

TruCluster Server

Cluster Hardware Configuration

Part Number: AA-RHGWD-TE

June 2001

Product Version: TruCluster Server Version 5.1A

Operating System and Version: Tru64 UNIX Version 5.1A

This manual describes how to configure the hardware for a TruCluster Server environment. TruCluster Server Version 5.1A runs on the Tru64 UNIX operating system.

© 2001 Compaq Computer Corporation

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

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States and other countries. UNIX and The Open Group are trademarks of The Open Group in the United States and other countries. All other product names mentioned herein may be trademarks of their respective companies.

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

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

Contents

About This Manual

1 Introduction

1.1	The TruCluster Server Product	1-1
1.2	Memory Requirements	1-2
1.3	Minimum Disk Requirements	1-2
1.3.1	Disks Needed for Installation	1-2
1.3.1.1	Tru64 UNIX Operating System Disk	1-3
1.3.1.2	Clusterwide Disks	1-3
1.3.1.3	Member Boot Disk	1-4
1.3.1.4	Quorum Disk	1-4
1.4	Generic Two-Node Cluster	1-5
1.5	Growing a Cluster from Minimum Storage to an NSPOF Cluster	1-7
1.5.1	Two-Node Clusters Using an UltraSCSI BA356 Storage Shelf and Minimum Disk Configurations	1-8
1.5.2	Two-Node Clusters Using UltraSCSI BA356 Storage Units with Increased Disk Configurations	1-11
1.5.3	Two-Node Configurations with UltraSCSI BA356 Storage Units and Dual SCSI Buses	1-12
1.5.4	Using Hardware RAID to Mirror the Quorum and Member System Boot Disks	1-14
1.5.5	Creating an NSPOF Cluster	1-15
1.6	Eight-Member Clusters	1-20
1.7	Overview of Setting Up the TruCluster Server Hardware Configuration	1-20

2 Hardware Requirements and Restrictions

2.1	TruCluster Server Member System Requirements	2-1
2.2	Memory Channel Restrictions	2-3
2.3	Host Bus Adapter Restrictions	2-6
2.3.1	Fibre Channel Requirements and Restrictions	2-6
2.3.2	KZPSA-BB SCSI Adapter Restrictions	2-9
2.3.3	KZPBA-CB SCSI Bus Adapter Restrictions	2-10
2.4	Disk Device Restrictions	2-10

2.5	RAID Array Controller Restrictions	2-10
2.6	SCSI Signal Converters	2-12
2.7	DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI Hubs	2-13
2.8	SCSI Cables	2-14
2.9	SCSI Terminators and Trilink Connectors	2-16

3 Shared SCSI Bus Requirements and Configurations Using UltraSCSI Hardware

3.1	Shared SCSI Bus Configuration Requirements	3-2
3.2	SCSI Bus Performance	3-4
3.2.1	SCSI Bus Versus SCSI Bus Segments	3-4
3.2.2	Transmission Methods	3-4
3.2.3	Data Path	3-5
3.2.4	Bus Speed	3-5
3.3	SCSI Bus Device Identification Numbers	3-6
3.4	SCSI Bus Length	3-7
3.5	Terminating the Shared SCSI Bus When Using UltraSCSI Hubs	3-7
3.6	UltraSCSI Hubs	3-8
3.6.1	Using a DWZZH UltraSCSI Hub in a Cluster Configuration	3-9
3.6.1.1	DS-DWZZH-03 Description	3-9
3.6.1.2	DS-DWZZH-05 Description	3-10
3.6.1.2.1	DS-DWZZH-05 Configuration Guidelines	3-11
3.6.1.2.2	DS-DWZZH-05 Fair Arbitration	3-12
3.6.1.2.3	DS-DWZZH-05 Address Configurations	3-13
3.6.1.2.4	SCSI Bus Termination Power	3-15
3.6.1.2.5	DS-DWZZH-05 Indicators	3-15
3.6.1.3	Installing the DS-DWZZH-05 UltraSCSI Hub	3-16
3.7	Preparing the UltraSCSI Storage Configuration	3-16
3.7.1	Configuring Radially Connected TruCluster Server Clusters with UltraSCSI Hardware	3-17
3.7.1.1	Preparing an HSZ70 or HSZ80 for a Shared SCSI Bus Using Transparent Failover Mode	3-19
3.7.1.2	Preparing a Dual-Redundant HSZ70 or HSZ80 for a Shared SCSI Bus Using Multiple-Bus Failover	3-23
3.7.1.3	Preparing an RA3000 for Use on a Shared SCSI Bus with an UltraSCSI Hub	3-26

4 TruCluster Server System Configuration Using UltraSCSI Hardware

4.1	Planning Your TruCluster Server Hardware Configuration ...	4-2
4.2	Obtaining the Firmware Release Notes	4-4
4.3	TruCluster Server Hardware Installation	4-5
4.3.1	Installation of a KZPBA-CB Using Internal Termination for a Radial Configuration	4-7
4.3.2	Displaying KZPBA-CB Adapters with the show Console Commands	4-10
4.3.3	Displaying Console Environment Variables and Setting the KZPBA-CB SCSI ID	4-14
4.3.3.1	Displaying KZPBA-CB pk* or isp* Console Environment Variables	4-15
4.3.3.2	Setting the KZPBA-CB SCSI ID	4-17
4.3.3.3	KZPBA-CB Termination Resistors	4-17

5 Setting Up the Memory Channel Cluster Interconnect

5.1	Setting the Memory Channel Adapter Jumpers	5-2
5.1.1	MC1 and MC1.5 Hub Mode Jumper	5-2
5.1.2	MC2 Jumpers	5-3
5.2	Installing the Memory Channel Adapter	5-6
5.3	Installing the MC2 Optical Converter in the Member System	5-6
5.4	Installing the Memory Channel Hub	5-7
5.5	Installing the Memory Channel Cables	5-7
5.5.1	Installing the MC1 or MC1.5 Cables	5-7
5.5.1.1	Connecting MC1 or MC1.5 Link Cables in Virtual Hub Mode	5-8
5.5.1.2	Connecting MC1 Link Cables in Standard Hub Mode ..	5-8
5.5.2	Installing the MC2 Cables	5-9
5.5.2.1	Installing the MC2 Cables for Virtual Hub Mode Without Optical Converters	5-9
5.5.2.2	Installing MC2 Cables in Virtual Hub Mode Using Optical Converters	5-10
5.5.2.3	Connecting MC2 Link Cables in Standard Hub Mode (No Fiber Optics)	5-10
5.5.2.4	Connecting MC2 Cables in Standard Hub Mode Using Optical Converters	5-11
5.6	Running Memory Channel Diagnostics	5-12
5.7	Maintaining Memory Channel Interconnects	5-14
5.7.1	Adding a Memory Channel Interconnect	5-14

5.7.2	Upgrading Memory Channel Adapters	5-14
5.7.3	Upgrading a Virtual Hub Configuration to a Standard Hub Configuration	5-26

6 Using Fibre Channel Storage

6.1	Fibre Channel Overview	6-2
6.1.1	Basic Fibre Channel Terminology	6-2
6.1.2	Fibre Channel Topologies	6-4
6.1.2.1	Point-to-Point	6-5
6.1.2.2	Fabric	6-5
6.1.2.3	Arbitrated Loop Topology	6-6
6.2	Fibre Channel Topology Comparison	6-7
6.3	Example Fibre Channel Configurations Supported by TruCluster Server	6-8
6.3.1	Fibre Channel Cluster Configurations for Transparent Failover Mode	6-8
6.3.2	Fibre Channel Cluster Configurations for Multiple-Bus Failover Mode	6-10
6.4	QuickLoop	6-17
6.5	Zoning	6-18
6.5.1	Switch Zoning Versus Selective Storage Presentation	6-18
6.5.2	Types of Zoning	6-19
6.5.3	Zoning Example	6-20
6.6	Cascaded Switches	6-21
6.7	Procedure for Installation Using Fibre Channel Disks	6-24
6.8	Installing and Configuring Fibre Channel Hardware	6-26
6.8.1	Installing and Setting Up the Fibre Channel Switch	6-26
6.8.1.1	Installing the Switch	6-27
6.8.1.2	Managing the Fibre Channel Switches	6-29
6.8.1.2.1	Using the Switch Front Panel	6-29
6.8.1.2.2	Setting the Ethernet IP Address and Subnet Mask from the Front Panel	6-30
6.8.1.2.3	Setting the DS-DSGGB-AA and DS-DSGGC-AA/AB Ethernet IP Address and Subnet Mask from a PC or Terminal	6-31
6.8.1.2.4	Logging In to the Switch with a Telnet Connection	6-32
6.8.1.2.5	Setting the Switch Name via Telnet Session	6-33
6.8.2	Installing and Setting Up the DS-SWXHB-07 Hub	6-34
6.8.2.1	Installing the Hub	6-34
6.8.2.2	Determining the Hub Status	6-35
6.8.3	Installing and Configuring the KGPSA PCI-to-Fibre Channel Adapter Module	6-36

6.8.3.1	Installing the KGPSA PCI-to-Fibre Channel Adapter Module	6-36
6.8.3.2	Setting the KGPSA-BC or KGPSA-CA to Run on a Fabric	6-37
6.8.3.3	Setting the KGPSA-CA Adapter to Run in a Loop	6-39
6.8.3.4	Obtaining the Worldwide Names of KGPSA Adapters	6-41
6.8.4	Setting Up the HSG80 Array Controller for Tru64 UNIX Installation	6-42
6.8.4.1	Setting Up the HSG80 Array Controller for Arbitrated Loop	6-48
6.8.4.2	Obtaining the Worldwide Names of HSG80 Controller	6-50
6.9	Installing Tru64 UNIX and TruCluster Server Software Using Fibre Channel Storage	6-51
6.9.1	Before You Install	6-53
6.9.1.1	Configure the HSG80 Storagesets	6-53
6.9.1.2	Adding Units and Identifiers to the HSG80 Storagesets	6-56
6.9.1.3	Set the Device Unit Number	6-61
6.9.1.4	Displaying the Available Boot Devices	6-66
6.9.1.5	Set the bootdef_dev Console Environment Variable for Tru64 UNIX Installation	6-67
6.9.2	Install the Base Operating System	6-69
6.9.3	Determining /dev/disk/dskn to Use for a Cluster Installation	6-69
6.9.4	Label the Disks to Be Used to Create the Cluster	6-71
6.9.5	Install the TruCluster Server Software and Create the First Cluster Member	6-71
6.9.6	Reset the bootdef_dev Console Environment Variable	6-72
6.9.7	Add Additional Systems to the Cluster	6-74
6.10	Converting the HSG80 from Transparent to Multiple-Bus Failover Mode	6-76
6.10.1	Overview	6-76
6.10.2	Procedure to Convert from Transparent to Multiple-bus Failover Mode	6-77
6.11	Using the emx Manager to Display Fibre Channel Adapter Information	6-81
6.11.1	Using the emxmgr Utility to Display Fibre Channel Adapter Information	6-82
6.11.2	Using the emxmgr Utility in an Arbitrated Loop Topology	6-84
6.11.3	Using the emxmgr Utility Interactively	6-85

7 Using GS80, GS160, or GS320 Hard Partitions in a TruCluster Server Configuration

7.1	Overview	7-1
7.2	Hardware Requirements for a Hard Partition in a Cluster	7-2
7.3	Configuring Partitioned GS80, GS160, or GS320 Systems in a TruCluster Configuration	7-7
7.3.1	Repartitioning a Single-Partition AlphaServer GS80/160/320 as Two Partitions in a Cluster	7-7
7.4	Determining AlphaServer GS80/160/320 System Configuration	7-14
7.5	Updating GS80/160/320 Firmware	7-18
7.5.1	Updating AlphaServer GS80/160/320 Firmware	7-18

8 Configuring a Shared SCSI Bus for Tape Drive Use

8.1	Preparing the TZ88 for Shared Bus Usage	8-2
8.1.1	Setting the TZ88N-VA SCSI ID	8-2
8.1.2	Cabling the TZ88N-VA	8-4
8.1.3	Setting the TZ88N-TA SCSI ID	8-6
8.1.4	Cabling the TZ88N-TA	8-6
8.2	Preparing the TZ89 for Shared SCSI Usage	8-7
8.2.1	Setting the DS-TZ89N-VW SCSI ID	8-7
8.2.2	Cabling the DS-TZ89N-VW Tape Drives	8-9
8.2.3	Setting the DS-TZ89N-TA SCSI ID	8-10
8.2.4	Cabling the DS-TZ89N-TA Tape Drives	8-10
8.3	Compaq 20/40 GB DLT Tape Drive	8-11
8.3.1	Setting the Compaq 20/40 GB DLT Tape Drive SCSI ID ..	8-11
8.3.2	Cabling the Compaq 20/40 GB DLT Tape Drive	8-12
8.4	Compaq 40/80-GB DLT Drive	8-15
8.4.1	Setting the Compaq 40/80-GB DLT Drive SCSI ID	8-15
8.4.2	Cabling the Compaq 40/80-GB DLT Drive	8-16
8.5	Preparing the TZ885 for Shared SCSI Usage	8-18
8.5.1	Setting the TZ885 SCSI ID	8-18
8.5.2	Cabling the TZ885 Tape Drive	8-19
8.6	Preparing the TZ887 for Shared SCSI Bus Usage	8-21
8.6.1	Setting the TZ887 SCSI ID	8-21
8.6.2	Cabling the TZ887 Tape Drive	8-22
8.7	Preparing the TL891 and TL892 DLT MiniLibraries for Shared SCSI Usage	8-24
8.7.1	Setting the TL891 or TL892 SCSI ID	8-24
8.7.2	Cabling the TL891 or TL892 MiniLibraries	8-26

8.8	Preparing the TL890 DLT MiniLibrary Expansion Unit	8-30
8.8.1	TL890 DLT MiniLibrary Expansion Unit Hardware	8-30
8.8.2	Preparing the DLT MiniLibraries for Shared SCSI Bus Usage	8-30
8.8.2.1	Cabling the DLT MiniLibraries	8-31
8.8.2.2	Configuring a Base Module as a Slave	8-34
8.8.2.3	Powering Up the DLT MiniLibrary	8-35
8.8.2.4	Setting the TL890/TL891/TL892 SCSI ID	8-36
8.9	Preparing the TL894 DLT Automated Tape Library for Shared SCSI Bus Usage	8-38
8.9.1	TL894 Robotic Controller Required Firmware	8-38
8.9.2	Setting TL894 Robotics Controller and Tape Drive SCSI IDs	8-38
8.9.3	TL894 Tape Library Internal Cabling	8-40
8.9.4	Connecting the TL894 Tape Library to the Shared SCSI Bus	8-42
8.10	Preparing the TL895 DLT Automated Tape Library for Shared SCSI Bus Usage	8-44
8.10.1	TL895 Robotic Controller Required Firmware	8-45
8.10.2	Setting the TL895 Tape Library SCSI IDs	8-45
8.10.3	TL895 Tape Library Internal Cabling	8-46
8.10.4	Upgrading a TL895	8-48
8.10.5	Connecting the TL895 Tape Library to the Shared SCSI Bus	8-49
8.11	Preparing the TL893 and TL896 Automated Tape Libraries for Shared SCSI Bus Usage	8-49
8.11.1	Communications with the Host Computer	8-51
8.11.2	MUC Switch Functions	8-51
8.11.3	Setting the MUC SCSI ID	8-52
8.11.4	Tape Drive SCSI IDs	8-52
8.11.5	TL893 and TL896 Automated Tape Library Internal Cabling	8-53
8.11.6	Connecting the TL893 and TL896 Automated Tape Libraries to the Shared SCSI Bus	8-55
8.12	Preparing the TL881 and TL891 DLT MiniLibraries for Shared Bus Usage	8-58
8.12.1	TL881 and TL891 DLT MiniLibraries Overview	8-58
8.12.1.1	TL881 and TL891 DLT MiniLibrary Tabletop Model ..	8-58
8.12.1.2	TL881 and TL891 MiniLibrary Rackmount Components	8-59
8.12.1.3	TL881 and TL891 Rackmount Scalability	8-60
8.12.1.4	DLT MiniLibrary Part Numbers	8-61

8.12.2	Preparing a TL881 or TL891 MiniLibrary for Shared SCSI Bus Use	8-61
8.12.2.1	Preparing a Tabletop Model or Base Unit for Standalone Shared SCSI Bus Usage	8-62
8.12.2.1.1	Setting the Standalone MiniLibrary Tape Drive SCSI ID	8-62
8.12.2.1.2	Cabling the TL881 or TL891 DLT MiniLibrary	8-64
8.12.2.2	Preparing a TL881 or TL891 Rackmount MiniLibrary for Shared SCSI Bus Usage	8-67
8.12.2.2.1	Cabling the Rackmount TL881 or TL891 DLT MiniLibrary	8-67
8.12.2.2.2	Configuring a Base Unit as a Slave to the Expansion Unit	8-70
8.12.2.2.3	Powering Up the TL881/TL891 DLT MiniLibrary .	8-71
8.12.2.2.4	Setting the SCSI IDs for a Rackmount TL881 or TL891 DLT MiniLibrary	8-72
8.13	Compaq ESL9326D Enterprise Library	8-73
8.13.1	General Overview	8-73
8.13.2	ESL9326D Enterprise Library Overview	8-74
8.13.3	Preparing the ESL9326D Enterprise Library for Shared SCSI Bus Usage	8-75
8.13.3.1	ESL9326D Enterprise Library Robotic and Tape Drive Required Firmware	8-75
8.13.3.2	Library Electronics and Tape Drive SCSI IDs	8-75
8.13.3.3	ESL9326D Enterprise Library Internal Cabling	8-76
8.13.3.4	Connecting the ESL9326D Enterprise Library to the Shared SCSI Bus	8-77

9 Configuring Systems for External Termination or Radial Connections to Non-UltraSCSI Devices

9.1	TruCluster Server Hardware Installation Using PCI SCSI Adapters	9-1
9.1.1	Radial Installation of a KZPSA-BB or KZPBA-CB Using Internal Termination	9-3
9.1.2	Installing a KZPSA-BB or KZPBA-CB Using External Termination	9-6
9.1.3	Displaying KZPSA-BB and KZPBA-CB Adapters with the show Console Commands	9-9
9.1.4	Displaying Console Environment Variables and Setting the KZPSA-BB and KZPBA-CB SCSI ID	9-13
9.1.4.1	Displaying KZPSA-BB and KZPBA-CB pk* or isp* Console Environment Variables	9-13

9.1.4.2	Setting the KZPBA-CB SCSI ID	9-16
9.1.4.3	Setting KZPSA-BB SCSI Bus ID, Bus Speed, and Termination Power	9-17
9.1.4.4	KZPSA-BB and KZPBA-CB Termination Resistors	9-18
9.1.4.5	Updating the KZPSA-BB Adapter Firmware	9-18

10 Configurations Using External Termination or Radial Connections to Non-UltraSCSI Devices

10.1	Using SCSI Bus Signal Converters	10-2
10.1.1	Types of SCSI Bus Signal Converters	10-2
10.1.2	Using the SCSI Bus Signal Converters	10-3
10.1.2.1	DWZZA and DWZZB Signal Converter Termination ..	10-3
10.1.2.2	DS-BA35X-DA Termination	10-4
10.2	Terminating the Shared SCSI Bus	10-5
10.3	Overview of Disk Storage Shelves	10-8
10.3.1	BA350 Storage Shelf	10-9
10.3.2	BA356 Storage Shelf	10-10
10.3.2.1	Non-UltraSCSI BA356 Storage Shelf	10-10
10.3.2.2	UltraSCSI BA356 Storage Shelf	10-13
10.4	Preparing the Storage for Configurations Using External Termination	10-14
10.4.1	Preparing BA350, BA356, and UltraSCSI BA356 Storage Shelves for an Externally Terminated TruCluster Server Configuration	10-15
10.4.1.1	Preparing a BA350 Storage Shelf for Shared SCSI Usage	10-15
10.4.1.2	Preparing a BA356 Storage Shelf for Shared SCSI Usage	10-16
10.4.1.3	Preparing an UltraSCSI BA356 Storage Shelf for a TruCluster Configuration	10-17
10.4.2	Cabling a Single Storage Shelf	10-17
10.4.2.1	Cabling a Single BA350 Storage Shelf	10-18
10.4.2.2	Cabling a Single BA356 Storage Shelf	10-18
10.4.2.3	Cabling a Single UltraSCSI BA356 Storage Shelf	10-18
10.4.3	Connecting Storage Shelves Together	10-18
10.4.3.1	Connecting a BA350 and a BA356 for Shared SCSI Bus Usage	10-19
10.4.3.2	Connecting Two BA356s for Shared SCSI Bus Usage ..	10-21
10.4.3.3	Connecting Two UltraSCSI BA356s for Shared SCSI Bus Usage	10-23

10.4.4	Cabling a Non-UltraSCSI RAID Array Controller to an Externally Terminated Shared SCSI Bus	10-25
10.4.4.1	Cabling an HSZ40 or HSZ50 in a Cluster Using External Termination	10-26
10.4.4.2	Cabling an HSZ20 in a Cluster Using External Termination	10-29
10.4.5	Cabling an RA3000 Using External Termination	10-29
10.5	Cabling an HSZ40 or HSZ50 RAID Array Controller in a Radial Configuration with an UltraSCSI Hub	10-34

11 Configuring an Eight-Member Cluster Using Externally Terminated Shared SCSI Buses

11.1	Overview of an Eight-Node TruCluster Server Cluster	11-2
11.2	Configuring an Eight-Node Cluster Using an UltraSCSI BA356 and External Termination	11-5
11.2.1	Cabling the First Two Nodes on the First Externally Terminated Shared SCSI Cluster	11-5
11.2.2	Cabling the Second Externally Terminated Shared SCSI Bus	11-9
11.2.3	Cabling the Third Externally Terminated Shared SCSI Bus	11-12

A Worldwide ID-to-Disk Name Conversion Table

Index

Examples

4-1	Displaying Configuration on an AlphaServer DS20	4-10
4-2	Displaying Devices on an AlphaServer DS20	4-12
4-3	Displaying Configuration on an AlphaServer 8200	4-13
4-4	Displaying Devices on an AlphaServer 8200	4-13
4-5	Displaying the pk* Console Environment Variables on an AlphaServer DS20 System	4-15
4-6	Displaying Console Variables for a KZPBA-CB on an AlphaServer 8x00 System	4-16
4-7	Setting the KZPBA-CB SCSI Bus ID	4-17
5-1	Running the mc_cable Test	5-13
6-1	Determine HSG80 Connection Names	6-45
6-2	Setting Up the Mirrorset	6-54
6-3	Adding Units and Identifiers to the HSG80 Storagesets	6-57

6-4	Displaying the UDID and Worldwide Names of Devices Known to the Console	6-62
6-5	Setting the Device Unit Number with the wwidmgr quickset Command	6-64
6-6	Sample Fibre Channel Device Names	6-67
6-7	Setting Device Unit Number for Additional Member System ..	6-74
7-1	Defining Hard Partitions with SCM Environment Variables ..	7-10
7-2	Turning Partition Power On	7-12
7-3	Displaying AlphaServer GS160 System Information	7-14
7-4	Displaying Console Serial Bus Information	7-16
9-1	Displaying Configuration on an AlphaServer 4100	9-9
9-2	Displaying Devices on an AlphaServer 4100	9-10
9-3	Displaying Configuration on an AlphaServer 8200	9-11
9-4	Displaying Devices on an AlphaServer 8200	9-12
9-5	Displaying the pk* Console Environment Variables on an AlphaServer 4100 System	9-13
9-6	Displaying Console Variables for a KZPBA-CB on an AlphaServer 8x00 System	9-15
9-7	Displaying Console Variables for a KZPSA-BB on an AlphaServer 8x00 System	9-15
9-8	Setting the KZPBA-CB SCSI Bus ID	9-16
9-9	Setting KZPSA-BB SCSI Bus ID and Speed	9-17

Figures

1-1	Two-Node Cluster with Minimum Disk Configuration and No Quorum Disk	1-6
1-2	Generic Two-Node Cluster with Minimum Disk Configuration and Quorum Disk	1-7
1-3	Minimum Two-Node Cluster with UltraSCSI BA356 Storage Unit	1-9
1-4	Two-Node Cluster with Two UltraSCSI DS-BA356 Storage Units	1-12
1-5	Two-Node Configurations with UltraSCSI BA356 Storage Units and Dual SCSI Buses	1-13
1-6	Cluster Configuration with HSZ70 Controllers in Transparent Failover Mode	1-14
1-7	NSPOF Cluster Using HSZ70s in Multiple-Bus Failover Mode	1-16
1-8	NSPOF Fibre Channel Cluster Using HSG80s in Multiple-Bus Failover Mode	1-17
1-9	NSPOF Cluster Using LSM and UltraSCSI BA356s	1-19

2-1	PCI Backplane Slot Layout	2-3
3-1	VHDCI Trilink Connector (H8861-AA)	3-8
3-2	DS-DWZZH-03 Front View	3-10
3-3	DS-DWZZH-05 Rear View	3-14
3-4	DS-DWZZH-05 Front View	3-15
3-5	Shared SCSI Bus with HSZ70 Configured for Transparent Failover	3-21
3-6	Shared SCSI Bus with HSZ80 Configured for Transparent Failover	3-22
3-7	TruCluster Server Configuration with HSZ70 in Multiple-Bus Failover Mode	3-25
3-8	TruCluster Server Configuration with HSZ80 in Multiple-Bus Failover Mode	3-26
3-9	TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Passive Failover	3-29
3-10	TruCluster Server Configuration with an RA3000 Pedestal with Active/Passive Failover and a DWZZH-05 UltraSCSI Hub	3-30
3-11	TruCluster Server Configuration with an RA3000 Pedestal with Active/Active or Active/Passive Failover	3-31
3-12	TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Active or Active/Passive Failover	3-32
4-1	KZPBA-CB Termination Resistors	4-18
5-1	Connecting Memory Channel Adapters to Hubs	5-9
5-2	MC1-to-MC2 Virtual Hub Upgrade	5-20
5-3	MC1-to-MC2 Standard Hub Upgrade: Initial Configuration ..	5-21
5-4	MC1-to-MC2 Standard Hub Upgrade: First MC1 Module Replaced	5-22
5-5	MC1-to-MC2 Standard Hub Upgrade: Replace First MC1 Adapter in Second System	5-23
5-6	MC1-to-MC2 Standard Hub Upgrade: Replace Third System Memory Channel Adapters	5-24
5-7	MC1-to-MC2 Standard Hub Upgrade: Replace Second MC1 in Second System	5-25
5-8	MC1-to-MC2 Standard Hub Upgrade: Final Configuration ...	5-26
6-1	Point-to-Point Topology	6-5
6-2	Fabric Topology	6-6
6-3	Arbitrated Loop Topology	6-7
6-4	Fibre Channel Single Switch Transparent Failover Configuration	6-9
6-5	Arbitrated Loop Configuration with One Storage Array	6-10
6-6	Multiple-Bus NSPOF Configuration Number 1	6-12
6-7	Multiple-Bus NSPOF Configuration Number 2	6-13

6-8	A Configuration That Is Not Recommended	6-15
6-9	Another Configuration That Is Not Recommended	6-16
6-10	Arbitrated Loop Maximum Configuration	6-17
6-11	A Simple Zoned Configuration	6-21
6-12	Meshed Fabric with Three Cascaded Switches	6-22
6-13	Meshed Resilient Fabric with Four Cascaded Switches	6-23
7-1	Portion of QBB Showing I/O Riser Modules	7-3
7-2	Front View of Expansion and Primary PCI Drawers	7-6
7-3	Rear View of Expansion and Primary PCI Drawers	7-7
8-1	TZ88N-VA SCSI ID Switches	8-3
8-2	Shared SCSI Buses with SBB Tape Drives	8-5
8-3	DS-TZ89N-VW SCSI ID Switches	8-8
8-4	Compaq 20/40 GB DLT Tape Drive Rear Panel	8-12
8-5	Cabling a Shared SCSI Bus with a Compaq 20/40 GB DLT Tape Drive	8-14
8-6	Cabling a Shared SCSI Bus with a Compaq 40/80-GB DLT Drive	8-17
8-7	Cabling a Shared SCSI Bus with a TZ885	8-20
8-8	TZ887 DLT MiniLibrary Rear Panel	8-22
8-9	Cabling a Shared SCSI Bus with a TZ887	8-23
8-10	TruCluster Server Cluster with a TL892 on Two Shared SCSI Buses	8-29
8-11	TL890 and TL892 DLT MiniLibraries on Shared SCSI Buses .	8-33
8-12	TL894 Tape Library Four-Bus Configuration	8-41
8-13	Shared SCSI Buses with TL894 in Two-Bus Mode	8-43
8-14	TL895 Tape Library Internal Cabling	8-48
8-15	TL893 Three-Bus Configuration	8-54
8-16	TL896 Six-Bus Configuration	8-55
8-17	Shared SCSI Buses with TL896 in Three-Bus Mode	8-57
8-18	TL891 Standalone Cluster Configuration	8-66
8-19	TL891 DLT MiniLibrary Rackmount Configuration	8-69
8-20	ESL9326D Internal Cabling	8-76
9-1	KZPSA-BB Termination Resistors	9-18
10-1	Standalone SCSI Signal Converter	10-4
10-2	SBB SCSI Signal Converter	10-4
10-3	DS-BA35X-DA Personality Module Switches	10-5
10-4	BN21W-0B Y Cable	10-7
10-5	HD68 Trilink Connector (H885-AA)	10-8
10-6	BA350 Internal SCSI Bus	10-10
10-7	BA356 Internal SCSI Bus	10-12
10-8	BA356 Jumper and Terminator Module Identification Pins ...	10-13
10-9	BA350 and BA356 Cabled for Shared SCSI Bus Usage	10-20

10-10	Two BA356s Cabled for Shared SCSI Bus Usage	10-22
10-11	Two UltraSCSI BA356s Cabled for Shared SCSI Bus Usage ..	10-24
10-12	Externally Terminated Shared SCSI Bus with Mid-Bus HSZ50 RAID Array Controllers	10-27
10-13	Externally Terminated Shared SCSI Bus with HSZ50 RAID Array Controllers at Bus End	10-28
10-14	Externally Terminated TruCluster Server Configuration with an RA3000 Pedestal with Active/Passive Failover	10-31
10-15	Externally Terminated TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Passive Failover	10-32
10-16	Externally Terminated TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Active or Active/Passive Failover	10-32
10-17	Externally Terminated TruCluster Server Configuration with a Mid-bus RA3000 Controller Shelf with Active/Active or Active/Passive Failover	10-33
10-18	TruCluster Server Cluster Using DS-DWZZH-03, SCSI Adapter with Terminators Installed, and HSZ50	10-36
10-19	TruCluster Server Cluster Using KZPSA-BB SCSI Adapters, a DS-DWZZH-05 UltraSCSI Hub, and an HSZ50 RAID Array Controller	10-37
11-1	Block Diagram of an Eight-Node Cluster	11-3
11-2	First Two Nodes of an Eight-Node Cluster	11-8
11-3	Second Shared SCSI Bus of an Eight-Node Cluster	11-11
11-4	Third Shared SCSI Bus of an Eight-Node Cluster	11-14

Tables

2-1	AlphaServer Systems Supported for Fibre Channel	2-6
2-2	RAID Controller Minimum Required Array Controller Software	2-11
2-3	RAID Controller SCSI IDs	2-11
2-4	Supported SCSI Cables	2-14
2-5	Supported SCSI Terminators and Trilink Connectors	2-16
3-1	SCSI Bus Speeds	3-6
3-2	SCSI Bus Segment Length	3-7
3-3	DS-DWZZH UltraSCSI Hub Maximum Configurations	3-12
3-4	Hardware Components Shown in Figure 3-5 Through Figure 3-8	3-22
3-5	Installing Cables for RA3000 Radial Configuration with a DWZZH UltraSCSI Hub	3-28
3-6	Hardware Components Used in the Configurations Shown in Figure 3-9 through Figure 3-12	3-33

4-1	Planning Your Configuration	4-4
4-2	Configuring TruCluster Server Hardware	4-6
4-3	Installing the KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub	4-9
5-1	MC1 and MC1.5 J4 Jumper Configuration	5-3
5-2	MC2 Jumper Configuration	5-4
5-3	MC2 Linecard Jumper Configurations	5-5
5-4	Adding a Memory Channel Interconnect or Upgrading from a Dual, Redundant MC1 Interconnect to MC2 Interconnects	5-15
5-5	Upgrading from a Virtual Hub Configuration to a Standard Hub Configuration	5-27
6-1	Fibre Channel Fabric and Arbitrated Loop Comparison	6-8
6-2	Type of Zoning Supported by Switches	6-20
6-3	Telnet Session Default User Names for Fibre Channel Switches	6-33
6-4	Converting Storage Unit Numbers to Disk Names	6-60
8-1	TZ88N-VA Switch Settings	8-3
8-2	Hardware Components Used to Create the Configuration Shown in Figure 8-2	8-5
8-3	DS-TZ89N-VW Switch Settings	8-8
8-4	Hardware Components Used to Create the Configuration Shown in Figure 8-5	8-14
8-5	Hardware Components in the Configuration in Figure 8-6	8-17
8-6	Hardware Components Used to Create the Configuration Shown in Figure 8-7	8-20
8-7	Hardware Components Used to Create the Configuration Shown in Figure 8-10	8-29
8-8	Hardware Components Used to Create the Configuration Shown in Figure 8-11	8-34
8-9	TL894 Default SCSI ID Settings	8-38
8-10	Hardware Components Used to Create the Configuration Shown in Figure 8-13	8-43
8-11	TL895 Default SCSI ID Settings	8-45
8-12	MUC Switch Functions	8-51
8-13	MUC SCSI ID Selection	8-52
8-14	TL893 Default SCSI IDs	8-52
8-15	TL896 Default SCSI IDs	8-53
8-16	Hardware Components Used to Create the Configuration Shown in Figure 8-17	8-57
8-17	TL881 and TL891 MiniLibrary Performance and Capacity Comparison	8-60
8-18	DLT MiniLibrary Part Numbers	8-61

8-19	Hardware Components Used to Create the Configuration Shown in Figure 8-18	8-67
8-20	Hardware Components Used to Create the Configuration Shown in Figure 8-19	8-70
8-21	Shared SCSI Bus Cable and Terminator Connections for the ESL9326D Enterprise Library	8-78
9-1	Configuring TruCluster Server Hardware for Use with a PCI SCSI Adapter	9-2
9-2	Installing the KZPSA-BB or KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub	9-5
9-3	Installing a KZPSA-BB or KZPBA-CB for Use with External Termination	9-7
10-1	Hardware Components Used for Configuration Shown in Figure 10-9 and Figure 10-10	10-21
10-2	Hardware Components Used for Configuration Shown in Figure 10-11	10-25
10-3	Hardware Components Used for Configuration Shown in Figure 10-12 and Figure 10-13	10-28
10-4	Installing Cables for RA3000 Configuration Using External Termination and Y Cables	10-30
10-5	Hardware Components Used in the TruCluster Server Configuration Shown in Figure 10-14, Figure 10-15, and Figure 10-16	10-33
10-6	Hardware Components Used in the Configuration Shown in Figure 10-17	10-34
10-7	Hardware Components Used in Configuration Shown in Figure 10-18 and Figure 10-19	10-36
11-1	Hardware Components Used for Configuration Shown in Figure 11-2	11-9
11-2	Hardware Components Used for Configuration Shown Figure 11-3	11-12
11-3	Hardware Components Used for Configuration Shown in Figure 11-4	11-15
A-1	Converting Storage Unit Numbers to Disk Names	A-1

About This Manual

This manual describes how to set up and maintain the hardware configuration for a TruCluster™ Server cluster.

Audience

This manual is for system administrators who will set up and configure the hardware before installing the TruCluster Server software. The manual assumes that you are familiar with the tools and methods that are needed to maintain your hardware, operating system, and network.

New and Changed Features

The following changes have been made to this manual since the Version 5.1 release:

- *Section 1.6* is a new section that introduces eight-node clusters.
- *Section 3.7.1.3* is a new section that introduces the RA3000 and discusses how to configure a RAID array 3000 (RA3000) for use on a shared SCSI bus with UltraSCSI hubs.
- *Section 5.7.3* is a new section that explains how to upgrade a virtual hub Memory Channel configuration to a standard hub configuration.
- *Chapter 9* and *Chapter 10* were re-ordered and rewritten to provide better continuity between the chapters, and to make them flow better.
- *Section 10.4.2* is a new section that discusses how to cable a single storage shelf onto a shared SCSI bus.
- *Section 10.4.5* is a new section that discusses how to cable an RA3000 onto a shared SCSI bus using external termination.
- *Chapter 11* is a new chapter that discusses how to create an externally terminated eight-node cluster.

Organization

This manual is organized as follows:

Chapter 1 Introduces the TruCluster Server product and provides an overview of setting up TruCluster Server hardware.

Chapter 2 Describes hardware requirements and restrictions.

- Chapter 3* Contains information about setting up a shared SCSI bus, SCSI bus requirements, and how to connect storage to a shared SCSI bus using the latest UltraSCSI products (DS-DWZZH UltraSCSI hubs, and HSZ70 and HSZ80 RAID array controllers).
- Chapter 4* Describes how to prepare systems for a TruCluster Server configuration, and how to connect host bus adapters to shared storage using the DS-DWZZH UltraSCSI hubs and the HSZ70 and HSZ80 RAID array controllers.
- Chapter 5* Describes how to set up the Memory Channel cluster interconnect, and how to upgrade Memory Channel interconnects.
- Chapter 6* Provides an overview of Fibre Channel and describes how to set up Fibre Channel hardware.
- Chapter 7* Describes the use of AlphaServer™ GS80, GS160, or GS320 hardware partitions in a TruCluster Server configuration.
- Chapter 8* Describes how to configure a shared SCSI bus for tape drive, tape loader, or tape library usage.
- Chapter 9* Describes how to prepare systems and host bus adapters for the TruCluster Server configurations described in *Chapter 10*.
- Chapter 10* Describes the requirements for a shared SCSI bus using externally terminated configurations and radial configurations using non-UltraSCSI RAID array controllers.
- Chapter 11* Describes how to configure an externally terminated eight-node cluster.
- Appendix A* Provides a blank table to use to convert from the HSG80 unit numbers to `/dev/disk/dskn` and device names for an actual Fibre Channel TruCluster Server configuration.

Related Documents

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

- *TruCluster Server Software Product Description (SPD)* — The comprehensive description of the TruCluster Server Version 5.1A product. You can find the latest version of the SPD at: http://www.tru64unix.compaq.com/docs/pub_page/spds.html
- *Cluster Release Notes* — Provides important information about TruCluster Server Version 5.1A, including new features, known problems, and workarounds.
- *Cluster Technical Overview* — Provides an overview of the TruCluster Server technology.
- *Cluster Installation* — Describes how to install the TruCluster Server product.

- *Cluster Administration* — Describes cluster-specific administration tasks.
- *Cluster Highly Available Applications* — Describes how to deploy applications on a TruCluster Server cluster and how to write cluster-aware applications.
- *Cluster LAN Interconnect* — Describes how to install and configure LAN hardware for the cluster interconnect.

You can find the latest version of the TruCluster Server documentation at the following URL: http://www.tru64unix.com-paq.com/docs/pub_page/cluster_list.html

Consult the following AlphaServer GS80/160/320 documentation to assist you in configuring an AlphaServer GS80, GS160, or GS320 system in a TruCluster Server configuration:

- *Installation Guide*
- *System Management Console Installation and User's Guide*
- *User's Guide*
- *Firmware Reference Manual*

The StorageWorks™ *UltraSCSI Configuration Guidelines* document provides guidelines regarding UltraSCSI configurations.

For information about setting up a redundant array of independent disks (RAID) subsystem, see the following manuals as appropriate for your configuration:

- *DEC RAID Subsystem User's Guide*
- *HS Family of Array Controllers User's Guide*
- *RAID Array 310 Configuration and Maintenance Guide User's Guide*
- *Configuring Your StorageWorks Subsystem HSZ40 Array Controllers HSOF Version 3.0*
- *Getting Started RAID Array 450 V5.4 for Compaq Tru64™ UNIX Installation Guide*
- *HSZ70 Array Controller HSOF Version 7.0 Configuration Manual*
- *HSZ80 Array Controller ACS Version 8.2*
- *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*
- *RAID Array 3000 Controller Shelf Hardware User's Guide*
- *RAID Array 3000 Pedestal Storage Subsystem Hardware User's Guide*
- *RAID Array 3000 Subsystem Second Controller Option Installation Guide*
- *RAID Array 3000 Storage Subsystem Expansion Pedestal Option Installation Guide*
- *Command Console V2.2 for the RAID Array 3000 (Pedestal and Rack Mount Models) User's Guide*
- *Getting Started RAID Array 3000 for Tru64 UNIX Installation Guide*

Consult the following documentation for other Fibre Channel storage information:

- *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 SAN Switch Fabric Operating System Management Guide*

- *Compaq StorageWorks Heterogeneous Open SAN Design Reference Guide*
- *Fibre Channel Storage Hub 7 Installation Guide*
- *Fibre Channel Storage Hub 7 Rack Mounting Installation Card*

For information about the tape devices, see the following manuals:

- *TZ88 DLT Series Tape Drive Owner's Manual*
- *TZ89 DLT Series Tape Drive User's Guide*
- *TZ885 Model 100/200 GB DLT 5-Cartridge MiniLibrary Owner's Manual*
- *TZ887 Model 140/280 GB DLT 7-Cartridge MiniLibrary Owner's Manual*
- *TL881 MiniLibrary System User's Guide*
- *TL881 MiniLibrary Drive Upgrade Procedure*
- *Pass-Through Expansion Kit Installation Instructions*
- *TL891 MiniLibrary System User's Guide*
- *TL81X/TL894 Automated Tape Library for DLT Cartridges Facilities Planning and Installation Guide*
- *TL81X/TL894 Automated Tape Library for DLT Cartridges Diagnostic Software User's Manual*
- *TL895 DLT Tape Library Facilities Planning and Installation Guide*
- *TL895 DLT Library Operator's Guide*
- *TL895 DLT Tape Library Diagnostic Software User's Manual*
- *TL895 Drive Upgrade Instructions*
- *TL82X/TL893/TL896 Automated Tape Library for DLT Cartridges Facilities Planning and Installation Guide*
- *TL82X/TL893/TL896 Automated Tape Library for DLT Cartridges Operator's Guide*
- *TL82X/TL893/TL896 Automated Tape Library for DLT Cartridges Diagnostic Software User's Manual*
- *TL82X Cabinet-to-Cabinet Mounting Instructions*
- *TL82X/TL89X MUML to MUSL Upgrade Instructions*
- For more information on the ESL9326D Enterprise Library, see the following Compaq StorageWorks ESL9000 Series Tape Library documentation:
 - *Unpacking Guide*
 - *Reference Guide*
 - *Maintenance and Service Guide*

– *ESL9326 Tape Drive Upgrade Guide*

The Golden Eggs Visual Configuration Guide provides configuration diagrams of workstations, servers, storage components, and clustered systems. It is available on line in PostScript and Portable Document Format (PDF) formats at: <http://www.compaq.com/info/golden-eggs>

At this URL you will find links to individual system, storage, or cluster configurations. You can order the document through the Compaq Literature Order System (LOS) as order number EC-R026B-36.

In addition, have available the following manuals from the Compaq Tru64™ UNIX operating system software documentation set:

- *Installation Guide*
- *Release Notes for Version 5.1A*
- *System Administration*
- *Network Administration: Connections*
- *Network Administration: Services*

Also have available the hardware documentation for the systems, SCSI controllers, disk storage shelves or RAID controllers, and any other hardware you plan to install.

Documentation for the following optional software products will be useful if you intend to use these products with TruCluster Server:

- Compaq Analyze
- DECEvent
- Logical Storage Manager (LSM)
- NetWorker
- Advanced File System (AdvFS) Utilities
- Performance Manager

Icons on Tru64 UNIX Printed Manuals

The printed version of the Tru64 UNIX documentation uses letter icons on the spines of the manuals to help specific audiences quickly find the manuals that meet their needs. (You can order the printed documentation from Compaq.) The following list describes this convention:

- G Manuals for general users
- S Manuals for system and network administrators

- P Manuals for programmers
- R Manuals for reference page users

Some manuals in the documentation help meet the needs of several audiences. For example, the information in some system manuals is also used by programmers. Keep this in mind when searching for information on specific topics.

The *Documentation Overview* provides information on all of the manuals in the Tru64 UNIX documentation set.

Reader's Comments

Compaq welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32
- Internet electronic mail: `readers_comment@zk3.dec.com`

A Reader's Comment form is located on your system in the following location:

```
/usr/doc/readers_comment.txt
```

Please include the following information along with your comments:

- The full title of the manual and the order number. (The order number appears on the title page of printed and PDF versions of a manual.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate Compaq technical support office. Information provided with the software media explains how to send problem reports to Compaq.

Conventions

The following typographical conventions are used in this manual:

#	A number sign represents the superuser prompt.
% cat	Boldface type in interactive examples indicates typed user input.
<i>file</i>	Italic (slanted) type indicates variable values, placeholders, and function argument names.
⋮	A vertical ellipsis indicates that a portion of an example that would normally be present is not shown.
cat(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, <code>cat(1)</code> indicates that you can find information on the <code>cat</code> command in Section 1 of the reference pages.
cluster	Bold text indicates a term that is defined in the glossary.

1

Introduction

This chapter introduces the TruCluster Server product and some basic cluster hardware configuration concepts.

The chapter discusses the following topics:

- An overview of the TruCluster Server product (Section 1.1)
- TruCluster Server memory requirements (Section 1.2)
- TruCluster Server minimum disk requirements (Section 1.3)
- A description of a generic two-node cluster with the minimum disk layout (Section 1.4)
- How to grow a cluster to a no-single-point-of-failure (NSPOF) cluster (Section 1.5)
- An overview of eight-member clusters (Section 1.6)
- An overview of setting up the TruCluster Server hardware configuration (Section 1.7)

Subsequent chapters describe how to set up and maintain TruCluster Server hardware configurations. See the TruCluster Server *Cluster Installation* manual for information about software installation; see the *Cluster Administration* manual for detailed information about setting up member systems; see the *Cluster Highly Available Applications* manual for detailed information about setting up highly available applications.

1.1 The TruCluster Server Product

TruCluster Server extends single-system management capabilities to clusters. It provides a clusterwide namespace for files and directories, including a single root file system that all cluster members share. It also offers a cluster alias for the Internet protocol suite (TCP/IP) so that a cluster appears as a single system to its network clients.

TruCluster Server preserves the availability and performance features found in the earlier TruCluster products:

- Like the TruCluster Available Server Software and TruCluster Production Server products, TruCluster Server lets you deploy highly available applications that have no embedded knowledge that they are

executing in a cluster. They can access their disk data from any member in the cluster.

- Like the TruCluster Production Server Software product, TruCluster Server lets you run components of distributed applications in parallel, providing high availability while taking advantage of cluster-specific synchronization mechanisms and performance optimizations.

TruCluster Server augments the feature set of its predecessors by allowing all cluster members access to all file systems and all storage in the cluster, regardless of where they reside. From the viewpoint of clients, a TruCluster Server cluster appears to be a single system; from the viewpoint of a system administrator, a TruCluster Server cluster is managed as if it were a single system. Because TruCluster Server has no built-in dependencies on the architectures or protocols of its private cluster interconnect or shared storage interconnect, you can more easily alter or expand your cluster's hardware configuration as newer and faster technologies become available.

1.2 Memory Requirements

The base operation system sets a minimum requirement for the amount of memory required to install Tru64 UNIX. In a cluster, each member must have at least 64 MB more than this minimum requirement. For example, if the base operating system requires 128 MB of memory, each system used in a cluster must have at least 192 MB of memory.

1.3 Minimum Disk Requirements

This section provides an overview of the minimum file system or disk requirements for a two-node cluster. For more information on the amount of space required for each required cluster file system, see the *Cluster Installation* manual.

1.3.1 Disks Needed for Installation

You need to allocate disks for the following uses:

- 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.
- One or more disks on a shared SCSI bus to hold the clusterwide root (/), /usr, and /var Advanced File System (AdvFS) file systems.
- One disk per member, normally on a shared SCSI bus, to hold member boot partitions.

- Optionally, one disk on a shared SCSI bus to act as the quorum disk (see Section 1.3.1.4). For a more detailed discussion of the quorum disk, see the *Cluster Administration* manual.

The following sections provide more information about these disks. Figure 1–1 shows a generic two-member cluster with the required file systems.

1.3.1.1 Tru64 UNIX Operating System Disk

The Tru64 UNIX operating system is installed using AdvFS file systems on one or more disks that are accessible to the system that will become the first cluster member. For example:

```
dsk0a      root_domain#root
dsk0g      usr_domain#usr
dsk0h      var_domain#var
```

The operating system disk (Tru64 UNIX disk) cannot be used as a clusterwide disk, as a member boot disk, or as the quorum disk.

Because the Tru64 UNIX operating system will be available on the first cluster member, in an emergency, after shutting down the cluster, you have the option of booting the Tru64 UNIX operating system and attempting to fix the problem. See the *Cluster Administration* manual for more information.

1.3.1.2 Clusterwide Disks

When you create a cluster, the installation scripts copy the Tru64 UNIX root (/), /usr, and /var file systems from the Tru64 UNIX disk to the disk or disks you specify.

We recommend that the disk or disks that you use for the clusterwide file systems be placed on a shared SCSI bus so that all cluster members have access to these disks.

During the installation, you supply the disk device names and partitions that will contain the clusterwide root (/), /usr, and /var file systems. For example, dsk3b, dsk4c, and dsk3g:

```
dsk3b      cluster_root#root
dsk4c      cluster_usr#usr
dsk3g      cluster_var#var
```

The /var fileset cannot share the cluster_usr domain, but must be a separate domain, cluster_var. Each AdvFS file system must be a separate partition; the partitions do not have to be on the same disk.

If any partition on a disk is used by a clusterwide file system, only clusterwide file systems can be on that disk. A disk containing a clusterwide

file system cannot also be used as the member boot disk or as the quorum disk.

1.3.1.3 Member Boot Disk

Each member has a boot disk. A boot disk contains that member's boot, swap, and cluster-status partitions. For example, `dsk1` is the boot disk for the first member and `dsk2` is the boot disk for the second member:

```
dsk1      first member's boot disk  [pepicelli]
dsk2      second member's boot disk [polishham]
```

The installation scripts reformat each member's boot disk to contain three partitions: an `a` partition for that member's root (`/`) file system, a `b` partition for swap, and an `h` partition for cluster status information. (There are no `/usr` or `/var` file systems on a member's boot disk.)

A member boot disk cannot contain one of the clusterwide root (`/`), `/usr`, and `/var` file systems. Also, a member boot disk cannot be used as the quorum disk. A member disk can contain more than the three required partitions. You can move the swap partition off the member boot disk. See the *Cluster Administration* manual for more information.

1.3.1.4 Quorum Disk

The quorum disk allows greater availability for clusters consisting of two members. Its `h` partition contains cluster status and quorum information. See the *Cluster Administration* manual for a discussion of how and when to use a quorum disk.

The following restrictions apply to the use of a quorum disk:

- A cluster can have only one quorum disk.
- The quorum disk should be on a shared bus to which all cluster members are directly connected. If it is not, members that do not have a direct connection to the quorum disk may lose quorum before members that do have a direct connection to it.
- The quorum disk must not contain any data. The `clu_quorum` command will overwrite existing data when initializing the quorum disk. The integrity of data (or file system metadata) placed on the quorum disk from a running cluster is not guaranteed across member failures.

This means that the member boot disks and the disk holding the clusterwide root (`/`) cannot be used as quorum disks.

- The quorum disk can be small. The cluster subsystems use only 1 MB of the disk.

- A quorum disk can have either 1 vote or no votes. In general, a quorum disk should always be assigned a vote. You might assign an existing quorum disk no votes in certain testing or transitory configurations, such as a one-member cluster (in which a voting quorum disk introduces a single point of failure).
- You cannot use the Logical Storage Manager (LSM) on the quorum disk.

1.4 Generic Two-Node Cluster

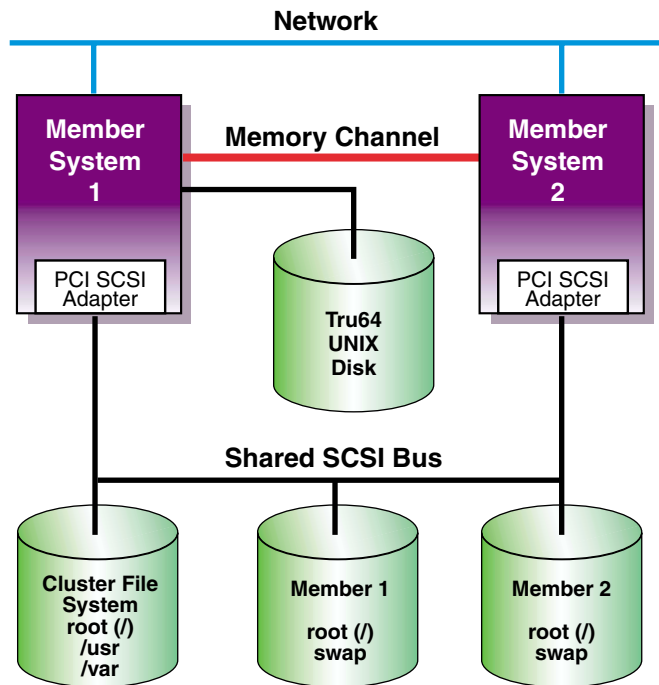
This section describes a generic two-node cluster with the minimum disk layout of four disks. Additional disks may be needed for highly available applications. In this section, and the following sections, the type of peripheral component interconnect (PCI) SCSI bus adapter is not significant. Also, although an important consideration, SCSI bus cabling, including Y cables or trilink connectors, termination, the use of UltraSCSI hubs, and the use of Fibre Channel are not considered at this time.

Figure 1–1 shows a generic two-node cluster with the minimum number of disks.

- Tru64 UNIX disk
- Clusterwide root (/), /usr, and /var
- Member 1 boot disk
- Member 2 boot disk

A minimum configuration cluster may have reduced availability due to the lack of a quorum disk. As shown, with only two-member systems, both systems must be operational to achieve quorum and form a cluster. If only one system is operational, it will loop, waiting for the second system to boot before a cluster can be formed. If one system crashes, you lose the cluster.

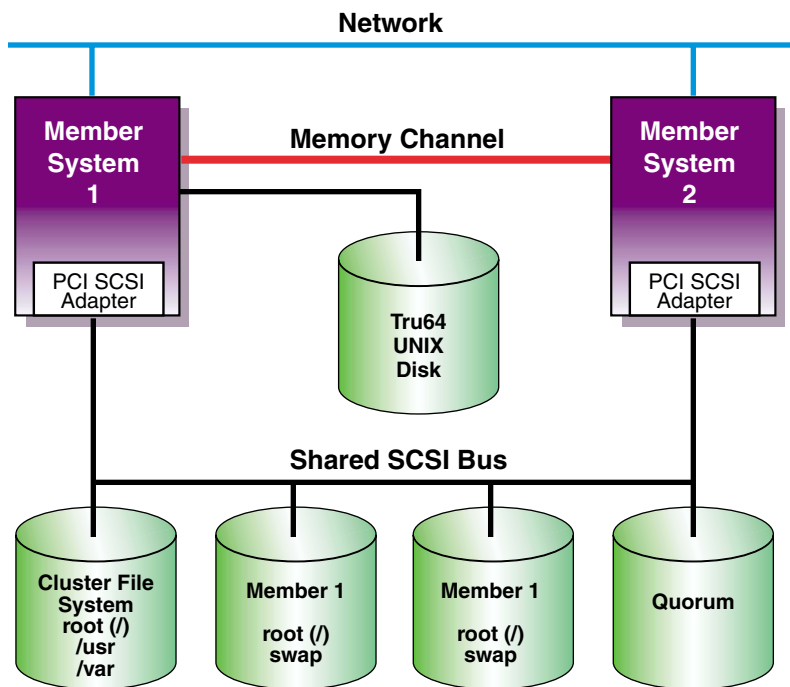
Figure 1–1: Two-Node Cluster with Minimum Disk Configuration and No Quorum Disk



ZK-1587U-AI

Figure 1–2 shows the same generic two-node cluster as shown in Figure 1–1, but with the addition of a quorum disk. By adding a quorum disk, a cluster may be formed if both systems are operational, or if either of the systems and the quorum disk is operational. This cluster has a higher availability than the cluster shown in Figure 1–1. See the *Cluster Administration* manual for a discussion of how and when to use a quorum disk.

Figure 1–2: Generic Two-Node Cluster with Minimum Disk Configuration and Quorum Disk



ZK-1588U-AI

1.5 Growing a Cluster from Minimum Storage to an NSPOF Cluster

The following sections take a progression of clusters from a cluster with minimum storage to a no-single-point-of-failure (NSPOF) cluster — a cluster where one hardware failure will not interrupt the cluster operation:

- The starting point is a cluster with minimum storage for highly available applications (Section 1.5.1).
- By adding a second storage shelf, you have a cluster with more storage for applications, but the single SCSI bus is a single point of failure (Section 1.5.2).
- Adding a second SCSI bus allows the use of LSM to mirror the clusterwide root (/), /usr, and /var file systems, the member system swap disks, and the data disks. However, because LSM cannot mirror the member system boot or quorum disks, full redundancy is not achieved (Section 1.5.3).

- Using a redundant array of independent disks (RAID) array controller in transparent failover mode allows the use of hardware RAID to mirror the disks. However, without a second SCSI bus, second Memory Channel, and redundant networks, this configuration is still not an NSPOF cluster (Section 1.5.4).
- By using an HSZ70, HSZ80, or HSG80 with multiple-bus failover enabled, you can use two shared SCSI buses to access the storage. Hardware RAID is used to mirror the root (/), /usr, and /var file systems, and the member system boot disks, data disks, and quorum disk (if used). A second Memory Channel, redundant networks, and redundant power must also be installed to achieve an NSPOF cluster (Section 1.5.5).

Note

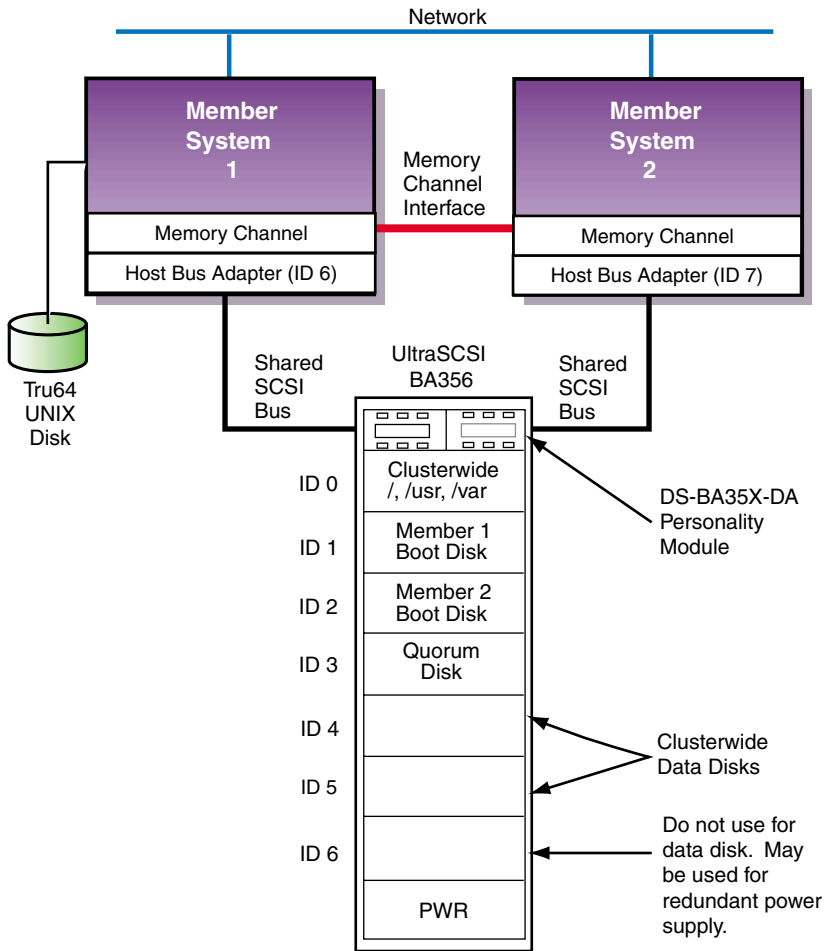
The figures in this section are generic drawings and do not show shared SCSI bus termination, cable names, and so forth.

1.5.1 Two-Node Clusters Using an UltraSCSI BA356 Storage Shelf and Minimum Disk Configurations

This section takes the generic illustrations of our cluster example one step further by depicting the required storage in storage shelves. The storage shelves can be BA350, BA356 (non-UltraSCSI), or UltraSCSI BA356s. The BA350 is the oldest model, and can only respond to SCSI IDs 0-6. The non-Ultra BA356 can respond to SCSI IDs 0-6 or 8-14 (see Section 3.2). The UltraSCSI BA356 also responds to SCSI IDs 0-6 or 8-14, but also can operate at UltraSCSI speeds (see Section 3.2).

Figure 1-3 shows a TruCluster Server configuration using an UltraSCSI BA356 storage unit. The DS-BA35X-DA personality module used in the UltraSCSI BA356 storage unit is a differential-to-single-ended signal converter, and therefore accepts differential inputs.

Figure 1–3: Minimum Two-Node Cluster with UltraSCSI BA356 Storage Unit



ZK-1591U-AI

The configuration shown in Figure 1–3 might represent a typical small or training configuration with TruCluster Server Version 5.1A required disks.

In this configuration, because of the TruCluster Server Version 5.1A disk requirements, only two disks are available for highly available applications.

Note

Slot 6 in the UltraSCSI BA356 is not available because SCSI ID 6 is generally used for a member system SCSI adapter. However,

this slot can be used for a second power supply to provide fully redundant power to the storage shelf.

With the use of the cluster file system (see the *Cluster Administration* manual for a discussion of the cluster file system), the clusterwide root (/), /usr, and /var file systems can be physically placed on a private bus of either of the member systems. But, if that member system is not available, the other member systems do not have access to the clusterwide file systems. Therefore, we do not recommend placing the clusterwide root (/), /usr, and /var file systems on a private bus.

Likewise, the quorum disk can be placed on the local bus of either of the member systems. If that member is not available, quorum can never be reached in a two-node cluster. We do not recommend placing the quorum disk on the local bus of a member system because it creates a single point of failure.

The individual member boot and swap partitions can also be placed on a local bus of either of the member systems. If the boot disk for member system 1 is on a SCSI bus internal to member 1, and the system is unavailable due to a boot disk problem, other systems in the cluster cannot access the disk for possible repair. If the member system boot disks are on a shared SCSI bus, they can be accessed by other systems on the shared SCSI bus for possible repair.

By placing the swap partition on a system's internal SCSI bus, you reduce total traffic on the shared SCSI bus by an amount equal to the system's swap volume.

TruCluster Server Version 5.1A 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.

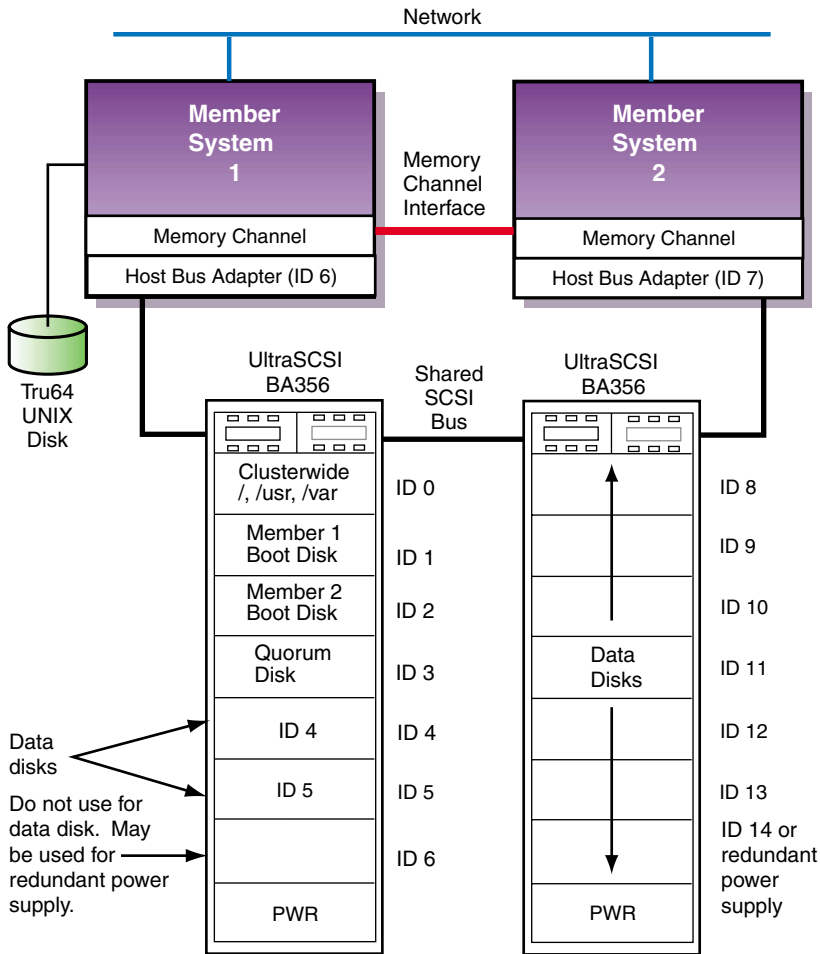
We recommend that you place the clusterwide root (/), /usr, and /var file systems, member boot disks, and quorum disk on a shared SCSI bus that is connected to all member systems. After installation, you have the option to reconfigure swap and can place the swap disks on an internal SCSI bus to increase performance. See the *Cluster Administration* manual for more information.

1.5.2 Two-Node Clusters Using UltraSCSI BA356 Storage Units with Increased Disk Configurations

The configuration shown in Figure 1–3 is a minimal configuration, with a lack of disk space for highly available applications. Starting with Tru64 UNIX Version 5.0, 16 devices are supported on a SCSI bus. Therefore, multiple BA356 storage units can be used on the same SCSI bus to allow more devices on the same bus.

Figure 1–4 shows the configuration in Figure 1–3 with a second UltraSCSI BA356 storage unit that provides an additional seven disks for highly available applications.

Figure 1–4: Two-Node Cluster with Two UltraSCSI DS-BA356 Storage Units



ZK-1590U-AI

This configuration, while providing more storage, has a single SCSI bus that presents a single point of failure. Providing a second SCSI bus can allow the use of the Logical Storage Manager (LSM) to mirror the clusterwide root (/), /usr, and /var file systems, and the data disks across SCSI buses, removing the single SCSI bus as a single point of failure for these file systems.

1.5.3 Two-Node Configurations with UltraSCSI BA356 Storage Units and Dual SCSI Buses

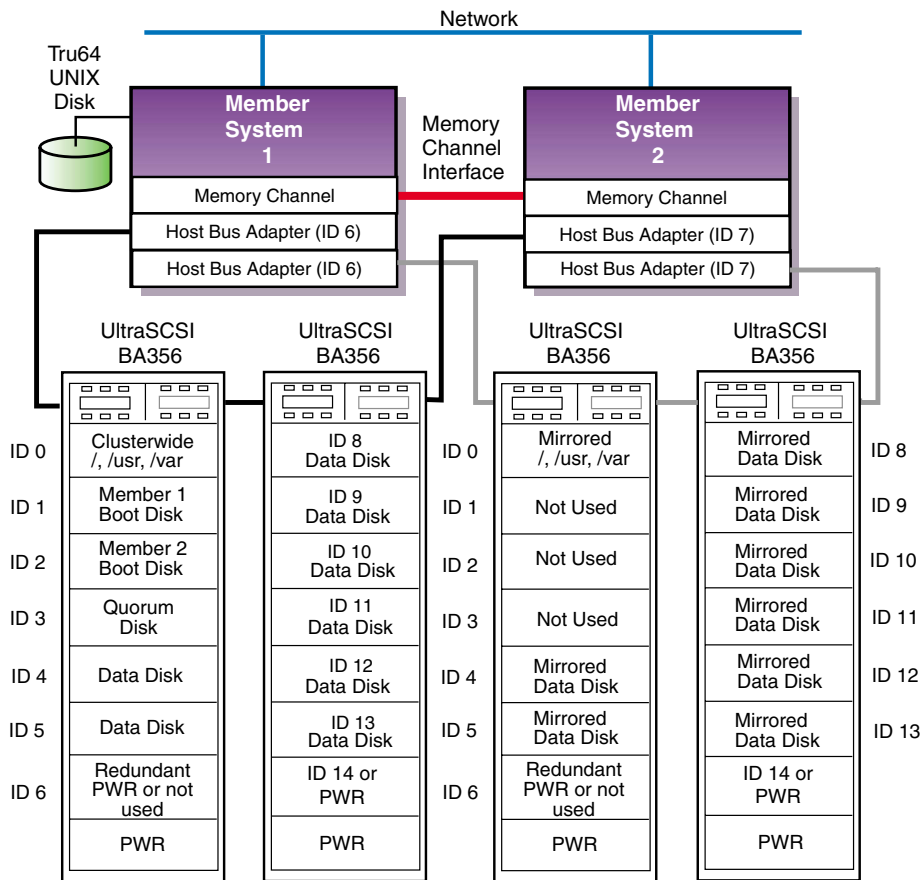
By adding a second shared SCSI bus, you now have the capability to use LSM to mirror data disks, and the clusterwide root (/), /usr, and /var file systems across SCSI buses.

Note

You cannot use LSM to mirror the member system boot or quorum disks, but you can use hardware RAID.

Figure 1–5 shows a small cluster configuration with dual SCSI buses using LSM to mirror the clusterwide root (/), /usr, and /var file systems and the data disks.

Figure 1–5: Two-Node Configurations with UltraSCSI BA356 Storage Units and Dual SCSI Buses



ZK-1811U-AI

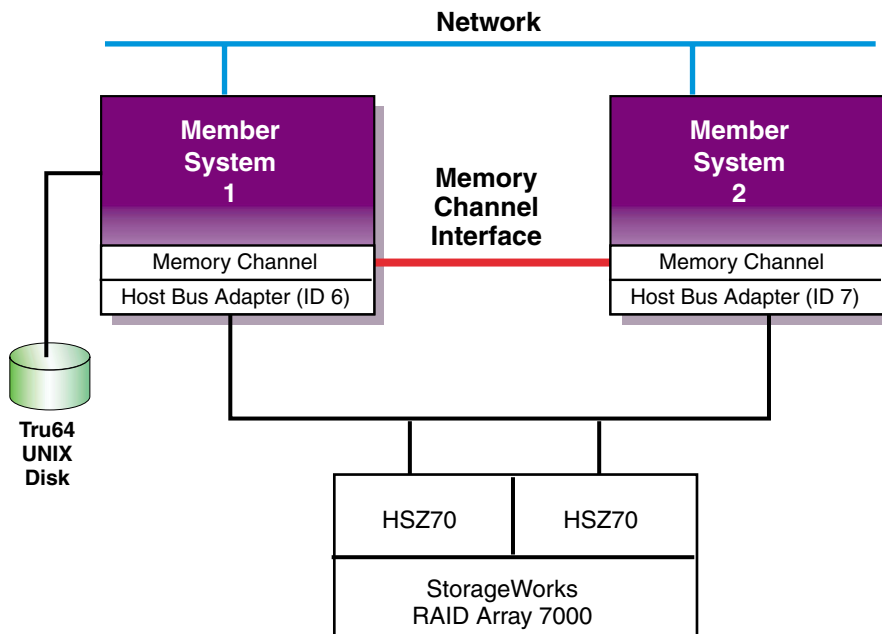
By using LSM to mirror the clusterwide root (/), /usr and /var file systems and the data disks, we have achieved higher availability. But, even if you have a second Memory Channel and redundant networks, because we cannot

use LSM to mirror the quorum or the member system boot disks, we do not have a no-single-point-of-failure (NSPOF) cluster.

1.5.4 Using Hardware RAID to Mirror the Quorum and Member System Boot Disks

You can use hardware RAID with any of the supported RAID array controllers to mirror the quorum and member system boot disks. Figure 1–6 shows a cluster configuration using an HSZ70 RAID array controller. An HSZ40, HSZ50, HSZ80, HSG60, or HSG80, or RAID array 3000 (with HSZ22 controller) can be used instead of the HSZ70. The array controllers can be configured as a dual redundant pair. If you want the capability to fail over from one controller to another controller, you must install the second controller. Also, you must set the failover mode.

Figure 1–6: Cluster Configuration with HSZ70 Controllers in Transparent Failover Mode



ZK-1589U-AI

In Figure 1–6 the HSZ40, HSZ50, HSZ70, HSZ80, HSG60, or HSG80 has transparent failover mode enabled (`SET FAILOVER COPY = THIS_CONTROLLER`). In transparent failover mode, both controllers are connected to the same shared SCSI bus and device buses. Both controllers service the entire group of storagesets, single-disk units, or other storage

devices. Either controller can continue to service all of the units if the other controller fails.

Note

The assignment of HSZ target IDs can be balanced between the controllers to provide better system performance. See the RAID array controller documentation for information on setting up storagesets.

In the configuration shown in Figure 1–6, there is only one shared SCSI bus. Even by mirroring the clusterwide root and member boot disks, the single shared SCSI bus is a single point of failure.

1.5.5 Creating an NSPOF Cluster

A no-single-point-of-failure (NSPOF) cluster can be achieved by:

- Using two shared SCSI buses and hardware RAID to mirror the cluster file system
- Using multiple shared SCSI buses with storage shelves and mirroring those file systems that can be mirrored with LSM, and by judicious placement of those file systems that cannot be mirrored with LSM.

To create an NSPOF cluster with hardware RAID or LSM and shared SCSI buses with storage shelves, you need to:

- Install a second Memory Channel interface for redundancy.
- Install redundant power supplies.
- Install redundant networks.
- Connect the systems and storage to an uninterruptible power supply (UPS).

Additionally, if you are using hardware RAID, you need to:

- Use hardware RAID to mirror the clusterwide root (`/`), `/usr`, and `/var` file systems, the member boot disks, quorum disk (if present), and data disks.
- Use at least two shared SCSI buses to access dual-redundant RAID array controllers set up for multiple-bus failover mode (HSZ70, HSZ80, HSG60, and HSG80).

Tru64 UNIX support for multipathing provides support for multiple-bus failover.

Notes

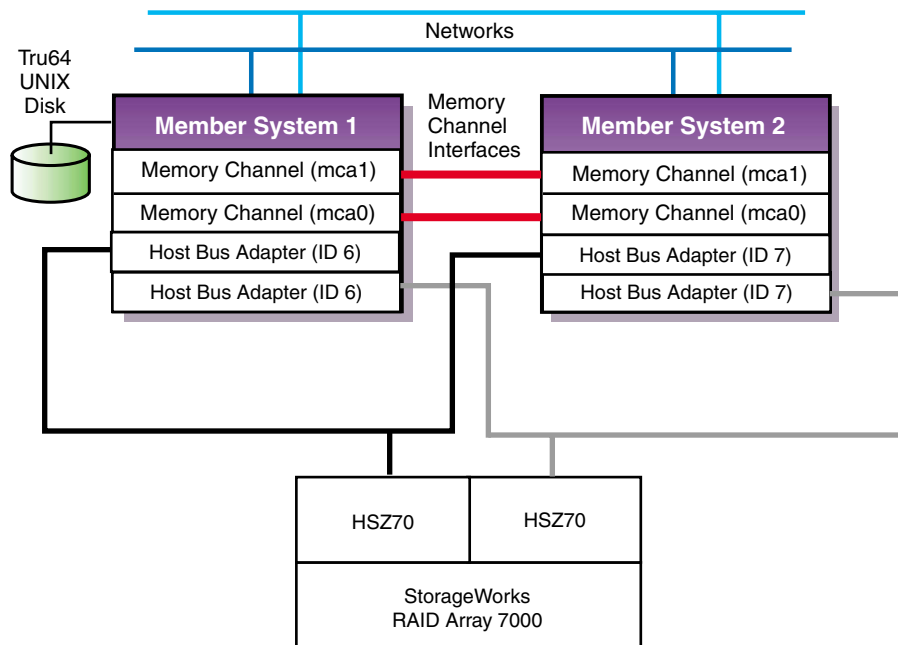
Only the HSZ70, HSZ80, HSG60, and HSG80 are capable of supporting multiple-bus failover (`SET MULTIBUS_FAILOVER COPY = THIS_CONTROLLER`).

