. Fill in the fields accordingly to prepare for renaming your connections.

!NEWCONxx	Worldwide Name	Host Name	Path Number

Figure 4-6. Host renaming worksheet

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

RENAME !NEWCONxx *TargetHostConnectionName*x

RENAME !NEWCONxx *TargetHostConnectionName*y

Example: RENAME !NEWCONxx hostA1

Example: RENAME !NEWCONxx hostA2

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

SHOW CONNECTIONS

You should see a display similar to that contained in “Example Display 13”.

Example Display 13

```

Connection Unit
Name      Operating system  Controller  Port  Address  Status Offset
-----
HostA1    VMS                THIS       1     210013  online  . . 0
          HOST_ID=nnnn-nnnn-nnnn-nnnn . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

HostA2    VMS                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

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

- Identify the worldwide name on the host bus adapters.
- Establish the name that you will assign to the initiator site. This name should be different from the one you assigned to the target.

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 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 cabinets are turned on, the controllers will boot if the PCMCIA cards are already installed. If there are no cards in the controller slots, insert them now, and depress the reset button. Refer to the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, for complete instructions on how to properly seat the controller cards.

7. Turn on the cabinets.
8. Establish a local connection to the controller. Refer to the *HSG80 Array Controller ACS Version 8.5 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 contained in “Example Display 14”.

Example Display 14

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 worldwide 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 worldwide name is valid. If there is no label there, contact your Compaq customer service representative for assistance. Refer to the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, for more information on worldwide names. Each subsystem's worldwide name begins with 5000 and ends in zero, for example 5000-1FE1-FFOC-EE00. The controller port IDs are derived from the worldwide name.

11. Verify that the subsystem worldwide name is set. If it is, go to step 15. If the worldwide name has not been assigned to the controller, you will need to obtain the name and set it before proceeding.

NOTE: Never set two subsystems to the same worldwide name or data corruption will occur.

12. Once the worldwide 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 contained in "Example Display 15".

Example Display 15

```
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. Restart the controller using the following CLI command:

```
RESTART THIS_CONTROLLER
```

NOTE: Once you have restarted the controller, you will see a series of %LFL, %CER, and %EVL prompts. These indicate a Last Failure Log, a CLI Event Report, and an Event Log, respectively. For a complete explanation of the event reports, refer to the *HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*.

14. Issue a **SHOW THIS** command to verify that the worldwide name has been set.

You will see a display similar to that contained 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-2
Not configured for dual-redundancy
    Controller misconfigured -- other controller present
Device Port SCSI address 7
Time: NOT SET
Command Console LUN is disabled
```

.
.

15. 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.

An %LFL and %EVL prompt will then be displayed. Refer to the *HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide* for more details on these reports.

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

```
SHOW THIS_CONTROLLER FULL
```

Check the display to verify that that the controllers have been configured to support multiple bus failover mode. You will see a display similar to that contained in “Example Display 17”.

Example Display 17

Controller:

```
HSG80 ZG8nnnnnnn Software V85P, Hardware E03
NODE_ID          = nnnn-nnnn-nnnn-nnnn
ALLOCATION_CLASS = 0
SCSI_VERSION     = SCSI-2
Configured for MULTIBUS_FAILOVER with ZG8nnnnnnn
      In dual-redundant configuration
Device Port SCSI address 7
Time: NOT SET
Command Console LUN is disabled
```

NOTE: These settings will automatically be 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 using 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> "

19. If you are working in an OpenVMS environment, proceed to step 23. If you are using Windows NT, set the SCSI version to SCSI-2 and check to see if the Command Console LUN (CCL) is disabled. Use the following CLI commands:

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

```
SHOW THIS_CONTROLLER
```

20. If the CCL is not disabled, issue the following CLI command:

```
SET THIS NOCOMMAND_CONSOLE_LUN
```

21. Verify that the CCL has been disabled by using the following CLI command:

```
SHOW THIS_CONTROLLER
```

You will see a display similar to that contained in “Example Display 18”.

Example Display 18

Controller:

```
HSG80 ZG8nnnnnnn Software V85P, Hardware E03
NODE_ID          = nnnn-nnnn-nnnn-nnnn
ALLOCATION_CLASS = 0
SCSI_VERSION     = SCSI-2
Configured for MULTIBUS_FAILOVER with ZG8nnnnnnn
    In dual-redundant configuration
Device Port SCSI address 7
Time: NOT SET
Command Console LUN is disabled
```

NOTE: For OpenVMS, use SCSI-3.

22. Verify that the settings you have established from controller A have been applied to controller B by using the following CLI command:

```
SHOW OTHER_CONTROLLER
```

Proceed to step 24 if you are working in a Windows NT environment.

23. **For OpenVMS**, set the SCSI version, alloclass, and identifier with the following commands:

```
SET THIS_CONTROLLER SCSI_VERSION = SCSI-3
```

```
SET THIS_CONTROLLER ALLOCATION_CLASS = 1
```

```
SET THIS_CONTROLLER IDENTIFIER = 99
```

Setting this switch causes the host to load the SYSS\$GDRIVER, which provides the GG devices. The value range is 1 - 99.

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

```
SHOW THIS_CONTROLLER
```

You will see a display similar to that contained in “Example Display 19”

Example Display 19

```
Mirrored Cache:  
Not enabled
```

.
.
.

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

```
SET THIS_CONTROLLER MIRRORRED_CACHE
```

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

NOTE: It may take up to five minutes after controller restart to check cache. The controllers will reject this command until cache check is complete. If this command is rejected, don't restart the controllers. Wait a few minutes, then retry.

25. 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 contained in "Example Display 20".

Example Display 20

.

```
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.

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

NOTE: You will be 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
```

27. Restart the controllers (in this order) with the following CLI commands:

```
RESTART OTHER_CONTROLLER
RESTART THIS_CONTROLLER
```

28. To ensure that fabric is up and running, issue the following CLI command:

```
SHOW THIS_CONTROLLER
```

You will see a display similar to that contained in “Example Display 21”.

Example Display 21

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

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

29. You are now ready to enable Data Replication Manager. Use the following CLI command:

```
SET THIS_CONTROLLER REMOTE_COPY=InitiatorName
```

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 as specified in the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*.

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

30. Use the following CLI command to verify that these settings are in place:

```
SHOW THIS_CONTROLLER
```

You will see a display similar to that contained in “Example Display 22”

Example Display 22

.

