Tru64 UNIX and TruCluster Server

Hardware Configuration Technical Update for Fibre Channel

February 2001

Product Version: TruCluster Server Version 5.1

Operating System and Version: Tru64 UNIX Version 5.1

This technical update provides Fibre Channel information to update Chapter 6 of the Version 5.1 TruCluster™ Server *Hardware Configuration* manual. This document describes how to set up shared Fibre Channel mass storage devices for an example Compaq Tru64™ UNIX Version 5.1 or TruCluster Server Version 5.1 installation.

Technical updates are used as the distribution media to keep Fibre Channel documentation up-to-date with changes in supported Fibre Channel hardware and changes in technology.

Compaq Computer Corporation
Houston, Texas

© 2001 Compaq Computer Corporation

Compaq, the Compaq logo, AlphaServer, StorageWorks, and TruCluster Registered in U.S. Patent and Trademark Office. Alpha and Tru64 are trademarks of Compaq Information Technologies Group, L.P. in the United States and other countries.

Windows and Windows NT are trademarks of Microsoft Corporation in the United States and other countries. UNIX is a trademark of The Open Group in the United States and other countries. Legato NetWorker is a registered trademark of Legato Systems, Inc., in the U.S. 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.

Contents

5 Installing Tru64 UNIX and TruCluster Server Software Using Fibre Channel Storage

Tables

About This Technical Update

This technical update provides updated information to Chapter 6 (Fibre Channel) of the TruCluster™ Server *Hardware Configuration* Version 5.1 manual. This technical update is provided to make the information more readily available and easier to update.

This information is valid for Fibre Channel fabric and arbitrated loop topologies with Tru64™ UNIX Version 5.1 and TruCluster Server Version 5.1.

This technical update provides an overview of Fibre Channel, Fibre Channel configuration examples, and information on Fibre Channel hardware installation and configuration in a Tru64 UNIX Version 5.1 or TruCluster Server Version 5.1 configuration.

The information includes how to determine the `/dev/disk/dskn` value that corresponds to the Fibre Channel storagesets that have been set up as the Tru64 UNIX boot disk, cluster root (`/`), cluster `/usr`, cluster `/var`, cluster member boot, and quorum disks, and how to set the `bootdef_dev` console environment variable to facilitate installation of Tru64 UNIX Version 5.1 and TruCluster Server Version 5.1.

Note

TruCluster Server Version 5.1 configurations require one or more disks to hold the Tru64 UNIX operating system. The disks are either private disks on the system that will become the first cluster member, or disks on a shared bus that the system can access.

Whether or not you install the base operating system on a shared disk, always shut down the cluster before booting the Tru64 UNIX disk.

Audience

This technical update is for system administrators who are responsible for TruCluster Server Version 5.1 configurations.

New and Changed Features

This manual incorporates the following changes since the last version of the manual:

- Removed the restriction of four cascaded switches from the requirements and restrictions section.
- Modified *Section 3.8* to provide a reference to the *Compaq StorageWorks Heterogeneous Open SAN Design Reference Guide*. See this document for information regarding SAN topologies and restrictions.
- Added support for the DS-DSGGC-AB Fibre Channel switch.
- Made minor corrections and additions in Chapter 4 and Chapter 5.

Organization

This technical update is organized as follows:

Chapter 1	Provides a procedure to follow when using Fibre Channel storage for the Tru64 UNIX Version 5.1 operating system or TruCluster Server Version 5.1 installation.
Chapter 2	Provides information about the supported Fibre Channel hardware and the configuration requirements and restrictions.
Chapter 3	Provides a basic overview of Fibre Channel terminology and topologies, plus examples of some Fibre Channel cluster configurations.
Chapter 4	Provides information about installing the Fibre Channel hardware needed for a TruCluster Server configuration accessing storage over the Fibre Channel hardware.
Chapter 5	Provides the steps necessary to install the Tru64 UNIX operating system and TruCluster Server software on Fibre Channel storage.
Chapter 6	Provides the steps needed to convert an HSG80 RAID array controller from transparent to multiple-bus failover mode.
Chapter 7	Describes how to use the `emxmgr` utility to display information about the Fibre Channel adapters and Fibre Channel topology.
Appendix A	Provides a blank table to use during base operating system and cluster installation to keep a record of device names, worldwide names, and `/dev/disk/dskn` names.

Purpose of This Document

This document outlines the procedures to install the Tru64 UNIX operating system and TruCluster Server software on Fibre Channel storage in a fabric or loop topology.

Related Documentation

Consult the following Tru64 UNIX Version 5.1 manual to assist you in installing Tru64 UNIX:

- *Installation Guide*

Consult the following TruCluster Server manuals for assistance in cluster installation, administration, and programming tasks:

- *TruCluster Server Software Product Description* (SPD) — The comprehensive description of the TruCluster Server product. You can find the latest version of the SPD at the following URL:

  `http://www.tru64unix.compaq.com/faqs/publications/pub_page/spds.html`

- *Release Notes* — Provides important information about TruCluster Server Version 5.1.
- *Technical Overview* — Provides an overview of the TruCluster Server technology.
- *Software Installation* — Describes how to install the TruCluster Server Version 5.1 product.
- *Cluster Administration* — Describes cluster-specific administration tasks.
- *Highly Available Applications* — Describes how to deploy applications on a TruCluster Server cluster.
- *Hardware Configuration* — Describes how to set up the processors that are to become cluster members, and how to configure cluster shared storage.

You can find the latest version of the TruCluster Server documentation at the following URL:

`http://www.tru64unix.compaq.com/faqs/publications/pub_page/cluster_list.html`

Consult the following documentation to assist you in Fibre Channel storage configuration and administration:

- *KGPSA-BC PCI-to-Optical Fibre Channel Host Adapter User Guide*
- *64-Bit PCI-to-Fibre Channel Host Bus Adapter User Guide*
- *Compaq StorageWorks Fibre Channel Storage Switch User's Guide*
- *Compaq StorageWorks SAN Switch 8 Installation and Hardware Guide*
- *Compaq StorageWorks SAN Switch 16 Installation and Hardware Guide*
- *Compaq StorageWorks Fibre Channel SAN Switch 8-EL Installation and Hardware Guide*

- *Compaq StorageWorks Fibre Channel SAN Switch 16-EL Installation and Hardware Guide*
- *Compaq StorageWorks Fibre Channel SAN Switch Management Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Configuration Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*
- *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Compaq Tru64 UNIX Installation and Configuration Guide*
- *Compaq StorageWorks HSG60/HSG80 Array Controller ACS Version 8.5 Maintenance and Service Guide*
- *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Compaq Tru64 UNIX Installation and Configuration Guide*
- *Compaq StorageWorks Release Notes RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX*
- *Compaq StorageWorks Modular Array Configuration Guide*
- *Model 2100 and 2200 Ultra SCSI Controller Enclosures User Guide*
- *Compaq StorageWorks Enclosure 4200 Family LVD Disk Enclosure User Guide*
- *Wwidmgr User's Manual*
- *Heterogeneous Storage Area Networks Application Note*
- *RA8000/ESA12000 FC-Switch Configurations for Tru64 UNIX Application Note*
- *Fibre Channel Storage Hub 7 Installation Guide*
- *Fibre Channel Storage Hub 7 Rack Mounting Installation Card*
- *Enterprise Backup Solution with Legato NetWorker User Guide*
- *Compaq StorageWorks Heterogeneous Open SAN Design Reference Guide*

1

Overview of Installation Using Fibre Channel Disks

This chapter provides a procedure for installing Tru64 UNIX Version 5.1 or TruCluster Server Version 5.1 using Fibre Channel disks in a fabric or loop topology.

1.1 Procedure for Installation Using Fibre Channel Disks

Use the following procedure to install Tru64 UNIX Version 5.1 and TruCluster Server Version 5.1 using Fibre Channel disks. If you are only installing Tru64 UNIX Version 5.1, complete the first eight steps. Complete all the steps for a TruCluster Server Version 5.1 installation. See the Tru64 UNIX *Installation Guide*, TruCluster Server *Software Installation* manual, and other hardware manuals as appropriate for the actual installation procedures.

1. Install the Fibre Channel switch or hub (Section 4.1.1 or Section 4.1.2).
2. Install the KGPSA PCI-to-Fibre Channel host bus adapter (Section 4.1.3).
3. Set up the HSG80 RAID array controllers for a fabric or loop configuration (Section 4.1.4).
4. Configure the HSG80 disks to be used for base operating system and cluster installation. Be sure to set the identifier for each storage unit you will use for operating system or cluster installation (Section 5.1.1 and Section 5.1.2).
5. If the system is not already powered on, power on the system where you will install Tru64 UNIX Version 5.1. If this is a cluster installation, this system will also be the first cluster member.

   Use the console WWID manager (`wwidmgr`) utility to set the device unit number for the Fibre Channel Tru64 UNIX Version 5.1 disk and first cluster member system boot disks (Section 5.1.3).
6. Use the `show wwid*` and `show n*` console commands to show the disk devices that are currently reachable, and the paths to the devices (Section 5.1.4).

7. Use the WWID manager to set the `bootdef_dev` console environment variable for the system where you will install the Tru64 UNIX operating system (Section 5.1.5).
8. See the Tru64 UNIX *Installation Guide* and install the base operating system from the CD-ROM. The installation procedure will recognize the disks for which you set the device unit number. Select the disk you have chosen as the Tru64 UNIX operating system installation disk from the list of disks provided (Section 5.2).

   After the new kernel has booted to multi-user mode, complete the operating system installation.

   If you will not be installing TruCluster Server software, reset the `bootdef_dev` console environment variable to provide multiple boot paths to the boot disk (Section 5.6), then boot the operating system.
9. Determine the `/dev/disk/dsk`*n* values to be used for cluster installation (Section 5.3).
10. Use the `disklabel` utility to label the disks used to create the cluster (Section 5.4).
11. See the TruCluster Server *Software Installation* manual and install the TruCluster Server software subsets then run the `clu_create` command to create the first cluster member. Do not allow `clu_create` to boot the system. Shut down the system to the console prompt (Section 5.5).
12. Reset the `bootdef_dev` console environment variable to provide multiple boot paths to the cluster member boot disk (Section 5.6). Boot the first cluster member.
13. See the TruCluster Server *Software Installation* manual and add subsequent cluster member systems (Section 5.7). As with the first cluster member, you will have to:
    - Use the `wwidmgr` command to set the device unit number for the member system boot disk.
    - Set the `bootdef_dev` environment variable.
    - Reset the `bootdef_dev` environment variable after building a kernel on the new cluster member system.

2

Fibre Channel Requirements and Restrictions

This chapter describes the Fibre Channel hardware requirements and restrictions for Tru64 UNIX Version 5.1 and TruCluster Server Version 5.1.

For the latest information about supported hardware, see the AlphaServer™ options list for your system at the following URL:
`http://www.compaq.com/alphaserver/products/options.html`

2.1 Fibre Channel Requirements and Restrictions

Table 2–1 lists the supported AlphaServer systems with Fibre Channel and the number of KGPSA-BC or KGPSA-CA PCI-to-Fibre Channel adapters supported on each system.

Table 2–1: AlphaServer Systems Supported for Fibre Channel

AlphaServer	Number of Adapters Supported in Fabric Topology	Number of Adapters Supported in Loop Topology
AlphaServer 800	2	—
AlphaServer 1200	4	—
AlphaServer 4000, 4000A, or 4100	4	—
AlphaServer DS10	2	2[a]
AlphaServer DS20 and DS20E	4	2[a]
AlphaServer ES40	4	2[a]
AlphaServer 8200 or 8400[b]	63[c], 32[d]	—
AlphaServer GS60, GS60E, and GS140[b]	63[c], 32[d]	—
AlphaServer GS80, GS160, and GS320[e]	26	—

[a] The arbitrated loop topology requires the KGPSA-CA adapter with V3.03 (or later) firmware and Version 5.8 or later of the SRM console.
[b] The KGPSA-BC/CA PCI-to-Fibre Channel adapters are only supported on the DWLPB PCIA option; they are not supported on the DWLPA.
[c] The 8200, 8400, GS60, GS60E, and GS140 AlphaServers support up to 63 KGPSAs in a standalone configuration.
[d] The 8200, 8400, GS60, GS60E, and GS140 AlphaServers support up to 32 KGPSAs in a cluster configuration.
[e] The GS80, GS160, and GS320 systems support only the KGPSA-CA Fibre Channel host bus adapter.

The following requirements and restrictions apply to the use of Fibre Channel with TruCluster Server Version 5.1:

- The HSG60 and HSG80 require Array Control Software (ACS) Version 8.5.
- Eight member systems may be connected to common storage over Fibre Channel in a fabric (switch) configuration. A maximum of two member systems is supported in arbitrated loop configurations.
- The Fibre Channel RAID Array 8000 (RA8000) midrange departmental storage subsystem and Fibre Channel Enterprise Storage Array 12000 (ESA12000) house two HSG80 dual-channel controllers. There are provisions for six UltraSCSI channels. A maximum of 72 disks is supported.
- The StorageWorks™ Modular Array 6000 (MA6000) supports dual-redundant HSG60 controllers and the 1-inch universal drives. For more information on configuring the MA6000, see the following manuals:
  - *MA6000 HSG60 Array Controller ACS Version 8.5 Solution Software for Compaq Tru64 UNIX Installation and Configuration Guide*
  - *Compaq StorageWorks Modular Array Configuration Guide*
  - *Model 2100 and 2200 Ultra SCSI Controller Enclosures User Guide*
  - *Compaq StorageWorks Enclosure 4200 Family LVD Disk Enclosure User Guide*
- The StorageWorks Modular Array 8000 (MA8000) and Enterprise Modular Array 12000 (EMA12000) support dual redundant HSG80 controllers and 1-inch universal drives.

  For more information on configuring the MA8000/EMA12000, see the following manuals:
  - *Compaq StorageWorks Release Notes RA8000/ESA12000 and MA8000/EMA12000 Solution Software V8.5b for Tru64 UNIX*
  - *Compaq StorageWorks Modular Array Configuration Guide*.
  - *Model 2100 and 2200 Ultra SCSI Controller Enclosures User Guide*
  - *Compaq StorageWorks Enclosure 4200 Family LVD Disk Enclosure User Guide*
- The HSG60 or HSG80 Fibre Channel array controller support only disk devices.
- The only supported Fibre Channel adapters are the KGPSA-BC and KGPSA-CA PCI-to-Fibre Channel host bus adapters. The KGPSA-BC adapter is supported in fabric configurations only; the KGPSA-CA adapter is supported in either fabric and arbitrated loop configurations.

- The KGPSA-BC/CA PCI-to-Fibre Channel adapters are only supported on the DWLPB PCIA option; they are not supported on the DWLPA.
- The only supported Fibre Channel hub is the 7-port DS-SWXHB-07. The DS-SWXHB-07 has clock and data recovery on each port. It also features Gigabit Interface Converter (GBIC) transceiver-based port connections for maximum application flexibility. The hub is hot pluggable and is unmanaged.
- Only single-hub arbitrated loop configurations are supported; that is, there are no cascaded hubs on any SCSI bus.
- The only Fibre Channel switches supported are the DS-DSGGA-AA/AB 8/16 port, DS-DSGGB-AA/AB 8/16 port, or DS-DSGGC-AA/AB 8/16 port Fibre Channel switches.
- The DSGGA, DSGGB, and DSGGC Fibre Channel switches and the DS-SWXHB-07 hub support both shortwave (GBIC-SW) and longwave (GBIC-LW) Gigabit Interface Converter (GBIC) modules. Seven of the eight DSGGC-AA ports are fixed shortwave optical transceivers. Only one DSGGC-AA port is configured as a removable GBIC. It may be shortwave or longwave.

  The GBIC-SW module supports 50-micron, multimode fiber cables with the standard subscriber connector (SC) connector in lengths up to 500 meters. It also supports 62.5-micron multimode fiber cables in lengths up to 200 meters. The GBIC-LW supports 9-micron, single-mode fiber cables with the SC connector in lengths up to 10 kilometers.

  The KGPSA-BC/CA PCI-to-Fibre Channel host bus adapters and the HSG60 and HSG80 RAID controller support the 50-micron Gigabit Link Module (GLM) for fiber connections. Therefore, only the 50-micron multimode fiber optical cable is supported between the KGPSA and switch (or hub) and the switch (or hub) and HSG60 or HSG80 for cluster configurations. You must install GBIC-SW GBICs in the Fibre Channel switches (or hub) for communication between the switches (or hub) and KGPSA or HSG60/HSG80.
- Tru64 UNIX Version 5.1 allows up to 255 Fibre Channel targets. An active host port or host bus adapter constitutes a target.
- Tru64 UNIX Version 5.1 allows up to 255 LUNs per target.
- The HSG60 and HSG80 supports transparent and multiple-bus failover mode when used in a TruCluster Server Version 5.1 configuration. Multiple-bus failover is recommended.
- A storage array with dual-redundant HSG60 or HSG80 controllers in transparent mode failover is two targets and consumes four ports on a switch. Transparent mode is recommended only while upgrading from

Tru64 UNIX Version 4.*x*. After the upgrade is complete, you should switch to multiple-bus failover.

- A storage array with dual-redundant HSG60 or HSG80 controllers in multiple-bus failover is four targets and consumes four ports on a switch.
- The HSG60 and HSG80 documentation refers to the controllers as Controllers A (top) and B (bottom). Each controller provides two ports (left and right). (The HSG60 and HSG80 documentation refers to these ports as Port 1 and 2, respectively.) In transparent failover mode, only one left port and one right port are active at any given time.

  With transparent failover enabled, assuming that the left port of the top controller and the right port of the bottom controller are active, if the top controller fails in such a way that it can no longer properly communicate with the switch, then its functions will fail over to the bottom controller (and vice versa).
- In transparent failover mode, you can configure which controller presents each HSG60 or HSG80 storage element (unit) to the cluster. Ordinarily, the connections on port 1 (left port) have a default unit offset of 0, and units designated D0 through D99 are accessed through port 1 of either controller. The connections on port 2 (right port) have a default unit offset of 100, and units designated D100 through D199 are accessed through port 2 of either controller.
- In multiple-bus failover mode, the connections on all ports have a default unit offset of 0, and all units (D0 through D199) are visible to all host ports, but accessible only through one controller at any specific time. The host can control the failover process by moving units from one controller to the other controller.
- The Fibre Channel Tape Controller, Fibre Channel Tape Controller II, TL891, TL895, and ESL9326D are supported on a Fibre Channel storage bus. For more information, see the *Enterprise Backup Solution with Legato NetWorker User Guide*. Legato Networker Version 6.0 is required for application failover.
- Tapes are single-stream devices. There is no load balancing of I/O requests over the available paths to the tape devices. The first available path to the tape devices is selected for I/O.

3

Fibre Channel Overview

This chapter provides an overview of Fibre Channel and Fibre Channel configuration examples.

3.1 Overview

Fibre Channel supports multiple protocols over the same physical interface. Fibre Channel is primarily a protocol-independent transport medium; therefore, it is independent of the function for which you use it.

TruCluster Server uses the Fibre Channel Protocol (FCP) for SCSI to use Fibre Channel as the physical interface.

Fibre Channel, with its serial transmission method, overcomes the limitations of parallel SCSI by providing:

- Data rates of 100 MB/sec, 200 MB/sec, and 400 MB/sec
- Support for multiple protocols
- Better scalability
- Improved reliability, serviceability, and availability

Fibre Channel uses an extremely high-transmit clock frequency to achieve the high data rate. Using optical fiber transmission lines allows the high-frequency information to be sent up to 40 km (24.85 mi), the maximum distance between transmitter and receiver. Copper transmission lines may be used for shorter distances.

3.2 Basic Fibre Channel Terminology

The following list describes the basic Fibre Channel terminology:

AL_PA The Arbitrated Loop Physical Address (AL_PA) is used to address nodes on the Fibre Channel loop. When a node is ready to transmit data, it transmits Fibre Channel primitive signals that include its own identifying AL_PA.

Arbitrated Loop A Fibre Channel topology in which frames are routed around a loop set up by the links between the nodes in the loop. All nodes in a loop share the

bandwidth, and bandwidth degrades slightly as nodes and cables are added.

Frame — All data is transferred in a packet of information called a frame. A frame is limited to 2112 bytes. If the information consists of more than 2112 bytes, it is divided up into multiple frames.

Node — The source and destination of a frame. A node may be a computer system, a redundant array of independent disks (RAID) array controller, or a disk device. Each node has a 64-bit unique node name (worldwide name) that is built into the node when it is manufactured.

N_Port — Each node must have at least one Fibre Channel port from which to send or receive data. This node port is called an N_Port. Each port is assigned a 64-bit unique port name (worldwide name) when it is manufactured. An N_Port is connected directly to another N_Port in a point-to-point topology. An N_Port is connected to an F_Port in a fabric topology.

NL_Port — In an arbitrated loop topology, information is routed around a loop. A node port that can operate on the loop is called an NL_Port (node loop port). The information is repeated by each NL_Port until it reaches its destination. Each port has a 64-bit unique port name (worldwide name) that is built into the node when it is manufactured.

Fabric — A switch, or multiple interconnected switches, that route frames between the originator node (transmitter) and destination node (receiver).

F_Port — The ports within the fabric (fabric port). This port is called an F_port. Each F_port is assigned a 64-bit unique node name and a 64-bit unique port name when it is manufactured. Together, the node name and port name make up the worldwide name.

FL_Port — An F_Port containing the loop functionality is called an FL_Port.

Link	The physical connection between an N_Port and another N_Port or an N_Port and an F_Port. A link consists of two connections, one to transmit information and one to receive information. The transmit connection on one node is the receive connection on the node at the other end of the link. A link may be optical fiber, coaxial cable, or shielded twisted pair.
E_Port interswitch expansion port	An expansion port on a switch used to make a connection between two switches in the fabric.

3.3 Fibre Channel Topologies

Fibre Channel supports three different interconnect topologies:

- Point-to-point (Section 3.3.1)
- Fabric (Section 3.3.2)
- Arbitrated loop (Section 3.3.3)

Note

Although you can interconnect an arbitrated loop with fabric, hybrid configurations are not supported at the present time, and therefore are not discussed in this manual.

3.3.1 Point-to-Point

The point-to-point topology is the simplest Fibre Channel topology. In a point-to-point topology, one N_Port is connected to another N_Port by a single link.

