

TruCluster Server

Hardware Configuration

Part Number: AA-RHGWB-TE

April 2000

Product Version: TruCluster Server Version 5.0A

Operating System and Version: Tru64 UNIX Version 5.0A

This manual describes how to configure the hardware for a TruCluster Server environment. TruCluster Server Version 5.0A runs on the Tru64™ UNIX® operating system.

© 2000 Compaq Computer Corporation

COMPAQ and the Compaq logo Registered in U.S. Patent and Trademark Office. TruCluster and Tru64 are trademarks of Compaq Information Technologies Group, L.P.

Microsoft and Windows are trademarks of Microsoft Corporation. UNIX and The Open Group are trademarks of The Open Group. All other product names mentioned herein may be trademarks or registered 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 publication is subject to change without notice and is provided "as is" without warranty of any kind. The entire risk arising out of the use of this information remains with recipient. In no event shall Compaq be liable for any direct, consequential, incidental, special, punitive, or other damages whatsoever (including without limitation, damages for loss of business profits, business interruption or loss of business information), even if Compaq has been advised of the possibility of such damages. The foregoing shall apply regardless of the negligence or other fault of either party and regardless of whether such liability sounds in contract, negligence, tort, or any other theory of legal liability, and notwithstanding any failure of essential purpose of any limited remedy.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Contents

About This Manual

1 Introduction

1.1	The TruCluster Server Product	1-1
1.2	Overview of the TruCluster Server Hardware Configuration ..	1-2
1.3	Memory Requirements	1-3
1.4	Minimum Disk Requirements	1-3
1.4.1	Disks Needed for Installation	1-3
1.4.1.1	Tru64 UNIX Operating System Disk	1-3
1.4.1.2	Clusterwide Disk(s)	1-4
1.4.1.3	Member Boot Disk	1-4
1.4.1.4	Quorum Disk	1-5
1.5	Generic Two-Node Cluster	1-5
1.6	Growing a Cluster from Minimum Storage to a NSPOF Cluster	1-7
1.6.1	Two-Node Clusters Using an UltraSCSI BA356 Storage Shelf and Minimum Disk Configurations	1-8
1.6.2	Two-Node Clusters Using UltraSCSI BA356 Storage Units with Increased Disk Configurations	1-10
1.6.3	Two-Node Configurations with UltraSCSI BA356 Storage Units and Dual SCSI Buses	1-12
1.6.4	Using Hardware RAID to Mirror the Clusterwide Root File System and Member System Boot Disks	1-13
1.6.5	Creating a NSPOF Cluster	1-15
1.7	Overview of Setting up the TruCluster Server Hardware Configuration	1-17

2 Hardware Requirements and Restrictions

2.1	TruCluster Server Member System Requirements	2-1
2.2	Memory Channel Restrictions	2-1
2.3	Fibre Channel Requirements and Restrictions	2-3
2.4	SCSI Bus Adapter Restrictions	2-6
2.4.1	KZPSA-BB SCSI Adapter Restrictions	2-6
2.4.2	KZPBA-CB SCSI Bus Adapter Restrictions	2-6
2.5	Disk Device Restrictions	2-7

2.6	RAID Array Controller Restrictions	2-7
2.7	SCSI Signal Converters	2-8
2.8	DS-DWZZH-03 and DS-DWZZH-05 UltraSCSI Hubs	2-9
2.9	SCSI Cables	2-9
2.10	SCSI Terminators and Trilink Connectors	2-11

3 Shared SCSI Bus Requirements and Configurations Using UltraSCSI Hardware

3.1	Shared SCSI Bus Configuration Requirements	3-2
3.2	SCSI Bus Performance	3-3
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	3-5
3.4	SCSI Bus Length	3-6
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	3-10
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	3-15
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	3-18
3.7.1.2	Preparing a Dual-Redundant HSZ70 or HSZ80 for a Shared SCSI Bus Using Multiple-Bus Failover	3-22

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 Jumpers	5-2
5.1.2	MC2 Jumpers	5-3
5.2	Installing the Memory Channel Adapter	5-5
5.3	Installing the MC2 Optical Converter in the Member System	5-6
5.4	Installing the Memory Channel Hub	5-6
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	5-10
5.6	Running Memory Channel Diagnostics	5-11

6 Using Fibre Channel Storage

6.1	Procedure for Installation Using Fibre Channel Disks	6-2
6.2	Fibre Channel Overview	6-4
6.2.1	Basic Fibre Channel Terminology	6-4
6.2.2	Fibre Channel Topologies	6-5
6.2.2.1	Point-to-Point	6-6

6.2.2.2	Fabric	6-6
6.2.2.3	Arbitrated Loop Topology	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	Zoning and Cascaded Switches	6-13
6.4.1	Zoning	6-13
6.4.2	Cascaded Switches	6-14
6.5	Installing and Configuring Fibre Channel Hardware	6-15
6.5.1	Installing and Setting Up the Fibre Channel Switch	6-15
6.5.1.1	Installing the Switch	6-16
6.5.1.2	Managing the Fibre Channel Switches	6-17
6.5.1.2.1	Using the Switch Front Panel	6-17
6.5.1.2.2	Setting the Ethernet IP Address and Subnet Mask from the Front Panel	6-18
6.5.1.2.3	Setting the DS-DSGGB-AA Ethernet IP Address and Subnet Mask from a PC or Terminal	6-20
6.5.1.2.4	Logging Into the Switch with a Telnet Connection	6-20
6.5.1.2.5	Setting the Switch Name via Telnet Session	6-21
6.5.2	Installing and Configuring the KGPSA PCI-to-Fibre Channel Adapter Module	6-22
6.5.2.1	Installing the KGPSA PCI-to-Fibre Channel Adapter Module	6-22
6.5.2.2	Setting the KGPSA-BC or KGPSA-CA to Run on a Fabric	6-23
6.5.2.3	Obtaining the Worldwide Names of KGPSA Adapters	6-25
6.5.3	Setting up the HSG80 Array Controller for Tru64 UNIX Installation	6-26
6.5.3.1	Obtaining the Worldwide Names of HSG80 Controller	6-31
6.6	Preparing to Install Tru64 UNIX and TruCluster Server on Fibre Channel Storage	6-33
6.6.1	Configuring the HSG80 Storagesets	6-33
6.6.2	Setting the Device Unit Number	6-40
6.6.3	Setting the bootdef_dev Console Environment Variable ...	6-46
6.7	Install the Base Operating System	6-48
6.8	Resetting the bootdef_dev Console Environment Variable ..	6-49
6.9	Determining /dev/disk/dskn to Use for a Cluster Installation .	6-51
6.10	Installing the TruCluster Server Software	6-53
6.11	Changing the HSG80 from Transparent to Multiple-Bus Failover Mode	6-54

6.12	Using the emx Manager to Display Fibre Channel Adapter Information	6-59
6.12.1	Using the emxmgr Utility to Display Fibre Channel Adapter Information	6-59
6.12.2	Using the emxmgr Utility Interactively	6-61
7	Preparing ATM Adapters	
7.1	ATM Overview	7-1
7.2	Installing ATM Adapters	7-3
7.3	Verifying ATM Fiber Optic Cable Connectivity	7-4
7.4	ATMworks Adapter LEDs	7-6
8	Configuring a Shared SCSI Bus for Tape Drive Use	
8.1	Preparing the TZ88 for Shared Bus Usage	8-1
8.1.1	Setting the TZ88N-VA SCSI ID	8-2
8.1.2	Cabling the TZ88N-VA	8-3
8.1.3	Setting the TZ88N-TA SCSI ID	8-4
8.1.4	Cabling the TZ88N-TA	8-4
8.2	Preparing the TZ89 for Shared SCSI Usage	8-5
8.2.1	Setting the DS-TZ89N-VW SCSI ID	8-5
8.2.2	Cabling the DS-TZ89N-VW Tape Drives	8-7
8.2.3	Setting the DS-TZ89N-TA SCSI ID	8-8
8.2.4	Cabling the DS-TZ89N-TA Tape Drives	8-8
8.3	Compaq 20/40 GB DLT Tape Drive	8-9
8.3.1	Setting the Compaq 20/40 GB DLT Tape Drive SCSI ID ..	8-9
8.3.2	Cabling the Compaq 20/40 GB DLT Tape Drive	8-10
8.4	Preparing the TZ885 for Shared SCSI Usage	8-13
8.4.1	Setting the TZ885 SCSI ID	8-13
8.4.2	Cabling the TZ885 Tape Drive	8-13
8.5	Preparing the TZ887 for Shared SCSI Bus Usage	8-15
8.5.1	Setting the TZ887 SCSI ID	8-15
8.5.2	Cabling the TZ887 Tape Drive	8-16
8.6	Preparing the TL891 and TL892 DLT MiniLibraries for Shared SCSI Usage	8-18
8.6.1	Setting the TL891 or TL892 SCSI ID	8-18
8.6.2	Cabling the TL891 or TL892 MiniLibraries	8-20
8.7	Preparing the TL890 DLT MiniLibrary Expansion Unit	8-23
8.7.1	TL890 DLT MiniLibrary Expansion Unit Hardware	8-24
8.7.2	Preparing the DLT MiniLibraries for Shared SCSI Bus Usage	8-24

8.7.2.1	Cabling the DLT MiniLibraries	8-24
8.7.2.2	Configuring a Base Module as a Slave	8-26
8.7.2.3	Powering Up the DLT MiniLibrary	8-28
8.7.2.4	Setting the TL890/TL891/TL892 SCSI ID	8-28
8.8	Preparing the TL894 DLT Automated Tape Library for Shared SCSI Bus Usage	8-30
8.8.1	TL894 Robotic Controller Required Firmware	8-30
8.8.2	Setting TL894 Robotics Controller and Tape Drive SCSI IDs	8-30
8.8.3	TL894 Tape Library Internal Cabling	8-33
8.8.4	Connecting the TL894 Tape Library to the Shared SCSI Bus	8-34
8.9	Preparing the TL895 DLT Automated Tape Library for Shared SCSI Bus Usage	8-36
8.9.1	TL895 Robotic Controller Required Firmware	8-37
8.9.2	Setting the TL895 Tape Library SCSI IDs	8-37
8.9.3	TL895 Tape Library Internal Cabling	8-38
8.9.4	Upgrading a TL895	8-40
8.9.5	Connecting the TL895 Tape Library to the Shared SCSI Bus	8-40
8.10	Preparing the TL893 and TL896 Automated Tape Libraries for Shared SCSI Bus Usage	8-40
8.10.1	Communications with the Host Computer	8-42
8.10.2	MUC Switch Functions	8-42
8.10.3	Setting the MUC SCSI ID	8-43
8.10.4	Tape Drive SCSI IDs	8-43
8.10.5	TL893 and TL896 Automated Tape Library Internal Cabling	8-44
8.10.6	Connecting the TL893 and TL896 Automated Tape Libraries to the Shared SCSI Bus	8-47
8.11	Preparing the TL881 and TL891 DLT MiniLibraries for Shared Bus Usage	8-48
8.11.1	TL881 and TL891 DLT MiniLibraries Overview	8-48
8.11.1.1	TL881 and TL891 DLT MiniLibrary Tabletop Model ..	8-49
8.11.1.2	TL881 and TL891 MiniLibrary Rackmount Components	8-49
8.11.1.3	TL881 and TL891 Rackmount Scalability	8-50
8.11.1.4	DLT MiniLibrary Part Numbers	8-51
8.11.2	Preparing a TL881 or TL891 MiniLibrary for Shared SCSI Bus Use	8-52
8.11.2.1	Preparing a Tabletop Model or Base Unit for Standalone Shared SCSI Bus Usage	8-52