```
Host PORT_2:
```

```
Reported PORT_ID = nnnn-nnnn-nnnn-nnnn
```

```
PORT_2_TOPOLOGY = FABRIC (offline)
```

```
REMOTE_COPY = BuildngA
```

Configure the Storage at the Initiator Site

Devices and StorageSets

Before you can configure the storage for Data Replication Manager, you need to add the disks, create the RAIDsets, and create units. Follow the instructions in the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, but note the restrictions listed at the beginning of this chapter.

NOTE: Keep in mind that the initiator site must have the same exact storageset and unit configuration as the target site.

Units

Refer to the *HSG80 Array Controller ACS Version 8.5 Configuration Guide*, for detailed information.

Once all of the units have been created, you can proceed to the following steps and configure the storage for the Initiator Site.

1. Disable access on all units with this command:

```
SET UnitName DISABLE_ACCESS_PATH=ALL
```

NOTE: Be sure to issue this command for all units.

2. Verify that the access on each unit is set to none by using the following CLI command:

```
SHOW UNITS FULL
```

You will see a display similar to that contained in “Example Display 23”

Example Display 23

LUN	Uses	Used by

D10DISK1000	
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)	
.		
.		

NOTE: For OpenVMS, all units must have an identifier set.

3. **For OpenVMS**, set device ID on all units with the following command:

```
SET UnitName IDENTIFIER = value
```

Example: SET D1 IDENTIFIER = 1

This becomes the VMS device ID for DGA1.

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

```
SET UnitName PREFERRED_PATH=THIS_CONTROLLER
```

```
SET UnitName PREFERRED_PATH=OTHER_CONTROLLER
```


Keep the busiest units on different host ports.

5. Restart the controllers after configuring the units. Otherwise, the preferred path settings will not go into effect.

Cabling the Initiator Site

Connect Fiber Optic Cables Between the Controllers and Switches

To better understand which ports you will be instructed to connect the cables to, Figure 4-7 shows the port locations on the switch.




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 illustrates what your cabling should look like. The numbered callouts reflect the steps that you just completed.



CXO7089A

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 them documented by issuing the *nbrStateShow* command at the switch. Issue the *topologyShow* command at the switch to reveal if you have more than one fibre 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

Prior to creating the remote copy set, create the connections between the initiator and target sites by typing the following CLI command:

```
ADD REMOTE RCS199 D199 TargetName\D199
```

Example: ADD REMOTE RCS199 D199 BuildngB\D199

NOTE: This command will fail, but creates and names the connections appropriately.

Create Connections From the Target Site

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

```
ADD REMOTE RCS199 D199 InitiatorName\D199
```

Example: ADD REMOTE RCS199 D199 BuildngA\D199

NOTE: This command will fail, but 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 will need to 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: *InitiatorC* and *InitiatorD* are port 2 on the controllers. Be sure to repeat this command for each *UnitName*.

Create Remote Copy Sets 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 will need to allow access to the controllers at the target site. Enable access with this CLI command:

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

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

NOTE: Be sure to repeat this command for each *UnitName*.

3. The CLI command below will create 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*.

Use the following CLI command to create remote copy sets:

```
ADD REMOTE RemoteCopySetName InitiatorUnitName
RemoteNodeName\ TargetUnitName
```

Example: ADD REMOTE RMT0 D1 BuildngBVD1

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

You will see an %EVL display that includes your remote copy set information. You will see a display similar to that contained in “Example Display 24”.

Example Display 24

```
%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: ZG8nnnnnnnn Hardware Version: Enn(2B)
Software Version: V85P
Informational Report
Target Controller Board Serial Number: "      ZG8nnnnnnnn"
```

```
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 will be allowed to the initiator member, and an error will be 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 will be rejected until the remote copy sets return to normal mode.

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

```
SET RemoteCopySetName ERROR_MODE=NORMAL
```

Example: SET RMT0 ERROR_MODE=NORMAL

This can also be used 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 MIRRORSET MIR_D1LOG DISK50100 DISK60100
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 contained in “Example Display 25”.

Example Display 25

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 contained in “Example Display 26”.

Example Display 26

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 can 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 contained in “Example Display 27”.

Example Display 27

```

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, use 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 contained in “Example Display 28”.

Example Display 28

Name	Association	Uses	Used by
AS_D1	association	RC_D1	

Switches:

NOFAIL_ALL

NOORDER_ALL

LOG_UNIT = D10 (No data logged)

Configure the Host from the Initiator Site

Install the Host Bus Adapters and Drivers

You should install two host bus adapters in each host system. If you are using Windows NT, you will need to change the default Topology value from *loop* to *switch*. Follow the procedures outlined. Refer to the *KGPSA-BC PCI-to-Optical Fibre Channel Host Bus Adapter User's Guide* for installation information.

1. To access REGEDT32, select START>RUN, and enter REGEDT32.
2. Find DriverSetting with the following sequence:
 - HKEY_LOCAL_MACHINE
 - SYSTEM
 - CurrentControlSet
 - Services
 - Lp6nds35
 - Parameters
 - Device
 - DriverParameters
3. Using the String Editor, make sure the value of *Topology* in the string is equal to *1*. This will tell the driver to operate in a Fibre Channel switched environment.
4. Exit the Registry Editor.

Install Secure Path (NT only)

For installation information, refer to the *StorageWorks Secure Path for Windows NT, A High Availability MultiPath Solution, Installation Guide*.

1. Verify that the Secure Path Agent is installed by going to the control panel and selecting *services*. The Secure Path Agent should be set for automatic start up. Verify that HSZDisk and RAIDisk are also installed by going to the control panel and selecting *devices*.
2. Use the *Secure Path Agent Configuration* to grant access to the client at both the initiator and target sites. This can be found by following these menus:
 - Program Files
 - StorageWorks
 - Secure Path Agent
 - Secure Path Cfg
3. You can set the password and allow client access via the *Secure Path Agent Configuration*.

NOTE: Compaq recommends that you set both the fully qualified and unqualified DNS names as valid, authorized clients.
4. Reboot the host.

Install SWCC (optional)

Detailed information about SWCC can be found in the *CompaqStorageWorks Command Console Getting Started Guide*.

Connecting Fibre Optic Cables Between the Hosts and the Switches

1. Connect a multi-mode, 50-micron fiber optic cable from port 0 of the top switch to the host.
2. Connect the final multi-mode, 50-micron fiber optic cable from port 0 of the bottom switch to the host.

The host is now connected to the initiator site via the multi-mode, 50-micron fiber optic cables. Your cabling should appear as it does in Figure 4-10.

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




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

3. 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.

You will see a display similar to that in Example Display 29.

Example Display 29

```

Connection Unit
Name      Operating system  Controller  Port  Address  Status  Offset
-----
!NEWCON00      WINNT           THIS        1     210013  online  0
HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

!NEWCON01      WINNT           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

```

4. **For OpenVMS** systems, change the operating system for each connection. Enter the following CLI command:

```
SET !NEWCONnn OPERATING_SYSTEM = VMS
```

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.

You should see a display similar to that contained in “Example Display 30”.

Example Display 30

```

Connection Unit
Name      Operating system  Controller  Port   Address  Status Offset
-----
!NEWCON00      VMS                THIS        1      210013  offline 0
  HOST_ID=nnnn-nnnn-nnnn-nnnn . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

!NEWCON01      VMS                OTHER        1      200013  offline.0
  HOST_ID=nnnn-nnnn-nnnn-nnnn . . . . ADAPTER_ID=nnnn-nnnn-nnnn-nnnn

```


Rename the Host Connections from the Initiator Site

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

Figure 4-12 is a helpful worksheet to use when renaming your hosts. Fill in the fields accordingly to keep an accurate record of connections and host names.

!NEWCONxx	Worldwide Name	Host Name	Path Number

Figure 4-12. Host renaming worksheet

When you have completed the worksheet, rename the !NEWCONxx prompt using the following CLI commands:

```

RENAME !NEWCONxx InitiatorHostConnectionNamex
RENAME !NEWCONxx InitiatorHostConnectionNamey

```

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