Partitioned storagesets and partitioned single-disk units cannot function in multiple-bus failover dual-redundant configurations with the HSZ70 or HSZ80. You must delete any partitions before configuring the controllers for multiple-bus failover.

Partitioned storagesets and partitioned single-disk units are supported with the HSG60 and HSG80 and ACS V8.5 or later.

Figure 1–7 shows a cluster configuration with dual-shared SCSI buses and a storage array with dual-redundant HSZ70s. If there is a failure in one SCSI bus, the member systems can access the disks over the other SCSI bus.

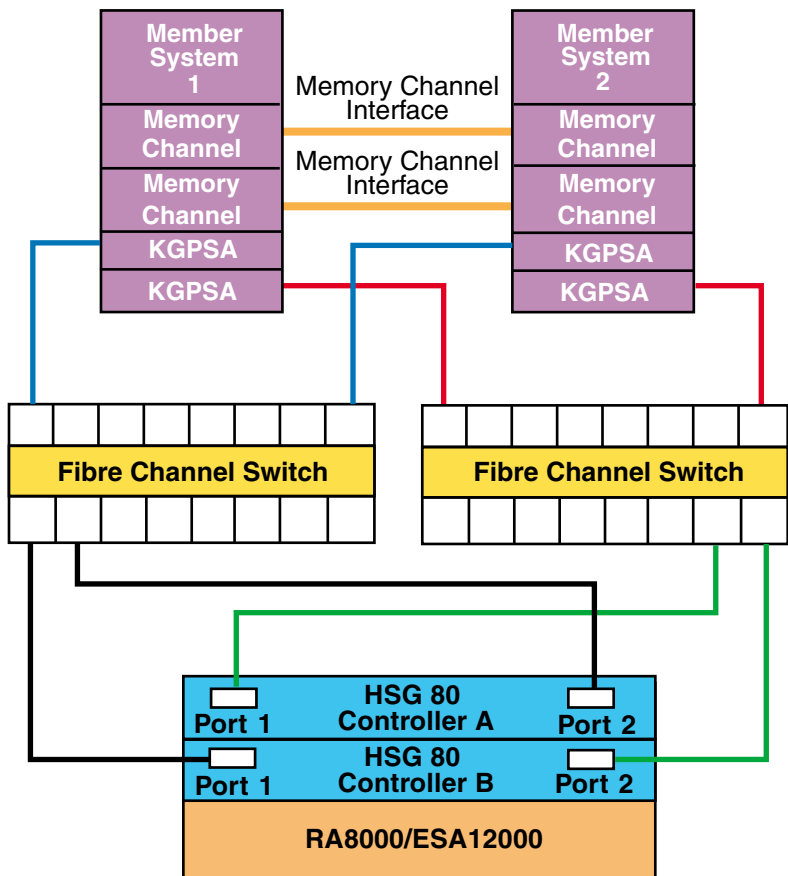
Figure 1–7: NSPOF Cluster Using HSZ70s in Multiple-Bus Failover Mode



ZK-1594U-AI

Figure 1–8 shows a cluster configuration with dual-shared Fibre Channel buses and a storage array with dual-redundant HSG80s configured for multiple-bus failover.

Figure 1–8: NSPOF Fibre Channel Cluster Using HSG80s in Multiple-Bus Failover Mode



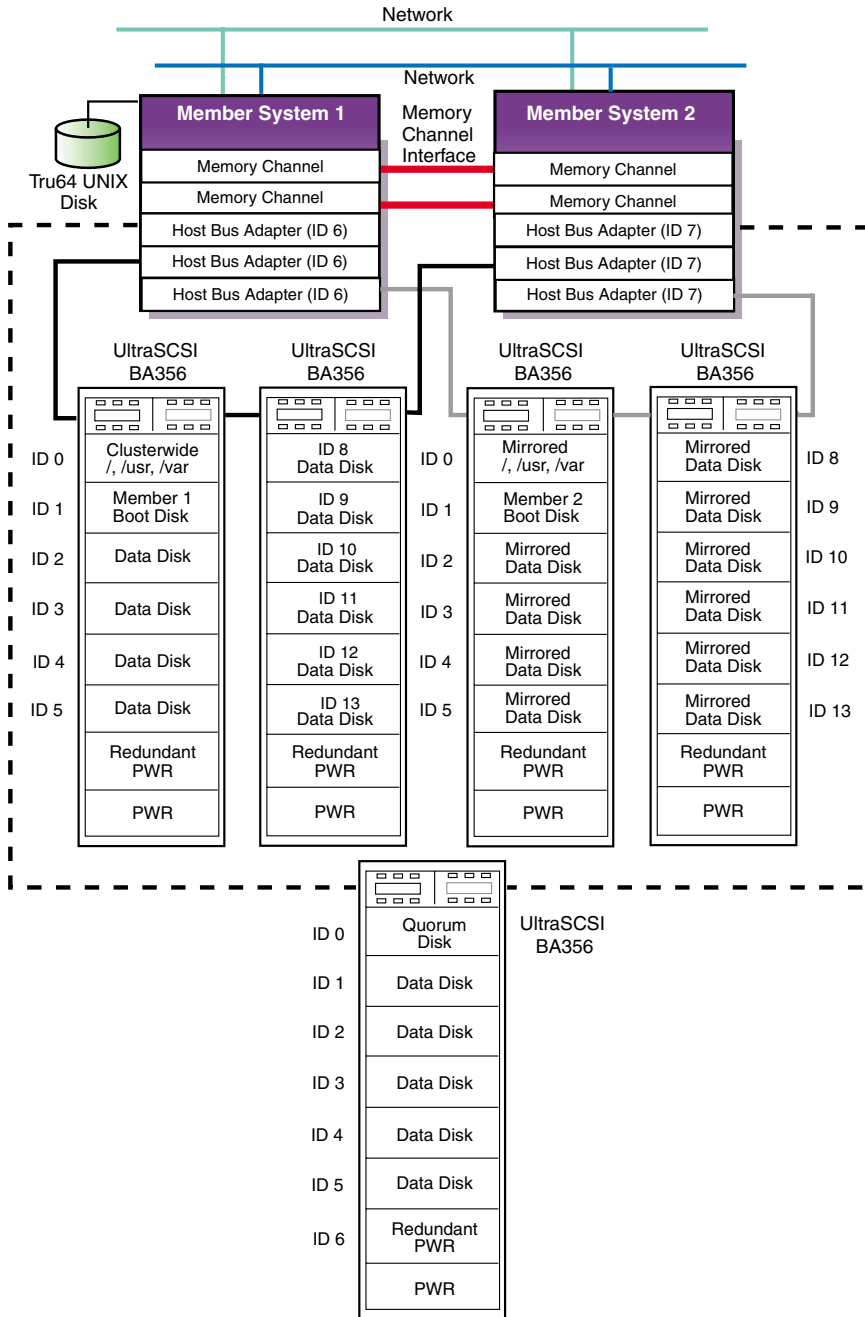
ZK-1765U-AI

If you are using LSM and multiple shared SCSI buses with storage shelves, you need to:

- Mirror the clusterwide root (/), /usr, and /var file systems across two shared SCSI buses.
- Place the boot disk for each member system on a separate shared SCSI bus.
- Provide another shared SCSI bus for the quorum disk.

Figure 1-9 shows a two-member cluster configuration with three shared SCSI buses. The clusterwide root (/), /usr, and /var file systems are mirrored across the first two shared SCSI buses. The boot disk for member system one is on the first shared SCSI bus. The boot disk for member system two is on the second shared SCSI bus. The quorum disk is on the third shared SCSI bus. You can lose one system, or any one shared SCSI bus, and still maintain a cluster.

Figure 1–9: NSPOF Cluster Using LSM and UltraSCSI BA356s



1.6 Eight-Member Clusters

TruCluster Server Version 5.1A supports eight-member cluster configurations as follows:

- **Fibre Channel:** Eight-member systems may be connected to common storage over Fibre Channel in a fabric (switch) configuration.
- **Parallel SCSI:** Only four of the member systems may be connected to any one SCSI bus, but you can have multiple SCSI buses connected to different sets of nodes, and the sets of nodes may overlap. We recommend you use a DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled when connecting four-member systems to a common SCSI bus using RAID array controllers.

An eight-member cluster using Fibre Channel can be extrapolated easily from the discussions in Chapter 6; just connect the systems and storage to your fabric.

An eight-member cluster using shared SCSI storage is more complicated than Fibre Channel, and requires considerable care to configure. One way to configure an eight-member cluster using external termination is discussed in Chapter 11.

1.7 Overview of Setting Up the TruCluster Server Hardware Configuration

To set up a TruCluster Server hardware configuration, follow these steps:

1. Plan your hardware configuration. (See Chapter 3, Chapter 4, Chapter 6, Chapter 9, Chapter 10, and Chapter 11.)
2. Draw a diagram of your configuration.
3. Compare your diagram with the examples in Chapter 3, Chapter 6, Chapter 10, and Chapter 11.
4. Identify all devices, cables, SCSI adapters, and so forth. Use the diagram that you just constructed.
5. Prepare the shared storage by installing disks and configuring any RAID controller subsystems. (See Chapter 3, Chapter 6, and Chapter 10 and the documentation for the StorageWorks enclosure or RAID controller.)
6. Install signal converters in the StorageWorks enclosures, if applicable. (See Chapter 3 and Chapter 10.)
7. Connect storage to the shared SCSI buses. Terminate each bus. Use Y cables or trilink connectors where necessary. (See Chapter 3 and Chapter 10.)

For a Fibre Channel configuration, connect the HSG60 or HSG80 controllers to the switches. You want the HSG60 or HSG80 to recognize the connections to the systems when the systems are powered on.

8. Prepare the member systems by installing:
 - Additional Ethernet or Asynchronous Transfer Mode (ATM) network adapters for client networks.
 - SCSI bus adapters. Ensure that adapter terminators are set correctly. Connect the systems to the shared SCSI bus. (See Chapter 4 or Chapter 9.)
 - The KGPSA host bus adapter for Fibre Channel configurations. Ensure that the KGPSA is operating in the correct mode (FABRIC or LOOP). Connect the KGPSA to the switch. (See Chapter 6.)
 - Memory Channel adapters. Ensure that jumpers are set correctly. (See Chapter 5.)
9. Connect the Memory Channel adapters to each other or to the Memory Channel hub as appropriate. (See Chapter 5.)
10. Turn on the Memory Channel hubs and storage shelves, then turn on the member systems.
11. Install the firmware, set SCSI IDs, and enable fast bus speed as necessary. (See Chapter 4 and Chapter 9.)
12. Display configuration information for each member system, and ensure that all shared disks are seen at the same device number. (See Chapter 4, Chapter 6, or Chapter 9.)

Hardware Requirements and Restrictions

This chapter describes the hardware requirements and restrictions for a TruCluster Server cluster. It includes lists of supported cables, trilink connectors, Y cables, and terminators.

The chapter discusses the following topics:

- Requirements for member systems in a TruCluster Server cluster (Section 2.1)
- Memory Channel requirements (Section 2.2)
- Host bus adapter restrictions (including KGPSA, KZPSA-BB, and KZPBA-CB) (Section 2.3)
- Disk device restrictions (Section 2.4)
- RAID array controller restrictions (Section 2.5)
- SCSI signal converters (Section 2.6)
- Supported DWZZH UltraSCSI hubs (Section 2.7)
- SCSI cables (Section 2.8)
- SCSI terminators and trilink connectors (Section 2.9)

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 TruCluster Server Member System Requirements

The requirements for member systems in a TruCluster Server cluster are as follows:

- Each supported member system requires a minimum firmware revision. See the *Release Notes Overview* supplied with the Alpha™ Systems Firmware Update CD-ROM.

You can also obtain firmware information from the Web at the following URL: <http://www.compaq.com>. Select software & drivers, in the support column, then select AlphaServer, in the servers column. Select the appropriate system.

- Alpha System Reference Manual (SRM) console firmware Version 5.7 or later must be installed on any cluster member that boots from a disk

behind an HSZ80, HSG60, or HSG80 controller. If the cluster member is using earlier firmware, the member may fail to boot, indicating "Reservation Conflict" errors.

- TruCluster Server Version 5.1A supports eight-member cluster configurations as follows:
 - Fibre Channel: Eight-member systems may be connected to common storage over Fibre Channel in a fabric (switch) configuration.
 - Parallel SCSI: Only four of the member systems may be connected to any one SCSI bus, but you can have multiple SCSI buses connected to different sets of nodes, and the sets of nodes may overlap. We recommend you use a DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled when connecting four-member systems to a common SCSI bus using RAID array controllers.

Illustrations of an externally terminated eight-node cluster are shown in Chapter 11. The cluster shown is more appropriate for high performance technical computing (HPTC) customers who are looking for performance instead of availability.

- The following items pertain to the AlphaServer GS80/160/320 systems:
 - High power peripheral component interconnect (PCI) modules (approximately 25 watts or greater) must be placed in PCI slots with a 1-inch module pitch; any slot except 0-5, 0-6, 1-5, and 1-6.


A primary or expansion PCI drawer contains two 3-slot PCI buses and two 4-slot PCI buses (see Figure 2–1):

- PCI0 for I/O riser 0: Slots 0-0/1, 0-2, and 0-3
- PCI1 for I/O riser 0: Slots 0-4, 0-5, 0-6, and 0-7
- PCI0 for I/O riser 1: Slots 1-1, 1-2, and 1-3
- PCI1 for I/O riser 1: Slots 1-4, 1-5, 1-6, and 1-7

Note

Slot 0-0/1 in a primary PCI drawer contains the standard I/O module.

Figure 2–1: PCI Backplane Slot Layout



- TruCluster Server does not support the XMI CIXCD on an AlphaServer 8x00, GS60, GS60E, or GS140 system.

2.2 Memory Channel Restrictions

The Memory Channel interconnect is one method used for cluster communications between the member systems.

There are currently three versions of the Memory Channel product: Memory Channel 1, Memory Channel 1.5, and Memory Channel 2. The Memory Channel 1 and Memory Channel 1.5 products are very similar (the PCI adapter for both versions is the CCMAA module) and are generally referred to as MC1 throughout this manual. The Memory Channel 2 product (CCMAB module) is referred to as MC2.

Ensure that you abide by the following Memory Channel restrictions:

- The DS10, DS20, DS20E, ES40, GS80, GS160, and GS320 systems only support MC2 hardware.
- If you configure a cluster with a single rail Memory Channel in standard hub mode and the hub fails, every cluster member panics. They panic because no member can see any of the other cluster members over the Memory Channel interface. A quorum disk does not help in this case, because no system is given the opportunity to obtain ownership of the quorum disk and survive.

To prevent this situation in standard hub mode, install a second Memory Channel rail. A hub failure on one rail will cause failover to the other rail.

When the Memory Channel is set up in standard hub mode, the Memory Channel hub must be visible to each member's Memory Channel adapter. If the hub is powered off, no system is able to boot.

A two-node cluster configured in virtual hub mode does not have these problems. In virtual hub mode, each system is always connected to the virtual hub. A loss of communication over the Memory Channel causes both members (if both members are still up) to attempt to obtain ownership of the quorum disk. The member that succeeds continues as a single-member cluster. The other member panics.

A single system of a two-node cluster that is configured in virtual hub mode will boot because a virtual hub is always present.

- If a TruCluster Server cluster configuration utilizes multiple Memory Channel adapters in standard hub mode, the Memory Channel adapters must be connected to separate Memory Channel hubs. The first Memory Channel adapter (mca0) in each system must be connected to one Memory Channel hub. The second Memory Channel adapter (mcb0) in each system must be connected to a second Memory Channel hub. Also, each Memory Channel adapter on one system must be connected to the same linecard in each Memory Channel hub.
- If redundant Memory Channel adapters are used with a DS10, they must be jumpered for 128 MB and not the default of 512 MB.
- If you have redundant MC2 modules on a GS80, GS160, or GS320 system jumpered for 512 MB, you cannot have any other modules except the CCMFB fiber-optic module on that PCI bus.
- Redundant Memory Channels are supported within a mixed Memory Channel configuration, as long as MC1 adapters are connected to other MC1 adapters and MC2 adapters are connected to MC2 adapters.

In a cluster with mixed revision Memory Channel rails, the MC2 adapter modules must be jumpered for 128 MB.

- A Memory Channel interconnect can use either virtual hub mode (two member systems connected without a Memory Channel hub) or standard hub mode (two or more systems connected to a hub). A TruCluster Server cluster with three or more member systems must be jumpered for standard hub mode and requires a Memory Channel hub.
- If Memory Channel modules are jumpered for virtual hub mode, all Memory Channel modules on a system must be jumpered in the same manner, either virtual hub 0 (VH0) or virtual hub 1 (VH1). You cannot have one Memory Channel module jumpered for VH0 and another jumpered for VH1 on the same system.
- The maximum length of an MC1 BC12N link cable is 3 meters (9.8 feet).

- The maximum length of an MC2 BN39B link cable is 10 meters (32.8 feet).
- In an MC2 configuration, you can use a CCMFB optical converter in conjunction with the MC2 CCMAB host bus adapter or a CCMLB hub line card to increase the distance between systems.
 - The BN34R fiber-optic cable, which is used to connect two CCMFB optical converters, is available in 10-meter (32.8-foot) (BN34R-10) and 31-meter (101.7-foot) (BN34R-31) lengths. Customers may provide their own fiber-optic cables to achieve greater separation of systems.
 - The Memory Channel fiber-optic connection may be up to 2 kilometers (1.24 miles) between two CCMFB optical converters connected to CCMAB host bus adapters in virtual hub mode.
 - The Memory Channel fiber-optic connection may be up to 3 kilometers (1.86 miles) between a CCMFB optical converter connected to a CCMAB host bus adapter and a CCMFB optical converter connected to a CCMLB hub line card in standard hub mode (providing a maximum separation of 6 kilometers (3.73 miles) between systems).
- Always examine a Memory Channel link cable for bent or broken pins. Be sure that you do not bend or break any pins when you connect or disconnect a cable.
- For AlphaServer 8200, 8400, GS60, GS60E, or GS140 systems, the Memory Channel adapter must be installed in slots 0-7 of a DWLPA PCIA option; there are no restrictions for a DWLPB.
- For AlphaServer 1000A systems, the Memory Channel adapter must be installed on the primary PCI (in front of the PCI-to-PCI bridge chip) in PCI slots 11, 12, or 13 (the top three slots).
- For AlphaServer 2000 systems, the B2111-AA module must be at Revision H or higher.

For AlphaServer 2100 systems, the B2110-AA module must be at Revision L or higher.

Use the `examine` console command to determine if these modules are at a supported revision as follows:

```
P00>>> examine -b econfig:20008
econfig: 20008 04
P00>>>
```

If a hexadecimal value of 04 or greater is returned, the I/O module supports Memory Channel.

If a hexadecimal value of less than 04 is returned, the I/O module is not supported for Memory Channel usage.

Order an H3095-AA module to upgrade an AlphaServer 2000 or an H3096-AA module to upgrade an AlphaServer 2100 to support Memory Channel.

- For AlphaServer 2100A systems, the Memory Channel adapter must be installed in PCI 4 through PCI 7 (slots 6, 7, 8, and 9), which are the bottom four PCI slots.

2.3 Host Bus Adapter Restrictions

To connect a member system to a shared SCSI bus, you must install a host bus adapter in an I/O bus slot.

The Tru64 UNIX operating system supports a maximum of 64 I/O buses. TruCluster Server supports a total of 32 shared I/O buses using KZPSA-BB host bus adapters, KZPBA-CB UltraSCSI host bus adapters, or KGPSA Fibre Channel host bus adapters.

The following sections describe the host bus adapter restrictions in more detail.

2.3.1 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 that are supported on each system at the time the TruCluster Server Version 5.1A product was shipped. 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>

Table 2–1: AlphaServer Systems Supported for Fibre Channel

AlphaServer System	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	—

Table 2–1: AlphaServer Systems Supported for Fibre Channel (cont.)

AlphaServer System	Number of Adapters Supported in Fabric Topology	Number of Adapters Supported in Loop Topology
AlphaServer GS60, GS60E, and GS140 ^b	63 ^c , 32 ^d	—
AlphaServer GS80, GS160, and GS320 ^e	62	—

^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 5.1A:

- The HSG60 and HSG80 require Array Control Software (ACS) Version 8.5 or later.
- 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.
- TheStorageWorks Modular Array 6000 (MA6000) supports dual-redundant HSG60 controllers and 1-inch universal drives.
- The StorageWorks Modular Array 8000 (MA8000) and Enterprise Modular Array 12000 (EMA12000) support dual redundant HSG80 controllers and 1-inch universal drives.
- 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 or 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 (1640.4 feet). It also supports 62.5-micron multimode fiber cables in lengths up to 200 meters (656.2 feet). The GBIC-LW supports 9-micron, single-mode fiber cables with the SC connector in lengths up to 10 kilometers (6.2 miles).

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 5.1A allows up to 255 Fibre Channel targets. An active host port or host bus adapter constitutes a target.
- Tru64 UNIX Version 5.1A allows up to 255 logical unit numbers (LUNs) per target.
- The HSG60 and HSG80 supports transparent and multiple-bus failover mode when used in a TruCluster Server Version 5.1A 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.

2.3.2 KZPSA-BB SCSI Adapter Restrictions

KZPSA-BB SCSI adapters have the following restrictions:

- The KZPSA-BB requires A12 firmware.
- If you have a KZPSA-BB adapter installed in an AlphaServer that supports the `bus_probe_algorithm` console variable (for example, the AlphaServer 800, 1000, 1000A, 2000, 2100, or 2100A systems), you must set the `bus_probe_algorithm` console variable to `new` by entering the following command:

```
>>> set bus_probe_algorithm new
```

Use the `show bus_probe_algorithm` console command to determine if your system supports the variable. If the response is null or an error, there is no support for the variable. If the response is anything other than `new`, you must set it to `new`.

- On AlphaServer 1000A and 2100A systems, updating the firmware on the KZPSA-BB SCSI adapter is not supported when the adapter is behind the PCI-to-PCI bridge.

2.3.3 KZPBA-CB SCSI Bus Adapter Restrictions

KZPBA-CB UltraSCSI adapters have the following restrictions:

- A maximum of four HSZ50, HSZ70, or HSZ80 RAID array controllers can be placed on a single KZPBA-CB UltraSCSI bus. Only two redundant pairs of array controllers are allowed on one SCSI bus.
- The KZPBA-CB requires ISP 1020/1040 firmware Version 5.57 or higher, which is available with the system SRM console firmware on the Alpha Systems Firmware 5.3 Update CD-ROM (or later).
- The maximum length of any differential SCSI bus segment is 25 meters (82 feet), including the length of the SCSI bus cables and SCSI bus internal to the SCSI adapter, hub, or storage device. A SCSI bus may have more than one SCSI bus segment (see Section 3.1).
- See the KZPBA-CB UltraSCSI Storage Adapter Module *Release Notes* for more information.

2.4 Disk Device Restrictions

The restrictions for disk devices are as follows:

- Disks on shared SCSI buses must be installed in external storage shelves or behind a RAID array controller.
- TruCluster Server does not support Prestoserve on any shared disk.

2.5 RAID Array Controller Restrictions

RAID array controllers provide high performance, high availability, and high connectivity access to SCSI devices through a shared SCSI bus.

RAID array controllers require the minimum Array Controller Software (ACS) listed in Table 2-2.

Table 2–2: RAID Controller Minimum Required Array Controller Software

RAID Controller	Minimum Required Array Controller Software
HSZ20	3.4
HSZ22 (RAID Array 3000)	D11x
HSZ40	3.7
HSZ50	5.7
HSZ70	7.7
HSZ80	8.3-1
HSG60	8.5
HSG80	8.5

RAID controllers can be configured with the number of SCSI IDs as listed in Table 2–3.

Table 2–3: RAID Controller SCSI IDs

RAID Controller	Number of SCSI IDs Supported
HSZ20	4
HSZ22 (RAID Array 3000)	2
HSZ40	4
HSZ50	4
HSZ70	8
HSZ80	15
HSG60	N/A
HSG80	N/A

The following restrictions are imposed for support of the StorageWorks RAID Array 3000 (RA3000) subsystem:

- The RAID Array 3000 (RA3000) with HSZ22 controller does not support multi-bus access or multiple-bus failover. You cannot achieve a no-single-point-of-failure (NSPOF) cluster using an RA3000.
- The KZPBA-CB UltraSCSI host adapter is the only SCSI bus host adapter supported with the RA3000 in a TruCluster Server cluster. The KZPBA-CB requires ISP 1020/1040 firmware Version 5.57 (or higher), which is available with the system SRM console firmware on the Alpha Systems Firmware 5.4 or later Update CD.

- Only RA3000 storage units visible to the host as LUN0 (storage units with a zero (0) as the last digit of the unit number such as D0, D100, D200, and so forth) can be used as a boot device.
- StorageWorks Command Console (SWCC) V2.2 is the only configuration utility that will work with the RA3000. SWCC V2.2 runs only on a Microsoft Windows NT or Windows 2000 PC.
- The controller will not operate without at least one 16-MB SIMM installed in its cache.
- The device expansion shelf (DS-SWXRA-GN) for the rackmount version must be at revision level B01 or higher.
- The single-ended personality module used in the DS-SWXRA-GN UltraSCSI storage expansion shelves must be at revision H01 or higher.
- The RA3000 order includes an uninterruptible power supply (UPS), which must be connected to the RA3000.

2.6 SCSI Signal Converters

If you are using a standalone storage shelf with a single-ended SCSI interface in your cluster configuration, you must connect it to a SCSI signal converter. SCSI signal converters convert wide, differential SCSI to narrow or wide, single-ended SCSI and vice versa. Some signal converters are standalone desktop units and some are StorageWorks building blocks (SBBs) that you install in storage shelves disk slots.

Note

UltraSCSI hubs logically belong in this section because they contain a DOC (DWZZA on a chip) chip, but they are discussed separately in Section 2.7.

The restrictions for SCSI signal converters are as follows:

- If you remove the cover from a standalone unit, be sure to replace the star washers on all four screws that hold the cover in place when you reattach the cover. If the washers are not replaced, the SCSI signal converter may not function correctly because of noise.
- If you want to disconnect a SCSI signal converter from a shared SCSI bus, you must turn off the signal converter before disconnecting the cables. To reconnect the signal converter to the shared bus, connect the cables before turning on the signal converter. Use the power switch to turn off a standalone SCSI signal converter. To turn off an SBB SCSI signal converter, pull it from its disk slot.

- If you observe any “bus hung” messages, your DWZZA signal converters may have the incorrect hardware. In addition, some DWZZA signal converters that appear to have the correct hardware revision may cause problems if they also have serial numbers in the range from CX444xxxxx through CX449xxxxx.

To upgrade a DWZZA-AA or DWZZA-VA signal converter to the correct revision, use the appropriate field change order (FCO), as follows:

- DWZZA-AA-F002
- DWZZA-VA-F001

2.7 DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI Hubs

The DS-DWZZH-03 and DS-DWZZH-05 series UltraSCSI hubs are the only hubs that are supported in a TruCluster Server configuration. They are SCSI-2- and draft SCSI-3-compliant SCSI 16-bit signal converters capable of data transfer rates of up to 40 MB/sec.

These hubs can be listed with the other SCSI bus signal converters, but because they are used differently in cluster configurations, they are discussed differently in this manual.

A DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub can be installed in:

- A StorageWorks UltraSCSI BA356 shelf (which has the required 180-watt power supply).
- The lower righthand device slot of the BA370 shelf within the RA7000 or ESA 10000 RAID array subsystems. This position minimizes cable lengths and interference with disks.
- A wide BA356 that has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.

A DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub:

- Improves the reliability of the detection of cable faults.
- Provides for bus isolation of cluster systems while allowing the remaining connections to continue to operate.
- Allows for more separation of systems and storage in a cluster configuration, because each SCSI bus segment can be up to 25 meters (82 feet) in length. This allows a total separation of nearly 50 meters (164 feet) between a system and the storage.

Note

The DS-DWZZH-03/05 UltraSCSI hubs cannot be connected to a StorageWorks BA35X storage shelf because the storage shelf does not provide termination power to the hub.

2.8 SCSI Cables

If you are using shared SCSI buses, you must determine if you need cables with connectors that are low-density 50-pins, high-density 50-pins, high-density 68-pins (HD68), or VHDCI (UltraSCSI). If you are using an UltraSCSI hub, you will need HD68-to-VHDCI and VHDCI-to-VHDCI cables. In some cases, you also have the choice of straight or right-angle connectors. In addition, each supported cable comes in various lengths. Use the shortest possible cables to adhere to the limits on SCSI bus length.

Table 2–4 describes each supported cable and the context in which you would use the cable. Some Compaq equivalent 6-3 part numbers are not provided.

Table 2–4: Supported SCSI Cables

Cable	Connector Density	Pins	Configuration Use
BN21W-0B	Three high	68-pin	A Y cable that can be attached to a KZPSA-BB or KZPBA-CB if there is no room for a trilink connector. It can be used with a terminator to provide external termination.
BN21M	One low, one high	50-pin LD to 68-pin HD	Connects the single-ended end of a DWZZA-AA or DWZZB-AA to a TZ885 or TZ887. ^a
BN21K, BN21L, or 328215-00X	Two HD68	68-pin	Connects BN21W Y cables or wide devices. For example, connects KZPBA-CBs, KZPSA-BBs, HSZ40s, HSZ50s, the differential sides of two SCSI signal converters, or a DWZZB-AA to a BA356.
BN37A	Two VHDCI	VHDCI to VHDCI	Connects two VHDCI trilinks to each other or an UltraSCSI hub to a trilink on an HSZ70, HSZ80, or an UltraSCSI hub to a RAID Array 3000.
BN38C or BN38D	One HD68, one VHDCI	HD68 to VHDCI	Connects a KZPBA-CB or KZPSA-BB to a port on an UltraSCSI hub.

Table 2–4: Supported SCSI Cables (cont.)

Cable	Connector Density	Pins	Configuration Use
BN38E-0B	Technology adapter cable	HD68 male to VHDCI female	May be connected to a BN37A cable and the combination used in place of a BN38C or BN38D cable
199629-002 or 189636-002	Two high	50-pin HD to 68-pin HD	Connect a Compaq 20/40 GB DLT Tape Drive to a DWZZB-AA
146745-003 or 146776-003	Two high	50-pin HD to 50-pin HD	Daisy-chain two Compaq 20/40 GB DLT Tape Drives
189646-001 or 189646-002	Two high	68-pin HD	Connect a Compaq 40/80 DLT Tape Drive to a DWZZB-AA or daisy-chain two Compaq 40/80 DLT Tape Drives

^a Do not use a KZPBA-CB with a DWZZA-AA or DWZZB-AA and a TZ885 or TZ887. The DWZZAs and DWZZBs cannot operate at UltraSCSI speed.

Always examine a SCSI cable for bent or broken pins. Be sure that you do not bend or break any pins when you connect or disconnect a cable.

2.9 SCSI Terminators and Trilink Connectors

Table 2–5 describes the supported trilink connectors and SCSI terminators and the context in which you use them.

Table 2–5: Supported SCSI Terminators and Trilink Connectors

Trilink Connector or Terminator	Density	Pins	Configuration Use
H885-AA	Three	68-pin	Trilink connector that attaches to high-density, 68-pin cables or devices, such as a KZPSA-BB, KZPBA-CB, HSZ40, HSZ50, or the differential side of a SCSI signal converter. Can be terminated with an H879-AA terminator to provide external termination.
H8574-A or H8860-AA	Low	50-pin	Terminates a TZ885 or TZ887 tape drive.
341102-001	High	50-pin	Terminates a Compaq 20/40 GB DLT Tape Drive
H879-AA or 330563-001	High	68-pin	Terminates an H885-AA trilink connector, BN21W-0B Y cable, or an ESL9326D Enterprise Library tape drive.
H8861-AA	VHDCI	68-pin	VHDCI trilink connector that attaches to VHDCI 68-pin cables, UltraSCSI BA356 JA1, and HSZ70 or HSZ80 RAID controllers. Can be terminated with an H8863-AA terminator if necessary.
H8863-AA	VHDCI	68-pin	Terminate a VHDCI trilink connector.
152732-001	VHDCI	68-pin	Low Voltage Differential terminator

The requirements for trilink connectors are as follows:

- If you connect a SCSI cable to a trilink connector, do not block access to the screws that mount the trilink, or you will be unable to disconnect the trilink from the device without disconnecting the cable.
- Do not install an H885-AA trilink if installing it will block an adjacent peripheral component interconnect (PCI) port. Use a BN21W-0B Y cable instead.

3

Shared SCSI Bus Requirements and Configurations Using UltraSCSI Hardware

A TruCluster Server cluster uses shared SCSI buses, external storage shelves or redundant array of independent disks (RAID) controllers, and supports disk mirroring and fast file system recovery to provide high data availability and reliability.

This chapter discusses the following topics:

- Shared SCSI bus configuration requirements (Section 3.1)
- SCSI bus performance (Section 3.2)
- SCSI bus device identification numbers (Section 3.3)
- SCSI bus length (Section 3.4)
- SCSI bus termination (Section 3.5)
- UltraSCSI hubs (Section 3.6)
- How to configure UltraSCSI hubs with RAID array controllers (Section 3.7)

This chapter:

- Introduces SCSI bus configuration concepts
- Describes requirements for the shared SCSI bus
- Provides procedures for cabling TruCluster Server radial configurations using UltraSCSI hubs and:
 - Dual-redundant HSZ70 or HSZ80 RAID array controllers that are enabled for simultaneous failover
 - Dual-redundant HSZ70 or HSZ80 RAID array controllers that are enabled for multiple-bus failover
 - RAID Array 3000 (RA3000) with HSZ22 controller configured for active/active or active/passive mode.
- Provides diagrams of TruCluster Server storage configurations using UltraSCSI hardware that is configured for radial connections

Note

Although the UltraSCSI BA356 might have been included in this chapter with the other UltraSCSI devices, it is not. The UltraSCSI BA356 is discussed in Chapter 10 with the configurations using external termination. It cannot be cabled directly to an UltraSCSI hub because it does not provide SCSI bus termination power (`termpwr`).

In addition to using only supported hardware, adhering to the requirements described in this chapter will ensure that your cluster operates correctly.

Chapter 10 contains additional information about using SCSI bus signal converters, and also contains diagrams of TruCluster Server configurations using UltraSCSI and non-UltraSCSI storage shelves and RAID array controllers. The chapter also covers the older method of using external termination and covers radial configurations with the DWZZH UltraSCSI hubs and non-UltraSCSI RAID array controllers.

3.1 Shared SCSI Bus Configuration Requirements

A shared SCSI bus must adhere to the following requirements:

- Only an external bus can be used for a shared SCSI bus.
- SCSI bus specifications set a limit of 8 devices on an 8-bit (narrow) SCSI bus. The limit is 16 devices on a 16-bit SCSI bus (wide). See Section 3.3 for more information.
- The length of each physical bus is strictly limited. See Section 3.4 for more information.
- You can directly connect devices only if they have the same transmission mode (differential or single-ended) and data path (narrow or wide). Use a SCSI signal converter to connect devices with different transmission modes. See Section 10.1 for information about the DWZZA (BA350) or DWZZB (BA356) signal converters or the DS-BA35X-DA personality module (which acts as a differential to single-ended signal converter for the UltraSCSI BA356).
- For each SCSI bus segment, you can have only two terminators, one at each end. A physical SCSI bus may be composed of multiple SCSI bus segments.
- If you do not use an UltraSCSI hub, you must use trilink connectors and Y cables to connect devices to a shared bus, so you can disconnect the devices without affecting bus termination. See Section 10.2 for more information.

- Be careful when performing maintenance on any device that is on a shared bus because of the constant activity on the bus. Usually, to perform maintenance on a device without shutting down the cluster, you must be able to isolate the device from the shared bus without affecting bus termination.
- All supported UltraSCSI host adapters support UltraSCSI disks at UltraSCSI speeds in UltraSCSI BA356 shelves, RA7000 or ESA10000 storage arrays (HSZ70 and HSZ80), RA8000 or ESA12000 storage arrays (HSZ80 and HSG80), or the RAID Array 3000 (RA3000, with HSZ22 controller). Older, non-UltraSCSI BA356 shelves are supported with UltraSCSI host adapters and host RAID controllers as long as they contain no UltraSCSI disks.

Note

The RA3000 is supported only with the KZPBA-CB UltraSCSI host bus adapter.

- UltraSCSI drives and fast wide drives can be mixed together in an UltraSCSI BA356 shelf. (See Chapter 10.)
- Differential UltraSCSI adapters may be connected to either (or both) a non-UltraSCSI BA356 shelf (via a DWZZB-VW) and the UltraSCSI BA356 shelf (via the DS-BA35X-DA personality module) on the same shared SCSI bus. The UltraSCSI adapter negotiates maximum transfer speeds with each SCSI device. (See Chapter 10.)
- The HSZ70 and HSZ80 UltraSCSI RAID controllers have a wide differential UltraSCSI host bus with a Very High Density Cable Interconnect (VHDCI) connector. HSZ70 and HSZ80 controllers will work with fast and wide differential SCSI adapters (for example, KZPSA-BB) at fast SCSI speeds.
- Fast, wide SCSI drives (green StorageWorks building blocks (SBBs) with part numbers ending in -VW) may be used in an UltraSCSI BA356 shelf.
- Do not use fast, narrow SCSI drives (green SBBs with part numbers ending in -VA) in any shelf that can assign the drive a SCSI ID greater than 7. They will not work.
- The UltraSCSI BA356 requires a 180-watt power supply (BA35X-HH). It will not function properly with the older, lower-wattage BA35X-HF universal 150-watt power supply. (See Chapter 10.)
- An older BA356 that has been retrofitted with a BA35X-HH 180-watt power supply and DS-BA35X-DA personality module is still only FCC certified for Fast 10 configurations. (See Chapter 10.)

3.2 SCSI Bus Performance

Before you set up a SCSI bus, it is important that you understand a number of issues that affect the viability of a bus and how the devices that are connected to it operate. Specifically, bus performance is influenced by the following factors:

- Transmission method (Section 3.2.2)
- Data path (Section 3.2.3)
- Bus speed (Section 3.2.4)

3.2.1 SCSI Bus Versus SCSI Bus Segments

An UltraSCSI bus may comprise multiple UltraSCSI bus segments. Each UltraSCSI bus segment comprises electrical conductors that may be in a cable or a backplane, and cable or backplane connectors. Each UltraSCSI bus segment must have a terminator at each end.

Up to two UltraSCSI bus segments may be coupled together with UltraSCSI hubs or signal converters, increasing the total length of the UltraSCSI bus.

3.2.2 Transmission Methods

Two transmission methods can be used in a SCSI bus:

- Single-ended — In a single-ended SCSI bus, one data lead and one ground lead are utilized for the data transmission. A single-ended receiver looks only at the signal wire as the input. The transmitted signal arrives at the receiving end of the bus on the signal wire somewhat distorted by signal reflections. The length and loading of the bus determine the magnitude of this distortion. This transmission method is economical, but is more susceptible to noise than the differential transmission method, and requires short cables. Devices with single-ended SCSI devices include the following:
 - BA350, BA356, and UltraSCSI BA356 storage shelves
 - Single-ended side of a SCSI signal converter or personality module
- Differential — Differential signal transmission uses two wires to transmit a signal. The two wires are driven by a differential driver that places a signal on one wire (+SIGNAL) and another signal that is 180 degrees out of phase (-SIGNAL) on the other wire. The differential receiver generates a signal output only when the two inputs are different. As signal reflections occur virtually the same on both wires, they are not seen by the receiver, because it only sees differences on the two wires.

This transmission method is less susceptible to noise than single-ended SCSI and enables you to use longer cables. Devices with differential SCSI interfaces include the following:

- KZPBA-CB
- KZPSA-BB
- HSZ40, HSZ50, HSZ70, and HSZ80 controllers
- Differential side of a SCSI signal converter or personality module

You cannot use the two transmission methods in the same SCSI bus segment. For example, a device with a differential SCSI interface must be connected to another device with a differential SCSI interface. If you want to connect devices that use different transmission methods, use a SCSI signal converter between the devices. The DS-BA35X-DA personality module is discussed in Section 10.1.2.2. See Section 10.1 for information about using the DWZZ* series of SCSI signal converters.

You cannot use a DWZZA or DWZZB signal converter at UltraSCSI speeds for TruCluster Server if there are any UltraSCSI disks on the bus, because the DWZZA or DWZZB will not operate correctly at UltraSCSI speed. The DS-BA35X-DA personality module contains a signal converter for the UltraSCSI BA356. It is the interface between the shared differential UltraSCSI bus and the UltraSCSI BA356 internal single-ended SCSI bus.

RAID array controller subsystems provide the function of a signal converter, accepting the differential input and driving the single-ended device buses.

3.2.3 Data Path

There are two data paths for SCSI devices:

- Narrow — Implies an 8-bit data path for SCSI-2. The performance of this mode is limited.
- Wide — Implies a 16-bit data path for SCSI-2 or UltraSCSI. This mode increases the amount of data that is transferred in parallel on the bus.

3.2.4 Bus Speed

Bus speeds vary depending upon the bus clocking rate and bus width, as listed in Table 3-1.

Table 3–1: SCSI Bus Speeds

SCSI Bus	Transfer Rate (MHz)	Bus Width (Bytes)	Bus Bandwidth (Speed) MB/sec
SCSI	5	1	5
Fast SCSI	10	1	10
Fast-Wide	10	2	20
UltraSCSI	20	2	40
UltraSCSI-II	40	2	80

3.3 SCSI Bus Device Identification Numbers

On a shared SCSI bus, each SCSI device uses a device address and must have a unique SCSI ID (from 0 through 15). For example, each SCSI bus adapter and each disk in a single-ended storage shelf uses a device address.

SCSI bus adapters have a default SCSI ID that you can change by using console commands or utilities. For example, a KZPSA adapter has an initial SCSI ID of 7.

Note

If you are using a DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled, SCSI ID numbering will change. (See Section 3.6.1.2.)

Use the following priority order to assign SCSI IDs to the SCSI bus adapters connected to a shared SCSI bus:

7-6-5-4-3-2-1-0-15-14-13-12-11-10-9-8

This order specifies that 7 is the highest priority, and 8 is the lowest priority. When assigning SCSI IDs, use the highest priority ID for member systems (starting at 7). Use lower priority IDs for disks.

Normal SCSI priority is not followed when using the DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled because the DS-DWZZH-05 determines the SCSI ID of the next device to use the SCSI bus.

The SCSI ID for a disk in a BA350 storage shelf corresponds to its slot location. The SCSI ID for a disk in a BA356 or UltraSCSI BA356 depends upon its slot location and the personality module SCSI bus address switch settings.

3.4 SCSI Bus Length

There is a limit to the length of the cables in a shared SCSI bus. The total cable length for a SCSI bus segment is calculated from one terminated end to the other.

If you are using devices that have the same transmission method and data path (for example, wide differential), a shared bus will consist of only one bus segment. If you have devices with different transmission methods, you will have both single-ended and differential bus segments, each of which must be terminated only at both ends and must adhere to the rules on bus length.

Table 3–2 describes the maximum cable length for a physical SCSI bus segment.

Table 3–2: SCSI Bus Segment Length

SCSI Bus	Bus Speed (MB/sec)	Maximum Cable Length
Narrow, single-ended	5	6 meters (19.7 feet)
Narrow, single-ended fast	10	3 meters (9.8 feet)
Wide differential, fast	20	25 meters (82 feet)
Differential UltraSCSI	40	25 meters ^a

^a The maximum separation between a host and storage in a TruCluster Server configuration is 50 meters (164 feet): 25 meters (82 feet) between any host and the UltraSCSI hub and 25 meters (82 feet) between the UltraSCSI hub and the RAID array controller.

Because of the cable length limit, you must plan your hardware configuration carefully, and ensure that each SCSI bus meets the cable limit guidelines. In general, you must place systems and storage shelves as close together as possible and choose the shortest possible cables for the shared bus.

3.5 Terminating the Shared SCSI Bus When Using UltraSCSI Hubs

You must properly connect devices to a shared SCSI bus. In addition, you can terminate only the beginning and end of each bus segment (either single-ended or differential).

There are two rules for SCSI bus termination:

- There are only two terminators for each SCSI bus segment. If you use an UltraSCSI hub, you only have to install one terminator.
- If you do not use an UltraSCSI hub, bus termination must be external. External termination is covered in Section 10.2.

Notes

With the exception of the TZ885, TZ887, TL890, TL891, and TL892, tape devices can only be installed at the end of a shared SCSI bus. These tape devices are the only supported tape devices that can be terminated externally.

We recommend that tape loaders be on a separate, shared SCSI bus to allow normal shared SCSI bus termination for those shared SCSI buses without tape loaders.


Whenever possible, connect devices to a shared bus so that they can be isolated from the bus. This allows you to disconnect devices from the bus for maintenance purposes, without affecting bus termination and cluster operation. You also can set up a shared SCSI bus so that you can connect additional devices at a later time without affecting bus termination.

Most devices have internal termination. For example, the UltraSCSI KZPBA-CB and the fast and wide KZPSA-BB host bus adapters have internal termination. When using a KZPBA-CB or KZPSA-BB with an UltraSCSI hub, ensure that the onboard termination resistor SIPs have not been removed.

You will need to provide termination at the storage end of one SCSI bus segment. You will install an H8861-AA trilinear connector on the HSZ70 or HSZ80 at the bus end. Connect an H8863-AA terminator to the trilinear connector to terminate the bus.

Figure 3-1 shows a VHDCI trilinear connector (UltraSCSI), which you may attach to an HSZ70 or HSZ80.

Figure 3-1: VHDCI Trilinear Connector (H8861-AA)



CX05744A

3.6 UltraSCSI Hubs

The DS-DWZZH series UltraSCSI hubs are UltraSCSI signal converters that provide radial connections of differential SCSI bus adapters and RAID array controllers. Each connection forms a SCSI bus segment with SCSI bus adapters or the storage unit. The hub provides termination for one end

of the bus segment. Termination for the other end of the bus segment is provided by the following components:

- Installed KZPBA-CB (or KZPSA-BB) termination resistor SIPs
- External termination on a tralink connector that is attached to an HSZ40, HSZ50, HSZ70, or HSZ80

Note

The DS-DWZZH-03/05 UltraSCSI hubs cannot be connected to a StorageWorks BA35X storage shelf because the storage shelf does not provide termination power to the hub.

3.6.1 Using a DWZZH UltraSCSI Hub in a Cluster Configuration

The DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI hubs are supported in a TruCluster Server cluster. They both provide radial connection of cluster member systems and storage, and are similar in the following ways:

- Contain internal termination for each port; therefore, the hub end of each SCSI bus segment is terminated.

Note

Do not put trilinks on a DWZZH UltraSCSI hub because it is not possible to remove the DWZZH internal termination.

- Require that termination power (`termpwr`) be provided by the SCSI bus host adapters on each SCSI bus segment.

Note

The UltraSCSI hubs are designed to sense loss of termination power (such as a cable pull or `termpwr` not enabled on the host adapter) and shut down the applicable port to prevent corrupted signals on the remaining SCSI bus segments.

3.6.1.1 DS-DWZZH-03 Description

The DS-DWZZH-03:


- Is an 8.9-centimeter (3.5-inch) StorageWorks building block (SBB).
- Can be installed in:

- A StorageWorks UltraSCSI BA356 storage shelf (which has the required 180-watt power supply).
- The lower righthand device slot of the BA370 shelf within the RA7000 or ESA 10000 RAID array subsystems. This position minimizes cable lengths and interference with disks.
- A non-UltraSCSI BA356 that has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.
- Uses the storage shelf only to provide its power and mechanical support. (It is not connected to the shelf internal SCSI bus.)
- Has three Very High Density Cable Interconnect (VHDCI) differential SCSI bus connectors.
- Does not use a SCSI ID.

DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI hubs may be housed in the same storage shelf with disk drives. Table 3–3 provides the supported configurations.

Figure 3–2 shows a front view of the DS-DWZZH-03 UltraSCSI hub.

Figure 3–2: DS-DWZZH-03 Front View



ZK-1412U-AI

The differential symbol (and the lack of a single-ended symbol) indicates that all three connectors are differential.

3.6.1.2 DS-DWZZH-05 Description

The DS-DWZZH-05:

- Is a 13.33-centimeter (5.25-inch) StorageWorks building block (SBB).
- Has five Very High Density Cable Interconnect (VHDCI) differential SCSI bus connectors.
- Uses SCSI ID 7 whether or not fair arbitration mode is enabled. Therefore, you cannot use SCSI ID 7 on the member systems' SCSI bus adapter.

The following section describes how to prepare the DS-DWZZH-05 UltraSCSI hub for use on a shared SCSI bus in more detail.

3.6.1.2.1 DS-DWZZH-05 Configuration Guidelines

The DS-DWZZH-05 UltraSCSI hub can be installed in:

- A StorageWorks UltraSCSI BA356 shelf (which has the required 180-watt power supply).
- A non-UltraSCSI BA356 that has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.

Note

Dual power supplies are recommended for any BA356 shelf containing a DS-DWZZH-05 UltraSCSI hub in order to provide a higher level of availability between cluster member systems and storage.

- The lower righthand device slot of the BA370 shelf within the RA7000 or ESA 10000 RAID array subsystems. This position minimizes cable lengths and interference with disks.

A DS-DWZZH-05 UltraSCSI hub uses the storage shelf only to provide its power and mechanical support (it is not connected to the shelf internal SCSI bus).

Note

When the DS-DWZZH-05 is installed, its orientation is rotated 90 degrees counterclockwise from what is shown in Figure 3-3 and Figure 3-4.

The maximum configurations with combinations of DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI hubs, and disks in the same storage shelf containing dual 180-watt power supplies, are listed in Table 3-3.

Note

With dual 180-watt power supplies installed, there are slots available for six 8.9-centimeter (3.5-inch) SBBs or two 13.33-centimeter (5.25-inch) SBBs.

Table 3–3: DS-DWZZH UltraSCSI Hub Maximum Configurations

DS-DWZZH-03	DS-DWZZH-05	Disk Drives ^a	Personality Module ^{b c}
5	0	0	Not Installed
4	0	0	Installed
3	0	3	Installed
2	0	4	Installed
1	0	5	Installed
0	2	0	Not Installed
3	1	0	Not Installed
2	1	1	Installed
1	1	2	Installed
0	1	3	Installed

^a DS-DWZZH UltraSCSI hubs and disk drives may coexist in a storage shelf. Installed disk drives are not associated with the DS-DWZZH UltraSCSI hub SCSI bus segments; they are on the SCSI bus that is connected to the personality module.

^b If the personality module is installed, you can install a maximum of four DS-DWZZH-03 UltraSCSI hubs.

^c The personality module must be installed to provide a path to any disks that are installed in the storage shelf.

3.6.1.2.2 DS-DWZZH-05 Fair Arbitration

Although each cluster member system and storage controller that are connected to an UltraSCSI hub are on separate SCSI bus segments, they all share a common SCSI bus and its bandwidth. As the number of systems accessing the storage controllers increases, the adapter with the highest priority SCSI ID will probably obtain a higher proportion of the UltraSCSI bandwidth.

The DS-DWZZH-05 UltraSCSI hub provides a fair arbitration feature that overrides the traditional SCSI bus priority. Fair arbitration applies only to the member systems, not to the storage controllers (which are assigned higher priority than the member system host adapters).

You enable fair arbitration by placing the switch on the front of the DS-DWZZH-05 UltraSCSI hub to the `Fair` position. (See Figure 3–4.)

Fair arbitration works as follows. The DS-DWZZH-05 UltraSCSI hub is assigned the highest SCSI ID, which is 7. During the SCSI arbitration phase, the hub, because it has the highest priority, captures the SCSI ID of all host adapters arbitrating for the bus. The hub compares the SCSI IDs of the host adapters requesting use of the SCSI bus, and then allows the device with the highest priority SCSI ID to take control of the SCSI bus. That SCSI ID is removed from the group of captured SCSI IDs prior to the next comparison.

After the host adapter has been serviced, if there are still SCSI IDs retained from the previous arbitration cycle, the next highest SCSI ID is serviced.

When all devices in the group have been serviced, the DS-DWZZH-05 repeats the sequence at the next arbitration cycle.

Fair arbitration is disabled by placing the switch on the front of the DS-DWZZH-05 UltraSCSI hub in the `Disable` position. (See Figure 3–4.) With fair arbitration disabled, the SCSI requests are serviced in the conventional manner; the highest SCSI ID asserted during the arbitration cycle obtains use of the SCSI bus.

Note

Host port SCSI ID assignments are not linked to the physical port when fair arbitration is disabled.

The DS-DWZZH-05 reserves SCSI ID 7 regardless of whether fair arbitration is enabled or not.

3.6.1.2.3 DS-DWZZH-05 Address Configurations

The DS-DWZZH-05 has two addressing modes: wide addressing mode and narrow addressing mode. With either addressing mode, if fair arbitration is enabled, each hub port is assigned a specific SCSI ID. This allows the fair arbitration logic in the hub to identify the SCSI ID of the device participating in the arbitration phase of the fair arbitration cycle.

Caution


If fair arbitration is enabled, the SCSI ID of the host adapter must match the SCSI ID assigned to the hub port. Mismatching or duplicating SCSI IDs will cause the hub to hang.

SCSI ID 7 is reserved for the DS-DWZZH-05 whether fair arbitration is enabled or not.

Jumper W1, which is accessible from the rear of the DS-DWZZH-05 (see Figure 3–3), determines which addressing mode is used. The jumper is installed to select narrow addressing mode. If fair arbitration is enabled, the SCSI IDs for the host adapters are 0, 1, 2, and 3. (See the port numbers not in parentheses in Figure 3–4.) The controller ports are assigned SCSI IDs 4 through 6, and the hub uses SCSI ID 7.


If jumper W1 is removed, the host adapter ports assume SCSI IDs 12, 13, 14, and 15. The controllers are assigned SCSI IDs 0 through 6. The DS-DWZZH-05 retains the SCSI ID of 7.

Figure 3-3: DS-DWZZH-05 Rear View



ZK-1448U-AI

Figure 3-4: DS-DWZZH-05 Front View



3.6.1.2.4 SCSI Bus Termination Power

Each host adapter that is connected to a DS-DWZZH-05 UltraSCSI hub port must supply termination power (`termpwr`) to enable the termination resistors on each end of the SCSI bus segment. If the host adapter is disconnected from the hub, the port is disabled. Only the UltraSCSI bus segment losing termination power is affected. The remainder of the SCSI bus operates normally.

3.6.1.2.5 DS-DWZZH-05 Indicators

The DS-DWZZH-05 has two indicators on the front panel. (See Figure 3-4.) The green LED indicates that power is applied to the hub. The yellow LED indicates that the SCSI bus is busy.

3.6.1.3 Installing the DS-DWZZH-05 UltraSCSI Hub

To install the DS-DWZZH-05 UltraSCSI hub, follow these steps:

1. Remove the W1 jumper to enable wide addressing mode. (See Figure 3–3.)
2. If fair arbitration is to be used, ensure that the switch on the front of the DS-DWZZH-05 UltraSCSI hub is in the `Fair` position.
3. Install the DS-DWZZH-05 UltraSCSI hub in a UltraSCSI BA356, non-UltraSCSI BA356 (if it has the required 180-watt power supply), or BA370 storage shelf.

3.7 Preparing the UltraSCSI Storage Configuration

A TruCluster Server cluster provides you with high data availability through the Cluster File System (CFS), the device request dispatcher, service failover through the cluster application availability (CAA) subsystem, disk mirroring, and fast file system recovery. TruCluster Server supports mirroring of the member-specific boot disks and the cluster quorum disk through hardware RAID only. You can mirror the clusterwide root (`/`), `/usr` and `/var` file systems, the data disks, and the swap disk using the Logical Storage Manager (LSM) technology. You must determine the storage configuration that will meet your needs. Mirroring disks across two shared buses provides the most highly available data.

To determine the supported storage shelves, disk devices, and RAID array controllers, see the AlphaServer options list for your system at the following URL: <http://www.compaq.com/alphaserver/products/options.html>

Disk devices that are used on the shared bus must be installed in a supported storage shelf or behind a RAID array controller. Before you connect a storage shelf to a shared SCSI bus, you must install the disks in the unit. Before connecting a RAID array controller to a shared SCSI bus, install the disks and configure the storagesets. For detailed information about installation and configuration, see your storage shelf (or RAID array controller) documentation.

Note

The following sections mention only the KZPBA-CB UltraSCSI host bus adapter because it is needed to obtain UltraSCSI speeds for UltraSCSI configurations. The KZPSA-BB host bus adapter may be used in any configuration in place of the KZPBA-CB without any cable changes providing that it is supported on

the member system and storage device. Be aware though, the KZPSA-BB is not an UltraSCSI device and therefore only works at fast-wide speed (20 MB/sec).

The following sections describe how to prepare and install cables for storage configurations on a shared SCSI bus using UltraSCSI hubs and the HSZ70 and HSZ80 RAID array controllers, or the RAID Array 3000.

3.7.1 Configuring Radially Connected TruCluster Server Clusters with UltraSCSI Hardware

Radial configurations with RAID array controllers allow you to take advantage of the benefits of hardware mirroring, and to achieve a no-single-point-of-failure (NSPOF) cluster. Typical RAID array storage subsystems used in TruCluster Server cluster configurations are:

- RA7000 or ESA10000 with HSZ70 controller
- RA7000 or ESA10000 with HSZ80 controller
- RA8000 or ESA12000 with HSZ80 controller
- RA3000 with HSZ22 controller

Note

You cannot achieve a NSPOF configuration with a RA3000.

When used with TruCluster Server, one advantage of using a RAID array controller is the ability to hardware mirror the clusterwide root (/) file system, member system boot disks, swap disk, and quorum disk.

HSZ70 or HSZ80

When used in a dual-redundant configuration, Tru64 UNIX Version 5.1A supports both transparent failover, which occurs automatically, without host intervention, and multiple-bus failover, which requires host intervention for some failures.

Note

Enable mirrored cache for dual-redundant configurations to further ensure the availability of unwritten cache data.

Use transparent failover if you only have one shared SCSI bus. Both controllers are connected to the same host and device buses, and either controller can service all of the units if the other controller fails.

Transparent failover compensates only for a controller failure, and not for failures of either the SCSI bus or host adapters and is therefore not a NSPOF configuration.

Note

Set each controller to transparent failover mode before configuring devices (SET FAILOVER COPY = THIS_CONTROLLER).

To achieve a NSPOF configuration, you need multiple-bus failover and two shared SCSI buses.

You may use multiple-bus failover (SET MULTIBUS_FAILOVER COPY = THIS_CONTROLLER) to help achieve a NSPOF configuration if each host has two shared SCSI buses to the array controllers. One SCSI bus is connected to one controller and the other SCSI bus is connected to the other controller. Each member system has a host bus adapter for each shared SCSI bus. The load can be distributed across the two controllers. In case of a host adapter or SCSI bus failure, the host can redistribute the load to the surviving controller. In case of a controller failure, the surviving controller will handle all units.

Notes

Multiple-bus failover does not support device partitioning with the HSZ70 or HSZ80.

Partitioned storagesets and partitioned single-disk units cannot function in multiple-bus failover dual-redundant configurations. Because they are not supported, you must delete your partitions before configuring the HSZ70 or HSZ80 controllers for multiple-bus failover.

Device partitioning is supported with HSG60 and HSG80 array controllers with ACS Version 8.5.

Multiple-bus failover does not support tape drives or CD-ROM drives.

RA3000

The RA3000 uses either active/active or active/passive mode and does not support transparent or multiple-bus failover.

In the active/active mode, the top controller sees one host port as active, while the other controller sees the other host port as active. The controllers

see their non-active host ports as passive. If one of the controllers fails, the surviving controller sees both host ports as active.

In the active/passive mode, the primary controller sees both host ports as active. The other controller sees both host ports as passive. If the primary controller fails, the remaining controller takes over and sees both host ports as active.

The following sections describe how to cable the HSZ70, HSZ80, or RA3000 for TruCluster Server configurations using an UltraSCSI hub. See Chapter 9 and Chapter 10 for information on configurations using external termination. See Chapter 6 for information regarding Fibre Channel storage.

3.7.1.1 Preparing an HSZ70 or HSZ80 for a Shared SCSI Bus Using Transparent Failover Mode

When using transparent failover mode:

- Both controllers of an HSZ70 are connected to the same shared SCSI bus
- For an HSZ80:
 - Port 1 of controller A and Port 1 of controller B are on the same SCSI bus.
 - If used, Port 2 of controller A and Port 2 of controller B are on the same SCSI bus.
 - HSZ80 targets assigned to Port 1 cannot be seen by Port 2.

To cable a dual-redundant HSZ70 or HSZ80 for transparent failover in a TruCluster Server configuration using a DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub, see Figure 3–5 (HSZ70) or Figure 3–6 (HSZ80) and follow these steps:

1. You will need two H8861-AA VHDCI trilink connectors. Install an H8863-AA VHDCI terminator on one of the trilinks.
2. Attach the trilink with the terminator to the controller that you want to be on the end of the shared SCSI bus. Attach an H8861-AA VHDCI trilink connector to:
 - HSZ70 controller A and controller B
 - HSZ80 Port 1 (2) of controller A and Port 1 (2) of controller B

Note

You must use the same port on each HSZ80 controller.

3. Install a BN37A cable between the trilinks on:

- HSZ70 controller A and controller B
- HSZ80 controller A Port 1 (2) and controller B Port 1 (2)

The BN37A-0C is a 30-centimeter (11.8-inch) cable and the BN37A-0E is a 50-centimeter (19.7-inch) cable.

4. Install the DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub in an UltraSCSI BA356, non-UltraSCSI BA356 (with the required 180-watt power supply), or BA370 storage shelf. (See Section 3.6.1.1 or Section 3.6.1.2.)
5. If you are using a:
 - DWZZH-03: Install a BN37A cable between any DWZZH-03 port and the open trilink connector on HSZ70 controller A (B) or HSZ80 controller A Port 1 (2) or controller B Port 1 (2).
 - DWZZH-05:
 - a. Verify that the fair arbitration switch is in the Fair position to enable fair arbitration. (See Section 3.6.1.2.2.)
 - b. Ensure that the W1 jumper is removed to select wide addressing mode. (See Section 3.6.1.2.3.)
 - c. Install a BN37A cable between the DWZZH-05 controller port and the open trilink connector on HSZ70 controller A (B) or HSZ80 controller A Port 1 (2) or controller B Port 1 (2).
6. When the KZPBA-CB host bus adapters in each member system are installed, connect each KZPBA-CB to a DWZZH port with a BN38C (or BN38D) HD68 to VHDCI cable. Ensure that the KZPBA-CB SCSI ID matches the SCSI ID that is assigned to the DWZZH-05 port it is cabled to (12, 13, 14, and 15).

Figure 3–5 shows a two-member TruCluster Server configuration with a radially connected dual-redundant HSZ70 RAID array controller configured for transparent failover.

Figure 3–5: Shared SCSI Bus with HSZ70 Configured for Transparent Failover

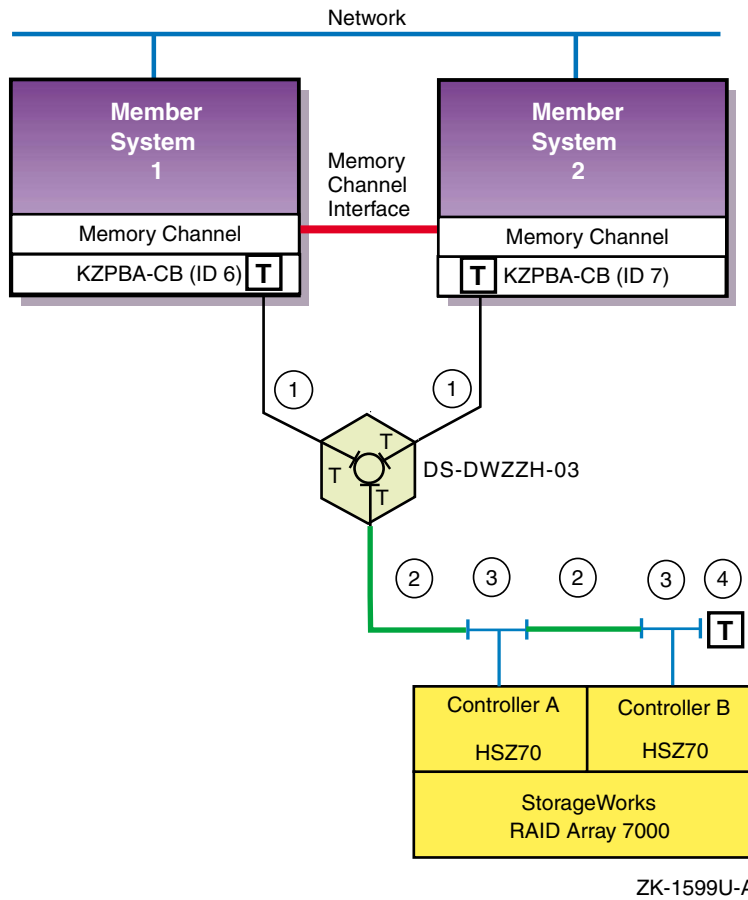


Table 3–4 lists the components that are used to create the clusters that are shown in Figure 3–5, Figure 3–6, Figure 3–7, and Figure 3–8.

Table 3–4: Hardware Components Shown in Figure 3–5 Through Figure 3–8

Callout Number	Description
1	BN38C cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI trilink connector
4	H8863-AA VHDCI terminator ^b

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum combined length of the BN37A cables must not exceed 25 meters (82 feet).

Figure 3–6 shows a two-member TruCluster Server configuration with a radially connected dual-redundant HSZ80 RAID array controller configured for transparent failover.

Figure 3–6: Shared SCSI Bus with HSZ80 Configured for Transparent Failover

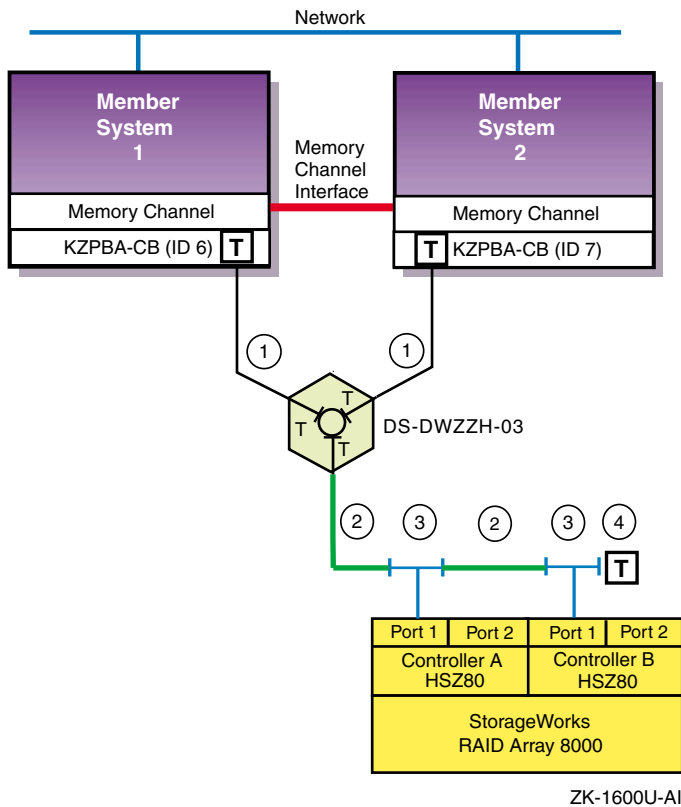


Table 3–4 lists the components that are used to create the cluster that is shown in Figure 3–6.

3.7.1.2 Preparing a Dual-Redundant HSZ70 or HSZ80 for a Shared SCSI Bus Using Multiple-Bus Failover

Multiple-bus failover is a dual-redundant controller configuration in which each host has two paths (two shared SCSI buses) to the array controller subsystem. The hosts have the capability to move LUNs from one controller (shared SCSI bus) to the other. If one host adapter or SCSI bus fails, the hosts can move all storage to the other path. Because both controllers can service all of the units, either controller can continue to service all of the units if the other controller fails. Therefore, multiple-bus failover can compensate for a failed host bus adapter, SCSI bus, or RAID array controller, and can, if the rest of the cluster has necessary hardware, provide a NSPOF configuration.

Note

Each host (cluster member system) requires at least two KZPBA-CB host bus adapters.

Although both the HSZ70 and HSZ80 have multiple-bus failover, they operate differently:

- **HSZ70:** Only one controller (or shared SCSI bus) is active for the units that are preferred (assigned) to it. If all units are preferred to one controller, then all units are accessed through one controller. If a controller detects a problem, all of its units fail over to the other controller. If the host detects a problem with the host bus adapter or SCSI bus, the host initiates the failover to the other controller (and SCSI bus).
- **HSZ80:** Both HSZ80 controllers can be active at the same time. If the host detects a problem with a host bus adapter or SCSI bus, the host initiates the failover to the other controller. If a controller detects a problem, all of its units fail over to the other controller.

Also, the HSZ80 has two ports on each controller. If multiple-bus failover mode is enabled, the targets assigned to any one port are visible to all ports unless access to a unit is restricted to a particular port (on a unit-by-unit basis).

To cable an HSZ70 or HSZ80 for multiple-bus failover in a TruCluster Server configuration using DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hubs (you need two hubs), see Figure 3–7 (HSZ70) and Figure 3–8 (HSZ80) and follow these steps:

1. Install an H8863-AA VHDCI terminator on each of two H8861-AA VHDCI trilinear connectors.

2. Install H8861-AA VHDCI trilink connectors (with terminators) on:
 - HSZ70 controller A and controller B
 - HSZ80 controller A Port 1 (2) and controller B Port 1 (2)

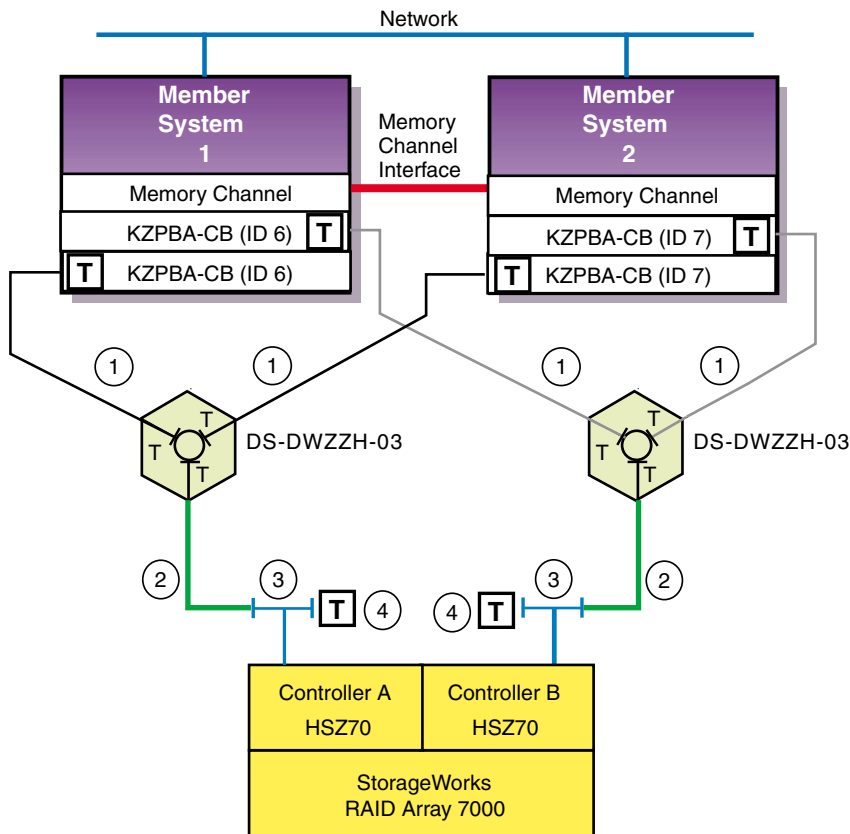
Note

You must use the same port on each HSZ80 controller.

3. Install the DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub in a DS-BA356, BA356 (with the required 180-watt power supply), or BA370 storage shelf. (See Section 3.6.1.1 or Section 3.6.1.2.)
4. If you are using a:
 - DS-DWZZH-03: Install a BN37A VHDCI to VHDCI cable between the trilink connector on controller A (HSZ70) or controller A Port 1 (2) (HSZ80) and any DS-DWZZH-03 port. Install a second BN37A cable between the trilink on controller B (HSZ70) or controller B Port 1 (2) (HSZ80) and any port on the second DS-DWZZH-03.
 - DS-DWZZH-05:
 - a. Verify that the fair arbitration switch is in the Fair position to enable fair arbitration. (See Section 3.6.1.2.2.)
 - b. Ensure that the W1 jumper is removed to select wide addressing mode. (See Section 3.6.1.2.3.)
 - c. Install a BN37A cable between the DWZZH-05 controller port and the open trilink connector on HSZ70 controller A or HSZ80 controller A Port 1 (2).
 - d. Install a second BN37A cable between the second DWZZH-05 controller port and the open trilink connector on HSZ70 controller B or HSZ80 controller B Port 1 (2).
5. When the KZPBA-CBs are installed, use a BN38C (or BN38D) HD68-to-VHDCI cable to connect the first KZPBA-CB on each system to a port on the first DWZZH hub. Ensure that the KZPBA-CB SCSI ID matches the SCSI ID that is assigned to the DWZZH-05 port it is cabled to (12, 13, 14, and 15).
6. Install BN38C (or BN38D) HD68-to-VHDCI cables to connect the second KZPBA-CB on each system to a port on the second DWZZH hub. Ensure that the KZPBA-CB SCSI ID matches the SCSI ID that is assigned to the DWZZH-05 port it is cabled to (12, 13, 14, and 15).

Figure 3–7 shows a two-member TruCluster Server configuration with a radially connected dual-redundant HSZ70 configured for multiple-bus failover.

Figure 3–7: TruCluster Server Configuration with HSZ70 in Multiple-Bus Failover Mode

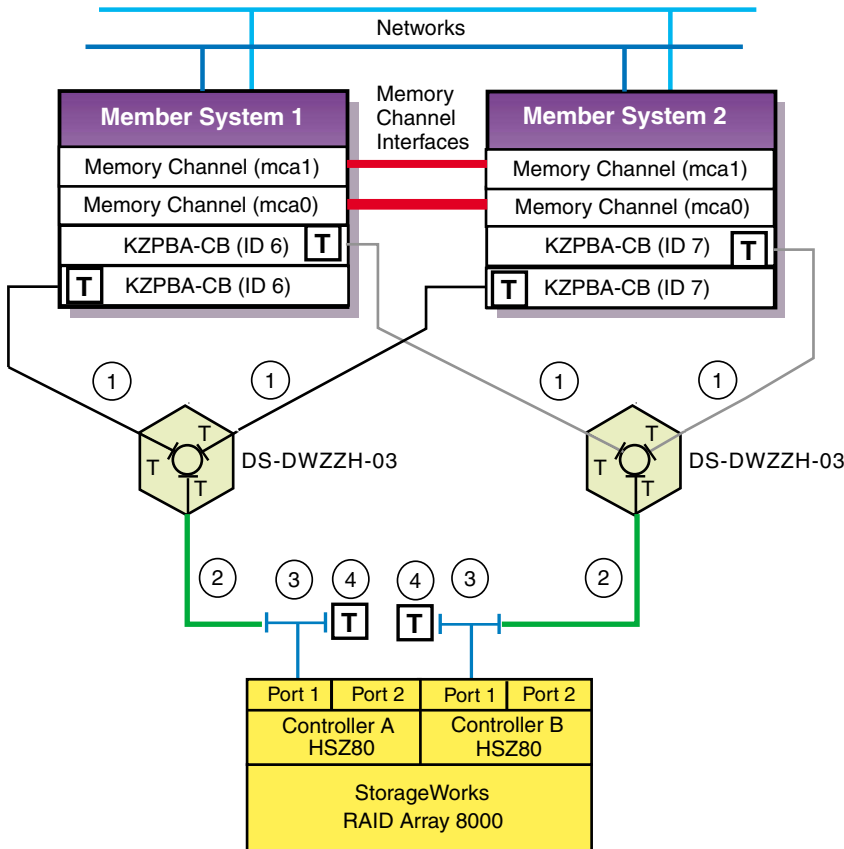


ZK-1601U-AI

Table 3–4 lists the components that are used to create the cluster that is shown in Figure 3–7.