8.11.2.1.1	Setting the Standalone MiniLibrary Tape Drive SCSI ID	8-53
8.11.2.1.2	Cabling the TL881 or TL891 DLT MiniLibrary	8-54
8.11.2.2	Preparing a TL881 or TL891 Rackmount MiniLibrary for Shared SCSI Bus Usage	8-58
8.11.2.2.1	Cabling the Rackmount TL881 or TL891 DLT MiniLibrary	8-58
8.11.2.2.2	Configuring a Base Unit as a Slave to the Expansion Unit	8-61
8.11.2.2.3	Powering Up the TL881/TL891 DLT MiniLibrary .	8-62
8.11.2.2.4	Setting the SCSI IDs for a Rackmount TL881 or TL891 DLT MiniLibrary	8-63
8.12	Compaq ESL9326D Enterprise Library	8-64
8.12.1	General Overview	8-64
8.12.2	ESL9326D Enterprise Library Overview	8-65
8.12.3	Preparing the ESL9326D Enterprise Library for Shared SCSI Bus Usage	8-65
8.12.3.1	ESL9326D Enterprise Library Robotic and Tape Drive Required Firmware	8-66
8.12.3.2	Library Electronics and Tape Drive SCSI IDs	8-66
8.12.3.3	ESL9326D Enterprise Library Internal Cabling	8-66
8.12.3.4	Connecting the ESL9326D Enterprise Library to the Shared SCSI Bus	8-68

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

9.1	Using SCSI Bus Signal Converters	9-2
9.1.1	Types of SCSI Bus Signal Converters	9-2
9.1.2	Using the SCSI Bus Signal Converters	9-3
9.1.2.1	DWZZA and DWZZB Signal Converter Termination ..	9-3
9.1.2.2	DS-BA35X-DA Termination	9-4
9.2	Terminating the Shared SCSI Bus	9-5
9.3	Overview of Disk Storage Shelves	9-8
9.3.1	BA350 Storage Shelf	9-9
9.3.2	BA356 Storage Shelf	9-10
9.3.2.1	Non-UltraSCSI BA356 Storage Shelf	9-10
9.3.2.2	UltraSCSI BA356 Storage Shelf	9-13
9.4	Preparing the Storage for Configurations Using External Termination	9-14

9.4.1	Preparing BA350, BA356, and UltraSCSI BA356 Storage Shelves for an Externally Terminated TruCluster Server Configuration	9-15
9.4.1.1	Preparing a BA350 Storage Shelf for Shared SCSI Usage	9-15
9.4.1.2	Preparing a BA356 Storage Shelf for Shared SCSI Usage	9-16
9.4.1.3	Preparing an UltraSCSI BA356 Storage Shelf for a TruCluster Configuration	9-17
9.4.2	Connecting Storage Shelves Together	9-17
9.4.2.1	Connecting a BA350 and a BA356 for Shared SCSI Bus Usage	9-18
9.4.2.2	Connecting Two BA356s for Shared SCSI Bus Usage ..	9-20
9.4.2.3	Connecting Two UltraSCSI BA356s for Shared SCSI Bus Usage	9-21
9.4.3	Cabling a Non-UltraSCSI RAID Array Controller to an Externally Terminated Shared SCSI Bus	9-24
9.4.3.1	Cabling an HSZ40 or HSZ50 in a Cluster Using External Termination	9-25
9.4.3.2	Cabling an HSZ20 in a Cluster using External Termination	9-27
9.4.4	Cabling an HSZ40 or HSZ50 RAID Array Controller in a Radial Configuration with an UltraSCSI Hub	9-28

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

10.1	TruCluster Server Hardware Installation Using PCI SCSI Adapters	10-1
10.1.1	Radial Installation of a KZPSA-BB or KZPBA-CB Using Internal Termination	10-2
10.1.2	Installing a KZPSA-BB or KZPBA-CB Using External Termination	10-6
10.1.3	Displaying KZPSA-BB and KZPBA-CB Adapters with the show Console Commands	10-9
10.1.4	Displaying Console Environment Variables and Setting the KZPSA-BB and KZPBA-CB SCSI ID	10-13
10.1.4.1	Displaying KZPSA-BB and KZPBA-CB pk* or isp* Console Environment Variables	10-13
10.1.4.2	Setting the KZPBA-CB SCSI ID	10-16
10.1.4.3	Setting KZPSA-BB SCSI Bus ID, Bus Speed, and Termination Power	10-17
10.1.4.4	KZPSA-BB and KZPBA-CB Termination Resistors	10-18

10.1.4.5	Updating the KZPSA-BB Adapter Firmware	10-18
----------	--	-------

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	6-29
6-2	Setting up the Mirrorset	6-34
6-3	Using the wwidmgr quickset Command to Set Device Unit Number	6-43
6-4	Sample Fibre Channel Device Names	6-45
10-1	Displaying Configuration on an AlphaServer 4100	10-9
10-2	Displaying Devices on an AlphaServer 4100	10-10
10-3	Displaying Configuration on an AlphaServer 8200	10-11
10-4	Displaying Devices on an AlphaServer 8200	10-12
10-5	Displaying the pk* Console Environment Variables on an AlphaServer 4100 System	10-13
10-6	Displaying Console Variables for a KZPBA-CB on an AlphaServer 8x00 System	10-15
10-7	Displaying Console Variables for a KZPSA-BB on an AlphaServer 8x00 System	10-15
10-8	Setting the KZPBA-CB SCSI Bus ID	10-16
10-9	Setting KZPSA-BB SCSI Bus ID and Speed	10-17

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	1-11
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
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	3-20
3-6	Shared SCSI Bus with HSZ80 Configured for Transparent Failover	3-21
3-7	TruCluster Server Configuration with HSZ70 in Multiple-Bus Failover Mode	3-24
3-8	TruCluster Server Configuration with HSZ80 in Multiple-Bus Failover Mode	3-25
4-1	KZPBA-CB Termination Resistors	4-18
5-1	Connecting Memory Channel Adapters to Hubs	5-9
6-1	Point-to-Point Topology	6-6
6-2	Fabric Topology	6-7
6-3	Arbitrated Loop Topology	6-8
6-4	Fibre Channel Single Switch Transparent Failover Configuration	6-9
6-5	Multiple-Bus NSPOF Configuration Number 1	6-11
6-6	Multiple-Bus NSPOF Configuration Number 2	6-12
6-7	Multiple-Bus NSPOF Configuration Number 3	6-13
6-8	A Simple Zoned Configuration	6-14
7-1	Emulated LAN Over an ATM Network	7-3
8-1	TZ88N-VA SCSI ID Switches	8-2
8-2	Shared SCSI Buses with SBB Tape Drives	8-4
8-3	DS-TZ89N-VW SCSI ID Switches	8-6
8-4	Compaq 20/40 GB DLT Tape Drive Rear Panel	8-10
8-5	Cabling a Shared SCSI Bus with a Compaq 20/40 GB DLT Tape Drive	8-12

8-6	Cabling a Shared SCSI Bus with a TZ885	8-15
8-7	TZ887 DLT MiniLibrary Rear Panel	8-16
8-8	Cabling a Shared SCSI Bus with a TZ887	8-17
8-9	TruCluster Server Cluster with a TL892 on Two Shared SCSI Buses	8-23
8-10	TL890 and TL892 DLT MiniLibraries on Shared SCSI Buses ..	8-26
8-11	TL894 Tape Library Four-Bus Configuration	8-33
8-12	Shared SCSI Buses with TL894 in Two-Bus Mode	8-35
8-13	TL895 Tape Library Internal Cabling	8-39
8-14	TL893 Three-Bus Configuration	8-45
8-15	TL896 Six-Bus Configuration	8-46
8-16	Shared SCSI Buses with TL896 in Three-Bus Mode	8-48
8-17	TL891 Standalone Cluster Configuration	8-57
8-18	TL881 DLT MiniLibrary Rackmount Configuration	8-60
8-19	ESL9326D Internal Cabling	8-67
9-1	Standalone SCSI Signal Converter	9-4
9-2	SBB SCSI Signal Converter	9-4
9-3	DS-BA35X-DA Personality Module Switches	9-5
9-4	BN21W-0B Y Cable	9-7
9-5	HD68 Trilink Connector (H885-AA)	9-8
9-6	BA350 Internal SCSI Bus	9-10
9-7	BA356 Internal SCSI Bus	9-12
9-8	BA356 Jumper and Terminator Module Identification Pins ...	9-13
9-9	BA350 and BA356 Cabled for Shared SCSI Bus Usage	9-19
9-10	Two BA356s Cabled for Shared SCSI Bus Usage	9-21
9-11	Two UltraSCSI BA356s Cabled for Shared SCSI Bus Usage ..	9-23
9-12	Externally Terminated Shared SCSI Bus with Mid-Bus HSZ50 RAID Array Controllers	9-26
9-13	Externally Terminated Shared SCSI Bus with HSZ50 RAID Array Controllers at Bus End	9-27
9-14	TruCluster Server Cluster Using DS-DWZZH-03, SCSI Adapter with Terminators Installed, and HSZ50	9-30
9-15	TruCluster Server Cluster Using KZPSA-BB SCSI Adapters, a DS-DWZZH-05 UltraSCSI Hub, and an HSZ50 RAID Array Controller	9-31
10-1	KZPSA-BB Termination Resistors	10-18

Tables

2-1	AlphaServer Systems Supported for Fibre Channel	2-4
2-2	RAID Controller SCSI IDs	2-7
2-3	Supported SCSI Cables	2-10

2-4	Supported SCSI Terminators and Trilink Connectors	2-11
3-1	SCSI Bus Speeds	3-5
3-2	SCSI Bus Segment Length	3-7
3-3	DS-DWZZH UltraSCSI Hub Maximum Configurations	3-11
3-4	Hardware Components Used in Configuration Shown in Figure 3-5 Through Figure 3-8	3-21
4-1	Planning Your Configuration	4-3
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 Jumper Configuration	5-2
5-2	MC2 Jumper Configuration	5-3
5-3	MC2 Linecard Jumper Configurations	5-5
6-1	Telnet Session Default User Names for Fibre Channel Switches	6-21
6-2	Converting StorageSet Unit Numbers to Disk Names	6-40
7-1	ATMworks Adapter LEDs	7-6
8-1	TZ88N-VA Switch Settings	8-3
8-2	DS-TZ89N-VW Switch Settings	8-6
8-3	Hardware Components Used to Create the Configuration Shown in Figure 8 — 5	8-12
8-4	TL894 Default SCSI ID Settings	8-30
8-5	TL895 Default SCSI ID Settings	8-37
8-6	MUC Switch Functions	8-43
8-7	MUC SCSI ID Selection	8-43
8-8	TL893 Default SCSI IDs	8-44
8-9	TL896 Default SCSI IDs	8-44
8-10	TL881 and TL891 MiniLibrary Performance and Capacity Comparison	8-51
8-11	DLT MiniLibrary Part Numbers	8-51
8-12	Hardware Components Used to Create the Configuration Shown in Figure 8-17	8-57
8-13	Hardware Components Used to Create the Configuration Shown in Figure 8-18	8-61
8-14	Shared SCSI Bus Cable and Terminator Connections for the ESL9326D Enterprise Library	8-68
9-1	Hardware Components Used for Configuration Shown in Figure 8-9 and Figure 8-10	9-20
9-2	Hardware Components Used for Configuration Shown in Figure 9-11	9-24
9-3	Hardware Components Used for Configuration Shown in Figure 8-12 and Figure 8-13	9-27

9-4	Hardware Components Used in Configuration Shown in Figure 9-14	9-30
10-1	Configuring TruCluster Server Hardware for Use with a PCI SCSI Adapter	10-2
10-2	Installing the KZPSA-BB or KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub	10-4
10-3	Installing a KZPSA-BB or KZPBA-CB for use with External Termination	10-7
A-1	Converting StorageSet 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 needed to maintain your hardware, operating system, and network.