```
SHOW CONNECTIONS
```

You will see a display similar to that contained in “Example Display 31”.

Example Display 31

Connection Unit

Name Offset	Operating system	. .Controller	Port	Address	Status
HOSTA1	WINNT	THIS	1	210013	online 0
HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn					
HOSTA2	WINNT	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		online 0
HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn					
BUILDNGBC	PPRC_INITIATOR	THIS	2		online 0
HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn					
BUILDNGBD	PPRC_INITIATOR	OTHER	2		online 0
HOST_ID=nnnn-nnnn-nnnn-nnnn ADAPTER_ID=nnnn-nnnn-nnnn-nnnn					

At the target site, the initiator host will appear as new connections, !NEWCONxx. These should be renamed as described above.

Enable Access to the Hosts at the Initiator Site

1. The initiator units will need to have access to the hosts. Enable access with this command:

```
SET UnitName ENABLE_ACCESS_PATH=  
(InitiatorHostConnectionName,InitiatorHostConnectionName)
```

Example: SET *UnitName* ENABLE_ACCESS_PATH=(HostA1,HostA2)

NOTE: Keep in mind that there should be two paths per host. You will need to repeat this sequence for each host. Be sure to reboot the host after you have enabled hosts to access to units.

2. Reboot the host.
3. Verify that Hszdisk has been installed. See *Secure Path* installation procedures, Step 1, in this chapter.
4. Run Disk Administrator to create partitions, format, and assign drive letters on the newly created storage. Once the storage has been configured with Disk Administrator, you need to save this configuration to a floppy for use during a failover or failback condition. To create a saved Windows NT configuration from the Disk Administrator, perform the following steps:
 - a. Open *Partition* and select *Configuration*.
 - b. Select *Save* and insert a floppy as requested.

NOTE: It is recommended that two copies of the saved configuration be created and maintained at both the initiator and target sites. This data will be used to restore a known configuration in the event of a site failover or failback.

5. You are now ready to run Secure Path Manager, which can be found in the *Start/Programs/StorageWorks/Secure Path Manager* path. Start the application, and specify the server and password that you prefer.
6. Select the “Save Password” box if you would like to use the same password each time you log in.
7. The *StorageWorks Secure Path Manager* screen appears, and now you must verify the drives. Go to the disk you want to check, right-click the mouse, and choose *Properties*. You will see the device properties.

Install Cluster Server for Windows NT (optional)

Windows NT Fibre Channel cluster software enables two host servers to share a Fibre Channel storage subsystem through a Fibre Channel switch. If a failure on the server occurs, the cluster software detects that failure, and a failover is initiated. The failed components can be warm-swapped or serviced while the functioning components remain active. This process requires minimal downtime and ensures high availability of data. If you are using Windows NT and wish to run the cluster option, you can safely install it now. Refer to the *Storageworks RAID Array8000/ESA12000 Fibre Channel Cluster Solutions for Windows NT Installation Guide*.

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, use the HyperTerminal emulator. Settings to be used are: 9600 baud, 8 bits, No parity, 1 stop bit, XON/XOFF.
3. Select *Capture Text* from the Transfer menu.
The *Capture Text* dialog box appears.
4. In the c:\field, type *initiator.txt* or *target.txt*.
5. 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 32”.

Example Display 32

Controller:

```
HSG80 ZG91412410 Software S050P-0, Hardware E05
NODE_ID           = 5000-1FE1-0001-3AE0
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 = 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 33” for each association set.

Example Display 33

```
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 contained in “Example Display 34” for each remote copy set.

Example Display 34

```

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

```

- 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 35” for each unit.

Example Display 35

```

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 )
    
```

- 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 36” for each connection.

Example Display 36

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

- Save this file for future reference.
- 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:

- “Power Up Data Replication Manager Systems” on page 5-2
- “Power Down Data Replication Manager Systems” on page 5-3
- “Site Failover Basic Description” on page 5-3
- “Failback Procedure Choices” on page 5-5
- “Data Replication Manager Configuration Basics” on page 5-6
- “Planning Considerations” on page 5-7
- “Planned Failover Procedures” on page 5-8
- “Simple Failback Procedure” on page 5-12
- “Unplanned Failover” on page 5-16
- “Full Failback Procedure” on page 5-18
- “New Hardware Failback Procedure” on page 5-25

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 Powerup 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 Powerup Procedures” .

Initiator Site Powerup 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 that there will be 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 will prevent 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 will help you understand 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.

The illustration below depicts a basic DRM configuration and will be 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 described will 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.

This is because a unit will 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 will not be able to access or alter the unit or its remote copy set in any way from the remaining controller. Access will not be reestablished until the failed controller is either repaired or replaced. Furthermore, you will 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 will indicate 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.

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 Failover 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.) Verify that your target controller configuration is the same as your initiator controller configuration.

2. The following steps will require actions relative to each operating system being used in your configuration.
 - a. **Windows NT-86:** If the operating system is up and running, shut down the operating system, and power off the hosts.
 - b. **OpenVMS:** If the operating system is up and running, and is being used exclusively for DRM operations, shutdown the operating system and power off the hosts. If the operating system is used for other applications, remove all I/O to the LUNs with remote copy sets that will be failed over and then dismount the volumes associated with these LUNs.
3. 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.
4. 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.
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. 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, type 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)

7. 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.
8. Shut down the initiator HSG80 controllers (in this order) with the following CLI commands:
SHUTDOWN OTHER_CONTROLLER
SHUTDOWN THIS_CONTROLLER
9. After the preceding command has completed and both controllers have shut down successfully, power off the controller cabinet. If the initiator site will be powered down for a long period of time, you may need to disable cache batteries.

Continue the failover process at the target site with the “Initiator 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 for information on how to create 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. The following step will require actions relative to each operating system being used in your configuration:
 - a. **Windows NT-86:** Allow hosts to recognize new units:

Reboot the server(s) at the target site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.
 - b. **OpenVMS:** Allow hosts to recognize new units:

If you have shut down the host, boot the host at this time. Booting the host enables OpenVMS to recognize the units.

If you did not shut down the host, use the following command from a privileged account to enable OpenVMS to recognize the units:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

This completes the failover 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

Initiator Site Failback Preparation Procedure

1. Shut down the initiator hosts if they are still up and running. (This may not be necessary for all operating systems)
2. 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” of this chapter.

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

3. If you made the decision to remove the remote copy set targets, use the "Full Failback Procedure" of this chapter.
4. Power up the controller cabinets. Once the connection between the initiator and the target has been re-established, the remote copy sets will begin to merge.

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 at this point. You should 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 will result in a full normalization:

```
SET RemoteCopySetName ADD=InitiatorRemoteCopyName\UnitNumber
```

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

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

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 host(s). This will not be necessary on all Operating Systems.

IMPORTANT: If host(s) are not shutdown, host access must be removed on all LUNs used with Remote Copy Sets.

3. Disable host access to the target units for all remote copy sets by using the following CLI command:

```
SET UnitName DISABLE=(TargetHostConnectionName,TargetHostConnectionName)
```

4. You may now boot hosts for non-remote copy set units.
5. Turn off write history logging, if enabled, with the following CLI command:

```
SET AssociationSetName NOLOG_UNIT
```

Repeat this procedure for each association set.

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