Because all frames transmitted by one N_Port are received by the other N_Port, and in the same order in which they were sent, frames require no routing.

Figure 3–1 shows an example point-to-point topology.

Figure 3–1: Point-to-Point Topology

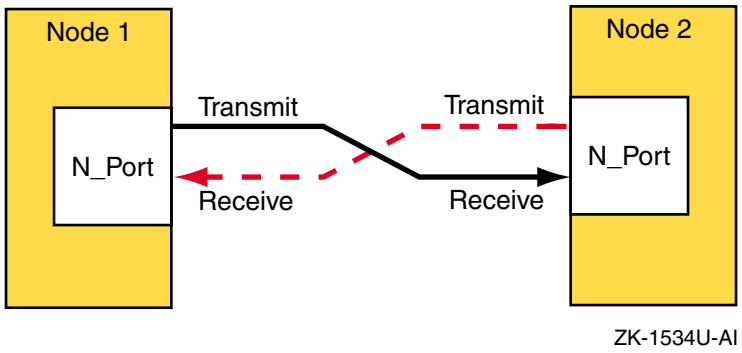

3.3.2 Fabric

The fabric topology provides more connectivity than point-to-point topology. The fabric topology can connect up to 2^{24} ports.

The fabric examines the destination address in the frame header and routes the frame to the destination node.

A fabric may consist of a single switch, or there may be several interconnected switches (up to three interconnected switches are supported). Each switch contains two or more fabric ports (F_Port) that are internally connected by the fabric switching function, which routes the frame from one F_Port to another F_Port within the switch. Communication between two switches is routed between two expansion ports (E_Ports).

When an N_Port is connected to an F_Port, the fabric is responsible for the assignment of the Fibre Channel address to the N_Port attached to the fabric. The fabric is also responsible for selecting the route a frame will take, within the fabric, to be delivered to the destination.

When the fabric consists of multiple switches, the fabric can determine an alternate route to ensure that a frame gets delivered to its destination.

Figure 3–2 shows an example fabric topology.

Figure 3–2: Fabric Topology

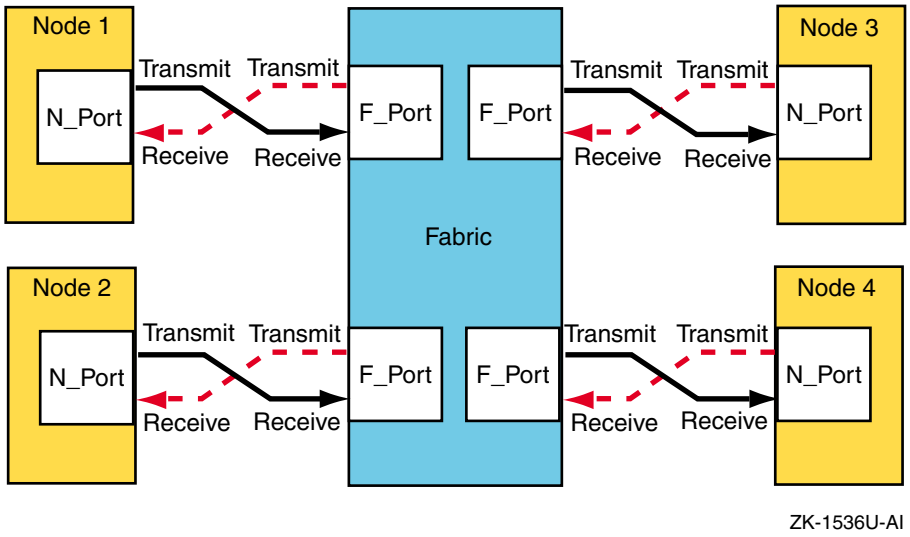

3.3.3 Arbitrated Loop Topology

In an arbitrated loop topology, frames are routed around a loop set up by the links between the nodes. The hub maintains loop continuity by bypassing a node when the node or its cabling fails, when the node is powered down, or when the node is removed for maintenance. The hub is transparent to the protocol. It does not consume any Fibre Channel arbitrated loop addresses so it is not addressable by a Fibre Channel arbitrated loop port.

The nodes arbitrate to gain control (become master) of the loop. After a node becomes master, the nodes select (by way of setting bits in a bitmask) their own Arbitrated Loop Physical Address (AL_PA). The AL_PA is used to address nodes on the loop. The AL_PA is dynamic and can change each time the loop is initialized, a node is added or removed, or at any other time that an event causes the membership of the loop to change. When a node is ready to transmit data, it transmits Fibre Channel primitive signals that include its own identifying AL_PA.

In the arbitrated loop topology, a node port is called an NL_Port (node loop port), and a fabric port is called an FL_Port (fabric loop port).

Figure 3–3 shows an example of an arbitrated loop topology.

Figure 3–3: Arbitrated Loop Topology

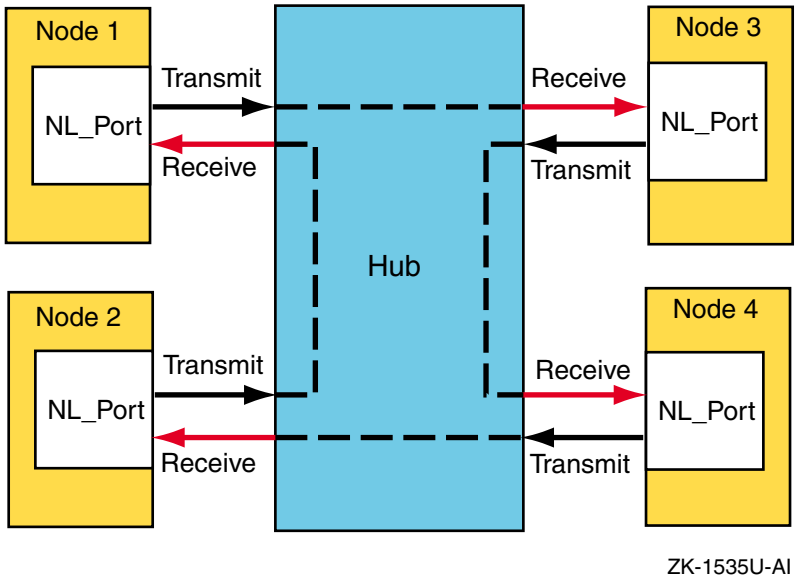

3.4 Fibre Channel Topology Comparison

This section compares and contrasts the fabric and arbitrated loop topologies and describes why you might choose to use them.

When compared with the fabric (switched) topology, arbitrated loop is a lower cost, and lower performance, alternative. Arbitrated loop reduces Fibre Channel cost by substituting a lower-cost, often nonintelligent and unmanaged hub, for a more expensive switch. The hub operates by collapsing the physical loop into a logical star. The cables, associated connectors, and allowable cable lengths are similar to those of a fabric. Arbitrated loop supports a theoretical limit of 127 nodes in a loop. Arbitrated loop nodes are self-configuring and do not require Fibre Channel address switches.

Arbitrated loop provides reduced cost at the expense of bandwidth; all nodes in a loop share the bandwidth (100 MB/sec per loop), and bandwidth degrades slightly as nodes and cables are added. Nodes on the loop see all traffic on the loop, including traffic between other nodes. The hub can include port-bypass functions that manage movement of nodes on and off the loop. For example, if the port bypass logic detects a problem, the hub can remove that node from the loop without intervention. Data availability is then preserved by preventing the down time associated with node failures, cable disconnections, and network reconfigurations. However, traffic caused by node insertion and removal, errors, and so forth, can cause temporary disruption on the loop.

Although the fabric topology is more expensive, it provides both increased connectivity and higher performance; switches provide a full-duplex 100 (200) MB/sec point-to-point connection to the fabric. Switches also provide improved performance and scaling because nodes on the fabric see only data destined for themselves, and individual nodes are isolated from reconfiguration and error recovery of other nodes within the fabric. Switches can provide management information about the overall structure of the Fibre Channel fabric, which may not be the case for an arbitrated loop hub.

Table 3–1 presents a comparison between the fabric and arbitrated loop topologies.

Table 3–1: Fibre Channel Fabric and Arbitrated Loop Comparison

When to use Arbitrated Loop	When to use Fabric
In clusters of up to two members	In clusters of more than two members
In applications where low total solution cost and simplicity are key requirements	In multinode cluster configurations when possible temporary traffic disruption due to reconfiguration or repair is a concern
In applications where the shared bandwidth of an arbitrated loop configuration is not a limiting factor	In high bandwidth applications where a shared arbitrated loop topology is not adequate
In configurations where expansion and scaling are not anticipated	In cluster configurations where expansion is anticipated and requires performance scaling

3.5 Example Fibre Channel Configurations Supported by TruCluster Server

This section provides diagrams of some of the configurations supported by TruCluster Server Version 5.1. Diagrams are provided for both transparent failover mode and multiple-bus failover mode.

3.5.1 Fibre Channel Cluster Configurations for Transparent Failover Mode

With transparent failover mode:

- The hosts do not know a failover has taken place (failover is transparent to the hosts).
- The units are divided between an HSG80 port 1 and port 2.
- If there are dual-redundant HSG80 controllers, controller A port 1 and controller B port 2 are normally active; controller A port 2 and controller B port 1 are normally passive.

- If one controller fails, the other controller takes control and both its ports are active.

Figure 3–4 shows a typical Fibre Channel cluster configuration using transparent failover mode.

Figure 3–4: Fibre Channel Single Switch Transparent Failover Configuration

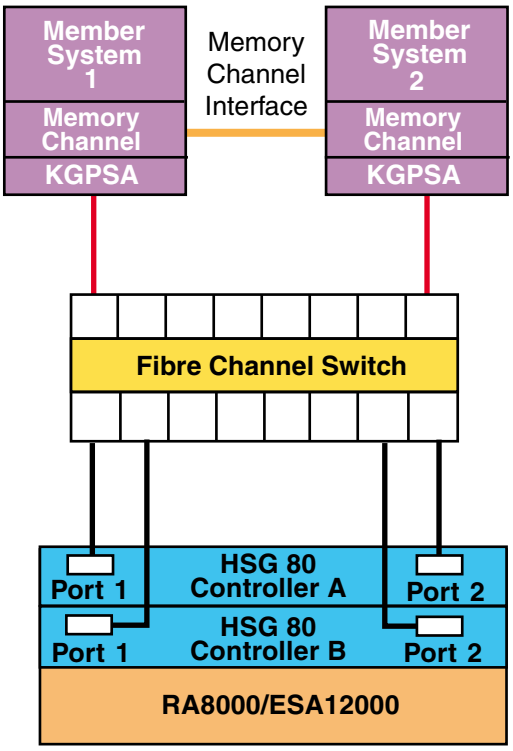

ZK-1531U-AI

In transparent failover, units D00 through D99 are accessed through port 1 of both controllers. Units D100 through D199 are accessed through port 2 of both HSG80 controllers.

You cannot achieve a no-single-point-of-failure (NSPOF) configuration using transparent failover. The host cannot initiate failover, and if you lose a host bus adapter, switch or hub, or a cable, you lose the units behind at least one port.

You can, however, add the hardware for a second bus (another KGPSA, switch, and RA8000/ESA12000 with associated cabling) and use LSM to mirror across the buses. However, because you cannot use LSM to mirror the cluster root (/) file system, member boot partitions, the quorum disk,

or swap partitions you cannot obtain an NSPOF transparent failover configuration, even though you have increased availability.

Figure 3–5 shows a two-node Fibre Channel cluster with a single RA8000 or ESA12000 storage array with dual-redundant HSG80 controllers and an DS-SWXHB-07 Fibre Channel hub.

Figure 3–5: Arbitrated Loop Configuration with One Storage Array

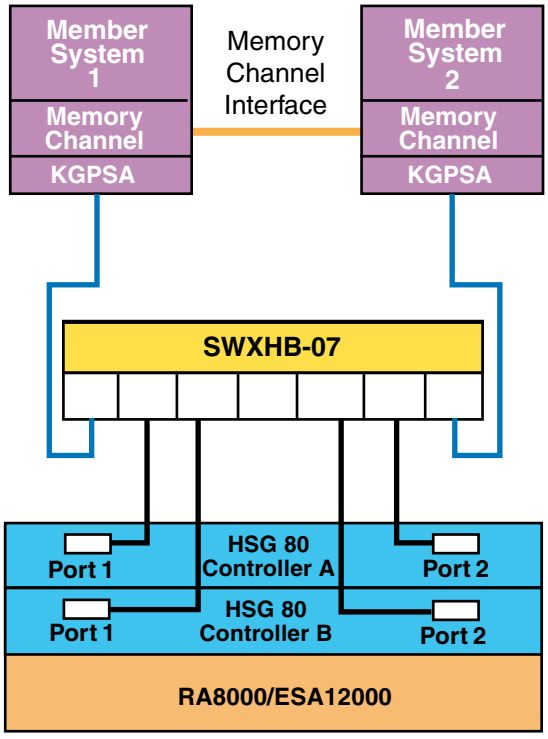

3.5.2 Fibre Channel Cluster Configurations for Multiple-Bus Failover Mode

With multiple-bus failover:

- The host controls the failover by accessing units over a different path or causing the access to the unit to be through the other HSG80 controller.
- An active controller causes a failover to the other controller if the controller recognizes the loss of the switch, hub, or cable to a controller port.
- Each cluster member system has two or more (fabric only) KGPSA host bus adapters (multiple paths to the storage units).

- Normally, all available units (D0 through D199) are available at all host ports. Only one HSG80 controller will be actively doing I/O for any particular storage unit.

  However, both controllers can be forced active by preferring units to one controller or the other (`SET` *`unit`* `PREFERRED_PATH=THIS`). By balancing the preferred units, you can obtain the best I/O performance using two controllers.

  Note

  If you have preferred units, and the HSG80 controllers restart because of an error condition or power failure, and one controller restarts before the other controller, the HSG80 controller restarting first will take all the units, whether they are preferred or not. When the other HSG80 controller starts, it will not have access to the preferred units, and will be inactive.

  Therefore, you want to ensure that both HSG80 controllers start at the same time under all circumstances so that the controller sees its own preferred units.

Figure 3–6 and Figure 3–7 show two different recommended multiple-bus NSPOF cluster configurations. The only difference is the fiber-optic cable connection path between the switch and the HSG80 controller ports.

There is no difference in performance between these two configurations. It may be easier to cable the configuration shown in Figure 3–6 because the cables from one switch (or switch zone) both go to the ports on the same side of both controllers (for example, port 1 of both controllers).

Figure 3–6: Multiple-Bus NSPOF Configuration Number 1

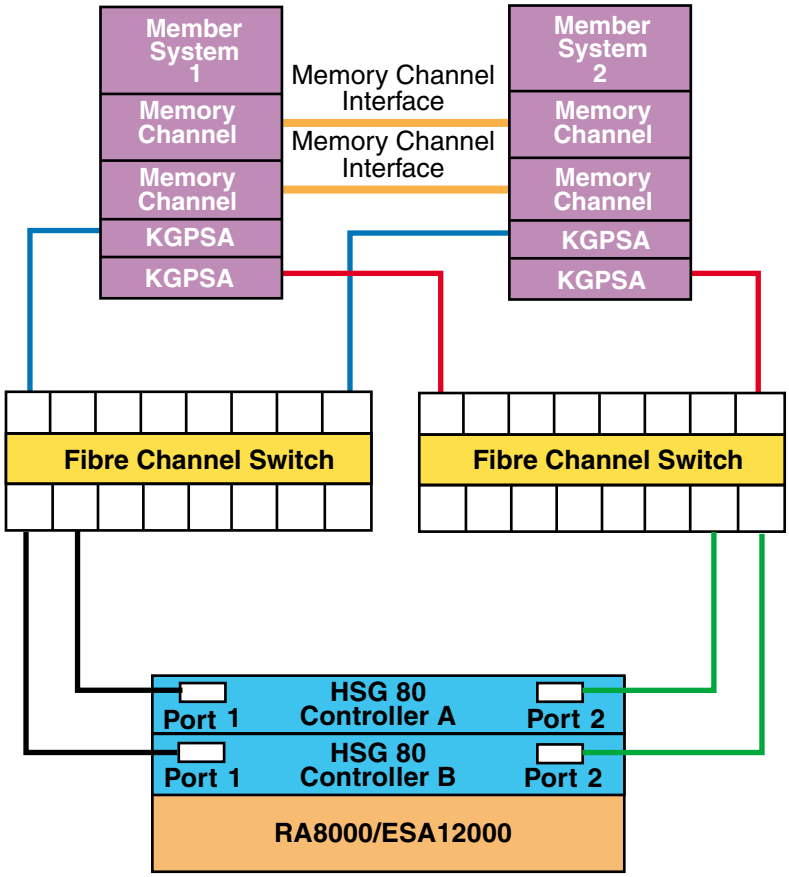

Figure 3–7: Multiple-Bus NSPOF Configuration Number 2

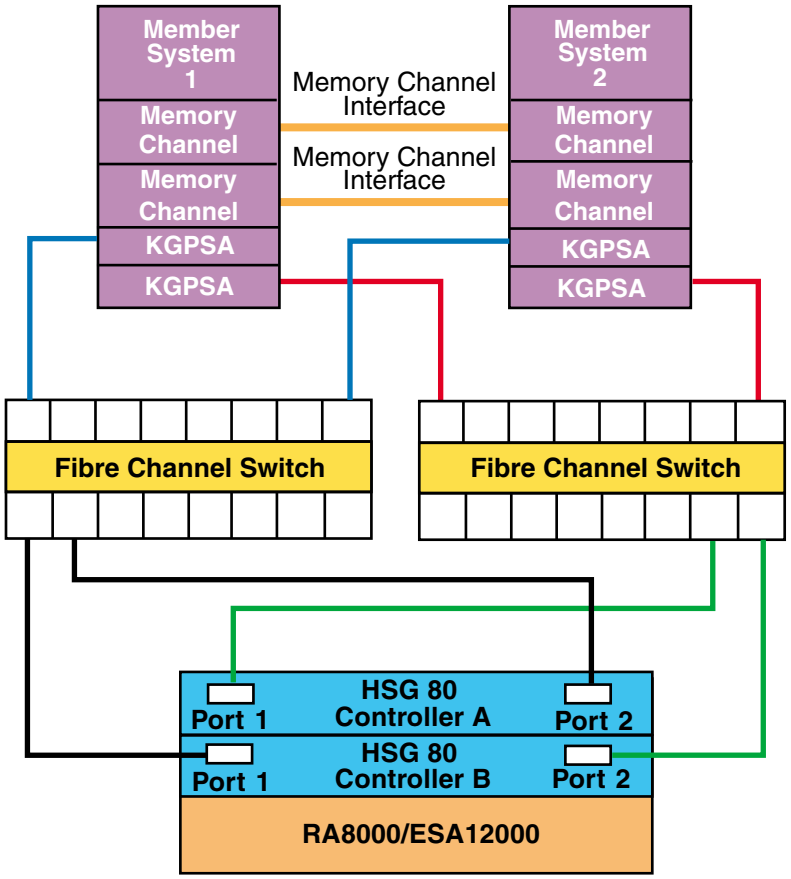

The configuration shown in Figure 3–8 is not a recommended cluster configuration because it has multiple single-points-of-failure that can cause a loss of access to storage.

Figure 3–8, in a single-system configuration, is not an NSPOF configuration.

Note

Previous documentation erroneously showed Figure 3–8 as an NSPOF configuration.

If you have a configuration like the one shown in Figure 3–8, change the switch to HSG80 cabling to match the configurations shown in Figure 3–6 or Figure 3–7.

The single-system configuration shown in Figure 3–9 is also a configuration that is not recommended. It also has a single point of failure.

Note

If you have a configuration like the one shown in Figure 3–9, convert it to a recommended configuration by adding the following two cables:

- From the left-hand switch to controller B Port 2.
- From the right-hand switch to controller A Port 2.

Figure 3–8: A Configuration That Is Not Recommended

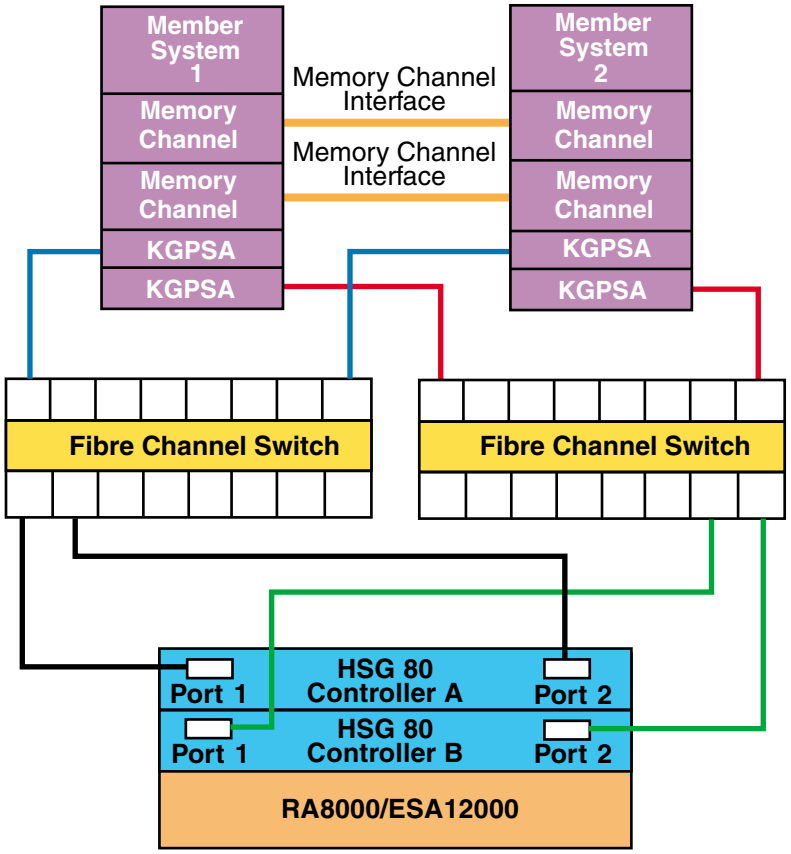

Figure 3–9: Another Configuration That Is Not Recommended

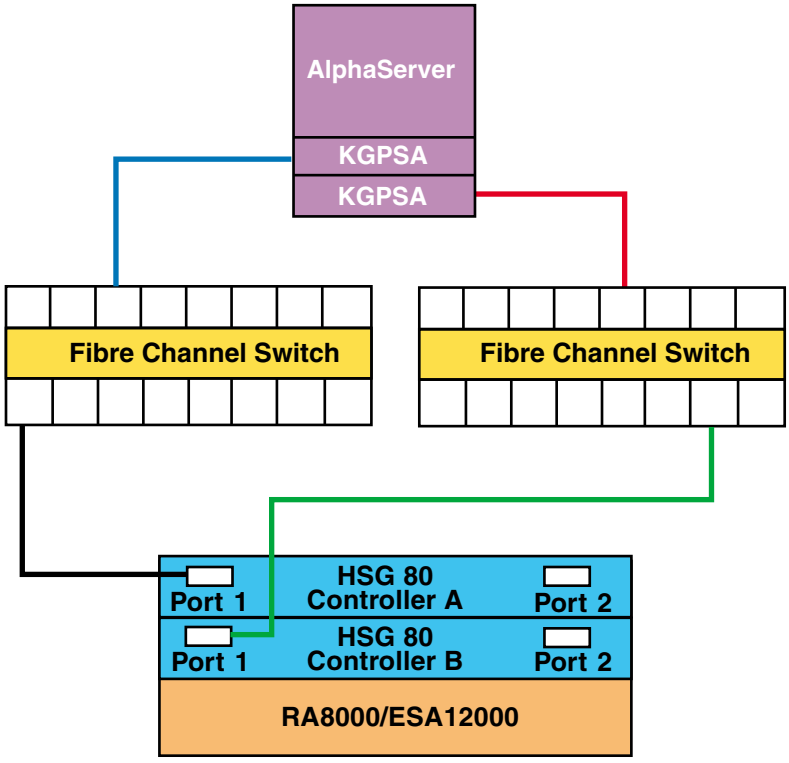

Figure 3–10 shows the maximum supported arbitrated loop configuration of a two-node Fibre Channel cluster with two RA8000 or ESA12000 storage arrays, each with dual-redundant HSG80 controllers and two DS-SWXHB-07 Fibre Channel hubs. This provides a NSPOF configuration.