Organization

This manual contains ten chapters and an index. The organization of this manual has been restructured to provide a more streamlined manual. Those chapters containing information on SCSI bus requirements and configuration, and configuring hardware have been split up into two sets of two chapters each. One set covers the UltraSCSI hardware and is geared towards radial configurations. The other set covers configurations using either external termination or radial connection to non-UltraSCSI devices. A brief description of the contents 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, 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 newest RAID array controllers (HSZ70 and HSZ80).
- Chapter 5* Describes how to set up the Memory Channel cluster interconnect.
- Chapter 6* Provides an overview of Fibre Channel and describes how to set up Fibre Channel hardware.
- Chapter 7* Provides information on the use of, and installation of, Asynchronous Transfer Mode (ATM) hardware.

- Chapter 8* Describes how to configure a shared SCSI bus for tape drive, tape loader, or tape library usage.
- Chapter 9* Contains information about setting up a shared SCSI bus, SCSI bus requirements, and how to connect storage to a shared SCSI bus using external termination or radial connections to non-UltraSCSI devices.
- Chapter 10* Describes how to prepare systems for a TruCluster Server configuration, and how to connect host bus adapters to shared storage using external termination or radial connection to non-UltraSCSI devices.

Related Documents

Users of the TruCluster Server product can consult the following manuals for assistance in cluster installation, administration, and programming tasks:

- *TruCluster Server Software Product Description (SPD)* — The comprehensive description of the TruCluster Server Version 5.0A product. You can find the latest version of the SPD and other TruCluster Server documentation at the following URL:
http://www.unix.digital.com/faqs/publications/pub_page/cluster_list.html
- *Release Notes* — Provides important information about TruCluster Server Version 5.0A.
- *Technical Overview* — Provides an overview of the TruCluster Server technology.
- *Software Installation* — Describes how to install the TruCluster Server product.
- *Cluster Administration* — Describes cluster-specific administration tasks.
- *Highly Available Applications* — Describes how to deploy applications on a TruCluster Server cluster.

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

For information about setting up a RAID subsystem, see the following documentation 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*
- *Wwidmgr User's Manual*

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

- *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*

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, you should have available the following manuals from the Tru64 UNIX documentation set:

- *Installation Guide*
- *Release Notes*
- *System Administration*
- *Network Administration*

You should also have 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 (DS20 and ES40)
- DECEvent™ (AlphaServers other than the DS20 and ES40)
- Logical Storage Manager (LSM)
- NetWorker
- Advanced File System (AdvFS) Utilities
- Performance Manager

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`

- Mail:
Compaq Computer Corporation

UBPG Publications Manager
ZKO3-3/Y32
110 Spit Brook Road
Nashua, NH 03062-2698

A Reader's Comment form is located in the back of each printed manual. The form is postage paid if you mail it in the United States.

Please include the following information along with your comments:

- The full title of the book and the order number. (The order number is printed on the title page of this book and on its back cover.)
- 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.

Introduction

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

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

1.1 The TruCluster Server Product

TruCluster Server, the newest addition to the Compaq Tru64 UNIX TruCluster Software products family, 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 Overview of the TruCluster Server Hardware Configuration

A TruCluster Server hardware configuration consists of a number of highly specific hardware components:

- TruCluster Server currently supports from one to eight member systems.
- There must be sufficient internal and external SCSI controllers, Fibre Channel host bus adapters, and disks to provide sufficient storage for the applications.
- The clusterwide root (/), /usr, and /var file systems should be on a shared SCSI bus. We recommend placing all member system boot disks on a shared SCSI bus. If you have a quorum disk, it must be on a shared SCSI bus.

Note

The clusterwide root (/), /usr, and /var file systems, the member system boot disks, and the quorum disk may be located behind a RAID array controller, including the HSG80 controller (Fibre Channel).

- You need to allocate a number of Internet Protocol (IP) addresses from one IP subnet to allow client access to the cluster. The IP subnet has to be visible to the clients directly or through routers. The minimum number of allocated addresses is equal to the number of cluster member systems plus one (for the cluster alias), depending on the type of cluster alias configuration.

For client access, TruCluster Server allows you to configure any number of monitored network adapters (using a redundant array of independent network adapters (NetRAIN) and Network Interface Failure Finder (NIFF) facilities of the Tru64 UNIX operating system).

- TruCluster Server requires at least one peripheral component interconnect (PCI) Memory Channel adapter on each system. The Memory Channel adapters comprise the cluster interconnect for TruCluster Server, providing host-to-host communications. For a cluster with two systems, a Memory Channel hub is optional; the Memory Channel adapters can be connected with a cable.

If there are more than two systems in the cluster, a Memory Channel hub is required. The Memory Channel hub is a PC-class enclosure that

contains up to eight linecards. The Memory Channel adapter in each system in the cluster is connected to the Memory Channel hub.

One or two Memory Channel adapters can be used with TruCluster Server. When dual Memory Channel adapters are installed, if the Memory Channel adapter being used for cluster communication fails, the communication will fail over to the other Memory Channel.

1.3 Memory Requirements

Cluster members require a minimum of 128 MB of memory.

1.4 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 TruCluster Server *Software Installation* manual.

1.4.1 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 disk(s) are either private disk(s) on the system that will become the first cluster member, or disk(s) 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 1.4.1.4, and for a more detailed discussion of the quorum disk, see the TruCluster Server *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.

1.4.1.1 Tru64 UNIX Operating System Disk

The Tru64 UNIX operating system is installed using AdvFS file systems on one or more disks on 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, 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 TruCluster Server *Cluster Administration* manual for more information.

1.4.1.2 Clusterwide Disk(s)

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 used 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.

1.4.1.3 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 TruCluster Server *Cluster Administration* manual for more information.

1.4.1.4 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 TruCluster Server *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 second point of failure).
- You cannot use the Logical Storage Manager (LSM) on the quorum disk.

1.5 Generic Two-Node Cluster

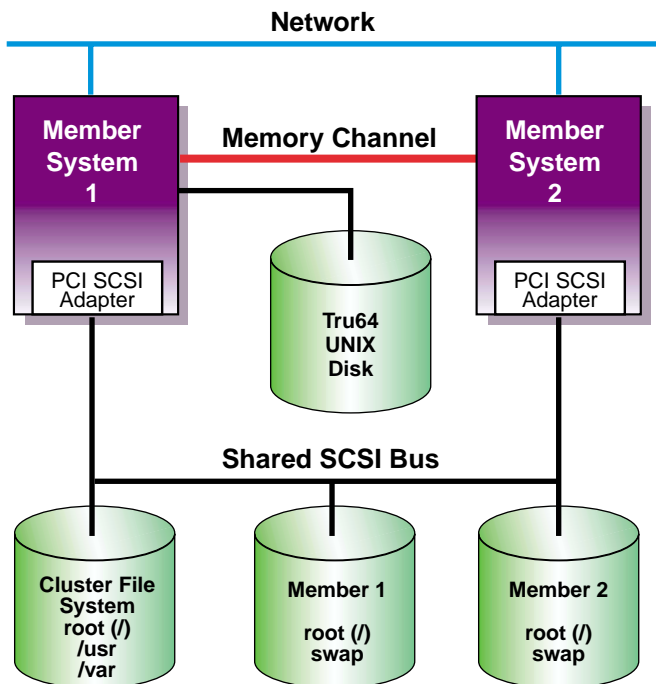
This section describes a generic two-node cluster with the minimum disk layout of four disks. Note that additional disks may be needed for highly available applications. In this section, and the following sections, the type of PCI SCSI bus adapter is not significant. Also, although an important consideration, SCSI bus cabling, including Y cables or trilink connectors, termination, and the use of UltraSCSI hubs is 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 TruCluster Server *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

1.6 Growing a Cluster from Minimum Storage to a 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:

- A cluster with minimum storage for highly available applications (Section 1.6.1).
- A cluster with more storage, but the single SCSI bus is a single point of failure (Section 1.6.2).
- Adding a second SCSI bus allows the use of LSM to mirror the `/usr` and `/var` file systems and data disks. However, as LSM cannot mirror the root (`/`), member system boot, swap, or quorum disks, so full redundancy is not achieved (Section 1.6.3).

- Using a 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 a NSPOF cluster (Section 1.6.4).
- 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 a NSPOF cluster (Section 1.6.5).