```
DELETE AssociationSetName
```

Repeat this procedure for each association set.

7. 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 will eventually succeed.

Repeat this step for all remote copy sets.

8. 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.

9. 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 the DRM Configuration Guide 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=(InitiatorHostConnectionName,InitiatorHostConnectionName)`
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 for information on how to create and 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. The following step will require actions relative to each operating system being used in your configuration:
 - a. **Windows NT-86:** Allow hosts to recognize new units:

Reboot the server(s) at the initiator site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.

b. **OpenVMS:** Allow hosts to recognize new units:

If you have shut down the host, boot the host at this time. Booting the host enables OpenVMS to recognize the units.

If you did not shut down the host, use the following command from a privileged account to enable OpenVMS to recognize the units:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

This completes the Simple Failback 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, the units are required to be preferred 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=buildingA\d1
Repeat this step for all remote copy sets.
5. Give the target site hosts access to the units in its storage subsystems with this command:
SET UnitName ENABLE = (TargetHostConnectionName,TargetHostConnectionName)
If you do not recall the target host name, use the SHOW CONNECTION command.
Repeat this step for all units.
6. To verify that all of the steps have been completed successfully, issue this CLI command:
SHOW REMOTE_COPY FULL
7. 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.
8. 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
9. (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

NOTE: The target will be added at failback.
10. The following step will require actions relative to each operating system being used in your configuration.
 - a. **Windows NT-86:** Allow hosts to recognize new units:
Reboot the server(s) at the target site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.
 - b. **OpenVMS:** Allow hosts to recognize new units:
If you have shut down the host, boot the host at this time. Booting the host enables OpenVMS to recognize the units.

If you did not shut down the host, use the following command from a privileged account to enable OpenVMS to recognize the units:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

This completes the failover procedure. When the problem that disabled the initiator site is remedied, refer to the matrix in "Failback Procedure Choices" for proper failback procedure.

Full Failback Procedure

Before performing 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. Power up the controllers if necessary.
2. 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.
3. 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.

4. The following steps will require actions relative to each operating system being used in your configuration.
 - a. **Windows NT-86:** If the operating system is up and running, shut down the operating system, and power off the hosts.
 - b. **OpenVMS:** If the operating system is up and running, and is being used exclusively for DRM operations, shutdown the operating system and power off the hosts. If the operating system is used for other applications, remove all I/O to the LUNs with remote copy sets that will be failed over and then dismount the volumes associated with these LUNs.
5. Disable access to all remote copy set units by issuing the following CLI command:
SET UnitName DISABLE=ALL
Repeat this step for all copy set units.
6. Optional: 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
7. For each remote copy set, use the following CLI command to set the error mode to normal if set to failsafe mode:
SET RemoteCopySetName ERROR_MODE = NORMAL
8. At the initiator site, the units are required to be preferred to one controller or the other using the following CLI command:
SET UnitName PREFERRED_PATH = THIS_CONTROLLER or OTHER_CONTROLLER
9. Set maximum cached transfer size back to 1 with the following CLI command:
SET UnitName MAXIMUM_CACHED_TRANSFER_SIZE = 1
Repeat this step for all association sets.
10. Disable Port 1 connections to fabric by using the following CLI commands:
SET THIS (and OTHER) PORT_1_TOPOLOGY=OFFLINE

11. Delete association sets and log disk by using the following CLI commands:

```
SET AssociationSetName NOLOG_UNIT
```

```
DELETE AssociationSetName
```

Repeat this step for all association sets.

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

```
DELETE RemoteCopySetName
```

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

```
SHOW REMOTE_COPY FULL
```

You should not see any remote copy sets.

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 will report as failed, but it creates and names the connections appropriately.
6. Continue with the full failover 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,TargetRemoteCopyNameB,
TargetRemoteCopyNameC,TargetRemoteCopyNameD)
Repeat this procedure for all copy set units.
2. Verify that you have enabled access with the following CLI command:
SHOW UnitName FULL

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 will be NORMAL.
4. Stop I/O from the target hosts to the remote copy set units.
5. Disable host access to the target units by using the following CLI command:
SETUnitName DISABLE=(TargetHostConnectionName,TargetHostConnectionName)
Repeat this step for all units.
6. Shut down the target HSG80 controllers (in this order) with the following CLI commands:
SHUTDOWN OTHER_CONTROLLER
SHUTDOWN THIS_CONTROLLER

Continue with the full failback procedure at the initiator site with “Initiator Site Return Control Procedure” .

Initiator Site Return Control Procedure

This section describes the returning of 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.

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

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 restoring of 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 for information on how to create 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:
SETUnitName 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
9. In the Access field of the display, all units should show that the initiator hosts are enabled.
10. The following step will require actions relative to each operating system being used in your configuration:
 - a. **Windows NT-86:** Allow hosts to recognize new units:

Reboot the server(s) at the initiator site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.
 - b. **OpenVMS:** Allow hosts to recognize new units:

If you have shut down the host, boot the host at this time. Booting the host enables OpenVMS to recognize the units.

If you did not shut down the host, use the following command from a privileged account to enable OpenVMS to recognize the units:
MCR SYSMAN IO AUTOCONFIGURE/LOG

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. The following steps will require actions relative to each operating system being used in your configuration.
 - a. **Windows NT-86:** If the operating system is up and running, shut down the operating system, and power off the hosts.
 - b. **OpenVMS:** If the operating system is up and running, and is being used exclusively for DRM operations, shutdown the operating system and power off the hosts. If the operating system is used for other applications, remove all I/O to the LUNs with remote copy sets that will be failed over and then dismount the volumes associated with these LUNs.
2. Manually reconfigure the controllers, but do not re-create the original remote copy sets. This procedure includes the following steps:

NOTE: Steps c, f, and g will 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 for information on the node ID and World Wide Name.

OpenVMS only: Set the Identifier to its previous value by using the following command:

```
SET THIS_CONTROLLER IDENTIFIER = value (
```

```
Example: SET THIS_CONTROLLER IDENTIFIER = 99
```

- 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.
```

NOTE: This CLI command may fail because internal cache diagnostics are running. This diagnostics can take up to 5 minutes to complete, so a retry of this command may be necessary.

- f. Enter the following command

```
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 will be part of remote copy sets must be identical to the corresponding units at the target site.

OpenVMS only: Set the device ID as they were prior to this hardware replacement on all units by using the following command:

```
SET UnitName IDENTIFIER = value
```

Example: SET D1 IDENTIFIER = 1

This becomes the VMS device ID for DGx1.

3. Disable access to all units by issuing the following CLI command:
SET UnitName DISABLE=ALL
Repeat this step for all units.
4. Optional: 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
OpenVMS only: Set the device ID on all new units by using the following command:
SET UNIT IDENTIFIER = value
Example: SET D1 IDENTIFIER = 1
This becomes the VMS device ID for DGx1.
5. At the initiator site, the units are required to be preferred 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 32 with the following CLI command:
SET UnitName MAXIMUM_CACHED_TRANSFER_SIZE = 32
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. A status comparison is accomplished by bringing up a terminal emulator session and entering a SHOW THIS command.

Continue with the new hardware 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. Remove the targets from the remote copy sets if necessary 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 commands:
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 will report as failed, but it creates and names the connections appropriately.

Continue with the new hardware failover 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. Issue the following CLI command:

```
ADD REMOTE_COPY_SETS RCS199 D199 TargetRemoteCopyName\D199
```

NOTE: This command will report 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,TargetRemoteCopyNameB,  
TargetRemoteCopyNameC,TargetRemoteCopyNameD)
```

Repeat this procedure for all units.

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 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 will be NORMAL.

4. Stop I/O from the target hosts to the remote copy set units.

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

```
SETUnitName DISABLE=(TargetHostConnectionName,TargetHostConnectionName)
```

Repeat this step for all units.