Figure 3–10: Arbitrated Loop Maximum Configuration

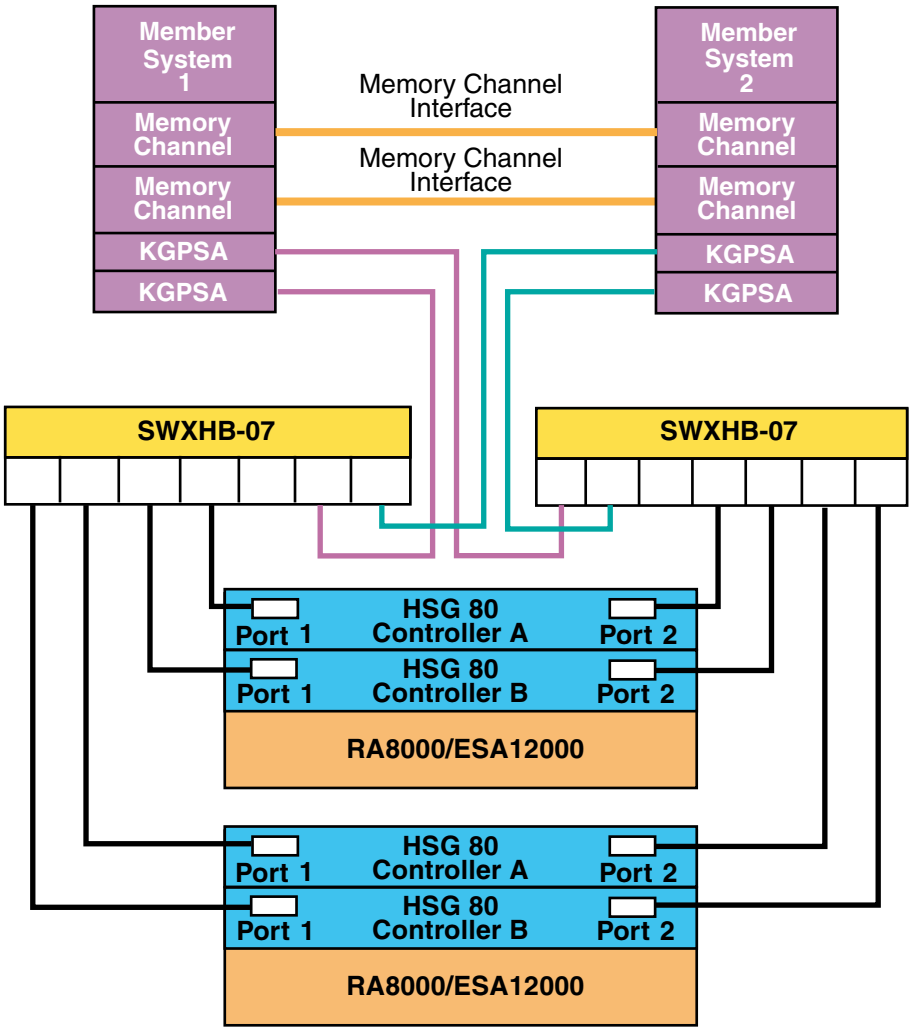

ZK-1814U-AI

3.6 QuickLoop

QuickLoop supports Fibre Channel arbitrated loop (FC-AL) devices within a fabric. This logical private loop fabric attach (PLFA) consists of multiple

private arbitrated loops (looplets) that are interconnected by a fabric. A private loop is formed by logically connecting ports on up to two switches.

Note

QuickLoop is not supported in a Tru64 UNIX Version 5.1 or TruCluster Server Version 5.1 configuration.

3.7 Zoning

This section provides a brief overview of zoning.

A zone is a logical subset of the Fibre Channel devices that are connected to the fabric. Zoning allows partitioning of resources for management and access control. In some configurations, it may provide for more efficient use of hardware resources by allowing one switch to serve multiple clusters or even multiple operating systems. Zoning entails splitting the fabric into zones, where each zone is essentially a virtual fabric.

Zoning may be used:

- When you want to set up barriers between systems of different operating environments or uses, for instance to allow two clusters to utilize the same switch.
- To create test areas that are separate from the rest of the fabric.
- To provide better utilization of a switch by reducing the number of unused ports.

Note

Any initial zoning must be made before connecting the host bus adapters and the storage to the switches, but after zoning is configured, changes can be made dynamically.

3.7.1 Switch Zoning Versus Selective Storage Presentation

Switch zoning and the selective storage presentation (SSP) feature of the HSG80 controllers have similar functions.

Switch zoning controls which servers can communicate with each other and each storage controller host port. SSP controls which servers will have access to each storage unit.

Switch zoning controls access at the storage system level, whereas SSP controls access at the storage unit level.

The following configurations require zoning or selective storage presentation:

- When you have a TruCluster Server cluster in a storage array network (SAN) with other stand-alone systems (UNIX or non-UNIX), or other clusters.
- Any time you have Windows NT or Windows 2000 in the same SAN with Tru64 UNIX. (Windows NT or Windows 2000 must be in a separate switch zone.)
- The SAN configuration has more than 64 connections to an RA8000, ESA12000, MA6000, MA8000, or EMA12000.

The use of selective storage presentation is the preferred way to control access to storage (so zoning is not required).

3.7.2 Types of Zoning

There are two types of zoning, soft and hard:

- Soft zoning is a software implementation based on the Simple Name Server (SNS) enforcing a zone. Zones are defined by either the node or port World Wide Names (WWN), or the domain and port numbers in the form of D,P, where D is the domain and P is the physical port number on the switch.

  A host system requests a list of all adapters and storage controllers that are connected to the fabric. The name service provides a list of all ports that are in the same zone or zones as the requesting host bus adapter.

  Soft zoning only works if all hosts honor it; it does not work if a host is not programmed to allow for soft zoning. For instance, if a host tries to access a controller that is outside the zone, the switch does not prevent the access.

  Tru64 UNIX honors soft zoning and does not attempt to access devices outside the zone.

  If you have used the WWN to define the zone and replace a KGPSA host bus adapter, you must modify the zone configuration and SSP because the node World Wide Name has changed.
- With hard zoning, zones are enforced at the physical level across all fabric switches by hardware blocking of Fibre Channel frames. Hardware zone definitions are in the form of D,P, where D is the domain and P is the physical port number on the switch. An example might be 1,2 for switch 1, port 2.

If a host attempts to access a port that is outside its zone, the switch hardware blocks the access.

You must modify the zone configuration if you move any cables from one port to another within the zone.

If you want to guarantee that there is no access outside any zone, either use hard zoning, or use operating systems that state that they support soft zoning.

Table 3–2 lists the types of zoning that are supported on each of the supported Fibre Channel switches.

Table 3–2: Type of Zoning Supported by Switches

Switch Type	Type of Zoning Supported
DS-DSGGA	Soft
DS-DSGGB	Soft and Hard
DS-DSGGC	Soft and Hard

3.7.3 Zoning Example

Figure 3–11 provides an example configuration using zoning. This configuration consists of two independent zones with each zone containing an independent cluster.

Figure 3–11: A Simple Zoned Configuration

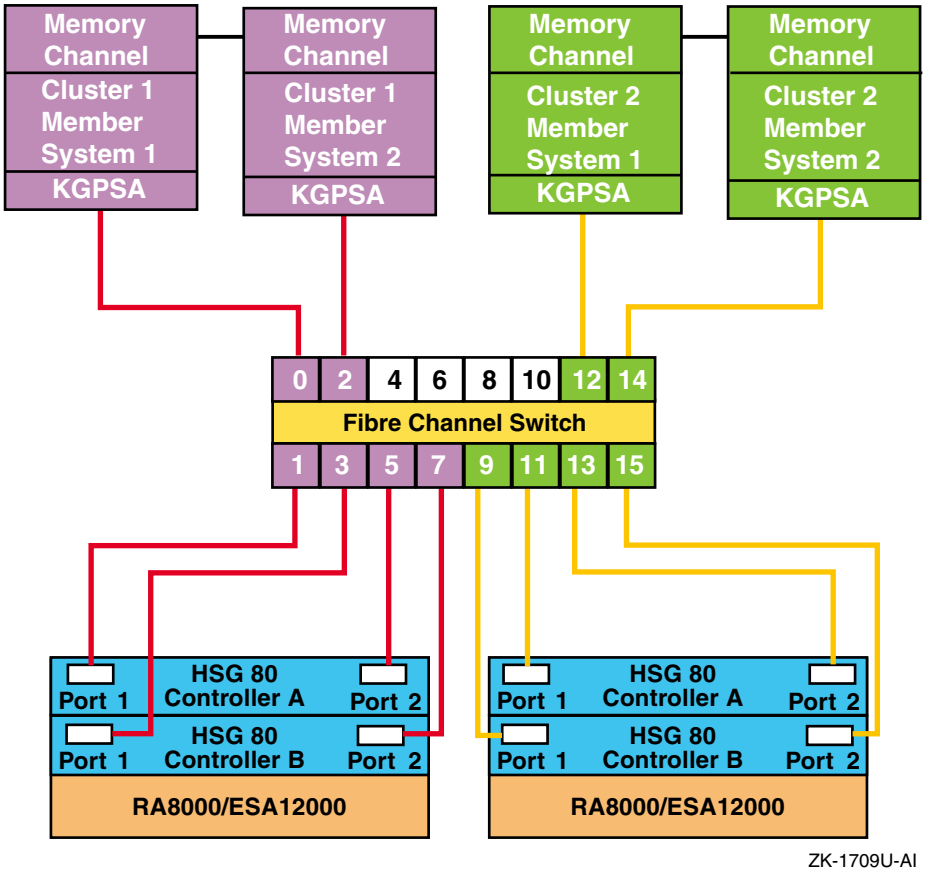

For information on setting up zoning, see the SAN Switch Zoning documentation that is provided with the switch.

3.8 Cascaded Switches

Multiple switches may be connected to each other to form a network of switches, or cascaded switches.

A cascaded switch configuration, which allows for network failures up to and including the switch without losing a data path to a SAN connected node, is called a mesh or meshed fabric. Figure 3–12 shows an example meshed fabric with three cascaded switches. This is not a no-single-point-of-failure (NSPOF) configuration.

Figure 3–12: Meshed Fabric with Three Cascaded Switches

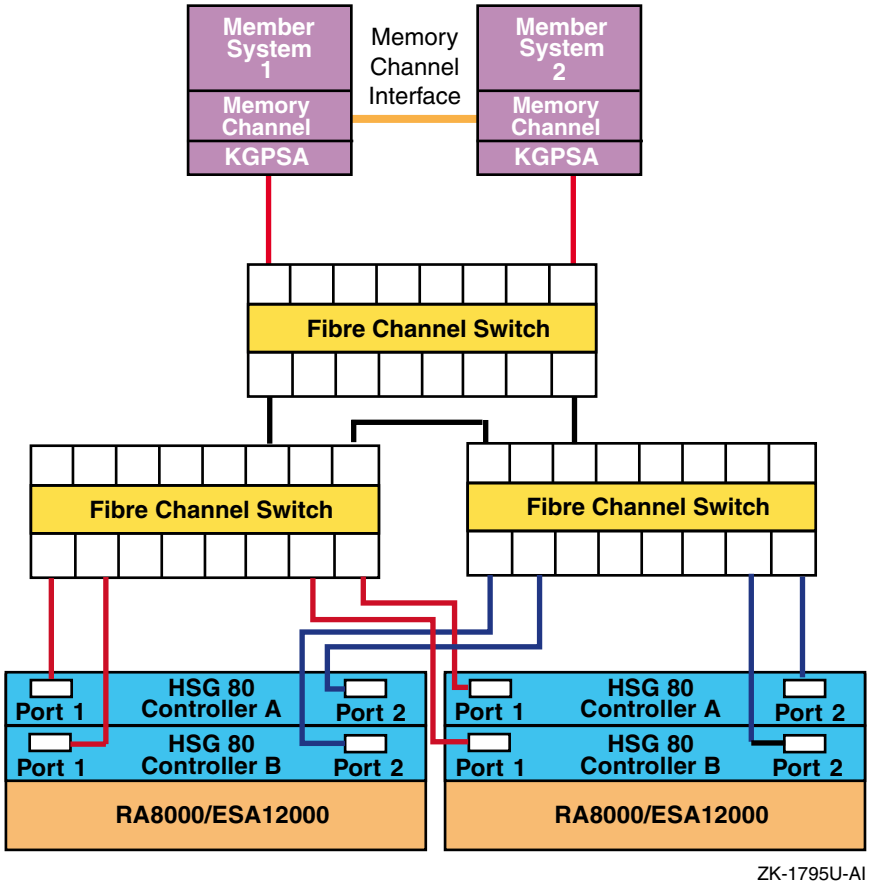

Figure 3–13 shows an example meshed resilient fabric with four cascaded interconnected switches. This configuration will tolerate multiple data path failures, and is an NSPOF configuration.

Figure 3–13: Meshed Resilient Fabric with Four Cascaded Switches

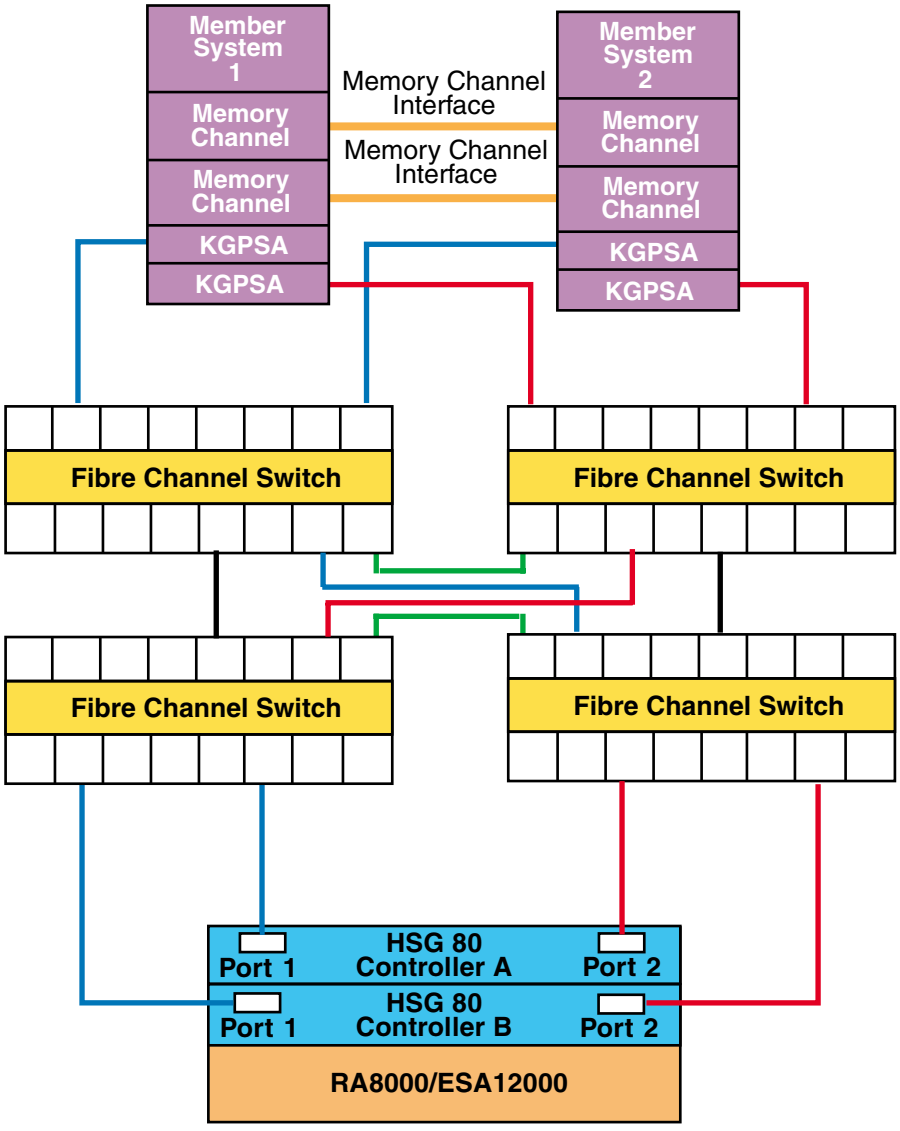

Note

If you lose an ISL, the communication can be routed through another switch to the same port on the other controller. This can constitute the maximum allowable two hops.

You can find the following information about storage array networks (SAN) in the *Compaq StorageWorks Heterogeneous Open SAN Design Reference Guide* located at:

http://www5.compaq.com/products/storageworks/techdoc/san/AA-RMPNA-TE.html

- Supported SAN topologies
- SAN fabric design rules
- SAN platform and operating system restrictions (including the number of switches supported)

4

Installing Fibre Channel Hardware

4.1 Installing and Configuring Fibre Channel Hardware

This chapter provides information about installing the Fibre Channel hardware needed to support Tru64 UNIX or a TruCluster Server configuration using Fibre Channel storage.

Ensure that the member systems, the Fibre Channel switches or hubs, and the HSG80 array controllers are placed within the lengths of the optical cables you will be using.

Note

The maximum length of the optical cable between the KGPSA and the switch (or hub), or the switch (or hub) and the HSG80 array controller, is 500 meters (1640.4 feet) via shortwave multimode Fibre Channel cable. The maximum distance between switches in a cascaded switch configuration is 10 kilometers (6.2 miles) using longwave single-mode fiber.

4.1.1 Installing and Setting Up the Fibre Channel Switch

The Fibre Channel switches support up to 8 (DS-DSGGA-AA, DS-DSGGB-AA, and DS-DSGGC-AA) or 16 (DS-DSGGA-AB, DS-DSGGB-AB, or DS-DSGGC-AB) full-duplex 1.6025 Gbits/sec ports. Each switch port can be connected to a KGPSA-BC or KGPSA-CA PCI-to-Fibre Channel host bus adapter, an HSG60 or HSG80 array controller, or another switch.

Each switch, except the DS-DSGGB-AA and DS-DSGGC-AA/AB, has a front panel display and four push buttons that you use to manage the switch. Four menus allow you to configure, operate, obtain status, or test the switch. The DS-DSGGB-AA and DS-DSGGC-AA/AB are managed through a telnet session after the IP address has been set (from a PC or terminal).

All switches have a 10Base-T Ethernet (RJ45) port, and after the IP address is set, the Ethernet connection allows you to manage the switch:

- Remotely using a telnet TCP/IP connection
- With the Simple Network Management Protocol (SNMP)

- Using Web management tools

Note

You have to set the IP address and subnet mask from the front panel (or from a PC or terminal with the DS-DSGGB-AA or DS-DSGGC-AA/AB) before you can manage the switch via a telnet session, SNMP, or the Web.

The DS-DSGGC-AA/AB Fibre Channel switches have a default IP address of 10.77.77.77. You may need to change this IP address before you connect the switch to the network.

The DSGGA switch has slots to accommodate up to four (DS-DSGGA-AA) or eight (DS-DSGGA-AB) plug-in interface modules. Each interface module in turn supports two Gigabit Interface Converter (GBIC) modules. The GBIC module is the electrical-to-optical converter.

Note

If you need to install additional interface modules in the DSGGA switch, do so before you place the switch in a relatively inaccessible location because you have to remove the top cover to install the interface modules.

The DSGGB switch accommodates up to 8 (DS-DSGGB-AA) or 16 (DS-DSBBG-AB) GBIC modules.

The DS-DSGGC-AA switch has seven fixed shortwave optical transceivers and one removable GBIC.

The DSGGC-AB switch accommodates up to 16 GBIC modules.

The shortwave GBIC supports 50-micron multimode fiber (MMF) using the standard subscriber connector (SC) connector. The longwave GBIC supports 9-micron, single-mode fiber optical cables. Only the 50-micron MMF optical cable is supported between the host bus adapters and switches or switches and HSG80 controllers for the TruCluster Server product. Longwave single-mode fiber optical cables are supported between switches in a cascaded switch configuration.

4.1.1.1 Installing the Switch

Place the switch within 500 meters (1640.4 feet) of the member systems (with KGPSA PCI-to-Fibre Channel adapter) and the HSG80 array controllers.

You can mount the switches in a 48.7-cm (19-in) rackmount installation or place the switch on a flat solid surface.

When you plan the switch location, ensure that you provide access to the front of the switch. All cables plug into the front of the switch. Also, for those switches with a control panel, the display and switches are on the front of the switch.