Figure 3–8 shows a two-member TruCluster Server configuration with a radially connected dual-redundant HSZ80 configured for multiple-bus failover.

Figure 3–8: TruCluster Server Configuration with HSZ80 in Multiple-Bus Failover Mode



ZK-1602U-AI

Table 3–4 lists the components that are used to create the cluster that is shown in Figure 3–8.

3.7.1.3 Preparing an RA3000 for Use on a Shared SCSI Bus with an UltraSCSI Hub

The RAID Array 3000 (RA3000) is a low-end, standalone UltraSCSI RAID subsystem. It supports RAID levels 0, 1, 0+1, 4, 5, and just a bunch of disks (JBOD) disks.

The RA3000 storage subsystem has fully redundant components to eliminate single points of failure. It comes with a standard uninterruptible power supply (UPS) for cache data protection during power outages.

The RA3000 uses the dual-ported HSZ22 controller. Optional dual redundant controllers with mirrored write-back cache provide maximum data integrity.

The StorageWorks Command Console (SWCC) V2.2 (or higher) client graphical user interface (GUI) runs on a Microsoft Windows NT V4.0 Service Pack 4 (or later) or Windows 2000 PC connected directly to the RA3000 by a serial line.

After the first virtual disk has been created, you can also communicate with your RAID Array 3000 over a TCP/IP network provided the V2.2 (or higher) SWCC Agent has been installed on the Tru64 UNIX member system.

The RA3000 is available as:

- DS-SWXRA-GH — A rackmount subsystem (standard RETMA or metric cabinet) containing a controller shelf with one HSZ22 controller, an uninterruptible power supply (UPS), two host I/O modules, a device I/O module, and one 6-slot device expansion shelf. Up to three additional expansion shelves (DS-SWXRA-GN) may be added to provide a maximum of 24 storage devices.
- DS-SWXRA-GA — A deskside pedestal subsystem that includes one HSZ22 controller. The base pedestal accommodates up to seven storage devices. The included battery backup subsystem is a free-standing UPS. An expansion pedestal option (DS-SWXRA-GD) increases the storage capacity of the subsystem to 14 storage devices.

A second HSZ22 controller option (DS-HSZ22-AA) can be added to either RA3000 subsystem.

To cable a RA3000 in a TruCluster Server configuration using a DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub, follow the steps in Table 3-5.

Notes

All configuration illustrations assume that a second, redundant HSZ22 controller is installed to achieve active/active or active/passive failover.

See the RA3000 documentation for information about configuring the storage devices.

Table 3–5: Installing Cables for RA3000 Radial Configuration with a DWZZH UltraSCSI Hub

Action	Refer to:
Install a BN38C HD68 to VHDCI cable between each KZPBA-CB UltraSCSI host adapter and a DWZZH port. The DWZZH accepts the VHDCI connector. You may use a BN38E-0B technology adapter cable with a BN37A cable instead of the BN38C cable. ^{a b}	Figure 3–9 through Figure 3–12
Install BN37A cables: ^c	—
RA3000 controller shelf with active/passive failover: Install a BN37A cable between any DWZZH-03 port or the DWZZH-05 controller port and the RA3000 controller shelf Host 0 I/O module Host In port.	Figure 3–9
RA3000 pedestal with active/passive failover: Install a BN37A cable between any DWZZH-03 port or the DWZZH-05 controller port and the RA3000 pedestal Host 0 port.	Figure 3–10
RA3000 pedestal with active/active or active/passive failover: Install a BN37A cable between the DWZZH-05 controller port and the RA3000 pedestal Host 0 port. Install a second BN37A cable between a DWZZH-05 host port and the RA3000 pedestal Host 1 port.	Figure 3–11
RA3000 controller shelf with active/active or active/passive failover: Install a BN37A cable between any DWZZH-03 port or the DWZZH-05 controller port and the RA3000 controller shelf Host 0 I/O module Host In port. Install a BN37A-0E 50-centimeter (19.7-inch) cable between the Host 0 I/O module Host Out port and the Host 1 I/O module Host In port.	Figure 3–12

Note

If you connect a DWZZH-05 host port to an RA3000 pedestal host port to provide active/active failover, you must disable fair arbitration on the DWZZH-05 by placing the fair arbitration switch in the `DISABLE` position.

^a The maximum length of the SCSI bus segment, including the combined length of BN38C cables and internal device length, must not exceed 25 meters (82 feet).

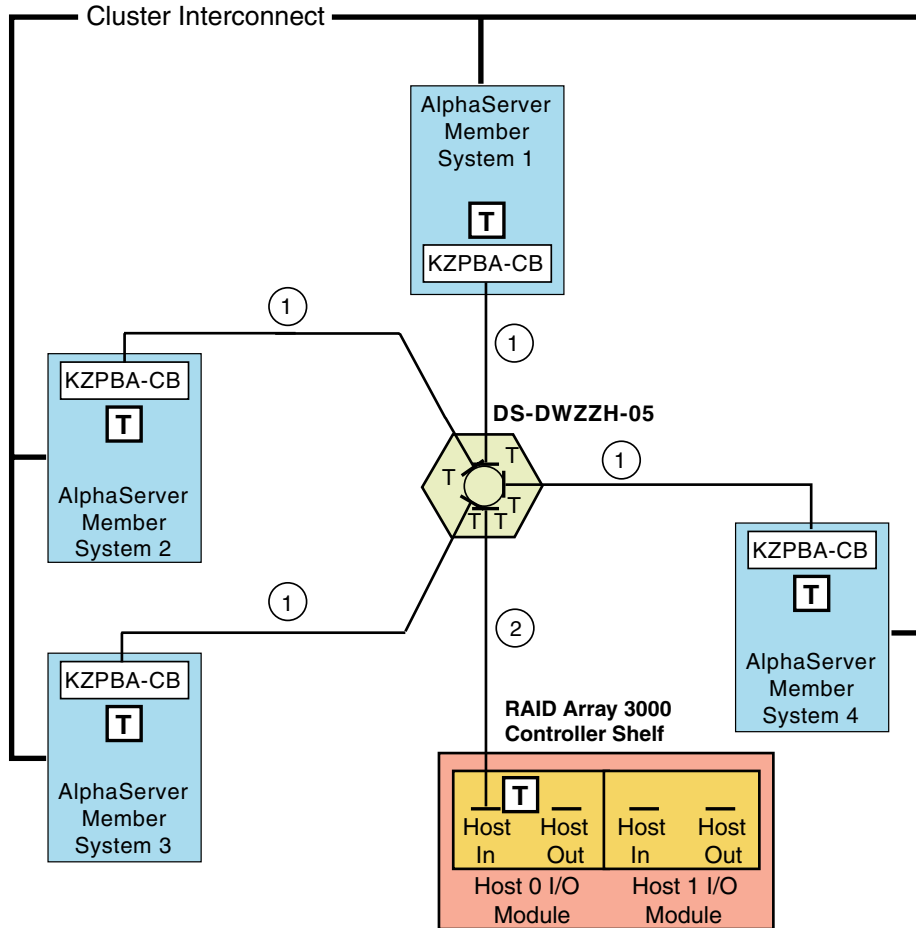
^b The maximum length of the SCSI bus segment, including the combined length of BN38E-0B and BN37A cables and internal device length, must not exceed 25 meters (82 feet).

^c The maximum length of the SCSI bus segment, including the BN37A cables and internal device length, must not exceed 25 meters (82 feet).

In the configuration shown in Figure 3–9 through Figure 3–12, it is assumed that the RA3000 contains two HSZ22 controllers.

Figure 3–9 shows a four-member TruCluster Server configuration and an RA3000 controller shelf with active/passive failover radially connected to a DS-DWZZH-05 UltraSCSI hub. Table 3–6 describes the callouts.

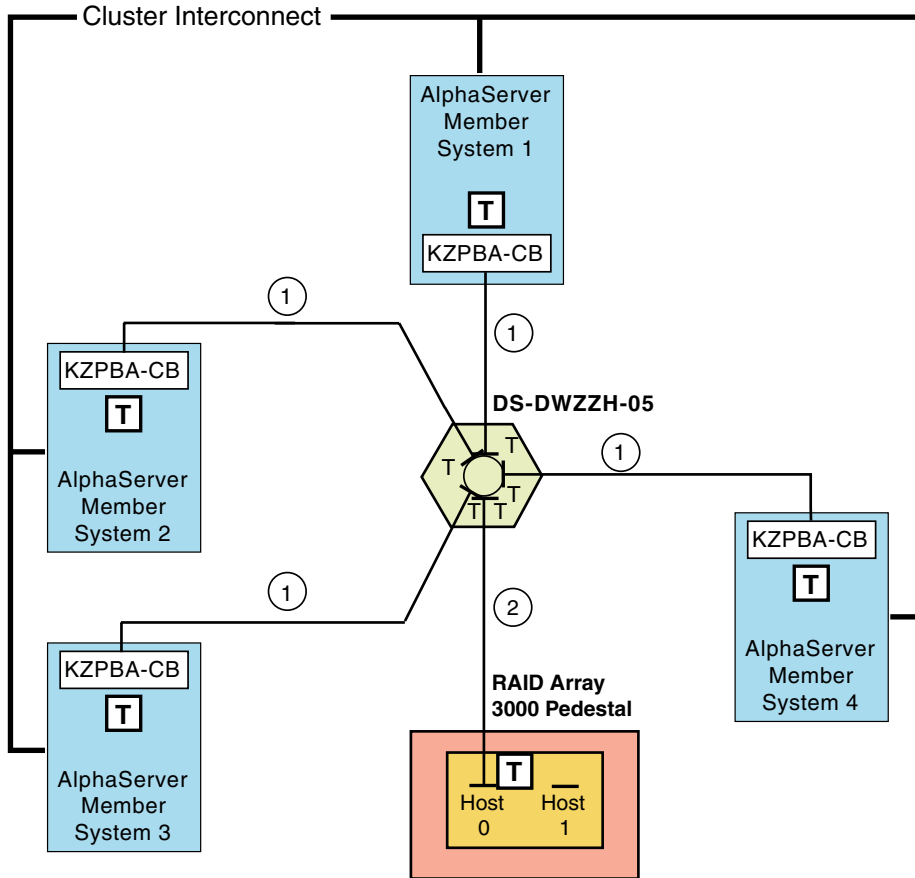
Figure 3–9: TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Passive Failover



ZK-1477U-AI

Figure 3–10 shows a four-member TruCluster Server configuration and an RA3000 pedestal with active/passive failover radially connected to a DS-DWZZH-05 UltraSCSI hub. Table 3–6 describes the callouts.

Figure 3–10: TruCluster Server Configuration with an RA3000 Pedestal with Active/Passive Failover and a DWZZH-05 UltraSCSI Hub



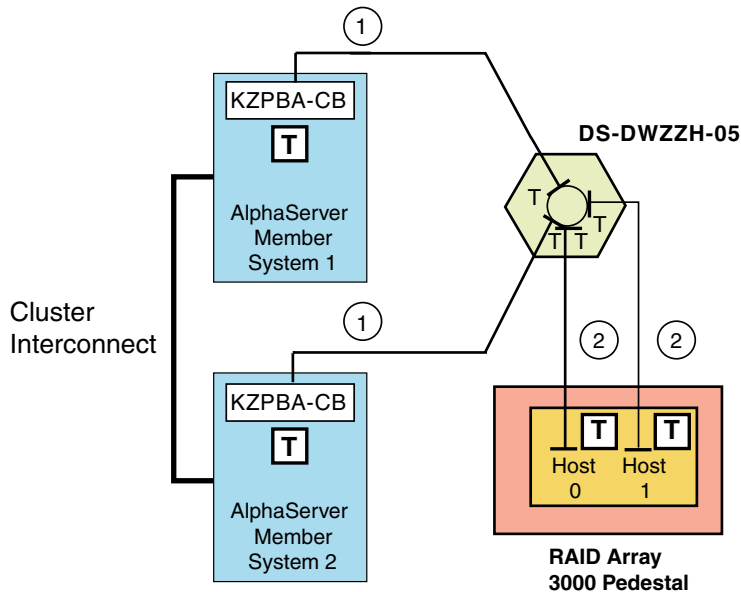
ZK-1478U-AI

Figure 3–11 shows a two-member TruCluster Server configuration and an RA3000 pedestal with active/active or active/passive failover radially connected to a DS-DWZZH-05 UltraSCSI hub. This configuration uses independent connections to the two pedestal host ports to increase the available bandwidth to the RA3000 controllers. Table 3–6 describes the callouts.

Note

If you connect a DWZZH-05 host port to an RA3000 pedestal host port to provide active/active failover, you must disable fair arbitration on the DWZZH-05 by placing the fair arbitration switch in the DISABLE position.

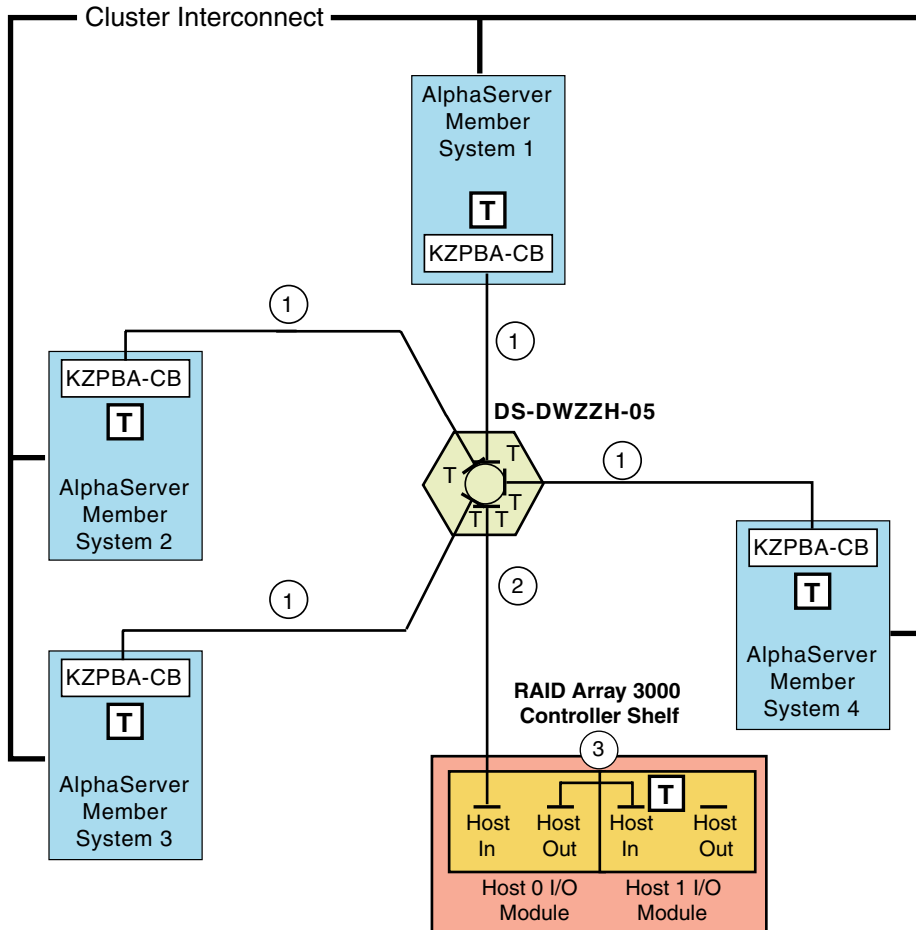
Figure 3–11: TruCluster Server Configuration with an RA3000 Pedestal with Active/Active or Active/Passive Failover



ZK-1480U-AI

Figure 3–12 shows a four-member TruCluster Server configuration and an RA3000 controller shelf with active/active or active/passive failover radially connected to a DS-DWZZH-05 UltraSCSI hub. Table 3–6 describes the callouts.

Figure 3–12: TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Active or Active/Passive Failover



ZK-1479U-AI

Table 3–6: Hardware Components Used in the Configurations Shown in Figure 3–9 through Figure 3–12

Callout Number	Description
1	BN38C HD68-to-VHDCI cable. ^a A BN38E-0B technology adapter cable may be connected to a BN37A cable and used in place of a BN38C cable. ^b
2	BN37A VHDCI cable ^c
3	BN37A-0E 50-centimeter (19.7-inch) VHDCI cable ^d

^a The maximum length of the SCSI bus segment, including the combined length of BN38C cables and internal device length, must not exceed 25 meters (82 feet).

^b The maximum length of the SCSI bus segment, including the combined length of BN38E-0B and BN37A cables and internal device length, must not exceed 25 meters (82 feet).

^c The maximum length of the SCSI bus segment, including the BN37A cables and internal device length, must not exceed 25 meters (82 feet).

^d This cable is used only in the configuration shown in Figure 3–12.

4

TruCluster Server System Configuration Using UltraSCSI Hardware

This chapter describes how to prepare systems for a TruCluster Server cluster, using UltraSCSI hardware and the preferred method of radial configuration, including how to connect devices to a shared SCSI bus for the TruCluster Server product. This chapter does not provide detailed information about installing devices; it describes only how to set up the hardware in the context of the TruCluster Server product. Therefore, you must have the documentation that describes how to install the individual pieces of hardware. That documentation typically arrives with the hardware.

The chapter discusses the following topics:

- Planning a TruCluster Server hardware configuration (Section 4.1)
- Obtaining the firmware release notes (Section 4.2)
- Installing hardware in a TruCluster Server configuration (Section 4.3)

All systems in the cluster must be connected via the Memory Channel cluster interconnect. Not all members must be connected to a shared SCSI bus.

You need to allocate disks for the following uses:

- 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.
- One or more disks on a shared SCSI bus to hold the clusterwide root (/), /usr, and /var AdvFS file systems.
- One disk per member, normally on a shared SCSI bus, to hold member boot partitions.
- Optionally, one disk on a shared SCSI bus to act as the quorum disk. See Section 1.3.1.4 and, for a more detailed discussion of the quorum disk, see the *Cluster Administration* manual.

All configurations that are covered in this manual are based on the use of a shared SCSI bus.

Note

If you are using Fibre Channel storage, see Chapter 6.

Before you connect devices to a shared SCSI bus, you must:

- Plan your hardware configuration, determining which devices will be connected to each shared SCSI bus, which devices will be connected together, and which devices will be at the ends of each bus.

Planning is especially critical if you will install tape devices on the shared SCSI bus. With the exception of the TZ885, TZ887, TL890, TL891, and TL892, tape devices can only be installed at the end of a shared SCSI bus. These tape devices are the only supported tape devices that can be terminated externally.

- Place the devices as close together as possible and ensure that shared SCSI buses will be within length limitations.
- Prepare the systems and storage shelves for the appropriate bus connection, including installing SCSI controllers, UltraSCSI hubs, tralink connectors, and SCSI signal converters.

After you install all necessary cluster hardware and connect the shared SCSI buses, be sure that the systems can recognize and access all the shared disks. (See Section 4.3.2.) You can then install the TruCluster Server software as described in the *Cluster Installation* manual.

4.1 Planning Your TruCluster Server Hardware Configuration

Before you set up a TruCluster Server hardware configuration, you must plan a configuration to meet your performance and availability needs. You must determine the following components for your configuration:

- Number and type of member systems and the number of shared SCSI buses

You can use two to eight member systems for TruCluster Server. A greater number of member systems connected to shared SCSI buses gives you better application performance and more availability. However, all the systems compete for the same buses to service I/O requests, so a greater number of systems decreases I/O performance.

Each member system must have a supported SCSI adapter for each shared SCSI bus connection. There must be enough PCI slots for the Memory Channel cluster interconnects and SCSI adapters. The number of available PCI slots depends on the type of AlphaServer system.

- Cluster interconnects

You need only one cluster interconnect in a cluster. For TruCluster Server Version 5.1A, the cluster interconnect can be the Memory Channel or a private LAN. (See *Cluster LAN Interconnect* for more information on using a private LAN as the cluster interconnect.) However, you can use redundant cluster interconnects to protect against an interconnect failure and for easier hardware maintenance. If you have more than two member systems, you must have one Memory Channel hub for each interconnect.

- Number of shared SCSI buses and the storage on each shared bus

Using shared SCSI buses increases storage availability. You can connect 32 shared SCSI buses to a cluster member. You can use any combination of KZPSA-BB, KZPBA-CB, or KGPSA-BC/CA host bus adapters.

In addition, RAID array controllers allow you to increase your storage capacity and protect against disk, controller, host bus adapter, and SCSI bus failures. Mirroring data across shared buses provides you with more reliable and available data. You can use Logical Storage Manager (LSM) host-based mirroring for all storage except the member-specific boot disks and the quorum disk.

- No single-point-of-failure (NSPOF) TruCluster Server cluster

You can use mirroring and multiple-bus failover with the HSZ70, HSZ80, and HSG80 RAID array controllers to create an NSPOF TruCluster Server cluster (provided that the rest of the hardware is installed).

- Tape loaders on a shared SCSI bus

Because of the length of the internal SCSI cables in some tape loaders (up to 3 meters (9.8 feet)), they cannot be externally terminated with a tralink/terminator combination. Therefore, in general, with the exception of the TL890, TL891, and TL892, tape loaders must be on the end of the shared SCSI bus. See Chapter 8 for information on configuring tape devices on a shared SCSI bus.

- You cannot use Prestoserve in a TruCluster Server cluster to cache I/O operations for any storage device, regardless of whether it is located on a shared bus or a bus local to a given system. Because data in the Prestoserve buffer cache of one member is not accessible to other member systems, TruCluster Server cannot provide correct failover when Prestoserve is being used.

Table 4–1 describes how to maximize performance, availability, and storage capacity in your TruCluster Server hardware configuration. For example, if you want greater application performance without decreasing I/O performance, you can increase the number of member systems or you can set up additional shared storage.

Table 4–1: Planning Your Configuration

To increase:	You can:
Application performance	Increase the number of member systems.
I/O performance	Increase the number of shared buses.
Member system availability	Increase the number of member systems.
Cluster interconnect availability	Use redundant cluster interconnects.
Disk availability	Mirror disks across shared buses. Use a RAID array controller.
Shared storage capacity	Increase the number of shared buses. Use a RAID array controller. Increase disk size.

4.2 Obtaining the Firmware Release Notes

You may be required to update the system or SCSI controller firmware during a TruCluster Server installation, so you may need the firmware release notes.

You can obtain the firmware release notes from:

- The Web at the following URL: <http://www.compaq.com/support/>
Select software & drivers, then select AlphaServer under the servers heading. Select the appropriate system.
- The current Alpha Systems Firmware Update CD-ROM.

Note

To obtain the firmware release notes from the Firmware Update Utility CD-ROM, your kernel must be configured for the ISO 9660 Compact Disk File System (CDFS).

To obtain the release notes for the firmware update, follow these steps:

1. At the console prompt, or using the system startup log if the Tru64 UNIX operating system is running, determine the drive number of the CD-ROM.
2. Boot the Tru64 UNIX operating system if it is not already running.
3. Log in as root.
4. Place the Alpha Systems Firmware Update CD-ROM for the installed (or to be installed) Tru64 UNIX version into the drive.

5. Mount the CD-ROM as follows (`/dev/disk/cdrom0c` is used as an example CD-ROM drive):

```
# mount -rt cdfs -o noversion /dev/disk/cdrom0cc /mnt
```

6. Copy the appropriate release notes to your system disk. In this example, obtain the firmware release notes for the AlphaServer DS20 from the Version 5.6 Alpha Firmware Update CD-ROM:

```
# cp /mnt/doc/ds20_v56_fw_relnote.txt ds20-rel-notes
```

7. Unmount the CD-ROM drive:

```
# umount /mnt
```

8. Print the release notes.

4.3 TruCluster Server Hardware Installation

Member systems may be connected to a shared SCSI bus with a peripheral component interconnect (PCI) SCSI adapter. Before you install a PCI SCSI adapter into a PCI slot on a member system, ensure that the module is at the correct hardware revision.

The qualification and use of the DS-DWZZH-series UltraSCSI hubs in TruCluster Server clusters allows the PCI host bus adapters to be cabled into a cluster in two different ways:

- Preferred method with radial connection to a DWZZH UltraSCSI hub and internal termination: The PCI host bus adapter internal termination resistor SIPs are not removed. The host bus adapters and storage subsystems are connected directly to a DWZZH UltraSCSI hub port. There can be only one member system connected to a hub port.

The use of a DWZZH UltraSCSI hub in a TruCluster Server cluster is preferred because it improves the reliability to detect cable faults.

- Old method with external termination: Shared SCSI bus termination is external to the PCI host adapters. This is the old method used to connect a PCI host adapter to the cluster; remove the adapter termination resistor SIPs and install a Y cable and an H879-AA terminator for external termination. This allows the removal of a SCSI bus cable from the host adapter without affecting SCSI bus termination.

This method (which is discussed in Chapter 10 and Chapter 9) may be used with or without a DWZZH UltraSCSI hub. When used with an UltraSCSI hub, there may be more than one member system on a SCSI bus segment attached to a DS-DWZZH-03 hub port.

The following sections describe how to install the KZPBA-CB PCI-to-UltraSCSI differential host adapter and configure them into

TruCluster Server clusters using the preferred method of radial connection with internal termination.

Note

The KZPSA-BB can be used in any configuration in place of the KZPBA-CB. The use of the KZPSA-BB is not mentioned in this chapter because it is not UltraSCSI hardware, and it cannot operate at UltraSCSI speeds.

The use of the KZPSA-BB (and the KZPBA-CB) with external termination is discussed in Chapter 9.

It is assumed that when you start to install the hardware necessary to create a TruCluster Server configuration, you have sufficient storage to install the TruCluster Server software, and that you have set up any RAID storagesets.

Follow the steps in Table 4–2 to start the procedure for TruCluster Server hardware installation. You can save time by installing the Memory Channel adapters, redundant network adapters (if applicable), and KZPBA-CB SCSI adapters all at the same time.

Follow the directions in the referenced documentation, or the steps in the referenced tables, returning to Table 4–2 when you have completed the steps in the referenced table.

Caution

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

Table 4–2: Configuring TruCluster Server Hardware

Step	Action	Refer to:
1	Install the Memory Channel modules, cables, and hubs (if a hub is required).	Chapter 5 ^a
2	Install Ethernet or FDDI network adapters. Install ATM adapters if using ATM.	User's guide for the applicable Ethernet or FDDI adapter, and the user's guide for the applicable system The documentation that came with the ATM adapter
3	Install a KZPBA-CB UltraSCSI adapter for each radially connected shared SCSI bus in each member system.	Section 4.3.1 and Table 4–3

Table 4–2: Configuring TruCluster Server Hardware (cont.)

Step	Action	Refer to:
4	Update the system SRM console firmware from the latest Alpha Systems Firmware Update CD-ROM.	The firmware update release notes (Section 4.2)

Note

The SRM console firmware includes the ISP1020/1040-based PCI option firmware, which includes the KZPBA-CB. When you update the SRM console firmware, you are enabling the KZPBA-CB firmware to be updated. On a powerup reset, the SRM console loads KZPBA-CB adapter firmware from the console system flash ROM into NVRAM for all Qlogic ISP1020/1040-based PCI options, including the KZPBA-CB PCI-to-Ultra SCSI adapter.

^a If you install additional KZPBA-CB SCSI adapters or an extra network adapter at this time, delay testing the Memory Channel until you have installed all of the hardware.

4.3.1 Installation of a KZPBA-CB Using Internal Termination for a Radial Configuration

Use this method of cabling member systems and shared storage in a TruCluster Server cluster if you are using a DWZZH UltraSCSI hub. You must reserve at least one hub port for shared storage.

The DWZZH-series UltraSCSI hubs are designed to allow more distance between member systems and shared storage. Using the UltraSCSI hub also improves the reliability of the detection of cable faults.

Another benefit is the ability to connect the member systems' SCSI adapter directly to a hub port without external termination. This simplifies the configuration by reducing the number of cable connections.

A DWZZH UltraSCSI hub can be installed in:

- A StorageWorks UltraSCSI BA356 shelf that has the required 180-watt power supply.
- The lower righthand device slot of the BA370 shelf within the RA7000 or ESA 10000 RAID array subsystems. This position minimizes cable lengths and interference with disks.
- A non-UltraSCSI BA356 that has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.

An UltraSCSI hub only receives power and mechanical support from the storage shelf. There is no SCSI bus continuity between the DWZZH and storage shelf.

The DWZZH contains a differential to single-ended signal converter for each hub port (which is sometimes referred to as a DWZZA on a chip, or DOC chip). The single-ended sides are connected together to form an internal single-ended SCSI bus segment. Each differential SCSI bus port is terminated internal to the DWZZH with terminators that cannot be disabled or removed.

Power for the DWZZH termination (`termpwr`) is supplied by the host SCSI bus adapter or RAID array controller connected to the DWZZH port. If the member system or RAID array controller is powered down, or the cable is removed from the KZPBA-CB, RAID array controller, or hub port, the loss of `termpwr` disables the hub port without affecting the remaining hub ports or SCSI bus segments. This result is similar to removing a Y cable when using external termination.

Note

The UltraSCSI BA356 DS-BA35X-DA personality module does not generate `termpwr`. Therefore, you cannot connect an UltraSCSI BA356 directly to a DWZZH hub. The use of the UltraSCSI BA356 in a TruCluster Server cluster is discussed in Chapter 10.

The other end of the SCSI bus segment is terminated by the KZPBA-CB onboard termination resistor SIPs, or by a tralink connector/terminator combination installed on the RAID array controller.

The KZPBA-CB UltraSCSI host adapter:

- Is a high-performance PCI option connecting the PCI-based host system to the devices on a 16-bit, ultrawide differential SCSI bus.
- Is installed in a PCI slot of the supported member system.
- Is a single-channel, ultrawide differential adapter.
- Operates at the following speeds:
 - 5 MB/sec narrow SCSI at slow speed
 - 10 MB/sec narrow SCSI at fast speed
 - 20 MB/sec wide differential SCSI
 - 40 MB/sec wide differential UltraSCSI

Note

Even though the KZPBA-CB is an UltraSCSI device, it has an HD68 connector.

Make sure that your storage shelves or RAID array subsystems are set up before completing this portion of an installation.

Use the steps in Table 4–3 to set up a KZPBA-CB for a TruCluster Server cluster that uses radial connection to a DWZZH UltraSCSI hub.

Table 4–3: Installing the KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub

Step	Action	Refer to:
1	Ensure that the eight KZPBA-CB internal termination resistor SIPs, RM1-RM8 are installed.	Section 4.3.1, Figure 4–1, and <i>KZPBA-CB PCI-to-Ultra SCSI Differential Host Adapter User's Guide</i>
2	Power down the system. Install a KZPBA-CB PCI-to-UltraSCSI differential host adapter in the PCI slot corresponding to the logical bus to be used for the shared SCSI bus. Ensure that the number of adapters are within limits for the system, and that the placement is acceptable.	TruCluster Server <i>Cluster Administration</i> , Section 2.3.3, and <i>KZPBA-CB PCI-to-Ultra SCSI Differential Host Adapter User's Guide</i>
3	Install a BN38C cable between the KZPBA-CB UltraSCSI host adapter and a DWZZH port.	—

Notes

The maximum length of a SCSI bus segment is 25 meters (82 feet), including the bus length internal to the adapter and storage devices.

One end of the BN38C cable is 68-pin high density. The other end is 68-pin VHDCI. The DWZZH accepts the 68-pin VHDCI connector.

The number of member systems in the cluster has to be one less than the number of DWZZH ports.

4	Power up the system and use the <code>show config</code> and <code>show device</code> console commands to display the installed devices and information about the KZPBA-CBs on the AlphaServer systems. Look for QLogic ISP1020 in the <code>show config</code> display and <code>isp</code> in the <code>show device</code> display to determine which devices are KZPBA-CBs.	Section 4.3.2 and Example 4–1 through Example 4–4
5	Use the <code>show pk*</code> or <code>show isp*</code> console commands to determine the KZPBA-CB SCSI bus ID, and then use the <code>set</code> console command to set the SCSI bus ID.	Section 4.3.3 and Example 4–5 through Example 4–7

Table 4-3: Installing the KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub (cont.)

Step	Action	Refer to:
	Notes	
	<p>Ensure that the SCSI ID that you use is distinct from all other SCSI IDs on the same shared SCSI bus. If you do not remember the other SCSI IDs, or do not have them recorded, you must determine these SCSI IDs.</p> <p>If you are using a DS-DWZZH-05, you cannot use SCSI ID 7 for a KZPBA-CB UltraSCSI adapter; SCSI ID 7 is reserved for DS-DWZZH-05 use.</p> <p>If you are using a DS-DWZZH-05 and fair arbitration is enabled, you must use the SCSI ID assigned to the hub port that the adapter is connected to.</p> <p>You will have problems if you have two or more SCSI adapters at the same SCSI ID on any one SCSI bus.</p>	
6	Repeat steps 1 through 5 for any other KZPBA-CBs to be installed on this shared SCSI bus on other member systems.	—
7	Connect a DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub to an:	Section 3.6
	HSZ70 or HSZ80 in transparent failover mode	Section 3.7.1.1
	HSZ70 or HSZ80 in multiple-bus failover mode	Section 3.7.1.2
	RAID Array 3000	Section 3.7.1.3

4.3.2 Displaying KZPBA-CB Adapters with the show Console Commands

Use the `show config` and `show device console` commands to display system configuration. Use the output to determine which devices are KZPBA-CBs, and to determine their SCSI bus IDs.

Example 4-1 shows the output from the `show config console` command on an AlphaServer DS20 system.

Example 4-1: Displaying Configuration on an AlphaServer DS20

```
P00>>> show config
                               AlphaServer DS20 500 MHz

SRM Console:   T5.4-15
PALcode:      OpenVMS PALcode V1.54-43, Tru64 UNIX PALcode V1.49-45

Processors
CPU 0         Alpha 21264-4 500 MHz   SROM Revision: V1.82
              Bcache size: 4 MB
```

Example 4–1: Displaying Configuration on an AlphaServer DS20 (cont.)

```
CPU 1          Alpha 21264-4 500 MHz   SR0M Revision: V1.82
              Bcache size: 4 MB

Core Logic
Cchip         DECchip 21272-CA Rev 2.1
Dchip         DECchip 21272-DA Rev 2.0
Pchip 0       DECchip 21272-EA Rev 2.2
Pchip 1       DECchip 21272-EA Rev 2.2

TIG           Rev 4.14
Arbiter       Rev 2.10 (0x1)

MEMORY
Array #       Size      Base Addr
-----
0             512 MB    000000000

Total Bad Pages = 0
Total Good Memory = 512 MBytes

PCI Hose 00
  Bus 00 Slot 05/0: Cypress 82C693          Bridge to Bus 1, ISA
  Bus 00 Slot 05/1: Cypress 82C693 IDE      dqa.0.0.105.0
  Bus 00 Slot 05/2: Cypress 82C693 IDE      dqb.0.1.205.0
  Bus 00 Slot 05/3: Cypress 82C693 USB
  Bus 00 Slot 07: DECchip 21152-AA          Bridge to Bus 2, PCI
  Bus 00 Slot 08: QLogic ISP1020
      pkc0.7.0.8.0          SCSI Bus ID 7
      dkc0.0.0.8.0          HSZ70
      dkc1.0.0.8.0          HSZ70
      dkc100.1.0.8.0        HSZ70
      dkc101.1.0.8.0        HSZ70CCL
      dkc2.0.0.8.0          HSZ70
      dkc3.0.0.8.0          HSZ70
      dkc4.0.0.8.0          HSZ70
      dkc5.0.0.8.0          HSZ70
      dkc6.0.0.8.0          HSZ70
      dkc7.0.0.8.0          HSZ70
  Bus 00 Slot 09: QLogic ISP1020
      pkd0.7.0.9.0          SCSI Bus ID 7
      dkd0.0.0.9.0          HSZ40
      dkd1.0.0.9.0          HSZ40
      dkd100.1.0.9.0        HSZ40
      dkd101.1.0.9.0        HSZ40
      dkd102.1.0.9.0        HSZ40
      :
      dkd5.0.0.9.0          HSZ40
      dkd6.0.0.9.0          HSZ40
      dkd7.0.0.9.0          HSZ40
  Bus 02 Slot 00: NCR 53C875
      pka0.7.0.2000.0        SCSI Bus ID 7
      dka0.0.0.2000.0        RZ1CB-CS
      dka100.1.0.2000.0      RZ1CB-CS
      dka200.2.0.2000.0      RZ1CB-CS
      dka500.5.0.2000.0      RRD46
```

Example 4–1: Displaying Configuration on an AlphaServer DS20 (cont.)

```
Bus 02 Slot 01: NCR 53C875
                    pkb0.7.0.2001.0      SCSI Bus ID 7
Bus 02 Slot 02: DE500-AA Network Controller
                    ewa0.0.0.2002.0      00-06-2B-00-0A-48

PCI Hose 01
Bus 00 Slot 07: DEC PCI FDDI
                    fwa0.0.0.7.1        08-00-2B-B9-0D-5D
Bus 00 Slot 08: DEC PCI MC
                    Rev: 22, mca0
Bus 00 Slot 09: DEC PCI MC
                    Rev: 22, mcb0

ISA
Slot  Device Name           Type           Enabled BaseAddr  IRQ   DMA
0
  0    MOUSE           Embedded       Yes       60       12
  1    KBD             Embedded       Yes       60       1
  2    COM1            Embedded       Yes       3f8      4
  3    COM2            Embedded       Yes       2f8      3
  4    LPT1           Embedded       Yes       3bc      7
  5    FLOPPY         Embedded       Yes       3f0      6    2
```

Example 4–2 shows the output from the `show device console` command entered on an AlphaServer DS20 system.

Example 4–2: Displaying Devices on an AlphaServer DS20

```
P00>>> show device
dka0.0.0.2000.0      DKA0          RZ1CB-CS 0656
dka100.1.0.2000.0   DKA100       RZ1CB-CS 0656
dka200.2.0.2000.0   DKA200       RZ1CB-CS 0656
dka500.5.0.2000.0   DKA500       RRD46 1337
dkc0.0.0.8.0       DKC0          HSZ70 V71Z
dkc1.0.0.8.0       DKC1          HSZ70 V71Z
:
:
dkc7.0.0.8.0       DKC7          HSZ70 V71Z
dkd0.0.0.9.0       DKD0          HSZ40 YA03
dkd1.0.0.9.0       DKD1          HSZ40 YA03
dkd100.1.0.9.0     DKD100       HSZ40 YA03
dkd101.1.0.9.0     DKD101       HSZ40 YA03
dkd102.1.0.9.0     DKD102       HSZ40 YA03
:
:
dkd7.0.0.9.0       DKD7          HSZ40 YA03
dva0.0.0.0.0       DVA0
ewa0.0.0.2002.0     EWA0          00-06-2B-00-0A-48
fwa0.0.0.7.1        FWA0          08-00-2B-B9-0D-5D
pka0.7.0.2000.0     PKA0          SCSI Bus ID 7
pkb0.7.0.2001.0     PKB0          SCSI Bus ID 7
pkc0.7.0.8.0        PKC0          SCSI Bus ID 7 5.57
pkd0.7.0.9.0        PKD0          SCSI Bus ID 7 5.57
```

Example 4–3 shows the output from the `show config` console command entered on an AlphaServer 8200 system.

Example 4–3: Displaying Configuration on an AlphaServer 8200

```
>>> show config
      Name                Type    Rev    Mnemonic
-----
TLSB
4++   KN7CC-AB            8014   0000   kn7cc-ab0
5+    MS7CC                 5000   0000   ms7cc0
8+    KFTIA                  2020   0000   kftia0

C0 Internal PCI connected to kftia0 pci0
0+   QLogic ISP1020 10201077 0001   isp0
1+   QLogic ISP1020 10201077 0001   isp1
2+   DECchip 21040-AA 21011   0023   tulip0
4+   QLogic ISP1020 10201077 0001   isp2
5+   QLogic ISP1020 10201077 0001   isp3
6+   DECchip 21040-AA 21011   0023   tulipl

C1 PCI connected to kftia0
0+   KZPAA                  11000   0001   kzpaa0
1+   QLogic ISP1020 10201077 0005   isp4
2+   KZPSA                  81011   0000   kzpsa0
3+   KZPSA                  81011   0000   kzpsa1
4+   KZPSA                  81011   0000   kzpsa2
7+   DEC PCI MC             181011  000B   mc0
```

Example 4–4 shows the output from the `show device` console command entered on an AlphaServer 8200 system.

Example 4–4: Displaying Devices on an AlphaServer 8200

```
>>> show device
polling for units on isp0, slot0, bus0, hose0...
polling for units on isp1, slot1, bus0, hose0...
polling for units on isp2, slot4, bus0, hose0...
polling for units on isp3, slot5, bus0, hose0...
polling for units kzpaa0, slot0, bus0, hose1...
pke0.7.0.0.1      kzpaa4          SCSI Bus ID 7
dke0.0.0.0.1      DKE0            RZ28      442D
dke200.2.0.0.1    DKE200          RZ28      442D
dke400.4.0.0.1    DKE400          RRD43     0064

polling for units isp4, slot1, bus0, hose1...
dkf0.0.0.1.1      DKF0            HSZ70     V70Z
dkf1.0.0.1.1      DKF1            HSZ70     V70Z
dkf2.0.0.1.1      DKF2            HSZ70     V70Z
dkf3.0.0.1.1      DKF3            HSZ70     V70Z
```

Example 4–4: Displaying Devices on an AlphaServer 8200 (cont.)

```
dkf4.0.0.1.1      DKF4          HSZ70      V70Z
dkf5.0.0.1.1      DKF5          HSZ70      V70Z
dkf6.0.0.1.1      DKF6          HSZ70      V70Z
dkf100.1.0.1.1    DKF100        RZ28M      0568
dkf200.2.0.1.1    DKF200        RZ28M      0568
dkf300.3.0.1.1    DKF300        RZ28       442D

polling for units on kzpsa0, slot 2, bus 0, hose1...
kzpsa0.4.0.2.1    dkg          TPwr 1 Fast 1 Bus ID 7  L01 A11
dkg0.0.0.2.1      DKG0          HSZ50-AX X29Z
dkg1.0.0.2.1      DKG1          HSZ50-AX X29Z
dkg2.0.0.2.1      DKG2          HSZ50-AX X29Z
dkg100.1.0.2.1    DKG100        RZ26N      0568
dkg200.2.0.2.1    DKG200        RZ28       392A
dkg300.3.0.2.1    DKG300        RZ26N      0568

polling for units on kzpsa1, slot 3, bus 0, hose1...
kzpsa1.4.0.3.1    dkh          TPwr 1 Fast 1 Bus ID 7  L01 A11
dkh100.1.0.3.1    DKH100        RZ28       442D
dkh200.2.0.3.1    DKH200        RZ26       392A
dkh300.3.0.3.1    DKH300        RZ26L      442D

polling for units on kzpsa2, slot 4, bus 0, hose1...
kzpsa2.4.0.4.1    dki          TPwr 1 Fast 1 Bus ID 7  L01 A10
dki100.1.0.3.1    DKI100        RZ26       392A
dki200.2.0.3.1    DKI200        RZ28       442C
dki300.3.0.3.1    DKI300        RZ26       392A
```

4.3.3 Displaying Console Environment Variables and Setting the KZPBA-CB SCSI ID

The following sections show how to use the `show console` command to display the `pk*` and `isp*` console environment variables, and set the KZPBA-CB SCSI ID on various AlphaServer systems. Use these examples as guides for your system.

Console environment variables that are used for the SCSI options vary from system to system. Also, a class of environment variables (for example, `pk*` or `isp*`) may show both internal and external options.

Compare the following examples with the devices shown in the `show config` and `show dev` examples to determine which devices are KZPSA-BBs or KZPBA-CBs on the shared SCSI bus.

4.3.3.1 Displaying KZPBA-CB pk* or isp* Console Environment Variables

To determine the console environment variables to use, execute the `show pk*` and `show isp*` console commands.

Example 4–5 shows the `pk` console environment variables for an AlphaServer DS20.

Example 4–5: Displaying the pk* Console Environment Variables on an AlphaServer DS20 System

```
P00>>>show pk*
pka0_disconnect      1
pka0_fast            1
pka0_host_id         7

pkb0_disconnect      1
pkb0_fast            1
pkb0_host_id         7

pkc0_host_id         7
pkc0_soft_term       on

pkd0_host_id         7
pkd0_soft_term       on
```

Comparing the `show pk*` command display in Example 4–5 with the `show config` command in Example 4–1, you determine that the first two devices shown in Example 4–5, `pka0` and `pkb0` are for NCR 53C875 SCSI controllers. The next two devices, `pkc0` and `pkd0`, shown in Example 4–1 as Qlogic ISP1020 devices, are KZPBA-CBs, which are really Qlogic ISP1040 devices (regardless of what the console indicates).

Our interest then, is in `pkc0` and `pkd0`.

Example 4–5 shows two `pk*0_soft_term` environment variables, `pkc0_soft_term` and `pkd0_soft_term`, both of which are `on`.

The `pk*0_soft_term` environment variable applies to systems using the QLogic ISP1020 SCSI controller, which implements the 16-bit wide SCSI bus and uses dynamic termination.

The QLogic ISP1020 module has two terminators, one for the 8 low bits and one for the high 8 bits. There are five possible values for `pk*0_soft_term`:

- `off` — Turns off both low 8 bits and high 8 bits
- `low` — Turns on low 8 bits and turns off high 8 bits
- `high` — Turns on high 8 bits and turns off low 8 bits

- `on` — Turns on both low 8 bits and high 8 bits
- `diff` — Places the bus in differential mode

The KZPBA-CB is a Qlogic ISP1040 module, and its termination is determined by the presence or absence of internal termination resistor SIPs RM1-RM8. Therefore, the `pk*0_soft_term` environment variable has no meaning and it may be ignored.

Example 4–6 shows the use of the `show isp` console command to display the console environment variables for KZPBA-CBs on an AlphaServer 8x00.

Example 4–6: Displaying Console Variables for a KZPBA-CB on an AlphaServer 8x00 System

```
P00>>>show isp*
isp0_host_id      7
isp0_soft_term    on

isp1_host_id      7
isp1_soft_term    on

isp2_host_id      7
isp2_soft_term    on

isp3_host_id      7
isp3_soft_term    on

isp5_host_id      7
isp5_soft_term    diff
```

Both Example 4–3 and Example 4–4 show five `isp` devices; `isp0`, `isp1`, `isp2`, `isp3`, and `isp4`. In Example 4–6, the `show isp*` console command shows `isp0`, `isp1`, `isp2`, `isp3`, and `isp5`.

The console code that assigns console environment variables counts every I/O adapter including the KZPAA, which is the device after `isp3`, and therefore logically `isp4` in the numbering scheme. The `show isp` console command skips over `isp4` because the KZPAA is not a QLogic 1020/1040 class module.

Example 4–3 and Example 4–4 show that `isp0`, `isp1`, `isp2`, and `isp3` are devices on the internal KFTIA PCI bus and not on a shared SCSI bus. Only `isp4`, the KZPBA-CB, is on a shared SCSI bus (and the `show isp` console command displays it as `isp5`). The other three shared SCSI buses use KZPSA-BBs. (Use the `show pk*` console command to display the KZPSA console environment variables.)

4.3.3.2 Setting the KZPBA-CB SCSI ID

After you determine the console environment variables for the KZPBA-CBs on the shared SCSI bus, use the `set` console command to set the SCSI ID. For a TruCluster Server cluster, you will most likely have to set the SCSI ID for all KZPBA-CB UltraSCSI adapters except one. And, if you are using a DS-DWZZH-05, you will have to set the SCSI IDs for all KZPBA-CB UltraSCSI adapters.

Cautions

You will have problems accessing storage if you have two or more SCSI adapters at the same SCSI ID on any one SCSI bus.

If you are using a DS-DWZZH-05, you cannot use SCSI ID 7 for a KZPBA-CB UltraSCSI adapter; SCSI ID 7 is reserved for DS-DWZZH-05 use.

If DS-DWZZH-05 fair arbitration is enabled, the SCSI ID of the host adapter must match the SCSI ID assigned to the hub port. Mismatching or duplicating SCSI IDs will cause the hub to hang.

SCSI ID 7 is reserved for the DS-DWZZH-05 whether fair arbitration is enabled or not.

Use the `set` console command as shown in Example 4-7 to set the SCSI ID. In this example, the SCSI ID is set for KZPBA-CB `pkc` on the AlphaServer DS20 shown in Example 4-5.


Example 4-7: Setting the KZPBA-CB SCSI Bus ID

```
P00>>> show pkc0_host_id
7
P00>>> set pkc0_host_id 6
P00>>> show pkc0_host_id
6
```

4.3.3.3 KZPBA-CB Termination Resistors

The KZPBA-CB internal termination is disabled by removing the termination resistors RM1-RM8, as shown in Figure 4-1.

Figure 4-1: KZPBA-CB Termination Resistors



ZK-1451U-AI

5

Setting Up the Memory Channel Cluster Interconnect

This chapter describes Memory Channel configuration restrictions, and describes how to set up the Memory Channel cluster interconnect, including setting up a Memory Channel hub and Memory Channel optical converter (MC2 only), and connecting link cables.

Two versions of the Memory Channel peripheral component interconnect (PCI) adapter are available: CCMAA and CCMAB (MC2).

Two variations of the CCMAA PCI adapter are in use: CCMAA-AA (MC1) and CCMAA-AB (MC1.5). Because the hardware used with these two PCI adapters is the same, this manual often refers to MC1 when referring to either of these variations.

See the *TruCluster Server Software Product Description (SPD)* for a list of the supported Memory Channel hardware. See the *Memory Channel User's Guide* for illustrations and more detailed information about installing jumpers, Memory Channel adapters, and hubs.

See Section 2.2 for a discussion on Memory Channel restrictions.

You can have two Memory Channel adapters with TruCluster Server, but only one rail is active at a time. This is referred to as a failover pair. If the active rail fails, cluster communications fails over to the formerly inactive rail.

If you use multiple Memory Channel adapters with the Memory Channel application programming interface (API) for high performance data delivery over Memory Channel, setting the `rm_rail_style` configuration variable to zero (`rm_rail_style = 0`) enables single-rail style with multiple active rails. The default is one, which selects failover pair.

For more information on the Memory Channel failover pair model, see the *Cluster Highly Available Applications* manual.

To set up the Memory Channel interconnects, follow these steps, referring to the appropriate section and the *Memory Channel User's Guide* as necessary:

1. Set the Memory Channel jumpers (Section 5.1).

2. Install the Memory Channel adapter into a PCI slot on each system (Section 5.2).
3. If you are using fiber optics with MC2, install the CCMFB fiber-optic module (Section 5.3).
4. If you have more than two systems in the cluster, install a Memory Channel hub (Section 5.4).
5. Connect the Memory Channel cables (Section 5.5).
6. After you complete steps 1 through 5 for all systems in the cluster, apply power to the systems and run Memory Channel diagnostics (Section 5.6).

Note

If you are installing SCSI or network adapters, you may want to complete all hardware installation before powering up the systems to run Memory Channel diagnostics.

Section 5.7.2 provides procedures for upgrading from redundant MC1 interconnects to MC2 interconnects.

5.1 Setting the Memory Channel Adapter Jumpers

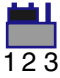


The meaning of the Memory Channel adapter module jumpers depends upon the version of the Memory Channel module.

5.1.1 MC1 and MC1.5 Hub Mode Jumper

The MC1 and MC1.5 modules (CCMAA-AA and CCMAA-AB, respectively) have an adapter jumper (J4) that designates whether the configuration is using standard or virtual hub mode. If virtual hub mode is being used, there can be only two systems. One system must be virtual hub 0 (VH0) and the other must be virtual hub 1 (VH1).

The Memory Channel adapter should arrive with the J4 jumper set for standard hub mode (pins 1 to 2 jumpered). Confirm that the jumper is set properly for your configuration. The jumper configurations in Table 5-1 are shown as if you are holding the module with the J4 jumper facing you, with the module end plate in your left hand. The jumper is next to the factory/maintenance cable connector.

Table 5–1: MC1 and MC1.5 J4 Jumper Configuration

If hub mode is:	Jumper:	Example:
Standard	J4 Pins 1 to 2	 1 2 3
Virtual: VH0	J4 Pins 2 to 3	 1 2 3
Virtual: VH1	None needed; store the jumper on J4 pin 1 or 3	 1 2 3

If you are upgrading from virtual hub mode to standard hub mode (or from standard hub mode to virtual hub mode), be sure to change the J4 jumper on all Memory Channel adapters on the rail.

5.1.2 MC2 Jumpers

The MC2 module (CCMAB) has multiple jumpers. They are numbered right to left, starting with J1 in the upper right corner (as you view the jumper side of the module with the endplate in your left hand). The leftmost jumpers are J11 and J10. J11 is above J10.

Most of the jumper settings are straightforward, but the window size jumper, J3, needs some explanation.

If a CCMAA adapter (MC1 or MC1.5) is installed, 128 MB of address space is allocated for Memory Channel use. If a CCMAB adapter (MC2) PCI adapter is installed, the memory space allocation for Memory Channel depends on the J3 jumper and can be 128 MB or 512 MB.

If two Memory Channel adapters are used as a failover pair to provide redundancy, the address space allocated for the logical rail depends on the smaller window size of the physical adapters.

During a rolling upgrade (see Section 5.7.2) from an MC1 failover pair to an MC2 failover pair, the MC2 modules can be jumpered for 128 MB or 512 MB. If jumpered for 512 MB, the increased address space is not achieved until all MC PCI adapters have been upgraded and the use of 512 MB is enabled. On one member system, use the `sysconfig` command to reconfigure the Memory Channel kernel subsystem to initiate the use of 512 MB address

space. The configuration change is propagated to the other cluster member systems by entering the following command:

```
# /sbin/sysconfig -r rm rm_use_512=1
```

See the *Cluster Administration* manual for more information on failover pairs.

The MC2 jumpers are described in Table 5–2.

Table 5–2: MC2 Jumper Configuration












Jumper:	Description:	Example:
J1: Hub Mode	Standard: Pins 1 to 2	
	VH0: Pins 2 to 3	
	VH1: None needed; store the jumper on pin 1 or pin 3	
J3: Window Size	512 MB: Pins 2 to 3	
	128 MB: Pins 1 to 2	
J4: Page Size	8-KB page size (UNIX): Pins 1 to 2	
	4-KB page size (not used): Pins 2 to 3	



Table 5–2: MC2 Jumper Configuration (cont.)

Jumper:	Description:	Example:
J5: AlphaServer 8x00 Mode	8x00 mode selected: Pins 1 to 2 ^a	 1 2 3
	8x00 mode not selected: Pins 2 to 3	 1 2 3
J10 and J11: Fiber-Optic Mode Enable	Fiber Off: Pins 1 to 2	 3 2 1
	Fiber On: Pins 2 to 3 pins	 3 2 1

^a Increases the maximum sustainable bandwidth for 8x00 systems. If the jumpers are in this position for other systems, the bandwidth is decreased.

The MC2 linecard (CCMLB) has two jumpers, J2 and J3, that are used to enable fiber-optic mode. The jumpers are located near the middle of the module (as you view the jumper side of the module with the endplate in your left hand). Jumper J2 is on the right. The MC2 linecard jumpers are described in Table 5–3.

Table 5–3: MC2 Linecard Jumper Configurations

Jumper:	Description:	Example:
J2 and J3: Fiber Mode	Fiber Off: Pins 2 to 3	 1 2 3
	Fiber On: Pins 1 to 2	 1 2 3

5.2 Installing the Memory Channel Adapter

Install the Memory Channel adapter in an appropriate peripheral component interconnect (PCI) slot. (See Section 2.2.) Secure the module at the backplane. Ensure that the screw is tight to maintain proper grounding.

The Memory Channel adapter comes with a straight extension plate. This fits most systems; however, you may have to replace the extender with an angled extender (AlphaServer 2100A, for instance), or for an AlphaServer 8200/8400, GS60, GS60E, or GS140, remove the extender completely.

If you are setting up a redundant Memory Channel configuration, install the second Memory Channel adapter immediately after installing the first Memory Channel adapter. Ensure that the jumpers are correct and are the same on both modules.

After you install the Memory Channel adapters, replace the system panels, unless you have more hardware to install.

5.3 Installing the MC2 Optical Converter in the Member System

If you plan to use a CCMFB optical converter along with the MC2 PCI adapter, install it at the same time that you install the MC2 CCMAB. To install a MC2 CCMFB optical converter in the member system, follow these steps. See Section 5.5.2.4 if you are installing an optical converter in an MC2 hub.

1. Remove the bulkhead blanking plate for the desired PCI slot.
2. Thread one end of the fiber-optic cable (BN34R) through the PCI bulkhead slot.
3. Thread the cable through the slot in the optical converter module (CCMFB) endplate (at the top of the endplate).
4. Remove the cable tip protectors and attach the keyed plug to the connector on the optical converter module. Tie-wrap the cable to the module.
5. Seat the optical converter module firmly into the PCI backplane and secure the module with the PCI card cage mounting screw.
6. Attach the 1-meter (3.3-foot) BN39B-01 cable from the CCMAB MC2 PCI adapter to the CCMFB optical converter.
7. Route the fiber-optic cable to the remote system or hub.
8. Repeat steps 1 through 7 for the optical converter on the second system. See Section 5.5.2.4 if you are installing an optical converter in an MC2 hub.

5.4 Installing the Memory Channel Hub

You may use a hub in a two-node TruCluster Server cluster, but the hub is not required. When there are more than two systems in a cluster, you must use a Memory Channel hub as follows:

- For use with the MC1 or MC1.5 CCMAA adapter, you must install the hub within 3 meters (9.8 feet) of each of the systems.

For use with the MC2 CCMAB adapter, the hub must be placed within 4 meters (13.1 feet) or 10 meters (32.8 feet) (the length of the BN39B link cables) of each system. If fiber optics is used in conjunction with the MC2 adapter, the hub may be placed up to 3000 meters (9842.5 feet) from the systems.

- Ensure that the voltage selection switch on the back of the hub is set to select the correct voltage for your location (115V or 230V).
- Ensure that the hub contains a linecard for each system in the cluster (the hub comes with four linecards) as follows:
 - CCMLA linecards for the CCMHA MC1 hub
 - CCMLB linecards for the CCMHB MC2 hub. The linecards cannot be installed in the `opto only` slot.
- If you have a four-node cluster, you may want to install an extra linecard for troubleshooting use.
- If you have an eight-node cluster, all linecards must be installed in the same hub.
- For MC2, if fiber-optic converters are used, they can only be installed in hub slots `opto only`, `0/opto`, `1/opto`, `2/opto`, and `3/opto`.
- If you have a five-node or greater MC2 cluster using fiber optics, you will need two or three CCMHB hubs, depending on the number of fiber-optic connections. You will need one hub for the CCMLB linecards (and possible optics converters) and up to two hubs for the CCMFB optic converter modules. The CCMHB-BA hub has no linecards.

5.5 Installing the Memory Channel Cables

Memory Channel cable installation depends on the Memory Channel module revision, and whether or not you are using fiber optics. The following sections describe how to install the Memory Channel cables for MC1 and MC2.

5.5.1 Installing the MC1 or MC1.5 Cables

To set up an MC1 or MC1.5 interconnect, use the BC12N-10 3-meter (9.8-foot) link cables to connect Memory Channel adapters and, optionally, Memory Channel hubs.

Note

Do not connect an MC1 or MC1.5 link cable to an MC2 module.

5.5.1.1 Connecting MC1 or MC1.5 Link Cables in Virtual Hub Mode

For an MC1 virtual hub configuration (two nodes in the cluster), connect the BC12N-10 link cables between the Memory Channel adapters that are installed in each of the systems.

Caution

Be very careful when installing the link cables. Insert the cables straight in.

Gently push the cable's connector into the receptacle, and then use the screws to pull the connector in tight. The connector must be tight to ensure a good ground contact.

If you are setting up redundant interconnects, all Memory Channel adapters in a system must have the same jumper setting, either VH0 or VH1.

Note

With the TruCluster Server Version 5.1A product and virtual hub mode, there is no longer a restriction requiring that `mca0` in one system be connected to `mca0` in the other system.

5.5.1.2 Connecting MC1 Link Cables in Standard Hub Mode


If there are more than two systems in a cluster, use a standard hub configuration. Connect a BC12N-10 link cable between the Memory Channel adapter and a linecard in the CCMHA hub, starting at the lowest numbered slot in the hub.

If you are setting up redundant interconnects, the following restrictions apply:

- Each adapter installed in a system must be connected to a different hub.
- Each Memory Channel adapter in a system must be connected to linecards that are installed in the same slot position in each hub. For example, if you connect one adapter to a linecard installed in slot 1 in one hub, you must connect the other adapter in that system to a linecard installed in slot 1 of the second hub.

Figure 5–1 shows Memory Channel adapters connected to linecards that are in the same slot position in the Memory Channel hubs.

Figure 5–1: Connecting Memory Channel Adapters to Hubs



ZK-1197U-AI

5.5.2 Installing the MC2 Cables

To set up an MC2 interconnect, use the BN39B-04 (4-meter; 13.1-foot) or BN39B-10 (10-meter; 32.8-foot) link cables for virtual hub or standard hub configurations without optical converters.

If optical converters are used, use the BN39B-01 (1-meter; 3.3-foot) link cable and the BN34R-10 (10-meter; 32.8-foot) or BN34R-31 (31-meter; 101.7-foot) fiber-optic cable.

5.5.2.1 Installing the MC2 Cables for Virtual Hub Mode Without Optical Converters

To set up an MC2 configuration for virtual hub mode, use BN39B-04 (4-meter; 13.1-foot) or BN39B-10 (10-meter; 32.8-foot) Memory Channel link cables to connect Memory Channel adapters to each other.

Notes

MC2 link cables (BN39B) are black cables.

Do not connect an MC2 cable to an MC1 or MC1.5 CCMAA module.

Gently push the cable's connector into the receptacle, and then use the screws to pull the connector in tight. The connector must be tight to ensure a good ground contact.

If you are setting up redundant interconnects, all Memory Channel adapters in a system must have the same jumper setting, either VH0 or VH1.

5.5.2.2 Installing MC2 Cables in Virtual Hub Mode Using Optical Converters

If you are using optical converters in an MC2 configuration, install an optical converter module (CCMFB) when you install the CCMAB Memory Channel PCI adapter in each system in the virtual hub configuration. Also connect the CCMAB Memory Channel adapter to the optical converter with a BN39B-01 cable. When you install the CCMFB optical converter module in the second system, you connect the two systems with the BN34R fiber-optic cable. Customer-supplied cables may be up to 2 kilometers (1.24 miles) in length. (See Section 5.3.)

5.5.2.3 Connecting MC2 Link Cables in Standard Hub Mode (No Fiber Optics)

If there are more than two systems in a cluster, use a Memory Channel standard hub configuration. Connect a BN39B-04 (4-meter; 13.1-foot) or BN39B-10 (10-meter; 32.8-foot) link cable between the Memory Channel adapter and a linecard in the CCMHB hub, starting at the lowest numbered slot in the hub.

If you are setting up redundant interconnects, the following restrictions apply:

- Each adapter installed in a system must be connected to a different hub.
- Each Memory Channel adapter in a system must be connected to linecards that are installed in the same slot position in each hub. For example, if you connect one adapter to a linecard installed in slot 0/opto in one hub, you must connect the other adapter in that system to a linecard installed in slot 0/opto of the second hub.

Note

You cannot install a CCMLB linecard in slot opto only.

5.5.2.4 Connecting MC2 Cables in Standard Hub Mode Using Optical Converters

If you are using optical converters in an MC2 configuration, install an optical converter module (CCMFB), with attached BN34R fiber-optic cable, when you install the CCMAB Memory Channel PCI adapter in each system in the standard hub configuration. Also connect the CCMAB Memory Channel adapter to the optical converter with a BN39B-01 cable.

Note

See Section 2.2 for restrictions on the lengths of Memory Channel fiber-optic cables.

Now you need to:

- Set the CCMLB linecard jumpers to support fiber optics
- Connect the fiber-optic cable to a CCMFB fiber-optic converter module
- Install the CCMFB fiber-optic converter module for each fiber-optic link

Note

If you have more than four fiber-optic links, you need two or more hubs. The CCMHB-BA hub has no linecards.

To set the CCMLB jumpers and install CCMFB fiber-optic converter modules in an MC2 hub, follow these steps:

1. Remove the appropriate CCMLB linecard and set the linecard jumpers to *Fiber On* (jumper pins 1 to 2) to support fiber optics. See Table 5–3.
2. Remove the CCMLB endplate and install the alternate endplate (with the slot at the bottom).
3. Remove the hub bulkhead blanking plate from the appropriate hub slot. Ensure that you observe the slot restrictions for the optical converter modules. Also keep in mind that all linecards for one Memory Channel interconnect must be in the same hub. (See Section 5.4.)
4. Thread the BN34R fiber-optic cable through the hub bulkhead slot. Make sure that the other end is attached to a CCMFB optics converter in the member system.
5. Thread the BN34R fiber-optic cable through the slot near the bottom of the endplate. Remove the cable tip protectors and insert the connectors into the transceiver until they click into place. Secure the cable to the module using the tie-wrap.

6. Install the CCMFB fiber-optic converter in slot `opto` only, `0/opto`, `1/opto`, `2/opto`, or `3/opto`, as appropriate.
7. Install a BN39B-01 1-meter (3.3-foot) link cable between the CCMFB optical converter and the CCMLB linecard.
8. Repeat steps 1 through 7 for each CCMFB module to be installed.

5.6 Running Memory Channel Diagnostics

After the Memory Channel adapters, hubs, link cables, fiber-optic converters, and fiber-optic cables have been installed, power up the systems and run the Memory Channel diagnostics.

There are two console level Memory Channel diagnostics, `mc_diag` and `mc_cable`:

- The `mc_diag` diagnostic:
 - Tests the Memory Channel adapters on the system running the diagnostic.
 - Runs as part of the initialization sequence when the system is powered up.
 - Runs on a standalone system or while connected to another system or a hub with the link cable.
- The `mc_cable` diagnostic:
 - Must be run on all systems in the cluster simultaneously (therefore, all systems must be at the console prompt).

Caution

If you attempt to run `mc_cable` on one cluster member while other members of the cluster are up, you may crash the cluster.

- Is designed to isolate problems to the Memory Channel adapter, BC12N or BN39B link cables, hub linecards, fiber-optic converters, BN34R fiber-optic cable, and, to some extent, to the hub.
- Indicates data flow through the Memory Channel by response messages.
- Runs continuously until terminated with Ctrl/C.
- Reports differences in connection state, not errors.
- Can be run in standard or virtual hub mode.

When the console indicates a successful response from all other systems being tested, the data flow through the Memory Channel hardware has been completed and the test may be terminated by pressing Ctrl/C on each system being tested.

Example 5–1 shows a sample output from node 1 of a standard hub configuration. In this example, the test is started on node 1, then on node 0. The test must be terminated on each system.

Example 5–1: Running the mc_cable Test

```
>>> mc_cable 1
To exit MC_CABLE, type <Ctrl/C>
mca0 node id 1 is online 2
No response from node 0 on mca0 2
mcb0 node id 1 is online 3
No response from node 0 on mcb0 3
Response from node 0 on mca0 4
Response from node 0 on mcb0 5
mcb0 is offline 6
mca0 is offline 6
Ctrl/C 7
>>>
```

- 1 The mc_cable diagnostic is initiated on node 1.
- 2 Node 1 reports that mca0 is on line but has not communicated with the Memory Channel adapter on node 0.
- 3 Node 1 reports that mcb0 is on line but has not communicated with the Memory Channel adapter on node 0.
- 4 Memory Channel adapter mca0 has communicated with the adapter on the other node.
- 5 Memory Channel adapter mcb0 has communicated with the adapter on the other node.
- 6 Typing a Ctrl/C on node 0 terminates the test on that node and the Memory Channel adapters on node 1 report off line.
- 7 Ctrl/C on node 1 terminates the test.

5.7 Maintaining Memory Channel Interconnects

The following sections contain information about maintaining Memory Channel interconnects. See other sections in this chapter or the Memory Channel *User's Guide* for detailed information about maintaining the Memory Channel hardware. Topics in this section include:

- Adding a Memory Channel interconnect (Section 5.7.1)
- Upgrading Memory Channel adapters (Section 5.7.2)
- Upgrading a virtual hub configuration to a standard hub configuration (Section 5.7.3)

5.7.1 Adding a Memory Channel Interconnect

If you want to change from a single Memory Channel interconnect to redundant Memory Channel interconnects without shutting down the cluster, follow the steps in Table 5–4, which covers adding a Memory Channel interconnect and rolling from a dual MC1 interconnect to a dual MC2 interconnect. Most of the steps are the same.

5.7.2 Upgrading Memory Channel Adapters

If you have a TruCluster Server configuration with redundant MC1 interconnects and want to upgrade to MC2 interconnects, you can do so without shutting down the entire cluster.

When performing an upgrade from MC1 interconnects, which use 128 MB Memory Channel address space, to MC2, which uses either 128 or 512 MB Memory Channel address space, all Memory Channel adapters must be operating at 128 MB Memory Channel address space (the default) until the last adapter has been changed. At that time the address space can be increased to 512 MB if all MC2 adapters are jumpered for 512 MB.

This section covers adding a Memory Channel interconnect and the following rolling upgrade situations:

- Dual, redundant MC1 interconnects in virtual hub mode (Table 5–4 and Figure 5–2)
- Dual, redundant MC1 interconnects in standard hub mode (Table 5–4 and Figure 5–3 through Figure 5–8)

The figures following Table 5–4 provide two sequences that you can follow while carrying out the steps of Table 5–4. Figure 5–2 shows a dual, redundant virtual hub configuration using MC1 hardware being upgraded to MC2. Figure 5–3 through Figure 5–8 show a three-node standard hub configuration being upgraded from MC1 to MC2.

Note

When you upgrade from dual, redundant MC1 hardware to dual, redundant MC2 hardware, you must replace all the MC1 hardware on one interconnect before you start on the second interconnect (except as described in step 4 of Table 5–4).

Memory Channel adapters jumpered for 512 MB may require a minimum of 512 MB physical RAM memory. Ensure that your system has enough physical memory to support the upgrade. For two MC2 Memory Channel adapters, you will need more than 1 GB of physical memory.

Table 5–4: Adding a Memory Channel Interconnect or Upgrading from a Dual, Redundant MC1 Interconnect to MC2 Interconnects

Step	Action	Refer to:
1	If desired, using the cluster application availability (CAA) <code>caa_relocate</code> command, manually relocate all applications from the cluster member that will be shut down.	TruCluster Server <i>Cluster Administration</i>
2	On the system having an MC1 adapter installed or replaced, log in as the root user and execute the <code>shutdown -h</code> utility to halt the system.	Tru64 UNIX <i>System Administration</i>
<hr/> <p style="text-align: center;">Note</p> <p>After the system is at the console prompt, use the console <code>set</code> command to set the <code>auto_action</code> console environment variable to <code>halt</code>. This halts the system at the console prompt when the system is turned on, ensuring that you are able to run the Memory Channel diagnostics.</p> <pre>>>> set auto_action halt</pre>		
3	Turn off the system.	—
4	Set the jumpers on the new Memory Channel module to be installed:	Section 5.1 and <i>Memory Channel User's Guide</i>

MC1:

Hub mode — Standard hub mode or virtual hub mode (VH0 or VH1)

- Virtual hub mode, VH0: Jumper pins 2 to 3
- Virtual hub mode, VH1: No jumper
- Standard hub mode: Jumper pins 1 to 2

Table 5–4: Adding a Memory Channel Interconnect or Upgrading from a Dual, Redundant MC1 Interconnect to MC2 Interconnects (cont.)

Step	Action	Refer to:
	<p>MC2:</p> <p>Hub mode — Standard hub mode or virtual hub mode (VH0 or VH1)</p> <ul style="list-style-type: none"> • Virtual hub mode, VH0: Jumper pins 2 to 3 • Virtual hub mode, VH1: No jumper • Standard hub mode: Jumper pins 1 to 2 <p>J3 — Memory Channel address space: Select 128 MB (jumper pins 1 to 2) or 512 MB (jumper pins 2 to 3) as required for your configuration</p> <hr/> <p style="text-align: center;">Note</p> <p>If you set the J3 jumpers for 128 MB because the other interconnect is MC1, and then later on decide to upgrade to dual, redundant MC2 hardware using 512 MB address space, you will have to reset the jumpers. If you set the jumpers to 512 MB now, the software will only allow the use of 128 MB address space for a mixed rail cluster (MC1 on one rail, MC2 on the other rail).</p> <hr/> <p>J4 — Page size: Jumper pins 1 to 3 to select 8 KB</p> <p>J5 — AlphaServer 8x00 Mode: Jumper pins 1 to 2 for AlphaServer 8200, 8400, GS60, GS60E, and GS140 systems and jumper pins 2 to 3 for all other AlphaServer systems</p> <p>J10 — Fiber Optics Mode Enable: Jumper pins 2 to 3 to enable the use of the fiber-optic modules. Jumper pins 1 to 2 to disable the use of fiber optics</p>	
5	<p>If adding a Memory Channel interconnect: Install the Memory Channel adapter module.</p> <p>If this is the second system in a virtual hub configuration, connect an MC1 or MC2 link cable between the MC1 or MC2 modules.</p> <p>For a standard hub configuration, use a link cable to connect the adapter to the Memory Channel hub linecard in the hub slot that corresponds to the existing Memory Channel linecard in the other hub.</p> <p>If upgrading from a dual, redundant MC1 interconnect to MC2 interconnects: Remove the MC1 adapter and install the MC2 adapter.</p>	Section 5.2 and <i>Memory Channel User's Guide</i>

Table 5–4: Adding a Memory Channel Interconnect or Upgrading from a Dual, Redundant MC1 Interconnect to MC2 Interconnects (cont.)

Step	Action	Refer to:
Virtual Hub:		
	If this is the first system in a virtual hub configuration, replace the MC1 adapter with an MC2 adapter.	Figure 5–2 (B)
	If this is the second system in a virtual hub configuration, replace both MC1 adapters with MC2 adapters. Use a BN39B-10 link cable to connect Memory Channel adapters between systems to form the first MC2 interconnect.	Figure 5–2 (C)
	If this is the second adapter on the first system in a virtual hub configuration, replace the MC1 adapter with an MC2 adapter. Connect the second set of MC2 adapters with a BN39B-10 link cable to form the second Memory Channel interconnect.	Figure 5–2 (D)
Standard Hub Configuration:		
	Remove the MC1 adapter and install the MC2 adapter in one system, and on one rail at a time. Use a BN39B-10 link cable to connect the new MC2 adapter to the linecard in the MC2 hub that corresponds to the same linecard that the MC1 module was connected to in the MC1 hub.	Figure 5–4 and Figure 5–5
	If this is the last system on this rail to receive an MC2 adapter (that is, all other member systems on this rail have one MC2 adapter) you can replace both MC1 adapters at the same time. Use a BN39B-10 link cable to connect the new MC2 adapters to the linecard in their respective MC2 hub that corresponds to the same linecard that the MC1 modules were connected to in the MC1 hubs.	Figure 5–6
6	Turn on the system and run the <code>mc_diag</code> Memory Channel diagnostic. Note that you cannot run <code>mc_cable</code> because this is the only system in the cluster that is shut down.	Section 5.6
7	Boot the system.	
8	Repeat steps 1 - 7 for all other systems in the cluster. When you have replaced both MC1 adapters in the last system, repeat steps 1 - 7 and replace the MC1 adapters on the other interconnect.	Figure 5–7 and Figure 5–8
9	If desired, enable increasing the address space to 512 MB after the following conditions have been met:	<code>sysconfig</code> reference pages

Table 5–4: Adding a Memory Channel Interconnect or Upgrading from a Dual, Redundant MC1 Interconnect to MC2 Interconnects (cont.)

Step	Action	Refer to:
	<ul style="list-style-type: none"> The last member system has had its second MC1 adapter replaced with an MC2 adapter. The cluster is operational. All MC2 adapters are jumpered for 512 MB (and you need to utilize 512 MB address space). 	

On one member system, use the `sysconfig` command to reconfigure the Memory Channel kernel subsystem to initiate the use of 512 MB address space. The configuration change is propagated to the other cluster member systems:
`/sbin/sysconfig -r rm rm_use_512=1`

Note

After the configuration change is propagated to the other member systems, you can reboot any member system and the 512 MB address space is still in effect.