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

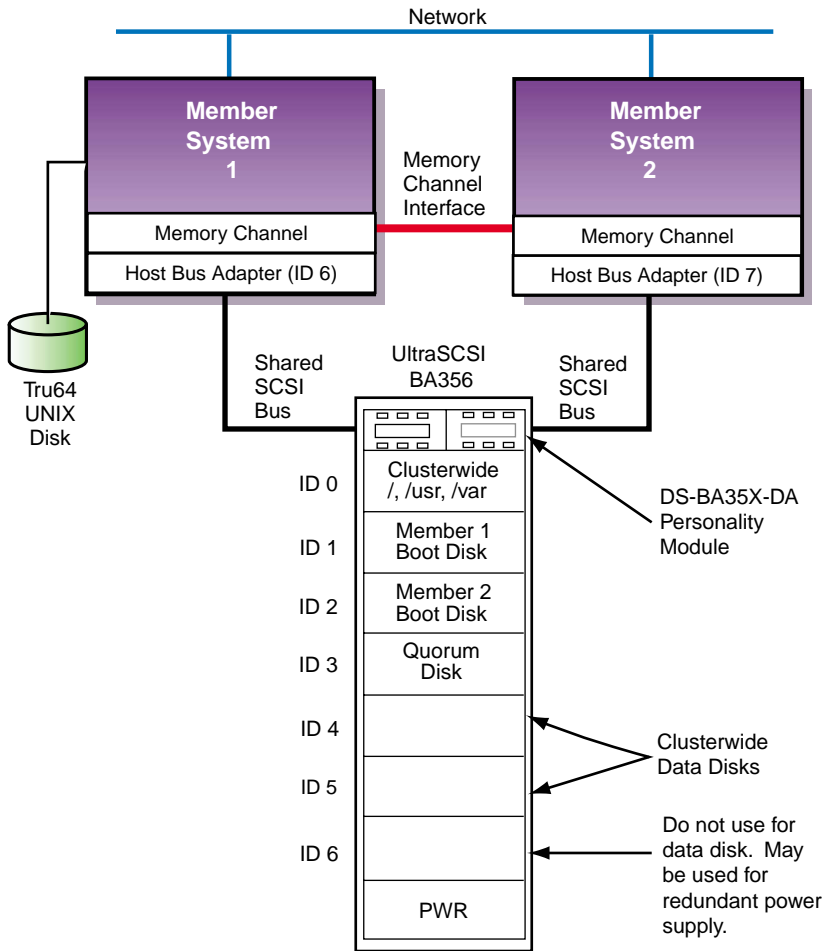
This section takes the generic illustrations of our cluster example one step further by depicting the required storage in storage shelves. The storage shelves could be BA350, BA356 (non-UltraSCSI), or UltraSCSI BA356s. The BA350 is the oldest model, and can only respond to SCSI IDs 0-6. The non-Ultra 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.

Note

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

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 5.0A required disks.

In this configuration, because of the TruCluster Server Version 5.0A disk requirements, there will only be two disks 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.

Note that with the use of the cluster file system (See the TruCluster Server *Cluster Administration* manual for a discussion of the cluster file system), the clusterwide root (/), /usr, and /var file systems could be physically placed on a private bus of either of the member systems. But, if that member system was not available, the other member system(s) would not have access to the clusterwide file systems. Therefore, placing the clusterwide root (/), /usr, and /var file systems on a private bus is not recommended.

Likewise, the quorum disk could be placed on the local bus of either of the member systems. If that member was not available, quorum could never be reached in a two-node cluster. Placing the quorum disk on the local bus of a member system is not recommended as it creates a single point of failure.

The individual member boot and swap partitions could also be placed on a local bus of either of the member systems. If the boot disk for member system 1 was on a SCSI bus internal to member 1, and the system was unavailable due to a boot disk problem, other systems in the cluster could not 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 5.0A configurations require one or more disks to hold the Tru64 UNIX operating system. The disk(s) are either private disk(s) on the system that will become the first cluster member, or disk(s) on a shared bus that the system can access.

We recommend that you place the /usr, /var, member boot disks, and quorum disk on a shared SCSI bus 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 TruCluster Server *Cluster Administration* manual for more information.

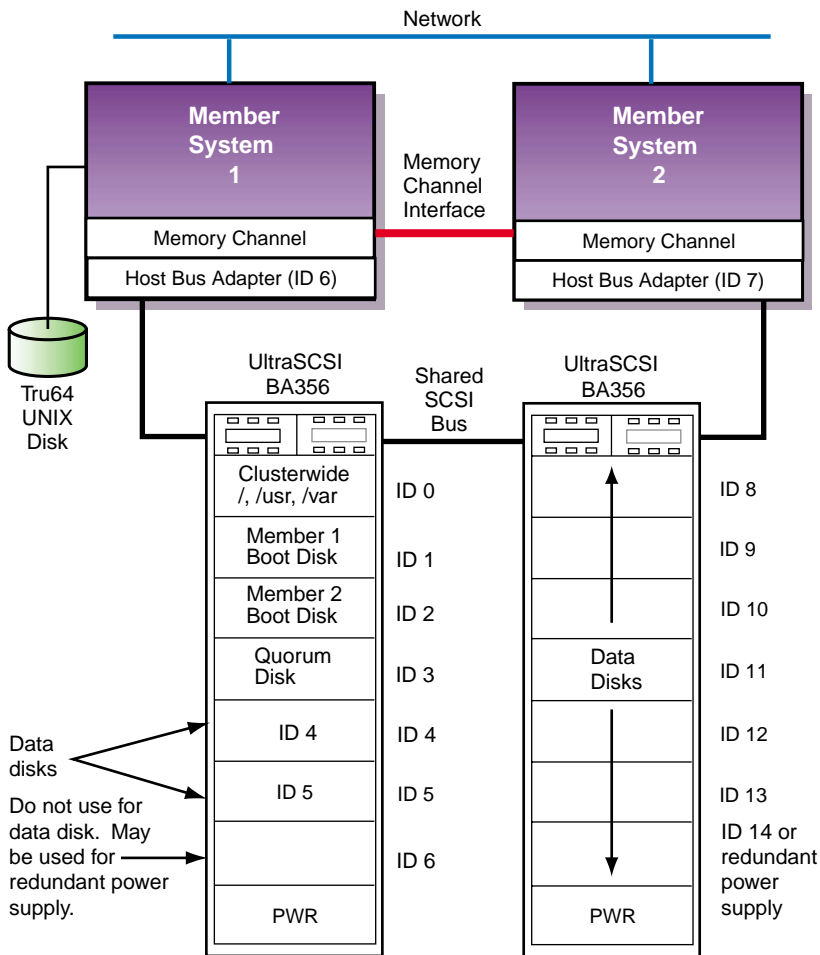
1.6.2 Two-Node Clusters Using UltraSCSI BA356 Storage Units with Increased Disk Configurations

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 would allow the use of the Logical Storage Manager (LSM) to mirror the /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.

1.6.3 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 the Logical Storage Manager (LSM) to mirror data disks, and the clusterwide `/usr` and `/var` file systems across SCSI buses.

Note

You cannot use LSM to mirror the clusterwide root (`/`), member system boot, swap, 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 `/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-1593U-AI


By using LSM to mirror the /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 clusterwide root (/), quorum, or the member boot disks, we do not have a no-single-point-of-failure (NSPOF) cluster.

1.6.4 Using Hardware RAID to Mirror the Clusterwide Root File System and Member System Boot Disks

You can use hardware RAID with any of the supported RAID array controllers to mirror the clusterwide root (/), quorum, and member boot disks. Figure 1–6 shows a cluster configuration using an HSZ70 RAID array controller. An HSZ40, HSZ50, HSZ80, or HSG80 could 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, 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.

Note that 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.

1.6.5 Creating a NSPOF Cluster

To create a no-single-point-of-failure (NSPOF) cluster:

- 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, and HSG80)
- Install a second Memory Channel interface for redundancy
- Install redundant power supplies
- Install redundant networks
- Connect the systems and storage to an uninterruptable power supply (UPS)

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

Notes


Only the HSZ70, HSZ80, 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 HSG80 and ACS V8.5.

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.

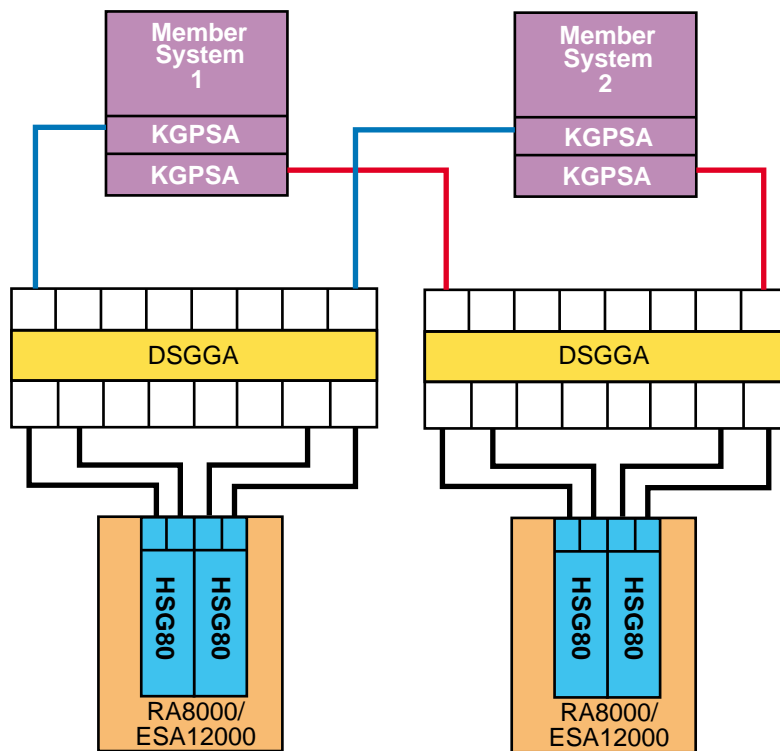
Figure 1–7: NSPOF Cluster using HSZ70s in Multiple-Bus Failover Mode



ZK-1594U-AI

Figure 1–8 shows a cluster configuration with dual-shared Fibre Channel SCSI 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-1533U-AI

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, and Chapter 10).
2. Draw a diagram of your configuration.
3. Compare your diagram with the examples in Chapter 3, Chapter 6, and Chapter 9.
4. Identify all devices, cables, SCSI adapters, and so forth. Use the diagram you just constructed.
5. Prepare the shared storage by installing disks and configuring any RAID controller subsystems (See Chapter 3, Chapter 6, and Chapter 9 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 9).
7. Connect storage to the shared SCSI buses. Terminate each bus. Use Y cables or trilink connectors where necessary (see Chapter 3 and Chapter 9).

For a Fibre Channel configuration, connect the HSG80 controllers to the switches. You want the 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 (see Chapter 7).
 - SCSI bus adapters. Ensure that adapter terminators are set correctly. Connect the systems to the shared SCSI bus (see Chapter 4 or Chapter 10).
 - 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 10).
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 10).

Hardware Requirements and Restrictions

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

See the TruCluster Server *Software Product Description* (SPD) for the latest information about supported hardware.

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 <http://www.compaq.com/support/>. Select Alpha Systems from the downloadable drivers & utilities menu. Check the release notes for the appropriate system type to determine the firmware version required.

- TruCluster Server Version 5.0A 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.
- 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 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, and ES40 systems only support MC2 hardware.
- 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 use the MC API library functions in a 2-node TruCluster Server configuration, you cannot use virtual hub mode, you must use a Memory Channel hub and standard hub mode.
- If you have redundant MC2 modules on a system jumpered for 512 MB, you cannot have two MC2 modules on the same PCI bus.
- The MC1 adapter cannot be cabled to a MC2 adapter.
 - Do not use the BC12N link cable with the CCMAB MC2 adapter.
 - Do not use the BN39B link cable with the CCMAA MC1 adapter.
- 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.
- A Memory Channel interconnect can use either virtual hub mode (two member systems connected without a Memory Channel hub) or standard 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 and MC2 BN39B link cable is 10 meters.
- Always check 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 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), the bottom four PCI slots.
- 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.
- 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.