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

```
SHUTDOWN OTHER_CONTROLLER
```

```
SHUTDOWN THIS_CONTROLLER
```

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 the returning of 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.

Continue with the new hardware failback procedure at the target site with the “Target Site Restore Procedure” .

Target Site Restore Procedure

Both controllers on the target site must be restarted after the site failover has taken place.

1. 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
```

Continue with the new hardware failback procedure at the initiator site with “Initiator Site Restoration of Target Connections” .

Initiator Site Restoration of Target Connections

This section describes the restoring of 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 for information on how to create 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. All the connections that were renamed set then to there appropriate operating system using the following CLI command.

```
SET !NEWCONXX OPERATING_SYSTEM = (VMS or WINNT)
```

Example: SET hostA1 OPERATING_SYSTEM=VMS

7. Enable host access to the units by using the following CLI command:
`SETUnitName ENABLE=(InitiatorHostConnectionNamex,InitiatorHostConnectionNamey)`
8. To verify that all of these steps have been completed successfully, issue this CLI command:
`SHOW REMOTE_COPY FULL`
9. 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.
10. Set maximum cached transfer size to the original value using the following CLI command:
`SET UnitName MAXIMUM_CACHE_TRANSFER_SIZE = initiator value`
11. 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.
12. The following step will require actions relative to each operating system being used in your configuration:
 - a. **Windows NT-86:** Allow hosts to recognize new units:
Reboot the server(s) at the initiator site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.
 - b. **OpenVMS:** Allow hosts to recognize new units:
If you have shut down the host, boot the host at this time. Booting the host enables OpenVMS to recognize the units.
If you did not shut down the host, use the following command from a privileged account to enable OpenVMS to recognize the units:
`MCR SYSMAN IO AUTOCONFIGURE/LOG`

This completes the New Hardware 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” on page 6–2
 - ❑ “Forced Errors Detected During Copy” on page 6–2
 - ❑ “Read Errors Detected During Full Copy” on page 6–2
 - ❑ “Dual Redundancy During Failback” on page 6–3
 - ❑ “Failsafe Lock Management” on page 6–3
 - ❑ “Link Failure Management” on page 6–3
 - ❑ “Remote Copy Set Member Failures” on page 6–3
 - ❑ “Remote Copy Set Worldwide LUN ID” on page 6–4
 - ❑ “Write History Logging” on page 6–4
 - ❑ “Failure Notification” on page 6–5
 - ❑ “HSG80 Array Controller Failure” on page 6–5
 - ❑ “SWCC Failure” on page 6–6

- ❑ “Failure of One Member in a Dual Redundant Controller Pair” on page 6-6
- “Failure Modes of a DT System in Normal Operation” on page 6-7
- ❑ “Failure of Both Fiber Optic Cables or Switch” on page 6-6
- ❑ “Failure at Target Site after Failover” on page 6-9

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 will be 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 qualifier should be restored from a known good backup.

Refer to the *HSG80 Array Controller ACS Version 8.5 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, 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. Therefore, the controller will terminate the copy and report the event.

Unrecoverable read errors on the source member will terminate the copy and send a fault management report to the host. Refer to the *HSG80 Array Controller ACS Version 8.5 Configuration and 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 it 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 will restart. 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 an indicator of a defective controller.

Remote Copy Set Member Failures

While most remote copy set members will be 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 will cause a write failure at the target. The target member will be removed, and the remote copy set will be 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 will become unavailable to the host. The target member of the remote copy set is not read and 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 will be 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 will 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, 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 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. See the *HSG80 Array Controller ACS Version 8.5 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, see the *HSG80 Array Controller ACS Version 8.5 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 will be notified via the GUI screen of any changes in status. The user is able to manipulate controllers and storage through the GUI and intervene in the DT process when there is a problem.

Refer to the *StorageWorks Command Console Getting Started Guide* and the on-line users help for more information.

Failure of One Member in a Dual Redundant Controller Pair

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 will automatically be 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, then I/O will continue through the initiator host, and the target will still be removed.

If you lose the fiber optic cable connection of a switch at either site, refer to Figure 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 <i>Loss of:</i>	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

Table 6-1 Failure Modes of a DT System with Normal Operation (Continued)

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 <i>Loss of:</i>	Action
									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 the same 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 <i>Loss Of</i>	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

Appendix **A**

Status Comparison

This appendix describes the procedure for performing a comparison of 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 from a local terminal
- Comparing the Results

Target Site 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 that is capable of capturing text. Settings to be used are: 9600 baud, 8 bits, No parity, 1 stop bit, XON/XOFF.

SHOW Commands

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

```
SHOW THIS_CONTROLLER FULL
```

You should 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 this file for future reference.

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

Name	Association	Uses	Used by
AS1	association	RC1 RC2 RC3	

Switches:

NOFAIL_ALL
 NOORDER_ALL
 NOLOG_UNIT

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

```
D2                                DISK10100          BuildingB\RC2
LUN ID: nnnnnnnnnnnnnnnnnnnnnnn
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

```
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.
```


Glossary

This glossary defines terms pertaining to the Data Replication Manager for the HSG80 running ACS V8.5P. It is not a comprehensive glossary of computer terms.

ACS	An abbreviation representing Array Controller software. <i>See</i> 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	A term used to express Arbitrated Loop Physical Address. A two-digit hexadecimal number that expresses the 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.
array controller	<i>See</i> controller.
array controller software	Also known by the abbreviation ACS. ACS is software that is contained on a removable PCMCIA program card that provides the operating system for the array controller.

association sets	<p>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:</p> <ul style="list-style-type: none">■ Shares the same log unit■ Has its host access removed from all members when one member fails■ Keeps I/O order across all members <p>CLI commands available are ADD ASSOCIATIONS and SET <i>associations</i>.</p>
asynchronous mode	<p>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.</p> <p>Asynchronous mode can provide greater performance and response time, but the data on all members at any one point in time cannot be assumed to be identical.</p> <p>)See also synchronous mode.</p>
ATM	<p>An abbreviation 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.</p>
autospare	<p>A controller feature that automatically replaces a failed disk drive. Autospare aids the controller in automatically replacing failed disk drives. You can enable the <i>AUTOSPARE</i> switch for the failedset causing physically replaced disk drives to be automatically placed into the spareset. Also called “autonewspare.”</p>
bad block	<p>A data block that contains a physical defect.</p>
bad block replacement	<p>Also known by the term BBR. BBR is 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.</p>
BBR	<p>See bad block replacement.</p>

block	<p>A stream of data stored on a 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.</p>
cache memory	<p>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.</p>
CBR	<p>An abbreviation used for Constant Bit Rate. CBR is a category of ATM service. This category supports a constant (guaranteed) data rate. CBR supports applications that require a highly-predictable transmission rate.</p>
chunk	<p>A block of data written by the host.</p> <p><i>See also</i> block, chunk size</p>
chunk size	<p>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.</p>
CLI	<p>An abbreviation used to express the Command Line Interpreter. Also known as Command Line Interface. The CLI is the configuration interface to operate the controller software.</p> <p><i>See</i> command line interpreter.</p>
connection	<p>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.</p> <p>CLI commands available are ADD CONNECTIONS, SET <i>connection-name</i>.</p> <p><i>See also</i> link.</p>
container	<ol style="list-style-type: none">1. Any entity that is capable of storing data, whether it is a physical device or a group of physical devices.2. A virtual internal controller structure representing either a single disk or a group of disk drives linked as a storageset (stripesets and mirrorsets are examples of storageset containers the controller uses to create units).