If you use the `sysconfig` command to promote address space to 512 MB and inadvertently leave an MC2 adapter jumpered for 128 MB, then reboot that system, it will not rejoin the cluster. When the system with the Memory Channel adapter jumpered for 128 MB is shut down, and the TruCluster software running on the remaining cluster member systems discover that all operational Memory Channel adapters are jumpered for 512 MB, because address space has been promoted to 512 MB, the active rail will use 512 MB address space. A system jumpered for 128 MB cannot join the cluster. The startup error message on the system jumpered for 128 MB follows:
`panic: MC2 adapter has too little memory`

If you have used the `sysconfig` command to promote Memory Channel address space to 512 MB, you may need to know the actual address space being used by a logical rail. Use the `dbx` debugger utility as follows to determine:

- Logical size (in 8-KB pages) of a rail
- Physical size (J3 jumper setting) for physical rails

```
# dbx -k /vmunix
(dbx) p rm_log_rail_to_ctx[0]->mgmt_page_va->size 1
16384 2
(dbx) p rm_adapters[0]->rmp_prail_va->rmc_size 3
{
    [0] 65536 4
    [1] 0
```

```

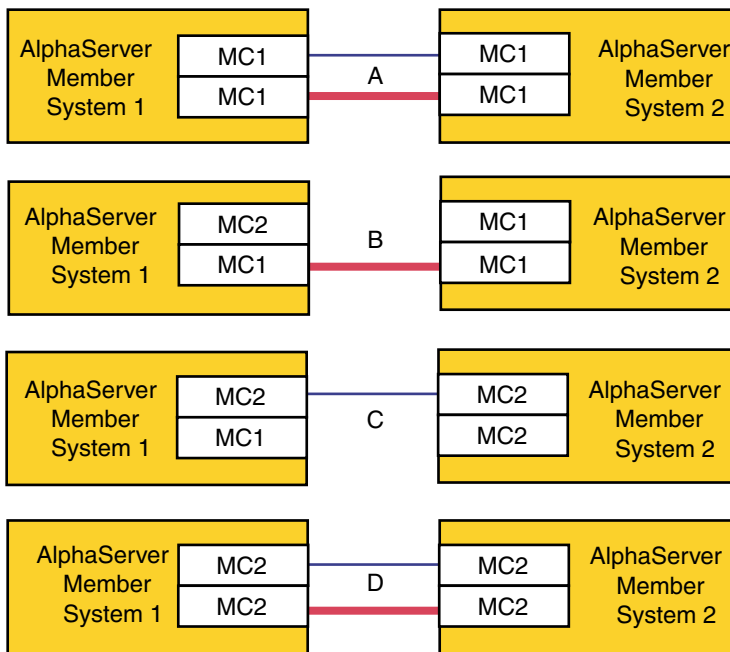
    [2] 65536 [4]
    [3] 0
    [4] 65536 [4]
    [5] 0
    [6] 0
    [7] 0
}
(dbx) p rm_adapters[1]->rmp_prail_va->rnc_size [5]
{
    [0] 16384 [6]
    [1] 0
    [2] 16384 [6]
    [3] 0
    [4] 16384 [6]
    [5] 0
    [6] 0
    [7] 0
}

```

- 1 Find the size of a logical rail.
- 2 The logical rail is operating at 128 MB (16384 eight-KB pages).
- 3 Verify the jumper settings for the member systems on the first physical rail.
- 4 The J3 jumper is set at 512 MB for nodes 0, 2, and 4 on the first physical rail (65536 eight-KB pages).
- 5 Verify the jumper settings for the member systems on the second physical rail.
- 6 The J3 jumper is set at 128 MB for nodes 0, 2, and 4 on the second physical rail (16384 eight-KB pages).

Figure 5–2 shows a dual, redundant virtual hub configuration using MC1 hardware being upgraded to MC2.

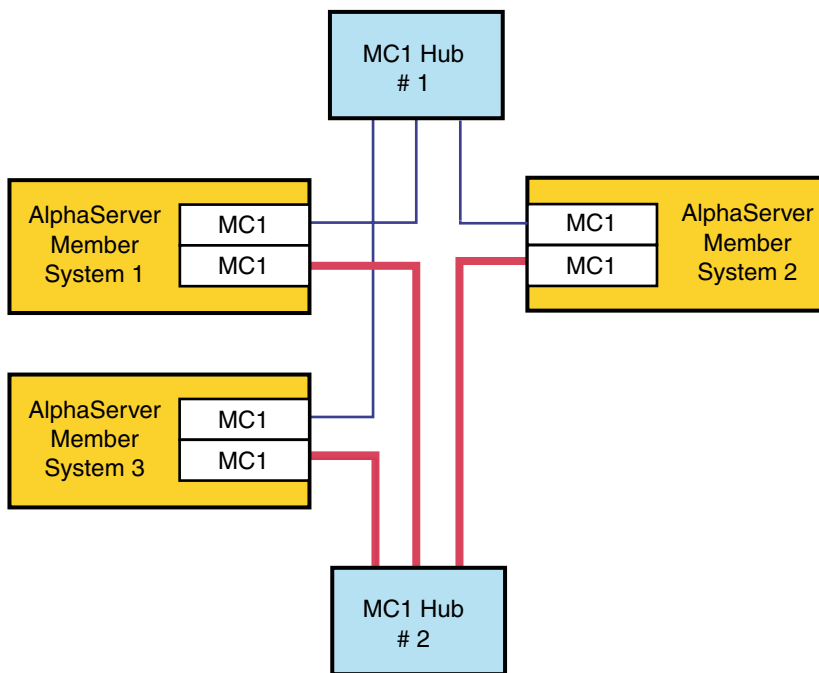
Figure 5–2: MC1-to-MC2 Virtual Hub Upgrade



ZK-1530U-AI

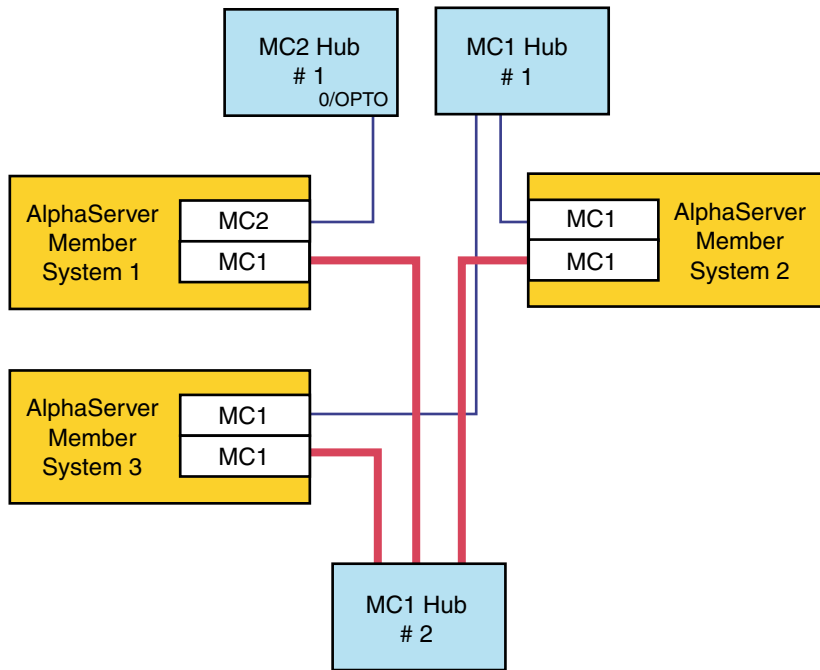
Figure 5-3 through Figure 5-8 show a three-node standard hub configuration being upgraded from MC1 to MC2.

Figure 5-3: MC1-to-MC2 Standard Hub Upgrade: Initial Configuration



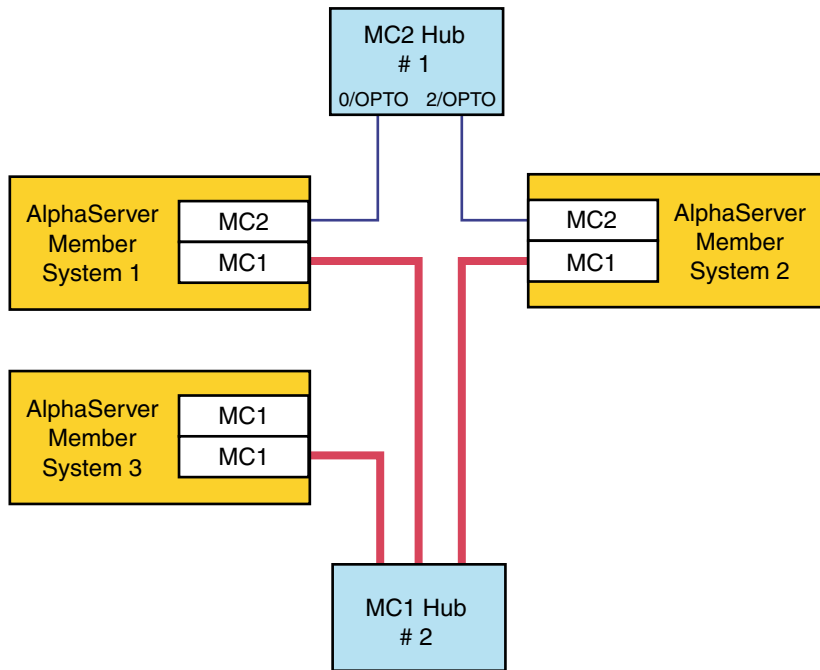
ZK-1522U-AI

Figure 5-4: MC1-to-MC2 Standard Hub Upgrade: First MC1 Module Replaced



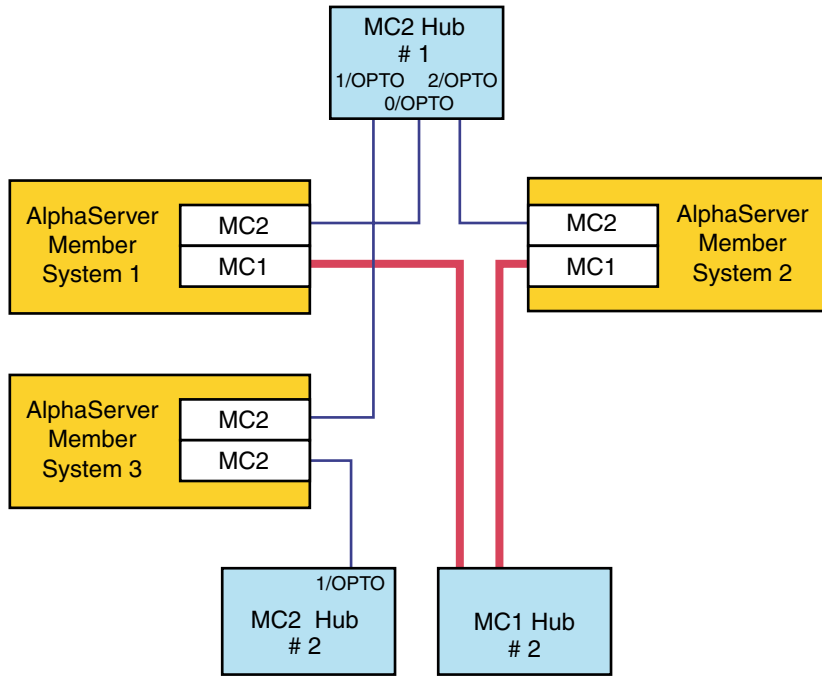
ZK-1523U-AI

Figure 5-5: MC1-to-MC2 Standard Hub Upgrade: Replace First MC1 Adapter in Second System



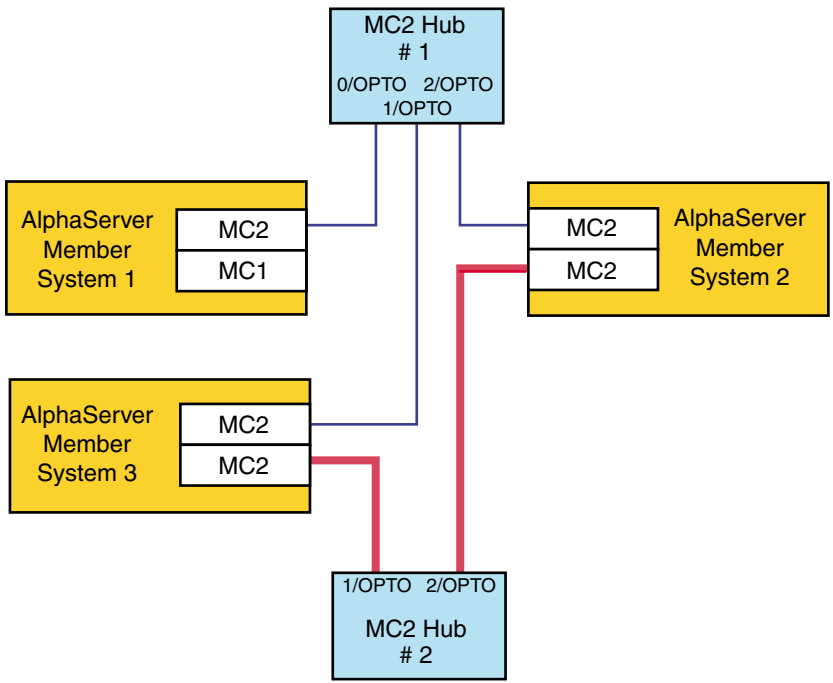
ZK-1524U-AI

Figure 5–6: MC1-to-MC2 Standard Hub Upgrade: Replace Third System Memory Channel Adapters



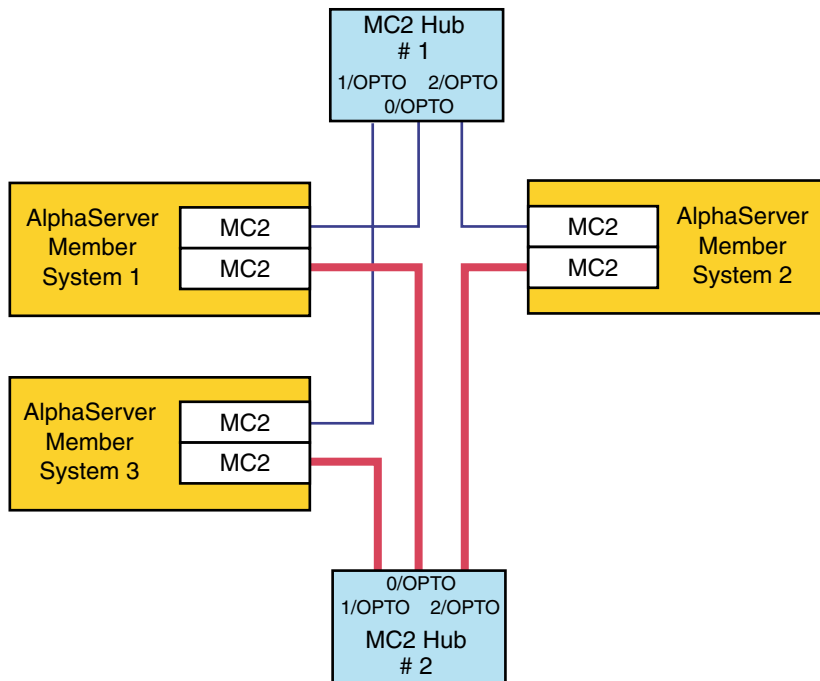
ZK-1525U-AI

Figure 5–7: MC1-to-MC2 Standard Hub Upgrade: Replace Second MC1 in Second System



ZK-1526U-AI

Figure 5–8: MC1-to-MC2 Standard Hub Upgrade: Final Configuration



ZK-1527U-AI

5.7.3 Upgrading a Virtual Hub Configuration to a Standard Hub Configuration

If your cluster is configured in virtual hub mode (two member systems with no Memory Channel hub), you must convert to standard hub mode in order to:

- Add another member system to the cluster.
- Add fiber optics to MC2 to provide more distance between the cluster systems.

Note

You need an additional PCI slot for each optical converter module to be installed in the system. The optical converter does not use PCI bandwidth, but it does take up a PCI slot.

You also need an available slot in the Memory Channel hub for an optical converter module for each member system.

There will be some cluster down time. During the procedure, you can maintain cluster operations except for the time it takes to shut down the second system and boot the first system as a single-node cluster.

Note

If you are not using a quorum disk, the first member you shut down must have zero votes for the cluster to survive its shutdown. Use the `clu_quorum` command to adjust quorum votes. See the `clu_quorum(8)` reference page and the *Cluster Administration* manual for more information.

To upgrade from a virtual hub configuration to a standard hub configuration, follow the steps in Table 5–5. In this procedure, `system1` is the member system that will be shut down first. Member system `system2` will be shut down last. The procedure is written with the assumption that you have dual-rail failover-pair Memory Channel adapter modules.

Table 5–5: Upgrading from a Virtual Hub Configuration to a Standard Hub Configuration

Step	Action	Refer to:
1	Install the Memory Channel hubs at an appropriate distance from the member systems.	Section 5.4
	If you are adding fiber optics, for each system you will have in the cluster you need to:	—
	Set the hub linecard J2 and J3 jumpers to enable fiber optics.	Section 5.1.2
	Install the optical converters in the hub, ensuring that you connect the optical cable to the optical converter when it is installed.	Section 5.5.2.4
	Connect the fiber-optic module in the hub to the linecard with a 1-meter (3.3-foot) BN39B-01 link cable.	Section 5.5.2.4
2	Manually relocate all applications from <code>system1</code> to <code>system2</code> . Use the cluster application availability (CAA) <code>caa_relocate</code> command.	<code>caa_relocate(8)</code> reference page and <i>Cluster Administration</i>
3	On <code>system1</code> log in as the root user and execute the <code>shutdown -h</code> command to halt the system.	Tru64 UNIX <i>System Administration</i>

Table 5–5: Upgrading from a Virtual Hub Configuration to a Standard Hub Configuration (cont.)

Step	Action	Refer to:
Note		
<p>When <code>system1</code> is at the console prompt, note the setting of the <code>auto_action</code> console environment variable, then use the console <code>set</code> command to set the <code>auto_action</code> variable to <code>halt</code>. This halts the system at the console prompt when the system is turned on, ensuring that you are able to run the Memory Channel diagnostics.</p> <pre>P00>>> show auto_action : P00>>> set auto_action halt</pre>		
4	Turn off <code>system1</code> power.	—
5	Disconnect the Memory Channel cables from <code>system1</code> .	—
6	Wearing an antistatic wrist strap, remove the Memory Channel adapter modules and place them on a grounded work surface.	—
7	On each Memory Channel adapter module, move the hub mode jumper (J4 for MC1 or MC1.5 and J1 for MC2) to pins 1 and 2 to select standard hub mode.	Section 5.1 and Memory Channel <i>User's Guide</i>
Note		
<p>If you are also adding Memory Channel fiber optics capabilities, ensure that Memory Channel adapter module J10 and J11 jumpers are set to enable fiber optics.</p>		
8	Reinstall the Memory Channel modules.	Section 5.2
9	If you are adding fiber optics, install the optical converters in the member system.	Section 5.3
Note		
<p>Install the fiber-optic cable in cable runs between the hub and member system. Connect the fiber-optic cable to the optical converter when you install the converter in the system.</p> <p>Connect the fiber-optic module to the Memory Channel adapter module with a 1-meter (3.3-foot) FN39B-01 link cable.</p>		
10	Connect the Memory Channel cables between the Memory Channel adapter module and the Memory Channel hub and turn on hub power. If you have multiple adapters, each adapter must be connected to a different hub, and be in the same linecard slot position in each hub.	Section 5.5

Table 5–5: Upgrading from a Virtual Hub Configuration to a Standard Hub Configuration (cont.)

Step	Action	Refer to:
Note		
If you are using fiber optics with Memory Channel, you have already installed the fiber-optic cable. Turn on hub power.		
11	Turn on <code>system1</code> system power and run the <code>mc_diag</code> Memory Channel diagnostic. (You cannot run <code>mc_cable</code> because this is the only system in the cluster that is at the console prompt and no other systems are connected to the hub.)	Section 5.6
Note		
Set the <code>auto_action</code> console environment variable to its previous value, <code>restart</code> or <code>boot</code> , for instance:		
<pre>>>> set auto_action restart</pre>		
12	Use the <code>shutdown -h</code> or <code>shutdown -c</code> command to shut down cluster member <code>system2</code> .	—
13	When <code>system2</code> is at the console prompt, boot <code>system1</code> , the system that is connected to the Memory Channel hub.	—
14	Repeat steps 4 - 9 for <code>system2</code> .	—
15	Connect the Memory Channel cables between the Memory Channel adapter module and the Memory Channel hub. If you have multiple adapters, each adapter must be connected to a different hub, and must be in the same linecard slot position in each hub.	Section 5.5
16	Turn on <code>system2</code> power and run the <code>mc_diag</code> Memory Channel diagnostic. (You cannot run <code>mc_cable</code> because the other system is at multi-user mode.)	Section 5.6
Note		
Reset the <code>auto_action</code> console environment variable to its previous value, <code>restart</code> or <code>boot</code> , for instance:		
<pre>>>> set auto_action restart</pre>		
17	Boot <code>system2</code> .	—

You can now connect a new system to the Memory Channel hub. After configuring the hardware, use the `clu_add_member` command to add each new system to the cluster. (See the `clu_add_member(8)` reference page and the *Cluster Installation* manual for more information.)

6

Using Fibre Channel Storage

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

This chapter discusses the following topics:

- An overview of Fibre Channel (Section 6.1).
- A comparison of Fibre Channel topologies (Section 6.2).
- Example cluster configurations using Fibre Channel storage (Section 6.3).
- A brief discussion of Quickloop (Section 6.4).
- A discussion of zoning (Section 6.5).
- A discussion of cascaded switches (Section 6.6).
- A procedure for Tru64 UNIX Version 5.1A or TruCluster Server Version 5.1A installation using Fibre Channel disks (Section 6.7).
- The steps necessary to install and configure the Fibre Channel hardware (Section 6.8).
- The steps necessary to install the base operating system and cluster software using disks accessible over the Fibre Channel hardware (Section 6.9).
- How to convert the HSG80 from transparent to multiple-bus failover mode (Section 6.10).
- A discussion on how you can use the `emx` manager (`emxmgr`) to display the presence of Fibre Channel adapters, target ID mappings for a Fibre Channel adapter, and the current Fibre Channel topology (Section 6.11).

The information includes an example storageset configuration, 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 up the `bootdef_dev` console environment variable to facilitate Tru64 UNIX Version 5.1A and TruCluster Server Version 5.1A installation.

Note

TruCluster Server Version 5.1A 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.

TruCluster Server requires a cluster interconnect, which can be the Memory Channel, or a private LAN. (See the *Cluster LAN Interconnect* manual for more information on the LAN interconnect.) All the configurations in this chapter are shown with the Memory Channel interconnect. The illustrations could represent multiple standalone systems by removing the Memory Channel.

6.1 Fibre Channel 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 kilometers (24.85 miles), which is the maximum distance between transmitter and receiver. Copper transmission lines may be used for shorter distances.

6.1.1 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.

6.1.2 Fibre Channel Topologies

Fibre Channel supports three different interconnect topologies:

- Point-to-point (Section 6.1.2.1)
- Fabric (Section 6.1.2.2)
- Arbitrated loop (Section 6.1.2.3)

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.

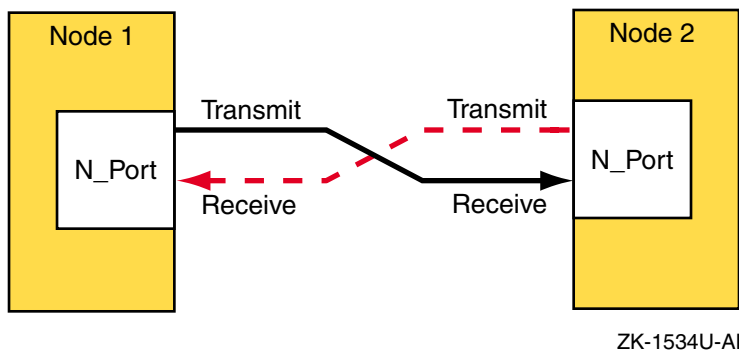
6.1.2.1 Point-to-Point

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 6–1 shows an example point-to-point topology.

Figure 6–1: Point-to-Point Topology



6.1.2.2 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 6–2 shows an example fabric topology.

Figure 6–2: Fabric Topology



ZK-1536U-AI

6.1.2.3 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 6–3 shows an example of an arbitrated loop topology.

Figure 6–3: Arbitrated Loop Topology



ZK-1535U-AI

6.2 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 6–1 compares the fabric and arbitrated loop topologies.

Table 6–1: 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

6.3 Example Fibre Channel Configurations Supported by TruCluster Server

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


6.3.1 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 6–4 shows a typical Fibre Channel cluster configuration using transparent failover mode.

Figure 6–4: Fibre Channel Single Switch Transparent Failover Configuration



ZK-1531U-AI

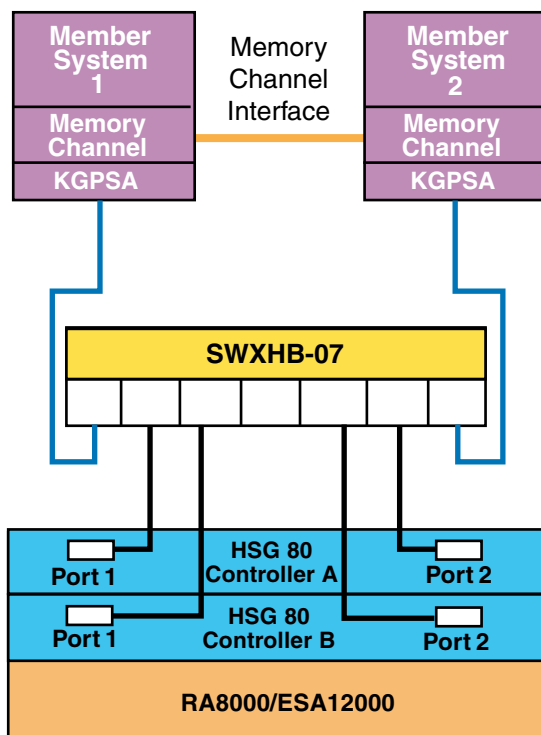
In transparent failover, units D00 through D99 are accessed through port 1 of both controllers. Units D100 through D199 are accessed through port 2 of both HSG80 controllers.

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 member boot partitions or the quorum disk you cannot obtain an NSPOF transparent failover configuration, even though you have increased availability.

Figure 6–5 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 6–5: Arbitrated Loop Configuration with One Storage Array



ZK-1697U-AI

6.3.2 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

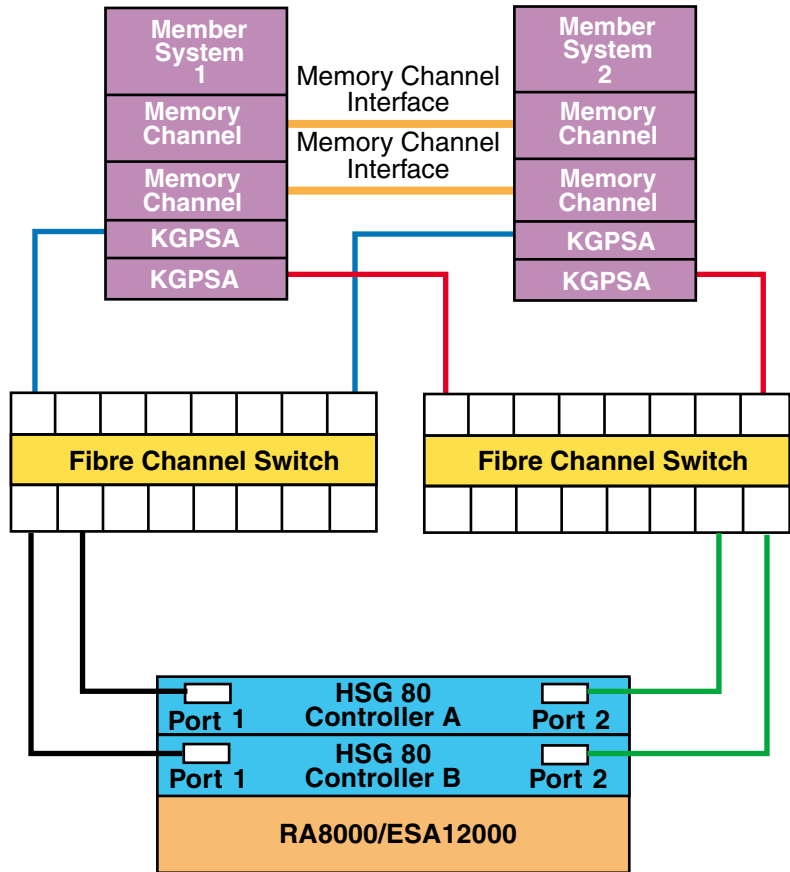
If you have preferred units, and the HSG80 controllers restart because of an error condition or power failure, and one controller restarts before the other controller, the HSG80 controller restarting first will take all the units, whether they are preferred or not. When the other HSG80 controller starts, 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 6–6 and Figure 6–7 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.

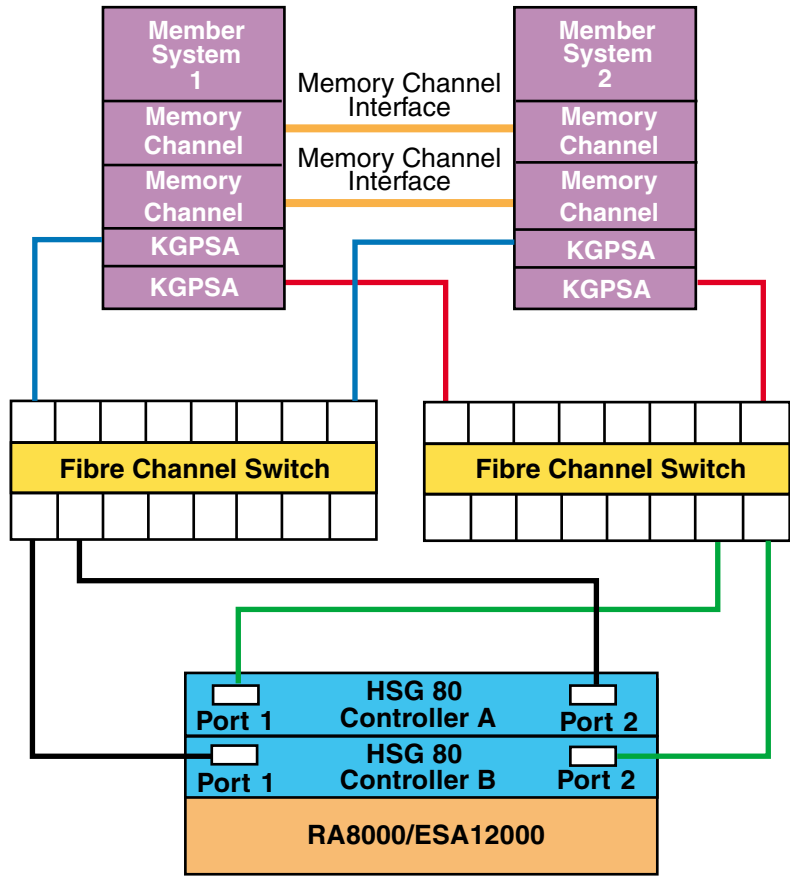
There is no difference in performance between these two configurations. It may be easier to cable the configuration shown in Figure 6–6 because the cables from one switch (or switch zone) both go to the ports on the same side of both controllers (for example, port 1 of both controllers).

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



ZK-1707U-AI

Figure 6–7: Multiple-Bus NSPOF Configuration Number 2



ZK-1765U-AI

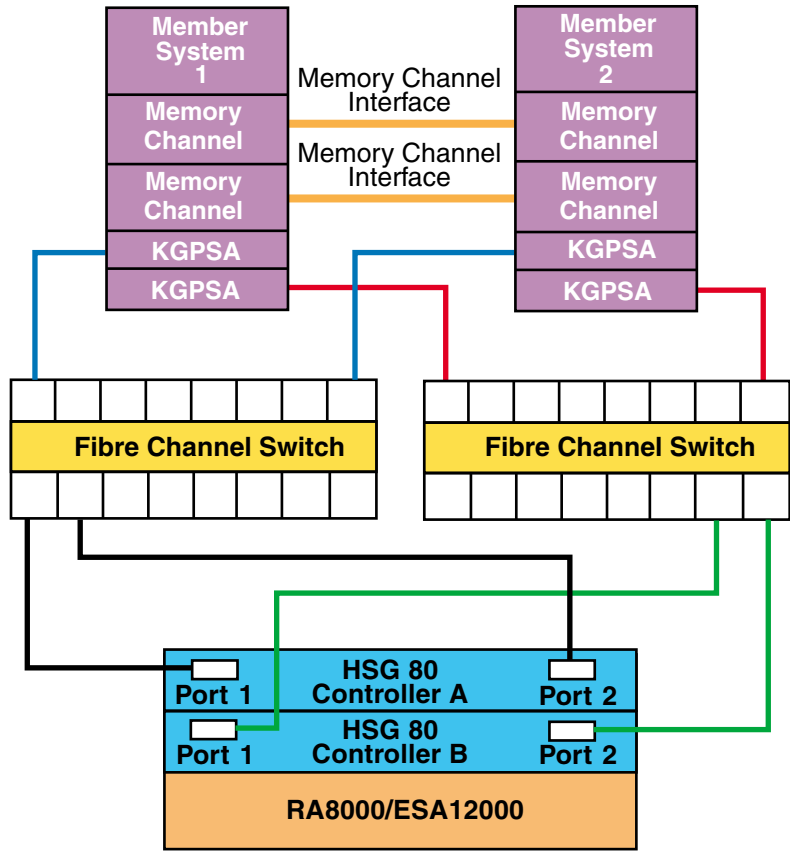
The configuration that is shown in Figure 6–8 is a NSPOF configuration, but is not a recommended cluster configuration because of the performance loss during failure conditions. If a switch or cable failure causes a failover to the other switch, access to the storage units has to be moved to the other controller, and that takes time. In the configurations shown in Figure 6–6 and Figure 6–7, the failure would cause access to the storage unit to shift to the other port of the same controller. This is faster than a change of controllers, providing better overall performance.

Note

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


The single-system configuration that is shown in Figure 6–9 is also a configuration that we do not recommend.

Figure 6–8: A Configuration That Is Not Recommended



ZK-1706U-AI

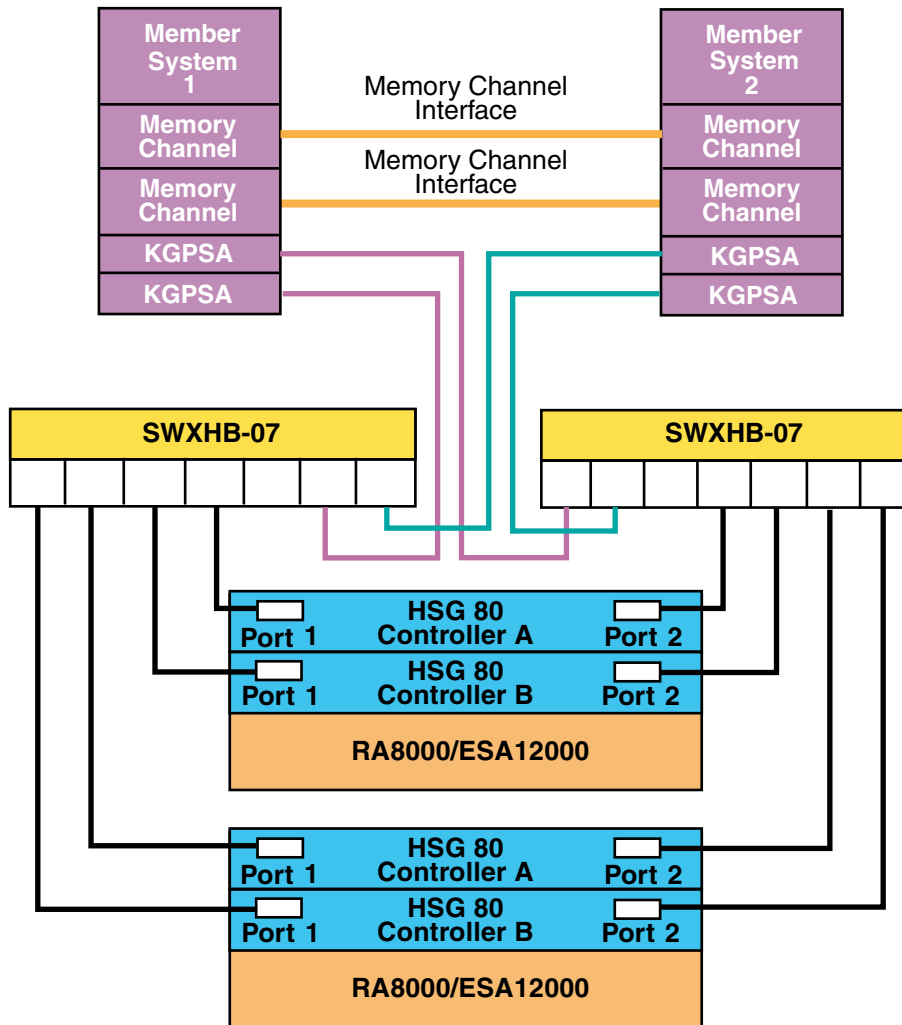
Figure 6–9: Another Configuration That Is Not Recommended



ZK-1806U-AI

Figure 6–10 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 6–10: Arbitrated Loop Maximum Configuration



ZK-1814U-AI

6.4 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 5.1A configuration or TruCluster Server Version 5.1A configuration.

6.5 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.

6.5.1 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).

6.5.2 Types of Zoning

There are two types of zoning, soft and hard:

- Soft zoning is a software implementation that is 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 when 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 6–2 lists the types of zoning that are supported on each of the supported Fibre Channel switches.

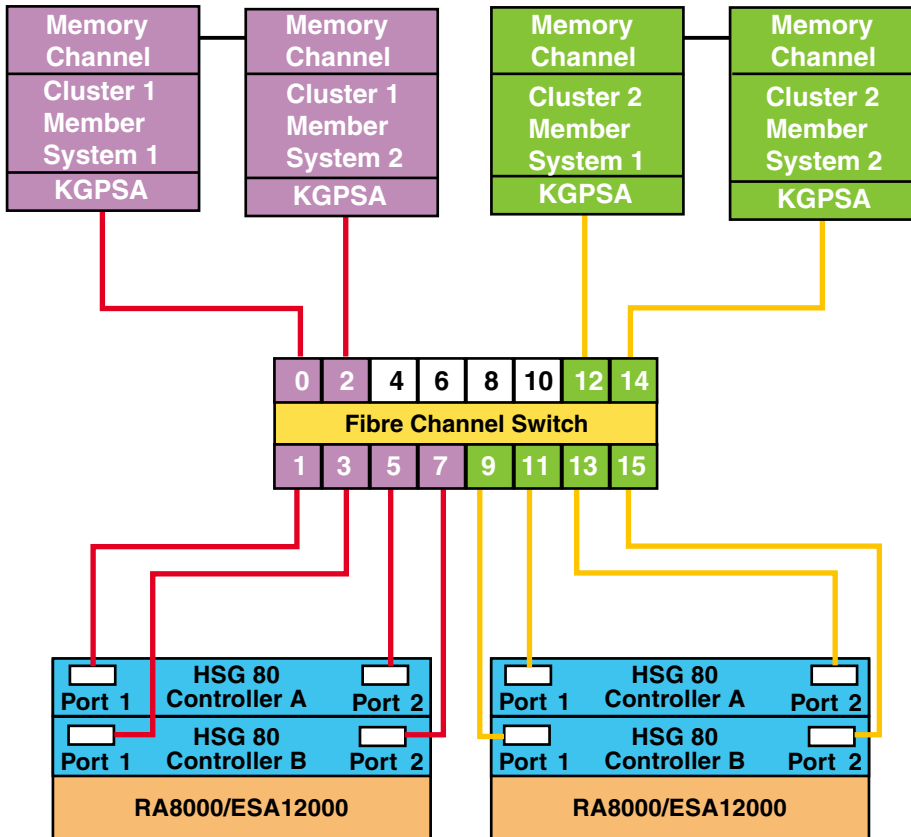
Table 6–2: Type of Zoning Supported by Switches

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

6.5.3 Zoning Example

Figure 6–11 provides an example configuration using zoning. This configuration consists of two independent zones with each zone containing an independent cluster.

Figure 6–11: A Simple Zoned Configuration



ZK-1709U-AI

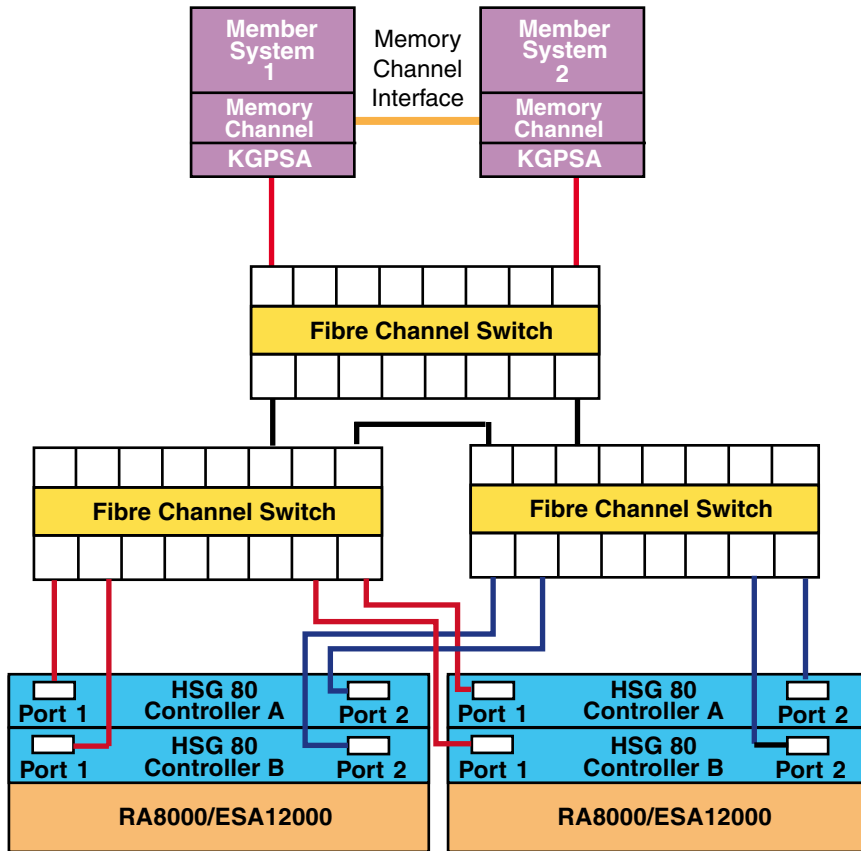
For information on setting up zoning, see the SAN Switch Zoning documentation that is provided with the switch.

6.6 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 6–12 shows an example meshed fabric with three cascaded switches. This is not a no-single-point-of-failure (NSPOF) configuration.


Figure 6–12: Meshed Fabric with Three Cascaded Switches



ZK-1795U-AI

Figure 6–13 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 6–13: Meshed Resilient Fabric with Four Cascaded Switches



ZK-1794U-AI

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)

6.7 Procedure for Installation Using Fibre Channel Disks

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

1. Install the Fibre Channel switch or hub (Section 6.8.1 or Section 6.8.2).
2. Install the KGPSA PCI-to-Fibre Channel host bus adapter (Section 6.8.3).
3. Set up the HSG80 RAID array controllers for a fabric or loop configuration (Section 6.8.4).
4. Configure the HSG80 disks to be used for installation of the base operating system and cluster. Be sure to set the identifier for each storage unit you will use for operating system or cluster installation (Section 6.9.1.1 and Section 6.9.1.2).
5. If the system is not already powered on, power on the system where you will install Tru64 UNIX Version 5.1A. 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 5.1A disk and first cluster member system boot disks (Section 6.9.1.3).

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 6.9.1.4).
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 6.9.1.5).
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 that you have chosen as the Tru64 UNIX operating system installation disk from the list of disks that is provided (Section 6.9.2).

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 6.9.6), then boot the operating system.

9. Determine the `/dev/disk/dskn` values to be used for cluster installation (Section 6.9.3).
10. Use the `disklabel` utility to label the disks that were used to create the cluster (Section 6.9.4).
11. See the TruCluster Server *Cluster 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 6.9.5).
12. Reset the `bootdef_dev` console environment variable to provide multiple boot paths to the cluster member boot disk (Section 6.9.6). Boot the first cluster member.
13. See the *Cluster Installation* manual and add subsequent cluster member systems (Section 6.9.7). 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.

6.8 Installing and Configuring Fibre Channel Hardware

This section provides information about installing the Fibre Channel hardware that is 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 that 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.

6.8.1 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 Gb/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.

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-optic 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-optic cables are supported between switches in a cascaded switch configuration.

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.

6.8.1.1 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, connect the switch to a terminal or PC (Section 6.8.1.2.3).
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 6.8.1.2.2 or Section 6.8.1.2.3). You can also set the switch name if desired (Section 6.8.1.2.5). 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 6.8.1.2.4).
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*.

6.8.1.2 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.

6.8.1.2.1 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.

6.8.1.2.2 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.

6.8.1.2.3 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.

6.8.1.2.4 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, which are listed from lowest security level to highest security level, are listed in Table 6-3.

Table 6–3: 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.

6.8.1.2.5 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>
```

6.8.2 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 Server products.

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

6.8.2.1 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-in) 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*.

6.8.2.2 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*.

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

The following sections discuss KGPSA installation and configuration.

6.8.3.1 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.

6.8.3.2 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.

6.8.3.3 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.

6.8.3.4 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 that 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.

6.8.4 Setting Up the HSG80 Array Controller for Tru64 UNIX Installation

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

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 6.8.4.1 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 they are not already installed, install the HSG80 controllers into the RA8000 or ESA12000 storage arrays.
2. If the external cache battery (ECB) is used, ensure that it 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 bits/sec
 - 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 6.8.4.1 for configuration differences.
- 8 Sets 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** Prompts you to restart both controllers after you set the SCSI version.

12. Enter the `show connection` command as shown in Example 6–1 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 6–1, 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 6–1: Determine HSG80 Connection Names

```
HSG80 show connection
```

Connection Name	Operating system	Controller	Port	Address	Status	Unit Offset
!NEWCON49	TRU64_UNIX HOST_ID=1000-0000-C920-DA01	THIS	2	230813 ADAPTER_ID=1000-0000-C920-DA01	OL this	0
!NEWCON50	TRU64_UNIX HOST_ID=1000-0000-C920-DA01	THIS	1	230813 ADAPTER_ID=1000-0000-C920-DA01	OL this	0
!NEWCON51	TRU64_UNIX HOST_ID=1000-0000-C920-EDEB	THIS	2	230913 ADAPTER_ID=1000-0000-C920-EDEB	OL this	0
!NEWCON52	TRU64_UNIX HOST_ID=1000-0000-C920-EDEB	THIS	1	230913 ADAPTER_ID=1000-0000-C920-EDEB	OL this	0
!NEWCON53	TRU64_UNIX HOST_ID=1000-0000-C920-EDEB	OTHER	1	230913 ADAPTER_ID=1000-0000-C920-EDEB	OL other	0

Example 6–1: 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 nine 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 6–1 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_system = TRU64_UNIX [2]
HSG80> restart other [3]
HSG80> restart this [3]
      :
HSG80> show connection [4]
```

- [1] Sets 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] Specifies that the host environment that is 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] Restarts both controllers to cause all changes to take effect.

- 4 Verifies 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 6.9.1.1.

6.8.4.1 Setting Up the HSG80 Array Controller for Arbitrated Loop

Section 6.8.4 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 6.8.4. 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 6.8.4.

6.8.4.2 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 6.8.4.2) and other configuration information, you can use the command-line interpreter (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
```

6.9 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 that is 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 section describes how to perform the following tasks:

- Before the installation:
 - a. Configure HSG80 storagesets — In this manual, example storagesets are configured for both Tru64 UNIX and TruCluster Server on Fibre Channel storage. Modify the storage configuration to meet your needs (Section 6.9.1.1).
 - b. 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 6.9.1.2).
 - c. 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 (which is 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 6.9.1.3).
 - d. Display available Fibre Channel boot devices — When you set the device unit number, you also set the `wwidn` and `Nn` 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 `wwidn` 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 6.9.1.4).
 - e. 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 6.9.1.5).
- Installing the base operating system and TruCluster Server software:
 - a. Install the base operating system (Section 6.9.2).
 - b. Determine the `/dev/disk/dskn` values to use as TruCluster Server disks (Section 6.9.3).
 - c. Label the disks that were used to create the cluster (Section 6.9.4).
 - d. Install the TruCluster Server software and create the first cluster member (Section 6.9.5).
 - e. Reset the `bootdef_dev` console environment variable (Section 6.9.6).

- f. Add additional systems to the cluster (Section 6.9.7).

If you are installing either the Tru64 UNIX operating system or TruCluster Server software, follow the procedure in Section 6.7.

6.9.1 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.

6.9.1.1 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 that are 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 four 4-GB disks that are 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 6–4 contains the necessary information to convert from the HSG80 unit numbers to /dev/disk/dskn 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 6–2.

Example 6–2: 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
  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
```


Example 6–2: Setting Up the Mirrorset (cont.)

```

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
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 Configures 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 Creates the BOOT-MIR mirrorset using disks 30200 and 30300 and the CROOT-MIR mirrorset using disks 40000 and 40100.

- 3 Initializes 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 Verifies the mirrorset configuration and switches. Ensure that the mirrorsets use the correct disks.
- 5 Creates 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 Creates 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 *Cluster Installation* manual to obtain partition sizes.
- 7 Verifies 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 Partitions for member system 1 boot disk.
- 9 Partitions for member system 2 boot disk.
- 10 Partitions for Tru64 UNIX operating system disk.
- 11 Verifies 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 Partitions for the quorum disk.
- 13 Partitions for cluster root (`/`) filesystem.
- 14 Partitions for cluster `/usr` filesystem.
- 15 Partitions for cluster `/var` filesystem.

6.9.1.2 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 6-3.

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 6–3 include:

1. Assigns a unit number to each storage unit and disables all access to the storage unit.
2. Sets an identifier for each storage unit.
3. Enable selective access to the storage unit.

Example 6–3: 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
              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.
```

Example 6–3: Adding Units and Identifiers to the HSG80 Stagesets (cont.)

```
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
-----
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 Assigns a unit number to each partition. When the unit is created by the ADD UNIT command, disables 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 6–4).

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 Sets 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 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/dskn` 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 6–4).

-
- 3 Enables 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 6–1 to determine the HSG80 connection names for the connection to the KGPSA host bus adapters. Many of the connections that are used here are shown in Example 6–1.

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 that are used here are shown in Example 6–1.

- 4 Using the `SHOW unit` command (where unit is D131 through D133 and D141 through 144 in the example), verifies 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 6–4 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 `dskn` numbers.

Table 6–4: Converting StorageSet Unit Numbers to Disk Names

File System or Disk	HSG80 Unit	Worldwide Name	UDID	Device Name	dskn
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

Table 6–4: Converting Stageset Unit Numbers to Disk Names (cont.)

File System or Disk	HSG80 Unit	Worldwide Name	UDID	Device Name	dsk _n
/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.

6.9.1.3 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 5.1A 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 stageset to be used as the Tru64 UNIX Version 5.1A 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 6–4, 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 `wwidn` and `Nn` 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 `wwidn` console environment variable has been set, and are accessible through an HSG80 N_Port as specified by the `Nn` 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 6-5.

3. Example 6-4 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 6-4: 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 Dn 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 6–2).
 - 4 The console environment variable set for this worldwide name. Only 4 `wwidn` console environment variables (`wwid0`, `wwid1`, `wwid2`, and `wwid3`) can be set. The console `show dev` command only shows those disk devices for which a `wwidn` console environment variable has been set using the `wwidmgr -quickset` or `wwidmgr -set` command. In this example, none of the `wwidn` environment variables is set.
4. Look through the `wwidmgr -show wwid display` (Example 6–4) 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 5.1A installation disk and the first cluster member system boot disk.

Example 6–5 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 6–5 shows:

- The use of the `wwidmgr -quickset` command to set the device unit number for the Tru64 UNIX Version 5.1A 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 6–5: 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
```

Example 6–5: Setting the Device Unit Number with the `wwidmgr quickset` Command (cont.)

	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 8x00, 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 6–5. You have to use the `quickset` command with the item number displayed by the `wwidmgr -show wwid` command (see Example 6–4). For instance, to set the device unit number of item number 17 of Example 6–4 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.

6.9.1.4 Displaying the Available Boot Devices

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

The devices that are shown in the reachability display of Example 6–5 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 `wwidn` 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 `wwidn` console environment variable set for the original storage unit. If you do not clear the `wwidn` console environment variable, the `wwidmgr` command used to quickset the UDID the second time sets a second `wwidn` console environment variable for the same UDID. The original `wwidn` 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 `wwidn` console environment variable for the same UDID, and the system has failed to boot, use the `wwidmgr` command to clear the `wwidn` console environment variable for the unreachable storage unit. For example, consider the following state of the `wwidn` 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 `wwidl` console environment variable as follows:

```
P00>>> wwidmgr -clear wwidl
```

Then, reboot the system.

Example 6–6 provides sample device names as displayed by the `show dev` command after using the `wwidmgr -quickset` command to set the device unit numbers.

Example 6–6: 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
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 that are displayed by the console `show dev` command are those devices that have been assigned to a `wwidn` environment variable.

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

6.9.1.5 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 6-5 with a `Yes` under the `connected` column. You can also use the `wwidmgr -show reachability` command to determine reachability. Example 6-6 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 6-5 and Example 6-6 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.

6.9.2 Install the Base Operating System

After you read the TruCluster Server *Cluster 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 5.1A operating system.

When the installation procedure displays the list of disks that are available for operating system installation as shown here, look for the identifier in the Location column. Verify the identifier from the table that you have been preparing (Table 6–4).

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 Name	Size in GB	Controller Type	Disk 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/dskn value (dsk15) for the Tru64 UNIX disk that matches the UDID (133) (Table 6–4).

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

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

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

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

1. With the Tru64 UNIX Version 5.1A 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/dskn` 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 :d05b003t00000100000"
  dev_base_name = cdrom0
```


⋮

For more information on the hardware manager, 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/dskn` 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/dskn` 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.

6.9.4 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/dskn` values to use for cluster installation (Section 6.9.3).

Initialize disklabels for all disks needed to create the cluster. The example (Table 6–4) uses disks `dsk18 (/var)`, `dsk19 (/usr)`, `dsk20 [cluster root (/)]`, and `dsk21 (Quorum)`. For instance:

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

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

After labeling the disks, use the TruCluster Server *Cluster 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 *Cluster 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 6.9.6). Then boot the first cluster member.

6.9.6 Reset the `bootdef_dev` Console Environment Variable

If you set the `bootdef_dev` console environment variable to multiple paths in Section 6.9.1.5, 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 6-4).
2. Examine the reachability display (Example 6-5) 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 6-5, 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_dev` 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  
P00>>> 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.

6.9.7 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 *Cluster 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 6–7. For member system 2 in the example configuration, it is the storage unit with UDID 132 (Table 6–4):

Example 6–7: 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 6–7) 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 6–4). Ensure that you choose the correct device name for the entity that you are booting (cluster member system 2, cluster member system 3, and so on).
 - b. Examine the reachability display (Example 6–7) that is 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 that 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 6–7, 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  
P00>>> 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.

6.10 Converting the HSG80 from Transparent to Multiple-Bus Failover Mode

If you are migrating from Tru64 UNIX Version 4.0F or Version 4.0G and TruCluster Software Products Version 1.6 to Tru64 UNIX Version 5.1A and TruCluster Server Version 5.1A, you may want to change from transparent failover to multiple-bus failover to take advantage of multibus support in Tru64 UNIX Version 5.1A 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 5.1A and TruCluster Server Version 5.1A and want to take advantage of the ability to create an NSPOF configuration, and the availability that multiple-bus failover provides over transparent failover mode.

6.10.1 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 that 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.

6.10.2 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 that you want to change from transparent failover to multiple-bus failover mode.

- 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 that you know has the good configuration information.

- 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
```

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

```
HSG80> SHOW CONNECTION
Connection
Name      Operating system  Controller  Port  Address      Status  Unit
      Name      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 that is 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 `wwidn` 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 `wwidn`). The `Nn` value is the HSG80 port being used to access the storage units (controller B, port 2).
- c. Clear the `wwidn` and `Nn` 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 6.9.6).
- i. Repeat steps a through h on each system accessing devices on the HSG80.

6.11 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 5.1A and, although it is not needed to maintain worldwide name to target ID mappings, you may use it with TruCluster Server Version 5.1A 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.

6.11.1 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) that 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.

6.11.2 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`.

6.11.3 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 5.1A 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

#

7

Using GS80, GS160, or GS320 Hard Partitions in a TruCluster Server Configuration

This chapter contains information about using AlphaServer GS80/160/320 hard partitions in a TruCluster Server Version 5.1A configuration with Tru64 UNIX Version 5.1A. The chapter discusses the following topics:

- An overview of the use of hard partitions in an AlphaServer GS80, GS160, or GS320 TruCluster Server configuration (Section 7.1).
- The hardware requirements for using an AlphaServer GS80, GS160, or GS320 hard partition in a cluster (Section 7.2).
- How to reconfigure a single partition AlphaServer GS80, GS160, or GS320 as multiple hard partitions in a TruCluster Server configuration (Section 7.3).
- How to determine an AlphaServer GS80, GS160, or GS320 system configuration (Section 7.4).
- How to update AlphaServer GS80, GS160, or GS320 firmware (Section 7.5).

7.1 Overview

An AlphaServer GS80/160/320 system provides the capability to define individual subsets of the system's computing resources. Each subset is capable of running an operating system.

The Tru64 UNIX Version 5.1A operating system supports hard partitions, which are partitions that are defined by a quad building block (QBB) boundary. All the CPUs, memory, and I/O resources in a QBB are part of a hard partition; you cannot split the components across multiple hard partitions, and resources cannot be shared between hard partitions. A partition can include multiple QBBs.

The TruCluster Server Version 5.1A product supports the use of AlphaServer GS80/160/320 hard partitions as a cluster member system. You can compose a cluster entirely of the partitions on a system, or of AlphaServer GS80/160/320 partitions and other AlphaServer systems. You can view an AlphaServer GS80/160/320 hard partition as a separate, standalone system.

The AlphaServer GS80/160/320 systems use the same switch technology, the same CPU, memory, and power modules, and the same I/O riser modules. The GS160 and GS320 systems house the modules in up to two system boxes, each with two QBBs, in a cabinet. The GS320 requires two cabinets for the system boxes.

The GS80 is a rack system with the system modules for each QBB in a drawer. An 8-processor GS80 uses two drawers for the CPU, memory, and I/O riser modules.

All the systems use the same type of PCI drawers for I/O. They are located in the GS160/GS320 power cabinet or in the GS80 RETMA cabinet. Additional PCI drawers are mounted in expansion cabinets.

7.2 Hardware Requirements for a Hard Partition in a Cluster

The TruCluster Server hardware requirements are the same for an AlphaServer GS80/160/320 hard partition as any other system in a cluster. You must have:


- A supported host bus adapter connected to shared storage. This may be a KZPBA-CB for parallel SCSI, or a KGPSA-CA for Fibre Channel.
- One or more network connections.
- A Memory Channel interface. The AlphaServer GS80/160/320 system supports only the MC2 products.

Each AlphaServer GS80/160/320 hard partition that is used in a cluster must contain at least one QBB with a minimum of one CPU and one memory module. Additionally, there must be:

- At least one local I/O riser module in the partition. Figure 7–1 shows a portion of an AlphaServer GS160 QBB with an I/O riser module with a BN39B cable that is connected to port 0.
- At least one I/O riser in the partition must be connected to a primary PCI drawer that provides the console terminal and operating system boot disk. For example, the portion of the cable on port 0 of the local I/O riser shown in Figure 7–1 could be connected to the I/O Riser 0 (0-R) connector in Figure 2–1 and Figure 7–3.

A primary PCI drawer contains a standard I/O module that provides both System Reference Manual (SRM) and system control manager (SCM) firmware. You can connect additional I/O risers in the partition to expansion PCI drawers.

Figure 7–1: Portion of QBB Showing I/O Riser Modules



Notes

You can have up to two I/O riser modules in a QBB, but you cannot split them across partitions.

Each I/O riser has two cable connections (Port 0 and Port 1). Ensure that both cables from one I/O riser are connected to the same PCI drawer (0-R and 1-R in Figure 2–1).

A QBB I/O riser (local) is connected to a PCI I/O riser (remote) by BN39B cables. These cables are the same cables that are used with MC2 hardware. Ensure that you connect the BN39B cable from a QBB I/O riser to the 0-R (I/O Riser 0) or 1-R (I/O Riser 1) connector in a PCI drawer and not to a Memory Channel module.

We recommend that you connect I/O riser 0 (local I/O riser ports 0 and 1) to the primary PCI drawer that will be the master system control manager (SCM).

The BA54A-AA PCI drawer (the bottom PCI drawer in Figure 7–2 and Figure 7–3) is a primary PCI drawer. See Figure 2–1 for PCI drawer slot layout. A primary PCI drawer contains:

- A standard I/O module in slot 0-0/1 that has EEPROMs for the system control manager (SCM) and system reference manual (SRM) firmware. The SCM is powered by the V_{aux} output of the PCI power supply whenever AC power is applied to the PCI drawer.

The master SCM uses the console serial bus (CSB) to:

- Control system power-up
 - Monitor and configure the system
 - Halt and reset the system
 - Update firmware
-
- Operating system disk
 - Two remote I/O riser modules (for connection to the QBB local I/O riser module)
 - Two PCI backplanes: Each PCI backplane (Figure 2–1) has two PCI buses. PCI bus 0 has three slots. PCI 1 has four slots. A primary PCI drawer has a standard I/O module in PCI bus 0 slot 0-0/1.
 - CD-ROM
 - Two power supplies (providing a redundant power supply)
 - Console serial bus (CSB) interface module: The console serial bus consists of a network of microprocessors that the master SCM controls in a master/slave relationship. Each node is programmed to control and monitor the subsystem in which it resides, in response to commands from, or when being polled, by the master SCM.

The CSB network consists of the following nodes:

- One to eight SCMs. The primary PCI drawer that is connected to the operator control panel (OCP), and, with the lowest node ID (usually 0), is the default master SCM upon initial power-up. The remaining SCMs are slaves. You can designate one slave SCM as a standby to the master. The primary PCI drawer with the slave SCM that you designate to be the standby must also be connected to the OCP. The OCP has two connectors for this purpose. The standby SCM must have a node ID (usually set to

1) that is higher than the master SCM. Both the master SCM and standby SCM must have the `scm_csb_master_eligible` SCM environment variable set.

Note

We recommend that you put the primary PCI drawers that contain the master and standby SCM in the power cabinet. They both must be connected to the OCP.

- One to eight power system managers (PSMs), one for each QBB
- One to 16 PCI backplane managers (PBMs), one for each PCI backplane
- A hierarchical switch power manager (HPMs), if the H-switch is present
- Local terminal/COM1 port (on the standard I/O module): Connect a cable from the local terminal port on the standard I/O module to the terminal server for each partition. The terminal server is connected to the system management console (PC) that provides a terminal emulator window for each console.
- Modem port (on the standard I/O module)
- Two universal serial bus (USB) ports (on standard I/O module)
- Keyboard port
- Mouse port
- Operator Control Panel (OCP) port
- Parallel port
- Communication port (COM2)

The BA54A-BA PCI drawer is an expansion PCI drawer (top PCI drawer in Figure 7–2 and Figure 7–3) and contains:

- Two I/O riser modules (for connection to a QBB I/O riser module)
- Two power supplies (which provides a redundant power supply)
- Two PCI backplanes. Each PCI backplane has 2 PCI buses, each with seven available slots.
- Console serial bus interface module

Figure 7–2 shows the front view of an expansion and a primary PCI drawer. The primary PCI drawer is on the bottom. You can easily recognize it because of the CD-ROM, keyboard and mouse ports, COM2 and parallel ports, and connection to the OCP. Figure 7–3 shows the rear view of both

types of PCI drawers. It is harder to distinguish the type of PCI drawer from the rear, but slot 1 provides the key. The primary PCI drawer has a standard I/O module in slot 1, and the console and modem ports and USB connections are visible on the module.

Figure 7–2: Front View of Expansion and Primary PCI Drawers





Figure 7-3: Rear View of Expansion and Primary PCI Drawers



7.3 Configuring Partitioned GS80, GS160, or GS320 Systems in a TruCluster Configuration

An AlphaServer GS80/160/320 system can be a member of a TruCluster Server configuration. Alternatively, any AlphaServer GS80/160/320 hard partition can participate as a member system, provided that the partition meets the hardware requirements that Section 7.2 describes.

The following section covers configuring a single partition AlphaServer GS80/160/320 system as multiple hard partitions in a TruCluster Server configuration. The description covers the case of a newly installed system that is to be used as two member systems in a TruCluster Server configuration.

7.3.1 Repartitioning a Single-Partition AlphaServer GS80/160/320 as Two Partitions in a Cluster

The information in this section assumes that this is a new AlphaServer GS80/160/320 system with hardware installed, the system management console is connected for the first partition, a terminal emulator window is open for the first partition, and that the system has been powered up and tested as a single partition. Also, this section assumes that you have determined which QBBs to use in each partition. Although the procedure specifies two hard partitions, the maximum for a GS80 system, it will work

equally well with any number of partitions (as supported by the system type) by modifying the amount and placement of hardware and the SCM environment variable values.

Notes

View each partition as a separate system.

Ensure that the system comes up as a single partition the first time that you turn power on. Do not turn the key switch on. Only turn on the AC circuit breakers. Use the SCM `set hp_count 0` command to ensure that the system comes up as a single partition. Then turn the key switch on to provide power to the system.

To repartition an AlphaServer GS80/160/320 system into two partitions to be used as TruCluster Server member systems, follow this procedure:

1. If necessary, install a primary PCI drawer for each additional hard partition beyond partition 0. Install any expansion PCI drawers as needed to provide additional PCI slots. Ensure that the system already has a primary PCI drawer for the first partition.

Note

We recommend that you install the primary PCI drawers that contain the master and standby SCM (if there is to be a standby SCM) in the power cabinet of a GS160 or GS320 or RETMA cabinet for a GS80; they both must be connected to the OCP.

2. Install the following hardware, as appropriate for your TruCluster Server configuration, in the primary (or expansion) PCI drawer of each partition and make all cable connections. Keep your configuration as symmetrical as possible to make troubleshooting and reconfiguration tasks easier.
 - Each system in a TruCluster Server configuration requires at least one Memory Channel adapter. Ensure that you abide by the restrictions described in Section 2.2, and that you connect the cables for Memory Channel interconnects to the Memory Channel modules and not to the I/O risers. The type of cables used, BN39B used for the Memory Channel interconnect are also used to connect the local I/O risers (on the QBB) to the remote I/O risers (on the PCI drawers).

- Shared storage that is connected to KZPBA-CB (parallel SCSI) or KGPSA-CA (Fibre Channel) host bus adapters.
 - Network controllers.
3. Install BN39B cables between the local I/O risers on the QBBs in the partition (see Figure 7–1) and the remote I/O risers in the primary and expansion PCI drawer (see Figure 2–1 and Figure 7–3). Use BN39B-01 cables (1-meter; 3.3-foot) for a PCI drawer in the GS80 RETMA cabinet. Use BN39B-04 cables (4-meter; 13.1-foot) if the PCI drawer is in a GS160 or GS320 power cabinet. Use BN39B-10 cables (10-meter; 32.8-foot) if the PCI drawer is in an expansion cabinet. Ensure that you connect the cables to the 0-R and 1-R (remote I/O riser) connections in the PCI drawer and not to a Memory Channel module.

Note

We recommend that you connect I/O riser 0 (local I/O riser ports 0 and 1) to the primary PCI drawer that will be the master system control manager (SCM).

If you require more than two PCI drawers in a hard partition, you need more than one QBB in the partition. Each QBB supports two PCI drawers (2 cables between a local I/O riser and a PCI drawer).

4. Set the PCI drawer node ID with the pushbutton up-down counter on the CSB node ID module at the rear of each PCI drawer (see Figure 7–3). Set the node ID of the primary PCI drawer with the master SCM to zero. Set the node ID of the primary PCI drawer with the standby SCM (if applicable) to one. Increment the PCI drawer node ID for successive PCI drawers.
5. Ensure that the primary PCI drawer that contains the master SCM is connected to the OCP. Connect the primary PCI drawer with the standby SCM (if applicable) to the OCP.
6. Connect an H8585-AA connector to the terminal port on the standard I/O module for the new partition. Connect a BN25G-07 cable between the H8585-AA connector and the terminal server to provide the console terminal connection to the system management console.
Use the system management console terminal emulator to create a new terminal window for the partition.
7. Turn on the AC circuit breakers for each of the QBBs. Doing so provides power to the console serial bus (CSB) and SCM. Do not turn on the OCP key switch; you do not have to go through the lengthy power-up sequence to partition the system.

Notes

If the OCP key switch is in the On or Secure position, the system will go through the power-up sequence.

In this case, when the power-up sequence terminates, power down the system with the `power off` SCM command, then partition the system.

If the `auto_quit_scm` SCM environment variable is set (equal 1), control will be passed to the SRM console firmware at the end of the power-up sequence. Use the escape sequence (`Esc Esc scm`) to transfer control to the SCM firmware. If the `auto_quit_scm` SCM environment variable is not set (equal 0), the SCM retains control.

If you execute the `power off` command at the master SCM, without designating a partition, power is turned off to the entire system. To turn power off to a partition, use the SCM `power off -par n`, where *n* is the partition number.

A slave SCM can only control power for its own partition.

- When the power-up self tests (POST) have completed, and the system has been powered down, use the master SCM to set the SCM environment variables to define the partitions.

The `hp_count` SCM environment variable defines the number of hard partitions. The `hp_qbb_mask n` SCM environment variables define which QBBs, by bit position, will be part of partition *n*. Example 7-1 shows how to set up two partitions, with each partition containing two QBBs. Partition 0 includes QBBs 0 and 1; partition 1 includes QBBs 2 and 3.

Use the `show nvr` SCM command to display the SCM environment variables.

Example 7-1: Defining Hard Partitions with SCM Environment Variables

```
scm_e0> set hp_count 2 1
scm_e0> set hp_qbb_mask0 3 2
scm_e0> set hp_qbb_mask1 c 3
scm_e0> show nvr 4
com1_print_en          1
hp_count                2 5
hp_qbb_mask0           3 5
hp_qbb_mask1          c 5
hp_qbb_mask2           0
hp_qbb_mask3           0
hp_qbb_mask4           0
```

Example 7–1: Defining Hard Partitions with SCM Environment Variables (cont.)

```
hp_qbb_mask5           0
hp_qbb_mask6           0
hp_qbb_mask7           0
srom_mask               ff f
xsrom_mask              ff ff ff ff ff ff ff ff ff 1 0 0
primary_cpu             ff
primary_qbb0           ff
auto_quit_scm           1 6
fault_to_sys            0
dimm_read_dis           0
scm_csb_master_eligible 1 7
perf_mon                20
scm_force_fsl           0
ocp_text                as gs160
auto_fault_restart      1
scm_sizing_time         c
```

- 1 Sets the number of hard partitions to 2.
 - 2 Sets bits 0 and 1 of the mask (0011) to select QBB 0 and QBB 1 for hard partition 0.
 - 3 Sets bits 2 and 3 of the mask (1100) to select QBB 2 and QBB 3 for hard partition 1.
 - 4 Displays the SCM environment variables (non-volatile ram) to verify that the hard partition variables are set correctly.
 - 5 Verifies that the hard partition environment variables are correct.
 - 6 Indicates that control will be transferred to the SRM console firmware at the end of a power-up sequence. If you want to execute SCM commands use the escape sequence (Esc Esc scm) to transfer control to the SCM firmware. If you want to ensure that control stays with the SCM at the end of a power-up sequence, set the `auto_quit_scm` SCM environment variable to zero.
 - 7 Indicates that the SCM on this primary PCI drawer is eligible to be selected as the master SCM on subsequent power-ups. It will be selected if it is connected to the OCP, its CSB node ID is the lowest of the SCMs that are eligible to become master, and the `scm_csb_master_eligible` environment variable is set.
9. Select one primary PCI drawer to be the master SCM and if desired, another primary PCI drawer to be a standby SCM by setting the

`scm_csb_master_eligible` environment variable. The master and standby SCM must be connected to the OCP. The master SCM must have the lowest node ID.

Use the node ID address obtained from the `show csb SCM` command (see Example 7-4). If multiple primary PCI drawers are eligible, the SCM on the PCI drawer with the lowest node ID is chosen as master. The other SCM will be a standby in case of a problem with the master SCM.

If the node ID switch is set to zero, the CSB node ID will be 10 (Example 7-4). If the node ID switch is set to one, the CSB node ID will be 11.

For example, the following command enables the SCMs in the primary PCI drawers at node IDs 10 and 11 (switch settings of 0 and 1) to be master (and standby) of the console serial bus.