2.3 Fibre Channel Requirements and Restrictions

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

Table 2–1: AlphaServer Systems Supported for Fibre Channel

AlphaServer	Number of KGPSA-BC Adapters Supported
AlphaServer 800	2
AlphaServer 1200	4
AlphaServer 4000, 4000A, or 4100	4
Compaq AlphaServer DS10	1
Compaq AlphaServer DS20 and DS20E	2
Compaq AlphaServer ES40	4
AlphaServer 8200 or 8400 ^a	32 (2 per DWLPB for throughput, 4 per DWLPB for connectivity)
Compaq AlphaServer GS60, GS60E, and GS140 ^a	32 (2 per DWLPB for throughput, 4 per DWLPB for connectivity)

^a 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 following requirements and restrictions apply to the use of Fibre Channel with the TruCluster Server:

- The HSG80 requires Array Control Software (ACS) Version 8.5.
- A maximum of four member systems is supported.
- The only supported Fibre Channel adapters are the KGPSA-BC and KGPSA-CA PCI-to-Fibre Channel host bus adapters.
- The only Fibre Channel switches supported are the 8/16 Port DSGGA or DSGGB Fibre Channel switches.
- The Fibre Channel switches support both shortwave (GBIC-SW) and longwave (GBIC-LW) Giga Bit Interface Converter (GBIC) modules. The GBIC-SW module supports 50-micron, multimode fibre cables with the standard subscriber connector (SC connector) in lengths up to 500 meters. The GBIC-LW supports 9-micron, single-mode fibre cables with the SC connector in lengths up to 10 kilometers.

The KGPSA-BC/CA PCI-to-Fibre Channel host bus adapters and the HSG80 RAID controller support the 50-micron Gigabit Link Module (GLM) for fibre connections. Therefore, only the 50-micron multimode fibre optical cable is supported between the KGPSA and switch and the switch and HSG80 for cluster configurations. You must install GBIC-SW GBICs in the Fibre Channel switches for communication between the switches and KGPSA or HSG80.

- A maximum of three cascaded switches is supported, with a maximum of two hops between switches. The maximum hop length is 10 km longwave single-mode or 500 meters via shortwave multimode Fibre Channel cable.

- 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.
- Only disk devices attached to the HSG80 Fibre Channel to Six Channel UltraSCSI Array controller are supported with the TruCluster Server product.
- No tape devices are supported.
- Tru64 UNIX Version 5.0A limits the number of Fibre Channel targets to 126.
- Tru64 UNIX Version 5.0A allows up to 255 LUNs per target.
- The HSG80 supports transparent and multiple-bus failover mode when used in a TruCluster Server Version 5.0A configuration. Multiple-bus failover is recommended for high availability in a cluster.
- A storage array with dual-redundant HSG80 controllers in transparent mode failover is two targets and consumes four ports on a switch.
- A storage array with dual-redundant HSG80 controllers in multiple-bus failover is four targets and consumes 4 ports on a switch.
- Each KGPSA is one target.
- The HSG80 documentation refers to the controllers as Controllers A (top) and B (bottom). Each controller provides two ports (left and right). (The 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 automatically fail over to the bottom controller (and vice versa).

- In transparent failover mode, you can configure which controller presents each HSG80 storage element (unit) to the cluster. Ordinarily, the left port of either controller serves the units designated D0 through D99, and the right port serves those designated D100 through D199.
- In multiple-bus failover mode, 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 unit(s) from one controller to the other controller.

2.4 SCSI Bus Adapter Restrictions

To connect a member system to a shared SCSI bus, you must install a SCSI 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 SCSI adapter restrictions in more detail.

2.4.1 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 support the variable), 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.

2.4.2 KZPBA-CB SCSI Bus Adapter Restrictions

KZPBA-CB UltraSCSI adapters have the following restrictions:

- Each system supporting the KZPBA-CB UltraSCSI host adapter limits the number of adapters that may be installed. The maximum number of KZPBA-CB UltraSCSI host adapters supported with TruCluster Server follow:
 - AlphaServer 800: 2
 - AlphaServer 1000A and 1200: 4
 - AlphaServer 4000: 8; only one KZPBA-CB is supported in IOD0 (PCI0).
 - AlphaServer 4100: 5; only one KZPBA-CB is supported in IOD0 (PCI0).

- AlphaServer 8200, 8400, GS60, GS60E, GS140: 32
The KZPBA-CB is supported on the DWLPB only; it is not supported on the DWLPA module.
- AlphaServer DS10: 2
- AlphaServer DS20/DS20E: 4
- AlphaServer ES40: 5
- 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, 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* (AA-R5XWD-TE) for more information.

2.5 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.

2.6 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 controllers can be configured with the number of SCSI IDs as shown in Table 2–2.

Table 2–2: RAID Controller SCSI IDs

RAID Controller	Number of SCSI IDs Supported
HSZ20	4
HSZ40	4
HSZ50	4
HSZ70	8

Table 2–2: RAID Controller SCSI IDs (cont.)

RAID Controller	Number of SCSI IDs Supported
HSZ80	15
HSG80	N/A

2.7 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

The UltraSCSI hubs could probably be listed here as they contain a DOC (DWZZA on a chip) chip, but they are covered separately in Section 2.8.

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 of CX444xxxxx to 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

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

The DS-DWZZH-03 and DS-DWZZH-05 series UltraSCSI hubs are the only hubs 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 as they are used differently in cluster configurations they will be 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 which 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 in length. This allows a total separation of nearly 50 meters 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.

2.9 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 2–3 describes each supported cable and the context in which you would use the cable. Note that there are cables with the Compaq 6-3 part number that are not listed, but are equivalent to the cables listed.

Table 2–3: 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.
BN38C or BN38D	One HD68, one VHDCI	HD68 to VHDCI	Connects a KZPBA-CB or KZPSA-BB to a port on an UltraSCSI hub.
BN37A	Two VHDCI	VHDCI to VHDCI	Connects two VHDCI trilinks to each other or an UltraSCSI hub to a trilink on an HSZ70 or HSZ80
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

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

Always check 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.

2.10 SCSI Terminators and Trilink Connectors

Table 2–4 describes the supported trilink connectors and SCSI terminators and the context in which you would use them.

Table 2–4: 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	High	68-pin	Terminates an H885-AA trilink connector or BN21W-0B Y cable.
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.

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.

Shared SCSI Bus Requirements and Configurations Using UltraSCSI Hardware

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

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 enabled for simultaneous failover
 - Dual-redundant HSZ70 or HSZ80 RAID array controllers enabled for multiple-bus failover
- Provides diagrams of TruCluster Server storage configurations using UltraSCSI hardware 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 covered in Chapter 9 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 9 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.

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)
- Configuring UltraSCSI hubs with RAID array controllers (Section 3.7)

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 9.1 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 tralink connectors and Y cables to connect devices to a shared bus, so you can disconnect the devices without affecting bus termination. See Section 9.2 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), or RA8000 or ESA12000 storage arrays (HSZ80 and HSG80). Older, non-UltraSCSI BA356 shelves are supported with UltraSCSI host adapters and host RAID controllers as long as they contain no UltraSCSI disks.
- UltraSCSI drives and fast wide drives can be mixed together in an UltraSCSI BA356 shelf (see Chapter 9).
- 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 9).
- 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 could assign the drive a SCSI ID greater than 7. It 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 9).
- 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 9).

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 connected to it operate. Specifically, bus performance is influenced by the following factors:

- Transmission method
- Data path
- Bus speed

The following sections describe these factors.

3.2.1 SCSI Bus Versus SCSI Bus Segments

An UltraSCSI bus may be comprised of multiple UltraSCSI bus segments. Each UltraSCSI bus segment is comprised of 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 of the bus segment.

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 9.1.2.2. See Section 9.1 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 shown in Table 3-1.

Table 3-1: SCSI Bus Speeds

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

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 to 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.

Note that you will not follow this general rule when using the DS-DWZZH-05 UltraSCSI hub with fair arbitration enabled.

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.

Note

In a TruCluster Server configuration, you always have single-ended SCSI bus segments since all of the storage shelves use a single-ended bus.

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	Maximum Cable Length
Narrow, single-ended	5 MB/sec	6 meters
Narrow, single-ended fast	10 MB/sec	3 meters
Wide differential, fast	20 MB/sec	25 meters
Differential UltraSCSI	40 MB/sec	25 meters ^a

^a The maximum separation between a host and the storage in a TruCluster Server configuration is 50 meters: 25 meters between any host and the UltraSCSI hub and 25 meters 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 9.2.

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 H8863-AA trilink connector on the HSZ70 or HSZ80 at the bus end. Connect an H8861-AA terminator to the trilink connector to terminate the bus.

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

Figure 3–1: VHDCI Trilink Connector (H8861-AA)



CXO5744A

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:

- Installed KZPBA-CB (or KZPSA-BB) termination resistor SIPs
- External termination on a trilink connector attached to an UltraSCSI BA356 personality module (DS-BA35X-DA), 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 as 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 a 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 which has been upgraded to the 180-watt power supply with the DS-BA35X-HH option.
- Has three Very High Density Cable Interconnect (VHDCI) differential SCSI bus connectors
- Does not use a SCSI ID
- Uses the storage shelf only to provide its power and mechanical support (it is not connected to the shelf internal SCSI bus).

- 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 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 shown in Table 3–3.

Note

With dual 180-watt power supplies installed, there are slots available for six 3.5-inch SBBs or two 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

Table 3–3: DS-DWZZH UltraSCSI Hub Maximum Configurations (cont.)

DS-DWZZH-03	DS-DWZZH-05	Disk Drives ^a	Personality Module ^{b c}
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 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 installed in the storage shelf.

3.6.1.2.2 DS-DWZZH-05 Fair Arbitration

Although each cluster member system and storage controller 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, it is likely that the adapter with the highest priority SCSI ID will 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, 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 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 (DRD), service failover through the cluster application availability (CAA) subsystem, disk mirroring, and fast file system recovery. TruCluster Server supports mirroring of the clusterwide root (`/`) file system, the member-specific boot disks, and the cluster quorum disk through hardware RAID only. You can mirror the clusterwide `/usr` and `/var` file systems and the data disks 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.

See the TruCluster Server *Software Product Description* (SPD) to determine the supported storage shelves, disk devices, and RAID array controllers.

Disk devices 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. 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 or HSZ80 RAID array controllers.

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

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. When used in a dual-redundant configuration, Tru64 UNIX Version 5.0A 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 storage sets 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 HSG80 array controllers with ACS Version 8.5.