For an installation, at a minimum, you have to complete the following steps. Some of the steps are explained in more detail in the following sections.

1. Place the switch or install it in the rack.
2. If you are using a DS-DSGGB-AA or DS-DSGGC Fibre Channel switch, connect the switch to a terminal or PC (Section 4.1.1.2.3).
3. Connect the Ethernet cable between the Fibre Channel switch and the Ethernet switch or hub.
4. Connect the fiber-optic cables between the switch and host bus adapters and RAID array controllers.
5. Connect power to the switch.
6. Turn on the switch power. The switch runs a series of power-on self test (POST) tests. The DS-DSGGC Fibre Channel switches have no power switch; power is applied when the unit is plugged in.
7. Set the switch IP address and subnet mask (Section 4.1.1.2.2 or Section 4.1.1.2.3). You can also set the switch name if desired (Section 4.1.1.2.5). The switch IP address and subnet mask must be set from the front panel, except for the DS-DSGGB-AA or DS-DSGGC-AA/AB Fibre Channel switches. In this case you have to set the IP address from the PC or terminal. You may have to use a telnet session to set the switch name (Section 4.1.1.2.4).
8. Reboot the switch to enable the change in IP address and subnet mask to take effect.

For more information on the individual switches, see the following documentation:

- *Compaq StorageWorks Fibre Channel Storage Switch User's Guide*
- *Compaq StorageWorks SAN Switch 8 Installation and Hardware Guide*
- *Compaq StorageWorks SAN Switch 16 Installation and Hardware Guide*
- *Compaq StorageWorks Fibre Channel SAN Switch 8-EL Installation and Hardware Guide*
- *Compaq StorageWorks Fibre Channel SAN Switch 16-EL Installation and Hardware Guide*

For more information on managing a Fibre Channel switch, see the *Compaq StorageWorks Fibre Channel SAN Switch Management Guide*.

For more information on the SAN switch fabric operating system, see the *Compaq StorageWorks SAN Switch Fabric Operating System Management Guide*.

4.1.1.2 Managing the Fibre Channel Switches

You can manage the DS-DSGGA-AA, DS-DSGGA-AB, and DS-DSGGB-AB switches, and obtain switch status from the front panel, by making a telnet connection or by accessing the Web. The DS-DSGGB-AA and DS-DSGGC-AA/AB Fibre Channel switches do not have a front panel, so you must use a telnet connection or use Web access.

Before you can make a telnet connection or access the switch via the Web, you must assign an IP address and subnet mask to the Ethernet connection using the front panel or from a PC or terminal (DS-DSGGB-AA and DS-DSGGC-AA/AB). You may have to modify the DS-DSGGC-AA/AB IP address from the default of 10.77.77.77.

4.1.1.2.1 Using the Switch Front Panel

The switch front panel consists of a display and four buttons. The display is normally not active, but it lights up when any of the buttons are pressed. The display has a timer. After approximately 30 seconds of inactivity, the display will go out.

The four front panel buttons are:

- Up — Upward triangle: Scrolls the menu up (which effectively moves down the list of commands) or increases the value that is being displayed.
- Down — Downward triangle: Scrolls the menu down (which effectively moves up the list of commands) or decreases the value that is being displayed.

Note

When the up or down buttons are used to increase or decrease a numerical display, the number changes slowly at first, but changes to fast mode if the button is held down. The maximum number displayed is 255. An additional increment at a count of 255 resets the count to 0.

- Tab/Esc — Leftward triangle: Allows you to tab through multiple optional functions, for example, the fields in an IP address. You can use

this button to abort an entry, which takes you to the previous menu item. If pressed repeatedly, the front panel display will turn off.

- Enter — Rightward triangle: Causes the switch to accept the input you have made and move to the next function.

4.1.1.2.2 Setting the Ethernet IP Address and Subnet Mask from the Front Panel

Before you telnet to the switch, you must connect the Ethernet cable and then set the Ethernet IP address and subnet mask.

To use the front panel to set the Ethernet address and subnet mask, follow these steps:

1. Press any of the switch front panel buttons to activate the display for the top-level menu. If the Configuration Menu is not displayed, press the down button repeatedly until it is displayed:

   ```
   Select Menu:
   Configuration Menu
   ```

   Note

   Pressing the down button selects the next lower top-level menu. The top-level menus are:

   ```
   Configuration Menu
   Operation Menu
   Status Menu
   Test Menu
   ```

2. Press Enter to display the first submenu item in the configuration menu, Ethernet IP address:

   ```
   Ethernet IP address:
   10.00.00.10
   --
   ```

   The underline cursor denotes the selected address field.

   Use the up or down button to increase or decrease the displayed number. Use the Tab/Esc button to select the next field. Modify the address fields until you have the address set correctly.

3. Use Enter to accept the value and step to the next submenu item (Ethernet Submask), and then repeat step 2 to set the Ethernet subnet mask.
4. Press Enter to accept the Ethernet subnet mask.
5. Press the Tab/Esc button repeatedly to get back to the top-level menu.

6. Press the down button to select the Operation Menu:

```
Select Menu:
Operation Menu
```

7. If the switch is operational, place the switch off line before rebooting or you will lose any transmission in progress.

   Press Enter to display the first submenu in the Operation Menu, Switch Offline:

```
Operation Menu:
Switch Offline
```

8. Press the down button until the Reboot submenu item is displayed:

```
Operation Menu:
Reboot
```

9. Press Enter. You can change your mind and not reboot:

```
Reboot
Accept?    Yes  No
```

10. Use the Tab/Esc button to select `Yes`. Press Enter to reboot the switch and execute the POST tests.

Note

After changing any configuration menu settings, you must reboot the switch for the change to take effect.

Refer to the switch documentation for information on other switch configuration settings.

4.1.1.2.3 Setting the DS-DSGGB-AA and DS-DSGGC-AA/AB Ethernet IP Address and Subnet Mask from a PC or Terminal

For the DS-DSGGB-AA and DS-DSGGC-AA/AB switches, which do not have a front panel, you must use a connection to a Windows 95/98/NT/2000 PC or video terminal to set the Ethernet IP address and subnet mask.

To set the Ethernet IP address and subnet mask for the DS-DSGGB-AA or DS-DSGGC-AA/AB switches, follow these steps:

1. Connect the switch serial port to a terminal or PC COM port with a standard serial cable with a DB9 connector. Note that the serial port is only used for initial power-on self-test (POST) verification, IP address configuration, or for resetting the factory/default settings. Ensure that you make the connection to the switch serial port and not the Ethernet port.

2. If you are using a PC, start a remote communication program, for example, HyperTerminal.
3. Set the port settings to 9600 bits per second, 8 bits per character, no parity, 1 stop bit, and no flow control.
4. Turn on power to the switch and log in. If the connection is correct, the self-test results will be displayed. It takes 2 to 3 minutes for self-tests to complete.
   - DS-DSGGB-AA: The switch automatically connects to the host and logs the user on to the switch as `admin` when the self-tests terminate. For subsequent logons, the default password is `password`.
   - DS-DSGGC-AA/AB: Plugging in the DS-DSGGC-AA/AB switch turns the power on. (There is no on/off power switch.) Log in as the `admin` user. The password is `password`.
5. Enter the `ipAddrSet` command, then enter the IP address, subnet mask, and gateway address (if necessary). For example:

```
admin> ipAddrSet
Ethernet IP Address [10.77.77.77]: 16.142.72.54
Ethernet Subnetmask [255.255.255.0]: 255.255.255.0
Fibre Channel IP Address [none]: none
Fibre Channel Subnetmask [none]: none
Gateway Address [172.17.1.1]: none
[ y = set now, n = next reboot ]: y
Fabric OS (tm) Release v2.1.7
login:
```

When you have completed setting the IP address and subnet mask, disconnect the PC or terminal from the DS-DSGGB-AA or DS-DSGGC-AA/AB switch.

4.1.1.2.4 Logging In to the Switch with a Telnet Connection

Before you telnet to a Fibre Channel switch, you must set the Ethernet IP address and subnet mask.

Note

A serial port connection and a telnet session cannot both be active (at the same time) with the DS-DSGGB-AA and DS-DSGGC-AA/AB switch. The telnet session takes precedence and the serial port session is aborted when the telnet session is started.

You can use a telnet session to log in to the switch at one of three security levels. The default user names, listed from lowest security level to highest security level, are listed in Table 4–1.

Table 4–1: Telnet Session Default User Names for Fibre Channel Switches

DSGGA	DSGGB or DSGGC	Description
other	n/a	Allows you to execute commands ending in Show, such as dateShow and portShow.
user	user	Allows you to execute all commands ending in Show, plus any commands from the help menu that do not change the state of the switch, for example, version and errDump. You can change the passwords for all users up to and including the current user's security level.
admin	admin	Provides access to all the commands that show up in the help menu. Most switch administration is done when logged in as admin.
n/a	root	Gives users access to an extensive command set that can significantly alter system performance. Use root commands only at the request of Compaq customer service.

You can set the user names and passwords for users at or below the security level of the present login level by executing the passwd command. Enter a new user name (if desired) and a new password for the user.

Notes

Use Ctrl/H to correct typing errors.

Use the logout command to log out from any telnet connection.

4.1.1.2.5 Setting the Switch Name via Telnet Session

After you set the IP address and subnet mask, you can use a telnet session to log in to the switch to complete other switch management functions or monitor switch status. For example, if a system's /etc/hosts file contains an alias for the switch's IP address, set the switch name to the alias. This allows you to telnet to the switch name from that system. Telnet from a system that has the IP address in its /etc/hosts file and set the switch address as follows:

```
# telnet 132.25.47.146
User admin
Passwd
```

```
:Admin> switchName fcsw1
:Admin> switchName
fcsw1
:Admin>
```

Note

When you telnet to the switch the next time, the prompt will include the switch name, for example:

```
fcsw1:Admin>
```

4.1.2 Installing and Setting Up the DS-SWXHB-07 Hub

The DS-SWXHB-07 hub supports up to seven 1.6025 Gb/sec ports. The ports can be connected to the KGPSA-CA PCI-to-Fibre Channel host bus adapter or to an HSG80 array controller.

Unlike the DSGGA switch, the DS-SWXHB-07 hub does not have any controls or even a power-on switch. Simply plug in the hub to power it on. The hub has a green power indicator on the front panel.

The DS-SWXHB-07 hub has slots to accommodate up to seven plug-in interface converters. Each interface converter in turn supports two 1-gigabit Gigabit Interface Converter (GBIC) modules. The GBIC module is the electrical-to-optical converter, and supports both 50-micron and 62.5-micron multi-mode fiber (MMF) using the standard SC connector. Only the 50-micron MMF optical cable is supported for the TruCluster Software Products.

The GBIC modules and MMF optical cables are not provided with the hub. To obtain them, contact your authorized Compaq Service Representative.

4.1.2.1 Installing the Hub

Ensure that you place the hub within 500 meters (1640.4 feet) of the member systems (with KGPSA-CA PCI-to-Fibre Channel adapter) and the HSG80 array controllers.

The DS-SWXHB-07 hub can be placed on a flat, solid surface or, when configured in the DS-SWXHX-07 rack mount kit, part number 242795-B21, the hub can be mounted in a 48.7-cm (19-inch) rackmount installation. (One rack kit holds two hubs.) The hub is shipped with rubber feet to prevent marring the surface.

When you plan the hub location, ensure that you provide access to the GBIC connectors on the back of the hub. All cables plug into the back of the hub.

Caution

Static electricity can damage modules and electronic components. We recommend using a grounded antistatic wrist strap and a grounded work surface when handling modules.

For an installation, at a minimum, you have to:

1. Place the hub on an acceptable surface or install it in the rackmount.
2. Install one or more GBIC modules. Gently push the GBIC module into an available port on the hub until you feel the GBIC module click into place. The GBIC module has a built-in guide key that prevents you from inserting it incorrectly. Do not use excessive force.
3. Connect the optical fiber cables. To do this, plug one end of an MMF cable into one of the GBIC modules installed in the hub. Attach an MMF cable for all active port connections. Unused ports or improperly seated GBIC modules remain in loop bypass and do not affect the operation of the loop.
4. Attach the other end of the MMF cable to either the KGPSA-CA adapter or to the HSG80.
5. Connect power to the hub using a properly grounded outlet. Look at the power indicator on the front of the hub to make sure that it powered on.

For more installation information, see the *Fibre Channel Storage Hub 7 Installation Guide*.

4.1.2.2 Determining the Hub Status

Because the DS-SWXHB-07 hub is not a manageable unit, examine the status of the LED indicators to make sure that the hub is operating correctly. The LED indicators will be particularly useful after you have connected the hub to the KGPSA-CA host adapters and the HSG80 controller. However, at this time you can use the LEDs to verify that the GBIC connectors are installed correctly.

At power on, with no optical cables attached, the green and amber LEDs should both be on, indicating that the port is active but that the connection is invalid. The other possible LED states are as follows:

- Both off: Not active. Make sure that the GBIC is installed correctly.
- Solid green: Indicates presence and proper functionality of a GBIC.
- Green off: Indicates a fault condition (GBIC transmitter fault, improperly seated GBIC, no GBIC installed, or other failed device). The port is in bypass mode. This is the normal status for ports without GBICs installed.

- Solid amber: Indicates that a loss of signal or poor signal integrity has put the port in bypass mode. Make sure that a GBIC is installed, that a cable is attached to the GBIC, and that the other end of the cable is attached to a KGPSA-CA or HSG80.
- Amber off (and green on): Indicates that the port and device are fully operational.

For more information on determining the hub status, see the *Fibre Channel Storage Hub 7 Installation Guide*.

4.1.3 Installing and Configuring the KGPSA PCI-to-Fibre Channel Adapter Module

The following sections discuss KGPSA installation and configuration.

4.1.3.1 Installing the KGPSA PCI-to-Fibre Channel Adapter Module

To install the KGPSA-BC or KGPSA-CA PCI-to-Fibre Channel adapter modules follow these steps. For more information, see the following documentation:

- *KGPSA-BC PCI-to-Optical Fibre Channel Host Adapter User Guide*
- *64-Bit PCI-to-Fibre Channel Host Bus Adapter User Guide*

Caution

Static electricity can damage modules and electronic components. We recommend using a grounded antistatic wrist strap and a grounded work surface when handling modules.

1. If necessary, install the mounting bracket on the KGPSA-BC module. Place the mounting bracket tabs on the component side of the board. Insert the screws from the solder side of the board.
2. The KGPSA-BC should arrive with the gigabit link module (GLM) installed. If not, close the GLM ejector mechanism. Then, align the GLM alignment pins, alignment tabs, and connector pins with the holes, oval openings, and board socket. Press the GLM into place.

   The KGPSA-CA does not use a GLM, it uses an embedded optical shortwave multimode Fibre Channel interface.
3. Install the KGPSA in an open 32-bit or 64-bit PCI slot.
4. Insert the optical cable SC connectors into the KGPSA-BC GLM or KGPSA-CA SC connectors. The SC connectors are keyed to prevent their being plugged in incorrectly. Do not use unnecessary force.

Remember to remove the transparent plastic covering on the extremities of the optical cable.

5. Connect the fiber-optic cables to the shortwave Gigabit Interface Converter (GBIC) modules in the DSGGA, DSGGB, or DSGGC Fibre Channel switch.

4.1.3.2 Setting the KGPSA-BC or KGPSA-CA to Run on a Fabric

The KGPSA host bus adapter defaults to the fabric mode, and can be used in a fabric without taking any action. However, if you install a KGPSA that has been used in the loop mode on another system, you will need to reformat the KGPSA nonvolatile RAM (NVRAM) and configure it to run on a Fibre Channel fabric configuration.

Use the `wwidmgr` utility to determine the mode of operation of the KGPSA host bus adapter, and to set the mode if it needs changing (for example from loop to fabric).

Notes

You must set the console to diagnostic mode to use the `wwidmgr` utility for the following AlphaServer systems: AS1200, AS4x00, AS8x00, GS60, GS60E, and GS140. Set the console to diagnostic mode as follows:

```
P00>>> set mode diag
Console is in diagnostic mode
P00>>>
```

The console remains in `wwid` manager mode (or diagnostic mode for the AS1200, AS4x00, AS8x00, GS60, GS60E, and GS140 systems), and you cannot boot until the system is re-initialized. Use the `init` command or a system reset to re-initialize the system after you have completed using the `wwid` manager.

If you try to boot the system and receive the following error, initialize the console to get out of WWID manager mode, then reboot:

```
P00>>> boot
warning -- main memory zone is not free
P00>>> init
       .
       .
       .
P00>>> boot
```

If you have initialized and booted the system, then shut down the system and try to use the `wwidmgr` utility, you may be prevented from doing so. If you receive the following error, initialize the system and retry the `wwidmgr` command:

```
P00>>> wwidmgr -show adapter
wwidmgr available only prior to booting.
Reinit system and try again.
P00>>> init
       .
       .
       .
P00>>> wwidmgr -show adapter
       .
       .
       .
```

For more information on the `wwidmgr` utility, see the Wwidmgr *User's Manual*, which is on the Alpha Systems Firmware Update CD-ROM in the `DOC` directory.

Use the worldwide ID manager to show all KGPSA adapters:

```
P00>>> wwidmgr -show adapter
Link is down.
item       adapter           WWN               Cur. Topo  Next Topo
pga0.0.0.3.1 - Nvram read failed
[ 0] pga0.0.0.3.1         1000-0000-c920-eda0    FABRIC     UNAVAIL
pgb0.0.0.4.0 - Nvram read failed
[ 1] pgb0.0.0.4.0         1000-0000-c920-da01    FABRIC     UNAVAIL
pgc0.0.0.5.1 - Nvram read failed.
[ 2] pgc0.0.0.5.1         1000-0000-c920-cd9c    FABRIC     UNAVAIL
[9999] All of the above.
```

The `Link is down` message indicates that one of the adapters is not available, probably due to its not being plugged into a switch. The warning message `Nvram read failed` indicates that the KGPSA NVRAM has not been initialized and formatted. The next topology will always be `UNAVAIL` for the host bus adapter that has an unformatted NVRAM. Both messages are benign and can be ignored for the fabric mode of operation. To correct the `Nvram read failed` situation, use the `wwidmgr -set adapter` command.

The previous display shows that all three KGPSA host bus adapters are set for fabric topology as the current topology, the default. When operating in a fabric, if the current topology is `FABRIC`, it does not matter if the next topology is `Unavail`, or that the NVRAM is not formatted (`Nvram read failed`).

If, however, the current topology is `LOOP`, you have to change the topology to `FABRIC` to operate in a fabric. You will never see the `Nvram read failed` message if the current topology is `LOOP`. The NVRAM has to have been formatted to change the current mode to `LOOP`.

Consider the case where the KGPSA current topology is `LOOP` as follows:

```
P00>>> wwidmgr -show adapter
item       adapter           WWN               Cur. Topo  Next Topo
[ 0] pga0.0.0.3.1         1000-0000-c920-eda0    LOOP       LOOP
[ 1] pgb0.0.0.4.0         1000-0000-c920-da01    LOOP       LOOP
[9999] All of the above.
```

If the current topology for an adapter is LOOP, set an individual adapter to FABRIC by using the item number for that adapter (for example, 0 or 1). Use 9999 to set all adapters:

```
P00>>> wwidmgr -set adapter -item 9999 -topo fabric
Reformatting nvram
Reformatting nvram
```

Note

The qualifier in the previous command is -topo and not -topology. You will get an error if you use -topology.

Displaying the adapter information again will show the topology that the adapters will assume after the next console initialization:

```
P00>>> wwidmgr -show adapter
item       adapter           WWN               Cur. Topo  Next Topo
[ 0]    pga0.0.0.4.1    1000-0000-c920-eda0       LOOP       FABRIC
[ 1]    pgb0.0.0.3.0    1000-0000-c920-da01       LOOP       FABRIC
[9999] All of the above.
P00>>> init
```

This display shows that the current topology for both KGPSA host bus adapters is LOOP, but will be FABRIC after the next initialization.

The system initialization configures the KGPSAs to run on a fabric.

4.1.3.3 Setting the KGPSA-CA Adapter to Run in a Loop

If you do not want to use the KGPSA-CA adapter in loop mode, you can skip this section.

Before you can use the KGPSA adapter in loop mode, you must set the link type of the adapter to LOOP. You use the wwidmgr to accomplish this task.

Version 5.8 of the SRM console is required for boot support.

The version of the wwidmgr utility included with the SRM console can set the KGPSA to run in arbitrated loop mode or in fabric mode. Specifically, the wwidmgr -set adapter command stores the selected topology into the nonvolatile random-access memory (NVRAM) storage on the KGPSA adapter. The adapter retains this setting even if the adapter is later moved to another system.

Link Type

If a KGPSA in loop mode is connected to a Fibre Channel switch, the results are unpredictable. The same is true for a KGPSA in

fabric mode that is connected to a loop. Therefore, determine the topology setting before using the adapter.

The wwidmgr utility is documented in the *Wwidmgr User's Manual*, which is located in the DOC subdirectory of the Alpha Systems Firmware CD-ROM.

The steps required to set the link type are summarized here; see the *Wwidmgr User's Manual* for complete information and additional examples.

Assuming that you have the required console firmware, use the wwidmgr utility to set the link type, as follows:

1. Display the adapter on the system to determine its configuration:

```
POO>>> wwidmgr -show adapter

item       adapter           WWN                Cur. Topo   Next Topo
kgpsaa0.0.0.4.6 - Nvram read failed.
[ 0]    kgpsaa0.0.0.4.6    1000-0000-c920-05ab    FABRIC     UNAVAIL
[9999] All of the above.
```

   The warning message Nvram read failed indicates that the NVRAM on the KGPSA adapter has not been initialized and formatted. This is expected and is corrected when you set the adapter link type.

2. Set the link type on the adapter using the following values:
   - loop : sets the link type to loop (FC-AL)
   - fabric : sets the link type to fabric (point to point)

   You use the item number to indicate which adapter you wanted to change. For example, to configure adapter 0 (zero) for loop, use the following command:

```
POO>>> wwidmgr -set adapter -item 0 -topo loop
```

   The item number 9999 refers to all adapters. If you have KGPSA adapters configured for both arbitrated loop and fabric topologies, selecting 9999 will set them all to loop mode.