```
SCM_E0> set scm_csb_master_eligible 10,11
```

Note

The system will hang if the master SCM is not connected to the OCP.

10. At the standby SCM, set the `hp_count` and `hp_qbb_maskn` SCM environment variables to match the setting at the master SCM:

```
SCM_E0> set hp_count 2
SCM_E0> set hp_qbb_mask0 3
SCM_E0> set hp_qbb_mask1 c
```

11. Turn the On/Off switch to the On or Secure position, then power on each of the partitions with the master SCM. After the power-up sequence completes, transfer control to the SRM console firmware as shown in Example 7-2.

Example 7-2: Turning Partition Power On

```
SCM_E0> power on -par 0 [1]
:
SCM_E0> power on -par 1 [2]
:
SCM_E0> quit [3]
```

- ❶ Turns on power to partition 0.
- ❷ Turns on power to partition 1.
- ❸ Transfers control from the SCM firmware to the SRM console firmware.

Note

If the `auto_quit_scm` SCM environment variable is set, control is passed to the SRM console firmware automatically at the end of the power-up sequence.

12. Obtain a copy of the latest firmware release notes for the AlphaServer system (see Section 7.5). Compare the present firmware revisions (see Example 7–4) with the required revisions that are indicated in the release notes. Update the firmware if necessary (see Section 7.5).
The SRM console firmware includes the ISP1020/1040-based PCI option firmware, which includes the KZPBA-CB. When you update the SRM console firmware, you are enabling the KZPBA-CB firmware to be updated. On a power-up reset, the SRM console loads PCI option firmware from the console system flash ROM into NVRAM for all Qlogic ISP1020/1040-based PCI options, including the KZPBA-CB PCI-to-Ultra SCSI adapter.
13. At the terminal emulator for each partition, access the SRM console firmware and complete each of the following as necessary:
 - a. If applicable, set the KZPBA-CB SCSI IDs and ensure that you have access to all the shared storage.
 - b. Run the Memory Channel diagnostics `mc_diag` and `mc_cable` to verify that the Memory Channel adapters are operational (Section 5.6).
 - c. Install the Tru64 UNIX operating system (see the Tru64 UNIX *Installation Guide*).
 - d. Install the TruCluster Server software (see the TruCluster Server *Cluster Installation* manual).
14. If you are using Fibre Channel storage, follow the procedures in Chapter 6, Using Fibre Channel Storage.
15. Set up highly available applications or services as required.

7.4 Determining AlphaServer GS80/160/320 System Configuration

You may be required to reconfigure an AlphaServer GS80/160/320 system that is not familiar to you. Before you start to reconfigure any system, you need to determine:

- The number of partitions in the system
- Which QBBs are in each partition
- Which PCI drawers are used by each partition
- Which PCI drawer is connected to each QBB
- The console serial bus (CSB) addresses

Determine the necessary information with the following system control manager (SCM) commands: `show nvr` (Example 7–1), `show system` (Example 7–3), and `show csb` (Example 7–4).

If you are at the SRM prompt, use the escape sequence (Esc Esc scm) to transfer control to the SCM firmware.

Example 7–3 shows the display for the `show system` SCM command for an AlphaServer GS160 system.

Example 7–3: Displaying AlphaServer GS160 System Information

```
SCM_E0> show system

System Primary QBB0 : 2
System Primary CPU  : 0 on QBB2


```

1	2	3	4	5	6	7	8	9	10
Par	hrd/sft	CPU	Mem	IOR3 IOR2 IOR1 IOR0	GP	QBB	Dir	PS	Temp
	QBB#	3210	3210	(pci_box.rio)	Mod	BP	Mod	321	(°C)
(0)	0/30	PPPP	--PP	--. - - - P0.1 P0.0	P	P	P	-PP	27.0
(0)	1/31	PPPP	--PP	--. - - - - - - -	P	P	P	-PP	26.0
(1)	2/32	PPPP	--PP	--. - - - P1.1 P1.0	P	P	P	PP-	26.0
(1)	3/33	PPPP	--PP	--. - - - - - - -	P	P	P	PP-	27.0

```

HSwitch  Type      Cables 7 6 5 4 3 2 1 0      Temp(°C)

HPM40    8-port          - - - - P P P P      29.0 11

```

12	13	14	15	16
PCI	Rise1-1	Rise1-0	Rise0-1	Rise0-0
Cab	7 6 5 4	3 2 1	7 6 5 4	3 2 1
10	L L L M	- M -	M L L L	L L S
11	L L L M	- M -	M L L L	L L S

- ❶ Hard partition number. There are two hard partitions in this example (0 and 1).
- ❷ QBB number and console serial bus (CSB) node ID. QBB 0 and 1 (CSB node IDs 30 and 31) are in partition 0. QBB 2 and 3 (CSB node IDs 32 and 33) are in partition 1.
- ❸ Status of the CPU module, which is present, powered up, and has passed self test (P). A dash (-) indicates an empty slot. An F indicates a self test failure. In this example, each QBB contains four CPU modules, each of which has passed self test.
- ❹ Status of the memory module, which is present, powered up, and has passed self test (P). A dash (-) indicates an empty slot. An F indicates a self test failure. In this example, each QBB contains two memory modules, both of which has passed self test.
- ❺ Status of the PCI drawer I/O risers that are plugged into the QBB I/O risers in the form of $Xm.n$. X can be a "P", "p", "F", or a dash (-). QBB local I/O risers are IOR0 (Port 0), IOR1 (Port 1), IOR2 (Port 2), and IOR3 (Port 3). A P (uppercase) indicates that power is on and self test passed. A p (lowercase) indicates that power is off and self test passed, and an F indicates a self test failure.

The $m.n$ numbers for each QBB indicate which PCI drawer ($m = 0$ through f) and which PCI drawer I/O riser ($n = 0, 1$) the local I/O riser is connected to. For example, QBB0 Port 0 (IOR0) is connected to PCI drawer 0 I/O riser 0 (P0.0); QBB0 Port 1 (IOR1) is connected to PCI drawer 0 I/O riser 1 (P0.1).

Dashes (-) in place of $m.n$ signify that the I/O riser module is not installed. The display always shows two sequences of - . - (for example - . - - . -) because there are two ports on a local I/O riser module.

The other sequence you may observe is $Px.x$, which indicates that the I/O riser module is installed, powered-up, and has passed self test, but a cable is not connected to the port. For example, a status of $Px.x$ P2.0 indicates that the local I/O riser is installed, but only one cable is connected.

- ❻ Status of the global port module, which passed self test.
- ❼ Status of the QBB backplane power system manager (PSM), which passed self test.
- ❽ Status of the QBB directory module, which passed self test.
- ❾ QBB power supply status. Each of these QBBs has two power supplies. A dash (-) indicates that there is no power supply in that position.
- ❿ QBB backplane temperature in degrees Celsius.

- 11** Hierarchical switch (H-switch) type, status, temperature, and a report of which QBBs are connected to the H-switch. In this example, QBBs 0, 1, 2, and 3 are connected to the H-switch.
- 12** Console serial bus node ID for PCI drawers. In this example, the first PCI drawer has node ID 10. The second PCI drawer has node ID 11. Note that in this case, the node ID switches are set to 0 and 1.
- 13** Status of each of the four PCI buses in a PCI drawer. An S indicates that a standard I/O module is present. Other modules present in a slot are identified by their power dissipation:
 - L: Lower power dissipation
 - M: Medium power dissipation
 - H: High power dissipation
 - Dash (-): There is no module in that slot.

In this example, the PCI modules with M (medium) power dissipation are Memory Channel and Fibre Channel-to-PCI host bus adapters.

- 14** An indication of the presence or absence of the I/O riser modules in the PCI drawer. An asterisk (*) indicates that a module is present.
- 15** Status of the PCI drawer power supplies as follows:
 - A P (uppercase) indicates that the power supply is powered on and passed self test.
 - A p (lowercase) indicates that the power supply passed self test but has been powered off.
 - An F (uppercase) indicates that the power supply is powered on and failed self test.
 - An f (lowercase) indicates that the power supply failed self test and has been powered off.
 - An asterisk (*) indicates that the SCM has detected the presence of the power supply, but that there has been no attempt to power on the power supply.
- 16** PCI drawer temperature in degrees Celsius.

Example 7–4 shows the display for the `show csb` SCM command for an AlphaServer GS160 system.

Example 7–4: Displaying Console Serial Bus Information

```

SCM_E0> show csb
 1  2          3          4          5          6
CSB Type      Firmware Revision    FSL Revision    Power State
10 PBM        T05.4 (03.24/01:14)  T4.2 (09.08)    ON
11 PBM        T05.4 (03.24/01:14)  T4.2 (09.08)    ON

```


Example 7-4: Displaying Console Serial Bus Information (cont.)

```

30 PSM          T05.4 (03.24/01:09) T4.0 (07.06) ON          SrvSw: NORMAL
30 XSROM       T05.4 (03.24/02:10)
C0 CPU0/SROM   V5.0-7
C1 CPU1/SROM   V5.0-7
C2 CPU2/SROM   V5.0-7
C3 CPU3/SROM   V5.0-7
C0 IOR0
C1 IOR1
31 PSM          T05.4 (03.24/01:09) T4.0 (07.06) ON          SrvSw: NORMAL
31 XSROM       T05.4 (03.24/02:10)
C4 CPU0/SROM   V5.0-7
C5 CPU1/SROM   V5.0-7
C6 CPU2/SROM   V5.0-7
C7 CPU3/SROM   V5.0-7
32 PSM          T05.4 (03.24/01:09) T4.0 (07.06) ON          SrvSw: NORMAL
32 XSROM       T05.4 (03.24/02:10)
C8 CPU0/SROM   V5.0-7
C9 CPU1/SROM   V5.0-7
CA CPU2/SROM   V5.0-7
CB CPU3/SROM   V5.0-7
C8 IOR0
C9 IOR1
33 PSM          T05.4 (03.24/01:09) T4.0 (07.06) ON          SrvSw: NORMAL
33 XSROM       T05.4 (03.24/02:10)
CC CPU0/SROM   V5.0-7
CD CPU1/SROM   V5.0-7
CE CPU2/SROM   V5.0-7
CF CPU3/SROM   V5.0-7
40 HPM          T05.4 (03.24/01:18) X4.1 (08.18) ON
E0 SCM MASTER  T05.4 (03.24/01:21) T4.2 (09.08) ON
E1 SCM SLAVE   T05.4 (03.24/01:21) T4.2 (09.08) ON          Ineligible

```

- ❶ Console serial bus (CSB) node ID, or in the case of a QBB, the CPU number in the QBB. The CSB node address ranges are as follows:
- 10 to 1f: PCI backplane manager (PBM) — The CSB node ID is based on the PCI drawer node ID setting.
 - e0 to e7: System control manager (SCM) — The CSB node ID is also based on the PCI drawer node ID setting.
 - 30 to 37: Power system manager (PSM) — Based on the hard QBB ID (QBB 0 - 7)
 - 40: Hierarchical switch power manager (HPM)
 - C0 to CF: In response to the SCM `show csb` command, the PSM provides CSB node addresses for the CPUs and I/O risers even though they are not on the console serial bus. This enables SCM commands to be directed at any specific CPU, for instance `power off -cpu c4`. The PSM responds to SCM commands and powers the CPU on or off.

- ❷ Type of CSB node:

- PBM (PCI backplane manager)
 - PSM (Power system manager)
 - HPM (Hierarchical switch power manager)
 - SCM master: This PCI primary drawer has the master SCM.
 - SCM slave: The SCM on this PCI primary drawer is a slave and has not been designated as a backup to the master.
 - CPU_n/SROM: Each CPU module has SROM firmware that is executed as part of the power-up sequence.
 - XROM: Each CPU executes this extended SROM firmware on the PSM module after executing the SROM firmware.
- 3 Revision level of the firmware and compilation date.
 - 4 Revision level of the fail-safe loader (FSL) firmware. Each microprocessor on the CSB has both a normal firmware image in its flash ROM and a fail-safe loader image in a backup ROM. The fail-safe loader firmware is executed when the system is reset. It performs a checksum on the normal firmware image, and then passes control to the normal firmware image.
 - 5 State of power for each CPU, I/O riser, and each node on the CSB.
 - 6 An indication that power is normal (NORMAL), or that the QBB power is off and can be serviced (SERVICE).

The `Ineligible` field for the slave SCM indicates that the SCM is not a backup to the master SCM.

7.5 Updating GS80/160/320 Firmware

Occasionally you must update the AlphaServer GS80/160/320, or PCI host bus adapter firmware. To determine the need for a firmware update, you compare the current firmware versions with the versions available on the latest AlphaServer firmware update CD-ROM. The firmware release notes for the system provide a list of current firmware versions.

See Section 4.2 for two methods of obtaining the firmware release notes.

The following section provides an overview of how to update the firmware.

7.5.1 Updating AlphaServer GS80/160/320 Firmware

You can update the AlphaServer GS80/160/320 firmware with the loadable firmware update (LFU) utility by booting the AlphaServer Firmware Update CD-ROM.

You can use the LFU to update the following firmware:

- System Reference Manual (SRM) flash ROM on the standard I/O module
- The flash ROMs for the following console serial bus (CSB) microprocessors:
 - SCM: One on the standard I/O module of each primary PCI drawer
 - Power system manager (PSM): One on the PSM module in each QBB
 - PCI backplane manager (PBM): One on each PCI backplane
 - Hierarchical switch power manager (HPM): One on the H-switch
- PCI host bus adapter EEPROMS

To update the AlphaServer GS80/160/320 firmware with the LFU utility, follow these steps:

1. At the console for each partition, shut down the operating system.
2. At the master SCM, turn power off to the system:

```
SCM_E0> power off
```

You can turn power off to individual partitions if you want. Ensure that power is turned off to all partitions.

```
SCM_E0> power off -par 0
SCM_E0> power off -par 1
```

3. Use the `show nvr SCM` command to display SCM environment variables. Record the `hp_count` and `hp_qbb_maskn` environment variables as a record of the hardware partition configuration. You do not change the `hp_qbb_maskn` environment variables, but record the variables anyway.

```
SCM_E0> show nvr
COM1_PRINT_EN          1
HP_COUNT                2
HP_QBB_MASK0           3
HP_QBB_MASK1           c
HP_QBB_MASK2           0
HP_QBB_MASK3           0
HP_QBB_MASK4           0
HP_QBB_MASK5           0
HP_QBB_MASK6           0
HP_QBB_MASK7           0
:
```

4. Remove all hardware partitions:

```
SCM_E0> set hp_count 0
```

Note

You do not need to zero the `hp_qbb_maskn` environment variables, only the `hp_count`.

5. Turn power on to the system to allow SRM console firmware execution. The SRM code is copied to memory on the partition primary QBB during the power-up initialization sequence. SRM code is executed out of memory, not the SRM EEPROM on the standard I/O module.

```
SCM_E0> power on
```

6. Transfer control from the SCM to SRM console firmware (if the `auto_quit_scm` SCM environment variable is not set):

```
SCM_E0> quit
P00>>>
```

7. Use the console `show device` command to determine which device is the CD-ROM.

8. Place the AlphaServer Firmware Update CD-ROM in the drive and boot:

```
P00>>> boot dqa0
```

The boot sequence provides firmware update overview information. Press Return to scroll the text, or press Ctrl/C to skip the text.

After the overview information has been displayed, the name of the default boot file is provided. If it is the correct boot file, press Return at the `Bootfile:` prompt. Otherwise, enter the name of the file from which you want to boot.

The LFU help message shown in the following example is displayed:

```
*****Loadable Firmware Update Utility*****
```

Function	Description
Display	Displays the system's configuration table.
Exit	Done exit LFU (reset).
List	Lists the device, revision, firmware name and update revision
Readme	Lists important release information.
Update	Replaces current firmware with loadable data image.
Verify	Compares loadable and hardware images.
? or Help	Scrolls this function table.

The `list` command indicates, in the `device` column, which devices it can update. It also shows the present firmware revision and the update revision on the CD-ROM.

Use the `update` command to update all firmware, or you can designate a specific device to update; for example, SRM console firmware:

```
UPD> update srm
```

Caution

Do not abort the update — doing so can cause a corrupt flash image in a firmware module.

A complete firmware update for a QBB can take from 5 minutes for a PCI with no updatable devices to over 30 minutes for a PCI with many updatable devices. The length of time increases proportionally with the number of PCI adapters that you have.

9. After you update the firmware, use the `verify` command to verify the firmware update, then transfer control back to the SCM and reset the system:

```
P00>>> EscEsc scm  
SCM_E0> reset
```

10. Set the hard partitions back to the original configuration:

```
SCM_E0> set hp_count 2
```

11. At the master SCM, turn system power on:

```
SCM_E0> power on
```

12. At the master SCM, transfer control to the SRM console firmware. Then, using the SRM at the console of each partition, boot the operating system.

8

Configuring a Shared SCSI Bus for Tape Drive Use

The topics in this section provide information on preparing the various tape devices for use on a shared SCSI bus with the TruCluster Server product. The topics discussed include preparing the following tape drives for shared SCSI bus usage:

- TZ88 (Section 8.1)
- TZ89 (Section 8.2)
- Compaq 20/40 GB DLT Tape Drive (Section 8.3)
- Compaq 40/80-GB DLT Drive (Section 8.4)
- TZ885 (Section 8.5)
- TZ887 (Section 8.6)
- TL891 and TL892 DLT MiniLibraries (Section 8.7)
- TL890 DLT MiniLibrary Expansion Unit (Section 8.8)
- TL894 DLT Automated Tape Library (Section 8.9)
- TL895 DLT Automated Tape Library (Section 8.10)
- TL893 and TL896 Automated Tape Libraries (Section 8.11)
- TL881 and TL891 DLT MiniLibraries (Section 8.12)
- Compaq ESL9326D Enterprise Library (Section 8.13)

Notes

Section 8.7 and Section 8.8 provide documentation for the TL890/TL891/TL892 MiniLibrary family as sold with the DS-TL891-NE/NG, DS-TL891-NT, DS-TL892-UA, DS-TL890-NE/NG part numbers.

The TL881, with a Compaq 6-3 part number was recently qualified in cluster configurations. The TL891 rackmount base unit has been provided with a Compaq 6-3 part number. The TL881 and TL891 only differ in the type of tape drive they use.

They both work with an expansion unit (previously called the DS-TL890-NE) and a new module called the data unit.

Section 8.12 covers the TL881 and TL891 with the common components as sold with the Compaq 6-3 part numbers.

As long as the TL89x MiniLibrary family is being sold with both sets of part numbers, the documentation will retain the documentation for both ways to configure the MiniLibrary.

8.1 Preparing the TZ88 for Shared Bus Usage

Two versions of the TZ88 are supported: the TZ88N-TA tabletop standalone enclosure, and the TZ88N-VA StorageWorks building blocks (SBB) 13.34-centimeter (5.25-inch) carrier.


As with any of the shared SCSI devices, the TZ88N-TA and TZ88N-VA SCSI IDs must be set to ensure that no two SCSI devices on the shared SCSI bus have the same SCSI ID.

The following sections describe preparing the TZ88 in more detail.

8.1.1 Setting the TZ88N-VA SCSI ID

You must set the TZ88N-VA switches before the tape drive is installed into the BA350 StorageWorks enclosure. The Automatic selection is normally used. The TZ88N-VA takes up three backplane slot positions. The physical connection is in the lower of the three slots. For example, if the tape drive is installed in slots 1, 2, and 3 with the switches in Automatic, the SCSI ID is 3. If the tape drive is installed in slots 3, 4, and 5 with the switches in Automatic, the SCSI ID is 5. The switch settings are listed in Table 8-1. Figure 8-1 shows the TZ88N-VA with the backplane interface connector and SCSI ID switch pack.

Figure 8–1: TZ88N-VA SCSI ID Switches



TZ88N-VA

Table 8–1: TZ88N-VA Switch Settings

SCSI ID	SCSI ID Selection Switches					
	1	2	3	4	5	6
Automatic ^a	Off	Off	Off	On	On	On
0	Off	Off	Off	Off	Off	Off
1	On	Off	Off	Off	Off	Off
2	Off	On	Off	Off	Off	Off
3	On	On	Off	Off	Off	Off
4	Off	Off	On	Off	Off	Off
5	On	Off	On	Off	Off	Off

Table 8–1: TZ88N-VA Switch Settings (cont.)

SCSI ID	SCSI ID Selection Switches					
6	Off	On	On	Off	Off	Off
7	On	On	On	Off	Off	Off

^a SBB tape drive SCSI ID is determined by the SBB physical slot.

8.1.2 Cabling the TZ88N-VA

There are no special cabling restrictions specific to the TZ88N-VA; it is installed in a BA350 StorageWorks enclosure. A DWZZA-VA installed in slot 0 of the BA350 provides the connection to the shared SCSI bus. The tape drive takes up three slots.

You must remove the DWZZA-VA differential terminators. Ensure that DWZZA-VA jumper J2 is installed to enable the single-ended termination. The BA350 jumper and terminator must be installed.

A tralink connector on the DWZZA-VA differential end allows connection to the shared bus. An H879-AA terminator is installed on the tralink for the BA350 on the end of the bus to provide shared SCSI bus termination.

Figure 8–2 shows a TruCluster Server cluster with three shared SCSI buses. One shared bus has a BA350 with a TZ88N-VA at SCSI ID 3.

Figure 8–2: Shared SCSI Buses with SBB Tape Drives

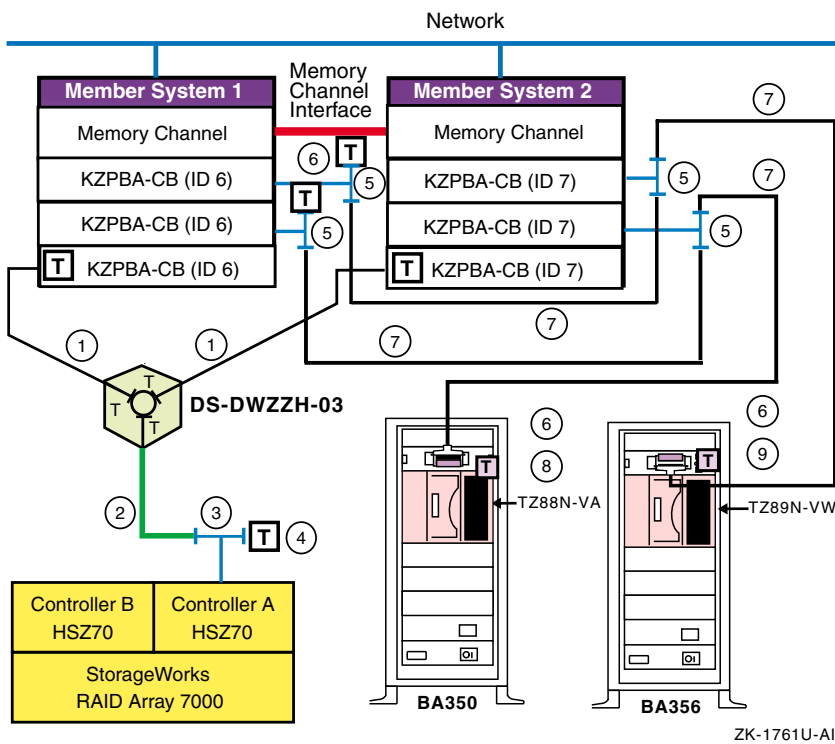


Table 8–2 lists the components that are used to create the cluster that is shown in Figure 8–2.

Table 8–2: Hardware Components Used to Create the Configuration Shown in Figure 8–2

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI tralink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

Table 8–2: Hardware Components Used to Create the Configuration Shown in Figure 8–2 (cont.)

Callout Number	Description
8	DWZZA-VA with H885-AA tralink connector
9	DWZZB-VW with H885-AA tralink connector

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.1.3 Setting the TZ88N-TA SCSI ID

The TZ88N-TA SCSI ID is set with a push-button counter switch on the rear of the unit. Push the button above the counter to increment the address; push the button below the counter to decrement the address until you have the desired SCSI ID selected.

8.1.4 Cabling the TZ88N-TA

You must connect the TZ88N-TA tabletop model to a single-ended segment of the shared SCSI bus. It is connected to a differential portion of the shared SCSI bus with a DWZZA-AA or DWZZB-AA. Figure 8–7 shows a configuration of a TZ885 for use on a shared SCSI bus. You can replace the TZ885 shown in the illustration with a TZ88N-TA. To configure the shared SCSI bus for use with a TZ88N-TA, follow these steps:

1. You will need one DWZZA-AA or DWZZB-AA for each TZ88N-TA.
Ensure that DWZZA jumper J2 or DWZZB jumpers W1 and W2 are installed to enable the single-ended termination.
Remove the termination from the differential end by removing the five 14-pin SIP resistors.
2. Attach a tralink connector to the differential end of the DWZZA or DWZZB.
3. Connect the single-ended end of a DWZZA to the TZ88N-TA with a BC19J cable.
Connect the single-ended end of a DWZZB to the TZ88N-TA with a BN21M cable.
4. Install a H8574-A or H8890-AA terminator on the other TZ88N-TA SCSI connector.
5. Connect a tralink or Y cable to the differential shared SCSI bus with BN21K, BN21L, or 328215-00X cables. Ensure that the tralink or Y cable at the end of the bus is terminated with an H879-AA terminator.

The single-ended SCSI bus may be daisy chained from one single-ended tape drive to another with BC19J cables as long as the SCSI bus maximum length is not exceeded. Ensure that the tape drive on the end of the bus is terminated with a H8574-A or H8890-AA terminator.

You can add additional TZ88N-TA tape drives to the differential shared SCSI bus by adding additional DWZZA or DWZZB/TZ88N-TA combinations.

Note

Ensure that there is no conflict with tape drive, system, and disk SCSI IDs.

8.2 Preparing the TZ89 for Shared SCSI Usage


Like the TZ88, the TZ89 comes in either a tabletop (DS-TZ89N-TA) or a StorageWorks building block (SBB) 13.34-centimeter (5.25-inch) carrier (DS-TZ89N-VW). The SBB version takes up three slots in a BA356 StorageWorks enclosure.

The following sections describe how to prepare the TZ89 in more detail.

8.2.1 Setting the DS-TZ89N-VW SCSI ID

The DS-TZ89N-VW backplane connector makes a connection with the backplane in the middle of the three slots occupied by the drive. If the switches are set to automatic to allow the backplane position to select the SCSI ID, the ID corresponds to the backplane position of the middle slot. For example, if the DS-TZ89N-VW is installed in a BA356 in slots 1, 2, and 3, the SCSI ID is 2. If it is installed in slots 3, 4, and 5, the SCSI ID is 4. Figure 8-3 shows a view of the DS-TZ89N-VW showing the backplane interface connector and SCSI ID switch pack.

Figure 8–3: DS-TZ89N-VW SCSI ID Switches



The SCSI ID is selected by switch positions, which must be selected before the tape drive is installed in the BA356. Table 8–3 lists the switch settings for the DS-TZ89N-VW.

Table 8–3: DS-TZ89N-VW Switch Settings

SCSI ID	SCSI ID Selection Switches							
	1	2	3	4	5	6	7	8
Automatic ^a	Off	Off	Off	Off	On	On	On	On
0	Off	Off	Off	Off	Off	Off	Off	Off
1	On	Off	Off	Off	Off	Off	Off	Off
2	Off	On	Off	Off	Off	Off	Off	Off
3	On	On	Off	Off	Off	Off	Off	Off
4	Off	Off	On	Off	Off	Off	Off	Off

Table 8–3: DS-TZ89N-VW Switch Settings (cont.)

SCSI ID	SCSI ID Selection Switches							
	1	2	3	4	5	6	7	8
5	On	Off	On	Off	Off	Off	Off	Off
6	Off	On	On	Off	Off	Off	Off	Off
7	On	On	On	Off	Off	Off	Off	Off
8	Off	Off	Off	On	Off	Off	Off	Off
9	On	Off	Off	On	Off	Off	Off	Off
10	Off	On	Off	On	Off	Off	Off	Off
11	On	On	Off	On	Off	Off	Off	Off
12	Off	Off	On	On	Off	Off	Off	Off
13	On	Off	On	On	Off	Off	Off	Off
14	Off	On	On	On	Off	Off	Off	Off
15	On	On	On	On	Off	Off	Off	Off

^a SBB tape drive SCSI ID is determined by the SBB physical slot.

8.2.2 Cabling the DS-TZ89N-VW Tape Drives

No special cabling is involved with the DS-TZ89N-VW because it is installed in a BA356 StorageWorks enclosure. A DWZZB-VA installed in slot 0 of the BA356 provides the connection to the shared SCSI bus.

You must remove the DWZZB-VW differential terminators. Ensure that jumpers W1 and W2 are installed to enable the single-ended termination. The BA356 jumper must be installed, and connector JB1 on the personality module must be left open to provide termination at the other end of the single-ended bus.

A tralink connector on the differential end of the DWZZB-VW allows connection to the shared bus. If the BA356 containing the DS-TZ89N-VW is on the end of the bus, install an H879-AA terminator on the tralink for that BA356 to provide termination for the shared SCSI bus.

Figure 8–2 shows a TruCluster Server cluster with three shared SCSI buses. One shared bus has a BA356 with a DS-TZ89N-VW at SCSI ID 2.

Table 8–2 lists the components that are used to create the cluster that is shown in Figure 8–2.

8.2.3 Setting the DS-TZ89N-TA SCSI ID

The DS-TZ89N-TA has a push-button counter switch on the rear panel to select the SCSI ID. It is preset at the factory to 15. Push the button above the counter to increment the SCSI ID (the maximum is 15); push the button below the switch to decrease the SCSI ID.

8.2.4 Cabling the DS-TZ89N-TA Tape Drives

You must connect the DS-TZ89N-TA tabletop model to a single-ended segment of the shared SCSI bus. It is connected to a differential portion of the shared SCSI bus with a DWZZB-AA. Figure 8–7 shows a configuration of a T885 for use on a shared SCSI bus. Just replace the TZ885 in the figure with a DS-TZ89N-TA and the DWZZA-AA with a DWZZB-AA. To configure the shared SCSI bus for use with a DS-TZ89N-TA, follow these steps:

1. You will need one DWZZB-AA for each DS-TZ89N-TA.
Ensure that the DWZZB jumpers W1 and W2 are installed to enable the single-ended termination.
Remove the termination from the differential end by removing the five 14-pin SIP resistors.
2. Attach a trilink connector to the differential end of the DWZZB-AA.
3. Connect the DWZZB-AA single-ended end to the DS-TZ89N-TA with a BN21K, BN21L, or 328215-00X cable.
4. Install an H879-AA terminator on the other DS-TZ89N-TA SCSI connector.
5. Connect the trilink to the differential shared SCSI bus with BN21K, BN21L, or 328215-00X cables. Ensure that the trilink at the end of the bus is terminated with an H879-AA terminator.

The wide, single-ended SCSI bus may be daisy chained from one single-ended tape drive to another with BN21K, BN21L, or 328215-00X cables as long as the SCSI bus maximum length is not exceeded. Ensure that the tape drive on the end of the bus is terminated with an H879-AA terminator.

You can add additional DS-TZ89N-TA tape drives to the differential shared SCSI bus by adding additional DWZZB-AA/DS-TZ89N-TA combinations.

Note

Ensure that there is no conflict with tape drive, system, and disk SCSI IDs.

8.3 Compaq 20/40 GB DLT Tape Drive

The Compaq 20/40 GB DLT Tape Drive is a Digital Linear Tape (DLT) tabletop cartridge tape drive that can hold up to 40 GB of data per CompacTape IV cartridge using 2:1 compression. It is capable of storing and retrieving data at a rate of up to 10.8 GB per hour (using 2:1 compression).

The Compaq 20/40 GB DLT Tape Drive uses CompacTape III, CompacTape IIIXT, or CompacTape IV media.

It is a narrow, single-ended SCSI device, and uses 50-pin, high-density connectors.

For more information on the Compaq 20/40 GB DLT Tape Drive, see the following Compaq documentation:

- *Compaq DLT User Guide*
- *DLT Tape Drive User Guide Supplement*

The following sections describe how to prepare the Compaq 20/40 GB DLT Tape Drive for shared SCSI bus usage in more detail.


8.3.1 Setting the Compaq 20/40 GB DLT Tape Drive SCSI ID

As with any of the shared SCSI devices, the Compaq 20/40 GB DLT Tape Drive SCSI ID must be set to ensure that no two SCSI devices on the shared SCSI bus have the same SCSI ID.

The Compaq 20/40 GB DLT Tape Drive SCSI ID is set with a push-button counter switch on the rear of the unit (see Figure 8–4). Push the button above the counter to increment the address; push the button below the counter to decrement the address until you have the desired SCSI ID selected.

Only SCSI IDs in the range of 0 through 7 are valid. Ensure that the tape drive SCSI ID does not conflict with the SCSI ID of the host bus adapters (usually 6 and 7) or other devices on this shared SCSI bus.

Figure 8–4: Compaq 20/40 GB DLT Tape Drive Rear Panel



8.3.2 Cabling the Compaq 20/40 GB DLT Tape Drive

The Compaq 20/40 GB DLT Tape Drive is connected to a single-ended segment of the shared SCSI bus. A DWZZB-AA signal converter is required to convert the differential shared SCSI bus to single-ended. Figure 8–5 shows a configuration with a Compaq 20/40 GB DLT Tape Drive on a shared SCSI bus.

To configure the shared SCSI bus for use with a Compaq 20/40 GB DLT Tape Drive, follow these steps:

1. You will need one DWZZB-AA for each shared SCSI bus with a Compaq 20/40 GB DLT Tape Drive.
Ensure that the DWZZB-AA jumpers W1 and W2 are installed to enable the single-ended termination.
Remove the termination from the differential end by removing the five 14-pin SIP resistors.
2. Attach an H885-AA trilink connector or BN21W-0B Y cable to the differential end of the DWZZB-AA.
3. Connect the single-ended end of the DWZZB-AA to the Compaq 20/40 GB DLT Tape Drive with cable part number 199629-002 or 189636-002 (1.8-meter (5.9-foot) cables).
4. Install terminator part number 341102-001 on the other tape drive SCSI connector.
5. Connect the trilink on the DWZZB-AA to another trilink or Y cable on the differential shared SCSI bus with a 328215-00X, BN21K, or BN21L cable. Keep the length of the differential segment below the 25-meter (82-foot) maximum length (cable part number 328215-004 is a 20-meter

(65.6-foot) cable). Ensure that the trilink or Y cable at both ends of the differential segment of the shared SCSI bus is terminated with an HD68 differential terminator such as an H879-AA.

The single-ended SCSI bus may be daisy chained from one single-ended tape drive to another with cable part number 146745-003 or 146776-003 (0.9-meter (2.95-foot) cables) as long as the SCSI bus maximum length of 3 meters (9.8 feet) (fast SCSI) is not exceeded. Ensure that the tape drive on the end of the bus is terminated with terminator part number 341102-001.

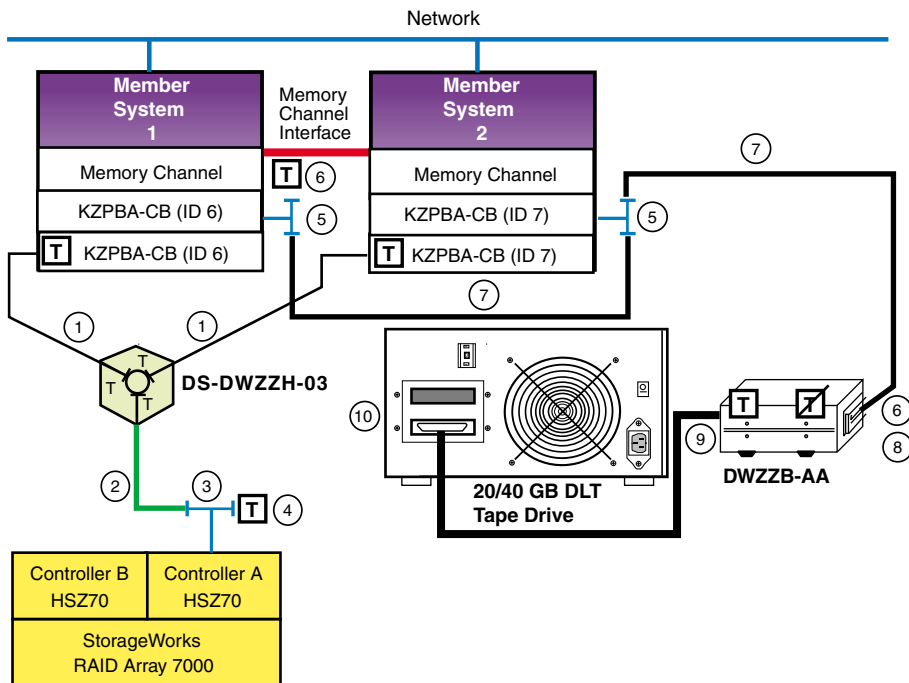
You can add additional shared SCSI buses with Compaq 20/40 GB DLT Tape Drives by adding additional DWZZB-AA/Compaq 20/40 GB DLT Tape Drive combinations.

Notes

Ensure that there is no conflict with tape drive and host bus adapter SCSI IDs.

To achieve system performance capabilities, we recommend placing no more than two Compaq 20/40 GB DLT Tape Drives on a SCSI bus, and also recommend that no shared storage be placed on the same SCSI bus with the tape drive.

Figure 8–5: Cabling a Shared SCSI Bus with a Compaq 20/40 GB DLT Tape Drive



ZK-1604U-AI

Table 8–4 lists the components that are used to create the cluster that is shown in Figure 8–5.

Table 8–4: Hardware Components Used to Create the Configuration Shown in Figure 8–5

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI tralink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c
8	H885-AA tralink connector

Table 8–4: Hardware Components Used to Create the Configuration Shown in Figure 8–5 (cont.)

Callout Number	Description
9	199629-002 or 189636-002 (68-pin high density to 50-pin high density 1.8-meter (5.9-foot) cables)
10	341102-001 50-pin high density terminator

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.4 Compaq 40/80-GB DLT Drive

The Compaq 40/80-DLT Drive is a Digital Linear Tape (DLT) table top cartridge tape drive that can hold up to 80 GB of data per DLT-IV cartridge using 2:1 compression. It can store and retrieve data at a rate of up to 43.2 GB per hour (using 2:1 compression).

The Compaq 40/80-DLT Drive uses DLT-III, DLT-IIIXT, or DLT-IV media.

This tape drive works with either a single-ended or low-voltage differential (LVD) host bus adapter using 68-pin high-density connectors. The single-ended voltages are either 0V or +5V. The low-voltage differential voltages range from -2.5V to +2.5V, a 5.0V differential. It will not work with high voltage differential (-5.0V to +5.0V) host bus adapters. Therefore, you cannot connect the KZPSA-BB or KZPBA-CB host bus adapters directly to the Compaq 40/80-DLT Drive. Instead, you must use a DWZZB-AA signal converter to convert the high voltage differential to single-ended SCSI.

For more information on the Compaq 40/80-DLT Drive, see the *Compaq StorageWorks 40/80 DLT Drive Reference Guide*.

The following sections describe how to prepare the Compaq 40/80-GB DLT Drive in more detail.

8.4.1 Setting the Compaq 40/80-GB DLT Drive SCSI ID

As with any of the shared SCSI devices, the Compaq 40/80-GB DLT Drive SCSI ID must be set to ensure that no two SCSI devices on the shared SCSI bus have the same SCSI ID.

You set the Compaq 40/80-GB DLT Drive SCSI ID with a pushbutton counter switch on the rear of the unit. Push the button above the counter to increment the address; push the button below the counter to decrement the address until you select the desired SCSI ID.

8.4.2 Cabling the Compaq 40/80-GB DLT Drive

The Compaq 40/80-GB DLT Drive is connected to a single-ended segment of the shared SCSI bus. Figure 8–6 shows a configuration with a Compaq 40/80-GB DLT Drive for use on a shared SCSI bus.

To configure the shared SCSI bus for use with a Compaq 40/80-GB DLT Drive, follow these steps:

1. You need one DWZZB-AA for each shared SCSI bus with a Compaq 40/80-GB DLT Drive.
Ensure that the DWZZB-AA jumpers W1 and W2 are installed to enable the single-ended termination.
Remove the termination from the differential end by removing the five 14-pin SIP resistors.
2. Connect an H885-AA trilink connector with attached H879-AA HD68 terminator to the differential end of the DWZZB-AA signal converter.
3. Connect a 328215-00X, BN21K, or BN21L cable between a Y cable or trilink connector on one of the system host bus adapters and the open connector on the trilink on the DWZZB-AA differential end. Keep the total length of the differential segment below the 25-meter (82-foot) maximum length. Ensure that the trilink or Y cable at both ends of the differential segment of the shared SCSI bus is terminated with an HD68 differential terminator such as an H879-AA.
4. Install a 152732-001 LVD terminator on one of the Compaq 40/80-GB DLT Drive SCSI connectors. The LVD terminator works with either a low voltage differential or a single-ended SCSI bus.
5. Connect one of the following HD68-to-HD68 cables between the other Compaq 40/80-GB DLT Drive SCSI connector and the single-ended end of the DWZZB-AA signal converter:
 - 189646-001 (0.9-meter; 2.95-foot) or 189646-002 (1.8-meter; 5.9-foot) cable
 - BN21K-01 or BN21L-01 (1-meter; 3.3-foot) cable
 - BN21K-02 or BN21L-02 (2-meter; 6.6-foot) cable

This single-ended SCSI bus segment is limited to a maximum length of 3 meters (9.8 feet) (fast SCSI).

The single-ended SCSI bus may be daisy chained from one tape drive to another with one of the cables listed in step 5 as long as you do not exceed the single-ended SCSI bus maximum length of 3 meters (9.8 feet). The distance includes the length of both cables on the single-ended bus segment. Ensure that the tape drive on the end of the bus is terminated with an LVD terminator.

Ensure that SCSI IDs for the tape drive and host bus adapter do not conflict.

To achieve system performance capabilities, we recommend that you place no more than two Compaq 40/80-GB DLT Drives on a SCSI bus, and that you place no shared storage on the same SCSI bus with the tape drive.

Figure 8–6: Cabling a Shared SCSI Bus with a Compaq 40/80-GB DLT Drive

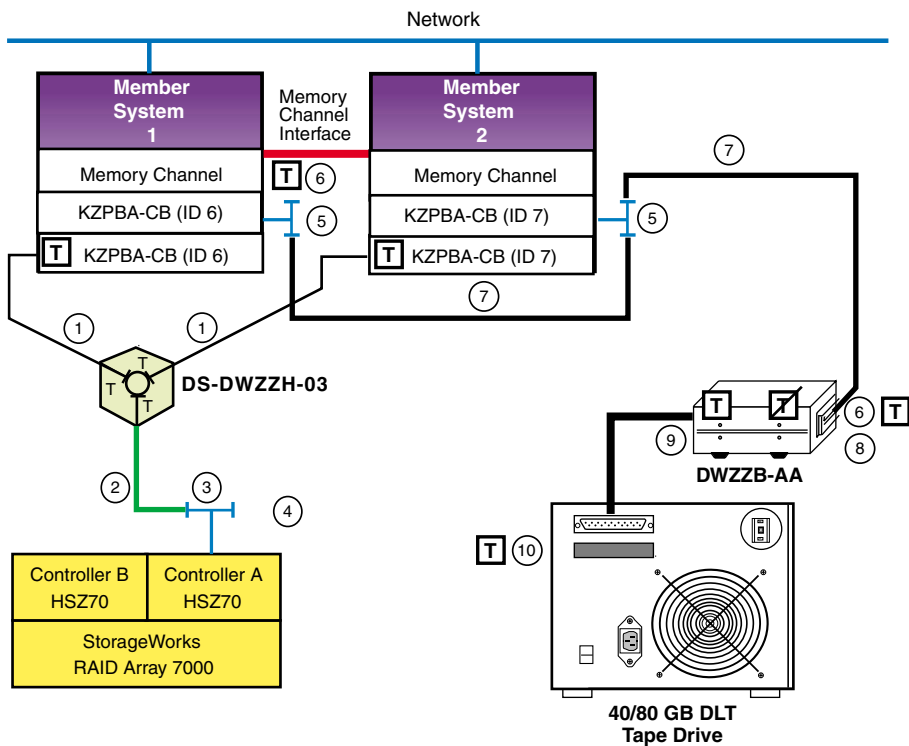


Table 8–5 lists the components that are used to create the cluster that is shown in Figure 8–6.

Table 8–5: Hardware Components in the Configuration in Figure 8–6

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI tralink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA HD68 terminator

Table 8–5: Hardware Components in the Configuration in Figure 8–6 (cont.)

Callout Number	Description
7	328215-00X, BN21K, or BN21L HD68 to HD68 cable ^c
8	H885-AA tralink connector
9	189646-001 (0.9 meter; 2.95-foot cable) or 189646-002 (1.8 meter; 5.9-foot cable) ^d BN21K-01 or BN21L-01 (1-meter; 3.3-foot cable) ^d BN21K-02 or BN21L-02 (2-meter; 6.6-foot cable) ^d
10	152732-001 LVD terminator

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

^d The maximum length of this cable, or of two cables if two tape drives are daisy-chained, is 3 meters (9.8 feet).

8.5 Preparing the TZ885 for Shared SCSI Usage

The TZ885 Digital Linear Tape subsystems combine a cartridge tape drive (TZ88) and an automatic cartridge loader. The TZ885 uses a removable magazine.

The TZ885 uses a five-cartridge (CompacTape IV) minitape library (magazine) with a 200-GB capacity (compressed). It is capable of reading/writing at approximately 10.8 GB per hour.

As with any of the shared SCSI devices, the TZ885 SCSI IDs must be set to ensure that no two SCSI devices on the shared SCSI bus have the same SCSI ID.

The following sections describe preparing the TZ885 in more detail.

8.5.1 Setting the TZ885 SCSI ID

To set the TZ885 SCSI ID from the Operators Control Panel (OCP), follow these steps:

1. Press and hold the Display Mode push button (for about 5 seconds) until the SCSI ID SEL message is displayed:

```
SCSI ID SEL  
SCSI ID 0
```

2. Press the Select push button until you see the desired SCSI ID number in the display.
3. Press the Display Mode push button again.

4. Issue a bus reset or turn the minilibrary power off and on again to cause the drive to recognize the new SCSI ID.

8.5.2 Cabling the TZ885 Tape Drive

The TZ885 is connected to a single-ended segment of the shared SCSI bus. It is connected to a differential portion of the shared SCSI bus with a DWZZA-AA or DWZZB-AA. Figure 8–7 shows a configuration of a TZ885 for use on a shared SCSI bus. To configure the shared SCSI bus for use with a TZ885, follow these steps:

1. You will need one DWZZA-AA or DWZZB-AA for each TZ885 tape drive. Ensure that the DWZZA jumper J2 or DWZZB jumpers W1 and W2 are installed to enable the single-ended termination.
Remove the termination from the differential end by removing the five 14-pin SIP resistors.
2. Attach a tralink connector to the differential end of the DWZZA or DWZZB.
3. Connect the single-ended end of a DWZZA to the TZ885 with a BC19J cable.
Connect the single-ended end of a DWZZB to the TZ885 with a BN21M cable.
4. Install an H8574-A or H8890-AA terminator on the other TZ885 SCSI connector.
5. Connect a tralink or Y cable to the differential shared SCSI bus with BN21K or BN21L cables. Ensure that the tralink or Y cable at the end of the bus is terminated with an H879-AA terminator.

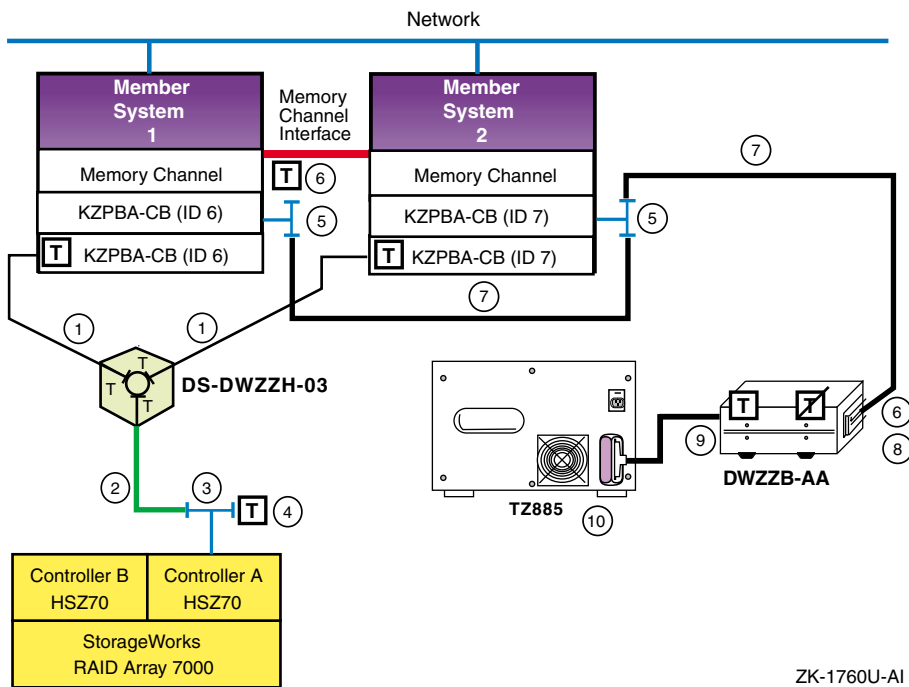
The single-ended SCSI bus may be daisy chained from one single-ended tape drive to another with BC19J cables as long as the SCSI bus maximum length is not exceeded. Ensure that the tape drive on the end of the bus is terminated with a H8574-A or H8890-AA terminator.

You can add additional TZ885 tape drives to the differential shared SCSI bus by adding additional DWZZA or DWZZB/TZ885 combinations.

Note

Ensure that there is no conflict with tape drive and host bus adapter SCSI IDs.

Figure 8–7: Cabling a Shared SCSI Bus with a TZ885



ZK-1760U-AI

Table 8–6 lists the components that are used to create the cluster that is shown in Figure 8–7.

Table 8–6: Hardware Components Used to Create the Configuration Shown in Figure 8–7

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI trilink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

Table 8–6: Hardware Components Used to Create the Configuration Shown in Figure 8–7 (cont.)

Callout Number	Description
8	H885-AA tralink connector
9	BN21M cable
10	H8574-A terminator

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.6 Preparing the TZ887 for Shared SCSI Bus Usage

The TZ887 Digital Linear Tape (DLT) MiniLibrary combines a cartridge tape drive (TZ88) and an automatic cartridge loader. It uses a seven-cartridge (CompacTape IV) removable magazine with a total capacity of nearly 280 GB compressed. It is capable of reading and writing at approximately 10.8 GB per hour.


As with any of the shared SCSI devices, the TZ887 SCSI IDs must be set to ensure that no two SCSI devices on the shared SCSI bus have the same SCSI ID.

The following sections describe how to prepare the TZ887 in more detail.

8.6.1 Setting the TZ887 SCSI ID

The TZ887 SCSI ID is set with a push-button counter switch on the rear of the unit (see Figure 8–8). Push the button above the counter to increment the address; push the button below the counter to decrement the address until you have the desired SCSI ID selected.

Figure 8–8: TZ887 DLT MiniLibrary Rear Panel



ZK-1461U-AI

8.6.2 Cabling the TZ887 Tape Drive

The TZ887 is connected to a single-ended segment of the shared SCSI bus. It is connected to a differential portion of the shared SCSI bus with a DWZZB-AA. Figure 8–9 shows a configuration with a TZ887 for use on a shared SCSI bus.

To configure the shared SCSI bus for use with a TZ887, follow these steps:

1. You will need one DWZZB-AA for each shared SCSI bus with a TZ887 tape drive.
Ensure that the DWZZB-AA jumpers W1 and W2 are installed to enable the single-ended termination.
Remove the termination from the differential end by removing the five 14-pin SIP resistors.
2. Attach an H885-AA trilink connector to the differential end of the DWZZB-AA.
3. Connect the single-ended end of the DWZZB-AA to the TZ887 with a BN21M cable.
4. Install an H8574-A or H8890-AA terminator on the other TZ887 SCSI connector.
5. Connect the trilink on the DWZZB-AA to another trilink or Y cable on the differential shared SCSI bus with BN21K or BN21L cables.
Ensure that the trilink or Y cable at both ends of the shared SCSI bus is terminated with an H879-AA terminator.

The single-ended SCSI bus may be daisy chained from one single-ended tape drive to another with BC19J cables, as long as the SCSI bus maximum

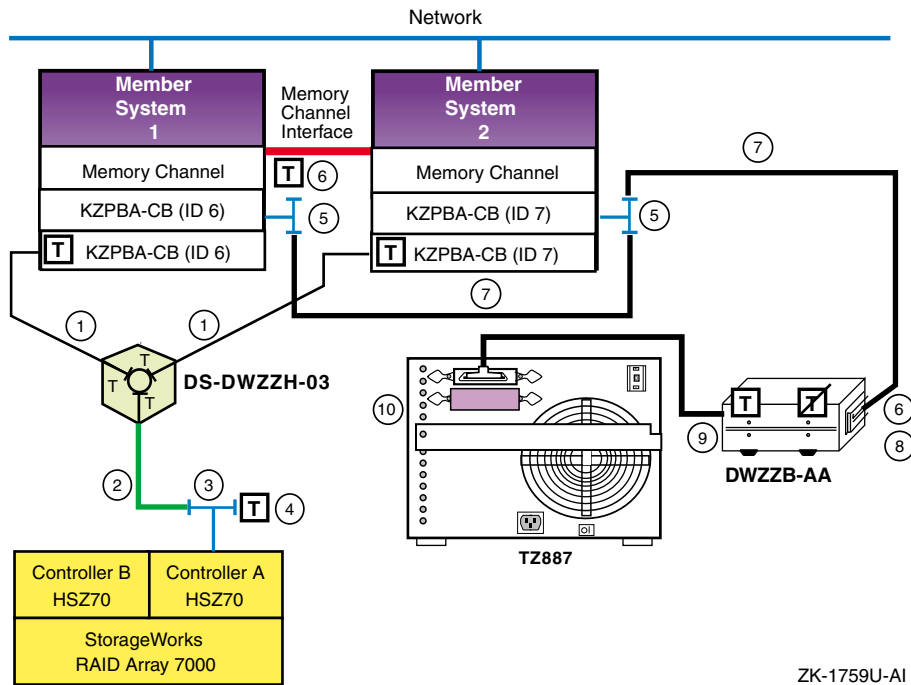
length is not exceeded and there are sufficient SCSI IDs available. Ensure that the tape drive on the end of the bus is terminated with an H8574-A or H8890-AA terminator.

You can add additional shared SCSI buses with TZ887 tape drives by adding additional DWZZB-AA/TZ887 combinations.

Note

Ensure that there is no conflict with tape drive and host bus adapter SCSI IDs.

Figure 8–9: Cabling a Shared SCSI Bus with a TZ887



ZK-1759U-AI

Table 8–6 lists the components that are used to create the cluster that is shown in Figure 8–9.

8.7 Preparing the TL891 and TL892 DLT MiniLibraries for Shared SCSI Usage

Note

To achieve system performance capabilities, we recommend placing no more than two TZ89 drives on a SCSI bus, and also recommend that no shared storage be placed on the same SCSI bus with a tape library.

The TL891 and TL892 MiniLibraries use one (TL891) or two (TL892) TZ89N-AV differential tape drives and a robotics controller, which access cartridges in a 10-cartridge magazine.

Each tape drive present, and the robotics controller, have individual SCSI IDs.

Six 68-pin, high-density SCSI connectors are located on the back of the MiniLibrary; two SCSI connectors for each drive and two for the robotics controller. The TL891 uses a 30-centimeter (11.8-inch) SCSI bus jumper cable (part of the TL891 package) to place the robotics controller and tape drive on the same SCSI bus. When upgrading to the TL892, you can place the second drive on the same SCSI bus (another 30-centimeter (11.8-inch) SCSI bus jumper cable is supplied with the DS-TL892-UA upgrade kit) or place it on its own SCSI bus.

The following sections describe how to prepare the TL891 and TL892 in more detail.

8.7.1 Setting the TL891 or TL892 SCSI ID

The control panel on the front of the TL891 and TL892 MiniLibraries is used to display power-on self-test (POST) status, display messages, and to set up MiniLibrary functions.

When power is first applied to a MiniLibrary, a series of POST diagnostics are performed. During POST execution, the MiniLibrary model number, current date and time, firmware revision, and the status of each test is displayed on the control panel.

After the POST diagnostics have completed, the default screen is shown:

```
DLT0 Idle
DLT1 Idle
Loader Idle
0> _ _ _ _ _ _ _ _ _ _ <9
```

The first and second lines of the default screen show the status of the two drives (if present). The third line shows the status of the library robotics, and the fourth line is a map of the magazine, with the numbers from 0 through 9 representing the cartridge slots. Rectangles on this line indicate cartridges that are present in the corresponding slot of the magazine.

For example, this fourth line (0> X X _ _ _ _ _ _ <9, where X represents rectangles) indicates that cartridges are installed in slots 0 and 1.

Note

There are no switches for setting a mechanical SCSI ID for the tape drives. The SCSI IDs default to 5. The MiniLibrary sets the electronic SCSI ID very quickly, before any device can probe the MiniLibrary, so the lack of a mechanical SCSI ID does not cause any problems on the SCSI bus.

To set the SCSI ID, follow these steps:

1. From the Default Screen, press the Enter button to enter the Menu Mode, displaying the Main Menu.

Note

When you enter the Menu Mode, the Ready light goes out, which indicates that the module is off line, and all media changer commands from the host return a SCSI not ready status until you exit the Menu Mode and the Ready light comes on again.

2. Press the down arrow button until the Configure Menu item is selected, then press the Enter button to display the Configure submenu.

Note

The control panel up and down arrows have an auto-repeat feature. When you press either button for more than one-half second, the control panel behaves as if you were pressing the button about four times per second. The effect stops when you release the button.

3. Press the down arrow button until the Set SCSI item is selected and press the Enter button.

4. Select the tape drive (DLT0 Bus ID: or DLT1 Bus ID:) or library robotics (LIB Bus ID:) whose SCSI bus ID you want to change. The default SCSI IDs are as follows:
 - Lib Bus ID: 0
 - DLT0 Bus ID: 4
 - DLT1 Bus ID: 5Use the up or down arrow button to select the item whose SCSI ID you want to change. Press the Enter button.
5. Use the up or down arrow button to scroll through the possible SCSI ID settings. Press the Enter button when the desired SCSI ID is displayed.
6. Repeat steps 4 and 5 to set other SCSI bus IDs as necessary.
7. Press the Escape button repeatedly until the default menu is displayed.

8.7.2 Cabling the TL891 or TL892 MiniLibraries

There are six 68-pin, high-density SCSI connectors on the back of the TL891. The two leftmost connectors are for the library robotics controller. The middle two are for tape drive 1. The two on the right are for tape drive 2 (if the TL892 upgrade has been installed).

Note

The tape drive SCSI connectors are labeled DLT1 (tape drive 1) and DLT2 (tape drive 2). The control panel designation for the drives is DLT0 (tape drive 1) and DLT1 (tape drive 2).

The default for the DLT MiniLibrary TL891 is to place the robotics controller and tape drive 1 on the same SCSI bus. A 30-centimeter (11.8-inch) SCSI jumper cable is provided with the unit. Plug this cable into the second connector (from the left) and the third connector. If the MiniLibrary has been upgraded to two drives, place the second drive on the same SCSI bus with another 30-centimeter (11.8-inch) SCSI bus jumper cable, or place it on its own SCSI bus.

Note

To achieve system performance capabilities, we recommend placing no more than two TZ89 tape drives on a SCSI bus.

The internal cabling of the TL891 and TL892 is too long to allow external termination with a tralink/H879-AA combination. Therefore, the TL891 or TL892 must be the last device on the shared SCSI bus. They may not be removed from the shared

SCSI bus without stopping all ASE services that generate activity on the bus.

For this reason, we recommend that tape devices be placed on separate shared SCSI buses, and that there be no storage devices on the SCSI bus.

The cabling depends on whether or not there are one or two drives, and for the two-drive configuration, if each drive is on a separate SCSI bus.

Note

It is assumed that the library robotics controller is on the same SCSI bus as tape drive 1.

To connect the library robotics and one drive to a single shared SCSI bus, follow these steps:

1. Connect a BN21K or BN21L between the last tralink connector on the bus to the leftmost connector (as viewed from the rear) of the TL891.
2. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost robotics connector (second connector from the left) and the left DLT1 connector (the third connector from the left).
3. Install an H879-AA terminator on the right DLT1 connector (the fourth connector from the left).

To connect the drive robotics and two drives to a single shared SCSI bus, follow these steps:

1. Connect a BN21K or BN21L between the last tralink connector on the bus to the leftmost connector (as viewed from the rear) of the TL892.
2. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost robotics connector (the second connector from the left) and the left DLT1 connector (the third connector from the left).
3. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost DLT1 connector (the fourth connector from the left) and the left DLT2 connector (the fifth connector from the left).
4. Install an H879-AA terminator on the right DLT2 connector (the rightmost connector).

To connect the drive robotics and one drive to one shared SCSI bus and the second drive to a second shared SCSI bus, follow these steps:

1. Connect a BN21K or BN21L between the last tralink connector on one shared SCSI bus to the leftmost connector (as viewed from the rear) of the TL892.
2. Connect a BN21K or BN21L between the last tralink connector on the second shared SCSI bus to the left DLT2 connector (the fifth connector from the left).
3. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost robotics connector (the second connector from the left) and the left DLT1 connector (the third connector from the left).
4. Install an H879-AA terminator on the right DLT1 connector (the fourth connector from the left) and install another H879-AA terminator on the right DLT2 connector (the rightmost connector).

Figure 8–10 shows an example of a TruCluster Server cluster with a TL892 connected to two shared SCSI buses.

8.8 Preparing the TL890 DLT MiniLibrary Expansion Unit

The topics in this section provide information on preparing the TL890 DLT MiniLibrary expansion unit with the TL891 and TL892 DLT MiniLibraries for use on a shared SCSI bus.

Note

To achieve system performance capabilities, we recommend placing no more than two TZ89 drives on a SCSI bus, and also recommend that no shared storage be placed on the same SCSI bus with a tape library.

8.8.1 TL890 DLT MiniLibrary Expansion Unit Hardware

The TL890 expansion unit is installed above the TL891/TL892 DLT MiniLibrary base units in a SW500, SW800, or RETMA cabinet. The expansion unit integrates the robotics in the individual modules into a single, coordinated library robotics system. The TL890 assumes control of the media, maintaining an inventory of all media present in the system, and controls movement of all media. The tape cartridges can move freely between the expansion unit and any of the base modules via the system's robotically controlled pass-through mechanism. The pass-through mechanism is attached to the back of the expansion unit and each of the base modules.

For each TL891/TL892 base module beyond the first module, the pass-through mechanism must be extended by 17.78 centimeters (7 inches), which is the height of each module with a DS-TL800-AA pass-through mechanism extension. A 17.78-centimeter (7-inch) gap may be left between base modules (providing there is sufficient space), but additional pass-through mechanism extensions must be used.

For complete hardware installation instructions, see the *DLT MiniLibrary (TL890) Expansion Unit User's Guide*.

The combination of the TL890 expansion unit and the TL891/TL892 MiniLibrary modules is referred to as a DLT MiniLibrary for the remainder of this discussion.

8.8.2 Preparing the DLT MiniLibraries for Shared SCSI Bus Usage

The following sections describe how to prepare the DLT MiniLibraries in more detail. The descriptions are based on the assumption that the expansion unit, base modules, and pass-through and motor mechanisms have been installed.

8.8.2.1 Cabling the DLT MiniLibraries

You must make the following connections to render the DLT MiniLibrary system operational:

- Expansion unit to the motor mechanism: The motor mechanism cable is about 1 meter (3.3 feet) long and has a DB-15 connector on each end. Connect it between the connector labeled Motor on the expansion unit to the motor on the pass-through mechanism.

Note

This cable is not shown in Figure 8–11 because the pass-through mechanism is not shown in the figure.

- Robotics control cables from each base module to the expansion unit: These cables have a DB-9 male connector on one end and a DB-9 female connector on the other end. Connect the male end to the Expansion Unit Interface connector on the base module and the female end to any Expansion Modules connector on the expansion unit.

Note

It does not matter which interface connector a base module is connected to.

- SCSI bus connection to the expansion unit robotics: Connect the shared SCSI bus that will control the robotics to one of the SCSI connectors on the expansion unit with a BN21K (or BN21L) cable. Terminate the SCSI bus with an H879-AA terminator on the other expansion unit SCSI connector.
- SCSI bus connection to each of the base module tape drives: Connect a shared SCSI bus to one of the DLT1 or DLT2 SCSI connectors on each of the base modules with BN21K (or BN21L) cables. Terminate the other DLT1 or DLT2 SCSI bus connection with an H879-AA terminator.

You can daisy chain between DLT1 and DLT2 (if present) with a 30-centimeter (11.8-inch) SCSI bus jumper (supplied with the TL891). Terminate the SCSI bus at the tape drive on the end of the shared SCSI bus with an H879-AA terminator.

Notes

Do not connect a SCSI bus to the SCSI connectors for the library connectors on the base modules.

We recommend that no more than two TZ89 tape drives be on a SCSI bus.

Figure 8–11 shows a MiniLibrary configuration with two TL892 DLT MiniLibraries and a TL890 DLT MiniLibrary expansion unit. The TL890 library robotics is on one shared SCSI bus, and the two TZ89 tape drives in each TL892 are on separate, shared SCSI buses. The pass-through mechanism and cable to the library robotics motor are not shown in this figure.

Figure 8–11: TL890 and TL892 DLT MiniLibraries on Shared SCSI Buses

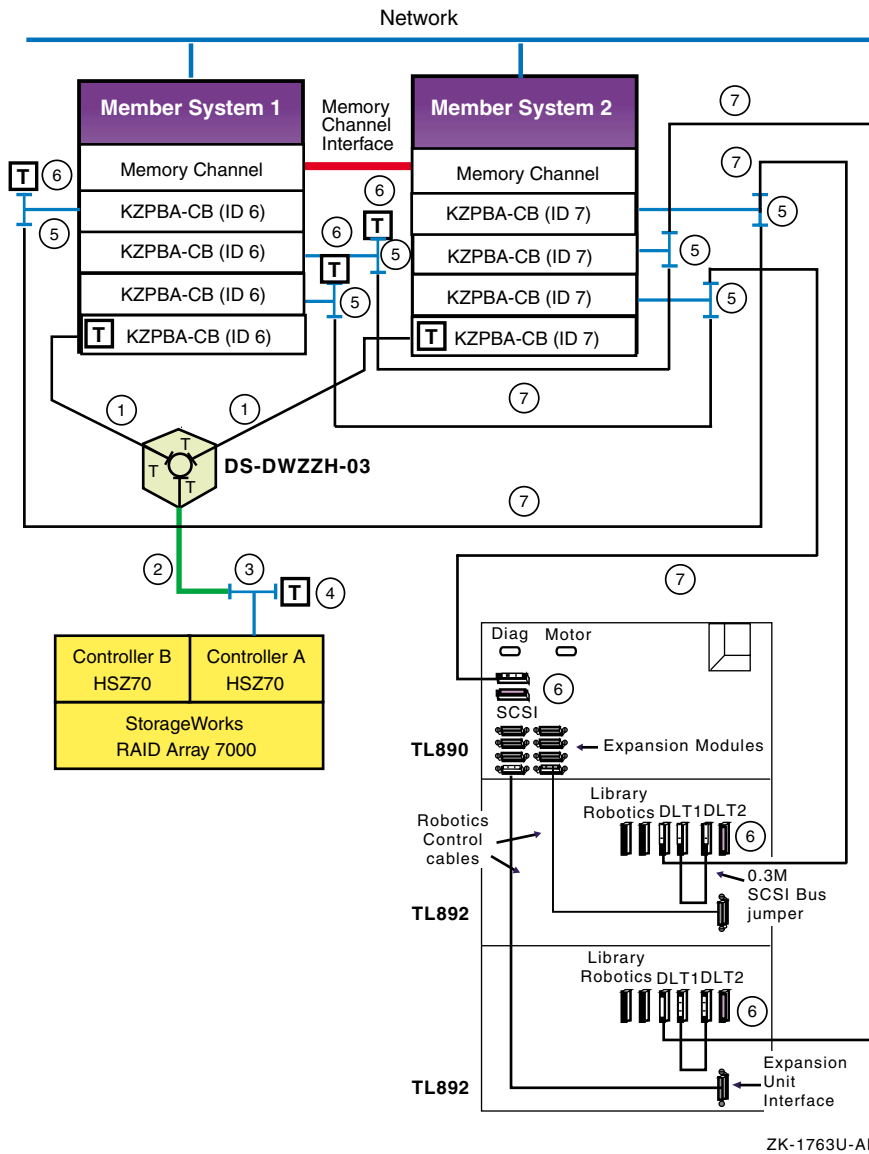


Table 8–8 lists the components that are used to create the cluster that is shown in Figure 8–11.

Table 8–8: Hardware Components Used to Create the Configuration Shown in Figure 8–11

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI trilink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.8.2.2 Configuring a Base Module as a Slave

The TL891/TL892 base modules are shipped configured as standalone systems. When they are used in conjunction with the TL890 DLT MiniLibrary expansion unit, the expansion unit must control the robotics of each of the base modules. Therefore, the base modules must be configured as a slave to the expansion unit.

After the hardware and cables are installed, but before you power up the expansion unit in a MiniLibrary system for the first time, you must reconfigure each of the base modules in the system as a slave. The expansion unit will not have control over the base module robotics when you power up the MiniLibrary system if you do not reconfigure the base modules as a slave.

To reconfigure a TL891/TL892 base module as a slave to the TL890 DLT MiniLibrary expansion unit, perform the following procedure on each base module in the system:

1. Turn on the power switch on the TL891/TL892 base module to be reconfigured.

Note

Do not power on the expansion unit. Leave it powered off until all base modules have been reconfigured as slaves.

After a series of power-on self-tests have executed, the default screen will be displayed on the base module control panel:

```
DLT0 Idle
DLT1 Idle
Loader Idle
0> _ _ _ _ _ _ _ _ _ _ <9
```

The default screen shows the state of the tape drives, loader, and number of cartridges present for this base module. A rectangle in place of the underscore indicates that a cartridge is present in that location.

2. Press the Enter button to enter the Menu Mode, displaying the Main Menu.
3. Press the down arrow button until the Configure Menu item is selected, then press the Enter button.

Note

The control panel up and down arrows have an auto-repeat feature. When you press either button for more than half a second, the control panel behaves as if you were pressing the button about four times per second. The effect stops when you release the button.

4. Press the down arrow button until the Set Special Config menu is selected and press the Enter button.
5. Press the down arrow button repeatedly until the Alternate Config item is selected and press the Enter button.
6. Press the down arrow button to change the alternate configuration from the default (Standalone) to Slave. Press the Enter button.
7. After the selection stops flashing and the control panel indicates that the change is not effective until a reboot, press the Enter button.
8. When the Special Configuration menu reappears, turn the power switch off and then on to cycle the power. The base module is now reconfigured as a slave to the TL890 expansion unit.
9. Repeat the steps for each TL891/TL892 base module present that is to be a slave to the TL890 expansion unit.

8.8.2.3 Powering Up the DLT MiniLibrary

When turning on power to the DLT MiniLibrary, power must be applied to the TL890 expansion unit simultaneously or after power is applied to the TL891/TL892 base modules. If the expansion unit is powered on first, its

inventory of modules may be incorrect and the contents of some or all of the modules will be inaccessible to the system and to the host.

When the expansion unit comes up, it will communicate with each base module through the expansion unit interface and inventory the number of base modules, tape drives, and cartridges present in each base module. After the MiniLibrary configuration has been determined, the expansion unit will communicate with each base module and indicate to the base module which cartridge group that base module contains. The cartridges slots are numbered by the expansion unit as follows:

- Expansion unit: 0 through 15
- Top TL891/TL892: 16 through 25
- Middle TL891/TL892: 26 through 35
- Bottom TL891/TL892: 36 through 45

When all initialization communication between the expansion module and each base module has completed, the base modules will display their cartridge numbers according to the remapped cartridge inventory.

For instance, the middle base module default screen would be displayed as follows:

```
DLT2 Idle
DLT3 Idle
Loader Idle
26> _ _ _ _ _ <35
```

8.8.2.4 Setting the TL890/TL891/TL892 SCSI ID

After the base modules have been reconfigured as slaves, each base module control panel still provides tape drive status and error information, but all control functions are carried out from the expansion unit control panel. This includes setting the SCSI ID for each of the tape drives present.

To set the SCSI IDs for the tape drives in a MiniLibrary configured with TL890/TL891/TL892 hardware, follow these steps:

1. Apply power to the MiniLibrary, ensuring that you power up the expansion unit after or at the same time as the base modules.
2. Wait until power-on self-tests (POST) have terminated and the expansion unit and each base module display the default screen.
3. At the expansion unit control panel, press the Enter button to display the Main Menu.

4. Press the down arrow button until the Configure Menu item is selected, and then press the Enter button to display the Configure submenu.
5. Press the down arrow button until the Set SCSI item is selected and press the Enter button.
6. Press the up or down arrow button to select the appropriate tape drive (DLT0 Bus ID:, DLT1 Bus ID:, DLT2 Bus ID:, and so on) or library robotics (Library Bus ID:) for which you want to change the SCSI bus ID. Assuming that each base module has two tape drives, the top base module contains DLT0 and DLT1. The next base module down contains DLT2 and DLT3. The bottom base module contains DLT4 and DLT5. The default SCSI IDs, after being reconfigured by the expansion unit, are as follows:
 - Library Bus ID: 0
 - DLT0 Bus ID: 1
 - DLT1 Bus ID: 2
 - DLT2 Bus ID: 3
 - DLT3 Bus ID: 4
 - DLT4 Bus ID: 5
 - DLT5 Bus ID: 6
7. Press Enter when you have the item selected for which you want to change the SCSI ID.
8. Use the up and down arrows to select the desired SCSI ID. Press the Enter button to save the new selection.
9. Press the Escape button once to return to the Set SCSI submenu to select another tape drive or the library robotics, and then repeat steps 6, 7, and 8 to set the SCSI ID.
10. If there are other items that you want to configure, press the Escape button until the Configure submenu is displayed, then select the item to be configured. Repeat this procedure for each item that you want to configure.
11. If there are no more items to be configured, press the Escape button until the Default window is displayed.

8.9 Preparing the TL894 DLT Automated Tape Library for Shared SCSI Bus Usage

The topics in this section provide information on preparing the TL894 DLT automated tape library for use on a shared SCSI bus in a TruCluster Server cluster.

Note

To achieve system performance capabilities, we recommend placing no more than two TZ89 drives on a SCSI bus segment. We also recommend that storage be placed on shared SCSI buses that do not have tape drives.

The TL894 midrange automated DLT library contains a robotics controller and four differential TZ89 tape drives.

The following sections describe how to prepare the TL894 in more detail.

8.9.1 TL894 Robotic Controller Required Firmware

Robotic firmware Version S2.20 is the minimum firmware revision that is supported in a TruCluster Server cluster. For information on upgrading the robotic firmware, see the Flash Download section of the *TL81X/TL894 Automated Tape Library for DLT Cartridges Diagnostic Software User's Manual*.

8.9.2 Setting TL894 Robotics Controller and Tape Drive SCSI IDs

The robotics controller, and each tape drive must have the SCSI ID set (unless the default is sufficient). Table 8–9 lists the default SCSI IDs.

Table 8–9: TL894 Default SCSI ID Settings

SCSI Device	SCSI Address
Robotics Controller	0
Tape Drive 0	2
Tape Drive 1	3
Tape Drive 2	4
Tape Drive 3	5

To set the SCSI ID for the TL894 robotics controller, follow these steps:

1. Press and release the Control Panel STANDBY button and verify that the SDA (Status Display Area) shows System Off-line.

2. Press and release SELECT to enter the menu mode.
3. Verify that the following information is displayed in the SDA:

```
Menu:
Configuration:
```
4. Press and release SELECT to choose the Configuration menu.
5. Verify that the following information is displayed in the SDA:

```
Menu: Configuration
Inquiry
```
6. Press and release the up or down arrow buttons to locate the SCSI Address submenu, and verify that the following information is displayed in the SDA:

```
Menu: Configuration
SCSI Address ..
```
7. Press and release the SELECT button to choose the SCSI Address submenu and verify that the following information is displayed in the SDA:

```
Menu: Configuration
Robotics
```
8. Press and release the SELECT button to choose the Robotics submenu and verify that the following information is displayed in the SDA:

```
Menu: SCSI Address
SCSI ID 0
```
9. Use the up and down arrow buttons to select the desired SCSI ID for the robotics controller.
10. When the desired SCSI ID is displayed on line 2, press and release the SELECT button.
11. Press and release the up or down button to clear the resulting display from the command.
12. Press and release the up or down button and the SELECT button simultaneously, and verify that System On-line or System Off-line is displayed in the SDA.

To set the SCSI ID for each tape drive if the desired SCSI IDs are different from those shown in Table 8–9, follow these steps:

1. Press and release the Control Panel STANDBY button and verify that the SDA (Status Display Area) shows System Off-line.
2. Press and release SELECT to enter the menu mode.
3. Verify that the following information is displayed in the SDA:

Menu:
Configuration:

4. Press and release SELECT to choose the Configuration menu.
5. Verify that the following information is displayed in the SDA:

Menu: Configuration
SCSI Address

6. Press and release the SELECT button again to choose SCSI Address and verify that the following information is shown in the SDA:

Menu: SCSI Address
Robotics

7. Use the down arrow button to bypass the Robotics submenu and verify that the following information is shown in the SDA:

Menu: SCSI Address
Drive 0

8. Use the up and down arrow buttons to select the drive number to set or change.
9. When you have the proper drive number displayed on line 2, press and release the SELECT button and verify that the following information is shown in the SDA:

Menu: Drive 0
SCSI ID 0

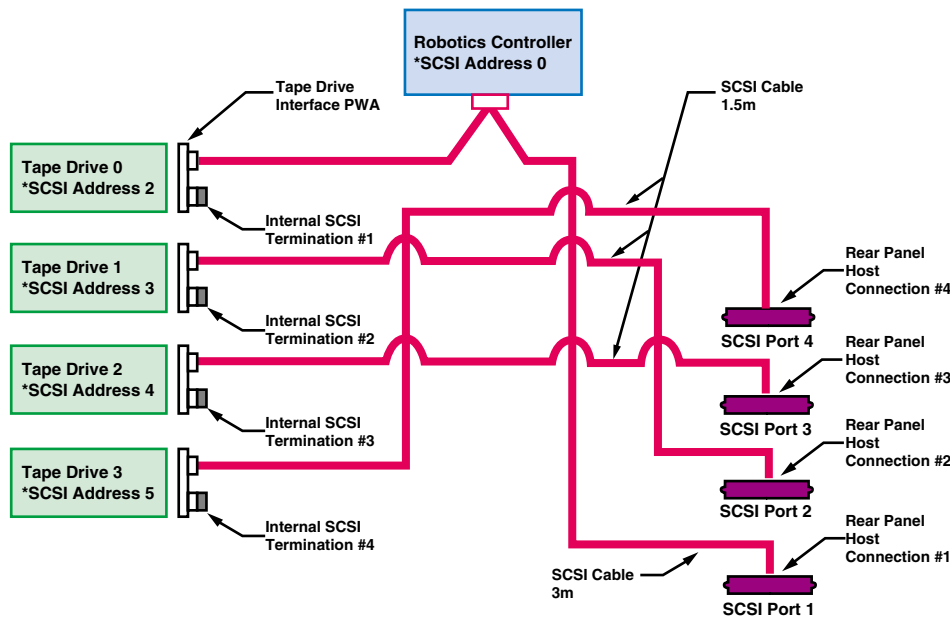
10. Use the up and down arrow buttons to select the desired SCSI ID for the selected drive.
11. When the desired SCSI ID is displayed on line 2, press and release the SELECT button.
12. Repeat steps 8 through 11 to set or change all other tape drive SCSI IDs.
13. Press and release the up or down button to clear the resulting display from the command.
14. Press and release the up or down button and the SELECT button simultaneously and verify that System On-line or System Off-line is displayed in the SDA.

8.9.3 TL894 Tape Library Internal Cabling

The default internal cabling configuration for the TL894 tape library has the robotics controller and top drive (drive 0) on SCSI bus port 1. Drive 1 is on SCSI bus port 2, drive 2 is on SCSI port 3, and drive 3 is on SCSI bus port 4. A terminator (part number 0415619) is connected to each of the drives to provide termination at that end of the SCSI bus.

This configuration, which is called the four-bus configuration, is shown in Figure 8–12. In this configuration, each of the tape drives, except SCSI bus drive 0 and the robotics controller, requires a SCSI address on a separate SCSI bus. The robotics controller and drive 0 use two SCSI IDs on their SCSI bus.

Figure 8–12: TL894 Tape Library Four-Bus Configuration



* - Indicates the "default" SCSI ID of the installed devices

ZK-1324U-AI

You can reconfigure the tape drives and robotics controller in a two-bus configuration by using the SCSI jumper cable (part number 6210567) that is supplied in the accessories kit that is shipped with each TL894 unit. Remove the terminator from one drive and remove the internal SCSI cable from the other drive to be daisy chained. Use the SCSI jumper cable to connect the two drives and place them on the same SCSI bus.

Notes

We recommend that you not place more than two TZ89 tape drives on any one SCSI bus in these tape libraries. We also recommend that storage be placed on shared SCSI buses that do not have tape drives.

Therefore, we recommend that you not reconfigure the TL894 tape library into the one-bus configuration.

Appendix B of the *TL81X/TL894 Automated Tape Library for DLT Cartridges Facilities Planning and Installation Guide* provides figures showing various bus configurations. In these figures, the configuration changes have been made by removing the terminators from both drives, installing the SCSI bus jumper cable on the drive connectors vacated by the terminators, then installing an HD68 SCSI bus terminator on the SCSI bus port connector on the cabinet exterior.

Doing this is not wrong, but by reconfiguring in this manner, the length of the SCSI bus is increased by 1.5 meters (4.9 feet), and may cause problems if SCSI bus length is of concern.

8.9.4 Connecting the TL894 Tape Library to the Shared SCSI Bus

The TL894 tape libraries have up to 3 meters (9.8 feet) of internal SCSI cabling per SCSI bus. Because of the internal SCSI cable lengths, you cannot use a tralink connector or Y cable to terminate the SCSI bus external to the library as is done with other devices on the shared SCSI bus. Each SCSI bus must be terminated internal to the tape library, at the tape drive itself with the installed SCSI terminators. Therefore, TruCluster Server clusters using the TL894 tape library must ensure that the tape library is on the end of the shared SCSI bus.

In a TruCluster Server cluster with a TL894 tape library, the member systems and StorageWorks enclosures or RAID subsystems may be isolated from the shared SCSI bus because they use tralink connectors or Y cables. However, the ASE must be shut down to remove a tape loader from the shared bus.

Figure 8–13 shows a sample TruCluster Server cluster using a TL894 tape library. In the sample configuration, the tape library has been connected in the two-bus mode by jumpering tape drive 0 to tape drive 1 and tape drive 2 to tape drive 3 (see Section 8.9.3 and Figure 8–12). The two SCSI buses are left at the default SCSI IDs and terminated at drives 1 and 3 with the installed terminators (part number 0415619).

To add a TL894 to a shared SCSI bus, select the member system or storage device that will be the next to last device on the shared SCSI bus. Connect a BN21K or BN21L cable between the Y cable on that device to the appropriate tape library port.

In Figure 8–13, one bus is connected to port 1 (robotics controller and tape drives 0 and 1) and the other bus is connected to port 3 (tape drives 2 and 3). Ensure that the terminators are present on the tape drives 1 and 3.

Figure 8–13: Shared SCSI Buses with TL894 in Two-Bus Mode

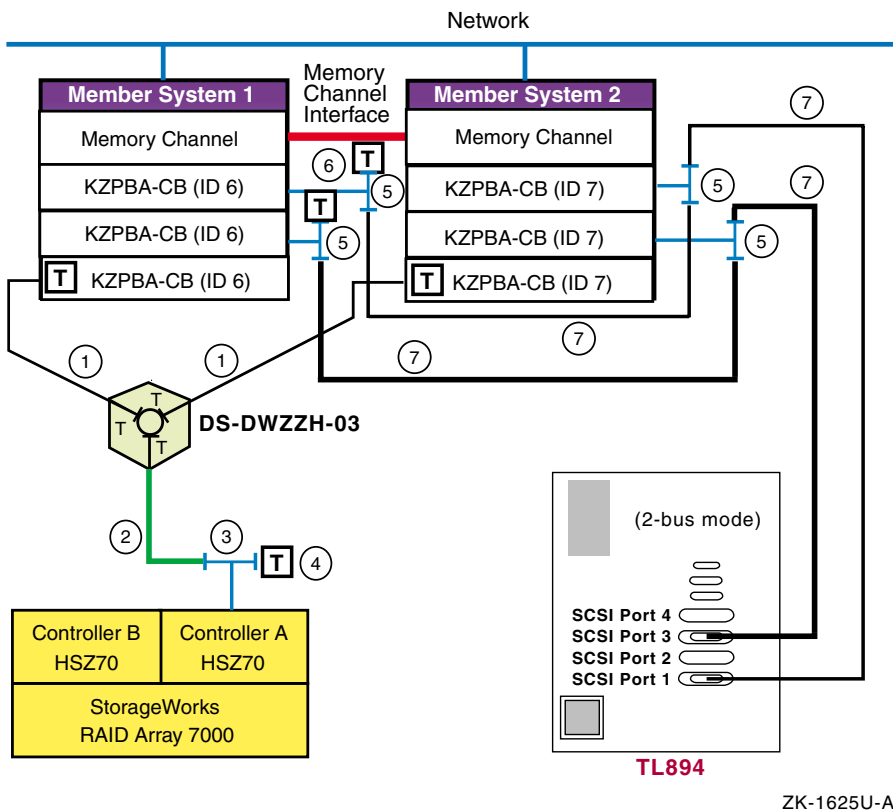


Table 8–10 lists the components that are used to create the cluster that is shown in Figure 8–13.

Table 8–10: Hardware Components Used to Create the Configuration Shown in Figure 8–13

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI trilinear connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable

Table 8–10: Hardware Components Used to Create the Configuration Shown in Figure 8–13 (cont.)

Callout Number	Description
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.10 Preparing the TL895 DLT Automated Tape Library for Shared SCSI Bus Usage

The topics in this section provide information on preparing the TL895 Digital Linear Tape (DLT) automated tape library for use on a shared SCSI bus.

Note

To achieve system performance capabilities, we recommend placing no more than two TZ89 drives on a SCSI bus segment. We also recommend that storage be placed on shared SCSI buses that do not have tape drives. This makes it easier to stop ASE services affecting the SCSI bus that the tape loaders are on.

The DS-TL895-BA automated digital linear tape library consists of five TZ89N-AV tape drives and 100 tape cartridge bins (96 storage bins in a fixed-storage array (FSA) and 4 load port bins). The storage bins hold CompacTape III, CompacTape IIIXT, or CompacTape IV cartridges. The maximum storage capacity of the library is 3500 GB uncompressed, based upon 100 CompacTape IV cartridges at 35 GB each. For more information on the TL895, see the following manuals:

- *TL895 DLT Tape Library Facilities Planning and Installation Guide*
- *TL895 DLT Library Operator's Guide*
- *TL895 DLT Tape Library Diagnostic Software User's Manual*

For more information on upgrading from five to six or seven tape drives, see the *TL895 Drive Upgrade Instructions* manual.

Note

Rotary switches on the library printed circuit board are used to set the library and tape drive SCSI IDs. The SCSI IDs set by these switches are used for the first 20 to 30 seconds after power

is applied, until the electronics is activated and able to set the SCSI IDs electronically.

The physical SCSI IDs should match the SCSI IDs set by the library electronics. Ensure that the SCSI IDs that are set by the rotary switch and from the control panel do not conflict with any SCSI bus controller SCSI ID.

The following sections describe how to prepare the TL895 for use on a shared SCSI bus in more detail.

8.10.1 TL895 Robotic Controller Required Firmware

Robotic firmware version N2.20 is the minimum firmware revision supported in a TruCluster Server cluster. For information on upgrading the robotic firmware, see the Flash Download section of the *TL895 DLT Tape Library Diagnostic Software User's Manual*.

8.10.2 Setting the TL895 Tape Library SCSI IDs

The library and each tape drive must have the SCSI ID set (unless the default is sufficient). Table 8–11 lists the TL895 default SCSI IDs.

Table 8–11: TL895 Default SCSI ID Settings

SCSI Device	SCSI ID
Library	0
Drive 0	1
Drive 1	2
Drive 2	3
Drive 3	4
Drive 4	5
Drive 5	1
Drive 6	2

The SCSI IDs must be set mechanically by the rotary switches, and electronically from the control panel. After you have set the SCSI IDs from the switches, power up the library and electronically set the SCSI IDs.

To electronically set the SCSI ID for the TL895 library and tape drives, follow these steps:

1. At the control panel, press the Operator tab.

2. On the Enter Password screen, enter the operator password. The default operator password is 1234. The lock icon is unlocked and shows an O to indicate that you have operator-level security clearance.
3. On the Operator screen, press the Configure Library button. The Configure Library screen displays the current library configuration.

Note

You can configure the library model number, number of storage bins, number of drives, library SCSI ID, and tape drive SCSI IDs from the Configure Library screen.

4. To change any of the configurations, press the Configure button.
5. Press the Select button until the item that you want to configure is highlighted. For the devices, select the desired device (library or drive) by scrolling through the devices with the arrow buttons. After the library or selected drive is selected, use the Select button to highlight the SCSI ID.
6. Use the arrow buttons to scroll through the setting choices until the desired setting appears.
7. When you have the desired setting, press the Change button to save the setting as part of the library configuration.
8. Repeat steps 5 through 7 to make additional changes to the library configuration.
9. Place the library back at the user level of security as follows:
 - a. Press the lock icon on the vertical bar of the control panel.
 - b. On the Password screen, press the User button.

A screen appears informing you that the new security level has been set.
 - c. Press the OK button. The lock icon appears as a locked lock and displays a U to indicate that the control panel is back at User level.
10. Power cycle the tape library to allow the new SCSI IDs to take effect.

8.10.3 TL895 Tape Library Internal Cabling

The default internal cabling configuration for the TL895 tape library has the library robotics controller and top drive (drive 0) on SCSI bus port 1. Drive 1 is on SCSI bus port 2, drive 2 is on SCSI bus port 3, and so on. A terminator (part number 0415619) is connected to each of the drives to provide termination at the tape drive end of the SCSI bus.

In this configuration each of the tape drives, except tape drive 0 and the robotics controller, require a SCSI ID on a separate SCSI bus. The robotics controller and tape drive drive 0 use two SCSI IDs on their SCSI bus.

You can reconfigure the tape drives and robotics controller to place multiple tape drives on the same SCSI bus with SCSI bus jumper (part number 6210567) included with the tape library.

Note

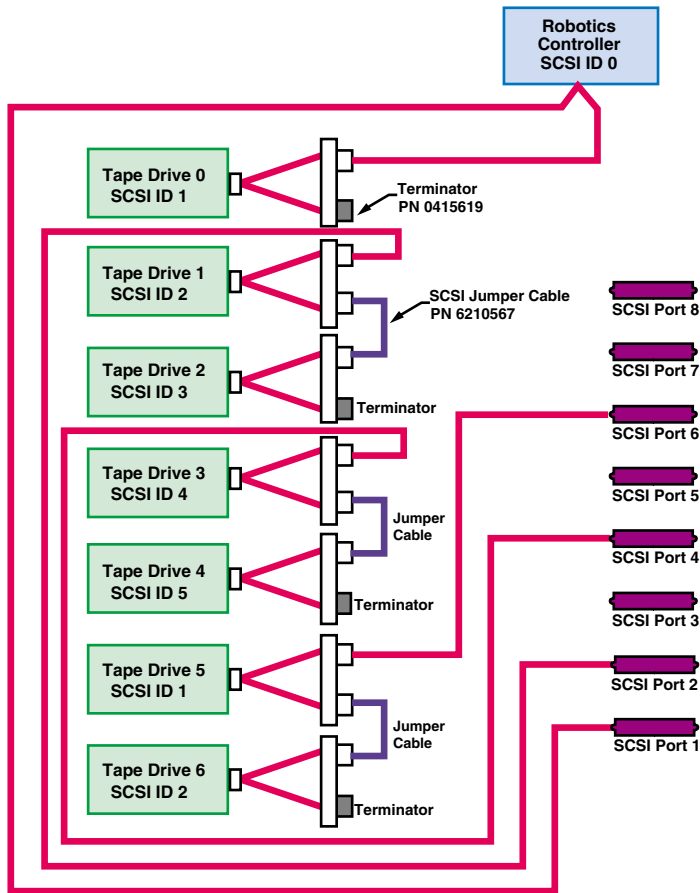
We recommend placing no more than two TZ89 drives on a SCSI bus segment. We also recommend that storage be placed on shared SCSI buses that do not have tape drives.

To reconfigure TL895 SCSI bus configuration, follow these steps:

1. Remove the SCSI bus cable from one drive to be daisy chained.
2. Remove the terminator from the other drive to be daisy chained.
3. Ensure that the drive that will be the last drive on the SCSI bus has a terminator installed.
4. Install a SCSI bus jumper cable (part number 6210567) on the open connectors of the two drives to be daisy chained.

Figure 8–14 shows an example of a TL895 that has tape drives 1, 3, and 5 daisy chained to tape drives 2, 4, and 6 respectively.

Figure 8–14: TL895 Tape Library Internal Cabling



ZK-1397U-AI

8.10.4 Upgrading a TL895

The TL985 DLT automated tape library can be upgraded from two or five tape drives to seven drives with multiple DS-TL89X-UA upgrade kits. Besides the associated documentation, the upgrade kit contains one TZ89N-AV tape drive, a SCSI bus terminator, a SCSI bus jumper (part number 6210567) so you can place more than one drive on the same SCSI bus, and other associated hardware.

Before the drive is physically installed, set the SCSI ID rotary switches (on the library printed circuit board) to the same SCSI ID that will be electronically set. After the drive installation is complete, set the

electronic SCSI ID using the Configure menu from the control panel (see Section 8.10.2).

The actual upgrade is beyond the scope of this manual. See the *TL895 Drive Upgrade Instructions* manual for upgrade instructions.

8.10.5 Connecting the TL895 Tape Library to the Shared SCSI Bus

The TL895 tape library has up to 3 meters (9.8 feet) of internal SCSI cabling per SCSI bus. Because of the internal SCSI cable lengths, you cannot use a tralink connector or Y cable to terminate the SCSI bus external to the library as is done with other devices on the shared SCSI bus. Each SCSI bus must be terminated internal to the tape library at the tape drive itself with the installed SCSI terminators. Therefore, TruCluster Server clusters using the TL895 tape libraries must ensure that the tape libraries are on the end of the shared SCSI bus.

In a TruCluster Server cluster with a TL895 tape library, the member systems and StorageWorks enclosures or RAID subsystems may be isolated from the shared SCSI bus because they use tralink connectors or Y cables. However, because the TL895 cannot be removed from the shared SCSI bus, all ASE services that use any shared SCSI bus attached to the TL895 must be stopped before the tape loader can be removed from the shared bus.

To add a TL895 tape library to a shared SCSI bus, select the member system or storage device that will be the next to last device on the shared SCSI bus. Connect a BN21K or BN21L cable between a tralink or Y cable on that device to the appropriate tape library port.

8.11 Preparing the TL893 and TL896 Automated Tape Libraries for Shared SCSI Bus Usage

The topics in this section provide information on preparing the TL893 and TL896 Automated Tape Libraries (ATLs) for use on a shared SCSI in a TruCluster Server cluster.

Note

To achieve system performance capabilities, we recommend placing no more than two TZ89 drives on a SCSI bus.

The TL893 and TL896 Automated Tape Libraries (ATLs) are designed to provide high-capacity storage and robotic access for the Digital Linear Tape (DLT) series of tape drives. They are identical except in the number of tape drives and the maximum capacity for tape cartridges.

Each tape library comes configured with a robotic controller and bar code reader (to obtain quick and accurate tape inventories).

The libraries have either three or six TZ89N-AV drives. The TL896, because it has a greater number of drives, has a lower capacity for tape cartridge storage.

Each tape library utilizes bulk loading of bin packs, with each bin pack containing a maximum of 11 cartridges. Bin packs are arranged on an eight-sided carousel that provides either two or three bin packs per face. A library with three drives has a carousel three bin packs high. A library with six drives has a carousel that is only two bin packs high. This provides for a total capacity of 24 bin packs (264 cartridges) for the TL893, and 16 bin packs (176 cartridges) for the TL896.

The tape library specifications are as follows:

- TL893 — The TL893 ATL is a high-capacity, 264-cartridge tape library providing up to 18.4 TB of storage. The TL893 uses three fast-wide, differential TZ89N-AV DLT tape drives. It has a maximum transfer rate of almost 10 MB per second (compressed) for each drive, or a total of about 30 MB per second.

The TL893 comes configured for three SCSI-2 buses (a three-bus configuration). The SCSI bus connector is high-density 68-pin, differential.

- TL896 — The TL896 ATL is a high-capacity, 176-cartridge tape library providing up to 12.3 TB of storage. The TL896 uses six fast-wide, differential TZ89N-AV DLT tape drives. It also has a maximum transfer rate of almost 10 MB per second per drive (compressed), or a total of about 60 MB per second.

The TL896 comes configured for six SCSI-2 buses (a six-bus configuration). The SCSI bus connector is also high-density 68-pin, differential.

Both the TL893 and TL896 can be extended by adding additional cabinets (DS-TL893-AC for the TL893 or DS-TL896-AC for the TL896). See the *TL82X Cabinet-to-Cabinet Mounting Instructions* manual for information on adding additional cabinets. Up to five cabinets are supported with the TruCluster Server.

For TruCluster Server, the tape cartridges in all the cabinets are combined into one logical unit, with consecutive numbering from the first cabinet to the last cabinet, by an upgrade from the multi-unit, multi-LUN (MUML) configuration to a multi-unit, single-LUN (MUSL) configuration. See the *TL82X/TL89X MUML to MUSL Upgrade Instructions* manual for information on the firmware upgrade.

These tape libraries each have a multi-unit controller (MUC) that serves two functions:

- It is a SCSI adapter that allows the SCSI interface to control communications between the host and the tape library.
- It permits the host to control up to five attached library units in a multi-unit configuration. Multi-unit configurations are not discussed in this manual. For more information on multi-unit configurations, see the *TL82X/TL893/TL896 Automated Tape Library for DLT Cartridges Facilities Planning and Installation Guide*.

The following sections describe how to prepare these tape libraries in more detail.

8.11.1 Communications with the Host Computer

Two types of communications are possible between the tape library and the host computer: SCSI and EIA/TIA-574 serial (RS-232 for nine-pin connectors). Either method, when used with the multi-unit controller (MUC), allows a single host computer to control up to five units.

A TruCluster Server cluster supports SCSI communications only between the host computer and the MUC. With SCSI communications, both control signals and data flow between the host computer and tape library use the same SCSI cable. The SCSI cable is part of the shared SCSI bus.

An RS-232 loopback cable must be connected between the Unit 0 and Input nine-pin connectors on the rear connector panel. The loopback cable connects the MUC to the robotic controller electronics.

Switch 7 on the MUC switch pack must be down to select the SCSI bus.

8.11.2 MUC Switch Functions

Switch pack 1 on the rear of the multi-unit controller (MUC) is located below the MUC SCSI connectors. The switches provide the functions shown in Table 8–12.

Table 8–12: MUC Switch Functions

Switch	Function
1, 2, and 3	MUC SCSI ID if Switch 7 is down ^a
4 and 5	Must be down, reserved for testing
6	Default is up, disable bus reset on power up

Table 8–12: MUC Switch Functions (cont.)

Switch	Function
7	Host selection: Down for SCSI, up for serial ^a
8	Must be down, reserved for testing

^a For a TruCluster Server cluster, switch 7 is down, allowing switches 1, 2, and 3 to select the MUC SCSI ID.

8.11.3 Setting the MUC SCSI ID

The multi-unit controller (MUC) SCSI ID is set with switch 1, 2, and 3, as shown in Table 8–13. Note that switch 7 must be down to select the SCSI bus and enable switches 1, 2, and 3 to select the MUC SCSI ID.

Table 8–13: MUC SCSI ID Selection

MUC SCSI ID	SW1	SW2	SW3
0	Down	Down	Down
1	Up	Down	Down
2	Down	Up	Down ^a
3	Up	Up	Down
4	Down	Down	Up
5	Up	Down	Up
6	Down	Up	Up
7	Up	Up	Up

^a This is the default MUC SCSI ID.

8.11.4 Tape Drive SCSI IDs

Each tape library arrives with default SCSI ID selections. The TL893 is shown in Table 8–14. The TL896 is shown in Table 8–15.

If you must modify the tape drive SCSI IDs, use the push-button up-down counters on the rear of the drive to change the SCSI ID.

Table 8–14: TL893 Default SCSI IDs

SCSI Port	Device	Default SCSI ID
C	MUC	2
	Drive 2 (top)	5
B	Drive 1 (middle)	4
A	Drive 0 (bottom)	3

Table 8–15: TL896 Default SCSI IDs

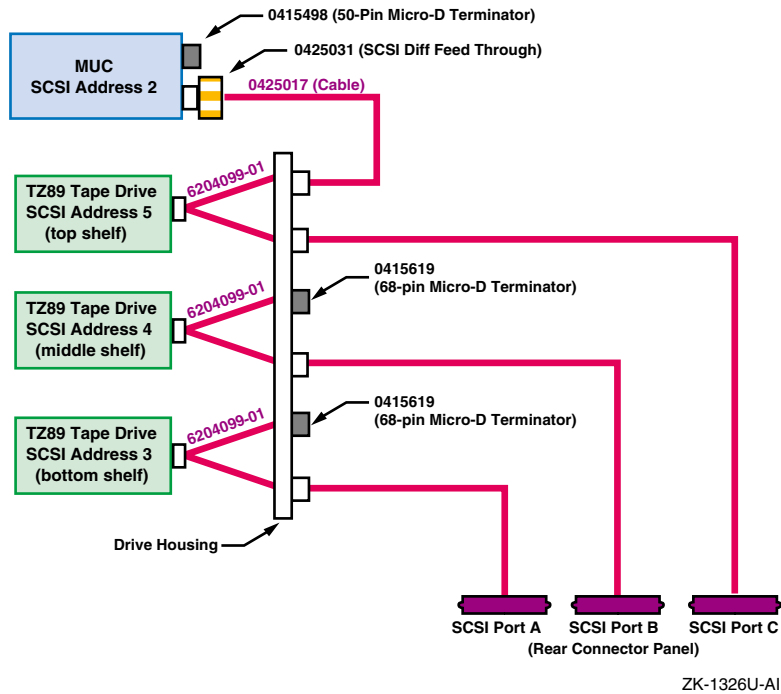
SCSI Port	Device	Default SCSI ID
D	MUC	2
	Drive 5 (top)	5
E	Drive 4	4
F	Drive 3	3
A	Drive 2	5
B	Drive 1	4
C	Drive 0 (bottom)	3

8.11.5 TL893 and TL896 Automated Tape Library Internal Cabling

The default internal cabling configurations for the TL893 and TL896 Automated Tape Libraries (ATLs) are as follows:

- The SCSI input for the TL893 is high-density, 68-pin differential. The default internal cabling configuration for the TL893 is a three-bus mode shown in Figure 8–15 as follows:
 - The top shelf tape drive (SCSI ID 5) and MUC (SCSI ID 2) are on SCSI Port C and are terminated on the MUC. To allow the use of the same MUC and terminator used with the TL822 and TL826, a 68-pin to 50-pin adapter is used on the MUC to connect the SCSI cable from the tape drive to the MUC. In Figure 8–15 it is shown as part number 0425031, the SCSI Diff Feed Through. This SCSI bus is terminated on the MUC with terminator part number 0415498, a 50-pin Micro-D terminator.
 - The middle shelf tape drive (SCSI ID 4) is on SCSI Port B and is terminated on the drive with a 68-pin Micro-D terminator, part number 0415619.
 - The bottom shelf tape drive (SCSI ID 3) is on SCSI Port A and is also terminated on the drive with a 68-pin Micro-D terminator, part number 0415619.

Figure 8–15: TL893 Three-Bus Configuration




ZK-1326U-AI

- The SCSI input for the TL896 is also high-density, 68-pin differential. The default internal cabling configuration for the TL896 is a six-bus configuration shown in Figure 8–16 as follows:
 - The upper bay top shelf tape drive (tape drive 5, SCSI ID 5) and MUC (SCSI ID 2) are on SCSI Port D. To allow the use of the same MUC and terminator used with the TL822 and TL826, a 68-pin to 50-pin adapter is used on the MUC to connect the SCSI cable from the tape drive to the MUC. In Figure 8–16 it is shown as part number 0425031, SCSI Diff Feed Through. This SCSI bus is terminated on the MUC with terminator part number 0415498, a 50-pin Micro-D terminator.
 - The upper bay middle shelf tape drive (tape drive 4, SCSI ID 4) is on SCSI Port E and is terminated on the tape drive.
 - The upper bay bottom shelf tape drive (tape drive 3, SCSI ID 3) is on SCSI Port F and is terminated on the tape drive.
 - The lower bay top shelf tape drive (tape drive 2, SCSI ID 5) is on SCSI Port A and is terminated on the tape drive.
 - The lower bay middle shelf tape drive (tape drive 1, SCSI ID 4) is on SCSI Port B and is terminated on the tape drive.

- The lower bay bottom shelf tape drive (tape drive 0, SCSI ID 3) is on SCSI Port C and is terminated on the tape drive.
- The tape drive terminators are 68-pin differential terminators (part number 0415619).

Figure 8-16: TL896 Six-Bus Configuration



8.11.6 Connecting the TL893 and TL896 Automated Tape Libraries to the Shared SCSI Bus

The TL893 and TL896 Automated Tape Libraries (ATLs) have up to 3 meters (9.8 feet) of internal SCSI cabling on each SCSI bus. Because of the internal SCSI cable lengths, you cannot use a trilink connector or Y cable to terminate the SCSI bus external to the library as is done with

other devices on the shared SCSI bus. Each SCSI bus must be terminated internal to the tape library at the tape drive itself with the installed SCSI terminators. Therefore, TL893 and TL896 tape libraries must be on the end of the shared SCSI bus.

In a TruCluster Server cluster with TL893 or TL896 tape libraries, the member systems and StorageWorks enclosures or RAID subsystems may be isolated from the shared SCSI bus because they use trilink connectors or Y cables. However, if there is disk storage and an ATL on the same shared SCSI bus, the ASE must be shut down to remove a tape library from the shared bus.

You can reconfigure the tape drives and robotics controller to generate other bus configurations by using the jumper cable (ATL part number 0425017) supplied in the accessories kit shipped with each TL893 or TL896 unit.

Remove the terminator from one drive and remove the internal SCSI cable from the other drive to be daisy chained. Use the jumper cable to connect the two drives and place them on the same SCSI bus.

Note

We recommend that you place not more than two drives on any one SCSI bus in these tape libraries.

Figure 8–17 shows a sample TruCluster Server cluster using a TL896 tape library in a three-bus configuration. In this configuration, tape drive 4 (Port E) has been jumpered to tape drive 5, tape drive 2 (Port A) has been jumpered to tape drive 3, and tape drive 1 (Port B) has been jumpered to tape drive 0.

To add a TL893 or TL896 tape library to a shared SCSI bus, select the member system that will be the next to the last device on the shared SCSI bus (the tape library always has to be the last device on the shared SCSI bus). Connect a BN21K, BN21L, or BN31G cable between the Y cable on the SCSI bus controller on that member system and the appropriate tape library port. In Figure 8–17, one shared SCSI bus is connected to port B (tape drives 0 and 1), one shared SCSI bus is connected to port A (tape drives 2 and 3), and a third shared SCSI bus is connected to port E (tape drives 4 and 5 and the MUC).

Figure 8–17: Shared SCSI Buses with TL896 in Three-Bus Mode

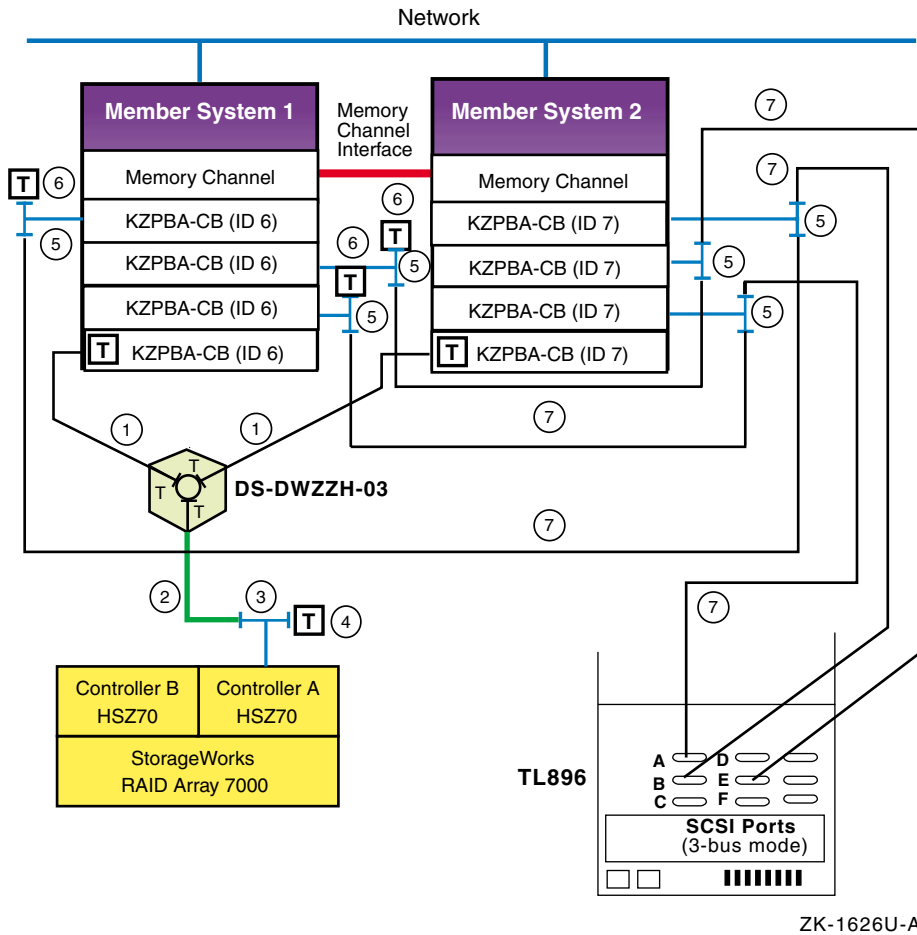


Table 8–16 lists the components that are used to create the cluster that is shown in Figure 8–17.

Table 8–16: Hardware Components Used to Create the Configuration Shown in Figure 8–17

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI tralink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable

Table 8–16: Hardware Components Used to Create the Configuration Shown in Figure 8–17 (cont.)

Callout Number	Description
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.12 Preparing the TL881 and TL891 DLT MiniLibraries for Shared Bus Usage

The topics in this section provide an overview of the Compaq StorageWorks TL881 and TL891 Digital Linear Tape (DLT) MiniLibraries and hardware configuration information for preparing the TL881 or TL891 DLT MiniLibrary for use on a shared SCSI bus.

8.12.1 TL881 and TL891 DLT MiniLibraries Overview

For more information on the TL881 or TL891 DLT MiniLibraries, see the following Compaq documentation:

- *TL881 MiniLibrary System User's Guide*
- *TL891 MiniLibrary System User's Guide*
- *TL881 MiniLibrary Drive Upgrade Procedure*
- *Pass-Through Expansion Kit Installation Instructions*

The TL881 and TL891 Digital Linear Tape (DLT) MiniLibraries are offered as standalone tabletop units or as expandable rackmount units.

The following sections describe these units in more detail.

8.12.1.1 TL881 and TL891 DLT MiniLibrary Tabletop Model

The TL881 and TL891 DLT MiniLibrary tabletop model consists of one unit with a removable 10-cartridge magazine, integral bar code reader, and either one or two DLT 20/40 (TL881) or DLT 35/70 (TL891) drives.

The TL881 DLT MiniLibrary tabletop model is available as either fast, wide differential or fast, wide single-ended. The single-ended model is not supported in a TruCluster Server configuration.

The TL891 DLT MiniLibrary tabletop model is only available as fast, wide differential.

8.12.1.2 TL881 and TL891 MiniLibrary Rackmount Components

A TL881 or TL891 base unit (which contains the tape drive) can operate as an independent, standalone unit, or in concert with an expansion unit and multiple data units.

A rackmount multiple-module configuration is expandable to up to six modules in a configuration. The configuration must contain at least one expansion unit and one base unit. The TL881 and TL891 DLT MiniLibraries may include various combinations of:

- MiniLibrary Expansion unit — the MiniLibrary expansion unit enables multiple TL881 or TL891 modules to share data cartridges and work as a single virtual library. The expansion unit also includes a 16-cartridge magazine.

The expansion unit integrates the robotics in the individual modules into a single coordinated library robotics system. The expansion unit assumes control of the media, maintaining an inventory of all media present in the system, and controls movement of all media. The tape cartridges can move freely between the expansion unit and any of the base units or data units via the system's robotically controlled pass-through mechanism.

The expansion unit can control up to five additional attached modules (base units and data units) to create a multimodule rackmount configuration. The expansion unit must be enabled to control the base unit by setting the base unit to slave mode. The data unit is a passive device and only works as a slave to the expansion unit. To create a multimodule rackmount system, there must be one expansion unit and at least one base unit. The expansion unit has to be the top module in the configuration.

The expansion unit works with either the TL881 or TL891 base unit.

- TL881 or TL891 base unit — includes library robotics, bar code reader, a removable 10-cartridge magazine, and one or two tape drives:
 - TL881 — DLT 20/40 (TZ88N-AV) drives
 - TL891 — DLT 35/70 (TZ89N-AV) drives

To participate in a MiniLibrary configuration, each base unit must be set up as a slave unit to pass control to the expansion unit. When the expansion unit has control over the base unit, the expansion unit controls tape-cartridge movement between the magazines and tape drives.

Note

You cannot mix TL881 and TL891 base units in a rackmount configuration because the tape drives use different formats.

- **Data unit** — This rackmount module contains a 16-cartridge magazine to provide additional capacity in a multi-module configuration. The data unit robotics works in conjunction with the robotics of the expansion unit and base units. It is under control of the expansion unit.

The data unit works with either the TL881 or TL891 base unit.

- **Pass through mechanism** — The pass-through mechanism is attached to the back of the expansion unit and each of the other modules and allows the transfer of tape cartridges between the various modules. It is controlled by the expansion unit.

For each base or data unit added to a configuration, the pass-through mechanism must be extended by 17.8 centimeters (7 inches), which is the height of each module. A 17.8-centimeter (7-inch) gap may be left between modules (providing there is sufficient space), but additional pass-through mechanism extensions must be used.

8.12.1.3 TL881 and TL891 Rackmount Scalability

The rackmount version of the TL881 and TL891 MiniLibraries provides a scalable tape library system that you can configure for maximum performance, maximum capacity, or various combinations between the extremes.

Either library uses DLT IV tape cartridges but can also use DLT III or DLT IIIxt tape cartridges. Table 8–17 lists the capacity and performance of a TL881 or TL891 MiniLibrary in configurations that are set up for either maximum performance or maximum capacity.

Table 8–17: TL881 and TL891 MiniLibrary Performance and Capacity Comparison

Configured for Maximum:	TL881 MiniLibrary			TL891 MiniLibrary		
	Number of Base Units ^{a b}	Number of Data Units ^c	Transfer Rate ^d	Storage Capacity ^e	Transfer Rate ^f	Storage Capacity ^g
Performance	5	0	15 MB/sec (54 GB/hr)	1.32 TB (66 cartridges)	50 MB/sec (180 GB/hr)	2.31 TB (66 cartridges)
Capacity	1	4	3 MB/sec (10.8 GB/hr)	1.8 TB (90 cartridges)	10 MB/sec (36 GB/hr)	3.15 TB (90 cartridges)

^a Using an expansion unit with a full 16-cartridge magazine.

^b Each base unit has a full 10-cartridge magazine and two tape drives.

^c Using a data unit with full 16-cartridge magazine.

^d Up to 1.5 MB/sec per drive.

^e Based on 20 GB/cartridge uncompressed. It can be up to 40 GB/cartridge compressed.

^f Up to 5 MB/sec per drive.

^g Based on 35 GB/cartridge uncompressed. It can be up to 70 GB/cartridge compressed.

By modifying the combinations of base units and data units, the performance and total capacity can be adjusted to meet the customers' needs.

8.12.1.4 DLT MiniLibrary Part Numbers

Table 8–18 lists the part numbers for the TL881 and TL891 DLT MiniLibrary systems. Part numbers are only shown for the TL881 fast, wide differential components.

Table 8–18: DLT MiniLibrary Part Numbers

DLT Library Component	Number of Tape Drives	Tabletop/Rackmount	Part Number
TL881 DLT Library	1	Tabletop	128667-B21
TL881 DLT Library	2	Tabletop	128667-B22
TL881 DLT MiniLibrary Base Unit	1	Rackmount	128669-B21
TL881 DLT MiniLibrary Base Unit	2	Rackmount	128669-B22
Add-on DLT 20/40 drive for TL881	1	N/A	128671-B21
TL891 DLT Library	1	Tabletop	120875-B21
TL891 DLT Library	2	Tabletop	120875-B22
TL891 DLT MiniLibrary Base Unit	1	Rackmount	120876-B21
TL891 DLT MiniLibrary Base Unit	2	Rackmount	120876-B22
Add-on DLT 35/70 drive for TL891	1	N/A	120878-B21
MiniLibrary Expansion Unit	N/A	Rackmount	120877-B21
MiniLibrary Data Unit	N/A	Rackmount	128670-B21

Note

The TL881 DLT MiniLibrary tabletop model is available as fast, wide differential or fast, wide single-ended. The single-ended model is not supported in a cluster configuration. The TL891 DLT MiniLibrary tabletop model is only available as fast, wide differential.

8.12.2 Preparing a TL881 or TL891 MiniLibrary for Shared SCSI Bus Use

The following sections describe how to prepare the TL881 and TL891 DLT MiniLibraries for shared SCSI bus use in more detail.

8.12.2.1 Preparing a Tabletop Model or Base Unit for Standalone Shared SCSI Bus Usage

A TL881 or TL891 DLT MiniLibrary tabletop model or a rackmount base unit may be used standalone. You may want to purchase a rackmount base unit for future expansion.

Note

To achieve system performance capabilities, we recommend placing no more than two tape drives on a SCSI bus, and also recommend that no shared storage be placed on the same SCSI bus with a tape library.

The topics in this section provide information on preparing the TL881 or TL891 DLT MiniLibrary tabletop model or rackmount base unit for use on a shared SCSI bus.

For complete hardware installation instructions, see the *TL881 MiniLibrary System User's Guide* or *TL891 MiniLibrary System User's Guide*.

8.12.2.1.1 Setting the Standalone MiniLibrary Tape Drive SCSI ID

The control panel on the front of the TL891 and TL892 MiniLibraries is used to display power-on self-test (POST) status, display messages, and to set up MiniLibrary functions.

When power is first applied to a MiniLibrary, a series of POST diagnostics are performed. During POST execution, the MiniLibrary model number, current date and time, firmware revision, and the status of each test is displayed on the control panel.

After the POST diagnostics have completed, the default screen is shown:

```
DLT0 Idle
DLT1 Idle
Loader Idle
0> _ _ _ _ _ <9
```

The first and second lines of the default screen show the status of the two (if present) drives. The third line shows the status of the library robotics, and the fourth line is a map of the magazine, with the numbers from 0 through 9 representing the cartridge slots. Rectangles on this line indicate cartridges that are present in the corresponding slot of the magazine.

For example, this fourth line (0> X X _ _ _ _ _ <9, where an X represents a rectangle) indicates that cartridges are installed in slots 0 and 1.

Note

There are no switches for setting a mechanical SCSI ID for the tape drives. The SCSI IDs default to five. The MiniLibrary sets the electronic SCSI ID very quickly, before any device can probe the MiniLibrary, so the lack of a mechanical SCSI ID does not cause any problems on the SCSI bus.

To set the SCSI ID, follow these steps:

1. From the Default Screen, press the Enter button to enter the Menu Mode, displaying the Main Menu.

Note

When you enter the Menu Mode, the Ready light goes out, an indication that the module is off line, and all medium changer commands from the host return a SCSI "not ready" status until you exit the Menu Mode and the Ready light comes on again.

2. Press the down arrow button until the Configure Menu item is selected, then press the Enter button to display the Configure submenu.

Note

The control panel up and down arrows have an auto-repeat feature. When you press either button for more than half a second, the control panel behaves as if you were pressing the button about four times per second. The effect stops when you release the button.

3. Press the down arrow button until the Set SCSI item is selected and press the Enter button.
4. Select the tape drive (DLT0 Bus ID: or DLT1 Bus ID:) or library robotics (LIB Bus ID:) whose SCSI bus ID you want to change. The default SCSI IDs are as follows:
 - Lib Bus ID: 0
 - DLT0 Bus ID: 4
 - DLT1 Bus ID: 5

Use the up or down arrow button to select the item for which you need to change the SCSI ID. Press the Enter button.

5. Use the up or down arrow button to scroll through the possible SCSI ID settings. Press the Enter button when the desired SCSI ID is displayed.
6. Repeat steps 4 and 5 to set other SCSI bus IDs as necessary.
7. Press the Escape button repeatedly until the default menu is displayed.

8.12.2.1.2 Cabling the TL881 or TL891 DLT MiniLibrary

There are six 68-pin, high-density SCSI connectors on the back of the TL881 or TL891 DLT MiniLibrary standalone model or rackmount base unit. The two leftmost connectors are for the library robotics controller. The middle two are for tape drive 1. The two on the right are for tape drive 2 (if the second tape drive is installed).

Note

The tape drive SCSI connectors are labeled DLT1 (tape drive 1) and DLT2 (tape drive 2). The control panel designation for the drives is DLT0 (tape drive 1) and DLT1 (tape drive 2).

The default for the TL881 or TL891 DLT MiniLibrary is to place the robotics controller and tape drive 1 on the same SCSI bus (Figure 8–18). A 30-centimeter (11.8-inch) SCSI jumper cable is provided with the unit. Plug this cable into the second connector (from the left) and the third connector. If the MiniLibrary has two drives, place the second drive on the same SCSI bus with another 30-centimeter (11.8-inch) SCSI bus jumper cable, or place it on its own SCSI bus.

Notes

The internal cabling of the TL881 and TL891 is too long to allow external termination with a trilink/terminator combination. Therefore, the TL881 or TL891 must be the last device on the shared SCSI bus. They may not be removed from the shared SCSI bus without stopping all ASE services that generate activity on the bus.

To achieve system performance capabilities, we recommend placing no more than two tape drives on a SCSI bus.

We recommend that tape devices be placed on separate shared SCSI buses, and that there be no storage devices on the SCSI bus.

The cabling depends on whether or not there are one or two drives, and for the two-drive configuration, if each drive is on a separate SCSI bus.

Note

It is assumed that the library robotics controller is on the same SCSI bus as tape drive 1.

To connect the library robotics and one drive to a single shared SCSI bus, follow these steps:

1. Connect a 328215-00X, BN21K, or BN21L between the last Y cable or tralink connector on the bus to the leftmost connector (as viewed from the rear) of the MiniLibrary. The 328215-004 is a 20-meter (65.6-foot) cable.
2. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost robotics connector (second connector from the left) and the left DLT1 connector (the third connector from the left).
3. Install an HD68 differential terminator (such as an H879-AA) on the right DLT1 connector (the fourth connector from the left).

To connect the drive robotics and two drives to a single shared SCSI bus, follow these steps:

1. Connect a 328215-00X, BN21K, or BN21L between the last tralink connector on the bus to the leftmost connector (as viewed from the rear) of the MiniLibrary.
2. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost robotics connector (the second connector from the left) and the left DLT1 connector (the third connector from the left).
3. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost DLT1 connector (the fourth connector from the left) and the left DLT2 connector (the fifth connector from the left).
4. Install an HD68 differential (H879-AA) terminator on the right DLT2 connector (the rightmost connector).

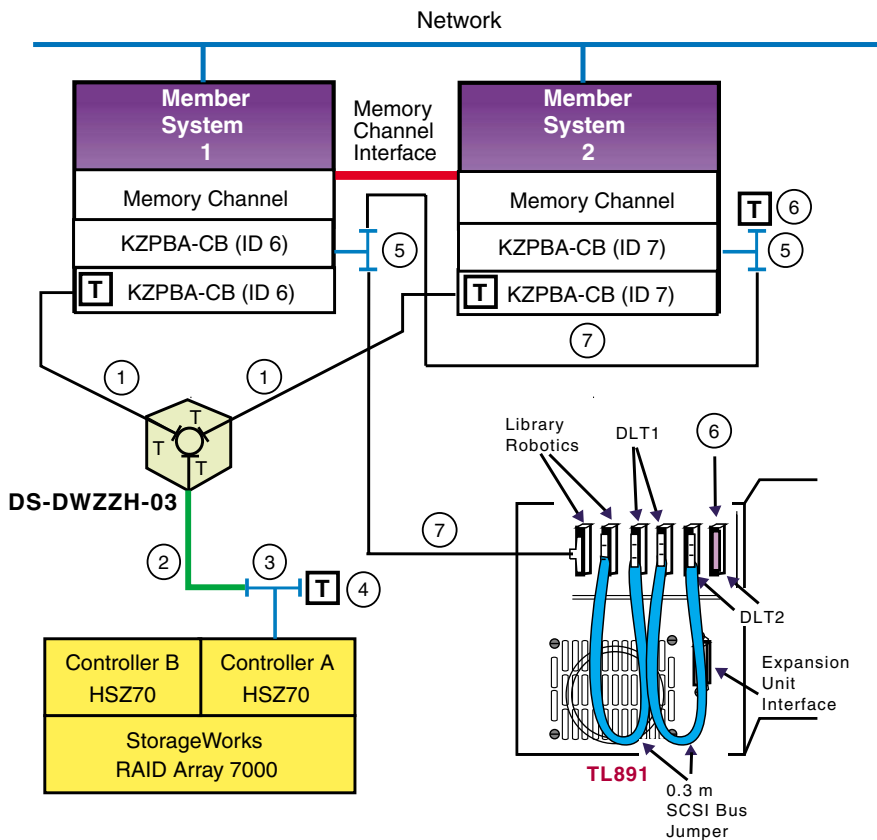
To connect the drive robotics and one drive to one shared SCSI bus and the second drive to a second shared SCSI bus, follow these steps:

1. Connect a 328215-00X, BN21K, or BN21L between the last tralink connector on one shared SCSI bus to the leftmost connector (as viewed from the rear) of the MiniLibrary.
2. Connect a 328215-00X, BN21K, or BN21L between the last tralink connector on the second shared SCSI bus to the left DLT2 connector (the fifth connector from the left).
3. Install a 30-centimeter (11.8-inch) SCSI bus jumper between the rightmost robotics connector (the second connector from the left) and the left DLT1 connector (the third connector from the left).

4. Install an HD68 differential (H879-AA) terminator on the right DLT1 connector (the fourth connector from the left) and install another HD68 differential terminator on the right DLT2 connector (the rightmost connector).

Figure 8–18 shows an example of a TruCluster configuration with a TL891 standalone MiniLibrary connected to two shared SCSI buses.

Figure 8–18: TL891 Standalone Cluster Configuration



ZK-1627U-AI

Table 8–19 lists the components that are used to create the cluster that is shown in Figure 8–18.

Table 8–19: Hardware Components Used to Create the Configuration Shown in Figure 8–18

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI trilink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.12.2.2 Preparing a TL881 or TL891 Rackmount MiniLibrary for Shared SCSI Bus Usage

A TL881 or TL891 MiniLibrary base unit may also be used in a rackmount configuration with an expansion unit, data units, and other base units, to add tape drive and/or cartridge capacity to the configuration.

The expansion unit is installed above the TL881 or TL891 DLT MiniLibrary base or data units in a SW500, SW800, or RETMA cabinet.

For complete hardware installation instructions, see the *TL881 MiniLibrary System User's Guide* or *TL891 MiniLibrary System User's Guide*.

The topics in this section provide information on preparing the rackmount TL881 or TL891 DLT MiniLibrary for use on a shared SCSI bus.

It is assumed that the expansion unit, base modules, and pass-through and motor mechanism have been installed.

8.12.2.2.1 Cabling the Rackmount TL881 or TL891 DLT MiniLibrary

You must make the following connections to render the DLT MiniLibrary system operational:

- Expansion unit to the pass-through motor mechanism: The motor mechanism cable is about 1 meter (3.3 feet) long and has a DB-15 connector on each end. Connect it between the connector labeled Motor on the expansion unit to the motor on the pass-through mechanism.

Note

This cable is not shown in Figure 8–19 because the pass-through mechanism is not shown in the figure.

- Robotics control cables from the expansion unit to each base unit or data unit: These cables have a DB-9 male connector on one end and a DB-9 female connector on the other end. Connect the male end to the Expansion Unit Interface connector on the base unit or Diagnostic connector on the data unit and the female end to any Expansion Modules connector on the expansion unit.

Note

It does not matter which interface connector you connect to a base unit or a data unit.

- SCSI bus connection to the expansion unit robotics: Connect the shared SCSI bus that will control the robotics to one of the SCSI connectors on the expansion unit with a 328215-00X, BN21K, or BN21L cable. Terminate the SCSI bus with an HD68 terminator (such as an H879-AA) on the other expansion unit SCSI connector.
- SCSI bus connection to each of the base module tape drives: Connect a shared SCSI bus to one of the DLT1 or DLT2 SCSI connectors on each of the base modules with 328215-00X, BN21K, or BN21L cables. Terminate the other DLT1 or DLT2 SCSI bus connection with an HD68 terminator (H879-AA).

You can daisy chain between DLT1 and DLT2 (if present) with a 30-centimeter (11.8-inch) SCSI bus jumper (supplied with the TL881 or TL891). Terminate the SCSI bus at the tape drive on the end of the shared SCSI bus with an HD68 terminator (H879-AA).

Notes

Do not connect a SCSI bus to the SCSI connectors for the library connectors on the base modules.

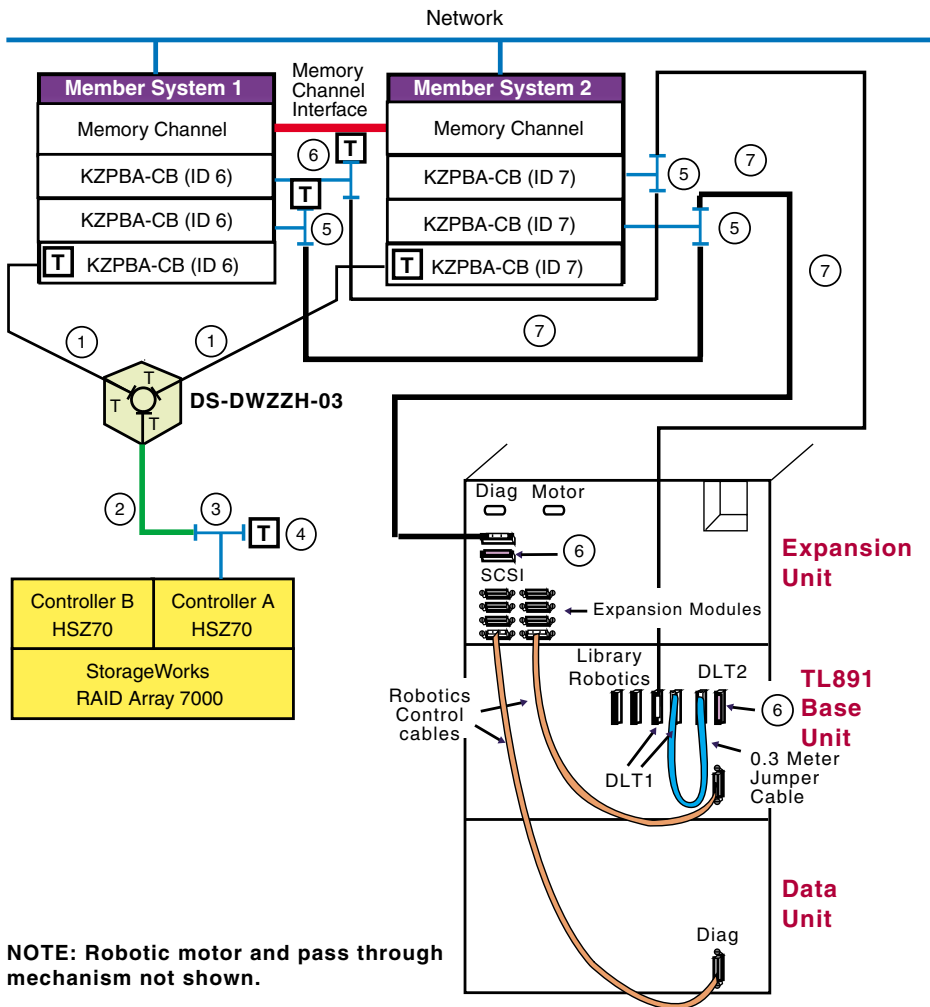
We recommend that no more than two tape drives be on a SCSI bus.

Figure 8–19 shows a TL891 DLT MiniLibrary configuration with an expansion unit, a base unit, and a data unit. The library robotics expansion unit is on one shared SCSI bus and the two tape drives in the base unit are on separate, shared SCSI buses. The data unit is not on a shared SCSI bus because it does not contain any tape drives but tape cartridges only. The

pass-through mechanism and cable to the library robotics motor are not shown in this figure.

For more information on cabling the units, see Section 8.12.2.1.2. With the exception of the robotics control on the expansion module, a rackmount TL881 or TL891 DLT MiniLibrary is cabled in the same manner as a tabletop unit.

Figure 8–19: TL891 DLT MiniLibrary Rackmount Configuration



NOTE: Robotic motor and pass through mechanism not shown.

ZK-1628U-AI

Table 8–20 list the components that are used to create the cluster that is shown in Figure 8–19.

Table 8–20: Hardware Components Used to Create the Configuration Shown in Figure 8–19

Callout Number	Description
1	BN38C or BN38D cable ^a
2	BN37A cable ^b
3	H8861-AA VHDCI trilink connector
4	H8863-AA VHDCI terminator
5	BN21W-0B Y cable
6	H879-AA terminator
7	328215-00X, BN21K, or BN21L cable ^c

^a The maximum length of the BN38C (or BN38D) cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum length of the BN37A cable must not exceed 25 meters (82 feet).

^c The maximum combined length of these cables must not exceed 25 meters (82 feet).

8.12.2.2.2 Configuring a Base Unit as a Slave to the Expansion Unit

The TL881/TL891 base units are shipped configured as standalone systems. When they are used in conjunction with the MiniLibrary expansion unit, the expansion unit must control the robotics of each of the base units. Therefore, the base units must be configured as slaves to the expansion unit.

After the hardware and cables are installed, but before you power up the expansion unit in a MiniLibrary system for the first time, you must reconfigure each of the base units in the system as a slave. The expansion unit will not have control over the base unit robotics when you power up the MiniLibrary system, if you do not reconfigure the base unit as a slave.

To reconfigure a TL881/TL891 base unit as a slave to the MiniLibrary expansion unit, perform the following procedure on each base unit in the system.

1. Turn on the power switch on the TL881/TL891 base unit to be reconfigured.

Note

Do not power on the expansion unit. Leave it powered off until all base units have been reconfigured as slaves.

After a series of self-tests have executed, the default screen will be displayed on the base module control panel:

```
DLT0 Idle
DLT1 Idle
Loader Idle
0> _ _ _ _ _ <9
```

The default screen shows the state of the tape drives, loader, and number of cartridges present for this base unit. A rectangle in place of the underscore indicates that a cartridge is present in that location.

2. Press the Enter button to enter the Menu Mode, displaying the Main Menu.
3. Press the down arrow button until the Configure Menu item is selected, then press the Enter button.

Note

The control panel up and down arrows have an auto-repeat feature. When you press either button for more than half a second, the control panel behaves as if you were pressing the button about four times per second. The effect stops when you release the button.

4. Press the down arrow button until the Set Special Config menu is selected and press the Enter button.
5. Press the down arrow button repeatedly until the Alternate Config item is selected and press the Enter button.
6. Press the down arrow button to change the alternate configuration from the default (Standalone) to Slave. Press the Enter button.
7. After the selection stops flashing and the control panel indicates that the change is not effective until a reboot, press the Enter button.
8. When the Special Configuration menu reappears, turn the power switch off and then on again to cycle the power. The base unit is now reconfigured as a slave to the expansion unit.
9. Repeat the steps for each TL881/TL891 base unit present that is to be a slave to the expansion unit.

8.12.2.2.3 Powering Up the TL881/TL891 DLT MiniLibrary

When turning on power to the TL881 or TL891 DLT MiniLibrary, power must be applied to the expansion unit simultaneously or after power is applied to the base units and data units. If the expansion unit is powered on first, its inventory of modules may be incorrect and the contents of some or all of the modules will be inaccessible to the system and to the host.

When the expansion unit comes up, it will communicate with each base and data unit through the expansion unit interface and inventory the number of base units, tape drives, data units, and cartridges present in each base and data unit. After the MiniLibrary configuration has been determined, the expansion unit will communicate with each base and data unit and indicate to the modules which cartridge group that base or data unit contains.

When all initialization communication between the expansion module and each base and data unit has completed, the base and data units will display their cartridge numbers according to the remapped cartridge inventory.

8.12.2.2.4 Setting the SCSI IDs for a Rackmount TL881 or TL891 DLT MiniLibrary

After the base units have been reconfigured as slaves, each base unit control panel still provides tape drive status and error information, but all control functions are carried out from the expansion unit control panel. This includes setting the SCSI ID for each of the tape drives present.

To set the SCSI IDs for the tape drives in a TL881 or TL891 DLT MiniLibrary rackmount configuration, follow these steps:

1. Apply power to the MiniLibrary, ensuring that you power up the expansion unit after or at the same time as the base and data units.
2. Wait until power-on self-tests (POST) have terminated and the expansion unit and each base and data unit display the default screen.
3. At the expansion unit control panel, press the Enter button to display the Main Menu.
4. Press the down arrow button until the Configure Menu item is selected, and then press the Enter button to display the Configure submenu.
5. Press the down arrow button until the Set SCSI item is selected and press the Enter button.
6. Press the up or down arrow button to select the appropriate tape drive (DLT0 Bus ID:, DLT1 Bus ID:, DLT2 Bus ID:, and so on) or library robotics (Library Bus ID:) for which you want to change the SCSI bus ID. In a configuration with three base units, and assuming that each base unit has two tape drives, the top base unit contains DLT0 and DLT1. The next base unit down contains DLT2 and DLT3. The next base unit contains DLT4 and DLT5. The default SCSI IDs, after being reconfigured by the expansion unit, are as follows:
 - Library Bus ID: 0
 - DLT0 Bus ID: 1
 - DLT1 Bus ID: 2
 - DLT2 Bus ID: 3

- DLT3 Bus ID: 4
 - DLT4 Bus ID: 5
 - DLT5 Bus ID: 6
7. Press Enter when you have the item selected for which you want to change the SCSI ID.
 8. Use the up and down arrows to select the desired SCSI ID. Press the Enter button to save the new selection.
 9. Press the Escape button once to return to the Set SCSI Submenu to select another tape drive or the library robotics, and then repeat steps 6, 7, and 8 to set the SCSI ID.
 10. If there are other items that you want to configure, press the Escape button until the Configure submenu is displayed, then select the item to be configured. Repeat this procedure for each item that you want to configure.
 11. If there are no more items to be configured, press the Escape button until the Default window is displayed.

Note

You do not have to cycle power to set the SCSI IDs.

8.13 Compaq ESL9326D Enterprise Library

The topics in this section provide an overview and hardware configuration information on preparing the ESL9326D Enterprise Library for use on a shared SCSI bus with the TruCluster Server.

8.13.1 General Overview

The Compaq StorageWorks ESL9326D Enterprise Library is the first building block of the Compaq ESL 9000 series tape library.

For more information on the ESL9326D Enterprise Library, see the following Compaq StorageWorks ESL9000 Series Tape Library documentation:

- *Unpacking Guide*
- *Reference Guide*
- *Maintenance and Service Guide*
- *ESL9326 Tape Drive Upgrade Guide*
- *Pass-Through Mechanism Installation Guide*

These tape devices have been qualified for use on shared SCSI buses with both the KZPSA-BB and KZPBA-CB host bus adapters. Ensure that the host bus adapter you use is supported on your system by searching the options list for your system at the following URL:
<http://www.compaq.com/alphaserver/products/options.html>

8.13.2 ESL9326D Enterprise Library Overview

The ESL9326D Enterprise Library is an enterprise Digital Linear Tape (DLT) automated tape library with from 6 to 16 fast-wide, differential tape drives. This tape library uses the 35/70 DLT (DS-TZ89N-AV) differential tape drives. The SCSI bus connectors are 68-pin, high-density.

The ESL9326D Enterprise Library has a capacity of 326 DLT cartridges in a fixed storage array (back wall, inside the left door, and inside the right door). This provides a storage capacity of 11.4 TB uncompressed for the ESL9326D Enterprise Library using DLT Tape IV cartridges. The library can also use DLT Tape III or IIIXT tape cartridges.

The ESL9326D Enterprise Library is available as seven different part numbers, based on the number of tape drives:

Order Number	Number of Tape Drives
146205-B21	0
146205-B23	6
146205-B24	8
146205-B25	10
146205-B26	12
146205-B27	14
146205-B28	16

You can upgrade a tape library with part number 146209-B21, which adds a 35/70 DLT tape drive. See the Compaq StorageWorks *ESL9326 Tape Drive Upgrade Guide* for more information.

Up to five ESL9000 series tape libraries can be connected together with pass-through mechanism, part number 161268-B21. See the ESL9326D QuickSpecs to determine the number of tape drives supported when the tape libraries are connected together. See the Compaq StorageWorks *Pass-Through Mechanism Installation Guide* for more information on the pass-through mechanism.

8.13.3 Preparing the ESL9326D Enterprise Library for Shared SCSI Bus Usage

The ESL9326D Enterprise Library contains library electronics (robotic controller) and from 6 to 16 35/70 DLT (DS-TZ89N-AV) fast-wide, differential DLT tape drives.

Tape devices are supported only on those shared SCSI buses that use the KZPSA-BB or KZPBA-CB host bus adapters.

Notes

The ESL9326D Enterprise Library is cabled internally for two 35/70 DLT tape drives on each SCSI bus. It arrives with the library electronics cabled to tape drives 0 and 1. Every other pair of tape drives is cabled together (2 and 3, 4 and 5, 6 and 7, and so on).

An extra SCSI bus jumper cable is provided with the ESL9326D Enterprise Library for those customers that do not have enough SCSI buses, and want to jumper two SCSI buses together and place four tape drives on the same SCSI bus.

We recommend that you place no more than two 35/70 DLT tape drives on a shared SCSI bus.

We also recommended that you do not place storage on shared SCSI buses that have tape drives.

The following sections describe how to prepare the ESL9326D Enterprise Library in more detail.

8.13.3.1 ESL9326D Enterprise Library Robotic and Tape Drive Required Firmware

Library electronics firmware V1.22 is the minimum firmware version that supports TruCluster Server.

The 35/70 DLT tape drives require V97 or later firmware. The firmware must be at V80 before you can upgrade to V97.

8.13.3.2 Library Electronics and Tape Drive SCSI IDs

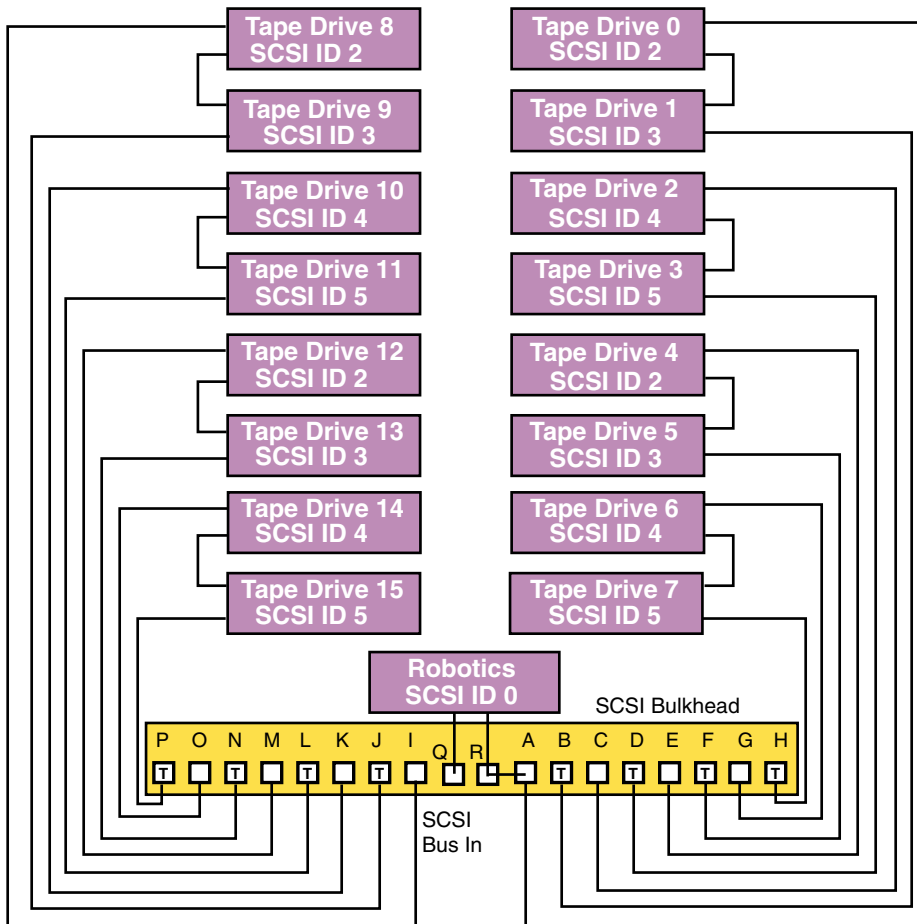
Figure 8–20 shows the default robotics and tape drive SCSI IDs. If these SCSI IDs are not acceptable for your configuration and you need to change them, follow the steps in the Compaq StorageWorks ESL9000 Series Tape Library *Reference Guide*.

8.13.3.3 ESL9326D Enterprise Library Internal Cabling

The default internal cabling for the ESL9326D Enterprise Library is to place two 35/70 DLT tape drives on one SCSI bus.

Figure 8–20 shows the default cabling for an ESL9326D Enterprise Library with 16 tape drives. Each pair of tape drives is cabled together internally to place two drives on a single SCSI bus. If your model has fewer drives, all internal cabling is supplied. The terminators for the drives that are not present are not installed on the SCSI bulkhead.

Figure 8–20: ESL9326D Internal Cabling



ZK-1705U-AI

Note

Each internal cable is up to 2.5 meters (8.2 feet) long. The length of the internal cables, two per SCSI bus, must be taken into consideration when ordering SCSI bus cables.

The maximum length of a differential SCSI bus segment is 25 meters (82 feet), and the internal tape drive SCSI bus length is 5 meters (16.4 feet). Therefore, you must limit the external SCSI bus cables to 20 meters (65.6 feet) maximum.

8.13.3.4 Connecting the ESL9326D Enterprise Library to the Shared SCSI Bus

The ESL9326D Enterprise Library has 5 meters (16.4 feet) of internal SCSI bus cabling for each pair of tape drives. Because of the internal SCSI bus lengths, you cannot use a trilink connector or Y cable to terminate the SCSI bus external to the tape library as is done with other devices on the shared SCSI bus. You must terminate each SCSI bus at the end of the SCSI bus by installing a terminator on the SCSI bulkhead SCSI connector. Therefore, you must ensure that the ESL9326D Enterprise Library is on the end of the shared SCSI bus in TruCluster Server configurations.

Note

We recommend that you put disk storage devices on separate shared SCSI buses.

Use 328215-001 (5-meter; 16.4 foot), 328215-002 (10-meter; 32.8-foot), 328215-003 (15-meter; 49.2-foot), 328215-004 (20-meter; 65.6-foot), or BN21K (BN21L) cables of the appropriate length to connect the ESL9326D Enterprise Library to a shared SCSI bus. The entire shared SCSI bus, including ESL9326 internal cables, cannot exceed 25 meters (82 feet). Terminate each SCSI bus with a 330563-001 (or H879-AA) HD-68 terminator. Connect the cables and terminator on the SCSI bulkhead SCSI connectors as indicated in Table 8–21 to form shared SCSI buses.

Table 8–21: Shared SCSI Bus Cable and Terminator Connections for the ESL9326D Enterprise Library

Tape Drives on Shared SCSI Bus	Connect SCSI Cable to Connector	Install HD68 Terminator on Connector
0, 1, and library electronics ^a	Q	B
2, 3	C	D
4, 5	E	F
6, 7	G	H
8, 9	I	J
10, 11	K	L
12, 13	M	N
14, 15	O	P

^a Install 30-centimeter (11.8-inch) jumper cable part number 330582-001 between SCSI connectors R and A to place the library electronics on the SCSI bus with tape drives 0 and 1.

Notes

Each ESL9326D Enterprise Library arrives with one 330563-001 HD68 terminator for each pair of tape drives (one SCSI bus). The kit also includes at least one 330582-001 jumper cable to connect the library electronics to tape drives 0 and 1.

Tape libraries with more than six tape drives include extra 330582-01 jumper cables in case you do not have enough host bus adapters and you want to place more than two tape drives on a single SCSI bus (a configuration that we do not recommend).

Configuring Systems for External Termination or Radial Connections to Non-UltraSCSI Devices

This chapter describes how to prepare the systems for a TruCluster Server cluster when there is a need to access:

- Shared SCSI storage using external termination.
- Non-UltraSCSI RAID array controllers (HSZ40 and HSZ50) using a radial connection.

This chapter does not provide detailed information about installing devices; it describes only how to set up the hardware in the context of the TruCluster Server product. Therefore, you must have the documentation that describes how to install the individual pieces of hardware. That documentation typically arrives with the hardware.

All systems in the cluster must be connected via the Memory Channel cluster interconnect. Not all members must be connected to a shared SCSI bus. We recommend placing the clusterwide root (`/`), `/usr`, and `/var` file systems, all member boot disks, and the quorum disk (if provided) on shared SCSI buses. All configurations covered in this manual assume the use of a shared SCSI bus.

Before proceeding further, review Section 4.1, Section 4.2, and the first two paragraphs of Section 4.3.

9.1 TruCluster Server Hardware Installation Using PCI SCSI Adapters

The following sections describe how to install the KZPSA-BB or KZPBA-CB host bus adapters and configure them into TruCluster Server clusters using both methods of termination — the preferred method of radial connection with internal termination used with the HSZ40 and HSZ50 RAID array controllers, and the old method of external termination.

The descriptions are based on the assumption that you will configure and cable your storage subsystems as described in Chapter 10 after installing the host bus adapters, Memory Channel adapters, and network adapters. When the system hardware (KZPSA-BB or KZPBA-CB host bus adapters, Memory

Channel adapters, hubs (if necessary), cables, and network adapters) have been installed, you can connect your host bus adapter to the UltraSCSI hub or storage subsystem.

Follow the steps in Table 9–1 to start the TruCluster Server hardware installation procedure. You can save time by installing the Memory Channel adapters, redundant network adapters (if applicable), and KZPSA-BB or KZPBA-CB SCSI adapters all at the same time.

Follow the directions in the referenced documentation, or the steps in the referenced tables for the particular SCSI host bus adapter, returning to the appropriate table when you have completed the steps in the referenced table.

After you complete installing host bus adapters, Memory Channel adapters, and network adapters, see Chapter 10 and connect your cluster member systems to shared storage.

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.

Table 9–1: Configuring TruCluster Server Hardware for Use with a PCI SCSI Adapter

Step	Action	Refer to:
1	Install the Memory Channel modules, cables, and hubs (if a hub is required).	Chapter 5 ^a
2	Install Ethernet or FDDI network adapters.	User's guide for the applicable Ethernet or FDDI adapter, and the user's guide for the applicable system
	Install ATM adapters if using ATM.	See the documentation that came with the ATM adapter
3	Install a KZPSA-BB PCI SCSI adapter or KZPBA-CB UltraSCSI adapter for each shared SCSI bus in each member system:	—
	Internally terminated host bus adapter for radial connection to DWZZH UltraSCSI hub:	Section 9.1.1 and Table 9–2
	Externally terminated host bus adapter:	Section 9.1.2 Table 9–3

^a If you install additional KZPSA-BB or KZPBA-CB SCSI adapters or an extra network adapter at this time, delay testing the Memory Channel until you have installed all hardware.

9.1.1 Radial Installation of a KZPSA-BB or KZPBA-CB Using Internal Termination

Use this method of cabling member systems and shared storage in a TruCluster Server cluster if you are using a DWZZH UltraSCSI hub. You must reserve at least one hub port for shared storage.

The DWZZH-series UltraSCSI hubs are designed to allow more separation between member systems and shared storage. Using the UltraSCSI hub also improves the reliability of the detection of cable faults.

A side benefit is the ability to connect the member systems' SCSI adapter directly to a hub port without external termination. This simplifies the configuration by reducing the number of cable connections.

A DWZZH UltraSCSI hub can be installed in:

- A StorageWorks UltraSCSI BA356 shelf (which has the required 180-watt power supply).
- A non-UltraSCSI BA356 that has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.

An UltraSCSI hub only receives power and mechanical support from the storage shelf. There is no SCSI bus continuity between the DWZZH and storage shelf.

The DWZZH contains a differential to single-ended signal converter for each hub port (sometimes referred to as a DWZZA on a chip, or DOC chip). The single-ended sides are connected together to form an internal single-ended SCSI bus segment. Each differential SCSI bus port is terminated internal to the DWZZH with terminators that cannot be disabled or removed.

Power for the DWZZH termination (`termpwr`) is supplied by the host bus adapter or RAID array controller connected to the DWZZH port. If the member system or RAID array controller is powered down, or the cable is removed from the host bus adapter, RAID array controller, or hub port, the loss of `termpwr` disables the hub port without affecting the remaining hub ports or SCSI bus segments. This is similar to removing a Y cable when using external termination.

Note

The DS-DWZZH-03/05 UltraSCSI hubs cannot be connected to a StorageWorks BA35X storage shelf because the storage shelf does not provide termination power to the hub.

The other end of the SCSI bus segment is terminated by the KZPSA-BB or KZPBA-CB onboard termination resistor SIPs, or a tralink connector/terminator combination installed on the HSZ40 or HSZ50.

The KZPSA-BB PCI-to-SCSI bus adapter:

- Is installed in a PCI slot of the supported member system (see Section 2.3.2).
- Is a fast, wide differential adapter with only a single port, so only one differential shared SCSI bus can be connected to a KZPSA-BB adapter.
- Operates at fast or slow speed and is compatible with narrow or wide SCSI. The fast speed is 10 MB/sec for a narrow SCSI bus and 20 MB/sec for a wide SCSI bus. The KZPSA-BB must be set to fast speed for TruCluster Server.

Note

You may have problems if the member system supports the `bus_probe_algorithm` console variable and it is not set to `new`. See Section 2.3.2.

The KZPBA-CB UltraSCSI host adapter:

- Is a high-performance PCI option connecting the PCI-based host system to the devices on a 16-bit, ultrawide differential SCSI bus.
- Is a single-channel, ultrawide differential adapter.
- Operates at the following speeds:
 - 5 MB/sec narrow SCSI at slow speed
 - 10 MB/sec narrow SCSI at fast speed
 - 20 MB/sec wide differential SCSI
 - 40 MB/sec wide differential UltraSCSI

Note

Even though the KZPBA-CB is an UltraSCSI device, it has an HD68 connector.

Use the steps in Table 9–2 to set up a KZPSA-BB or KZPBA-CB host bus adapter for a TruCluster Server cluster that uses radial connection to a DWZZH UltraSCSI hub with an HSZ40 or HSZ50 RAID array controller.

Table 9–2: Installing the KZPSA-BB or KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub

Step	Action	Refer to:
1	Ensure that the KZPSA-BB internal termination resistors, Z1, Z2, Z3, Z4, and Z5 are installed.	Section 9.1.4.4, Figure 9–1, and <i>KZPSA PCI-to-SCSI Storage Adapter Installation and User’s Guide</i>
	Ensure that the eight KZPBA-CB internal termination resistor SIPs, RM1-RM8 are installed.	Section 4.3.3.3, Figure 4–1, and <i>KZPBA-CB PCI-to-Ultra SCSI Differential Host Adapter User’s Guide</i>
2	Power down the system. Install a KZPSA-BB PCI-to-SCSI adapter or KZPBA-CB UltraSCSI host adapter in the PCI slot corresponding to the logical bus to be used for the shared SCSI bus. Ensure that the number of adapters is within limits for the system, and that the placement is acceptable.	<i>KZPSA PCI-to-SCSI Storage Adapter Installation and User’s Guide</i> and <i>KZPBA-CB PCI-to-Ultra SCSI Differential Host Adapter User’s Guide</i>
3	Install a BN38C or BN38D cable between the KZPSA-BB or KZPBA-CB UltraSCSI host adapter and a DWZZH port.	—

Notes

The maximum length of a SCSI bus segment is 25 meters (82 feet), including the bus length internal to the adapter and storage devices.

One end of the BN38C cable is 68-pin high density. The other end is 68-pin VHDCI. The DWZZH accepts the 68-pin VHDCI connector.

The number of member systems on the shared SCSI bus has to be one less than the number of DWZZH ports.

You may use a BN38E-0B technology adapter cable and a BN37A cable in place of the BN38C cable.

4	Power up the system, and update the system SRM console firmware and KZPSA-BB host bus adapter firmware from the latest Alpha Systems Firmware Update CD-ROM.	Firmware release notes for the system (Section 4.2) and Section 9.1.4.5
---	--	---

Note

The SRM console firmware includes the ISP1020/1040-based PCI option firmware, which includes the KZPBA-CB. When you update the SRM console firmware, you are enabling the KZPBA-CB firmware to be updated. On a powerup reset, the SRM console loads KZPBA-CB adapter firmware from the console system flash ROM into NVRAM for all Qlogic ISP1020/1040-based PCI options, including the KZPBA-CB PCI-to-Ultra SCSI adapter.

Table 9–2: Installing the KZPSA-BB or KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub (cont.)

Step	Action	Refer to:
5	Use the <code>show config</code> and <code>show device console</code> commands to display the installed devices and information about the KZPSA-BBs or KZPBA-CBs on the AlphaServer systems. Look for <code>KZPSA</code> or <code>pk*</code> in the display to determine which devices are KZPSA-BBs. Look for <code>QLogic ISP1020</code> in the <code>show config</code> display and <code>isp</code> in the <code>show device</code> display to determine which devices are KZPBA-CBs.	Section 9.1.3 and Example 9–1 through Example 9–4
6	Use the <code>show pk*</code> or <code>show isp*</code> console commands to determine the status of the KZPSA-BB or KZPBA-CB console environment variables, and then use the <code>set console</code> command to set the KZPSA-BB bus speed to fast, termination power to on, and the KZPSA or KZPBA-CB SCSI bus ID.	Section 9.1.4.1 through Section 9.1.4.3 and Example 9–6 through Example 9–9
Notes		
<p>Ensure that the SCSI ID that you use is distinct from all other SCSI IDs on the same shared SCSI bus. If you do not remember the other SCSI IDs, or do not have them recorded, you must determine these SCSI IDs.</p> <p>If you are using a DS-DWZZH-05, you cannot use SCSI ID 7 for a KZPSA-BB or KZPBA-CB host bus adapter; SCSI ID 7 is reserved for DS-DWZZH-05 use.</p> <p>If you are using a DS-DWZZH-05 and fair arbitration is enabled, you must use the SCSI ID assigned to the hub port the adapter is to be connected to.</p> <p>You will have problems, such as not being able to access devices, if you have two or more SCSI adapters at the same SCSI ID on any one SCSI bus.</p>		
7	Repeat steps 1 through 6 for any other KZPSA-BBs — or KZPBA-CBs to be installed on this shared SCSI bus on other member systems.	
8	Connect a DS-DWZZH-03 or DS-DWZZH-05 to an HSZ40 or HSZ50	Section 10.5

9.1.2 Installing a KZPSA-BB or KZPBA-CB Using External Termination

Use the steps in Table 9–3 to set up a KZPSA-BB or KZPBA-CB for a TruCluster Server cluster using the old method of external termination and Y cables.

Table 9–3: Installing a KZPSA-BB or KZPBA-CB for Use with External Termination

Step	Action	Refer to:
1	Remove the KZPSA-BB internal termination resistors, Z1, Z2, Z3, Z4, and Z5. Remove the eight KZPBA-CB internal termination resistor SIPs, RM1-RM8.	Section 9.1.4.4, Figure 9–1, and <i>KZPSA PCI-to-SCSI Storage Adapter Installation and User’s Guide</i> Section 4.3.3.3, Figure 4–1, and <i>KZPBA-CB PCI-to-Ultra SCSI Differential Host Adapter User’s Guide</i>
2	Power down the member system. Install a KZPSA-BB PCI-to-SCSI bus adapter or KZPBA-CB UltraSCSI host adapter in the PCI slot corresponding to the logical bus to be used for the shared SCSI bus. Ensure that the number of adapters is within limits for the system, and that the placement is acceptable.	<i>KZPSA PCI-to-SCSI Storage Adapter Installation and User’s Guide</i> and <i>KZPBA-CB PCI-to-Ultra SCSI Differential Host Adapter User’s Guide</i>
3	Install a BN21W-0B Y cable on each KZPSA-BB or KZPBA-CB host adapter.	—
4	Install an H879-AA terminator on one leg of the BN21W-0B Y cable of the member system that will be on the end of the shared SCSI bus.	—
5	Power up the system, and update the system SRM console firmware and KZPSA-BB host bus adapter firmware from the latest Alpha Systems Firmware Update CD-ROM.	Firmware release notes for the system (Section 4.2) and Section 9.1.4.5
Note		
The SRM console firmware includes the ISP1020/1040-based PCI option firmware, which includes the KZPBA-CB. When you update the SRM console firmware, you are enabling the KZPBA-CB firmware to be updated. On a powerup reset, the SRM console loads KZPBA-CB adapter firmware from the console system flash ROM into NVRAM for all Qlogic ISP1020/1040-based PCI options, including the KZPBA-CB PCI-to-Ultra SCSI adapter.		
6	Use the <code>show config</code> and <code>show device console</code> commands to display the installed devices and information about the KZPSA-BBs or KZPBA-CBs on the AlphaServer systems. Look for KZPSA or pk* in the display to determine which devices are KZPSA-BBs. Look for QLogic ISP1020 in the <code>show config</code> display and <code>isp</code> in the <code>show device</code> display to determine which devices are KZPBA-CBs.	Section 9.1.3 and Example 9–1 through Example 9–4

Table 9–3: Installing a KZPSA-BB or KZPBA-CB for Use with External Termination (cont.)

Step	Action	Refer to:
7	Use the <code>show pk*</code> or <code>show isp*</code> console commands to determine the status of the KZPSA-BB or KZPBA-CB console environment variables, and then use the <code>set console</code> command to set the KZPSA-BB bus speed to fast, termination power to on, and the KZPSA or KZPBA-CB SCSI bus ID.	Section 9.1.4.1 through Section 9.1.4.3 and Example 9–6 through Example 9–9
Notes		
Ensure that the SCSI ID that you use is distinct from all other SCSI IDs on the same shared SCSI bus. If you do not remember the other SCSI IDs, or do not have them recorded, you must determine these SCSI IDs. You will have problems, such as not being able to access devices, if you have two or more SCSI adapters at the same SCSI ID on any one SCSI bus.		
8	Repeat steps 1 through 7 for any other KZPSA-BBs or KZPBA-CBs to be installed on this shared SCSI bus on other member systems.	—
9	Install the remaining SCSI bus hardware needed for storage (DWZZA(B), RAID array controllers, storage shelves, cables, and terminators).	Section 10.4
	BA350 storage shelf	Section 10.3.1, Section 10.4.1.1, and Section 10.4.3.1
	Non-UltraSCSI BA356 storage shelf	Section 10.3.2.1, Section 10.4.1.2, and Section 10.4.3.2
	UltraSCSI BA356 storage shelf	Section 10.3.2.2, Section 10.4.1.3, and Section 10.4.3.3
	HSZ40 or HSZ50 RAID array controller	Section 10.4.4
	RAID Array 3000	Section 10.4.5
10	Install the tape device hardware and cables on the shared SCSI bus as follows:	Chapter 8
	TZ88	Section 8.1
	TZ89	Section 8.2
	Compaq 20/40 GB DLT Tape Drive	Section 8.3
	TZ885	Section 8.5
	TZ887	Section 8.6

Table 9–3: Installing a KZPSA-BB or KZPBA-CB for Use with External Termination (cont.)

Step	Action	Refer to:
	TL891/TL892 MiniLibrary	Section 8.7
	TL890 with TL891/TL892	Section 8.8
	TL894	Section 8.9
	TL895	Section 8.10
	TL893/TL896	Section 8.11
	TL881/TL891 DLT MiniLibraries	Section 8.12
	Compaq ESL9326D Enterprise Library	Section 8.13

Notes

If you install tape devices on the shared SCSI buses, ensure that you understand how the particular tape devices affect the shared SCSI bus. The TL893, TL894, TL895, TL896, and ESL9326D have long internal SCSI cables; therefore, they cannot be externally terminated with a tralink/terminator combination.

These tape libraries must be on the end of the shared SCSI bus.

We recommend that tape devices be placed on a separate shared SCSI bus.

9.1.3 Displaying KZPSA-BB and KZPBA-CB Adapters with the show Console Commands

Use the `show config` and `show device console` commands to display system configuration. Use the output to determine which devices are KZPSA-BBs or KZPBA-CBs, and to determine their SCSI bus IDs.

Example 9–1 shows the output from the `show config` console command on an AlphaServer 4100 system.

Example 9–1: Displaying Configuration on an AlphaServer 4100

```
P00>>> show config
                               Compaq Computer Corporation
                               AlphaServer 4x00

Console V5.1-3  OpenVMS PALcode V1.19-14, Tru64 UNIX PALcode V1.21-22

Module                Type      Rev      Name
System Motherboard    0        0000    mthrbrd0
Memory 64 MB SYNC     0        0000    mem0
Memory 64 MB SYNC     0        0000    mem1
```

Example 9-1: Displaying Configuration on an AlphaServer 4100 (cont.)