controller	A hardware device that 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	<p>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.</p> <p>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.</p>
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	<i>See</i> node and peripheral device.
disaster tolerance	<p>Disaster tolerance provides the ability for rapid recovery of user data from a remote location when a significant event (or disaster) at the primary computing site occurs.</p> <p><i>See also</i> remote copy sets, DT</p>
DT	<p>An abbreviation for Disaster Tolerance. The Data Replication manager is an example of a database that is made disaster-tolerant.</p> <p><i>See also</i> disaster tolerance</p>
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. <i>See also</i> failover, failback.
ECB	A abbreviation 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	<i>See</i> ECB.
F_Port	A port in a fabric where an N_Port or NL_Port may attach. <i>See</i> N_Port, NL_Port, and 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. <i>See also</i> 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. CLI commands available are: SITE_FAILOVER <i>See also</i> 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 occur, the remote copy set goes into the inoperative (offline) state and the failsafe error mode is “Failsafe Locked.” The CLI command SET <i>remote-copy-set-name</i> ERROR_MODE=FAILSAFE enables this error mode.

fast failback	<p>A term representing the synchronization of the initiator site with the target during a planned failover of the initiator subsystem.</p> <p>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.</p> <p><i>See also</i> mini-merge, unplanned failover, planned failover, write history logging.</p>
FC-AL or FCAL	<p>A term used to express 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 is attached to the receive input of the next.</p>
FC-ATM	<p>A term expressing ATM AAL5 over Fibre Channel.</p>
FC-FG	<p>A term expressing Fibre Channel Fabric Generic Requirements.</p>
FG-FP	<p>A term expressing Fibre Channel Framing Protocol</p> <p><i>See</i> HIPPI on FC.</p>
FC-GS-1	<p>A term expressing Fibre Channel Generic Services-1.</p>
FC-GS-2	<p>A term expressing Fibre Channel Generic Services-2.</p>
FC-IG	<p>A term expressing Fibre Channel Implementation Guide.</p>
FC-LE	<p>A term expressing Fibre Channel Link Encapsulation (ISO 8802.2).</p>
FC-PH	<p>A term expressing the Fibre Channel Physical and Signaling Standard.</p>
FC-SB	<p>A term expressing Fibre Channel Single Byte Command Code Set.</p>
FC-SW	<p>A term expressing 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.</p>
FCC	<p>A term expressing the Federal Communications Commission. The federal agency responsible for establishing standards and approving electronic devices within the United States.</p>
FCC Class A	<p>This certification label appears on electronic devices that can only be used in a commercial environment within the United States.</p>

FCC Class B	This certification label appears on electronic devices that can be used in either a home or a commercial environment within the United States.
FCP	The mapping of SCSI-3 operations to Fibre Channel.
FDDI	Fiber Distributed Data Interface. An ANSI standard for 100 megabaud transmission over fiber optic cable.
FD SCSI	The fast, narrow, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s. <i>See also</i> FWD SCSI and SCSI.
fiber	A fiber or optical strand used in fiber optic cable. Spelled <i>fib</i> re when used in “Fibre Channel” protocol. <i>See also</i> 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. <i>See</i> N_Port, NL_Port, and F_Port. <i>See also</i> fabric.
frame	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. <i>See also</i> Initiator.
GBIC	Gigabit Interface Converter. The hardware devices inserted into the ports of the Fibre Channel switch that hold the Fibre Channel cables. Converts fiber optic cable connections to Fibre Channel switch connections

GLM	An abbreviation for the Gigabit Link Modules used in Fibre Channel long distance applications. The GLMs provide the ability to increase the fiber optic cable transmission distances from 10 KM to ???
hard address	The ALPA which an NL_Port attempts to acquire during loop initialization.
heterogeneous host support	Also called <i>noncooperating host support</i> . This term expresses the ability to share storage between two similar (or dis-similar) hosts by way of storage partitioning.
HIPPI-FC	An acronym expressing 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.
initiator	<p>A term that is defined as:</p> <ol style="list-style-type: none">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-tolerant Data Replication Manager solution, 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. <p><i>See also</i> target.</p>
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.
L_port	A node or fabric port capable of performing arbitrated loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.
Link	<p>A connection between two adjacent Fibre Channel ports consisting of a transmit fibre and a receive fibre. An example would be the connection between the Fibre Channel switch port and the HSG80 controller.</p> <p><i>See also</i> connection.</p>
Logical Block Number	<i>See</i> LBN.

logical unit	A physical or virtual device addressable through a target ID number. LUNs use their target's bus connection to communicate on the SCSI bus.
Logical Unit Number	Abbreviated LUN. A value that identifies a specific logical unit belonging to a SCSI target ID number. A number associated with a physical device unit during a task's I/O operations. Each task in the system must establish its own correspondence between logical unit numbers and physical devices.
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). <i>See also</i> write history logging.
long distance mirroring	Also known as peer-to-peer remote copy. <i>See also</i> remote copy sets
loop	<i>See</i> arbitrated loop.
loop_ID	A seven-bit value numbered contiguously from zero to 126-decimal and represent the 127 legal AL_PA values on a loop (not all of the 256 hex values are allowed as AL_PA values per FC-AL).
loop tenancy	The period of time between the following 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.
mini-merge	As applied to the Data Replication manager: a term representing the data transfers to be made whenever a target becomes inaccessible. Inaccessibility would be both links or both target controllers going 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. <i>See</i> fast failback, write history logging.
mirroring	The act of creating an exact copy or image of data.
N_port	A port attached to a node for use with point-to-point topology or fabric topology.

NL_port	A port attached to a node for use in all three topologies.
network	A data communication, a configuration in which two or more terminals or devices are connected to enable information transfer.
Non-L_Port	A Node of Fabric port that is not capable of performing the Arbitrated Loop functions and protocols. N_Ports and F_Ports loop-capable ports.
non-participating mode	A mode within an L_Port that inhibits the port from participating in loop activities. L_Ports in this mode continue to retransmit received transmission words but are not permitted to arbitrate or originate frames. An L_Port in non-participating mode may or may not have an AL_PA. <i>See also</i> participating mode.
node	<ol style="list-style-type: none">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 one N_Port or NL_Port.
normal member	A mirrorset member that, block-for-block, contains the same data as other normal members within the mirrorset. Read requests from the host are always satisfied by normal members.
normalizing	Normalizing is a state in which, block-for-block, data written by the host to a mirrorset member is consistent with the data on other normal and normalizing members. The normalizing state exists only after a mirrorset is initialized. Therefore, no customer data is on the mirrorset.
normalizing member	A mirrorset member whose contents is the same as all other normal and normalizing members for data that has been written since the mirrorset was created or lost cache data was cleared. A normalizing member is created by a normal member when either all of the normal members fail or all of the normal members are removed from the mirrorset. <i>See also</i> copying member.
OC-3	A term used to express 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. <i>See also</i> 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 to be in participating mode.