3. Verify the adapter settings:

```
POO>>> wwidmgr -show adapter
  item       adapter           WWN                Cur. Topo  Next Topo
 [ 0]    kgpsaa0.0.0.4.6   1000-0000-c920-05ab     FABRIC     LOOP
```

4. After making the change, reinitialize the console:

```
POO>>> init
```

5. Boot the system. The emx driver (Version 1.12 or higher is required) displays a message at boot when it recognizes the console setting, and configures the link accordingly.

6. Repeat this process for the other cluster member if this is a two-node TruCluster configuration.

4.1.3.4 Obtaining the Worldwide Names of KGPSA Adapters

A worldwide name is a unique number assigned to a subsystem by the Institute of Electrical and Electronics Engineers (IEEE) and set by the manufacturer prior to shipping. The worldwide name assigned to a subsystem never changes. We recommend that you obtain and record the worldwide names of Fibre Channel components in case you need to verify their target ID mappings in the operating system.

Fibre Channel devices have both a node name and a port name worldwide name, both of which are 64-bit numbers. Most commands you use with Fibre Channel only show the port name.

There are multiple ways to obtain the KGPSA port name worldwide name:

- You can obtain the worldwide name from a label on the KGPSA module before you install it.
- You can use the `show dev` command as follows:

```
P00>>> show dev
       .
       .
       .
pga0.0.0.1.0        PGA0     WWN 1000-0000-c920-eda0
pgb0.0.0.2.0        PGB0     WWN 1000-0000-c920-da01
```

- You can use the `wwidmgr -show adapter` command as follows:

```
P00>>> wwidmgr -show adapter
item       adapter           WWN               Cur. Topo  Next Topo
[ 0]    pga0.0.0.4.1    1000-0000-c920-eda0       FABRIC     FABRIC
[ 1]    pgb0.0.0.3.0    1000-0000-c920-da01       FABRIC     FABRIC
[9999] All of the above.
```

- If the operating system is installed, the worldwide name of a KGPSA adapter is also displayed in the boot messages generated when the `emx` driver attaches to the adapter when the adapter's host system boots. Or, you can use the `grep` utility and obtain the worldwide name from the `/var/adm/messages` file as follows:

```
# grep wwn /var/adm/messages
F/W Rev 2.20X2(1.12): wwn 1000-0000-c920-eda0
F/W Rev 2.20X2(1.12): wwn 1000-0000-c920-eda0
F/W Rev 2.20X2(1.12): wwn 1000-0000-c920-eda0
       .
       .
       .
```

Record the worldwide name of each KGPSA adapter for later use.

4.1.4 Setting Up the HSG80 Array Controller for Tru64 UNIX Installation

This section covers setting up the HSG80 controller for operation with Tru64 UNIX Version 5.1 and TruCluster Server Version 5.1.

The steps described here apply to both fabric and arbitrated loop configurations. However, arbitrated loop requires specific settings for the port topology and AL_PA values. If this is an arbitrated loop configuration, follow the steps described here, taking note of the difference in the port topology setting. Then see Section 4.1.4.1 for additional information.

For more information on installing the HSG80, see the *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Configuration Guide* or *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

To set up an HSG80 for TruCluster Server operation, follow these steps:

1. If not already installed, install the HSG80 controllers into the RA8000 or ESA12000 storage arrays.
2. If used, ensure that the external cache battery (ECB) is connected to the controller cache modules.
3. Install the fiber-optic cables between the KGPSA and the switch or hub.
4. Set the power verification and addressing (PVA) ID. Use PVA ID 0 for the enclosure that contains the HSG80 controllers. Set the PVA ID to 2 and 3 on expansion enclosures (if present).

   Note

   Do not use PVA ID 1.

   With Port-Target-LUN (PTL) addressing, the PVA ID is used to determine the target ID of the devices on ports 1 through 6 (the LUN is always zero). Valid target ID numbers are 0 through 15, excluding numbers 4 through 7. Target IDs 6 and 7 are reserved for the controller pair, and target IDs 4 and 5 are never used.

   The enclosure with PVA ID 0 will contain devices with target IDs 0 through 3; with PVA ID 2, target IDs 8 through 11; with PVA ID 3, target IDs 12 through 15. Setting a PVA ID of an enclosure to 1 would set target IDs to 4 through 7, generating a conflict with the target IDs of the controllers.

5. Remove the program card ESD cover and insert the controller's program card. Replace the ESD cover.

6. Install disks into storage shelves.
7. Connect a terminal to the maintenance port on one of the HSG80 controllers. You need a local connection to configure the controller for the first time. The maintenance port supports serial communication with the following default values:
   - 9600 BPS
   - 8 data bits
   - 1 stop bit
   - No parity
8. Connect the RA8000 or ESA12000 to the power source and apply power.

   Note

   For the HSG80 to see the connection to the KGPSAs, the KGPSA host bus adapters must be cabled to the switch or hub, with the system power applied before you turn power on to the RA8000/ESA12000.

9. If an uninterruptible power supply (UPS) is used instead of the external cache battery, to prevent the controller from periodically checking the cache batteries after power is applied, enter the following command:

   ```
   HSG80> set this CACHE_UPS
   ```

   Note

   Setting the controller variable `CACHE_UPS` for one controller sets it for both controllers.

10. From the maintenance terminal, use the `show this` and `show other` commands to verify that controllers have the current firmware version. See the *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 CLI Reference Guide* for information on upgrading the firmware.
11. To ensure proper operation of the HSG80 with Tru64 UNIX and TruCluster Server, set the the controller values as follows. Note that the port topology setting of `fabric` is specific to fabric configurations.

    ```
    HSG80> set nofailover 1
    HSG80> clear cli 2
    HSG80> set this mirrored_cache 3
    HSG80> set multibus copy = this 4
    HSG80> clear cli 5
    HSG80> set this port_1_topology = offline 6
    HSG80> set this port_2_topology = offline 6
    ```

```
HSG80> set other port_1_topology = offline [6]
HSG80> set other port_2_topology = offline [6]
HSG80> set this port_1_topology = fabric [7]
HSG80> set this port_2_topology = fabric [7]
HSG80> set other port_1_topology = fabric [7]
HSG80> set other port_2_topology = fabric [7]
HSG80> set this time=dd-mmm-yyyy:hh:mm:ss [8]
HSG80> set this scsi_version = scsi-3 [9]
HSG80> set other scsi_version = scsi-3 [9]
HSG80> restart other [10]
HSG80> restart this [10]
```

[1] Removes any failover mode that may have been previously configured.

[2] Prevents the command line interpreter (CLI) from reporting a misconfiguration error resulting from not having a failover mode set.

[3] Sets up mirrored cache, if desired, for the controller pair.

[4] Puts the controller pair into multiple-bus failover mode. Ensure that you copy the configuration information from the controller known to have a good array configuration.

Note

Use the `set failover copy = this_controller` command to set transparent failover mode.

[5] When the command is entered to set multiple-bus failover and copy the configuration information to the other controller, the other controller will restart. The restart may set off the audible alarm (which is silenced by pressing the button on the EMU). The CLI will display an event report, and continue reporting the condition until cleared with the `clear cli` command.

[6] Takes the ports off line and resets the topology to prevent an error message when setting the port topology.

[7] Sets fabric as the switch topology. If this is an arbitrated loop configuration, see Section 4.1.4.1 for configuration differences.

[8] Set the date and time on this controller. In a dual-redundant configuration, the command sets the time on both controllers. The value takes effect immediately.

[9] Specifies the host protocol to use. You can use either SCSI-2 or SCSI-3.

Setting the SCSI_VERSION to SCSI-2 allows a disk unit to be at LUN 0, and specifies that the command console LUN (CCL) is not fixed at a particular location, but floats to the first available LUN.

If SCSI_VERSION is set to SCSI-3, the CCL is presented at LUN 0 for all connection offsets. Do not assign unit 0 at any connection offset because the unit would be masked by the CCL at LUN 0 and would not be available.

Setting SCSI_VERSION to SCSI-3 is preferred because the CCL is fixed and it is much easier to manage a fixed CCL than a CCL that can change.

10 The HSG80 prompts you to restart both controllers after you set the SCSI version.

12. Enter the `show connection` command as shown in Example 4–1 to determine the HSG80 connection names for the connections to the KGPSA host bus adapters. For an RA8000/ESA12000 with dual-redundant HSG80s in multiple-bus failover mode, there will be four connections for each KGPSA in the cluster (as long as all four HSG80 ports are connected to the same fabric).

    For example, in a two-node cluster with two KGPSAs in each member system, and an RA8000 or ESA12000 with dual-redundant HSG80s, there will be 16 connections for the cluster. If you have other systems or clusters connected to the switches in the fabric, there will be other connections for the other systems. (Arbitrated loop configurations support a maximum of two systems.) In Example 4–1, note that the ! (exclamation mark) is part of the connection name. The `HOST_ID` is the KGPSA host name worldwide name and the `ADAPTER_ID` is the port name worldwide name.

Example 4–1: Determine HSG80 Connection Names

```
HSG80 show connection
Connection                                                          Unit
   Name      Operating system    Controller  Port   Address  Status      Offset


!NEWCON49      TRU64_UNIX           THIS       2    230813   OL this       0
 HOST_ID=1000-0000-C920-DA01                   ADAPTER_ID=1000-0000-C920-DA01

!NEWCON50      TRU64_UNIX           THIS       1    230813   OL this       0
           HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01

!NEWCON51      TRU64_UNIX           THIS       2    230913   OL this       0
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON52    TRU64_UNIX           THIS         1    230913   OL this       0
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON53      TRU64_UNIX           OTHER      1    230913   OL other      0
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB
```

Example 4–1: Determine HSG80 Connection Names (cont.)

```
!NEWCON54      TRU64_UNIX           OTHER      1    230813   OL other      0
          HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01

!NEWCON55      TRU64_UNIX           OTHER      2    230913   OL other      0
          HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON56      TRU64_UNIX           OTHER      2    230813   OL other      0
          HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01
.
.

!NEWCON61      TRU64_UNIX           THIS       2      210513   OL this       0
          HOST_ID=1000-0000-C921-086C         ADAPTER_ID=1000-0000-C921-086C

!NEWCON62      TRU64_UNIX           OTHER      1      210513   OL other      0
          HOST_ID=1000-0000-C921-086C         ADAPTER_ID=1000-0000-C921-086C

!NEWCON63      TRU64_UNIX           OTHER      1               offline       0
          HOST_ID=1000-0000-C921-0943         ADAPTER_ID=1000-0000-C921-0943

!NEWCON64      TRU64_UNIX           OTHER      1      210413   OL other      0
          HOST_ID=1000-0000-C920-EDA0         ADAPTER_ID=1000-0000-C920-EDA0

!NEWCON65      TRU64_UNIX           OTHER      2      210513   OL other      0
          HOST_ID=1000-0000-C921-086C         ADAPTER_ID=1000-0000-C921-086C
.
.

!NEWCON74      TRU64_UNIX           THIS       2      210413   OL this       0
          HOST_ID=1000-0000-C920-EDA0         ADAPTER_ID=1000-0000-C920-EDA0

!NEWCON75      TRU64_UNIX           THIS       2               offline       0
          HOST_ID=1000-0000-C921-0A75         ADAPTER_ID=1000-0000-C921-0A75

!NEWCON76      TRU64_UNIX           THIS       1      210413   OL this       0
          HOST_ID=1000-0000-C920-EDA0         ADAPTER_ID=1000-0000-C920-EDA0

!NEWCON77      TRU64_UNIX           THIS       1    210513   OL this       0
          HOST_ID=1000-0000-C921-086C         ADAPTER_ID=1000-0000-C921-086C

!NEWCON78      TRU64_UNIX           THIS       2             offline       0
          HOST_ID=1000-0000-C920-CB77         ADAPTER_ID=1000-0000-C920-CB77

!NEWCON79      TRU64_UNIX           OTHER      1             offline       0
          HOST_ID=1000-0000-C920-CB77         ADAPTER_ID=1000-0000-C920-CB77
.
.
```

Note

You can change the connection name with the HSG80 CLI RENAME command. The new connection name is limited to 9 characters. You cannot use a comma (,) or backslash (\) in

the connection name, and you cannot rename the connection to a name of the form used by the HSG80 (!NEWCON50). For example, assume that member system pepicelli has two KGPSA Fibre Channel host bus adapters, and that the worldwide name for KGPSA pga is 1000-0000-C920-DA01. Example 4–1 shows that the connections for pga are !NEWCON49, !NEWCON50, !NEWCON54, and !NEWCON56. You can change the name of !NEWCON49 to indicate that it is the first connection (of four) to pga on member system pepicelli as follows:

```
HSG80> rename !NEWCON49 pep_pga_1
```

13. For each connection to your cluster, verify that the operating system is TRU64_UNIX and the unit offset is 0. Search the show connection display for the worldwide name of each of the KGPSA adapters in your cluster member systems. If the operating system and offsets are incorrect, set them, then restart both controllers as follows:

```
HSG80> set !NEWCON49 unit_offset = 0 [1]
HSG80> set !NEWCON49 operating_ TRU64_UNIX [2]
HSG80> restart other [3]
HSG80> restart this [3]
       .
       .
       .
HSG80> show connection [4]
```

[1] Set the relative offset for LUN numbering to 0. You can set the unit_offset to nonzero values, but use caution. You may not be able to access storage units if you set the unit_offset improperly.

[2] Specify that the host environment connected to the Fibre Channel port is TRU64_UNIX. You must change each connection to TRU64_UNIX. This is very important.

Caution

Failure to set this to TRU64_UNIX will prevent your system from booting correctly, recovering from run-time errors, or from booting at all. (In this release, arbitrated loop configurations cannot boot from storage connected via Fibre Channel.) The default operating system is Windows NT, which uses a different SCSI dialect to talk to the HSG80 controller.

[3] Restart both controllers to cause all changes to take effect.

4 Enter the `show connection` command once more and verify that all connections have the offsets set to 0 and the operating system is set to `TRU64_UNIX`.

Note

If the fiber-optic cables are not properly installed, there will be inconsistencies in the connections shown.

14. Set up the storage sets as required for the applications to be used. An example is provided in Section 5.1.1.

4.1.4.1 Setting Up the HSG80 Array Controller for Arbitrated Loop

Section 4.1.4 describes settings that are common to both fabric and arbitrated loop configurations. This section describes settings that are unique to setting up the HSG80 controller for the arbitrated loop topology.

For more information on installing the HSG80 in an arbitrated loop topology, see the *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Configuration Guide*.

To set up an HSG80 for TruCluster arbitrated loop operation, follow steps 1 through 10 in Section 4.1.4. Then, in step 11, use the maintenance terminal to set the controller values as follows:

1. Set the `PORT_x_TOPOLOGY` value to `LOOP_HARD`. For example:

```
HSG80> set nofailover
HSG80> clear cli
HSG80> set multibus copy = this
HSG80> clear cli
HSG80> set this port_1_topology = offline
HSG80> set this port_2_topology = offline
HSG80> set other port_1_topology = offline
HSG80> set other port_2_topology = offline
HSG80> set this port_1_topology = LOOP_HARD
HSG80> set this port_2_topology = LOOP_HARD
HSG80> set other port_1_topology = LOOP_HARD
HSG80> set other port_2_topology = LOOP_HARD
```

   The `PORT_x_TOPOLOGY` value of `LOOP_HARD` enables arbitrated loop operation. Although the HSG80 controller also permits a topology setting of `LOOP_SOFT`, this is not supported in Tru64 UNIX.

2. Set `PORT_x_AL_PA` to unique values. `PORT_x _AL_PA` specifies the hexadecimal arbitrated loop physical address (AL_PA) for the HSG80 host ports.

This is the preferred address, but the HSG80 controller is free to use whatever AL_PA it obtains during loop initialization. However, the address you specify must be valid and must not be used by another port. If the controller is unable to obtain the address you specify (for example, because two ports are configured for the same address), the controller cannot come up on the loop.

In particular, if you do not set PORT_*x*_AL_PA, multiple ports might attempt to use the default address, thus causing a conflict.

The valid AL_PA addresses are within the range of 0-EF (hexadecimal), but not all addresses within this range are valid; the default value is 69 (hexadecimal).

The list of valid AL_PA addresses is as follows:

```
0x01, 0x02, 0x04, 0x08, 0x0F, 0x10, 0x17, 0x18, 0x1B, 0x1D,
0x1E, 0x1F, 0x23, 0x25, 0x26, 0x27, 0x29, 0x2A, 0x2B, 0x2C,
0x2D, 0x2E, 0x31, 0x32, 0x33, 0x34, 0x35, 0x36, 0x39, 0x3A,
0x3C, 0x43, 0x45, 0x46, 0x47, 0x49, 0x4A, 0x4B, 0x4C, 0x4D,
0x4E, 0x51, 0x52, 0x53, 0x54, 0x55, 0x56, 0x59, 0x5A, 0x5C,
0x63, 0x65, 0x66, 0x67, 0x69, 0x6A, 0x6B, 0x6C, 0x6D, 0x6E,
0x71, 0x72, 0x73, 0x74, 0x75, 0x76, 0x79, 0x7A, 0x7C, 0x80,
0x81, 0x82, 0x84, 0x88, 0x8F, 0x90, 0x97, 0x98, 0x9B, 0x9D,
0x9E, 0x9F, 0xA3, 0xA5, 0xA6, 0xA7, 0xA9, 0xAA, 0xAB, 0xAC,
0xAD, 0xAE, 0xB1, 0xB2, 0xB3, 0xB4, 0xB5, 0xB6, 0xB9, 0xBA,
0xBC, 0xC3, 0xC5, 0xC6, 0xC7, 0xC9, 0xCA, 0xCB, 0xCC, 0xCD,
0xCE, 0xD1, 0xD2, 0xD3, 0xD4, 0xD5, 0xD6, 0xD9, 0xDA, 0xDC,
0xE0, 0xE1, 0xE2, 0xE4, 0xE8, 0xEF
```

In multiple-bus failover mode, each port must have a unique AL_PA address because all of the ports can be active at the same time.

The convention in transparent failover mode is to use the same AL_PA address for Port 1 on both controllers and the same AL_PA address for Port 2 on both controllers. This allows the standby port on the alternate controller to have the same AL_PA address in the event of a failover. Because the ports are not active at the same time, the AL_PA addresses do not conflict. Make sure that the AL_PA address assigned to Port 1 is not the same as that assigned to Port 2, because they are distinct devices on the Fibre Channel loop.

The following example sets the PORT_*x*_AL_PA value for ports on two HSG80 controllers in multiple-bus failover mode:

```
HSG80> set this PORT_1_AL_PA  = 01
HSG80> set this PORT_2_AL_PA  = 02
HSG80> set other PORT_1_AL_PA  = 04
HSG80> set other PORT_2_AL_PA  = 08
```

The following example sets the PORT_*x*_AL_PA value for ports on two HSG80 controllers in transparent failover mode:

```
HSG80> set this PORT_1_AL_PA  = 01
HSG80> set this PORT_2_AL_PA  = 02
HSG80> set other PORT_1_AL_PA  = 01
```

```
HSG80> set other PORT_2_AL_PA  = 02
```

After you have done this, continue with steps 12 through 14 in Section 4.1.4.

4.1.4.2 Obtaining the Worldwide Names of HSG80 Controller

The RA8000 or ESA12000 is assigned a worldwide name when the unit is manufactured. The worldwide name (and checksum) of the unit appears on a sticker placed above the controllers. The worldwide name ends in zero (0), for example, 5000-1FE1-0000-0D60. You can also use the SHOW THIS_CONTROLLER Array Controller Software (ACS) command.

For HSG80 controllers, the controller port IDs are derived from the RA8000/ESA12000 worldwide name as follows:

- In a subsystem with two controllers in transparent failover mode, the controller port IDs increment as follows:
  - Controller A and controller B, port 1 — worldwide name + 1
  - Controller A and controller B, port 2 — worldwide name + 2

  For example, using the worldwide name of 5000-1FE1-0000-0D60, the following port IDs are automatically assigned and shared between the ports as a REPORTED PORT_ID on each port:
  - Controller A and controller B, port 1 — 5000-1FE1-0000-0D61
  - Controller A and controller B, port 2 — 5000-1FE1-0000-0D62
- In a configuration with dual-redundant controllers in multiple-bus failover mode, the controller port IDs increment as follows:
  - Controller A port 1 — worldwide name + 1
  - Controller A port 2 — worldwide name + 2
  - Controller B port 1 — worldwide name + 3
  - Controller B port 2 — worldwide name + 4

  For example, using the worldwide name of 5000-1FE1-0000-0D60, the following port IDs are automatically assigned and shared between the ports as a REPORTED PORT_ID on each port:
  - Controller A port 1 — 5000-1FE1-0000-0D61
  - Controller A port 2 — 5000-1FE1-0000-0D62
  - Controller B port 1 — 5000-1FE1-0000-0D63
  - Controller B port 2 — 5000-1FE1-0000-0D64

Because the HSG80 controller's configuration information and worldwide name is stored in nonvolatile random-access memory (NVRAM) on the

controller, there are different procedures for replacing HSG80 controllers in an RA8000 or ESA12000:

- If you replace one controller of a dual-redundant pair, the NVRAM from the remaining controller retains the configuration information (including worldwide name). When you install the replacement controller, the existing controller transfers configuration information to the replacement controller.
- If you have to replace the HSG80 controller in a single controller configuration, or if you must replace both HSG80 controllers in a dual-redundant configuration simultaneously, you have two options:
  - If the configuration has been saved to disk (with the `INITIALIZE` *`DISKnnnn`* `SAVE_CONFIGURATION` or `INITIALIZE` *`storageset-name`* `SAVE_CONFIGURATION` option), you can restore it from disk with the `CONFIGURATION RESTORE` command.
  - If you have not saved the configuration to disk, but the label containing the worldwide name and checksum is still intact, or you have recorded the worldwide name and checksum (Section 4.1.4.2) and other configuration information, you can use the command-line interface (CLI) commands to configure the new controller and set the worldwide name. Set the worldwide name as follows:

    `SET THIS NODEID=`*`nnnn-nnnn-nnnn-nnnn checksum`*

5

Installing Tru64 UNIX and TruCluster Server Software Using Fibre Channel Storage

After the hardware has been installed and configured, there are preliminary steps that must be completed before you install Tru64 UNIX and TruCluster Server on Fibre Channel disks.