```
Memory 64 MB SYNC          0      0000  mem2
Memory 64 MB SYNC          0      0000  mem3
CPU (4MB Cache)           3      0000  cpu0
CPU (4MB Cache)           3      0000  cpu1
Bridge (IOD0/IOD1)        600    0021  iod0/iod1
PCI Motherboard            8      0000  saddle0
```

Bus 0 iod0 (PCI0)

Slot	Option Name	Type	Rev	Name
1	PCEB	4828086	0005	pceb0
2	S3 Trio64/Trio32	88115333	0000	vga0
3	DECchip 21040-AA	21011	0024	tulip0
4	DEC KZPSA	81011	0000	pks1
5	DEC PCI MC	181011	000B	mc0

Bus 1 pceb0 (EISA Bridge connected to iod0, slot 1)

Slot	Option Name	Type	Rev	Name
------	-------------	------	-----	------

Bus 0 iod1 (PCI1)

Slot	Option Name	Type	Rev	Name
1	NCR 53C810	11000	0002	ncr0
2	NCR 53C810	11000	0002	ncr1
3	QLogic ISP1020	10201077	0005	isp0
4	QLogic ISP1020	10201077	0005	isp1
5	DEC KZPSA	81011	0000	pks0

Example 9-2 shows the output from the `show device console` command entered on an AlphaServer 4100 system.

Example 9-2: Displaying Devices on an AlphaServer 4100

```
P00>>> show device
polling ncr0 (NCR 53C810) slot 1, bus0 PCI, hose 1 SCSI Bus ID 7
dka500.5.0.1.1 Dka500 RRD45 1645
polling ncr1 (NCR 53C810) slot 2, bus0 PCI, hose 1 SCSI Bus ID 7
dkb0.0.0.2.1 DKb0 RZ29B 0007
dkb100.1.0.2.1 DKb100 RZ29B 0007
polling isp0 (QLogic ISP1020) slot 3, bus 0 PCI, hose 1 SCSI Bus ID 7
dkc0.0.0.3.1 DKc0 HSZ70 V70Z
dkc1.0.0.3.1 DKc1 HSZ70 V70Z
dkc2.0.0.3.1 DKc2 HSZ70 V70Z
dkc3.0.0.3.1 DKc3 HSZ70 V70Z
dkc4.4.0.3.1 DKc4 HSZ70 V70Z
dkc5.0.0.3.1 DKc5 HSZ70 V70Z
dkc6.0.0.3.1 DKc6 HSZ70 V70Z
dkc100.1.0.3.1 DKc100 RZ28M 0568
dkc200.2.0.3.1 DKc200 RZ28M 0568
dkc300.3.0.3.1 DKc300 RZ28 442D
polling isp1 (QLogic ISP1020) slot 4, bus 0 PCI, hose 1 SCSI Bus ID 7
dkd0.0.0.4.1 DKd0 HSZ50-AX X29Z
dkd1.0.0.4.1 DKd1 HSZ50-AX X29Z
```

Example 9–2: Displaying Devices on an AlphaServer 4100 (cont.)

```
dkd2.0.0.4.1      DKd2          HSZ50-AX X29Z
dkd100.1.0.4.1   DKd100       RZ26N   0568
dkd200.1.0.4.1   DKd200       RZ26    392A
dkd300.1.0.4.1   DKd300       RZ26N   0568
polling kzpsa0 (DEC KZPSA) slot 5, bus 0 PCI, hose 1 TPwr 1 Fast 1 Bus ID 7
kzpsa0.7.0.5.1   dke TPwr 1 Fast 1 Bus ID 7 L01 A11
dke100.1.0.5.1   DKe100       RZ28    442D
dke200.2.0.5.1   DKe200       RZ26    392A
dke300.3.0.5.1   DKe300       RZ26L   442D
polling floppy0 (FLOPPY) pceb IBUS hose 0
dva0.0.0.1000.0 DVA0         RX23
polling kzpsa1 (DEC KZPSA) slot 4, bus 0 PCI, hose 0 TPwr 1 Fast 1 Bus ID 7
kzpsa1.7.0.4.1   dkf TPwr 1 Fast 1 Bus ID 7 E01 A11
dkf100.1.0.5.1   DKf100       RZ26    392A
dkf200.2.0.5.1   DKf200       RZ28    442D
dkf300.3.0.5.1   DKf300       RZ26    392A
polling tulip0 (DECchip 21040-AA) slot 3, bus 0 PCI, hose 0
ewa0.0.0.3.0     00-00-F8-21-0B-56 Twisted-Pair
```

Example 9–3 shows the output from the `show config` console command entered on an AlphaServer 8200 system.

Example 9–3: Displaying Configuration on an AlphaServer 8200

```
>>> show config
      Name                Type    Rev    Mnemonic
-----
TLSB
4++   KN7CC-AB             8014    0000    kn7cc-ab0
5+    MS7CC                  5000    0000    ms7cc0
8+    KFTIA                   2020    0000    kftia0

C0 Internal PCI connected to kftia0 pci0
0+   QLogic ISP1020 10201077 0001    isp0
1+   QLogic ISP1020 10201077 0001    isp1
2+   DECchip 21040-AA 21011 0023    tulip0
4+   QLogic ISP1020 10201077 0001    isp2
5+   QLogic ISP1020 10201077 0001    isp3
6+   DECchip 21040-AA 21011 0023    tulip1

C1 PCI connected to kftia0
0+   KZPAA                   11000   0001    kzpaa0
1+   QLogic ISP1020 10201077 0005    isp4
2+   KZPSA                   81011   0000    kzpsa0
3+   KZPSA                   81011   0000    kzpsa1
4+   KZPSA                   81011   0000    kzpsa2
7+   DEC PCI MC              181011  000B    mc0
```

Example 9–4 shows the output from the `show device` console command entered on an AlphaServer 8200 system.

Example 9-4: Displaying Devices on an AlphaServer 8200

```
>>> show device
polling for units on isp0, slot0, bus0, hose0...
polling for units on isp1, slot1, bus0, hose0...
polling for units on isp2, slot4, bus0, hose0...
polling for units on isp3, slot5, bus0, hose0...
polling for units kzpsaa0, slot0, bus0, hose1...
pke0.7.0.0.1      kzpsaa4      SCSI Bus ID 7
dke0.0.0.0.1      DKE0          RZ28      442D
dke200.2.0.0.1    DKE200        RZ28      442D
dke400.4.0.0.1    DKE400        RRD43     0064

polling for units isp4, slot1, bus0, hose1...
dkf0.0.0.1.1      DKF0          HSZ70     V70Z
dkf1.0.0.1.1      DKF1          HSZ70     V70Z
dkf2.0.0.1.1      DKF2          HSZ70     V70Z
dkf3.0.0.1.1      DKF3          HSZ70     V70Z
dkf4.0.0.1.1      DKF4          HSZ70     V70Z
dkf5.0.0.1.1      DKF5          HSZ70     V70Z
dkf6.0.0.1.1      DKF6          HSZ70     V70Z
dkf100.1.0.1.1    DKF100        RZ28M     0568
dkf200.2.0.1.1    DKF200        RZ28M     0568
dkf300.3.0.1.1    DKF300        RZ28      442D

polling for units on kzpsa0, slot 2, bus 0, hose1...
kzpsa0.4.0.2.1    dkg          TPwr 1 Fast 1 Bus ID 7 L01 A11
dkg0.0.0.2.1      DKG0          HSZ50-AX X29Z
dkg1.0.0.2.1      DKG1          HSZ50-AX X29Z
dkg2.0.0.2.1      DKG2          HSZ50-AX X29Z
dkg100.1.0.2.1    DKG100        RZ26N     0568
dkg200.2.0.2.1    DKG200        RZ28      392A
dkg300.3.0.2.1    DKG300        RZ26N     0568

polling for units on kzpsa1, slot 3, bus 0, hose1...
kzpsa1.4.0.3.1    dkh          TPwr 1 Fast 1 Bus ID 7 L01 A11
dkh100.1.0.3.1    DKH100        RZ28      442D
dkh200.2.0.3.1    DKH200        RZ26      392A
dkh300.3.0.3.1    DKH300        RZ26L     442D

polling for units on kzpsa2, slot 4, bus 0, hose1...
kzpsa2.4.0.4.1    dki          TPwr 1 Fast 1 Bus ID 7 L01 A10
dki100.1.0.3.1    DKI100        RZ26      392A
dki200.2.0.3.1    DKI200        RZ28      442C
dki300.3.0.3.1    DKI300        RZ26      392A
```

9.1.4 Displaying Console Environment Variables and Setting the KZPSA-BB and KZPBA-CB SCSI ID

The following sections show how to use the `show console` command to display the `pk*` and `isp*` console environment variables and set the KZPSA-BB and KZPBA-CB SCSI ID on various AlphaServer systems. Use these examples as guides for your system.

Note that the console environment variables used for the SCSI options vary from system to system. Also, a class of environment variables (for example, `pk*` or `isp*`) may show both internal and external options.

Compare the following examples with the devices shown in the `show config` and `show dev` examples to determine which devices are KZPSA-BBs or KZPBA-CBs on the shared SCSI bus.

9.1.4.1 Displaying KZPSA-BB and KZPBA-CB `pk*` or `isp*` Console Environment Variables

To determine the console environment variables to use, execute the `show pk*` and `show isp*` console commands.

Example 9–5 shows the `pk` console environment variables for an AlphaServer 4100.

Example 9–5: Displaying the `pk*` Console Environment Variables on an AlphaServer 4100 System

```
P00>>>show pk*
pka0_disconnect      1
pka0_fast            1
pka0_host_id         7

pkb0_disconnect      1
pkb0_fast            1
pkb0_host_id         7

pkc0_host_id         7
pkc0_soft_term       diff

pkd0_host_id         7
pkd0_soft_term       on

pke0_fast            1
pke0_host_id         7
pke0_termpr         1
```

Example 9–5: Displaying the pk* Console Environment Variables on an AlphaServer 4100 System (cont.)

```
pkf0_fast          1
pkf0_host_id       7
pkf0_termprw      1
```

Compare the `show pk*` command display in Example 9–5 with the `show config` command in Example 9–1 and the `show dev` command in Example 9–2. Note that there are no `pk*` devices in either display. Example 9–2 shows:

- The NCR 53C810 SCSI controllers as `ncr0` and `ncr1` with disk `DKa` and `DKb` (`pka` and `pkb`)
- The Qlogic ISP1020 devices (KZPBA-CBs) as `isp0` and `isp1` with disks `DKc` and `DKd` (`pkc` and `pkd`)
- The KZPSA-BBs with disks `DKe` and `DKf` (`pke` and `pkf`)

Example 9–5 shows two `pk*0_soft_term` environment variables; `pkc0_soft_term` which is `on`, and `pkd0_soft_term` which is `diff`.

The `pk*0_soft_term` environment variable applies to systems using the QLogic ISP1020 SCSI controller, which implements the 16-bit wide SCSI bus and uses dynamic termination.

The QLogic ISP1020 module has two terminators, one for the 8 low bits and one for the high 8 bits. There are five possible values for `pk*0_soft_term`:

- `off` — Turns off both the low 8 bits and high 8 bits
- `low` — Turns on the low 8 bits and turns off the high 8 bits
- `high` — Turns on the high 8 bits and turns off the low 8 bits
- `on` — Turns on both the low 8 bits and high 8 bits
- `diff` — Places the bus in differential mode

The KZPBA-CB is a Qlogic ISP1040 module, and its termination is determined by the presence or absence of internal termination resistor SIPs RM1-RM8. Therefore, the `pkb0_soft_term` environment variable has no meaning and it may be ignored.

Example 9–6 shows the use of the `show isp` console command to display the console environment variables for KZPBA-CBs on an AlphaServer 8x00.

Example 9–6: Displaying Console Variables for a KZPBA-CB on an AlphaServer 8x00 System

```
P00>>> show isp*
isp0_host_id      7
isp0_soft_term    on

isp1_host_id      7
isp1_soft_term    on

isp2_host_id      7
isp2_soft_term    on

isp3_host_id      7
isp3_soft_term    on

isp5_host_id      7
isp5_soft_term    diff
```

Both Example 9–3 and Example 9–4 show five `isp` devices; `isp0`, `isp1`, `isp2`, `isp3`, and `isp4`. In Example 9–6, the `show isp*` console command shows `isp0`, `isp1`, `isp2`, `isp3`, and `isp5`.

The console code that assigns console environment variables counts every I/O adapter including the `KZPAA`, which is the device after `isp3`, and therefore logically `isp4` in the numbering scheme. The `show isp` console command skips over `isp4` because the `KZPAA` is not a QLogic 1020/1040 class module.

Example 9–3 and Example 9–4 show that `isp0`, `isp1`, `isp2`, and `isp3` are on the internal `KFTIA` PCI bus and not on a shared SCSI bus. Only `isp5`, the `KZPBA-CB`, is on a shared SCSI bus. The other three shared SCSI buses use `KZPSA-BBs`.

Example 9–7 shows the use of the `show pk` console command to display the console environment variables for `KZPSA-BBs` on an AlphaServer 8x00.

Example 9–7: Displaying Console Variables for a KZPSA-BB on an AlphaServer 8x00 System

```
P00>>> show pk*
pka0_fast         1
pka0_host_id      7
pka0_termpr       on

pkb0_fast         1
pkb0_host_id      7
pkb0_termpr       on
```

Example 9–7: Displaying Console Variables for a KZPSA-BB on an AlphaServer 8x00 System (cont.)

```
pkc0_fast      1
pkc0_host_id   7
pkc0_termprw  on
```

9.1.4.2 Setting the KZPBA-CB SCSI ID

After you determine the console environment variables for the KZPBA-CBs on the shared SCSI bus, use the `set console` command to set the SCSI ID. For a TruCluster Server cluster, you will most likely have to set the SCSI ID for all KZPBA-CB UltraSCSI adapters except one. If you are using a DS-DWZZH-05 with fair arbitration enabled, you will have to set the SCSI IDs for all KZPBA-CB UltraSCSI adapters.

Note

You will have problems if you have two or more SCSI adapters at the same SCSI ID on any one SCSI bus.

If you are using a DS-DWZZH-05, you cannot use SCSI ID 7 for a KZPBA-CB UltraSCSI adapter; SCSI ID 7 is reserved for DS-DWZZH-05 use.

If DS-DWZZH-05 fair arbitration is enabled, the SCSI ID of the host adapter must match the SCSI ID assigned to the hub port. Mismatching or duplicating SCSI IDs will cause the hub to hang.

Use the `set console` command as shown in Example 9–8 to set the KZPBA-CB SCSI ID. In this example, the SCSI ID is set for KZPBA-CB `pkc` on the AlphaServer 4100 shown in Example 9–5.

Example 9–8: Setting the KZPBA-CB SCSI Bus ID

```
P00>>> show pkc0_host_id
7
P00>>> set pkc0_host_id 6
P00>>> show pkc0_host_id
6
```

9.1.4.3 Setting KZPSA-BB SCSI Bus ID, Bus Speed, and Termination Power

If the KZPSA-BB SCSI ID is not correct, or if it was reset to 7 by the firmware update utility, or you need to change the KZPSA-BB speed, or enable termination power, use the `set` console command.

Note

All KZPSA-BB host bus adapters should be enabled to generate termination power.

Set the SCSI bus ID with the `set` command as shown in the following example:

```
>>> set pkn_0_host_id #
```

The *n* specifies which KZPSA-BB the environment variables apply to. You obtain the *n* value from the `show device` and `show pk*` console commands. The number sign (#) is the SCSI bus ID for the KZPSA.

Set the bus speed with the `set` command as shown in the following example:

```
>>> set pkn0_fast #
```

The number sign (#) specifies the bus speed. Use a 0 for slow and a 1 for fast.

Enable SCSI bus termination power with the `set` command as shown in the following example:

```
>>> set pkn0_termpr on
```

Example 9–9 shows how to determine the present SCSI ID, bus speed, and the status of termination power, and then set the KZPSA-BB SCSI ID to 6 and bus speed to fast for `pkb0`.


Example 9–9: Setting KZPSA-BB SCSI Bus ID and Speed

```
P00>>> show pkb*
pkb0_fast 0
pkb0_host_id 7
pkb0_termpr on
P00>>> set pkb0_host_id 6
P00>>> set pkb0_fast 1
P00>>> show pkb0_host_id
6
P00>>> show pkb0_fast
1
```

9.1.4.4 KZPSA-BB and KZPBA-CB Termination Resistors

The KZPSA-BB internal termination is disabled by removing termination resistors Z1 through Z5, as shown in Figure 9-1.

Figure 9-1: KZPSA-BB Termination Resistors



The KZPBA-CB internal termination is disabled by removing the termination resistors RM1-RM8 as shown in Figure 4-1.

9.1.4.5 Updating the KZPSA-BB Adapter Firmware

You must check, and update as necessary, the system and host bus adapter firmware. The firmware may be out-of-date. Read the firmware release notes from the AlphaSystems Firmware Update CD-ROM for the applicable system/SCSI adapter.

If the System Reference Manual (SRM) console or KZPSA-BB firmware is not current, boot the Loadable Firmware Update (LFU) utility from the Alpha Systems Firmware Update CD-ROM. Choose the update entry from the list of LFU commands. LFU can update all devices or any particular device you select.

When you boot the Systems Firmware Update CD-ROM, you can read the firmware release notes. After booting has completed, enter `read_rel_notes` at the `UPD>` prompt. You can also copy and print the release notes as shown in Section 4.2.

To update the firmware, boot the LFU utility from the Alpha Systems Firmware Update CD-ROM.

It is not necessary to use the `-flag` option to the `boot` command. Insert the Alpha Systems Firmware Update CD-ROM and boot. For example, to boot from `dka600`:

```
P00>>> boot dka600
```

The boot sequence provides firmware update overview information. Use Return to scroll the text, or press Ctrl/C to skip the text.

After the overview information has been displayed, the name of the default boot file is provided. If it is the correct boot file, press Return at the Bootfile: prompt. Otherwise, enter the name of the file you want to boot from.

The firmware images are copied from the CD-ROM and the LFU help message shown in the following example is displayed:

```
*****Loadable Firmware Update Utility*****

-----
Function      Description
-----
Display      Displays the system's configuration table.
Exit         Done exit LFU (reset).
List         Lists the device, revision, firmware name and
            update revision
Readme       Lists important release information.
Update       Replaces current firmware with loadable data
            image.
Verify       Compares loadable and hardware images.
? or Help    Scrolls this function table.
```

The list command indicates, in the device column, which devices it can update.

Use the update command to update all firmware, or you can designate a specific device to update; for example, KZPSA-BB pkb0:

```
UPD> update pkb0
```

After updating the firmware and verifying this with the verify command, reset the system by cycling the power.

10

Configurations Using External Termination or Radial Connections to Non-UltraSCSI Devices

This chapter describes the requirements for the shared SCSI bus using:

- Externally terminated TruCluster Server configurations
- Radial configurations with non-UltraSCSI RAID array controllers

In addition to using only the supported hardware, adhering to the requirements described in this chapter will ensure that your cluster operates correctly.

This chapter discusses the following topics:

- SCSI bus signal converters (Section 10.1)
- SCSI bus termination in externally terminated TruCluster Server configurations (Section 10.2)
- Overview of the BA350, BA356, and UltraSCSI BA356 disk storage shelves (Section 10.3)
- How to configure shared storage for external termination using Y cables and trilinks (Section 10.4):
 - Preparing shared storage for an externally terminated TruCluster Server configuration (Section 10.4.1)
 - Cabling a single storage shelf (Section 10.4.2)
 - Connecting multiple storage shelves, for instance a BA350 and a BA356, two BA356s, or two UltraSCSI BA356s (Section 10.4.3)
 - Using the HSZ20, HSZ40, or HSZ50 RAID array controllers (Section 10.4.4)
 - Cabling a RAID Array 3000 (RA3000) with HSZ22 controller using external termination (Section 10.4.5)
- How to configure shared storage for radial configurations using the HSZ40 or HSZ50 RAID array controllers (Section 10.5)

Introductory information covering SCSI bus configuration concepts (SCSI bus speed, data path, and so on) and SCSI bus configuration requirements can be found in Chapter 3.

10.1 Using SCSI Bus Signal Converters

A SCSI bus signal converter allows you to couple a differential bus segment to a single-ended bus segment, allowing the mixing of differential and single-ended devices on the same SCSI bus to isolate bus segments for maintenance purposes.

Each SCSI signal converter has a single-ended side and a differential side as follows:

- DWZZA — 8-bit data path
- DWZZB — 16-bit data path
- DS-BA35X-DA 16-bit personality module

Note

Some UltraSCSI documentation uses the UltraSCSI term bus expander when referring to the DWZZB and UltraSCSI signal converters. Other UltraSCSI documentation refers to some UltraSCSI products as bus extender/converters.

For TruCluster Server there are no supported standalone UltraSCSI bus expanders (DWZZC).

In this manual, any device that converts a differential signal to a single-ended signal is referred to as a signal converter (the DS-BA35X-DA personality module contains a DWZZA-on-a-chip or DOC chip).

A SCSI signal converter is required when you want to connect devices with different transmission modes.

10.1.1 Types of SCSI Bus Signal Converters

Signal converters can be standalone units or StorageWorks building blocks (SBBs) that are installed in a storage shelf disk slot. You must use the signal converter module that is appropriate for your hardware configuration.

For example, use a DWZZA-VA signal converter to connect a wide, differential host bus adapter to a BA350 (single-ended and narrow) storage shelf, but use a DWZZB-VW signal converter to connect a wide, differential

host bus adapter to a non-UltraSCSI BA356 (single-ended and wide) storage shelf. The DS-BA35X-DA personality module is used in an UltraSCSI BA356 to connect an UltraSCSI host bus adapter to the single-ended disks in the UltraSCSI BA356. You can install a DWZZB-VW in an UltraSCSI BA356, but you will waste a disk slot and it will not work with a KZPBA-CB if there are any UltraSCSI disks in the storage shelves.

The following sections discuss the DWZZA and DWZZB signal converters and the DS-BA35X-DA personality module.

10.1.2 Using the SCSI Bus Signal Converters

The DWZZA and DWZZB signal converters are used in the BA350 and BA356 storage shelves. They have removable termination. The DS-BA35X-DA personality module is used in the UltraSCSI BA356. It has switch selectable differential termination. The single-ended termination is active termination.

The following sections describe termination for these signal converters in more detail.


10.1.2.1 DWZZA and DWZZB Signal Converter Termination

Both the single-ended side and the differential side of each DWZZA and DWZZB signal converter have removable termination. To use a signal converter, you must remove the termination in the differential side and attach a tralink connector to this side. To remove the differential termination, remove the five 14-pin termination resistor SIPs (located near the differential end of the signal converter). You can attach a terminator to the tralink connector to terminate the differential bus. If you detach the tralink connector from the signal converter, the shared SCSI bus is still terminated (provided there is termination power).

You must keep the termination in the single-ended side to provide termination for one end of the BA350 or BA356 single-ended SCSI bus segment. Verify that the termination is active. A DWZZA should have jumper J2 installed. Jumpers W1 and W2 should be installed in a DWZZB.

Figure 10–1 shows the status of internal termination for a standalone SCSI signal converter that has a tralink connector attached to the differential side.


Figure 10–1: Standalone SCSI Signal Converter



ZK-1050U-AI

Figure 10–2 shows the status of internal termination for an SBB SCSI signal converter that has a trilink connector attached to the differential side.

Figure 10–2: SBB SCSI Signal Converter



ZK-1576U-AI

10.1.2.2 DS-BA35X-DA Termination

The UltraSCSI BA356 shelf uses a 16-bit differential UltraSCSI personality module (DS-BA35X-DA) as the interface between the UltraSCSI differential bus and the UltraSCSI single-ended bus in the UltraSCSI BA356.

The personality module controls termination for the external differential UltraSCSI bus segment, and for both ends of the internal single-ended bus segment.

For normal cluster operation, the differential termination must be disabled since a trilink connector will be installed on personality module connector JA1, allowing the use of the UltraSCSI BA356 (or two UltraSCSI BA356s) in the middle of the bus or external termination for an UltraSCSI BA356 on the end of the bus.

Switch pack 4 switches S4-1 and S4-2 are set to ON to disable the personality module differential termination. The switches have no effect on the BA356 internal, single-ended UltraSCSI bus termination.


Notes

S4-3 and S4-4 have no function on the DS-BA35X-DA personality module.

See Section 10.3.2.2 for information on how to select the device SCSI IDs in an UltraSCSI BA356.

Figure 10–3 shows the relative positions of the two DS-BA35X-DA switch packs.

Figure 10–3: DS-BA35X-DA Personality Module Switches



ZK-1411U-AI

10.2 Terminating the Shared SCSI Bus

You must properly connect devices to a shared SCSI bus. In addition, you can terminate only the beginning and end of each SCSI bus segment (either single-ended or differential).

There are two rules for SCSI bus termination:

- There are only two terminators for each SCSI bus segment.
- If you do not use an UltraSCSI hub, bus termination must be external. You may use external termination with an UltraSCSI hub, but we do not recommend it.

Whenever possible, connect devices to a shared bus so that they can be isolated from the bus. This allows you to disconnect devices from the bus for maintenance purposes without affecting bus termination and cluster operation. You also can set up a shared SCSI bus so that you can connect additional devices at a later time without affecting bus termination.

Notes

With the exception of the TZ885, TZ887, TL890, TL891, and TL892, tape devices can only be installed at the end of a shared SCSI bus. These tape devices are the only supported tape devices that can be terminated externally.

We recommend that tape loaders be on a separate shared SCSI bus to allow normal shared SCSI bus termination for those shared SCSI buses without tape loaders.

Most devices have internal termination. For example, the KZPSA and KZPBA host bus adapters, BA350 and BA356 storage shelves, and the DWZZA and DWZZB SCSI bus signal converters have internal termination. Depending on how you set up a shared bus, you may have to enable or disable device termination.

Unless you are using an UltraSCSI hub, if you use a device's internal termination to terminate a shared bus, and you disconnect the bus cable from the device, the bus will not be terminated and cluster operation will be impaired. Therefore, unless you use an UltraSCSI hub, you must use external termination, enabling you to detach the device without affecting bus termination. The use of UltraSCSI hubs with UltraSCSI devices is discussed in Section 3.5 and Section 3.6. The use of a DS-DWZZH-03 UltraSCSI hub with externally terminated host bus adapters is discussed in Section 10.4.4.

To be able to externally terminate a bus and connect and disconnect devices without affecting bus termination, remove the device termination and use Y cables or tralink connectors to connect a device to a shared SCSI bus.

By attaching a Y cable or tralink connector to an unterminated device, you can locate the device in the middle or at the end of the shared bus. If the device is at the end of a bus, attach an H879-AA terminator to the BN21W-0B Y cable or H885-AA tralink connector to terminate the bus. For UltraSCSI devices, attach an H8863-AA terminator to the H8861 tralink connector. If you disconnect the Y cable or tralink connector from the device, the shared bus is still terminated and the shared SCSI bus is still operable.

In addition, you can attach a Y cable or a tralink connector to a properly terminated shared bus without connecting the Y cable or tralink connector to a device. If you do this, you can connect a device to the Y cable or tralink

connector at a later time without affecting bus termination. This allows you to expand your configuration without shutting down the cluster.

Figure 10–4 shows a BN21W-0B Y cable, which you may attach to a KZPSA-BB or KZPBA-CB SCSI adapter that has had its onboard termination removed. You can also use the BN21W-0B Y cable with a HSZ40 or HSZ50 controller or the unterminated differential side of a SCSI signal converter.

Note

You will normally use a Y cable on a KZPSA-BB or KZPBA-CB host bus adapter where there is not room for an H885-AA tralink, and a tralink connector elsewhere.

Figure 10–4: BN21W-0B Y Cable





Figure 10–5 shows an HD68 tralink connector (H885-AA), which you may attach to a KZPSA-BB or KZPBA-CB adapter that has its onboard termination removed, an HSZ40 or HSZ50 controller, or the unterminated differential side of a SCSI signal converter.

Figure 10–5: HD68 Trilink Connector (H885-AA)



Note

If you connect a trilink connector to a SCSI bus adapter, you may block access to an adjacent PCI slot. If this occurs, use a Y cable instead of the trilink connector. This is the case with the KZPBA-CB and KZPSA-BB SCSI adapters on some AlphaServer systems.

Use the H879-AA terminator to terminate one leg of a BN21W-0B Y cable or H885-AA trilink.

Use an H8861-AA VHDCI trilink connector (see Figure 3–1) with a DS-BA35X-DA personality module to daisy chain two UltraSCSI BA356s or to terminate external to the UltraSCSI BA356 storage shelf. Use the H8863-AA VHDCI terminator with the H8861-AA trilink connector.

10.3 Overview of Disk Storage Shelves

The following sections provide an introduction to the BA350, BA356, and UltraSCSI BA356 disk storage shelves.

10.3.1 BA350 Storage Shelf

Up to seven narrow (8-bit) single-ended StorageWorks building blocks (SBBs) can be installed in the BA350. Their SCSI IDs are based upon the slot they are installed in. For instance, a disk installed in BA350 slot 0 has SCSI ID 0, a disk installed in BA350 slot 1 has SCSI ID 1, and so forth.

Note

Do not install disks in the slots corresponding to the host SCSI IDs (usually SCSI ID 6 and 7 for a two-node cluster).

You use a DWZZA-VA as the interface between the wide, differential shared SCSI bus and the BA350 narrow, single-ended SCSI bus segment.

Note


Do not use a DWZZB-VW in a BA350. The use of the wide DWZZB-VW on the narrow single-ended bus will result in unterminated data lines in the DWZZB-VW, which will cause SCSI bus errors.

The BA350 storage shelf contains internal SCSI bus termination and a SCSI bus jumper. The jumper is not removed during normal operation.

The BA350 can be set up for two-bus operation, but that option is not very useful for a shared SCSI bus and is not covered in this manual.

Figure 10–6 shows the relative locations of the BA350 SCSI bus terminator and SCSI bus jumper. They are accessed from the rear of the box. For operation within a TruCluster Server cluster, both the J jumper and T terminator must be installed.

Figure 10–6: BA350 Internal SCSI Bus



ZK-1338U-AI

10.3.2 BA356 Storage Shelf

There are two variations of the BA356 used in TruCluster Server clusters: the BA356 (non-UltraSCSI BA356) and the UltraSCSI BA356.

An example of the non-UltraSCSI BA356 is the BA356-KC, which has a wide, single-ended internal SCSI bus. It has a BA35X-MH 16-bit personality module (only used for SCSI ID selection) and a 150-watt power supply. It is referred to as the non-UltraSCSI BA356 or BA356 in this manual. You use a DWZZB-VW as the interface between the wide, differential shared SCSI bus and the BA356 wide, single-ended SCSI bus segment.

10.3.2.1 Non-UltraSCSI BA356 Storage Shelf

The non-UltraSCSI BA356, like the BA350, can hold up to seven StorageWorks building blocks (SBBs). However, unlike the BA350, these SBBs are wide devices and can therefore support up to 16 disks (in two BA356 shelves). Also, like the BA350, the SBB SCSI IDs are based upon the slot they are installed in. The switches on the personality module (BA35X-MH) determine whether the disks respond to SCSI IDs 0 through 6 (slot 7 is the power supply) or 8 through 14 (slot 15 is the power supply). To

select SCSI IDs 0 through 6, set the personality module address switches 1 through 7 to off. To select SCSI IDs 8 through 14, set personality module address switches 1 through 3 to on and switches 4 through 7 to off.

Figure 10–7 shows the relative location of the BA356 SCSI bus jumper, BA35X-MF. The jumper is accessed from the rear of the box. For operation within a TruCluster Server cluster, you must install the J jumper in the normal position, behind slot 6. The SCSI bus jumper is not in the same position in the BA356 as in the BA350.

Termination for the BA356 single-ended bus is on the personality module, and is active unless a cable is installed on JB1 to daisy chain the single-ended SCSI bus in two BA356 storage shelves together. In this case, when the cable is connected to JB1, the personality module terminator is disabled.

Daisy chaining the single-ended bus between two BA356s is not used in clusters. We use DWZZB-VWs (with an attached H885-AA tralink connector) in each BA356 to connect the wide-differential connection from the host adapters to both BA356s in parallel. The switches on the personality module of one BA356 are set for SCSI IDs 0 through 7 and the switches on the personality module of the other BA356 are set for SCSI IDs 8 through 14.

Note


Do not install a narrow disk in a BA356 that is enabled for SCSI IDs 8 through 14. The SCSI bus will not operate correctly because the narrow disks cannot recognize wide addresses.

Like the BA350, you can set up the BA356 for two-bus operation by installing a SCSI bus terminator (BA35X-ME) in place of the SCSI bus jumper. However, like the BA350, two-bus operation in the BA356 is not very useful for a TruCluster Server cluster.

You can use the position behind slot 1 in the BA356 to store the SCSI bus terminator or jumper.

Figure 10–7 shows the relative locations of the BA356 SCSI bus jumper and the position for storing the SCSI bus jumper, if you do install the terminator. For operation within a TruCluster Server cluster, you must install the J jumper.

Figure 10–7: BA356 Internal SCSI Bus




ZK-1339U-AI

JA1 and JB1 are located on the personality module (in the top of the box when it is standing vertically). JB1, on the front of the module, is visible. JA1 is on the left side of the personality module as you face the front of the BA356, and is hidden from the normal view.

To determine if a jumper module or terminator module is installed in a BA356, remove the devices from slots 1 and 6 and note the following pin locations (see Figure 10–8):

- The identification pin on a jumper module aligns with the top hole in the backplane.
- The identification pin on a terminator module aligns with the bottom hole in the backplane.

Figure 10–8: BA356 Jumper and Terminator Module Identification Pins



10.3.2.2 UltraSCSI BA356 Storage Shelf

The UltraSCSI BA356 (DS-BA356-JF or DS-BA356-KH) has a single-ended, wide UltraSCSI bus. The DS-BA35X-DA personality module provides the interface between the internal, single-ended UltraSCSI bus segment and the shared, wide, differential UltraSCSI bus. The UltraSCSI BA356 uses a 180-watt power supply.

An older, non-UltraSCSI BA356 that has been retrofitted with a BA35X-HH 180-watt power supply and DS-BA35X-DA personality module is still only FCC certified for Fast 10 configurations. (See Section 3.2.4 for a discussion on bus speed.)

The UltraSCSI BA356 can hold up to seven StorageWorks building blocks (SBBs). These SBBs are UltraSCSI single-ended wide devices. The disk SCSI IDs are based upon the slot they are installed in. The S3 switches on the personality module (DS-BA35X-DA) determine whether the disks respond to SCSI IDs 0 through 6 (slot 7 is the power supply) or 8 through 14 (slot 15 is the power supply). To select SCSI IDs 0 through 6, set switches S3-1 through S3-7 to off. To select SCSI IDs 8 through 14, set personality module address switches S3-1 through S3-3 to on and switches S3-4 through S3-7 to off.

The jumper module is positioned behind slot 6 as with the non-UltraSCSI BA356 shown in Figure 10–7. For operation within a TruCluster Server cluster, you must install the J jumper. You verify the presence or absence of the jumper or terminator modules in the same manner as for the

non-UltraSCSI BA356, as shown in Figure 10–8. With proper lighting you will be able to see a J or T near the hole where the pin sticks through.

Termination for both ends of the UltraSCSI BA356 internal, single-ended bus is on the personality module, and is always active. Termination for the differential UltraSCSI bus is also on the personality module, and is controlled by the SCSI bus termination switches, switch pack S4. DS-BA35X-DA termination is discussed in Section 10.1.2.2.

10.4 Preparing the Storage for Configurations Using External Termination

A TruCluster Server cluster provides you with high data availability through the cluster file system (CFS), the device request dispatcher, service failover through the cluster application availability (CAA) subsystem, disk mirroring, and fast file system recovery. TruCluster Server supports mirroring of the member-specific boot disks and the cluster quorum disk through hardware RAID only. You can mirror the clusterwide root (/), /usr and /var file systems, the data disks, and the swap file system using the Logical Storage Manager (LSM) technology. You must determine the storage configuration that will meet your needs. Mirroring disks across two shared buses provides the most highly available data.

Disk devices used on the shared bus must be located in a supported storage shelf. Before you connect a storage shelf to a shared SCSI bus, you must install the disks in the unit. Before connecting a RAID array controller to a shared SCSI bus, install the disks and configure the storagesets. For detailed information about installation and configuration, see your storage shelf (or RAID array controller) documentation.

The following sections describe how to prepare storage for a shared SCSI bus and external termination for:

- A BA350, a BA356, and an UltraSCSI BA356
- Two BA356s
- Two UltraSCSI BA356s
- An HSZ20, HSZ40, or HSZ50 RAID array controller
- A RAID Array 3000 with HSZ22 controller

If you need to use a BA350 or non-UltraSCSI BA356 with an UltraSCSI BA356 storage shelf, extrapolate the needed information from Section 10.4.1 and Section 10.4.3.

Later sections describe how to install cables to configure an HSZ20, HSZ40, or HSZ50 in a TruCluster Server configuration with two member systems.

10.4.1 Preparing BA350, BA356, and UltraSCSI BA356 Storage Shelves for an Externally Terminated TruCluster Server Configuration

You may be using the BA350, BA356, or UltraSCSI BA356 storage shelves in your TruCluster Server configuration as follows:

- A BA350 storage shelf provides access to SCSI devices through an 8-bit, single-ended, and narrow SCSI-2 interface. It can be used with a DWZZA-VA and connected to a differential shared SCSI bus.
- A non-Ultra BA356 storage shelf provides access to SCSI devices through a 16-bit, single-ended, and wide SCSI-2 interface. In a cluster configuration, you connect a non-Ultra BA356 to the shared SCSI bus using DWZZB-VW.
- An UltraSCSI BA356 storage shelf provides access to UltraSCSI devices through a 16-bit, single-ended, wide UltraSCSI interface. In a cluster configuration, you connect an UltraSCSI BA356 to the shared SCSI bus through the DS-BA35X-DA personality module.

The following sections discuss the steps necessary to prepare the individual storage shelves, and then connect two storage shelves together to provide the additional storage.

Note

This material has been written with the premise that there are only two member systems in any TruCluster Server configuration using direct connect disks for storage. Using this assumption, and further assuming that the member systems use SCSI IDs 6 and 7, the storage shelf housing disks in the range of SCSI IDs 0 through 6 can only use SCSI IDs 0 through 5.

If there are more than two member systems, additional disk slots will be needed to provide the additional member system SCSI IDs.

10.4.1.1 Preparing a BA350 Storage Shelf for Shared SCSI Usage

To prepare a BA350 storage shelf for usage on a shared SCSI bus, follow these steps:

1. Ensure that the BA350 storage shelf's internal termination and jumper is installed (see Section 10.3.1 and Figure 10–6).
2. You will need a DWZZA-VA signal converter for the BA350. Ensure that the DWZZA-VA single-ended termination jumper, J2, is installed.

Remove the termination from the differential end by removing the five 14-pin differential terminator resistor SIPs.

3. Attach an H885-AA trilink connector to the DWZZA-VA 68-pin high-density connector.
4. Install the DWZZA-VA in slot 0 of the BA350.

10.4.1.2 Preparing a BA356 Storage Shelf for Shared SCSI Usage

To prepare a BA356 storage shelf for shared SCSI bus usage, follow these steps:

1. You need either a DWZZB-AA or DWZZB-VW signal converter. The DWZZB-VW is more commonly used. Verify signal converter termination as follows:
 - Ensure that the DWZZB W1 and W2 jumpers are installed to enable the single-ended termination at one end of the bus. The other end of the BA356 single-ended SCSI bus is terminated on the personality module.
 - Remove the termination from the differential side of the DWZZB by removing the five 14-pin differential terminator resistor SIPs. The differential SCSI bus will be terminated external to the DWZZB.
2. Attach an H885-AA trilink connector to the DWZZB 68-pin high-density connector.
3. Set the switches on the BA356 personality module as follows:
 - If the BA356 is to house disks with SCSI IDs in the range of 0 through 6, set the BA356 personality module address switches 1 through 7 to off.
 - If the BA356 is to house disks with SCSI IDs in the range of 8 through 14, set BA356 personality module address switches 1 through 3 to on and switches 4 through 7 to off.

If you are using a DWZZB-AA, do not replace the personality module until you attach the cable in the next step.

4. If you are using a DWZZB-AA signal converter, use a BN21K-01 (1-meter; 3.3-foot) or BN21L-01 (1-meter; 3.3-foot) cable to connect the single-ended side of the DWZZB-AA to the BA356 input connector, JA1, on the personality module. Connector JA1 is on the left side of the personality module as you face the front of the BA356, and is hidden from normal view. This connection forms a single-ended bus segment that is terminated by the DWZZB single-ended termination and the BA356 termination on the personality module. The use of a 1-meter (3.3-foot) cable keeps the single-ended SCSI bus (cable and BA356) under the 3-meter (9.8-foot) limit to still allow high speed operation.

If you are using a DWZZB-VW, install it in slot 0 of the BA356.

10.4.1.3 Preparing an UltraSCSI BA356 Storage Shelf for a TruCluster Configuration

An UltraSCSI BA356 storage shelf is connected to a shared UltraSCSI bus, and provides access to UltraSCSI devices on the internal, single-ended and wide UltraSCSI bus. The interface between the buses is the DS-BA35X-DA personality module installed in the UltraSCSI BA356.

To prepare an UltraSCSI BA356 storage shelf for usage on a shared SCSI bus, follow these steps:

1. Ensure that the BA35X-MJ jumper module is installed behind slot 6 (see Section 10.3.2.1, Figure 10–7, and Figure 10–8).
2. Set the SCSI bus ID switches on the UltraSCSI BA356 personality module (DS-BA35X-DA, Figure 10–3) as follows:
 - If the UltraSCSI BA356 is to house disks with SCSI IDs in the range of 0 through 6, set the personality module address switches S3-1 through S3-7 to OFF.
 - If the UltraSCSI BA356 is to house disks with SCSI IDs in the range of 8 through 14, set personality module address switches S3-1 through S3-3 to ON and switches S3-4 through S3-7 to OFF.
3. Disable the UltraSCSI BA356 differential termination. Ensure that personality module (DS-BA35X-DA) switch pack 4 switches S4-1 and S4-2 are ON (see Figure 10–3).

Note

S4-3 and S4-4 are not used on the DS-BA35X-DA.

4. Install an H8861-AA VHDCI trilinear connector on the UltraSCSI BA356 personality module.

10.4.2 Cabling a Single Storage Shelf

A cluster with a single storage shelf is generally of little use because of the lack of disk space. But, for those customers whose primary use of the cluster is the Memory Channel application programming interface (API), storage is not an issue; they only need a shared SCSI bus to fulfill the TruCluster Server requirements. Therefore, there are no illustrations in this section showing clusters with single storage shelves. The illustrations in Section 10.4.3 can be used with the descriptions in this section.

10.4.2.1 Cabling a Single BA350 Storage Shelf

To cable a single BA350 storage shelf into a cluster, install a BN21K, BN21L, or 328215-00X HD68 cable between the BN21W-0B Y cable on the host bus adapter of each system and the H885-AA tralink connector installed on the DWZZA-VA installed in slot 0 of the BA350. See the left-half of Figure 10–9.

10.4.2.2 Cabling a Single BA356 Storage Shelf

To cable a single BA356 storage shelf into a cluster, install a BN21K, BN21L, or 328215-00X HD68 cable between the BN21W-0B Y cable on the host bus adapter of each system and the H885-AA tralink connector installed on the DWZZB-VW installed in slot 0 of the BA356. See Figure 10–10.

10.4.2.3 Cabling a Single UltraSCSI BA356 Storage Shelf

To cable a single UltraSCSI BA356 storage shelf into a cluster, connect a BN38C, BN38D, or combination of a BN38E-0B technology adapter cable and a BN37A cable between the BN21W-0B Y cable on each system and H8861-AA VHDCI tralink connector on the UltraSCSI BA356 personality module. See Figure 10–11.

10.4.3 Connecting Storage Shelves Together

Section 10.4.1 covered the steps necessary to prepare the BA350, BA356, and UltraSCSI BA356 storage shelves for use on a shared SCSI bus. However, you will probably need more storage than one storage shelf can provide, so you will need two storage shelves on the shared SCSI bus.

Note

Because the BA350 contains a narrow (8-bit), single-ended SCSI bus, it only supports SCSI IDs 0 through 7. Therefore, a BA350 must be used with a BA356 or UltraSCSI BA356 if more than five disks are required.

The following sections provide the steps needed to connect two storage shelves and two member systems on a shared SCSI bus:

- BA350 and BA356 (Section 10.4.3.1)
- Two BA356s (Section 10.4.3.2)
- Two UltraSCSI BA356s (Section 10.4.3.3)

10.4.3.1 Connecting a BA350 and a BA356 for Shared SCSI Bus Usage


When you use a BA350 and a BA356 for storage on a shared SCSI bus in a TruCluster Server configuration, the BA356 must be configured for SCSI IDs 8 through 14.

To prepare a BA350 and BA356 for shared SCSI bus usage (see Figure 10–9), follow these steps:

1. Complete the steps in Section 10.4.1.1 and Section 10.4.1.2 to prepare the BA350 and BA356. Ensure that the BA356 is configured for SCSI IDs 8 through 14.
2. If either storage shelf will be at the end of the shared SCSI bus, attach an H879-AA terminator to the H885-AA tralink on the DWZZA or DWZZB for the storage shelf that will be at the end of the bus. You can choose either storage shelf to be on the end of the bus.
3. Connect a BN21K or BN21L between the H885-AA tralink on the DWZZA (BA350) and the H885-AA tralink on the DWZZB (BA356)
4. When the KZPSA-BB or KZPBA-CB host bus adapters have been installed:
 - If the storage shelves are on the end of the shared SCSI bus, connect a BN21K (or BN21L) cable between the BN21W-0B Y cables on the host bus adapters. Connect another BN21K (or BN21L) cable between the BN21W-0B Y cable with an open connector and the H885-AA tralink (on the storage shelf) with an open connector.
 - If the storage shelves are in the middle of the shared SCSI bus, connect a BN21K (or BN21L) cable between the BN21W-0B Y cable on each host bus adapter and the H885-AA tralink on a storage shelf.

Figure 10–9 shows a two-member TruCluster Server configuration using a BA350 and a BA356 for storage.

Figure 10–9: BA350 and BA356 Cabled for Shared SCSI Bus Usage



ZK-1595U-AI

Table 10–1 lists the components that are used to create the cluster that is shown in Figure 10–9 and Figure 10–10.

Table 10–1: Hardware Components Used for Configuration Shown in Figure 10–9 and Figure 10–10

Callout Number	Description
1	BN21W-0B Y cable
2	H879-AA terminator
3	BN21K, BN21L, or 328215-00X cable ^a
4	H885-AA trilink connector

^a The maximum combined length of the BN21K, BN21L, or 328215-00X cables must not exceed 25 meters (82 feet).

10.4.3.2 Connecting Two BA356s for Shared SCSI Bus Usage

When you use two BA356 storage shelves on a shared SCSI bus in a TruCluster configuration, one BA356 must be configured for SCSI IDs 0 through 6 and the other configured for SCSI IDs 8 through 14.

To prepare two BA356 storage shelves for shared SCSI bus usage (see Figure 10–10), follow these steps:

1. Complete the steps of Section 10.4.1.2 for each BA356. Ensure that the personality module address switches on one BA356 are set to select SCSI IDs 0 through 6, and that the address switches on the other BA356 personality module are set to select SCSI IDs 8 through 14.
2. If either of the BA356 storage shelves will be on the end of the SCSI bus, attach an H879-AA terminator to the H885-AA trilink on the DWZZB for the BA356 that will be on the end of the bus.
3. Connect a BN21K or BN21L cable between the H885-AA trilinks.
4. When the KZPSA-BB or KZPBA-CB host bus adapters have been installed:
 - If the BA356 storage shelves are on the end of the shared SCSI bus, connect a BN21K (or BN21L) cable between the BN21W-0B Y cables on the host bus adapters. Connect another BN21K (or BN21L) cable between the BN21W-0B Y cable with an open connector and the H885-AA trilink (on the BA356) with an open connector.
 - If the BA356s are in the middle of the shared SCSI bus, connect a BN21K (or BN21L) cable between the BN21W-0B Y cable on each host bus adapter and the H885-AA trilink on a BA356 storage shelf.

Figure 10–10 shows a two member TruCluster Server configuration using two BA356s for storage.

Figure 10–10: Two BA356s Cabled for Shared SCSI Bus Usage

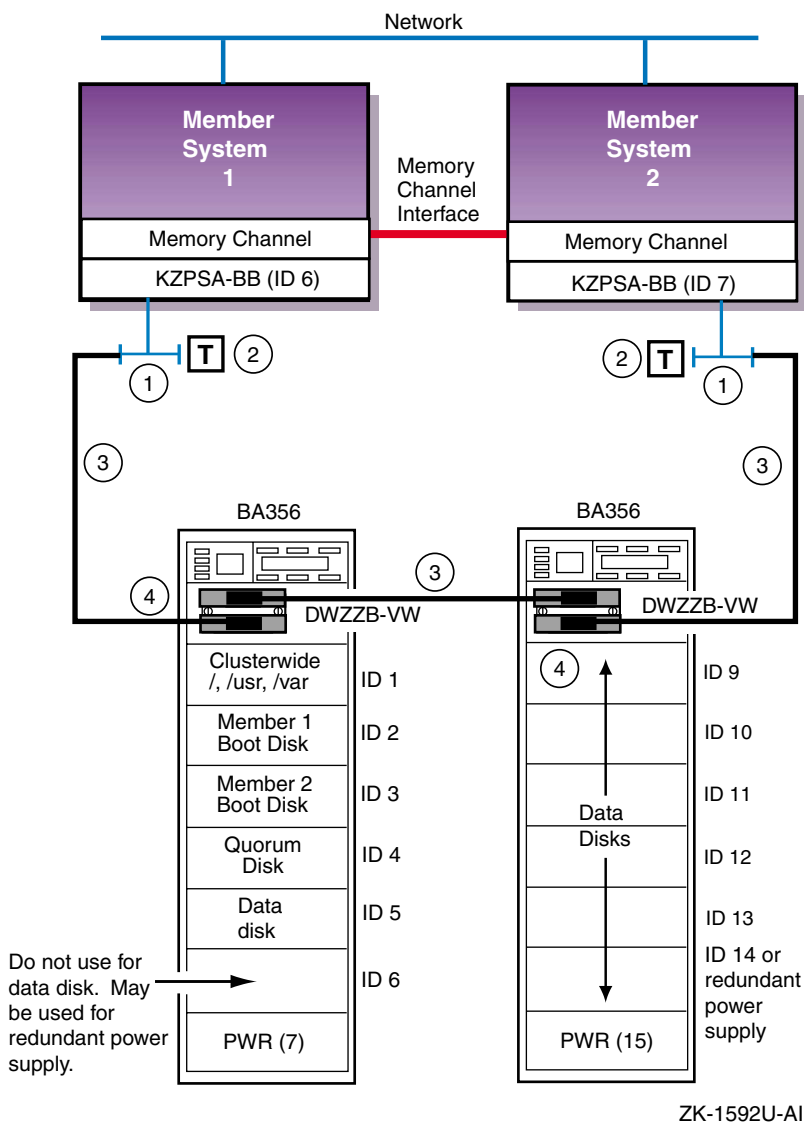


Table 10–1 lists the components that are used to create the cluster that is shown in Figure 10–10.

10.4.3.3 Connecting Two UltraSCSI BA356s for Shared SCSI Bus Usage

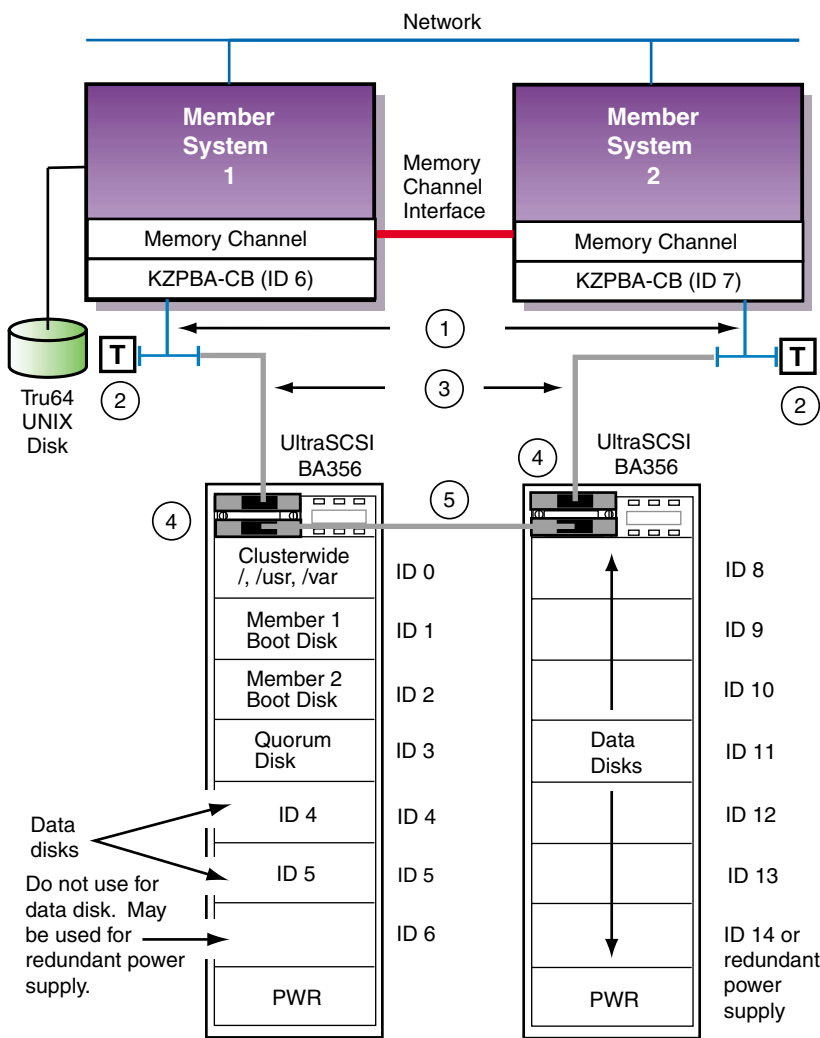
When you use two UltraSCSI BA356 storage shelves on a shared SCSI bus in a TruCluster configuration, one storage shelf must be configured for SCSI IDs 0 through 6 and the other configured for SCSI IDs 8 through 14.

To prepare two UltraSCSI BA356 storage shelves for shared SCSI bus usage (see Figure 10–11), follow these steps:

1. Complete the steps of Section 10.4.1.3 for each UltraSCSI BA356. Ensure that the personality module address switches on one UltraSCSI BA356 are set to select SCSI IDs 0 through 6 and the address switches on the other UltraSCSI BA356 personality module are set to select SCSI IDs 8 through 14.
2. You will need two H8861-AA VHDCI tralink connectors. If either of the UltraSCSI BA356 storage shelves will be on the end of the SCSI bus, attach an H8863-AA terminator to one of the H8861-AA tralink connectors. Install the tralink with the terminator on connector JA1 of the DS-BA35X-DA personality module of the UltraSCSI BA356 that will be on the end of the SCSI bus. Install the other H8861-AA tralink on JA1 of the DS-BA35X-DA personality module of the other UltraSCSI BA356.
3. Connect a BN37A VHDCI to VHDCI cable between the H8861-AA tralink connectors on the UltraSCSI BA356s.
4. When the KZPSA-BBs or KZPBA-CBs are installed:
 - If one of the UltraSCSI BA356s is on the end of the SCSI bus, install a BN38C (or BN38D) HD68 to VHDCI cable between one of the BN21W-0B Y cables (on the host bus adapters) and the open connector on the H8861-AA tralink connector on the DS-BA35X-DA personality module. Connect the BN21W-0B Y cables on the two member system host adapters together with a BN21K (or BN21L) cable.
 - If the UltraSCSI BA356s are in the middle of the SCSI bus, install a BN38C (or BN38D) HD68 to VHDCI cable between the BN21W-0B Y cable on each host bus adapter and the open connector on the H8861-AA tralink connector on the DS-BA35X-DA personality modules.

Figure 10–11 shows a two member TruCluster Server configuration using two UltraSCSI BA356s for storage.

Figure 10–11: Two UltraSCSI BA356s Cabled for Shared SCSI Bus Usage



ZK-1598U-AI

Table 10–2 lists the components that are used to create the cluster that is shown in Figure 10–11.

Table 10–2: Hardware Components Used for Configuration Shown in Figure 10–11

Callout Number	Description
1	BN21W-0B Y cable
2	H879-AA HD68 terminator
3	BN38C (or BN38D) cable ^{a b}
4	H8861-AA VHDCI tralink connector
5	BN37A cable ^b

^a A BN38E-0B technology adapter cable may be connected to a BN37A cable and used in place of a BN38C or BN38D cable.

^b The maximum combined length of the BN38C (or BN38D) and BN37A cables on one SCSI bus segment must not exceed 25 meters (82 feet).

10.4.4 Cabling a Non-UltraSCSI RAID Array Controller to an Externally Terminated Shared SCSI Bus

A RAID array controller provides high performance, high availability, and high connectivity access to SCSI devices through the shared SCSI buses.

Before you connect a RAID controller to a shared SCSI bus, you must install and configure the disks that the controller will use, and ensure that the controller has a unique SCSI ID on the shared bus.

You can configure the HSZ20, HSZ40, and HSZ50 RAID array controllers with one to four SCSI IDs.

Because the HSZ20, HSZ40, and HSZ50 have a wide differential connection on the host side, you connect them to one of the following differential devices:

- KZPSA-BB host bus adapter
- KZPBA-CB host bus adapter
- Another HSZ20, HSZ40, or HSZ50

Note

The HSZ20, HSZ40, and HSZ50 cannot operate at UltraSCSI speeds when used with the KZPBA-CB.

You can also use a DS-DWZZH-03 or DS-DWZZH-05 UltraSCSI hub with one of these RAID array controllers and either the KZPSA-BB or KZPBA-CB host bus adapters. UltraSCSI cables are required to make the connection to the hub. UltraSCSI speed is not supported with these RAID array controllers when used with a hub and the KZPBA-CB host bus adapter.

10.4.4.1 Cabling an HSZ40 or HSZ50 in a Cluster Using External Termination

To connect an HSZ40 or HSZ50 controller to an externally terminated shared SCSI bus, follow these steps:

1. If the HSZ40 or HSZ50 will be on the end of the shared SCSI bus, attach an H879-AA terminator to an H885-AA tralink connector.
2. Attach an H885-AA tralink connector to each RAID controller port. Attach the H885-AA tralink connector with the terminator to the controller that will be on the end of the shared SCSI bus.
3. If you are using dual-redundant RAID array controllers, install a BN21-K or BN21L cable (a BN21L-0B is a 15-centimeter (5.9-inch) cable) between the H885-AA tralink connectors on the controllers.
4. When the host bus adapters (KZPSA-BB or KZPBA-CB) have been installed, connect the host bus adapters and RAID array controllers together with BN21K or BN21L cables as follows:
 - Both member systems are on the ends of the bus: Attach a BN21K or BN21L cable from the BN21W-0B Y cable on each host bus adapter to the RAID array controllers.
 - RAID array controller is on the end of the bus: Connect a BN21K (or BN21L) cable from the BN21W-0B Y cable on one host bus adapter to the BN21W-0B Y cable on the other host bus adapter. Attach another BN21K (or BN21L) cable from the open BN21W-0B Y cable connector to the open H885-AA connector on the RAID array controller.

Figure 10–12 shows two AlphaServer systems in a TruCluster Server configuration with dual-redundant HSZ50 RAID controllers in the middle of the shared SCSI bus. Note that the SCSI bus adapters are KZPSA-BB PCI-to-SCSI adapters. They could be KZPBA-CB host bus adapters without changing any cables.

Figure 10–12: Externally Terminated Shared SCSI Bus with Mid-Bus HSZ50 RAID Array Controllers

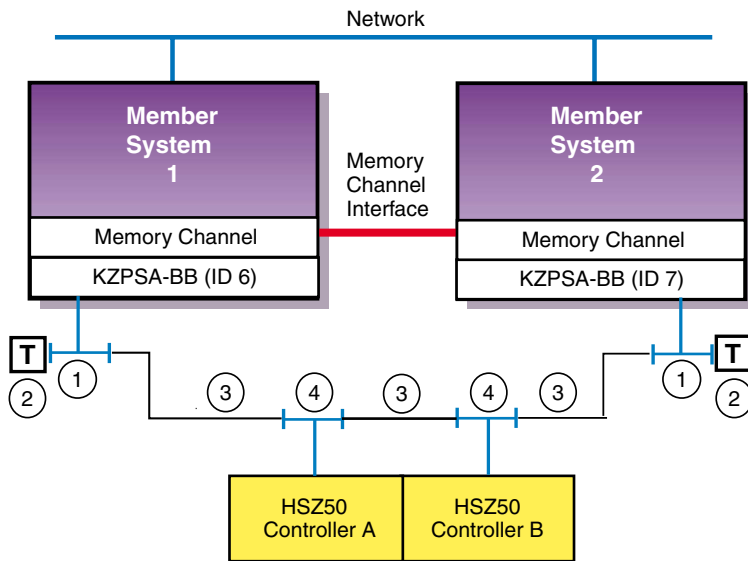
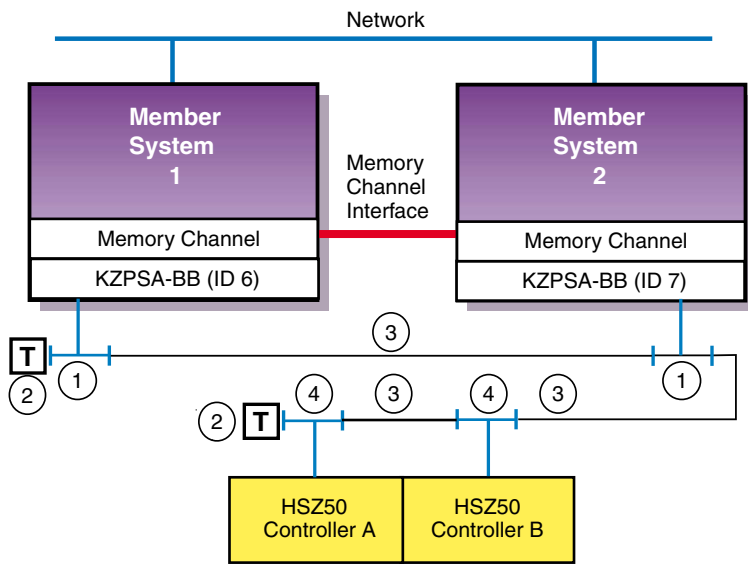


Table 10–3 lists the components that are used to create the cluster that is shown in Figure 10–12 and Figure 10–13.

Figure 10–13 shows two AlphaServer systems in a TruCluster Server configuration with dual-redundant HSZ50 RAID controllers at the end of the shared SCSI bus. As with Figure 10–12, the SCSI bus adapters are KZPSA-BB PCI-to-SCSI adapters. They could be KZPBA-CB host bus adapters without changing any cables.

Figure 10–13: Externally Terminated Shared SCSI Bus with HSZ50 RAID Array Controllers at Bus End



ZK-1597U-AI

Table 10–3 lists the components that are used to create the cluster that is shown in Figure 10–12 and Figure 10–13.

Table 10–3: Hardware Components Used for Configuration Shown in Figure 10–12 and Figure 10–13

Callout Number	Description
1	BN21W-0B Y cable
2	H879-AA terminator
3	BN21K (or BN21L) cable ^{a b}
4	H885-AA trilink connector

^a The maximum combined length of the BN21K (or BN21L) cables must not exceed 25 meters (82 feet).

^b The cable between the H885-AA trilink connectors on the HSZ50s could be a BN21L-0B, a 15-centimeter (5.9-inch) cable.

10.4.4.2 Cabling an HSZ20 in a Cluster Using External Termination

To connect a SWXRA-Z1 (HSZ20 controller) to a shared SCSI bus, follow these steps:

1. Referring to the *RAID Array 310 Deskside Subsystem (SWXRA-ZX) Hardware User's Guide*, open the SWXRA-Z1 cabinet, locate the SCSI bus converter board, and:
 - Remove the five differential terminator resistor SIPs.
 - Ensure that the W1 and W2 jumpers are installed to enable the single-ended termination on one end of the bus.

Note

The RAID Array 310 SCSI bus converter board is the same logic board used in the DWZZB signal converter.

2. Attach an H885-AA tralink connector to the SCSI input connector (on the back of the cabinet).
3. Use a BN21K or BN21L cable to connect the tralink connector to a tralink connector or BN21W-0B Y cable attached to a differential SCSI controller, another storage shelf, or the differential end of a signal converter.
4. Terminate the differential bus by attaching an H879-AA terminator to the H885-AA tralink connector or BN21W-0B Y cable at each end of the shared SCSI bus.

Ensure that all devices that make up the shared SCSI bus are connected, and that there is a terminator at each end of the shared SCSI bus.

10.4.5 Cabling an RA3000 Using External Termination

The RAID Array 3000 (RA3000) is a low-end, standalone UltraSCSI RAID subsystem. It supports RAID levels 0, 1, 0+1, 4, 5, and JBOD disks.

The RA3000 storage subsystem has fully redundant components to eliminate single points of failure. It comes with a standard uninterruptible power supply (UPS) for cache data protection during power outages.

The RA3000 uses the dual-ported HSZ22 controller. Optional dual redundant controllers with mirrored write-back cache provide maximum data integrity.

For more information on the RA3000, see Section 3.7.1.3.

Note

The RA3000 is supported on a shared SCSI bus only with the KZPBA-CB UltraSCSI host bus adapter.

Table 10–4 provides the steps necessary to connect TruCluster Server member systems to an RA3000 storage subsystem using external termination and Y cables.

Table 10–4: Installing Cables for RA3000 Configuration Using External Termination and Y Cables

Action	Refer to:
Install SCSI bus cables:	
RA3000 pedestal with active/passive failover: Install a BN38C HD68 to VHDCI cable between the BN21W-0B Y cable of one member system and the RA3000 Host 0 port. ^a	Figure 10–14
Install a BN21K, BN21L, or BN31G cable between the BN21W-0B Y cables of all other member systems. ^b	—
RA3000 controller shelf with active/passive failover: Install a BN38C HD68 to VHDCI cable between the BN21W-0B Y cable of one member system and the RA3000 Host 0 I/O module Host In port. ^a	Figure 10–15
Install a BN21K, BN21L, or BN31G cable between the BN21W-0B Y cables of all other member systems. ^b	—
RA3000 controller shelf with active/active or active/passive failover: Install a BN38C HD68 to VHDCI cable between the BN21W-0B Y cable of one member system and the RA3000 Host 0 I/O module Host In connection. ^a	Figure 10–16
Install a BN37A-0E 50-centimeter (19.7-inch) VHDCI cable between the RA3000 controller shelf Host 0 I/O module Host Out port and the Host 1 I/O module Host In port.	—
Install a BN21K, BN21L, or BN31G cable between the BN21W-0B Y cables of all other member systems. ^b	—
RA3000 mid-bus controller shelf with active/active or active/passive failover: Install a BN38C HD68 to VHDCI cable between the BN21W-0B Y cable of one member system and the RA3000 Host 0 I/O module Host In connection. Install a second BN38C HD68 to VHDCI cable between the BN21W-0B Y cable of another member system and the RA3000 Host 1 I/O module Host Out connection. This disables the termination on the Host 1 I/O module. ^a	Figure 10–17

Table 10–4: Installing Cables for RA3000 Configuration Using External Termination and Y Cables (cont.)

Action	Refer to:
Install a BN37A-0E 50-centimeter (19.7-inch) VHDCI cable between the RA3000 controller shelf Host 0 I/O module Host Out port and the Host 1 I/O module Host In port. The connection to Host 0 I/O module Host Out port disables the termination on that Host I/O module.	—
Install a BN21K, BN21L, or BN31G cable between the BN21W-0B Y cables of any other member systems. ^b	—

Note

You cannot create a mid-bus configuration using a RA3000 pedestal. The member systems on one SCSI bus segment connected to the Host 0 port would see some devices. The member systems on the other SCSI bus segment connected to the Host 1 port would not be able to see the same devices.

^a A BN38E-0B technology adapter cable may be connected to a BN37A cable and used in place of a BN38C or BN38D cable.

^b The maximum length of the SCSI bus segment, including the combined length of the BN38C (or BN38E-0B/BN37A combination), BN37A, and BN21K (or BN21L or BN31G) cables and internal device length, must not exceed 25 meters (82 feet).

Figure 10–14 (pedestal) and Figure 10–15 (controller shelf) show an externally terminated TruCluster Server configuration using an RA3000. The RA3000 controller shelf and pedestal contains internal termination. Table 10–5 lists the components used to create the clusters shown in Figure 10–14, Figure 10–15, and Figure 10–16.

Figure 10–14: Externally Terminated TruCluster Server Configuration with an RA3000 Pedestal with Active/Passive Failover

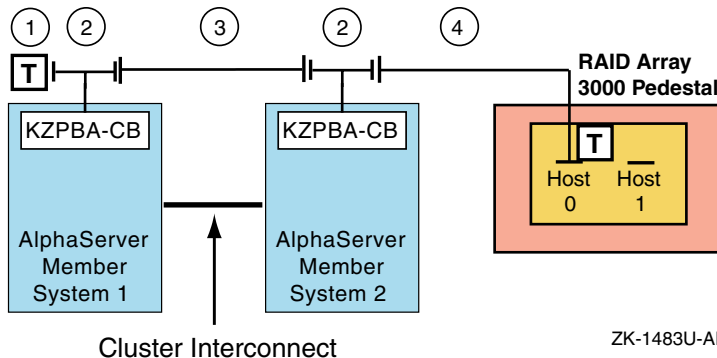
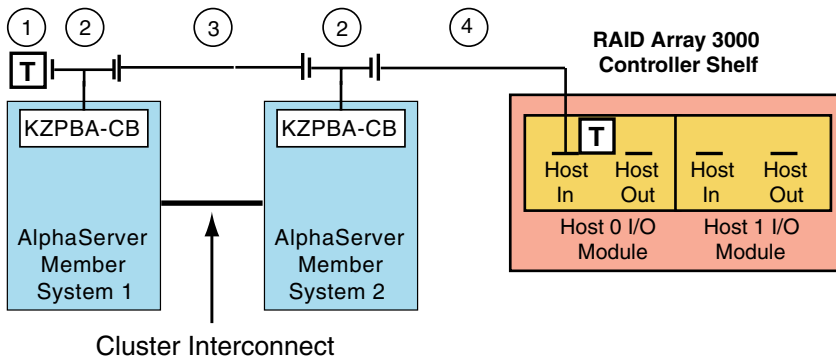


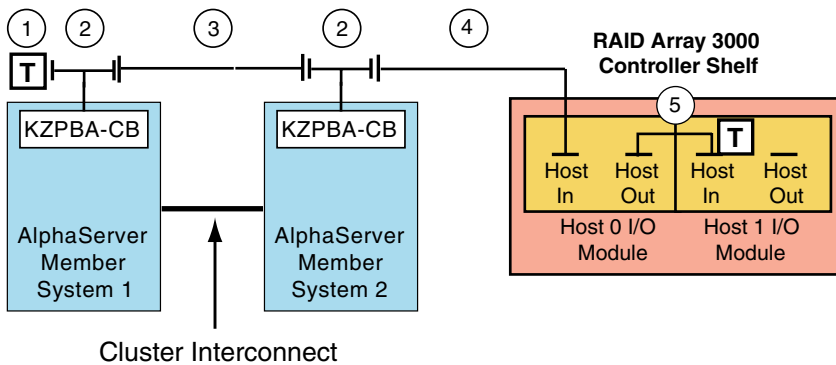
Figure 10–15: Externally Terminated TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Passive Failover



ZK-1481U-AI

Figure 10–16 shows an externally terminated TruCluster Server configuration using an RA3000 controller shelf. In this configuration, because the Host 0 I/O module is daisy-chained to Host 1 I/O module, dual HSZ22 controllers could use active/active or active/passive failover.

Figure 10–16: Externally Terminated TruCluster Server Configuration with an RA3000 Controller Shelf with Active/Active or Active/Passive Failover



ZK-1482U-AI

Table 10–5 lists the components used to create the clusters shown in Figure 10–14, Figure 10–15, and Figure 10–16.

Table 10–5: Hardware Components Used in the TruCluster Server Configuration Shown in Figure 10–14, Figure 10–15, and Figure 10–16

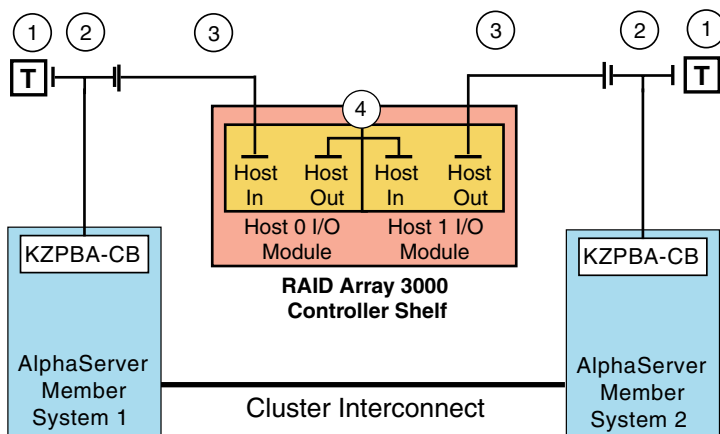
Callout Number	Description
1	H879-AA terminator
2	BN21W-0B Y cable
3	BN21K (BN21L or BN31G) HD68 cable ^a
4	BN38C HD68 to VHDCI cable ^a
5	BN37A-0E 50-centimeter (19.7-inch) VHDCI cable ^b

^a The maximum length of the SCSI bus segment, including the combined length of BN38C (or BN38E-0B/BN37A combination), BN21K (BN21L or BN31G), and BN37A cables and internal device length, must not exceed 25 meters (82 feet).

^b The BN37A-0E cable is used only in the TruCluster Server configuration shown in Figure 10–16. It is not used in the configurations shown in Figure 10–14 or Figure 10–15.

Figure 10–17 shows an externally terminated TruCluster Server configuration with a RA3000 controller shelf in the middle of the shared SCSI bus. In this configuration, because the Host 0 I/O module is daisy-chained to the Host 1 I/O module, dual HSZ22 controllers could use active/active or active/passive failover.

Figure 10–17: Externally Terminated TruCluster Server Configuration with a Mid-bus RA3000 Controller Shelf with Active/Active or Active/Passive Failover



ZK-1484U-AI

Table 10–6 lists the components used to create the cluster shown in Figure 10–17.

Table 10–6: Hardware Components Used in the Configuration Shown in Figure 10–17

Callout Number	Description
1	H879-AA terminator
2	BN21W-0B Y cable
3	BN38C HD68 to VHDCI cable ^a
4	BN37A-0E 50-centimeter (19.7-inch) VHDCI cable ^a

^a The maximum length of the SCSI bus segment, including the combined length of BN38C (or BN38E-0B/BN37A combination), BN37A-0E, and any BN21K (BN21L or BN31G) cables that might be connecting adjacent systems (which are not shown in the figure) and internal device length, must not exceed 25 meters (82 feet).

10.5 Cabling an HSZ40 or HSZ50 RAID Array Controller in a Radial Configuration with an UltraSCSI Hub

You may have an HSZ40 or HSZ50 that you want to keep when you upgrade to a newer AlphaServer system. You can connect an HSZ40 or HSZ50 to an UltraSCSI hub in a radial configuration, but even if the host bus adapter is a KZPBA-CB, it will not operate at UltraSCSI speed with the HSZ40 or HSZ50.

To configure a dual-redundant HSZ40 or HSZ50 RAID array controller and an UltraSCSI hub in a radial configuration, follow these steps:

1. You will need two H885-AA tralink connectors. Install an H879-AA terminator on one of the trilinks.
2. Attach the tralink with the terminator to the controller that you want to be on the end of the shared SCSI bus. Attach an H885-AA tralink connector to the other controller.
3. Install a BN21K or BN21L cable between the H885-AA tralink connectors on the two controllers. The BN21L-0B is a 15-centimeter (5.9-inch) cable.
4. If you are using a DS-DWZZH-05:
 - Verify that the fair arbitration switch is in the **Fair** position to enable fair arbitration (see Section 3.6.1.2.2)
 - Ensure that the W1 jumper is removed to select wide addressing mode (see Section 3.6.1.2.3)
5. Install the UltraSCSI hub in:
 - A StorageWorks UltraSCSI BA356 shelf (which has the required 180-watt power supply).
 - A non-UltraSCSI BA356 that has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.

6. If you are using a:
 - DS-DWZZH-03: Install a BN38C (or BN38D) HD to VHDCI cable between any DS-DWZZH-03 port and the open connector on the H885-AA trilink connector (on the RAID array controller).
 - DS-DWZZH-05: Install a BN38C (or BN38D) cable between the DWZZH-05 controller port and the open trilink connector on HSZ40 or HSZ50 controller.

Note

Ensure that the HSZ40 or HSZ50 SCSI IDs match the DS-DWZZH-05 controller port IDs (SCSI IDs 0-6).

7. When the host bus adapters (KZPSA-BB or KZPBA-CB) have been installed in the member systems, for a:
 - DS-DWZZH-03: Install a BN38C (or BN38D) HD68 to VHDCI cable between the KZPBA-CB or KZPSA-BB host bus adapter to each of the other two DS-DWZZH-03 ports.
 - DS-DWZZH-05: Install a BN38C (or BN38D) HD68 to VHDCI cable between the KZPBA-CB or KZPSA-BB host bus adapter on each system to a port on the DWZZH hub. Ensure that the host bus adapter SCSI ID matches the SCSI ID assigned to the DWZZH-05 port it is cabled to (12, 13, 14, and 15).

Figure 10–18 shows a sample configuration with radial connection of KZPSA-BB PCI-to-SCSI adapters, DS-DWZZH-03 UltraSCSI hub, and an HSZ50 RAID array controller. Note that the KZPSA-BBs could be replaced with KZPBA-CB UltraSCSI adapters without any changes in cables.

Figure 10–18: TruCluster Server Cluster Using DS-DWZZH-03, SCSI Adapter with Terminators Installed, and HSZ50

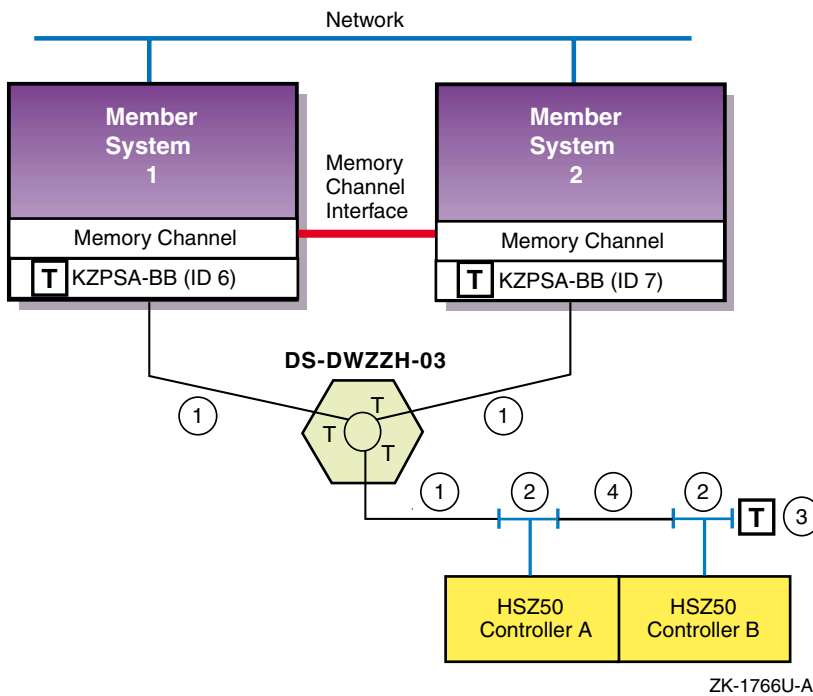


Table 10–7 lists the components that are used to create the cluster that is shown in Figure 10–18 and Figure 10–19.

Table 10–7: Hardware Components Used in Configuration Shown in Figure 10–18 and Figure 10–19

Callout Number	Description
1	BN38C cable ^a ^b
2	H885-AA HD68 trilink connector
3	H879-AA HD68 terminator
4	BN21K, BN21L, or 328215-00X cable ^b

^a The maximum length of the BN38C cable on one SCSI bus segment must not exceed 25 meters (82 feet).

^b The maximum combined length of the BN38C and BN21K, BN21L, or 321215-00X cables on the storage SCSI bus segment must not exceed 25 meters (82 feet).

Figure 10–19 shows a sample configuration that uses KZPSA-BB SCSI adapters, a DS-DWZZH-05 UltraSCSI hub, and an HSZ50 RAID array controller.

11

Configuring an Eight-Member Cluster Using Externally Terminated Shared SCSI Buses

This chapter discusses the following topics:

- Overview of an eight-node cluster (Section 11.1)
- How to configure an eight-node cluster using an UltraSCSI BA356 and external termination (Section 11.2)

TruCluster Server Version 5.1A supports eight-member cluster configurations as follows:

- **Fibre Channel:** Eight-member systems may be connected to common storage over Fibre Channel in a fabric (switch) configuration.
- **Parallel SCSI:** Only four of the member systems may be connected to any one SCSI bus, but you can have multiple SCSI buses connected to different sets of nodes, and the sets of nodes may overlap. We recommend you use a DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled when connecting four-member systems to a common SCSI bus using RAID array controllers.

Note

The DS-DWZZH-03/05 UltraSCSI hubs cannot be connected to a StorageWorks BA35X storage shelf because the storage shelf does not provide termination power to the hub.

Configuring an eight-member cluster using Fibre Channel is straightforward; connect the member systems to the Memory Channel hub and to the Fibre Channel switches. (See Chapter 6 for more information on configuring Fibre Channel).

Configuring an eight-member cluster using shared SCSI buses is more complex because you can only have four member systems on a single shared SCSI bus.

The primary focus of this chapter is on an eight-node cluster that uses externally terminated shared SCSI buses with minimal storage. This type of cluster is of primary interest to high-performance technical computing (HPTC) cluster customers. It is also of importance to those customers who use Tru64 UNIX Versions 4.0F or 4.0G with the TruCluster Software Products Memory Channel Software Version 1.6 product who want to upgrade to Tru64 UNIX Version 5.1A and TruCluster Server Version 5.1A.

11.1 Overview of an Eight-Node TruCluster Server Cluster

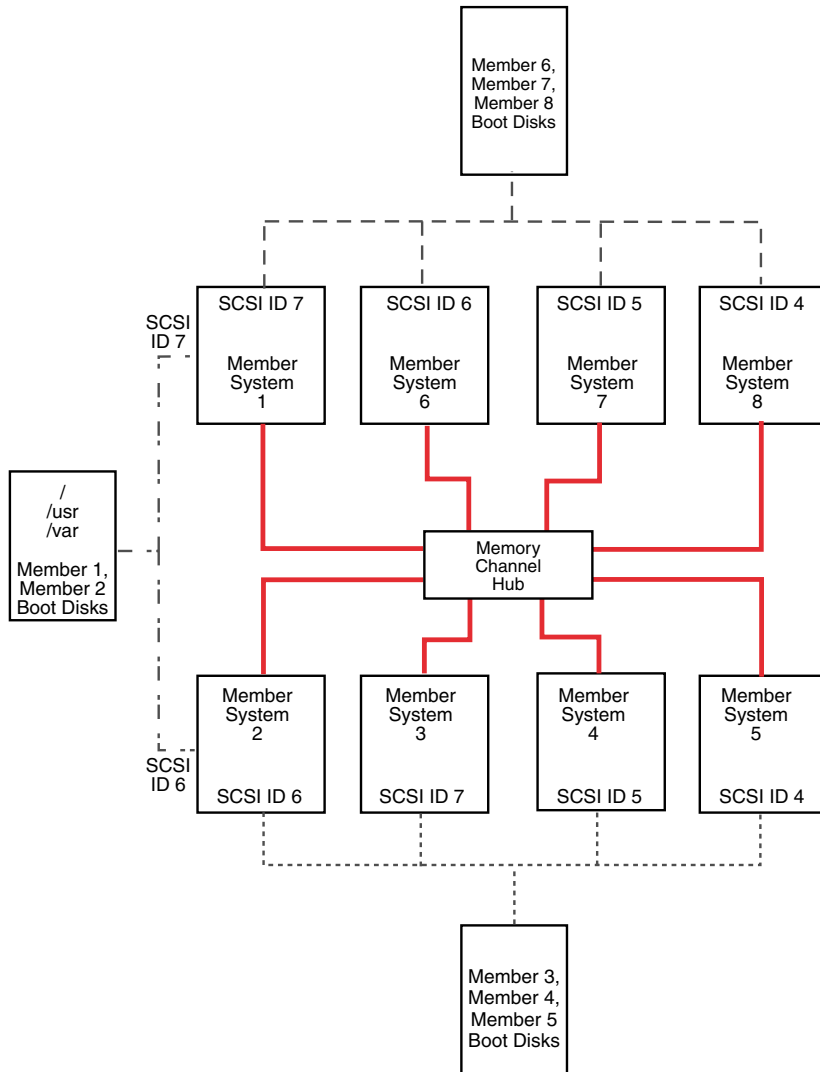
Figure 11–1 shows a basic block diagram of an eight-node cluster.

Note

The public network is not shown in Figure 11–1, or in any other illustration in this chapter. Ensure that you have network adapters for your public network.

This is just one of many ways to configure an eight-node cluster. You must choose a configuration that best fits your applications and needs.

Figure 11–1: Block Diagram of an Eight-Node Cluster



ZK-1847U-AI

Figure 11–1 shows the following:

- All member systems are connected via the Memory Channel at the Memory Channel hub.
- There are three shared SCSI buses with shared storage on each shared bus:
 - Member systems 1 and 2 are on the first shared SCSI bus.

The Tru64 UNIX Version 5.1A operating system is installed on member system 1. It can be installed on an internal disk, as is the case in Figure 11–1, or on a shared disk.

Member system 1 is used to create the cluster with the `clu_create` command.

Member system 2 is added to the cluster with the `clu_add_member` command.

The shared storage for member systems 1 and 2 contains the root (`/`), `/usr`, `/var` file systems for the cluster, and the boot disks for member systems 1 and 2. (See the Tru64 UNIX *Installation Guide* and the TruCluster Server *Cluster Installation* manual for information about installing the Tru64 UNIX and TruCluster Server software.)

Give member systems 1 and 2 one vote.

- Member systems 2, 3, 4, and 5 share storage on the second shared SCSI bus. Four systems is the maximum number of cluster members that may be on a shared SCSI bus.

The shared storage on this bus contains the member system boot disks for member systems 3, 4, and 5.

Use member system 1 or 2 to add member systems 3, 4, and 5 to the cluster.

Give member system 3 one vote.

- Member Systems 1, 6, 7, and 8 form the third shared SCSI bus.

The shared storage on this bus contains the member system boot disks for member systems 6, 7, and 8.

Use member system 1 or 2 to add member systems 6, 7, and 8 to the cluster.

Section 11.2.1 and Figure 11–2 provide the details about cabling member systems 1 and 2 for the first shared SCSI bus.

Section 11.2.2 and Figure 11–3 provide the details about cabling the second shared SCSI bus and member systems 2, 3, 4, and 5 into the cluster.

Section 11.2.3 and Figure 11–4 provide the details about cabling the third shared SCSI bus and member systems 1, 6, 7, and 8 into the cluster.

Note

You can install Tru64 UNIX and TruCluster Server Version 5.1A after you complete member system 1 and 2 hardware installation, or you can wait until the hardware for all the systems is installed.

11.2 Configuring an Eight-Node Cluster Using an UltraSCSI BA356 and External Termination

Configuring an eight-node cluster is carried out in three distinct stages, one stage for each shared SCSI bus:

1. Install member systems 1 and 2 and all associated cluster hardware needed to place these two systems on a shared SCSI bus.
2. Install member systems 3, 4, and 5 and all associated cluster hardware needed to place these two systems on a shared SCSI bus with member system 2.
3. Install member systems 6, 7, and 8 and all associated cluster hardware needed to place these two systems on a shared SCSI bus with member system 1.

Note

You can switch steps 2 and 3 around and install member systems 6, 7, and 8 before member systems 3, 4, and 5.

11.2.1 Cabling the First Two Nodes on the First Externally Terminated Shared SCSI Cluster

This section provides installation instructions for the cluster hardware for the first two nodes of an eight-node shared SCSI bus cluster. Complete the steps in order. When you are referred to another section or table, complete those steps completely before returning to this section.

Notes

If you are upgrading from Tru64 UNIX Version 4.0F or V4.0G and TruCluster Memory Channel Software Version 1.6 to Tru64 UNIX Version 5.1A and TruCluster Server Version 5.1A, you are required to have shared storage, but you do not have to change your Memory Channel or public network hardware. Therefore, you can skip those steps except as follows.

If you are using the Memory Channel adapters in multiple-active rail mode with the TruCluster Memory Channel Software product, after you have installed the Tru64 UNIX and TruCluster Server Version 5.1A software, you will have to reset the `rm` kernel subsystem configuration `rm_rail_style` variable to zero. The default for `rm_rail_style` for TruCluster Server Version 5.1A is

1, which enables failover pair. See the *Cluster Highly Available Applications* manual for more information.

Figure 11–2 provides a detailed illustration of the first two systems in an 8-node shared SCSI cluster. Table 11–1 lists the components that are used to create the portion of the cluster that is shown in Figure 11–2.

To install the cluster hardware for the first two member systems of an eight-node cluster, follow these steps:

1. Install Memory Channel adapters on member systems 1 and 2. See Chapter 5 for installation and jumper information on the Memory Channel adapters. Delay testing the Memory Channel until you have installed all hardware.
2. Install a Memory Channel hub within 10 meters (32.8 feet) of all eight member systems.
3. Use BN39B-04 (4 meters; 13.1 feet) or BN39B-10 (10 meters; 32.8 feet) to connect the Memory Channel adapters of member systems 1 and 2 to the Memory Channel hub.
4. Refer to the hardware manuals and install the network adapters for the public network on member systems 1 and 2. The public network is not shown in the illustrations in this chapter.
5. Refer to Table 9–3 and install two KZPBA-CB host bus adapters on member system 1 and 2 for the shared SCSI buses that they will use:
 - A shared SCSI bus for member system 1 and 2.
 - A shared SCSI bus for member system 2 with member systems 3, 4, and 5.
 - A shared SCSI bus for member system 1 with member systems 5, 6, and 7.

Ensure that you set the host bus adapter SCSI IDs as follows:

- Member system 1: SCSI bus ID 7 (for both host bus adapters)
 - Member system 2: SCSI bus ID 6 (for both host bus adapters)
6. Ensure that each system (member system 1 and 2) has a BN21W-0B Y cable attached to each KZPBA-CB host bus adapter and an H879-AA HD68 terminator attached to one leg of each BN21W-0B Y cable. Member systems 1 and 2 will be at one end of each of the two SCSI buses they share.

7. Prepare the UltraSCSI BA356 for TruCluster Server use (see Section 10.4.1.3). Ensure that you have installed an H8861-AA VHDCI trilink connector on the UltraSCSI BA356 personality module.

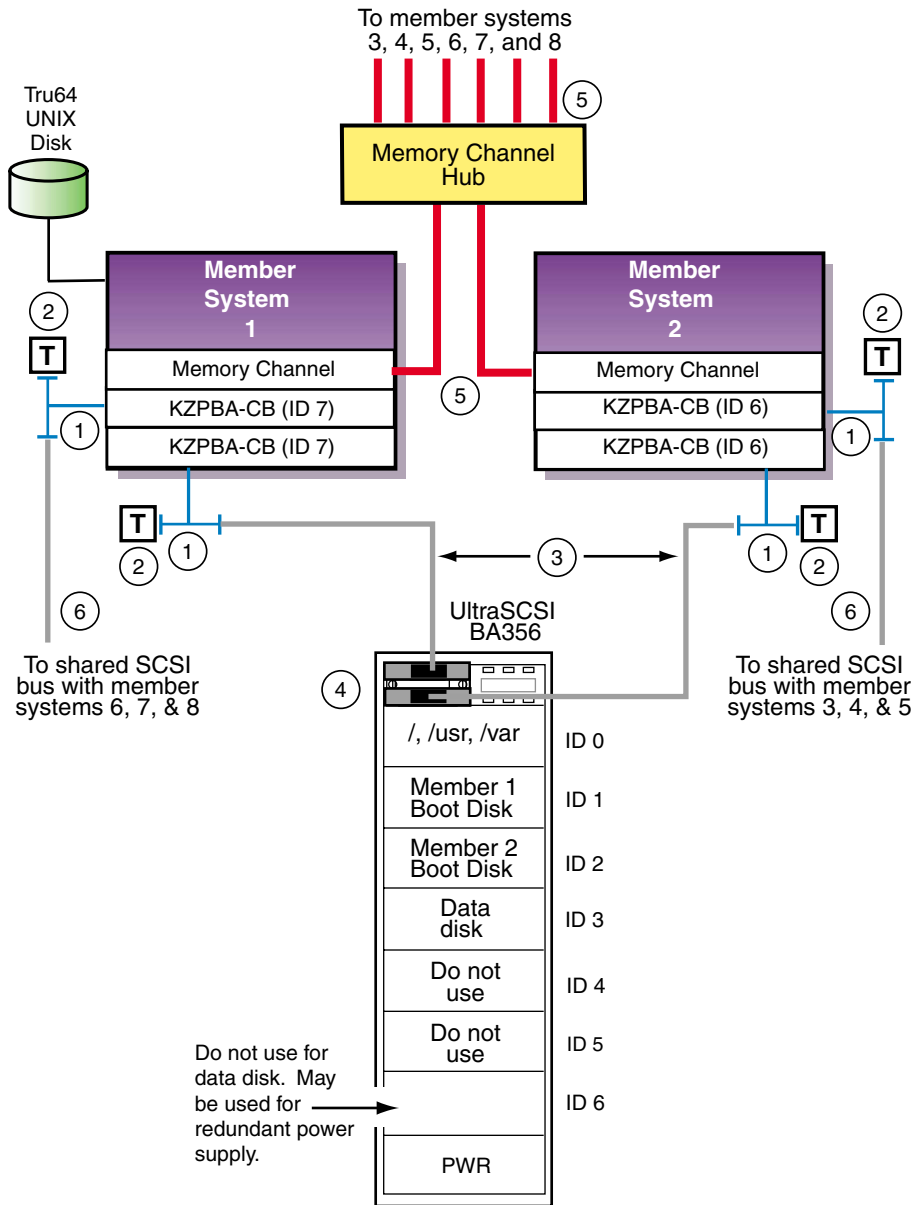
Note

If you need more storage than one UltraSCSI BA356 provides, you can daisy-chain two of them together. See Section 10.4.3.3 for more information.

8. Select one KZPBA-CB host bus adapter on each system. Connect a BN38C, BN38D, or a combination of a BN38E-0B technology adapter cable and a BN37A cable between the open leg of the BN21W-0B Y cable on each system to the H8861-AA VHDCI trilink connector on the UltraSCSI BA356 personality module. This creates the shared SCSI bus between member systems 1 and 2.

The remaining KZPBA-CB on each system has an open leg on its BN21W-0B Y cable. These connections will be used for the other shared SCSI buses.

Figure 11–2: First Two Nodes of an Eight-Node Cluster



ZK-1844U-AI

Table 11–1: Hardware Components Used for Configuration Shown in Figure 11–2

Callout Number	Description
1	BN21W-0B HD68 Y cable
2	H879-AA HD68 terminator
3	BN38C or BN38D HD68 to VHDCI cable ^{a b}
4	H8861-AA VHDCI trilink connector
5	BN39B-04 or BN39B-10 Memory Channel cable
6	BN21K, BN21L, or 328215-00X HD68 to HD68 cable

^a A BN38E-0B technology adapter cable may be connected to a BN37A cable and used in place of a BN38C or BN38D cable.

^b The maximum combined length of the BN38C (or BN38D) and BN37A cables on one SCSI bus segment must not exceed 25 meters (82 feet).

If you have performed each step correctly, each of the first two member systems is prepared to be added to three other member systems on a shared SCSI bus.

You can install Tru64 UNIX and TruCluster Server Version 5.1A software at this time, or you can wait until all cluster hardware is installed.

You need to configure two four-node shared SCSI buses to create your eight-node shared SCSI cluster. The next two sections cover the steps needed to configure member systems 3, 4, and 5 on a shared SCSI bus with member system 2, and member systems 6, 7, and 8 on a shared SCSI bus with member system 1.

11.2.2 Cabling the Second Externally Terminated Shared SCSI Bus

So far, you have configured a two-node externally terminated shared SCSI bus made up of member systems 1 and 2. This section covers the steps needed to configure member systems 3, 4, and 5 on a shared SCSI bus with member system 2.

Figure 11–3 shows a detailed illustration of member systems 2, 3, 4, and 5 on the second shared SCSI bus. Table 11–2 shows the components needed to configure the systems shown in Figure 11–3 into the cluster.

To configure member systems 2, 3, 4, and 5 on the second four-node shared SCSI bus, follow these steps:

1. Install the Memory Channel adapters on member systems 3, 4, and 5. See Chapter 5 for installation and jumper information on the Memory Channel adapters. Delay testing the Memory Channel until you have installed all hardware.

Note

If member systems 1 and 2 are running cluster software, do not run `mc_cable` Memory Channel diagnostics. Shut all systems down to the console level to run the `mc_cable` diagnostic.

2. Use BN39B-04 (4 meters; 13.1 feet) or BN39B-10 (10 meters; 32.8 feet) to connect the Memory Channel adapters of member systems 3, 4, and 5 to the Memory Channel hub.
3. Refer to the hardware manuals and install the network adapters for the public network on member systems 3, 4, and 5. The public network is not shown in the illustrations in this chapter.
4. Referring to Table 9–3, install a KZPBA-CB host bus adapter on member systems 3, 4, and 5. These host bus adapters will be used to form a shared SCSI bus with member system 2.

Ensure that you set the host bus adapter SCSI IDs as follows:

- Member system 2: SCSI ID 6 (which should already be set)
 - Member system 3: SCSI ID 7
 - Member system 4: SCSI ID 5
 - Member system 5: SCSI ID 4
5. Ensure that each system (member system 3, 4, and 5) has a BN21W-0B Y cable attached to the KZPBA-CB host bus adapter.
 6. Ensure that there is an H879-AA terminator attached to one leg of the BN21W-0B on member system 5. Member systems 2 and 5 will be at the end of this shared SCSI bus.
 7. Prepare the UltraSCSI BA356 for TruCluster Server use (see Section 10.4.1.3). Ensure that you have installed an H8861-AA VHDCI trilink connector on the UltraSCSI BA356 personality module.

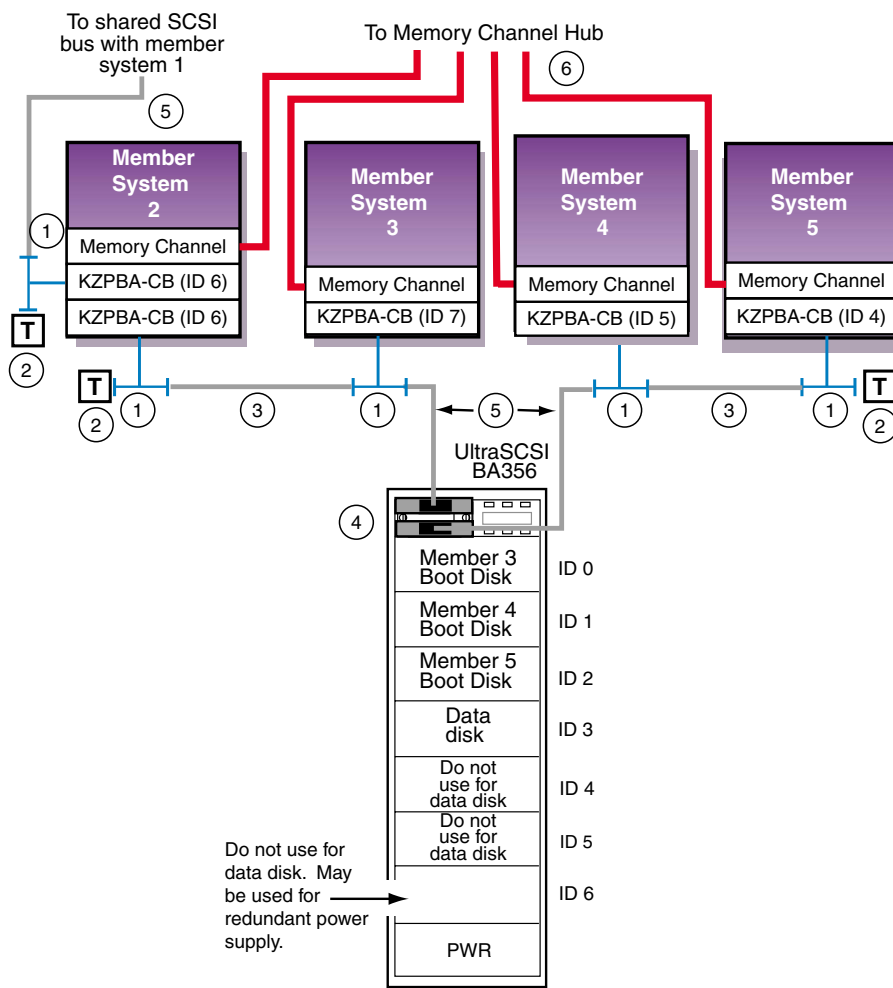
Note

If you need more storage than one UltraSCSI BA356 provides, you can daisy-chain two of them together. See Section 10.4.3.3 for more information.

8. Connect a BN21K, BN21L, or 328215-00X cable between the BN21W-0B Y cables on member system 2 and member system 3.

9. Connect a BN21K, BN21L, or 328215-00X cable between the BN21W-0B Y cables on member system 4 and member system 5.
10. Connect a BN38C, BN38D, or a combination of a BN38E-0B technology adapter cable and a BN37A cable between the open leg of the BN21W-0B on member systems 3 and 4 to the H8861-AA VHDCI tralink connector on the UltraSCSI BA356 personality module.

Figure 11–3: Second Shared SCSI Bus of an Eight-Node Cluster



ZK-1845U-AI

Table 11–2 lists the components that are used to create the cluster shown in Figure 11–3.

Table 11–2: Hardware Components Used for Configuration Shown Figure 11–3

Callout Number	Description
1	BN21W-0B HD68 Y cable
2	H879-AA HD68 terminator
3	BN21K, BN21L, or 328215-00X HD68 to HD68 cable ^a
4	H8861-AA VHDCI tralink connector
5	BN38C or BN38D HD68 to VHDCI cable ^{a b}
6	BN39B-04 or BN39B-10 Memory Channel cable

^a The maximum combined length of the BN21K, BN21L, 328215-00X, BN38C, BN38D, BN38E-0B, and BN37A cables on one SCSI bus segment must not exceed 25 meters (82 feet).

^b A BN38E-0B technology adapter cable may be connected to a BN37A cable and used in place of a BN38C or BN38D cable.

11.2.3 Cabling the Third Externally Terminated Shared SCSI Bus

So far, you have configured a two-node externally terminated shared SCSI bus made up of member systems 1 and 2, and an externally terminated four-node shared SCSI bus with member systems 2, 3, 4, and 5. You need to configure a third externally terminated four-node shared SCSI bus to complete your eight-node shared SCSI cluster.

This section covers the steps needed to configure member systems 1, 6, 7, and 8 on an externally terminated shared SCSI bus.

Figure 11–4 shows a detailed illustration of member systems 1, 6, 7, and 8 on a shared SCSI bus. Table 11–3 lists the components needed to configure the systems shown in Figure 11–4.

To configure member systems 1, 6, 7, and 8 on a four-node shared SCSI bus, follow these steps:

1. Install the Memory Channel adapters on member systems 6, 7, and 8. See Chapter 5 for installation and jumper information on the Memory Channel adapters. Delay testing the Memory Channel until you have installed all hardware.

Note

If member systems 1 and 2 are running cluster software, you should not run `mc_cable` Memory Channel diagnostics. Shut all systems down to the console level to run the `mc_cable` diagnostic.

2. Use BN39B-04 (4 meters; 13.1 feet) or BN39B-10 (10 meters; 32.8 feet) to connect the Memory Channel adapters of member systems 6, 7, and 8 to the Memory Channel hub.
3. Refer to the hardware manuals and install the network adapters for the public network on member systems 6, 7, and 8. The public network is not shown in the illustrations in this chapter.
4. Referring to Table 9–3, install a KZPBA-CB host bus adapter on member system 6, 7, and 8. These host bus adapters will be used to form a shared SCSI bus with member system 1.

Ensure that you set the host bus adapter SCSI IDs as follows:

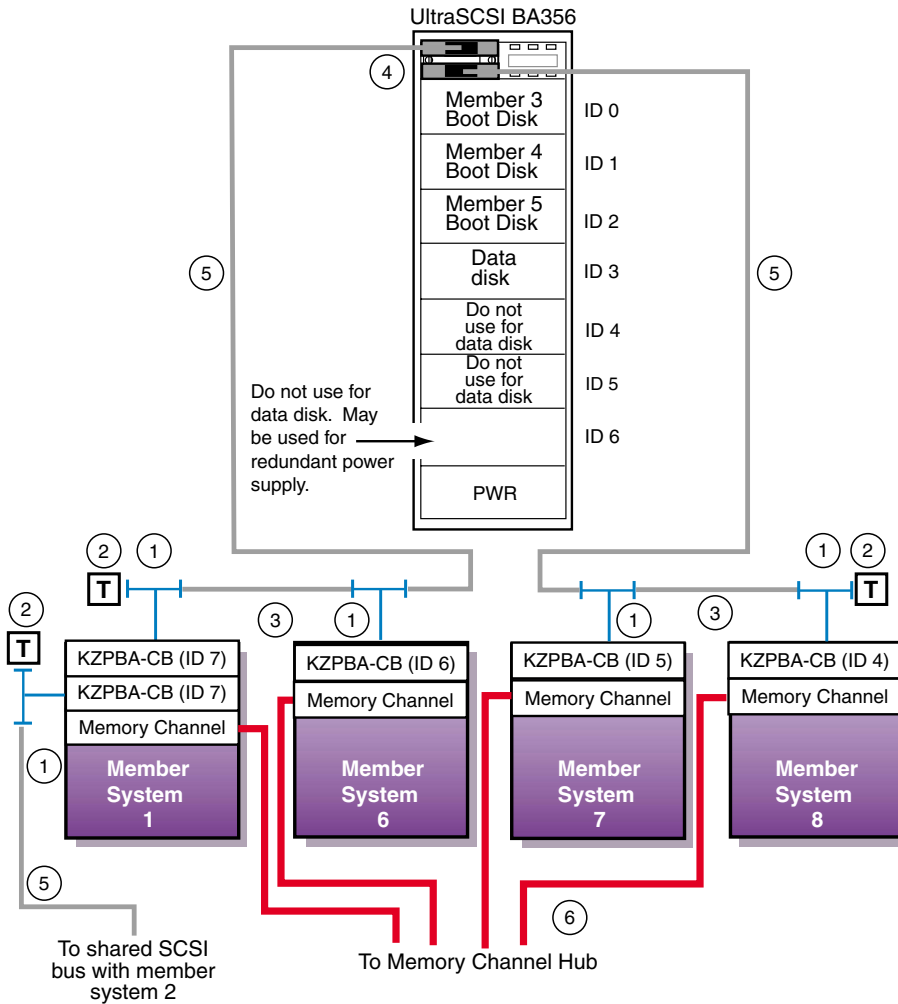
- Member system 1: SCSI bus ID 7 (which should already be set)
 - Member system 6: SCSI bus ID 6
 - Member system 7: SCSI bus ID 5
 - Member system 8: SCSI bus ID 4
5. Ensure that each system (member system 6, 7, and 8) has a BN21W-0B Y cable attached to the KZPBA-CB host bus adapter.
 6. Ensure that there is an H879-AA terminator attached to one leg of the BN21W-0B on member system 8. Member systems 1 and 8 will be at the end of this shared SCSI bus.
 7. Prepare the UltraSCSI BA356 for TruCluster Server use (See Section 10.4.1.3.) Ensure that you have installed an H8861-AA VHDCI tralink connector on the UltraSCSI BA356 personality module.

Note

If you need more storage than one UltraSCSI BA356 provides, you can daisy-chain two of them together. See Section 10.4.3.3 for more information.

8. Connect a BN21K, BN21L, or 328215-00X cable between the BN21W-0B Y cables on member system 1 and member system 6.
9. Connect a BN21K, BN21L, or 328215-00X cable between the BN21W-0B Y cables on member system 7 and member system 8.
10. Connect a BN38C, BN38D, or a combination of a BN38E-0B technology adapter cable and a BN37A cable between the open leg of the BN21W-0B on member systems 6 and 7 to the H8861-AA VHDCI tralink connector on the UltraSCSI BA356 personality module.

Figure 11–4: Third Shared SCSI Bus of an Eight-Node Cluster



ZK-1846U-AI

Table 11–3 lists the components that are used to create the cluster shown in Figure 11–4.

Table 11–3: Hardware Components Used for Configuration Shown in Figure 11–4

Callout Number	Description
1	BN21W-0B HD68 Y cable
2	H879-AA HD68 terminator
3	BN21K, BN21L, or 328215-00X HD68 to HD68 cable
4	H8861-AA VHDCI trilink connector
5	BN38C or BN38D HD68 to VHDCI cable ^{a b}
6	BN39B-04 or BN39B-10 Memory Channel cable

^a A BN38E-0B technology adapter cable may be connected to a BN37A cable and used in place of a BN38C or BN38D cable.

^b The maximum combined length of the BN21K, BN21L, 328215-00X, BN38C, BN38D, BN38E-0B, and BN37A cables on one SCSI bus segment must not exceed 25 meters (82 feet).

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 (/)					
/usr					
/var					
Member 1 boot disk					
Member 2 boot disk					
Member 3 boot disk					
Member 4 boot disk					
Quorum disk					

Index

Numbers and Special Characters

20/40-GB DLT Tape Drive, 8–11

- cabling, 8–12
- capacity, 8–11
- cartridges, 8–11
- connectors, 8–11
- setting SCSI ID, 8–11

40/80-GB DLT Drive, 8–15

- cabling, 8–16
- capacity, 8–15
- cartridges, 8–15
- connectors, 8–15
- setting SCSI ID, 8–15

A

ACS V8.5, 2–7

arbitrated loop

- AL_PA, 6–3
- characteristics, 6–6
- compared with fabric topology, 6–7
- defined, 6–3
- setting port_n_topology, 6–48
- use of wwidmgr -set, 6–40

Array Control Software

(See ACS V8.5)

ATL

- TL893 tape library, 8–49
- TL896 tape library, 8–49

Automated Tape Library

(See ATL)

availability

- increasing, 4–3

B

BA350 storage shelf, 10–9

- preparing for shared SCSI usage, 10–15
- termination, 10–3, 10–15

BA356 storage shelf, 10–9

- DS-DWZZH-03 installed in, 2–13, 3–10, 4–7, 9–3, 10–34
- jumper, 10–9, 10–11
- personality module address switches, 10–11
- preparing, 10–15, 10–17
- preparing for shared SCSI usage, 10–16
- SCSI ID selection, 10–16
- selecting SCSI IDs, 10–11
- termination, 10–3, 10–9, 10–11

BA370 storage shelf

- DS-DWZZH-03 installed in, 2–13, 3–10

BC12N-10 Memory Channel link cable

- length, 2–4

BC39B-10 Memory Channel link cable

- length, 2–5

bootdef_dev console environment variable, 6–67, 6–72, 6–73, 6–76

- resetting, 6–72
- setting, 6–67, 6–74, 6–76

bus hung message, 2–12

bus_probe_algorithm console environment variable, 2–10

buses

- data paths, 3–5

- extending differential, 10–2
- narrow data path, 3–5
- speed, 3–5
- terminating, 3–7, 10–5, 10–8
- wide data path, 3–5

C

caa_relocate command, 5–15t, 5–27t

cable length restrictions

- shared SCSI buses, 3–7t

cables

- BC12N-10 Memory Channel link cable, 2–4, 5–7
- BN39B-01 Memory Channel link cable, 5–7, 5–9
- BN39B-04 Memory Channel link cable, 5–7, 5–9
- BN39B-10 Memory Channel link cable, 2–5, 5–7, 5–9
- ESL9326D enterprise library, 8–77 supported, 2–14

cablings

- 20/40-GB DLT Tape Drive, 8–12
- 40/80-GB DLT Drive, 8–16
- DS-TZ89N-TA tabletop tape drive, 8–10
- DS-TZ89N-VW SBB tape drive, 8–9
- ESL9326D enterprise library, 8–74, 8–76, 8–77
- RA3000 using external termination, 10–30
- RA3000 using UltraSCSI hubs, 3–27
- TL881/891 DLT MiniLibrary, 8–64, 8–67
- TL890 tape library, 8–31
- TL891 tape library, 8–26, 8–31
- TL892 tape library, 8–26, 8–31
- TL893 tape library, 8–56
- TL894 tape library, 8–42
- TL895 tape library, 8–49

- TL896 tape library, 8–56
- TZ885 minilibrary, 8–19
- TZ887 minilibrary, 8–22
- TZ88N-TA tabletop tape drive, 8–6
- TZ88N-VA SBB tape drive, 8–4

changing HSG80 failover modes, 6–77

clu_create command, 6–71, 11–4

cluster

- expanding, 3–8, 10–6
- increasing availability, 4–3
- planning, 4–2

cluster interconnects

- increasing availability, 4–3

CONFIGURATION RESTORE

command, 6–51

configuration restrictions, 2–8

configuration variable

- rm_rail_style, 5–1

configuring an eight-node SCSI cluster

- cabling first two nodes, 11–6

configuring base unit as slave, 8–34, 8–70

connections to HSG80, 6–77

connectors supported, 2–16

console environment variable

- bootdef_dev, 6–68, 6–72, 6–74, 6–76
- bus_probe_algorithm, 2–10
- ffauto, 6–73
- ffnext, 6–73
- hp_count, 7–10
- hp_qbb_maskn, 7–10
- N, 6–62
- scm_csb_master_eligible, 7–5, 7–12
- wwid, 6–62

console serial bus

(See CSB)

console terminal

- need for, 7–9
- terminal emulator, 7–9
- terminal server, 7–9

CSB, 7-4
nodes, 7-5
purpose, 7-4

D

data path for buses, 3-5

default SCSI IDs

ESL9326D enterprise library, 8-75
TL881/TL891 DLT MiniLibrary,
8-63
TL890 tape library, 8-37
TL891 tape library, 8-37
TL892 tape library, 8-37
TL893 tape library, 8-52
TL894 tape library, 8-38
TL895 tape library, 8-45
TL896 tape library, 8-52

device name, 6-61

device unit number

console uses, 6-61
setting, 6-61

diagnostics

Memory Channel, 5-12

differential SCSI buses, 3-4

differential transmission, 3-4

Digital Linear Tape

(See DLT)

disk devices

restrictions, 2-10
setting up, 3-16, 10-14

disk placement

clusterwide /usr, 1-10
clusterwide /var, 1-10
clusterwide root, 1-10
member boot, 1-10
quorum, 1-10

disklabel command, 6-71

displaying device information

KZPBA-CB, 4-9t, 4-15, 9-6t,
9-7t
KZPSA-BB, 9-6t, 9-7t

DLT

20/40-GB DLT Tape Drive, 8-11
40/80-GB DLT Drive, 8-15
TZ885 minilibrary, 8-18
TZ887 minilibrary, 8-21

DLT MiniLibrary

configuring TL881/TL891 as slave,
8-70
configuring TL891 tape library as
slave, 8-34
TL881 tape library, 8-58
TL891 tape library, 8-58

DS-BA356 storage unit

(See UltraSCSI BA356 storage
unit)

DS-BA35X-DA personality module,

3-3, 3-5, 4-8, 10-2, 10-3

DS-DWZZH-03 UltraSCSI hub, 3-9

bus connectors, 3-10
bus isolation, 2-13
description, 2-13
installed in, 2-13, 3-10, 4-7, 9-3,
10-34
internal termination, 3-9
radial disconnect, 2-13
SBB, 3-9
SCSI ID, 3-10
support on, 3-9
termpwr, 3-9
transfer rate, 2-13

DS-DWZZH-05 UltraSCSI hub,

3-10
bus connectors, 3-10
bus isolation, 2-13
configurations, 3-16
description, 2-13
fair arbitration, 3-10
installed in, 2-13, 3-11, 4-7, 9-3,
10-34
internal termination, 3-9
radial disconnect, 2-13

- SBB, 3–10
- SCSI ID, 3–10
- termpwr, 3–9
- transfer rate, 2–13
- DS-TZ89N-TA tabletop tape drive**
 - cabling, 8–10
 - setting SCSI ID, 8–10
- DS-TZ89N-VW SBB tape drive**
 - cabling, 8–9
 - setting SCSI ID, 8–7
- dual-redundant controllers, 1–14**
- DWZZA signal converter**
 - incorrect hardware revision, 2–12
 - termination, 10–3, 10–16
 - upgrade, 2–13
- DWZZB signal converter**
 - termination, 10–3, 10–16
- DWZZH-03**
 - (See DS-DWZZH-03 UltraSCSI hub)
- DWZZH-05**
 - (See DS-DWZZH-05 UltraSCSI hub)

E

- eight-node cluster, 1–20**
 - cabling first two nodes, 11–5
 - cabling second shared SCSI bus, 11–9
 - cabling third shared SCSI bus, 11–12
 - requirements, 11–3
 - using Fibre Channel, 11–1
 - using shared SCSI, 11–1
- EMA12000 Enterprise Modular Array**
 - configuring, 2–9
 - port configuration, 2–9
 - transparent failover mode, 2–9
 - unit configuration, 2–9
- emxmgr command**
 - d, 6–82
 - displaying adapters, 6–82

- displaying target ID mapping, 6–82
- displaying topology, 6–82, 6–84
 - m, 6–82
 - t, 6–82, 6–84
- use, 6–82, 6–83, 6–85
- using interactively, 6–85
- enterprise library**
 - (See ESL9326D Enterprise Library tape drive)
- ESA12000 Enterprise Storage Array**
 - configuring, 2–9
 - port configuration, 2–9
 - replacing controllers of, 6–51
 - transparent failover mode, 2–9
 - unit configuration, 2–9
- ESL9000 series tape library**
 - (See ESL9326D Enterprise Library tape drive)
- ESL9326D Enterprise Library tape drive**
 - cables, 8–77
 - cabling, 8–74, 8–77
 - capacity, 8–74
 - firmware, 8–75
 - internal cabling, 8–76
 - number of drives, 8–74, 8–75
 - part numbers, 8–74
 - SCSI connectors, 8–77
 - setting SCSI ID, 8–75
 - tape cartridges, 8–74
 - tape drives, 8–74, 8–75
 - termination, 8–77
 - upgrading, 8–74

F

- F_Port fabric port, 6–4**
- fabric, 6–3, 6–5**
- failover mode**
 - changing, 6–77
 - multiple-bus, 6–77
 - set nofailover, 6–78
 - transparent, 6–77

FCP, 6–2
ffauto console environment variable, 6–73
ffnext console environment variable, 6–73
fiber-optic cable
Fibre Channel, 2–8, 6–27
Memory Channel, 2–5, 5–6, 5–7, 5–9, 5–11
Fibre Channel
AL_PA arbitrated loop physical address, 6–3
arbitrated loop, 6–3, 6–6
configurations supported, 6–8, 6–10, 6–17
data rates, 6–2
distance, 6–2
F_Port fabric port, 6–4
fabric, 6–3, 6–5
FL_Port fabric loop port, 6–4
frame, 6–3
hub installation, 6–34
N_Port node port, 6–3
NL_Port node loop port, 6–3
point-to-point, 6–5
restrictions, 2–6, 2–7
supported switches, 2–8
switch installation, 6–26
table of known connections, 6–77
terminology, 6–2
topology, 6–4, 6–83, 6–85
Fibre Channel Protocol
(*See* FCP)
Fibre Channel tape controller, 2–9
Fibre Channel tape controller II, 2–9
file
sysconfigtab, 5–18
/var/adm/messages, 6–41
firmware
35/70 DLT tape drives, 8–75

ESL9326D enterprise library, 8–75
fail-safe loader, 7–18
KZPBA-CB, 2–10, 4–7, 7–13, 9–5
KZPSA-BB, 2–9, 9–5t, 9–18
release notes, 4–4
reset system for update, 7–21, 9–19
SRM console, 4–7t, 7–4, 9–5t, 9–18
SROM, 7–18
TL894 tape library, 8–38
TL895 tape library, 8–45
update CD-ROM, 4–4
updating, 9–18
updating GS80/160/320, 7–18
FL_Port fabric loop port, 6–4

G

GBIC
GBIC-LW, 2–8, 6–27
GBIC-SW, 2–8, 6–27, 6–37
Gigabit Interface Converter
(*See* GBIC)
Gigabit Link Module
(*See* GLM)
GLM, 2–8

H

hard partition
defined, 7–1
repartitioning, 7–8
requirements, 7–2
hardware components
Fibre Channel, 2–6
SCSI adapters, 2–6
SCSI cables, 2–14
SCSI signal converters, 2–12
storage shelves, 10–9
terminators, 2–16

- trilink connectors, 2–16
- hardware configuration**
 - bus termination, 3–7, 10–5
 - cables supported, 2–1
 - disk devices, 3–16, 10–14
 - hardware requirements, 2–1
 - hardware restrictions, 2–1
 - requirements, 3–1, 10–2
 - SCSI bus adapters, 2–6
 - SCSI bus speed, 3–5
 - SCSI cables, 2–14
 - SCSI signal converters, 10–2
 - storage shelves, 3–16, 10–14
 - terminators, 2–16
 - terminators supported, 2–1
 - trilink connectors, 2–16
 - trilinks supported, 2–1
 - Y cables supported, 2–1
- hierarchical switch power manager**
 - (See HPM)
- host bus adapters**
 - (See KGPSA, KZPBA-CB, KZPSA-BB)
- hp_count console environment variable**, 7–10
- hp_qbb_mask console environment variable**, 7–10
- HPM**, 7–5
- HSG60 controller**, 1–14
 - configuring, 2–9
 - port configuration, 2–9
 - SCSI-2, 6–45
 - SCSI-3, 6–45
 - transparent failover mode, 2–9
 - unit configuration, 2–9
- HSG80 controller**, 1–14
 - ACS, 2–7
 - changing failover modes, 6–77
 - configuring, 2–9, 6–42
 - configuring for loop, 6–48
 - multiple-bus failover, 6–44
 - obtaining the worldwide name of, 6–50
 - port configuration, 2–9
 - port_n_topology, 6–44
 - port_n_topology in loop, 6–48
 - replacing, 6–51
 - resetting offsets, 6–77
 - SCSI-2, 6–45
 - SCSI-3, 6–45
 - setting controller values, 6–42, 6–44
 - setting controller values for loop, 6–48
 - transparent failover mode, 2–9
 - unit configuration, 2–9
- HSZ failover**
 - multiple-bus, 1–15
 - transparent, 1–15
- HSZ20 controller**
 - shared SCSI bus and, 10–25
- HSZ22**
 - (See RA3000)
- HSZ40 controller**, 1–14
 - shared SCSI bus and, 10–25
- HSZ50 controller**, 1–14
 - shared SCSI bus and, 10–25
- HSZ70 controller**, 1–14
 - fast wide differential SCSI and, 3–3
- HSZ80 controller**, 1–14
- hubs**, 3–9
 - (See also DS-DWZZH-03 UltraSCSI hub; DS-DWZZH-05 UltraSCSI hub)
 - GBIC, 6–34
 - installing, 6–34
 - interface module, 6–34
 - overview, 6–34
 - UltraSCSI, 3–9
- hwmgr command**, 6–69

I

I/O buses

number of, 2–6

I/O risers

cables, 7–3, 7–9

local, 7–3, 7–9

remote, 7–3, 7–9

init command, 6–37, 6–65, 6–68,
6–73, 6–76

initialize

after setting bootdef_dev console
environment variable, 6–68,
6–73, 6–76

after using wwidmgr, 6–37, 6–65

installation, 3–16

(*See also* hardware
configuration)

hub, 6–34

KGPSA, 6–36

KZPBA-CB, 4–6t, 9–2t

KZPSA-BB, 9–2t, 9–4

MC2, 5–10

MC2 cables, 5–9

Memory Channel, 5–6

Memory Channel cables, 5–7

Memory Channel hub, 5–7

optical converter, 5–6

optical converter cables, 5–10

switch, 6–26

internal cabling

ESL9326D enterprise library, 8–76

TL893 tape library, 8–53

TL896 tape library, 8–53

J

jumpers

MC1 and MC1.5 (CCMAA), 5–2

MC2 (CCMAB), 5–4

MC2 (CCMLB), 5–5

K

KGPSA Fibre Channel host bus adapter

GLM, 6–36

installing, 6–36

mounting bracket, 6–36

obtaining the worldwide name of,
6–41

setting to run on a loop, 6–39

setting to run on fabric, 6–37

KZPBA-CB UltraSCSI host bus adapter

displaying device information,

4–9t, 4–16, 9–6t, 9–7t

restrictions, 2–10

setting SCSI ID, 4–17

termination resistors, 4–9t, 4–17,
9–5t, 9–7t

using in cluster, 4–6, 9–2

KZPSA-BB SCSI bus adapter

displaying device information,

9–6t, 9–7t

installation, 9–4

restrictions, 2–9

setting bus speed, 9–17

setting SCSI ID, 9–17

setting termination power, 9–17

termination resistors, 9–5t, 9–7t

updating firmware, 9–18

using in cabling, 9–2

L

Legato Networker, 2–9

LFU, 9–18

booting, 9–18

updating firmware, 9–18

updating firmware with, 9–18

link cable

installation, 5–7

Loadable Firmware Update utility

(*See* LFU)

Logical Storage Manager

(*See* LSM mirroring)

loop topology

AL_PA, 6-3
characteristics, 6-6
defined, 6-3
setting controller values, 6-48
setting port_n_topology, 6-48
use of wwidmgr -set, 6-40

LSM mirroring

across SCSI buses, 1-12
clusterwide /usr, 1-13
clusterwide /var, 1-13
clusterwide data disks, 1-13

M

MA6000 modular array

configuring, 2-9
port configuration, 2-9
transparent failover mode, 2-9
unit configuration, 2-9

MA8000 modular array

configuring, 2-9
port configuration, 2-9
transparent failover mode, 2-9
unit configuration, 2-9

mc_cable command, 5-12

mc_diag command, 5-12

member systems

improving performance, 4-2
increasing availability, 4-2
requirements, 2-1

Memory Channel

diagnostics, 5-12
failover pair, 5-1
installation, 5-2, 5-6
interconnect, 2-4
jumpers, 5-2
optical converter installation, 5-6,
5-10, 5-11
versions, 2-3

Memory Channel diagnostics

mc_cable, 5-12

mc_diag, 5-12

Memory Channel fiber optics

BN34R fiber-optic cable, 2-5
cable length, 2-5
optical converter, 2-5
optical converter installation, 5-6,
5-10, 5-11

Memory Channel hub

installation, 5-7

Memory Channel interconnects

adding, 5-14
maintaining, 5-14
restrictions, 2-3
rolling upgrade, 5-14
setting up, 5-1

Memory Channel jumpers

MC1 and MC1.5 (CCMAA), 5-2
MC2 (CCMAB), 5-4
MC2 (CCMLB), 5-5

messages

bus hung, 2-12
LFU help, 7-20
MC2 adapter insufficient memory,
5-18

MiniLibrary

TL881 tape library, 8-58
TL891 tape library, 8-58

minimum cluster configuration, 1-5

MUC, 8-50

setting SCSI ID, 8-52

MUC switch functions

TL893 tape library, 8-51
TL896 tape library, 8-51

multi-unit controller

(*See* MUC)

multimode fiber, 6-27, 6-34

multiple-bus failover, 3-18, 3-23, 6-44

changing from transparent failover,
6-77

example configurations, 6-11

NSPOF, 3-18

setting, 6–44, 6–78

N

N_Port node port, 6–3

NL_Port node loop port, 6–3

no single point of failure

(*See* NSPOF)

node name, 6–50

non-Ultra BA356 storage shelf

preparing, 10–15

NSPOF, 1–14, 3–18

O

optical cable, 6–27, 6–34

optical converter

cable connection, 5–6

installation, 5–6

options list, 2–6, 3–16

P

partitioned storagesets, 3–18

PBM, 7–5

PCI backplane manager

(*See* PBM)

PCI drawer

contents, 7–4, 7–5

primary, 7–4, 7–5

slot layout, 2–2

performance

improving, 4–2

personality module, 3–3, 10–13

(*See also* signal converters)

planning the hardware

configuration, 4–2

point-to-point, 6–5

port name, 6–50

power system manager

(*See* PSM)

preparing storage shelves

BA350, 10–15

BA350 and BA356, 10–19

BA356, 10–16, 10–21

UltraSCSI BA356, 10–17, 10–23

Prestoserve

using in a cluster, 4–3

PSM, 7–5

Q

quorum disk

configuring, 1–4

LSM and, 1–5

number of votes, 1–5

R

RA3000, 1–14

active/active mode, 3–19

active/passive mode, 3–19

cabling for use with UltraSCSI

hubs, 3–27

cabling using external termination,

10–30

device restrictions, 2–11

part numbers, 3–27

RAID levels supported, 3–26,

10–29

restrictions, 2–11

UPS required, 3–27, 10–29

using DWZZH-05 host port with,

3–31

RA8000 storage array

configuring, 2–9

port configuration, 2–9

replacing controllers of, 6–51

transparent failover mode, 2–9

unit configuration, 2–9

radial connection

bus termination, 3–9

UltraSCSI hub, 3–9

RAID, 1–14
RAID Array 3000
 (*See* RA3000)
RAID array controllers
 advantages, 3–17
 preparing, 10–25
 shared SCSI bus and, 10–25
 using in ASE, 10–25
Redundant Array of Independent Disks
 (*See* RAID)
repartitioning procedure, 7–8
replacing
 HSG80 controller, 6–51
reset
 (*See* system reset)
resetting offsets, 6–77
restrictions, 2–8
 disk devices, 2–10
 KZPBA-CB adapters, 2–10
 KZPSA adapters, 2–9
 Memory Channel interconnects,
 2–3
 RA3000, 2–11
 SCSI bus adapters, 2–6
rm_rail_style, 5–1
rolling upgrade
 MC1 to MC2, 5–14

S

SAVE_CONFIGURATION

command, 6–51

SBB

DS-DWZZH-03, 3–9
 DS-DWZZH-05, 3–10
 DS-TZ89N-VW, 8–7
 fast, narrow SCSI drives, 3–3
 fast, wide SCSI drives, 3–3
 installed in BA350, 10–9
 installed in BA356, 10–11
 installed in UltraSCSI BA356,
 10–13
 internal termination, 10–4

part number ends in -VA, 3–3
 part number ends in -VW, 3–3
 shared SCSI bus with SBB tape
 drives, 8–5
 signal converter, 2–12, 10–2
 slots available for with DWZZH,
 3–11
 tape drive SCSI ID, 8–3t, 8–8t
 TZ88N-VA, 8–2

SC connector

, 6–27

SCM

, 7–4

master, 7–5
 master-designate, 7–5
 powered by, 7–4
 purpose, 7–4
 slave, 7–5

SCM command

power on, 7–12
 set hp_count, 7–10
 set hp_qbb_mask, 7–10
 show nvr, 7–10

scm_csb_master_eligible_console environment variable

, 7–5,
 7–12

SCSI bus

, 3–2

(*See also* shared SCSI bus
 requirements; shared SCSI
 buses)
 number of devices supported, 3–2
 with BA350 and BA356, 10–19
 with two BA356s, 10–21
 with two UltraSCSI BA356s, 10–23

SCSI cables

, 2–14

(*See also* cables)
 requirement, 2–15

SCSI controllers

bus speed for, 3–5

SCSI IDs

BA350, 10–9
 BA350 storage shelves, 10–15
 BA356, 10–11, 10–16
 HSZ20 controller, 10–25
 HSZ40 controller, 10–25

- HSZ50 controller, 10–25
- in BA356, 10–11
- in UltraSCSI BA356, 10–13
- KZPBA-CB, 4–17
- RAID subsystem controllers, 10–25
- requirement, 3–6
- setting, 4–9t, 4–17, 8–25, 9–17, 10–17
- UltraSCSI BA356, 10–13, 10–17
- SCSI targets**
 - number of, 2–8
- SCSI terminators supported**, 2–16
- SCSI-2 bus**, 3–6
- SCSI_VERSION**
 - SCSI-2, 6–45
 - SCSI-3, 6–45
- selecting BA356 disk SCSI IDs**, 10–11
- selecting UltraSCSI BA356 disk SCSI IDs**, 10–13
- set bootdef_dev command**, 6–68, 6–73, 6–76
- SET FAILOVER COPY = THIS_CONTROLLER command**, 1–15
- set ffauto command**, 6–73
- set ffnext command**, 6–73
- set hp_count command**, 7–10
- set hp_qbb_maskn command**, 7–10
- SET MULTIBUS_FAILOVER COPY = THIS_CONTROLLER command**, 3–18
- set scm_csb_master_eligible command**, 7–12
- setting bus speed**
 - KZPSA-BB, 9–17
- setting SCSI ID**
 - 20/40-GB DLT Tape Drive, 8–11
 - 40/80-GB DLT Drive, 8–15
 - DS-TZ89N-TA tabletop tape drive, 8–10
 - DS-TZ89N-VW SBB tape drive, 8–7
 - ESL9326D enterprise library, 8–75
 - KZPBA-CB, 4–17, 9–16
 - KZPSA-BB, 9–17
 - MUC, 8–52
 - TL881/891 DLT MiniLibrary, 8–62, 8–72
 - TL891 tape library, 8–24
 - TL892 tape library, 8–24
 - TL893 tape library, 8–52
 - TL894 tape library, 8–38
 - TL896 tape library, 8–52
 - TZ885 minilibrary, 8–18
 - TZ887 minilibrary, 8–21
 - TZ88N-TA tabletop tape drive, 8–6
 - TZ88N-VA SBB tape drive, 8–2
- shared SCSI bus requirements**
 - externally terminated configurations, 10–2
 - radial configurations with non-UltraSCSI hardware, 10–2
 - radial configurations with UltraSCSI hardware, 3–1
- shared SCSI buses**, 4–3
 - adding devices, 10–6
 - assigning SCSI IDs, 3–6
 - cable length restrictions, 3–7
 - connecting devices, 3–8, 10–6
 - device addresses, 3–6
 - differential, 3–4
 - number of, 2–6, 4–3
 - requirements, 3–2
 - single-ended, 3–4
 - using tralink connectors, 10–6
 - using Y cables, 10–6
- shared storage**
 - BA350 storage shelf, 10–15
 - increasing capacity, 4–2, 4–3
 - non-UltraSCSI BA356 storage shelf, 10–15
 - RAID subsystem array controller, 10–25

- UltraSCSI BA356 storage shelf, 10-15, 10-17
- show config command**, 4-10, 4-13, 9-6t, 9-7t, 9-9, 9-11
- show csb command**, 7-16
- show device command**, 4-9t, 4-12, 4-13, 9-6t, 9-7t, 9-11
- show nvr command**, 7-10, 7-19
- show system command**, 7-14
- SHOW THIS_CONTROLLER command**, 6-50
- signal converters**, 10-2
 - creating differential bus, 10-2
 - differential I/O module, 10-2
 - differential termination, 10-3
 - DS-BA35X-DA personality module, 3-5, 10-4
 - extending differential bus length, 10-2
 - fast SCSI bus speed, 10-2
 - requirement, 10-2
 - restrictions, 2-12
 - SBB, 10-2
 - single-ended termination, 10-3
 - standalone, 10-2
 - termination, 10-2, 10-4
- single-ended SCSI buses**, 3-4
- single-ended transmission**, 3-4
- single-mode fiber**, 6-27
- SRM console firmware**
 - control transferred to, 7-11, 7-13
 - executed out of primary QBB memory, 7-20
 - includes ISP1020/1040-based PCI option firmware, 7-13
 - obtaining firmware release notes, 4-4
 - standard I/O module and, 7-4
 - updating, 4-7t, 7-21, 9-5t, 9-7t
 - updating with LFU, 9-18
- storage shelves**, 10-8, 10-9, 10-13
 - attaching to shared SCSI bus, 10-8, 10-13

- overview, 10-8, 10-13
- setting up, 3-16, 10-14
- StorageWorks building block** (*See* SBB)
- subscriber connector** (*See* SC connector)
- supported options** (*See* options list)
- switch**
 - 10Base-T Ethernet connection, 6-26
 - changing password, 6-33
 - changing user names, 6-33
 - front panel, 6-26, 6-29
 - GBIC, 6-27
 - installing, 6-26
 - interface module, 6-27
 - overview, 6-26
 - setting Ethernet IP address, 6-30
 - setting switch name, 6-33
 - supported Fibre Channel, 2-8
 - telnet session, 6-33
- sysconfig command**, 5-18
- sysconfigtab file**, 5-18
- system control manager** (*See* SCM)
- system reference manual firmware** (*See* SRM console firmware)
- system reset**
 - after using wwidmgr, 6-37, 6-65

T

- table of known connections**, 6-77
- terminal emulator**, 7-9
- terminal server**, 7-9
- termination**, 10-13
 - BA356, 10-11
 - DWZZA, 10-16
 - DWZZB, 10-16
 - ESL9326D enterprise library, 8-78
 - terminating the shared bus, 3-7, 10-5

- UltraSCSI BA356, 10–14
- termination resistors**
 - KZPBA-CB, 4–9t, 4–17, 9–5t, 9–7t
 - KZPSA-BB, 9–5t, 9–7t
- terminators supported**, 2–16
- TL881 tape library**, 8–58
- TL881/891 DLT MiniLibrary**
 - cabling, 8–64, 8–67
 - capacity, 8–58, 8–60
 - components, 8–59
 - configuring base unit as slave, 8–70
 - models, 8–58
 - performance, 8–60
 - powering up, 8–71
 - setting the SCSI ID, 8–62, 8–72
- TL890 tape library**
 - cabling, 8–31
 - default SCSI IDs, 8–37
 - powering up, 8–36
 - setting SCSI ID, 8–36
- TL891 tape library**, 8–24, 8–58
 - cabling, 8–26, 8–31
 - configuring as slave, 8–34
 - default SCSI IDs, 8–26, 8–37
 - setting SCSI ID, 8–24, 8–25, 8–36
 - shared SCSI usage, 8–24
- TL892 tape library**, 8–24
 - cabling, 8–26, 8–31
 - configuring as slave, 8–34
 - default SCSI IDs, 8–26, 8–37
 - setting SCSI ID, 8–24, 8–25, 8–36
 - shared SCSI usage, 8–24
- TL893 tape library**, 8–50
 - cabling, 8–53, 8–56
 - MUC switch functions, 8–51
 - setting SCSI ID, 8–52
- TL894 tape library**
 - cabling, 8–42
 - setting SCSI ID, 8–38
- TL895 tape library**
 - cabling, 8–49
 - default SCSI IDs, 8–45
 - setting SCSI ID, 8–45
- TL896 tape library**, 8–50
 - cabling, 8–53, 8–56
 - MUC switch functions, 8–51
 - setting SCSI ID, 8–52
- transparent failover**, 1–15, 3–17
 - changing to multiple-bus failover, 6–77
 - example configuration, 6–9
- trilink connectors**
 - connecting devices with, 10–6
 - requirement, 2–16
 - supported, 2–16
- TZ885 minilibrary**, 8–18
 - cabling, 8–19
 - setting SCSI ID, 8–18
- TZ887 minilibrary**, 8–21
 - cabling, 8–22
 - setting SCSI ID, 8–21
- TZ88N-TA tabletop tape drive**, 8–2
 - cabling, 8–6
 - setting SCSI ID, 8–6
- TZ88N-VA SBB tape drive**, 8–2
 - cabling, 8–4
 - setting SCSI ID, 8–2
- TZ89 tape drive**, 8–7

U

- UltraSCSI BA356**
 - disable termination, 10–17
 - DS-BA35X-DA personality module, 3–3
 - DS-DWZZH-03 installed in, 2–13, 3–10, 4–7, 9–3, 10–34
 - DS-DWZZH-05 installed in, 2–13, 3–10, 3–11, 4–7, 9–3, 10–34
 - fast narrow SCSI drives, 3–3
 - fast wide SCSI drives, 3–3
 - jumper, 10–14

personality module address
switches, 10–13
power supply, 3–3
preparing, 10–15, 10–17
preparing for shared SCSI usage,
10–17
SCSI ID selection, 10–13, 10–17
termination, 10–13

UltraSCSI host adapter

host input connector, 3–3
with non-UltraSCSI BA356, 3–3
with UltraSCSI BA356, 3–3

UltraSCSI hubs, 3–9

(*See also* DS-DWZZH-03
UltraSCSI hub;
DS-DWZZH-05 UltraSCSI
hub)

upgrade

DWZZA, 2–13

upgrading ESL9326D, 8–74

V

/var/adm/messages file, 6–41
variable

(*See* console environment
variable)

**Very High Density Cable
Interconnect**

(*See* VHDCI)

VHDCI, 3–3

W

WorldWide ID Manager

(*See* wwidmgr command)

worldwide name

description, 6–41

wwidmgr command

-clear, 6–62
-quickset, 6–61, 6–63
-set, 6–40
-set in loop mode, 6–40
-show, 6–37, 6–63, 6–66

Y

Y cables

connecting devices with, 10–6
supported, 2–14