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

The following sections describe how to cable the HSZ70 or HSZ80 for TruCluster Server configurations. 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 0.3-meter cable and the BN37A-0E is a 0.5-meter 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:
 - 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)
 - 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 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

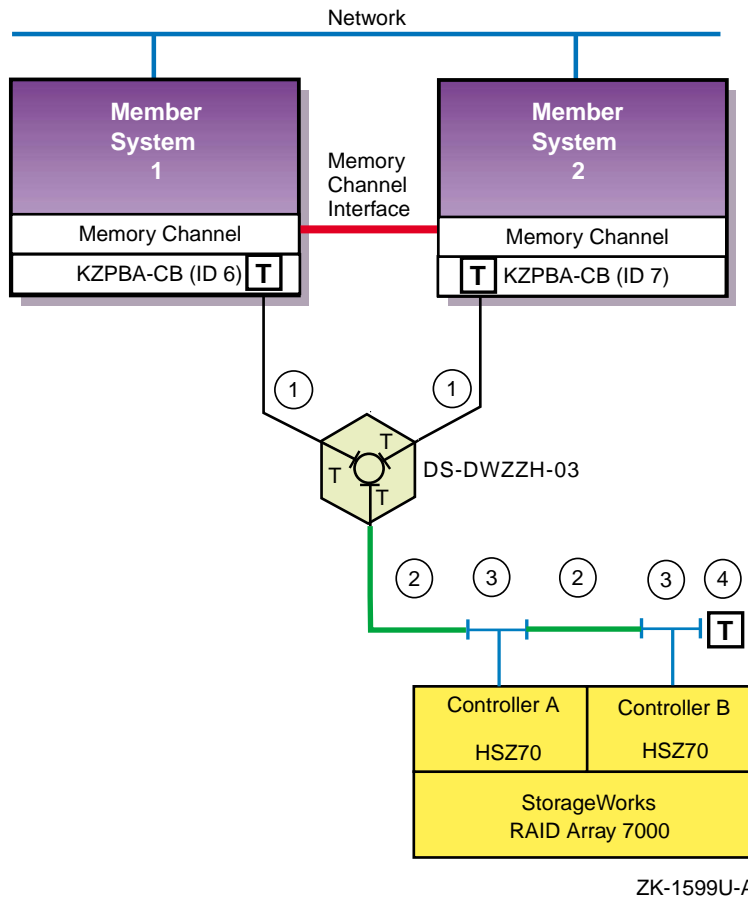


Table 3–4 shows the components used to create the clusters shown in Figure 3–5, Figure 3–6, Figure 3–7, and Figure 3–8.

Table 3–4: Hardware Components Used in Configuration 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.

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

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 shows the components used to create the cluster 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 host(s) have the capability to move LUNs from one controller (shared SCSI bus) to the other. If one host adapter or SCSI bus fails, the host(s) 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, it operates 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 are failed 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 are failed 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 tralink connectors.
2. Install H8861-AA VHDCI tralink connectors (with terminators) on:

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

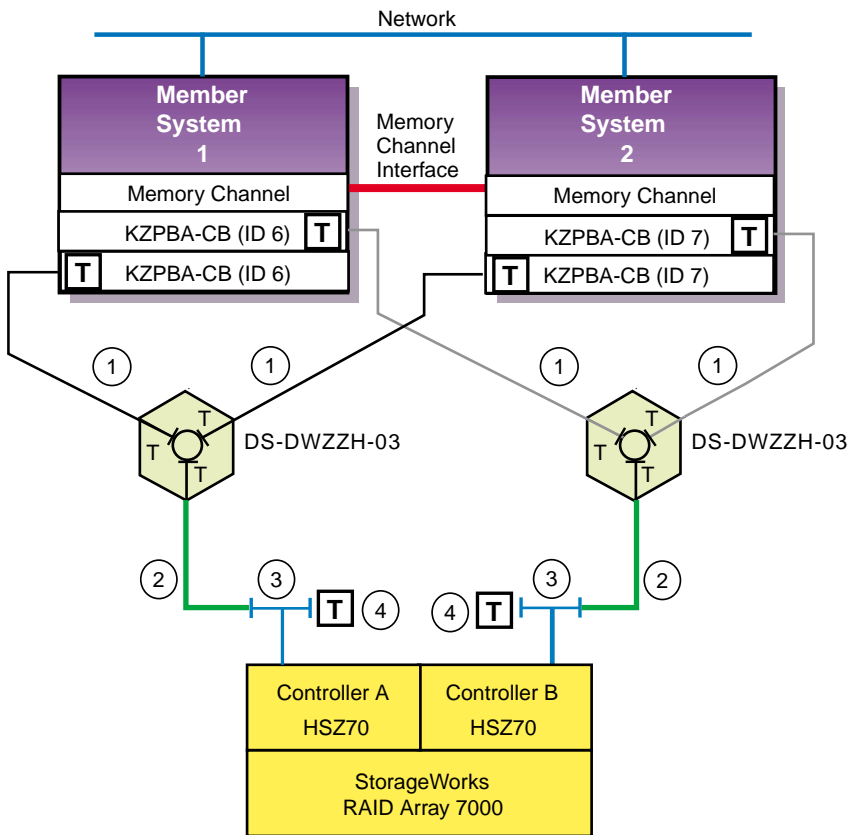
Note

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:
 - 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)
 - 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)
 - 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 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 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

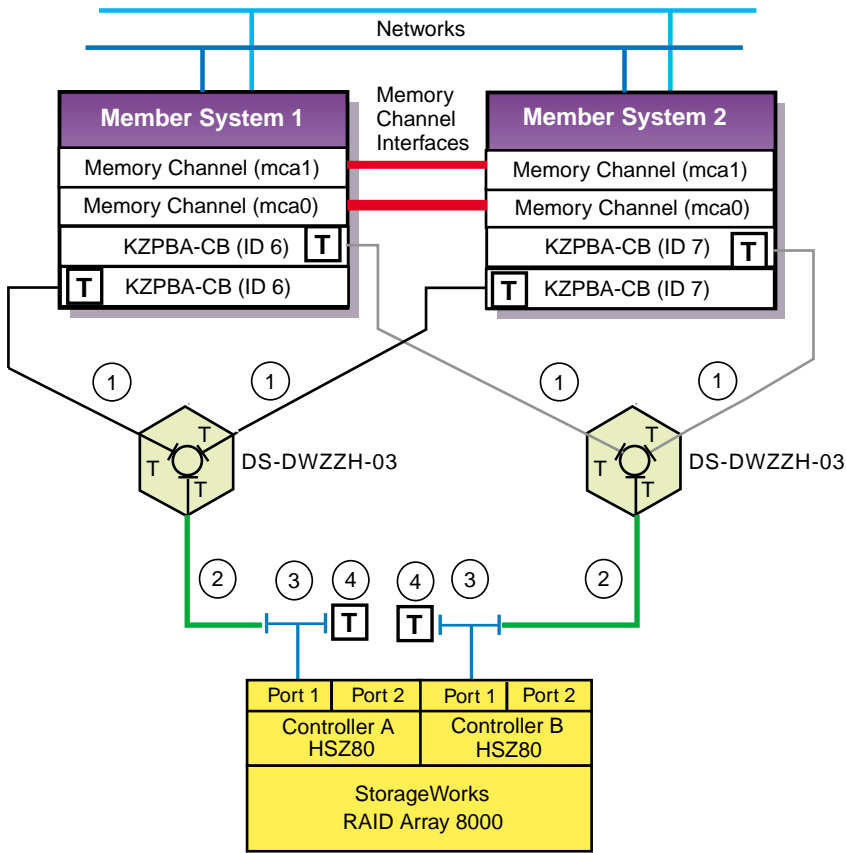


ZK-1601U-AI

Table 3–4 shows the components used to create the cluster 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 shows the components used to create the cluster shown in Figure 3–8.

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. This documentation should arrive 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.

You need to allocate disks for the following uses:

- One or more disks to hold the Tru64 UNIX operating system. The disk(s) are either private disk(s) on the system that will become the first cluster member, or disk(s) 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 1.4.1.4, and for a more detailed discussion of the quorum disk, see the TruCluster Server *Cluster Administration* manual.

All configurations covered in this manual assume 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.

This 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 TruCluster Server *Software 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 interconnect(s) 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 5.0A, the cluster interconnect is the Memory Channel. 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 clusterwide root (/) file system, the member-specific boot disks, and the swap and 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 a NSPOF TruCluster Server cluster (providing 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), 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.

Table 4–1: Planning Your Configuration (cont.)

To increase:	You can:
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 Alpha Systems from the downloadable drivers & utilities menu. Then 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 applicable to the Tru64 UNIX version installed (or to be installed) 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 (discussed in Chapter 9 and Chapter 10) 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 covered in Chapter 10.

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 the appropriate table 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 handling modules.

Table 4-2: Configuring TruCluster Server Hardware

Step	Action	Refer to:
1	Install the Memory Channel module(s), cables, and hub(s) (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.	Chapter 7 and <i>ATMworks 350 Adapter Installation and Service</i>
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.	Use 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 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 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 (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 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 9.

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.

Your storage shelves or RAID array subsystems should be 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 2.4.2, 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, 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 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 6 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

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

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


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

```
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
Bus 02 Slot 01: NCR 53C875
                        pkb0.7.0.2001.0   SCSI Bus ID 7
```