You may already know the physical devices on which you want to install Tru64 UNIX and TruCluster Server. That is, you probably know the location of the devices within the storage cabinet, and are aware that the HSG80 `show unit` command displays the unit number of each of these devices. For example, `DISK10200` might be unit number `D1`.

You need to equate the unit number identified by the HSG80 controller with device names that the AlphaServer console can use. That is, the AlphaServer console must know about the devices before it can boot from, or dump to them, and it must have a valid Fibre Channel connection to each of those devices. For example, to boot from storage unit `D1` as presented by the HSG80 controller, the AlphaServer console requires a device name such as `dga133.1002.0.1.0` that identifies the storage unit. In addition, `dga133.1002.0.1.0` must be reachable via a valid Fibre Channel connection.

This chapter describes how to perform the following tasks:

- Configure HSG80 storagesets — In this document, example storagesets are configured for both Tru64 UNIX and TruCluster Server on Fibre Channel storage. Modify the storage configuration to meet your needs (Section 5.1.1).
- Set a user defined identifier (UDID) for each storage unit — Although Tru64 UNIX does not use this identifier directly, you use the UDID as input to the `wwidmgr -quickset` command in a subsequent step (Section 5.1.2).
- Use the UDID as input to the `wwidmgr -quickset` command to set the device unit number — The device unit number is a subset of the device name (as shown in a `show device` display). For example, in the device name `DKA100.1001.0.1.0`, the device unit number is 100 (`DKA100`). The Fibre Channel worldwide name (often referred to as the worldwide

ID or WWID) is too long (64 bits) to be used as the device unit number. Therefore, you set a device unit number that is an alias for the Fibre Channel worldwide name (Section 5.1.3).

- Display available Fibre Channel boot devices — When you set the device unit number, you also set the `wwid`*n* and `N`*n* console environment variables. These variables indicate which Fibre channel devices the console can access and which HSG80 ports can be used to access the devices. The `wwid`*n* variables also show which devices are displayed by the `show dev` console command, indicating that the devices can be used for booting or dumping (Section 5.1.4).
- Set the `bootdef_dev` console environment variable — Before you install the operating system (or cluster software), you must set the `bootdef_dev` console environment variable to ensure that you boot from the correct disk over a reachable path (Section 5.1.5).

If you are installing the Tru64 UNIX operating system or TruCluster Server software, follow the procedure in Chapter 1.

5.1 Before You Install

The following sections cover the preliminary steps that must be completed before you install Tru64 UNIX and TruCluster Server on Fibre Channel disks.

5.1.1 Configure the HSG80 Storagesets

After the hardware has been installed and configured, storagesets must be configured for software installation. The following disks/disk partitions are needed for base operating system and cluster installation:

- Tru64 UNIX disk
- Cluster root (`/`)
- Cluster `/usr`
- Cluster `/var`
- Member boot disk (one for each cluster member system)
- Quorum disk (if used)

If you are installing only the operating system, you need only the Tru64 UNIX disk (and any disks needed for applications). In this document we assume that both the base operating system and cluster software are to be installed on Fibre Channel disks.

If you are installing a cluster, you need one or more disks to hold the Tru64 UNIX operating system. The disks are either private disks on the system

that will become the first cluster member, or disks on a shared bus that the system can access. Whether the Tru64 UNIX disk is on a private disk or a shared disk, shut down the cluster before booting a cluster member system standalone from the Tru64 UNIX disk.

An example configuration will show the procedure necessary to set up disks for base operating system and cluster installation. Modify the procedure according to your own disk needs. You can use any supported RAID level.

The example is based on the use of four 4-GB disks used to create two mirrorsets (RAID level 1) to provide reliability. The mirrorsets are partitioned to provide partitions of appropriate sizes. Disks 30200, 30300, 40000, and 40100 are used for the mirrorsets.

Table 5–1 contains the necessary information to convert from the HSG80 unit numbers to `/dev/disk/dsk`*n* and device names for the example configuration. A blank table (Table A–1) is provided in Appendix A for use in an actual installation.

One mirrorset, the `BOOT-MIR` mirrorset, is used for the Tru64 UNIX and cluster member system boot disks. The other mirrorset, `CROOT-MIR`, is used for the cluster root (`/`), cluster `/usr`, cluster `/var`, and quorum disks.

To set up the example disks for operating system and cluster installation, follow the steps in Example 5–1.

Example 5–1: Setting Up the Mirrorset

```
HSG80> RUN CONFIG [1]
Config Local Program Invoked

Config is building its table and determining what devices exist
on the system. Please be patient.

add disk DISK30200  3 2 0
add disk DISK30300  3 3 0
add disk DISK40000  4 0 0
add disk DISK40100  4 1 0
...

Config - Normal Termination

HSG80> ADD MIRRORSET BOOT-MIR DISK30200 DISK40000 [2]
HSG80> ADD MIRRORSET CROOT-MIR DISK30300 DISK40100 [2]
HSG80> INITIALIZE BOOT-MIR [3]
HSG80> INITIALIZE CROOT-MIR [3]
HSG80> SHOW BOOT-MIR [4]
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------

BOOT-MIR      mirrorset                     DISK30200
                                            DISK40000
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
```

Example 5–1: Setting Up the Mirrorset (cont.)

```
          MEMBERSHIP = 2, 2 members present
        State:
          UNKNOWN -- State only available when configured as a unit
        Size:             8378028 blocks
HSG80> SHOW CROOT-MIR [4]
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------

CROOT-MIR     mirrorset                     DISK30300
                                            DISK40100
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 2, 2 members present
        State:
          UNKNOWN -- State only available when configured as a unit
        Size:             8378028 blocks
HSG80> CREATE_PARTITION BOOT-MIR SIZE=25 [5]
HSG80> CREATE_PARTITION BOOT-MIR SIZE=25 [5]
HSG80> CREATE_PARTITION BOOT-MIR SIZE=LARGEST [5]
HSG80> CREATE_PARTITION CROOT-MIR SIZE=5 [6]
HSG80> CREATE_PARTITION CROOT-MIR SIZE=15 [6]
HSG80> CREATE_PARTITION CROOT-MIR SIZE=40 [6]
HSG80> CREATE_PARTITION CROOT-MIR SIZE=LARGEST [6]
HSG80> SHOW BOOT-MIR [7]
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------

BOOT-MIR      mirrorset                     DISK30200
                                            DISK40000
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 2, 2 members present
        State:
          UNKNOWN -- State only available when configured as a unit
        Size:             8378028 blocks
        Partitions:
          Partition number        Size               Starting Block     Used by
          ------------------------------------------------------------------------
            1                  2094502 (   1072.38 MB)            0 [8]
            2                  2094502 (   1072.38 MB)      2094507 [9]
            3                  4189009 (   2144.77 MB)      4189014 [10]
HSG80>


HSG80> SHOW CROOT-MIR [11]
Name          Storageset                    Uses             Used by
------------------------------------------------------------------------------

CROOT-MIR     mirrorset                     DISK30300
                                            DISK40100
        Switches:
          POLICY (for replacement) = BEST_PERFORMANCE
          COPY (priority) = NORMAL
          READ_SOURCE = LEAST_BUSY
          MEMBERSHIP = 2, 2 members present
        State:
          UNKNOWN -- State only available when configured as a unit
```

Example 5–1: Setting Up the Mirrorset (cont.)

```
Size:             8378028 blocks
Partitions:
  Partition number        Size                Starting Block     Used by
  ------------------------------------------------------------------------
    1                   418896 (    214.47 MB)              0  12
    2                  1256699 (    643.42 MB)         418901  13
    3                  3351206 (   1715.81 MB)        1675605  14
    4                  3351207 (   1715.81 MB)        5026816  15
```

1 Use the CONFIG utility to configure the devices on the device side buses and add them to the controller configuration. The CONFIG utility takes about 2 minutes to complete. You can use the ADD DISK command to add disk drives to the configuration manually.

2 Create the BOOT-MIR mirrorset using disks 30200 and 30300 and the CROOT-MIR mirrorset using disks 40000 and 40100.

3 Initialize the BOOT-MIR and CROOT-MIR mirrorsets. If you want to set any initialization switches, you must do so in this step. The BOOT-MIR mirrorset will be used for the Tru64 UNIX and cluster member system boot disks. The CROOT-MIR mirrorset will be used for the cluster root (/), cluster /usr and cluster /var file systems, and the quorum disk.

4 Verify the mirrorset configuration and switches. Ensure that the mirrorsets use the correct disks.

5 Create appropriately sized partitions in the BOOT-MIR mirrorset using the percentage of the storageset that each partition will use. These partitions will be used for the two member system boot disks (25 percent or 1 GB each) and the Tru64 UNIX disk. For the last partition, the controller assigns the largest free space available to the partition (which will be close to 50 percent, or 2 GB).

6 Create appropriately sized partitions in the CROOT-MIR mirrorset using the percentage of the storageset that each partition will use. These partitions will be used for the quorum disk (5 percent), cluster root partition (15 percent), /usr (40 percent), and /var file systems. For the last partition, /var, the controller assigns the largest free space available to the partition (which will be close to 40 percent). See the TruCluster Server *Software Installation* manual to obtain partition sizes.

7 Verify the BOOT-MIR mirrorset partitions. Ensure that the partitions are of the desired size. The partition number is in the first column, followed by the partition size and starting block.

8 Partition for member system 1 boot disk.

9 Partition for member system 2 boot disk.

10 Partition for Tru64 UNIX operating system disk.

11 Verify the CROOT-MIR mirrorset partitions. Ensure that the partitions are of the desired size. The partition number is in the first column, followed by the partition size and starting block.

12 Partition for the quorum disk.

13 Partition for cluster root (/) filesystem.

14 Partition for cluster /usr filesystem.

15 Partition for cluster /var filesystem.

5.1.2 Adding Units and Identifiers to the HSG80 Storagesets

After you have created the storagesets and partitions, assign a unit number to each partition and set a unique identifier as shown in Example 5–2.

Note

A storageset must reside on one controller or the other. All the partitions of a storageset must be on the same controller because all the partitions of a storageset fail over as a unit.

The steps performed in Example 5–2 include:

1. Assign a unit number to each storage unit and disable all access to the storage unit.
2. Set an identifier for each storage unit.
3. Enable selective access to the storage unit.

Example 5–2: Adding Units and Identifiers to the HSG80 Storagesets

```
HSG80> ADD UNIT D131 BOOT-MIR PARTITION=1 DISABLE_ACCESS_PATH=ALL [1]
HSG80> ADD UNIT D132 BOOT-MIR PARTITION=2 DISABLE_ACCESS_PATH=ALL
HSG80> ADD UNIT D133 BOOT-MIR PARTITION=3 DISABLE_ACCESS_PATH=ALL
HSG80> ADD UNIT D141 CROOT-MIR PARTITION=1 DISABLE_ACCESS_PATH=ALL
HSG80> ADD UNIT D142 CROOT-MIR PARTITION=2 DISABLE_ACCESS_PATH=ALL
HSG80> ADD UNIT D143 CROOT-MIR PARTITION=3 DISABLE_ACCESS_PATH=ALL
HSG80> ADD UNIT D144 CROOT-MIR PARTITION=4 DISABLE_ACCESS_PATH=ALL

HSG80> SET D131 IDENTIFIER=131 [2]
HSG80> SET D132 IDENTIFIER=132
HSG80> SET D133 IDENTIFIER=133
HSG80> SET D141 IDENTIFIER=141
HSG80> SET D142 IDENTIFIER=142
HSG80> SET D143 IDENTIFIER=143
HSG80> SET D144 IDENTIFIER=144
HSG80> set d131 ENABLE_ACCESS_PATH = !NEWCON49,!NEWCON50,!NEWCON51,!NEWCON52 [3]
HSG80> set d131 ENABLE_ACCESS_PATH = !NEWCON53,!NEWCON54,!NEWCON55,!NEWCON56
Warning 1000: Other host(s) in addition to the one(s) specified can still
```

Example 5–2: Adding Units and Identifiers to the HSG80 Storagesets (cont.)

```
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified.
HSG80> set d131 ENABLE_ACCESS_PATH = !NEWCON61,!NEWCON62,!NEWCON64,!NEWCON65
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified.
HSG80> set d131 ENABLE_ACCESS_PATH = !NEWCON68,!NEWCON74,!NEWCON76,!NEWCON77
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified.
HSG80> set d132 ENABLE_ACCESS_PATH = !NEWCON49,!NEWCON50,!NEWCON51,!NEWCON52
.
.

HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON49,!NEWCON50,!NEWCON51,!NEWCON52
HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON53,!NEWCON54,!NEWCON55,!NEWCON56
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified.
HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON61,!NEWCON62,!NEWCON64,!NEWCON65
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified.
HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON68,!NEWCON74,!NEWCON76,!NEWCON77
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified.
HSG80> show d131 [4]
    LUN                                     Uses             Used by
------------------------------------------------------------------------------

  D131                                      BOOT-MIR     (partition)
        LUN ID:      6000-1FE1-0000-0D60-0009-8080-0434-002F
        IDENTIFIER = 131
        Switches:
          RUN                    NOWRITE_PROTECT        READ_CACHE
          READAHEAD_CACHE        WRITEBACK_CACHE
          MAXIMUM_CACHED_TRANSFER_SIZE = 32
        Access:
          !NEWCON49, !NEWCON50, !NEWCON51, !NEWCON52, !NEWCON53, !NEWCON54,
          !NEWCON55, !NEWCON56, !NEWCON61, !NEWCON62, !NEWCON64, !NEWCON65,
          !NEWCON68, !NEWCON74, !NEWCON76, !NEWCON77
        State:
          ONLINE to the other controller
          NOPREFERRED_PATH
        Size:             2094502 blocks
        Geometry (C/H/S): ( 927 / 20 / 113 )
.
.

HSG80> show d144 [4]
    LUN                                     Uses             Used by
------------------------------------------------------------------------------
```

Example 5–2: Adding Units and Identifiers to the HSG80 Storagesets (cont.)

```
D144                                        CROOT-MIR    (partition)
      LUN ID:      6000-1FE1-0000-0D60-0009-8080-0434-0028
      IDENTIFIER = 144
      Switches:
        RUN                     NOWRITE_PROTECT        READ_CACHE
        READAHEAD_CACHE         WRITEBACK_CACHE
        MAXIMUM_CACHED_TRANSFER_SIZE = 32
      Access:
        !NEWCON49, !NEWCON50, !NEWCON51, !NEWCON52, !NEWCON53, !NEWCON54,
        !NEWCON55, !NEWCON56, !NEWCON61, !NEWCON62, !NEWCON64, !NEWCON65,
        !NEWCON68, !NEWCON74, !NEWCON76, !NEWCON77
      State:
        ONLINE to the other controller
        NOPREFERRED_PATH
      Size:             3351207 blocks
      Geometry (C/H/S): ( 1483 / 20 / 113 )
```

1 Assign a unit number to each partition. When the unit is created by the `ADD UNIT` command, disable access to all hosts. This allows selective access in case there are other systems or clusters that are connected to the same switch as the cluster.

Record the unit name of each partition with the intended use for that partition (Table 5–1).

Note

In a multiple-bus failover, dual-redundant configuration, you can balance the I/O load by specifying the controller through which the unit will be accessed. Use the `PREFERRED_PATH` switch to the `add unit` or `set unit` ACL commands (`PREFERRED_PATH = this_controller` or `PREFERRED_PATH = other_controller`).

All partitions of a container must be addressed through the same controller. If you set `PREFERRED_PATH` for one partition, all partitions of that container inherit the same path.

2 Set an identifier for each storage unit. Use any number between 1 and 9999.

To keep your storage naming as consistent and simple as possible, consider using the unit number of the unit as its UDID. For instance, if the unit number is `D133`, use `133` as the UDID. Note, however, that the identifier must be unique. If you have multiple RA8000 or ESA12000

storage arrays, a unit number such as `D133` might not be unique across the storage arrays.

The identifier you select shows up as the user defined identifier (UDID) in the `wwidmgr -show wwid` display. The WWID manager uses it when setting the device unit number and `boot_def_dev` console environment variable. The identifier also shows up during the Tru64 UNIX installation to allow you to select the Tru64 UNIX installation disk.

The identifier is also used with the hardware manager view devices command (`hwmgr -view devices`) to locate the `/dev/disk/dsk`*n* value.

Note

We recommend that you set the identifier for all Fibre Channel storagesets. It provides a sure method of identifying the storagesets. Make the identifiers unique numbers within the domain (or within the cluster at a minimum). In other words, do not use the same identifier on more than one HSG80. The identifiers should be easily recognized. Ensure that you record the identifiers (Table 5–1).

3 Enable access to each unit for those hosts that you want to be able to access this unit. Because access was initially disabled to all hosts, you can ensure selective access to the units. If you do not remember the connection names, use the HSG80 `show connection` command as shown in Example 4–1 to determine the HSG80 connection names for the connection to the KGPSA host bus adapters. Many of the connections used here are shown in Example 4–1.

Use the connection name for each connection to the KGPSA host bus adapter on the host for which you want access enabled. Many of the connections used here are shown in Example 4–1.

4 Use the `SHOW` *`unit`* command (where unit is D131 through D133 and D141 through 144 in the example) to verify the identifier and that access to each unit is correct. Ensure that there is no connection to an unwanted system. Record the identifier and worldwide name for later use. Table 5–1 is a sample table filled in for the example. Table A–1 in Appendix A is a blank table for your use in an actual installation.

Note

At this point, even though the table is filled in, we do not yet know the device names or dsk*n* numbers.

Table 5–1: Converting Storageset Unit Numbers to Disk Names

File System or Disk	HSG80 Unit	Worldwide Name	UDID	Device Name	dsk*n*
Member 1 boot disk	D131	6000-1FE1-0000-0D60-0009-8080-0434-002F	131	dga131.1001.0.1.0	dsk17
Member 2 boot disk	D132	6000-1FE1-0000-0D60-0009-8080-0434-0030	132	dga132.1001.0.1.0	dsk16
Tru64 UNIX disk	D133	6000-1FE1-0000-0D60-0009-8080-0434-002E	133	dga133.1001.0.1.0	dsk15
Quorum disk	D141	6000-1FE1-0000-0D60-0009-8080-0434-0029	141	N/A[a]	dsk21
Cluster root (/)	D142	6000-1FE1-0000-0D60-0009-8080-0434-002A	142	N/A[a]	dsk20
/usr	D143	6000-1FE1-0000-0D60-0009-8080-0434-002B	143	N/A[a]	dsk19
/var	D144	6000-1FE1-0000-0D60-0009-8080-0434-0028	144	N/A[a]	dsk18

[a] These units are not assigned an alias for the device unit number by the WWID manager command; therefore, they do not get a device name and will not show up in a console show dev display.

5.1.3 Set the Device Unit Number

The device unit number is a subset of the device name as shown in a show device console display. For example, in the device name dga133.1001.0.1.0, the device unit number is 133. The console uses this device unit number to identify a storage unit. When you set a device unit number, you are really setting an alias for the device worldwide name. The 64-bit worldwide name is too large to be used as the device unit number, so an alias is used instead.

This section describes how to use the wwidmgr -quickset command to set the device unit number for the Fibre Channel disks to be used as the Tru64 UNIX Version 5.1 installation disk or cluster member system boot disks. Setting the device unit number allows the installation scripts to recognize a Fibre Channel disk.

To set the device unit number for a Fibre Channel device, follow these steps:

1. Use the HSG80 show unit command to obtain the user-defined identifier (UDID) for the HSG80 storageset to be used as the Tru64

UNIX Version 5.1 installation disk or cluster member system boot disks. Record the UDIDs. The `show unit` command also displays the HSG80 disk name, so you can be sure you are using the correct device.

In the example in Table 5–1, the Tru64 UNIX disk is unit D133 with a UDID 133. The UDID for the cluster member 1 boot disk is 131, and the cluster member 2 boot disk is 132.

2. From the AlphaServer console, use the `wwidmgr -clear all` command to clear the stored Fibre Channel `wwid1`, `wwid2`, `wwid3`, `wwid4`, `N1`, `N2`, `N3`, and `N4` console environment variables. You want to start with all `wwid`*n* and `N`*n* variables clear.

```
P00>>> wwidmgr -clear all
P00>>> show wwid*
wwid0
wwid1
wwid2
wwid3
P00>>> show n*
N1
N2
N3
N4
```

Note

The console only creates devices for which the `wwid`*n* console environment variable has been set, and are accessible through an HSG80 N_Port as specified by the `N`*n* console environment variable also being set. These console environment variables are set with the `wwidmgr -quickset` or `wwidmgr -set wwid` commands. The use of the `wwidmgr -quickset` command is shown later in Example 5–4.

3. Example 5–3 shows the use of the `wwidmgr -show wwid` command to display the UDID and worldwide names of all devices known to the console. At this time, if you issue the `show device` console command, it will not detect the Fibre Channel devices connected to the HSG80.

Example 5–3: Displaying the UDID and Worldwide Names of Devices Known to the Console

```
P00>>> wwidmgr -show wwid
[0] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-0008 (ev:none)
[1] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-0007 (ev:none)
[2] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-0009 (ev:none)
[3] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-000a (ev:none)
[4] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-000b (ev:none)
[5] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-000c (ev:none)
[6] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-000d (ev:none)
[7] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-000e (ev:none)
[8] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-000f (ev:none)
[9] UDID:-1 WWID:01000010:6000-1fe1-0001-4770-0009-9171-3579-0010 (ev:none)
[10] UDID:131 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002f (ev:none)
[11] UDID:132 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0030 (ev:none)
[12] UDID:133 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002e (ev:none)
[13] UDID:141 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0029 (ev:none)
[14] UDID:142 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002a (ev:none)
[15] UDID:143 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002b (ev:none)
[16] UDID:144 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0028 (ev:none)
[17] UDID:-1 WWID:01000010:6000-1fe1-0000-0ca0-0009-8090-0708-002b (ev:none)
[18] UDID:-1 WWID:01000010:6000-1fe1-0000-0ca0-0009-8090-0708-002c (ev:none)
[19] UDID:-1 WWID:01000010:6000-1fe1-0000-0ca0-0009-8090-0708-002d (ev:none)
[20] UDID:-1 WWID:01000010:6000-1fe1-0000-0ca0-0009-8090-0708-002e (ev:none)
1  2    3                                                          4
```

1 The number within the brackets ([ ]) is the item number of the device shown on any particular line.

2 The UDID is assigned at the HSG80 with the `set D`*n* `IDENTIFIER = `*xxx* command, and is not used by the Tru64 UNIX operating system, but may be set (as we have done with the `SET D131 IDENTIFIER=131` group of commands). When the identifier is not set at the HSG80, a value of -1 is displayed.