PCM	Polycenter Console Manager.
PCMCIA	<p>An abbreviation 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 is about the size of a credit card. It is used in the HSJ80 to load the controller software.</p> <p><i>See also</i> program card, ACS.</p>
PCR	An abbreviation used to express peak cell rate. PCR is the maximum transmission speed of a virtual connection. PCR is a required parameter for the CBR service category.
peer-to-peer remote copy	<i>See</i> remote copy sets.
planned failover	<p>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.</p> <p><i>See also</i> synchronous operation mode</p>
PL_DA or PLDA	A term used to express 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.
port	<ul style="list-style-type: none">■ In general terms, the port is:<ol style="list-style-type: none">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:<ol style="list-style-type: none">1) The logical route for data in and out of a controller that can contain one or more channels, all of which contain the same type of data.2) The hardware and software that connects a controller to a SCSI device.

port_name	A 64-bit unique identifier assigned to each Fibre Channel port. The Port_Name is communicated during the 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. <i>See also</i> PCMCIA.
PTL	An abbreviation for Port-Target-LUN. The controller's method of locating a device on the controller device bus: <ul style="list-style-type: none">■ 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	Power Verification and Addressing module.
PVC	An abbreviation used to express 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.

QoS	<p>A term used to express 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).</p> <p>Performance issues include data, rate, cell loss rate, cell delay, and delay variation (jitter).</p> <p>Categories of ATM service are:</p> <ul style="list-style-type: none"> ■ 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).
redundancy	<p>The provision of multiple interchangeable components to perform a single function in order to cope with failures and errors. A RAIDset is considered to be redundant when user data is recorded directly to one member and all of the other members include associated parity information.</p>
remote copy sets	<p>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.</p> <p>CLI commands available are: <code>ADD REMOTE_COPY_SETS</code>, <code>SET remote-copy-set-name</code>, <code>SET controller REMOTE_COPY</code>.</p> <p><i>See also</i> disaster tolerance.</p>
remote copy set metadata	<p>Remote copy set metadata describes the remote copy set membership and state. This metadata is located in the mirrored write-back cache on the controller where each member resides to assist with site failover. 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.</p>
replacement policy	<p>The policy specified by a CLI command switch (<code>SET FAILEDSET</code> 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 <i>AUTOSPARE</i> and <i>NOAUTOSPARE</i>.</p>

SCSI	<p>An acronym for Small Computer System Interface:</p> <ol style="list-style-type: none">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	<p>A host computer adapter, a peripheral controller, or an intelligent peripheral that can be attached to the SCSI bus.</p> <p>Any physical unit that can communicate on a SCSI bus.</p>
SCSI device ID number	<p>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.</p>
SCSI ID number	<p>The representation of the SCSI address that refers to one of the signal lines numbered 0 through 15.</p>
storage array	<p>An integrated set of storage devices.</p>
storage unit	<p>The general term that refers to storage sets, single-disk units, and all other storage devices that are installed in your subsystem and accessed by the host. A storage unit can be any entity that is capable of storing data, whether it is a physical device or a group of physical devices.</p>
storage set	<p>A group of devices configured with RAID techniques to operate as a single container.</p> <p>Any collection of containers, such as stripe sets, mirror sets, striped mirror sets, JBODs, and RAID sets.</p>
subnet mask	<p>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.</p>
surviving controller	<p>The controller in a dual-redundant configuration pair that serves its companion's devices when the companion controller fails.</p>

synchronous mode	<p>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.</p> <p><i>See also</i> asynchronous mode.</p>
this controller	<p>The controller that is serving your current CLI session through a local or remote terminal.</p> <p><i>See also</i> other controller.</p>
TILX	<p>Tape inline exerciser. The controller's diagnostic software to test the data transfer capabilities of tape drives in a way that simulates a high level of user activity.</p>
UBR	<p>An abbreviation used to express an unspecified bit rate. The UBR is a category of ATM service that supports connections that have no specified performance requirements.</p>
ULP	<p>Upper Layer Protocol.</p>
ULP process	<p>A function executing within a Fibre Channel node which conforms to the Upper Layer Protocol (ULP) requirements when interacting with other ULP processes.</p>
UltraNet Wizard	<p>A term used to express 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.</p>
unit	<p>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.</p>
unplanned failover	<p>As applied to the Data Replication Manager: unplanned failover is a term used to express the 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.</p> <p><i>See also</i> planned failover.</p>

VCI	An abbreviation used to express virtual channel identifier. The VCI is the field of the cell header that stores the virtual channel address.
VPI	An abbreviation used to express virtual path identifier. The VCI is the field of the cell header that stores the virtual path address.
Worldwide name or Worldwide ID	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	<p>As applied to the Data Replication Manager: write history logging is a term describing 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.</p> <p><i>See mini-merge and fast-failback</i></p>
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.

Index

Switch. See Fibre Channel gigabit switch

Symbols

%CER. See CLI Event Report

%EVL. See Event Log

%LFL. See Last Failure Log

A

ACS. See Array Controller Software

Array Controller Software 1-9

Association Sets

characteristics 2-6

FAIL_ALL switch 2-7

LOG_UNIT switch 2-9

ORDER_ALL switch 2-9

write history logging 2-8

ATM Gateway hardware requirements 1-11

autospare

definition I-2

B

BA370 enclosure 1-2, 1-3, 1-6, 4-6, 4-25, 5-2, 5-3

bad block replacement, definition I-2

BBR I-2

block, defined

see also chunk I-3

C

Cabling. See Fiber optic cable

Caution, defined xv

CCL. See Command Console LUN

chunk - definition I-3

CLI Event Report 4-9, 4-28

Cluster server 4-51

Command Console LUN 4-10, 4-29

Compaq Website xiii

Component

failures 6-5

precaution xiv

Configuring

at the initiator site 4-25

at the target site 4-6

Data Replication Manager 4-1

devices and storage sets at initiator site 4-33

devices and storage sets at target site 4-15

log units and association sets 4-41

overview 4-4

preparatory steps 4-6

saving to disk 4-40

Connections

defined 3-5

host-to-switch 3-4

switch-to-controller 3-4

target site to external fiber link 4-18, 4-37

containers

defining I-3

Controller

assigning worldwide name 4-8

- changing prompt at initiator site 4-29
- changing prompt at target site 4-10
- configuring at the initiator site 4-25
- configuring at the target site 4-6
- failure 6-5
- failure of one dual redundant member 6-6
- forced errors during copy 6-2
- operating characteristics 6-2
- read errors during copy 6-2
- setting fabric topology at initiator site 4-31
- setting fabric topology at target site 4-13
- setting mirrored write-back cache 4-12
- setting up 4-6
- status comparison 5-18

controller

- “this” and “other” defined xvi

D

- data block I-3
- Data Replication Manager
 - component failures 6-5
 - components 1-5
 - configuring 4-1
 - defined 4-2
 - enabling at initiator site 4-32
 - enabling at target site 4-14
 - required hardware and software 1-10
 - switch settings 2-11
 - troubleshooting 6-1
- Devices
 - configuring at initiator site 4-33
 - configuring at target site 4-15
- Disaster Tolerance
 - configuring overview 4-4
 - defined 2-2
 - failure modes in normal operation 6-7
 - failure notification 6-5
- Disk Administrator 4-50
- Disk drives 1-2, 1-4
- Documentation, related xvii
- Documenting your configuration 4-51
- DT. See Disaster Tolerance

- Dual redundancy 6-3
 - failure of one member 6-6
- dual-redundant controllers
 - definition I-4