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

```
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.

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.

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 says).

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.

Notes

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.

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, Memory Channel optical converter (MC2 only), and connecting link cables.

Two versions of the Memory Channel 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). As 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.

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

See Section 2.2 for a discussion on Memory Channel restrictions.

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 optics 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 wish to complete all hardware installation before powering up the systems to run Memory Channel diagnostics.

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 Jumpers

The MC1 and MC1.5 modules (CCMAA-AA and CCMAA-AB respectively) have adapter jumpers that designate 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 jumpers set for standard hub mode (pins 1 to 2 jumpered). Confirm that the jumpers are set properly for your configuration. The jumper configurations are shown as if you were holding the module with the jumpers facing you, with the module end plate in your left hand. The jumpers are right next to the factory/maintenance cable connector, and are described in Table 5-1.

Table 5-1: MC1 and MC1.5 Jumper Configuration

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



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











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



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

Jumper:	Description:	Example:
J10 and J11: Fiber Optics Mode Enable	Fiber Off: Pins 1 to 2	
	Fiber On: Pins 2 to 3 pins	

^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 optics 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	
	Fiber On: Pins 1 to 2	

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 right 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 adapter(s), replace the system panels unless you have more hardware to install.

5.3 Installing the MC2 Optical Converter in the Member System

If you are going 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 optics cable (BN34R) through the PCI bulkhead slot.
3. Thread the optics 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 BN39B-01 cable from the CCMAB Memory Channel 2 PCI adapter to the CCMFB optical converter.
7. Route the fiber optics 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 10 meters of each of the systems.

For use with the MC2 CCMAB adapter, the hub must be placed within 4 or 10 meters (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 31 meters 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. Note that 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 optics 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 optics connections. You will need one hub for the CCMLB linecards (and possible optics converters) and up to two hubs for the CCMFB optics 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 10-meter 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 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 5.0A 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) or BN39B-10 (10-meter) link cables for virtual hub or standard hub configurations without optical converters.

If optical converters are used, use the BN39B-01 1-meter link cable and the BN34R-10 (10-meter) or BN34R-31 (31-meter) fiber optics 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) or BN39B-10 (10-meter) 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. You should 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 optics cable (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) or BN39B-10 (10-meter) 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 optics cable, when you install the CCMAB Memory Channel PCI adapter in each system in the standard hub configuration. You should also connect the CCMAB Memory Channel adapter to the optical converter with a BN39B-01 cable.

Now you need to:

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

Note

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

To set the CCMLB jumpers and install CCMFB optics 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 optics cable through the hub bulkhead slot. The other end should be attached to a CCMFB optics converter in the member system.
5. Thread the BN34R fiber optics 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 optics converter in slot `opto only`, `0/opto`, `1/opto`, `2/opto`, or `3/opto` as appropriate.
7. Install a BN39B-01 1-meter 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 optics converters, and fiber optics 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 adapter(s) 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 optics converters, BN34R fiber optics 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 A Ctrl/C on node 1 terminates the test.

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 5.0A configuration.

The information includes how to determine the `/dev/disk/dskn` value that corresponds to the Fibre Channel storage sets 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 5.0A and TruCluster Server Version 5.0A installation.

Note

TruCluster Server Version 5.0A configurations require one or more disks to hold the Tru64 UNIX operating system. The disk(s) are either private disk(s) on the system that will become the first cluster member, or disk(s) 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.

This chapter discusses the following topics:

- The procedure for Tru64 UNIX Version 5.0A or TruCluster Server Version 5.0A installation using Fibre Channel disks (Section 6.1).
- Fibre Channel overview (Section 6.2).
- Example cluster configurations using Fibre Channel storage (Section 6.3).
- A brief discussion of zoning and cascaded switches (Section 6.4).
- The steps necessary to install the Fibre Channel hardware, base operating system, and cluster software using disks accessible over the Fibre Channel hardware (Section 6.5 through Section 6.10).
- Changing the HSG80 from transparent to multiple-bus failover mode.

- 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 6.12).

6.1 Procedure for Installation Using Fibre Channel Disks

Use the following procedure to install Tru64 UNIX Version 5.0A or TruCluster Server Version 5.0A using Fibre Channel disks. If you are only installing Tru64 UNIX Version 5.0A, complete the first eight steps. Complete all the steps for a TruCluster Server Version 5.0A installation. Refer to the Tru64 UNIX *Installation Guide*, TruCluster Server *Software Installation* manual, and other hardware manuals as appropriate for the actual installation procedures.

1. Install the Fibre Channel switch (Section 6.5.1).
2. Install the KGPSA PCI-to-Fibre Channel host bus adapter (Section 6.5.2).
3. Set up the HSG80 RAID array controllers for a fabric configuration (Section 6.5.3).
4. Configure the HSG80 disks to be used for base operating system and cluster installation. Be sure to set the identifier for each storage unit you will use for operating system or cluster installation (Section 6.6.1).
5. Power on the system where you will install Tru64 UNIX Version 5.0A. 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 5.0A disk and cluster member system boot disks (Section 6.6.2).

6. Use the WWID manager to set the `bootdef_dev` console environment variable (Section 6.6.3).
7. Refer to the Tru64 UNIX *Installation Guide* and install the base operating system from CD-ROM. The installation procedure will recognize the disks for which you set the device unit number. Select the disk you have chosen as the base operating system installation disk from the list of disks provided (Section 6.7).
8. After the new kernel has booted to multi-user mode, shut down the operating system and reset the `bootdef_dev` console environment variable to provide multiple boot paths (Section 6.8).

Boot the system to multi-user mode and complete the operating system installation.

9. Determine the `/dev/disk/dskn` values to be used for cluster installation (Section 6.9).

10. Use the `disklabel` utility to label the disks used to create the cluster (Section 6.10).
11. Refer to the TruCluster Server *Software Installation* manual and install the TruCluster Server software subsets and 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 6.10).
12. Reset the `bootdef_dev` console environment variable to provide multiple boot paths (Section 6.8). Boot the first cluster member.
13. Run `clu_add_member` on a cluster member system 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 (Section 6.6.2).
 - b. Set the `bootdef_dev` console environment variable to one reachable path to the member system boot disk (Section 6.6.3)
 - c. Boot `genvmunix`.
14. Create a new kernel, but do not reboot. Shut the system down and reset the `bootdef_dev` console environment variable to provide multiple boot paths to the member system boot disk (Section 6.8). Boot the new cluster member system.
 15. Repeat steps 13 and 14 to add other cluster member systems.

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

- *KGPSA-BC PCI-to-Optical Fibre Channel Host Adapter User Guide (AA-RF2JB-TE)*
- *Compaq StorageWorks Fibre Channel Storage Switch User's Guide (AA-RHBYA-TE)*
- *Compaq StorageWorks SAN Switch 8 Installation and Hardware Guide (EK-BCP24-IA/161355-001)*
- *Compaq StorageWorks SAN Switch 16 Installation and Hardware Guide (EK-BCP28-IA/161356-001)*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Configuration Guide*
- *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*

- *Wwidmgr User's Manual*

6.2 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 that it is used for.

The 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 fibre transmission lines allows the high-frequency information to be sent up to 40 km, the maximum distance between transmitter and receiver. Copper transmission lines may be used for shorter distances.

6.2.1 Basic Fibre Channel Terminology

The following list describes the basic Fibre Channel terminology:

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. The information is repeated by each intermediate port until it reaches its destination. The N_Port that contains this additional loop functionality is an NL_Port.
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 fibre, coaxial cable, or shielded twisted pair.
E_Port	An expansion port on a switch used to make a connection between two switches in the fabric.

6.2.2 Fibre Channel Topologies

Fibre Channel supports three different interconnect topologies:

- Point-to-point (Section 6.2.2.1)
- Fabric (Section 6.2.2.2)

- Arbitrated loop (Section 6.2.2.3)

Note

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


6.2.2.1 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



ZK-1534U-AI

6.2.2.2 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 is 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


6.2.2.3 Arbitrated Loop Topology

In an arbitrated loop topology, frames are routed around a loop created by the links between the nodes.

In an 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 arbitrated loop topology.

Figure 6–3: Arbitrated Loop Topology



ZK-1535U-AI

Note

The arbitrated loop topology is not supported by the Tru64 UNIX or TruCluster Server products.

When support for Fibre Channel arbitrated loop is announced in the TruCluster Server *Software Product Description (SPD)*, the technical update version of this information will be modified to include arbitrated loop. The SPD will provide a pointer to the technical update.

6.3 Example Fibre Channel Configurations Supported by TruCluster Server

This section provides diagrams of some of the configurations supported by TruCluster Server Version 5.0A. Diagrams are provided for both transparent failover mode and multiple-bus failover mode.

6.3.1 Fibre Channel Cluster Configurations for Transparent Failover Mode

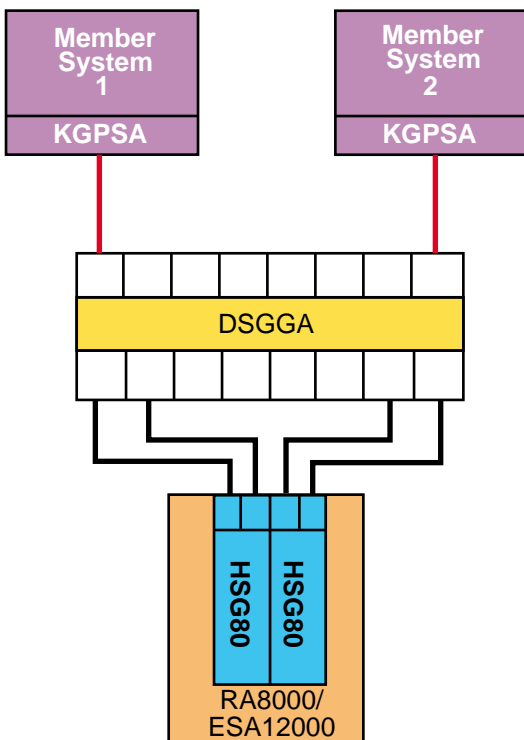
With transparent failover mode:

- 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 (with the limit of a total of 128 storage units).

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 a cable, you lose the units behind at least one port.

You can, however, add the hardware for a second bus (another KGPSA, switch, and RA8000/ESA12000 with associated cabling) and use LSM to mirror across the buses. However, because you cannot use LSM to mirror the cluster root (/) file system, member boot partitions, the quorum disk, or swap partitions you cannot obtain an NSPOF transparent failover configuration, even though you have increased availability.

6.3.2 Fibre Channel Cluster Configurations for Multiple-Bus Failover Mode

With multiple-bus failover:

- The host controls the fail over by accessing units over a different path or causing the access to the unit to be through the other HSG80 controller (one controller does not fail over to the other controller of its own accord).
- Each cluster member system has two or more KGPSA host bus adapters (multiple paths to the storage units).
- Normally, all available units (D0 through D199, with a limit of 128 storage units) 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

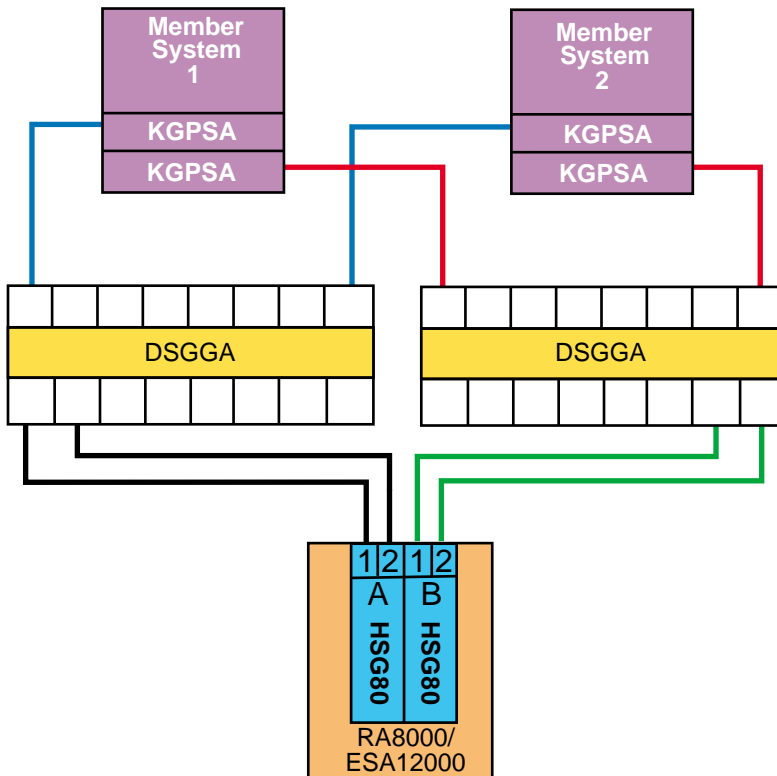
If you have preferred units, and the HSG80 controllers restart because of an error condition or power failure, and one controller restarts before the other controller, the HSG80 controller restarting first will take all the units, whether they are preferred or not. When the other HSG80 controller starts,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

Figure 6–5, Figure 6–6, and Figure 6–7 show three different multiple-bus NSPOF cluster configurations. The only difference is the fibre-optic cable connection path between the switch and the HSG80 controller ports.

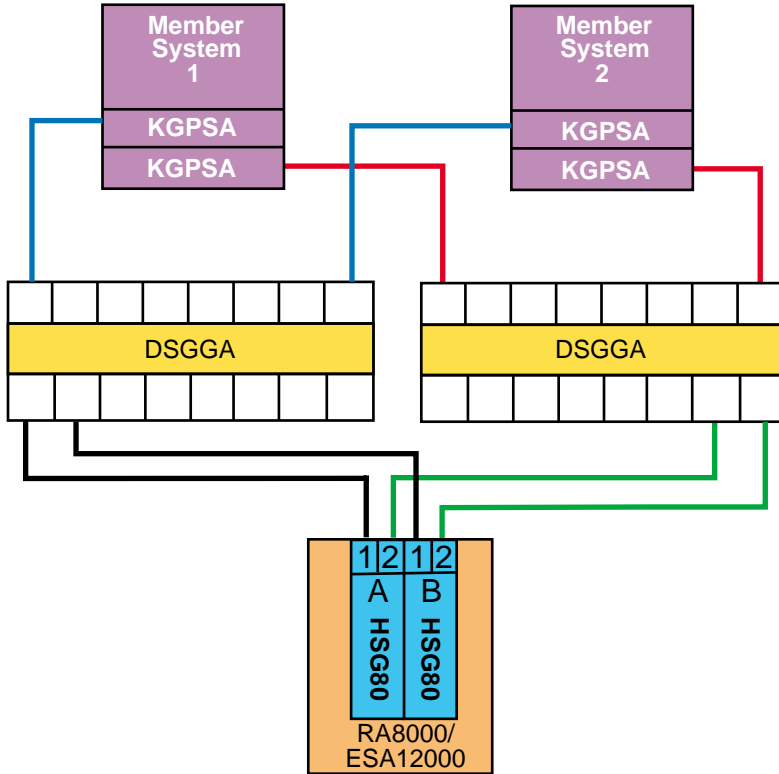
If you consider the loss of a host bus adapter or switch, the configurations in Figure 6-6 and Figure 6-7 will provide better throughput than Figure 6-5 because you still have access to both controllers. With Figure 6-5, if you lose a host bus adapter or switch, you lose the use of a controller.

Figure 6-5: Multiple-Bus NSPOF Configuration Number 1



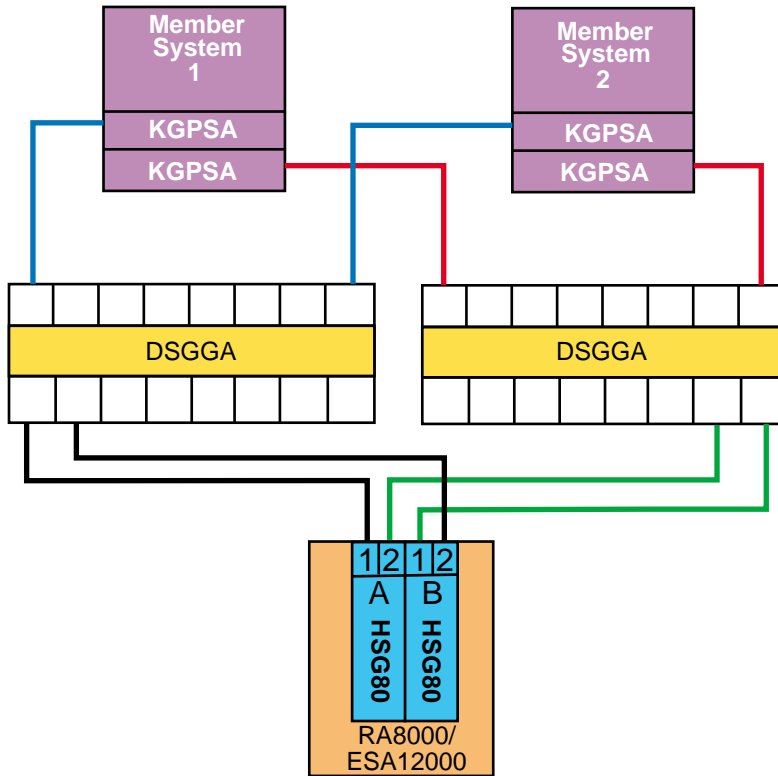
ZK-1706U-AI

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



ZK-1707U-AI

Figure 6–7: Multiple-Bus NSPOF Configuration Number 3



ZK-1708U-AI

6.4 Zoning and Cascaded Switches


This section provides a brief overview of zoning and cascaded switches.

6.4.1 Zoning

A zone is a logical subset of the Fibre Channel devices 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.

Figure 6–8 provides an example configuration using zoning. This configuration consists of two independent zones with each zone containing an independent cluster.

Figure 6–8: A Simple Zoned Configuration



ZK-1709U-AI

Note

Only static zoning is supported; zones can only be changed when all connected systems are shut down.

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

6.4.2 Cascaded Switches

Multiple switches may be connected to each other. When cascading switches, a maximum of three switches is supported, with a maximum of two hops between switches. The maximum hop length is 10 km longwave single-mode or 500 meters shortwave multimode Fibre Channel cable.

6.5 Installing and Configuring Fibre Channel Hardware

This section provides information about installing the Fibre Channel hardware needed for a TruCluster Server configuration accessing storage over the Fibre Channel.

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

Note

The maximum length of the optical cable between the KGPSA and the switch or switch and the HSG80 array controller is 500 meters via shortwave multimode Fibre Channel cable. The maximum distance between switches in a cascaded switch configuration is 10 kilometers using longwave single-mode fibre.

6.5.1 Installing and Setting Up the Fibre Channel Switch

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

Each switch, except the DS-DSGGB-AB, has a front panel display and four push buttons that you use to manage the switch. There are four menus that allow you to configure, operate, obtain status, or test the switch. The DS-DSGGB-AB is managed by way of a telnet session once the IP address has been set (from a PC or terminal).

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

- Remotely using a telnet TCP/IP connection
- With 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) before

you can manage the switch by way of a telnet session, SNMP, or the Web.

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 Giga Bit Interface Converter modules (GBIC). The GBIC module is the electrical-to-optical converter.

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

Note

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

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

6.5.1.1 Installing the Switch

Place the switch within 500 meters 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:

1. Place the switch or install it in the rack.
2. Connect the Ethernet cable.
3. Connect the fibre-optic cables.
4. Connect power to the switch.

5. Turn on the power. The switch runs a series of power-on self test (POST) tests.
6. Set the switch IP address and subnet mask (see Section 6.5.1.2.2). You can also set the switch name if desired (see Section 6.5.1.2.5). The switch IP address and subnet mask must initially be set from the front panel, except for the DS-DSGGB-AA 8-port Fibre Channel switch. In this case you have to connect a PC or terminal to the switch. You must use a telnet session to set the switch name.
7. 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* (AA-RHBYA-TE)
- *Compaq StorageWorks SAN Switch 8 Installation and Hardware Guide* (EK-BCP24-IA/161355-001)
- *Compaq StorageWorks SAN Switch 16 Installation and Hardware Guide* (EK-BCP28-IA/161356-001)

For more information on the DSGGB command set, see the *Compaq StorageWorks SAN Switch Fabric Operating System Management Guide* (EK-P20FF-GA/161358-001).

6.5.1.2 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 does 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).

6.5.1.2.1 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 being displayed.

Down — Downward triangle: Scrolls the menu down (which effectively moves up the list of commands) or decreases the value 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.

6.5.1.2.2 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.

6.5.1.2.3 Setting the DS-DSGGB-AA Ethernet IP Address and Subnet Mask from a PC or Terminal

For the DS-DSGGB-AA switch, which does not have a front panel, you must use a connection to a Windows 95/98/NT 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 switch, 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.
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, and no parity.
4. Turn on power to the switch. The switch automatically connects to the host and logs the user on to the switch as `admin`.
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]: 
Fibre Channel IP Address [none]: 
Fibre Channel Subnetmask [none]: 
Gateway Address [none]: 
admin> logout
```