3 The worldwide name for the device. It is prefixed with the value `WWID:01000010:`. The most significant 64 bits of the worldwide name resembles the HSG80 worldwide name, and is assigned when the unit is manufactured. The least significant 64 bits is a volume serial number generated by the HSG80. You can use the HSG80 `SHOW` *unit* command to determine the worldwide name for each storage unit (as shown in Example 5–1).

4 The console environment variable set for this worldwide name. Only 4 `wwid`*n* console environment variables (*wwid0*, *wwid1*, *wwid2*, and *wwid3*) can be set. The console `show dev` command only shows those disk devices for which a `wwid`*n* console environment variable has been set using the `wwidmgr -quickset` or `wwidmgr -set` command. In this example, none of the `wwid`*n* environment variables is set.

4. Look through the `wwidmgr -show wwid` display (Example 5–3) and locate the UDID for the Tru64 UNIX disk (133) and each member system boot disks (131, 132) to ensure the storage unit is seen. As a second check, compare the worldwide name values.

5. Use the `wwidmgr` command with the `-quickset` option to set a device unit number for the Tru64 UNIX Version 5.1 installation disk and the first cluster member system boot disk.

   Example 5–4 shows the use of the `wwidmgr` command with the `-quickset` option to define a device unit number, based on the UDID, as an alias for the worldwide name for the Tru64 UNIX installation disk and the first cluster member system boot disk. The `wwidmgr -quickset` utility sets the device unit number and also provides a display of the device names and how the disk is reachable (reachability display).

   The `wwidmgr -quickset` command generates multiple device unit numbers based on the UDID. The device unit number is an alias for the worldwide name, in a format the console can use to identify the storage unit. The `wwidmgr -quickset` command may generate multiple device names for a given device unit number because each possible path to a storage unit is given its own device name.

   Example 5–4 shows:

   - The use of the `wwidmgr -quickset` command to set the device unit number for the Tru64 UNIX Version 5.1 installation disk to 133, and the first cluster member system boot disk to 131.
   - The `wwidmgr -quickset` command provides a reachability display equivalent to execution of the `wwidmgr reachability` command. The reachability part of the display provides the following:
     - The worldwide name for the storage unit that is to be accessed.
     - The new device name for the storage unit.
     - The KGPSA adapters through which a connection to the storage unit is potentially available.
     - The WWID of the HSG80 port(s) (N_Ports) that will be used to access the storage unit.
     - The `connected` column indicates whether the storage unit is current available through the KGPSA to HSG80 controller port connection. The HSG80 controllers are in multiple-bus failover so each storage unit is presented by only one controller at a time.

Example 5–4: Setting the Device Unit Number with the wwidmgr quickset Command

```
P00>>> wwidmgr -quickset -udid 133

Disk assignment and reachability after next initialization:

6000-1fe1-0000-0d60-0009-8080-0434-002e
                          via adapter:         via fc nport:        connected:
dga133.1001.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d64      No
dga133.1002.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d62      Yes
dga133.1003.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d63      No
dga133.1004.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d61      Yes
dgb133.1001.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d64      No
dgb133.1002.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d62      Yes
dgb133.1003.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d63      No
dgb133.1004.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d61      Yes

P00>>> wwidmgr -quickset -udid 131

Disk assignment and reachability after next initialization:

6000-1fe1-0000-0d60-0009-8080-0434-002e
                          via adapter:         via fc nport:        connected:
dga133.1001.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d64      No
dga133.1002.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d62      Yes
dga133.1003.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d63      No
dga133.1004.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d61      Yes
dgb133.1001.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d64      No
dgb133.1002.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d62      Yes
dgb133.1003.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d63      No
dgb133.1004.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d61      Yes

6000-1fe1-0000-0d60-0009-8080-0434-002f
                          via adapter:         via fc nport:        connected:
dga131.1001.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d64      No
dga131.1002.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d62      Yes
dga131.1003.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d63      No
dga131.1004.0.1.0         pga0.0.0.1.0         5000-1fe1-0000-0d61      Yes
dgb131.1001.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d64      No
dgb131.1002.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d62      Yes
dgb131.1003.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d63      No
dgb131.1004.0.2.0         pgb0.0.0.2.0         5000-1fe1-0000-0d61      Yes
P00>>> init
```

Notes

The `wwidmgr -quickset` command can take up to a minute to complete on the AlphaServer 8*x*00, GS60, GS60E, and GS140 systems.

You must reinitialize the console after running the WWID manager (`wwidmgr`), and keep in mind that the AS1200, AS4x00, AS8x00, GS60, GS60E, and GS140 consoles are in diagnostic mode. The disks are not reachable and you cannot boot until after the system is initialized.

If you have not set the UDID, you cannot set the device unit number as shown in Example 5–4. You have to use the `quickset` command with the item number displayed by the `wwidmgr -show wwid` command (see Example 5–3). For instance, to set the device unit number of item number 17 of Example 5–3 to 770, use the following command:

```
P00>>> wwidmgr -quickset -item 17 -unit 770
```

The device names have also been set for the Tru64 UNIX disk and first cluster member system boot disks. Record these device names for later use.

In the reachability portion of the display, the storagesets are reachable from KGPSA `pga` through two HSG80 ports and from KGPSA `pgb` through two HSG80 ports. Also, the device unit number has been set for each KGPSA to HSG80 controller port connection, even if the storage unit is not currently reachable via that connection.

5.1.4 Displaying the Available Boot Devices

The only Fibre Channel devices displayed by the console `show dev` command are those devices that have been assigned to a `wwid`*n* environment variable with the `wwidmgr -quickset` command.

The devices shown in the reachability display of Example 5–4 are available for booting and the setting of the `bootdef_dev` console environment variable during normal console mode.

If you execute the `show wwid*` console command now, it will show that the environment variable `wwid`*n* is set for two disks. Also, the `show n*` command shows that the units are accessible through four HSG80 N_Ports as follows:

```
P00>>> show wwid*
wwid0  133 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002e
wwid1  131 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002f
wwid2
wwid3
P00>>> show n*
N1                      50001fe100000d64
N2                      50001fe100000d62
N3                      50001fe100000d63
N4                      50001fe100000d61
```

Caution

If you reuse a UDID, ensure that you clear the `wwid`*n* console environment variable set for the original storage unit. If you do not clear the `wwid`*n* console environment variable, the `wwidmgr`

command used to quickset the UDID the second time sets a second wwid*n* console environment variable for the same UDID. The original wwid*n* console environment variable has a lower number, and therefore is the console environment variable that will be used, but it points to a storage unit that is not reachable. Therefore, you will not be able to boot the system.

If you have already used the wwidmgr -quickset command the second time on the UDID and set a second wwid*n* console environment variable for the same UDID, and the system has failed to boot, use the wwidmgr command to clear the wwid*n* console environment variable for the unreachable storage unit. For example, consider the following state of the wwid*n* console environment variables where wwid1 and wwid2 are both for UDID 150, and wwid1 is the original console environment variable:

```
P00>>> show wwid*
wwid0    5 1 WWID:01000010:6000-1fe1-0006-3f10-0009-0270-0619-0006
wwid1    150 1 WWID:01000010:6000-1fe1-0006-3f10-0009-0270-0619-000d
wwid2    150 1 WWID:01000010:6000-1fe1-0006-3f10-0009-0270-0619-0011
wwid3
```

Clear the wwid1 console environment variable as follows:

```
P00>>> wwidmgr -clear wwid1
```

Then, reboot the system.

Example 5–5 provides sample device names as displayed by the show dev command after using the wwidmgr -quickset command to set the device unit numbers.

Example 5–5: Sample Fibre Channel Device Names

```
 P00>>> show dev
dga131.1001.0.1.0          $1$DGA131                      HSG80  V8.5F
dga131.1002.0.1.0          $1$DGA131                      HSG80  V8.5F
dga131.1003.0.1.0          $1$DGA131                      HSG80  V8.5F
dga131.1004.0.1.0          $1$DGA131                      HSG80  V8.5F
dga133.1001.0.1.0          $1$DGA133                      HSG80  V8.5F
dga133.1002.0.1.0          $1$DGA133                      HSG80  V8.5F
dga133.1003.0.1.0          $1$DGA133                      HSG80  V8.5F
dga133.1004.0.1.0          $1$DGA133                      HSG80  V8.5F
dgb131.1001.0.2.0          $1$DGA131                      HSG80  V8.5F
dgb131.1002.0.2.0          $1$DGA131                      HSG80  V8.5F
dgb131.1003.0.2.0          $1$DGA131                      HSG80  V8.5F
dgb131.1004.0.2.0          $1$DGA131                      HSG80  V8.5F
dgb133.1001.0.2.0          $1$DGA133                      HSG80  V8.5F
dgb133.1002.0.2.0          $1$DGA133                      HSG80  V8.5F
dgb133.1003.0.2.0          $1$DGA133                      HSG80  V8.5F
dgb133.1004.0.2.0          $1$DGA133                      HSG80  V8.5F
dka0.0.0.1.1               DKA0              COMPAQ BB00911CA0  3B05
dqa0.0.0.15.0              DQA0                COMPAQ CDR-8435  0013
dva0.0.0.1000.0            DVA0
ewa0.0.0.5.1               EWA0              08-00-2B-C4-61-11
pga0.0.0.1.0               PGA0        WWN 1000-0000-c920-eda0
```

Example 5–5: Sample Fibre Channel Device Names (cont.)

```
pgb0.0.0.2.0              PGB0        WWN 1000-0000-c920-da01
pka0.7.0.1.1              PKA0                 SCSI Bus ID 7  5.57
```

Note

The only Fibre Channel devices displayed by the console `show dev` command are those devices that have been assigned to a `wwid`*n* environment variable.

Before you start the Tru64 UNIX installation, you must set the `bootdef_dev` console environment variable.

5.1.5 Set the bootdef_dev Console Environment Variable for Tru64 UNIX Installation

When booting from Fibre Channel devices, you must set the `bootdef_dev` console environment variable to ensure that the installation procedure is able to boot the system after building the new kernel.

Note

The `bootdef_dev` environment variable values must point to the same HSG80.

To set the `bootdef_dev` console environment variable for the Tru64 UNIX installation when booting from a Fibre Channel device, follow these steps:

1. Obtain the device name for the Fibre Channel storage unit where you will install the Tru64 UNIX operating system. The device name shows up in the reachability display as shown in Example 5–4 with a `Yes` under the `connected` column. You can also use the `wwidmgr -show reachability` command to determine reachability. Example 5–5 provides the display for a `show dev` command, which shows the device names of devices that may be assigned to the `bootdef_dev` console environment variable. Example 5–4 and Example 5–5 show that the following device names can be used in the `bootdef_dev` console environment variable as possible boot devices:
   - `dga131.1002.0.1.0`
   - `dga131.1004.0.1.0`
   - `dga133.1002.0.1.0`

- dga133.1004.0.1.0
- dgb131.1002.0.2.0
- dgb131.1004.0.2.0
- dgb133.1002.0.2.0
- dgb133.1004.0.2.0

Each of the storage units are reachable through four different paths, two for each host bus adapter (the Yes in the connected column).

2. Set the bootdef_dev console environment variable to one of the boot paths that show up as connected to the disk you will use for the Tru64 UNIX installation. For the example disk configuration on the system where you are installing the Tru64 UNIX operating system (which will eventually be the first cluster member), set bootdef_dev as follows:

```
P00>>> set bootdef_dev dga133.1002.0.1.0
```

3. You must initialize the system to use any of the device names in the bootdef_dev variable:

```
P00>>> init
       .
       .
       .
```

After the initialization, the bootdef_dev will show up as follows:

```
P00>>> show bootdef_dev
bootdef_dev        dga133.1002.0.1.0
```

You are now ready to install the Tru64 UNIX operating system.

5.2 Install the Base Operating System

After you read the TruCluster Server *Software Installation* manual, and use the Tru64 UNIX *Installation Guide* as a reference, boot from the CD-ROM and perform a full installation of the Tru64 UNIX Version 5.1 operating system.

When the installation procedure displays the list of disks available for operating system installation as shown here, look for the identifier in the Location column. Verify the identifier from the table you have been preparing (Table 5–1).

```
To visually locate a disk, enter "ping <disk>",
where <disk> is the device name (for example, dsk0) of the disk you
want to locate.  If that disk has a visible indicator light, it will
blink until you are ready to continue.

      Device    Size  Controller  Disk
      Name     in GB  Type        Model       Location
1)    dsk0       4.0  SCSI        RZ2CA-LA    bus-0-targ-0-lun-0
2)    dsk15      1.0  SCSI        HSG80       IDENTIFIER=133
3)    dsk17      2.0  SCSI        HSG80       IDENTIFIER=131
```

If you flash the light on a storage unit (logical disk) that is a mirrorset, stripeset, or RAIDset, the lights on all disks in the storageset will blink.

Record the /dev/disk/dsk*n* value (dsk15) for the Tru64 UNIX disk that matches the UDID (133) (Table 5–1).

Complete the installation, following the instructions in the Tru64 UNIX *Installation Guide*.

5.3 Determining /dev/disk/dskn to Use for a Cluster Installation

Before you can install the TruCluster Server software, you must determine which /dev/disk/dsk*n* to use for the various TruCluster Server disks.

To determine the /dev/disk/dsk*n* to use for the cluster disks, follow these steps:

1. With the Tru64 UNIX Version 5.1 operating system at single-user or multi-user mode, use the hardware manager (hwmgr) utility with the -view devices option to display all devices on the system. Pipe the command through the grep utility to search for any items with the IDENTIFIER qualifier.

```
# hwmgr -view dev | grep IDENTIFIER
 HWID: Device Name          Mfg      Model            Location
 ------------------------------------------------------------------------
   62: /dev/disk/dsk15c     DEC      HSG80            IDENTIFIER=133
   63: /dev/disk/dsk16c     DEC      HSG80            IDENTIFIER=132
   64: /dev/disk/dsk17c     DEC      HSG80            IDENTIFIER=131
   65: /dev/disk/dsk18c     DEC      HSG80            IDENTIFIER=141
   66: /dev/disk/dsk19c     DEC      HSG80            IDENTIFIER=142
   67: /dev/disk/dsk20c     DEC      HSG80            IDENTIFIER=143
   68: /dev/disk/dsk21c     DEC      HSG80            IDENTIFIER=144
```

   If you know that you have set the UDID for a large number of disks, you can also grep for the UDID:

```
# hwmgr -view dev | grep IDENTIFIER | grep 131
```

   Note

   If you have not set the UDID, you can use hwmgr to determine the /dev/disk/dsk*n* name by using the hardware manager to display device attributes and searching for the storage unit worldwide name as follows:

```
# hwmgr -get attribute -a name -a dev_base_name | more
```

   Use the more search utility (/) to search for the worldwide name of the storageset you have set up for the particular disk in question. The following example shows the format of the command output:

```
# hwmgr -get attribute -a name -a dev_base_name
1:
  name = Compaq AlphaServer ES40
2:
  name = CPU0
       .
       .
       .
62:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002e
  dev_base_name = dsk15
63:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0030
  dev_base_name = dsk16
64:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002f
  dev_base_name = dsk17
65:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0028
  dev_base_name = dsk18
66:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002b
  dev_base_name = dsk19
67:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002a
  dev_base_name = dsk20
68:
  name = SCSI-WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0029
  dev_base_name = dsk21
69:
  name = SCSI-WWID:0710002c:"COMPAQ CDR-8435 :d05b003t00000l00000"
  dev_base_name = cdrom0
       .
       .
       .
```

For more information on the hardware manager (`hwmgr`), see `hwmgr`(8).

---

2. Search the display for the UDIDs (or worldwide names) for each of the cluster installation disks and record the `/dev/disk/dsk`*n* values.

   If you used the `grep` utility to search for a specific UDID, for example `hwmgr -view dev | grep "IDENTIFIER=131"` repeat the command to determine the `/dev/disk/dsk`*n* for each of the remaining cluster disks. Record the information for use when you install the cluster software.

You must label the disks before you install cluster software.

5.4 Label the Disks to Be Used to Create the Cluster

Before you run `clu_create` to create the first cluster member or `clu_add_member` to add subsequent cluster members, you must label the disks to be used for cluster software.

On the system where you installed the Tru64 UNIX operating system, if you have not already done so, boot the system. Determine the `/dev/disk/dsk`*n* values to use for cluster installation (Section 5.3).

Initialize disklabels for all disks needed to create the cluster. The example (Table 5–1) uses disks dsk18 (/var), dsk19 (/usr), dsk20 [cluster root (/)], and dsk21 (Quorum). For instance:

```
# disklabel -rw dsk20 HSG80
```

5.5 Install the TruCluster Server Software and Create the First Cluster Member

After labeling the disks, use the TruCluster Server *Software Installation* procedures and install the TruCluster Server software on the first cluster member (the system where you just installed Tru64 UNIX).

After installing the TruCluster Server software subsets, run the clu_create command to create the first cluster member using the procedures in the TruCluster Server *Software Installation* manual.

When clu_create terminates, do not reboot the system. Shut down the system and reset the bootdef_dev console environment variable to provide multiple boot paths to the member system boot disk before booting (Section 5.6). Then boot the first cluster member.

5.6 Reset the bootdef_dev Console Environment Variable

If you set the bootdef_dev console environment variable to multiple paths in Section 5.1.5, the base operating system installation or clu_create procedures modify the variable and you should reset it to provide multiple boot paths.

To reset the bootdef_dev console environment variable, follow these steps:

1. Obtain the device name and worldwide name for the Fibre Channel unit from where you will boot cluster member system 1 (Table 5–1).
2. Examine the reachability display (Example 5–4) provided by the wwidmgr -quickset or the wwidmgr -reachability commands for the device names that can access the storage unit from which you are booting. Examine the show dev command output to ensure the device name may be assigned to the bootdef_dev console environment variable.

Notes

Choose device names that show up as both Yes and No in the reachability display connected column. Keep in mind, that for multiple-bus failover, only one controller is normally active for a storage unit. You must ensure that the unit is reachable if the controllers have failed over.

If you have multiple Fibre Channel host bus adapters, use device names for at least two host bus adapters.

For example, to ensure that you have a connected boot path in case of a failed host bus adapter or controller failover, choose device names for multiple host bus adapters and each controller port. For example, if you use the reachability display shown in Example 5–4, you can choose the following device names when setting the `bootdef_dev` console environment variable for the first cluster member system:

```
dga131.1004.0.1.0    1
dgb131.1002.0.2.0    2
dgb131.1003.0.2.0    3
dga131.1001.0.1.0    4
```

1 Path from host bus adapter A to controller A port 1

2 Path from host bus adapter B to controller A port 2

3 Path from host bus adapter B to controller B port 1

4 Path from host bus adapter A to controller B port 2

You can set units preferred to a specific controller, in which case both controllers will be active.

If the `bootdef_dev` console environment variable ends up with all boot paths in an unconnected state, you can use the `ffauto` or `ffnext` console environment variables to force a boot device from a `not connected` to a `connected` state.

The `ffauto` console environment variable is effective only during autoboots (boots other than manual boots). Use the `set ffauto on` console command to enable `ffauto`. (The default for `ffauto` is `off`.) It is stored in non-volatile memory and therefore persists across system resets and power cycles.

During an autoboot, the console attempts to boot from each connected device listed in the `bootdef_dev` console environment variable. If `ffauto` is `on`, and if the end of devices listed in `bootdef_dev` is reached without successfully booting, the console starts again at the beginning of devices listed in the `bootdef_def` console environment variable. This time, devices that are not connected are changed to `connected` and an attempt is made to boot from that device.

The `ffnext` console environment variable is a one-time variable. It does not persist across a system reset, power

cycle, or reboot. This variable may be used (`set ffnext on`) to cause the next command to a `not connected` device to change the state to `connected`. After the command has been executed, the `ffnext` variable is automatically set to `off`, so it has no further effect.

For more information on using the `ffauto` and `ffnext` console environment variables, see the *Wwidmgr User's Manual*.

---

3. Set the `bootdef_dev` console environment variable for member system 1 boot disk to a comma-separated list of several of the boot paths that show up as connected in the reachability display (`wwidmgr -quickset` or `wwidmgr -show reachability`). You must initialize the system to use any of the device names in the `bootdef_dev` variable as follows:

```
P00>>> set bootdef_dev \
dga131.1001.0.1.0,dga131.1004.0.1.0,\
dgb131.1002.0.2.0,dgb131.1003.0.2.0
POO>>> init
```

Note

The console system reference manual (SRM) software guarantees that you can set the `bootdef_dev` console environment variable to a minimum of four device names. You may be able to set it to five, but only four are guaranteed.

5.7 Add Additional Systems to the Cluster

To add additional systems to the cluster, follow this procedure:

1. On the system where you installed the Tru64 UNIX operating system, boot the system into the cluster (single-member cluster), and referring to the TruCluster Server *Software Installation* procedures, use `clu_add_member` to add subsequent cluster members.

   Before you boot the system being added to the cluster, on the newly added cluster member:

   a. Use the `wwidmgr` utility with the `-quickset` option to set the device unit number for the member system boot disk as shown in Example 5–6. For member system 2 in the example configuration, it is the storage unit with UDID 132 (Table 5–1):

Example 5–6: Setting Device Unit Number for Additional Member System

```
P00>>> wwidmgr -quickset -udid 132

Disk assignment and reachability after next initialization:

6000-1fe1-0000-0d60-0009-8080-0434-0030
                    via adapter:     via fc nport:       connected:
dga132.1001.0.1.0    pga0.0.0.1.0     5000-1fe1-0000-0d64      No
dga132.1002.0.1.0    pga0.0.0.1.0     5000-1fe1-0000-0d62      Yes
dga132.1003.0.1.0    pga0.0.0.1.0     5000-1fe1-0000-0d63      No
dga132.1004.0.1.0    pga0.0.0.1.0     5000-1fe1-0000-0d61      Yes
dgb132.1001.0.2.0    pgb0.0.0.2.0     5000-1fe1-0000-0d64      No
dgb132.1002.0.2.0    pgb0.0.0.2.0     5000-1fe1-0000-0d62      Yes
dgb132.1003.0.2.0    pgb0.0.0.2.0     5000-1fe1-0000-0d63      No
dgb132.1004.0.2.0    pgb0.0.0.2.0     5000-1fe1-0000-0d61      Yes

P00>>> init
```

b. Set the `bootdef_dev` console environment variable to one reachable path (`Yes` in the connected column of Example 5–6) to the member system boot disk:

```
P00>>> set bootdef_dev dga132.1002.0.1.0
```

c. Boot `genvmunix` on the newly added cluster member system. Each installed subset will be configured and a new kernel will be built and installed. After the new kernel is built, do not reboot the new cluster member system.