E

- ECB. See External Cache Battery
- Electrostatic discharge precautions xiv
- EMU. See Environmental Monitoring Unit
- Environmental Monitoring Unit 1-3
- Error mode
 - failsafe 2-6
 - normal 2-6
- ESA12000 cabinet 1-2, 1-3, 1-4, 1-5, 1-6
- Ethernet 3-2
- Event Log 4-9, 4-12, 4-14, 4-28, 4-31, 4-33
- External Cache Battery 1-3

F

- Fabric topology 4-13, 4-31
- Failback
 - dual redundancy 6-3
 - failback procedures 5-5
 - full failback procedure 5-18
 - new hardware failback procedure 5-25
 - simple failback procedure 5-12
- failback - definition I-5
- failedset - definition I-5
- Failover
 - defined 5-4
 - failure at target site after failover 6-9
 - planned 2-10
 - planned failover procedure 5-8
 - scenarios 5-4
 - unplanned 2-10
 - unplanned failover procedure 5-16
- failover
 - definition I-5
- Failover mode. See Multiple bus failover
- Failsafe 4-40
 - lock management 6-3
 - locked condition 4-40

Failures

- at target site after failover 6-9
- both fiber optic cables or switches 6-6
- component 6-5
- controller 6-5
- network 6-7
- notification 6-5
- StorageWorks Command Console 6-6

Fiber optic cable 3-3

- 50 micron 3-4, 4-21, 4-45
- 9 micron 3-4
- connecting between initiator controllers and switches 4-35
- connecting between target controllers and switches 4-17
- connecting hosts and switches at initiator site 4-45
- connecting hosts and switches at target site 4-20
- connecting target site to external fiber link 4-18, 4-37
- failure of both fiber optic cables 6-6
- multi-mode 1-6
- single-mode 1-6

Fibre Channel gigabit switch 1-6

- defined 1-6
- failure of both switches 6-6
- setting up 3-2
- switch-to-controller connection 3-4

Forced errors 6-2**Fully-redundant power 1-7****G**

GBIC. See Gigabit Interface Converter

Getting Help xiii

Gigabit Interface Converter 1-6

- long-wave 1-3, 1-6
- short-wave 1-3, 1-6

H**Hardware**

- components 1-2

Hardware redundancy 1-2**Host**

- configuring at initiator site 4-44
- configuring at target site 4-19
- enabling access at initiator site 4-50
- host-to-switch connection 3-4
- renaming connections at initiator site 4-48
- renaming connections at target site 4-23

Host bus adapters 1-3, 1-7

- installing at initiator site 4-44
- installing at target site 4-19
- requirements 3-2
- worldwide name 3-2, 4-6

I**Initiator site**

- assigning worldwide name 4-27
- configuring controllers 4-25
- configuring devices and storage sets 4-33
- configuring host 4-44
- configuring LUNs 4-33
- connecting controllers and switches 4-35
- connecting hosts and switches 4-45
- creating remote copy sets 4-38
- disabling CCL 4-29
- enabling access to hosts 4-50
- enabling Data Replication Manager 4-32
- failure modes in normal operation 6-7
- installing host bus adapters and drivers 4-44
- installing Secure Path 4-44
- installing StorageWorks Command Console 4-45
- naming 4-25
- renaming host connections 4-48
- setting failsafe 4-40
- setting the fabric topology 4-31

L

Last Failure Log 4-9, 4-12, 4-14, 4-28, 4-31, 4-33

Link failure management 6-3

Log unit 2-9

Long distance transport modes 4-37
Long wave GBICs 4-37
LUNs
 configuring at initiator site 4-33
 configuring at target site 4-15

M
Mirrored write-back cache 4-12, 4-30
Multi-mode fiber optic cable 1-6
Multiple bus failover 4-9

O
Operation modes
 asynchronous 2-3
 synchronous 2-3
ORDER_ALL 2-9
other controller - definition I-10
Outstanding I/O settings
 asynchronous 2-4
 high outstanding I/O values 2-5
 low outstanding I/O values 2-5
 outstanding write operations 2-5
 synchronous 2-4

P
PCMCIA
 card I-1
 definition I-11
PDU. See Power Distribution Unit
Peer-to-peer remote copy 1-2
Planned failover procedure 5-8
Power
 fully-redundant 1-7
 power down 5-3
 powering up (after configuration) 5-2
Power Distribution Unit 1-6, 4-6, 4-25
Power Verification and Addressing module 1-3
Precautions
 component xiv
 electrostatic discharge xiv
program card I-1
Publication Revision History xix
Publications, related xvii

PVA. See Power Verification and Addressing module

R
Read errors 6-2
Related publications xvii
Remote copy sets
 creating 4-38
 error mode 2-6
 member failure 6-3
 operation modes 2-3
 outstanding I/O settings 2-4
 resume switch 2-5
 suspend switch 2-5
 worldwide LUN id 6-4
Renaming host connections 4-23, 4-48

S
Save Configuration 4-40
Saving controller information 4-51
SBB. See Storage Building Block
Secure Path 1-9
 installing at initiator site 4-44
 installing at target site 4-20
SET THIS_CONTROLLER_REMOTE_COPY
4-14, 4-32
SHOW_REMOTE_FULL 4-38, 4-39
Single-mode fiber optic cable 1-6
Software
 components 1-9
 requirements 1-10
Storage Building Block 1-3
storageset
 definition I-14
StorageSets
 configuring at target site 4-15, 4-33
StorageWorks Command Console 1-9
 failure 6-6
 installing at initiator site 4-45
 installing at target site 4-20
Subsystem
 final assembly 1-8

worldwide name location 4-8, 4-27
 SWCC. See StorageWorks Command Console
 Switch. See Fibre Channel gigabit switch

T

Target site

- assigning worldwide name 4-8
- configuring controllers 4-6
- configuring devices and storagesets 4-15
- configuring host 4-19
- configuring LUNs 4-15
- connecting controllers and switches 4-17
- connecting hosts and switches 4-20
- connecting to external fiber link 4-18, 4-37
- disabling CCL 4-10
- enabling Data Replication Manager 4-14
- failure after failover 6-9
- failure modes in normal operation 6-7
- installing host bus adapters and drivers 4-19
- installing Secure Path 4-20
- installing StorageWorks Command Console 4-20
- naming 4-6
- renaming host connections 4-23
- setting the fabric topology 4-13

Telephone numbers xiii

Terminal emulator session 5-18

This controller, defined xvi

this controller, defined I-15

Tip, defined xv

Troubleshooting 6-1

- component failures 6-5
- controller failure 6-5

- dual redundancy during failback 6-3
- failsafe lock management 6-3
- failure at target site after failover 6-9
- failure modes in normal operation 6-7
- failure notification 6-5
- failure of both fiber optic cables or switches 6-6
- failure of one dual redundant member 6-6
- forced errors during copy 6-2
- link failure management 6-3
- network failure 6-7
- read errors during copy 6-2
- remote copy set failure 6-3
- remote copy set worldwide LUN id 6-4
- StorageWorks Command Console failure 6-6

U

unit, defined I-15

W

Warning, defined xv

Windows NT 4-51

Worldwide LUN id

- for remote copy sets 6-4

Worldwide name

- assigning subsystem worldwide name to controller 4-8, 4-27
- location on host bus adapter 3-2
- location on subsystem 4-8, 4-27

Write history logging

- fast-failback 2-8
- log unit restrictions 2-8
- mini-merge 2-8