2. Shut down the system and reset the `bootdef_dev` console environment variable to provide multiple boot paths to the member system boot disk as follows:

a. Obtain the device name and worldwide name for the Fibre Channel unit from where you will boot (Table 5–1). Ensure that you choose the correct device name for the entity you are booting (cluster member system 2, cluster member system 3, and so on).

b. Examine the reachability display (Example 5–6) provided by the `wwidmgr -quickset` command for the device names that can access the storage unit from where you are booting. Examine the `show dev` command output to ensure the device name may be assigned to the `bootdef_dev` console environment variable.

Notes

Choose device names that show up as both `Yes` and `No` in the reachability display `connected` column. Keep in mind, that for multiple-bus failover, only one controller is

normally active for a storage unit. You must ensure that the unit is reachable if the controllers have failed over.

If you have multiple Fibre Channel host bus adapters, use device names for at least two host bus adapters.

For example, to ensure that you have a connected boot path in case of a failed host bus adapter or controller failover, choose device names for multiple host bus adapters and each controller port. For example, if you use the reachability display shown in Example 5–6, you can choose the following device names when setting the `bootdef_dev` console environment variable for the second cluster member system:

```
dga132.1004.0.1.0    1
dgb132.1002.0.2.0    2
dgb132.1003.0.2.0    3
dga132.1001.0.1.0    4
```

1 Path from host bus adapter A to controller A port 1

2 Path from host bus adapter B to controller A port 2

3 Path from host bus adapter B to controller B port 1

4 Path from host bus adapter A to controller B port 2

---

c. Set the `bootdef_dev` console environment variable for member system 2 boot disk to a comma-separated list of several of the boot paths that show up as connected in the reachability display (`wwidmgr -quickset` or `wwidmgr -show reachability`). You must initialize the system to use any of the device names in the `bootdef_dev` variable as follows:

```
P00>>> set bootdef_dev \
dga132.1001.0.1.0,dga132.1004.0.1.0,\
dgb132.1002.0.2.0,dgb132.1003.0.2.0
POO>>> init
```

3. Boot the new cluster member system into the cluster and complete the cluster installation.
4. Repeat steps 1 through 3 for other cluster member systems.

6

Converting the HSG80 from Transparent to Multiple-Bus Failover Mode

If you are upgrading from Tru64 UNIX Version 4.0F or Version 4.0G and TruCluster Software Products Version 1.6 to Tru64 UNIX Version 5.1 and TruCluster Server Version 5.1, you may want to change from transparent failover to multiple-bus failover to take advantage of multibus support in Tru64 UNIX Version 5.1 and multiple-bus failover mode and the ability to create a no-single-point-of-failure (NSPOF) cluster.

Or, you may be using transparent failover mode with Tru64 UNIX Version 5.1 and TruCluster Server Version 5.1 and want to take advantage of the ability to create a NSPOF configuration, and the availability that multiple-bus failover provides over transparent failover mode.

6.1 Overview

The change in failover modes cannot be accomplished with a simple `SET MULTIBUS COPY=THIS` HSG80 CLI command because:

- Unit offsets are not changed by the HSG80 `SET MULTIBUS_FAILOVER COPY=THIS` command.

  Each path between a Fibre Channel host bus adapter in a host computer and an active host port on an HSG80 controller is a connection. During Fibre Channel initialization, when a controller becomes aware of a connection to a host bus adapter through a switch or hub, it adds the connection to its table of known connections. The unit offset for the connection depends on the failover mode in effect at the time the connection is discovered. In transparent failover mode, host connections to port 1 default to an offset of 0; host connections on port 2 default to an offset of 100. Host connections on port 1 can see units 0 through 99; host connections on port 2 can see units 100 through 199.

  In multiple-bus failover mode, host connections on either port 1 or 2 can see units 0 through 199. In multiple-bus failover mode, the default offset for both ports is 0.

  If you change the failover mode from transparent failover to multiple-bus failover, the offsets in the table of known connections remain the same as if they were for transparent failover mode; the offset on port 2 remains 100. With an offset of 100 on port 2, a host cannot see units 0 through 99

on port 2. This reduces the availability. Also, if you have only a single HSG80 controller and lose the connection to port 1, you lose access to units 0 through 99.

Therefore, if you want to change from transparent failover to multiple-bus failover mode, you must change the offset in the table of known connections for each connection that has a nonzero offset.

Note

Disconnecting and then reconnecting the cables does no good because a connection that is added to the table remains in the table until you delete the connection.

- The system can access a storage device through only one HSG80 port. The system's view of the storage device is not changed when the HSG80 is placed in multiple-bus failover mode.

  In transparent failover mode, the system accesses storage units D0 through D99 through port 1 and units D100 through D199 through port 2. In multiple-bus failover mode, you want the system to be able to access all units through all four ports.

6.2 Procedure to Convert from Transparent to Multiple-bus Failover Mode

To change from transparent failover to multiple-bus failover mode by resetting the unit offsets and modifying the systems' view of the storage units, follow these steps:

1. Shut down the operating systems on all host systems that are accessing the HSG80 controllers you want to change from transparent failover to multiple-bus failover mode.
2. At the HSG80, set multiple-bus failover as follows. Before putting the controllers in multiple-bus failover mode, you must remove any previous failover mode:

```
HSG80> SET NOFAILOVER
HSG80> SET MULTIBUS_FAILOVER COPY=THIS
```

Note

Use the controller known to have the good configuration information.

3. If this HSG80 is being used in an arbitrated loop topology (port topology is set to LOOP_HARD), you need to set a unique AL_PA address for each port because all of the ports can be active at the same time. (The convention in transparent failover mode is to use the same AL_PA address for Port 1 on both controllers and the same AL_PA address for Port 2 on both controllers.)

   The following example sets the ports on two HSG80 controllers off line, sets the PORT_*x*_AL_PA value for multiple-bus failover mode, and sets the ports on line.

```
HSG80> set this port_1_topology = offline
HSG80> set this port_2_topology = offline
HSG80> set other port_1_topology = offline
HSG80> set other port_2_topology = offline
HSG80> set this PORT_1_AL_PA  = 01
HSG80> set this PORT_2_AL_PA  = 02
HSG80> set other PORT_1_AL_PA  = 04
HSG80> set other PORT_2_AL_PA  = 08
```

4. Execute the SHOW CONNECTION command to determine which connections have a nonzero offset as follows:

```
HSG80> SHOW CONNECTION
Connection                                                              Unit
   Name      Operating system    Controller  Port    Address    Status   Offset

!NEWCON49      TRU64_UNIX           THIS       2      230813   OL this     100
           HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01

!NEWCON50      TRU64_UNIX           THIS       1      230813   OL this       0
           HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01

!NEWCON51      TRU64_UNIX           THIS       2      230913   OL this     100
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON52      TRU64_UNIX           THIS       1      230913   OL this       0
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON53      TRU64_UNIX           OTHER      1      230913   OL other      0
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON54      TRU64_UNIX           OTHER      1      230813   OL other      0
           HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01

!NEWCON55      TRU64_UNIX           OTHER      2      230913   OL other    100
           HOST_ID=1000-0000-C920-EDEB         ADAPTER_ID=1000-0000-C920-EDEB

!NEWCON56      TRU64_UNIX           OTHER      2      230813   OL other    100
           HOST_ID=1000-0000-C920-DA01         ADAPTER_ID=1000-0000-C920-DA01

!NEWCON57      TRU64_UNIX           THIS       2               offline     100
           HOST_ID=1000-0000-C921-09F7         ADAPTER_ID=1000-0000-C921-09F7

!NEWCON58      TRU64_UNIX           OTHER      1               offline       0
           HOST_ID=1000-0000-C921-09F7         ADAPTER_ID=1000-0000-C921-09F7

!NEWCON59      TRU64_UNIX           THIS       1               offline       0
           HOST_ID=1000-0000-C921-09F7         ADAPTER_ID=1000-0000-C921-09F7
```

```
!NEWCON60      TRU64_UNIX           OTHER      2               offline     100
           HOST_ID=1000-0000-C921-09F7        ADAPTER_ID=1000-0000-C921-09F7

!NEWCON61      TRU64_UNIX           THIS       2      210513   OL this     100
           HOST_ID=1000-0000-C921-086C        ADAPTER_ID=1000-0000-C921-086C

!NEWCON62      TRU64_UNIX           OTHER      1      210513   OL other      0
           HOST_ID=1000-0000-C921-086C        ADAPTER_ID=1000-0000-C921-086C

!NEWCON63      TRU64_UNIX           OTHER      1               offline       0
           HOST_ID=1000-0000-C921-0943        ADAPTER_ID=1000-0000-C921-0943

!NEWCON64      TRU64_UNIX           OTHER      1      210413   OL other      0
           HOST_ID=1000-0000-C920-EDA0        ADAPTER_ID=1000-0000-C920-EDA0

!NEWCON65      TRU64_UNIX           OTHER      2      210513   OL other    100
           HOST_ID=1000-0000-C921-086C        ADAPTER_ID=1000-0000-C921-086C
     .
     .
     .
```

The following connections are shown to have nonzero offsets: !NEWCON49, !NEWCON51, !NEWCON55, !NEWCON56, !NEWCON57, !NEWCON60, !NEWCON61, and !NEWCON65

5. Set the unit offset to 0 for each connection that has a nonzero unit offset:

```
HSG80> SET !NEWCON49 UNIT_OFFSET = 0
HSG80> SET !NEWCON51 UNIT_OFFSET = 0
HSG80> SET !NEWCON55 UNIT_OFFSET = 0
HSG80> SET !NEWCON56 UNIT_OFFSET = 0
HSG80> SET !NEWCON57 UNIT_OFFSET = 0
HSG80> SET !NEWCON60 UNIT_OFFSET = 0
HSG80> SET !NEWCON61 UNIT_OFFSET = 0
HSG80> SET !NEWCON65 UNIT_OFFSET = 0
```

6. At the console of each system accessing storage units on this HSG80, follow these steps:

Note

The remaining steps apply only to fabric configurations. In this release, you cannot boot from storage connected via a Fibre Channel arbitrated loop.

a. Use the wwid manager to show the Fibre Channel environment variables and determine which units are reachable by the system. This is the information the console uses, when not in wwidmgr mode, to find Fibre Channel devices:

```
P00>>> wwidmgr -show ev
wwid0    133 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002e
wwid1    131 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002f
wwid2    132 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0030
wwid3
N1       50001fe100000d64
N2
N3
N4
```

Note

You must set the console to diagnostic mode to use the `wwidmgr` command for the following AlphaServer systems: AS1200, AS4x00, AS8x00, GS60, GS60E, and GS140. Set the console to diagnostic mode as follows:

```
P00>>> set mode diag
Console is in diagnostic mode
P00>>>
```

b. For each `wwid`*n* line, record the unit number (131, 132, and 133) and worldwide name for the storage unit. The unit number is the first field in the display (after `wwid`*n*). The `N`*n* value is the HSG80 port being used to access the storage units (controller B, port 2).

c. Clear the `wwid`*n* and `N`*n* environment variables:

```
P00>>> wwidmgr -clear all
```

d. Initialize the console:

```
P00>>> init
```

e. Use the `wwid` manager with the `-quickset` option to set up the device and port path information for the storage units from where each system will need to boot. Each system may need to boot from the base operating system disk. Each system will need to boot from its member system boot disk. Using the storage units from the example, cluster member 1 will need access to the storage units with UDIDs 131 (member 1 boot disk) and 133 (Tru64 UNIX disk). Cluster member 2 will need access to the storage units with UDIDs 132 (member 2 boot disk) and 133 (Tru64 UNIX disk). Set up the device and port path for cluster member 1 as follows:

```
P00>>> wwidmgr -quickset -udid 131
       .
       .
P00>>> wwidmgr -quickset -udid 133
       .
       .
```

f. Initialize the console:

```
P00>>> init
```

g. Verify that the storage units and port path information is set up, and then reinitialize the console. The following example shows the information for cluster member 1:

```
P00>>> wwidmgr -show ev
wwid0    133 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002e
wwid1    131 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-002f
wwid2
wwid3
N1       50001fe100000d64
N2       50001fe100000d62
N3       50001fe100000d63
N4       50001fe100000d61
P00>>> init
```

h. Set the `bootdef_dev` console environment variable to the member system boot device. Use the paths shown in the reachability display of the `wwidmgr -quickset` command for the appropriate device (Section 5.6).

i. Repeat steps a through h on each system accessing devices on the HSG80.

7

Using the emx Manager to Display Fibre Channel Adapter Information

The emx manager (emxmgr) utility was written for the TruCluster Software Product Version 1.6 products to be used to modify and maintain emx driver worldwide name to target ID mappings. It is included with Tru64 UNIX Version 5.1 and, although not needed to maintain worldwide name to target ID mappings, it may be used with TruCluster Server Version 5.1 to:

- Display the presence of KGPSA Fibre Channel adapters
- Display the target ID mappings for a Fibre Channel adapter
- Display the current Fibre Channel topology for a Fibre Channel adapter

See emxmgr(8) for more information on the emxmgr utility.

7.1 Using the emxmgr Utility to Display Fibre Channel Adapter Information

The primary use of the emxmgr utility for TruCluster Server is to display Fibre Channel information.

Use the emxmgr -d command to display the presence of KGPSA Fibre Channel adapters on the system. For example:

```
# /usr/sbin/emxmgr -d
emx0 emx1 emx2
```

Use the emxmgr -m command to display an adapter's target ID mapping. For example:

```
# /usr/sbin/emxmgr -m emx0
emx0 SCSI target id assignments:
    SCSI tgt id   0 : portname 5000-1FE1-0000-0CB2
                      nodename 5000-1FE1-0000-0CB0
    SCSI tgt id   5 : portname 1000-0000-C920-A7AE
                      nodename 1000-0000-C920-A7AE
    SCSI tgt id   6 : portname 1000-0000-C920-CD9C
                      nodename 1000-0000-C920-CD9C
    SCSI tgt id   7 : portname 1000-0000-C921-0D00      (emx0)
                      nodename 1000-0000-C921-0D00
```

The previous example shows four Fibre Channel devices on this SCSI bus. The Fibre Channel adapter in question, `emx0`, at SCSI ID 7, is denoted by the presence of the `emx0` designation.

Use the `emxmgr -t` command to display the Fibre Channel topology for the adapter. For example:

```
# emxmgr -t emx1

emx1 state information: [1]
  Link :  connection is UP
          Point to Point
          Fabric attached
          FC DID 0x210413
  Link is SCSI bus 3 (e.g. scsi3)
          SCSI target id 7
          portname is 1000-0000-C921-07C4
          nodename is 1000-0000-C921-07C4
  N_Port at FC DID 0x210013 - SCSI tgt id 5: [2]
    portname 5000-1FE1-0001-8932
    nodename 5000-1FE1-0001-8930
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x210113 - SCSI tgt id 1: [2]
    portname 5000-1FE1-0001-8931
    nodename 5000-1FE1-0001-8930
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x210213 - SCSI tgt id 2: [2]
    portname 5000-1FE1-0001-8941
    nodename 5000-1FE1-0001-8940
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x210313 - SCSI tgt id 4: [2]
    portname 5000-1FE1-0001-8942
    nodename 5000-1FE1-0001-8940
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x210513 - SCSI tgt id 6: [2]
    portname 1000-0000-C921-07F4
    nodename 2000-0000-C921-07F4
    Present, Logged in, FCP Initiator, FCP Target, FCP Logged in,
  N_Port at FC DID 0xfffffc - SCSI tgt id -1: [3]
    portname 20FC-0060-6900-5A1B
    nodename 1000-0060-6900-5A1B
    Present, Logged in, Directory Server,
  N_Port at FC DID 0xfffffe - SCSI tgt id -1: [3]
    portname 2004-0060-6900-5A1B
    nodename 1000-0060-6900-5A1B
    Present, Logged in, F_PORT,
```

[1] Status of the `emx1` link. The connection is a point-to-point fabric (switch) connection, and the link is up. The adapter is on SCSI bus 3 at SCSI ID 7. Both the port name and node name of the adapter (the

worldwide name) are provided. The Fibre Channel DID number is the physical Fibre Channel address being used by the N_Port.

2 A list of all other Fibre Channel devices on this SCSI bus, with their SCSI ID, port name, node name, physical Fibre Channel address and other items such as:

- Present — The adapter indicates that this N_Port is present on the fabric.
- Logged in — The adapter and remote N_Port have exchanged initialization parameters and have an open channel for communications (nonprotocol-specific communications).
- FCP Target — This N_Port acts as a SCSI target device (it receives SCSI commands).
- FCP Logged in — The adapter and remote N_Port have exchanged FCP-specific initialization parameters and have an open channel for communications (Fibre Channel protocol-specific communications).
- Logged Out — The adapter and remote N_Port do not have an open channel for communication.
- FCP Initiator — The remote N_Port acts as a SCSI initiator device (it sends SCSI commands).
- FCP Suspended — The driver has invoked a temporary suspension on SCSI traffic to the N_Port while it resolves a change in connectivity.
- F_PORT — The fabric connection (F_Port) allows the adapter to send Fibre Channel traffic into the fabric.
- Directory Server — The N_Port is the FC entity queried to determine who is present on the Fibre Channel fabric.

3 A target ID of -1 (or -2) shows up for remote Fibre Channel devices that do not communicate using Fibre Channel protocol, the directory server, and F_Port.

Note

You can use the `emxmgr` utility interactively to perform any of the previous functions.

7.1.1 Using the emxmgr Utility in an Arbitrated Loop Topology

The following example shows the results of the `emxmgr -t` command in an arbitrated loop topology.

```
# emxmgr -t emx0

emx0 state information:
  Link :  connection is UP
          FC-AL (Loop) [1]
          FC DID 0x000001
  Link is SCSI bus 2 (e.g. scsi2)
          SCSI target id 7
          portname is 1000-0000-C920-5F0E
          nodename is 1000-0000-C920-5F0E
  N_Port at FC DID 0x000002 - SCSI tgt id 6 :
    portname 1000-0000-C920-043C
    nodename 1000-0000-C920-043C
    Present, Logged in, FCP Initiator, FCP Target, FCP Logged in,
  N_Port at FC DID 0x00006b - SCSI tgt id 2 :
    portname 2200-0020-3704-846F
    nodename 2000-0020-3704-846F
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x00006c - SCSI tgt id 3 :
    portname 2200-0020-3704-A822
    nodename 2000-0020-3704-A822
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x00002d - SCSI tgt id 1 :
    portname 2200-0020-3703-146B
    nodename 2000-0020-3703-146B
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x00002e - SCSI tgt id 0 :
    portname 2200-0020-3703-137D
    nodename 2000-0020-3703-137D
    Present, Logged in, FCP Target, FCP Logged in,
  N_Port at FC DID 0x00006e - SCSI tgt id 4 :
    portname 2200-0020-3700-55CB
    nodename 2000-0020-3700-55CB
    Present, Logged in, FCP Target, FCP Logged in,
```

[1] Status of the `emx0` link. The connection is a Fibre Channel arbitrated loop (FC-AL) connection, and the link is up. The adapter is on SCSI bus 2 at SCSI ID 7. The port name and node name of the adapter are provided.

The Fibre Channel DID number is the physical Fibre Channel address being used by the N_Port.

7.2 Using the emxmgr Utility Interactively

Start the `emxmgr` utility without any command-line options to enter the interactive mode to:

- Display the presence of KGPSA Fibre Channel adapters
- Display the target ID mappings for a Fibre Channel adapter
- Display the current Fibre Channel topology for a Fibre Channel adapter

You have already seen how you can perform these functions from the command line. The same output is available using the interactive mode by selecting the appropriate option (shown in the following example).

When you start the `emxmgr` utility with no command-line options, the default device used is the first Fibre Channel adapter it finds. If you want to perform functions for another adapter, you must change the targeted adapter to the correct adapter. For instance, if `emx0` is present, when you start the `emxmgr` interactively, any commands executed to display information will provide the information for `emx0`.

Notes

The `emxmgr` has an extensive help facility in the interactive mode.

Option 3, "Change Target ID Mappings" is a hold-over from the Tru64 UNIX Version 4.0F product and has no use in the Tru64 UNIX Version 5.1 product. Do not use this option.

An example using the `emxmgr` in the interactive mode follows:

```
# emxmgr

Now issuing commands to : "emx0"

Select Option (against "emx0"):

  1.  View adapter's current Topology
  2.  View adapter's Target Id Mappings
  3.  Change Target ID Mappings

  d.  Display Attached Adapters
  a.  Change targeted adapter
  x.  Exit

  ----> 2

emx0 SCSI target id assignments:
    SCSI tgt id   0 : portname 5000-1FE1-0000-0CB2
                      nodename 5000-1FE1-0000-0CB0
    SCSI tgt id   5 : portname 1000-0000-C920-A7AE
                      nodename 1000-0000-C920-A7AE
    SCSI tgt id   6 : portname 1000-0000-C920-CD9C
                      nodename 1000-0000-C920-CD9C
    SCSI tgt id   7 : portname 1000-0000-C921-0D00      (emx0)
                      nodename 1000-0000-C921-0D00

Select Option (against "emx0"):
```

```
  1.  View adapter's current Topology
  2.  View adapter's Target Id Mappings
  3.  Change Target ID Mappings

  d.  Display Attached Adapters
  a.  Change targeted adapter
  x.  Exit

  ----> x
#
```

A

Worldwide ID-to-Disk Name Conversion Table

Table A–1: Converting Storageset Unit Numbers to Disk Names

File System or Disk	HSG80 Unit	WWID	UDID	Device Name	dsk*n*
Tru64 UNIX disk					
Cluster root (/)					
Member 1 boot disk					
Member 2 boot disk					
Quorum disk					
/var					
/usr					

Index

G

H

I

K

L

M

N

O

P

R

S

T

V

W