6.5.1.2.4 Logging Into 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 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, shown from lowest security level to highest security level, are shown in Table 6-1.

Table 6–1: Telnet Session Default User Names for Fibre Channel Switches

DSGGA	DSGGB	Description
other	n/a	Allows you to execute commands ending in Show, such as <code>dateShow</code> and <code>portShow</code> .
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, <code>version</code> and <code>errDump</code> . 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 <code>admin</code> .
n/a	root	Gives users access to an extensive command set that can significantly alter system performance. Root commands should only be used 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.

6.5.1.2.5 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 Return
User admin Return
Passwd Return
:Admin> switchName fcswl Return
:Admin> switchName Return
fcswl
:Admin>
```

Note

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

```
fcswl:Admin>
```

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

The following sections discuss KGPSA installation and configuration.

6.5.2.1 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 (AA-RF2JB-TE)*
- *64-Bit PCI-to-Fibre Channel Host Bus Adapter User Guide (AA-RKPDA-TE/173941-001)*

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- 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. Do not forget to remove the transparent plastic covering on the extremities of the optical cable.
5. Connect the fibre-optic cables to the shortwave gigabit interface converter modules (GBICs) in the DSGGA or DSGGB Fibre Channel switch.

6.5.2.2 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 -topology fabric
Reformatting nvram
Reformatting nvram
```

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.
```

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

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

6.5.2.3 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. You should obtain and record the worldwide names of Fibre Channel components in case you need to verify their target ID mappings in the operating system.

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

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

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

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

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

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

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

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

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

6.5.3 Setting up the HSG80 Array Controller for Tru64 UNIX Installation

This section covers setting up the HSG80 controller for operation with Tru64 UNIX Version 5.0A and TruCluster Server Version 5.0A. For more information on installing the HSG80, see the *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 Configuration Guide* or *Compaq StorageWorks HSG80 Array Controller ACS Version 8.5 CLI Reference Guide*.

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

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

Note

Do not use PVA ID 1:

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

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

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

Note

The KGPSA host bus adapters must be cabled to the switch, with the system power applied before you turn power on to the RA8000/ESA12000, in order for the HSG80 to see the connection to the KGPSAs.

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:

```
> 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:

```
set nofailover           1
clear cli                2
set multibus copy = this 3
clear cli                4
set this port_1_topology = offline 5
set this port_2_topology = offline 5
set other port_1_topology = offline 5
set other port_2_topology = offline 5
set this port_1_topology = fabric 6
set this port_2_topology = fabric 6
set other port_1_topology = fabric 6
set other port_2_topology = fabric 6
```

- 1 Remove 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 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 command `set failover copy = this_controller` to set transparent failover mode.

- 4 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.

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

6 Sets fabric as the switch topology.

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. 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        THIS        2    230813  OL this  0
HOST_ID=1000-0000-C920-DA01 ADAPTER_ID=1000-0000-C920-DA01
!NEWCON50      TRU64_UNIX        THIS        1    230813  OL this  0
HOST_ID=1000-0000-C920-DA01 ADAPTER_ID=1000-0000-C920-DA01
!NEWCON51      TRU64_UNIX        THIS        2    230913  OL this  0
HOST_ID=1000-0000-C920-EDEB ADAPTER_ID=1000-0000-C920-EDEB
!NEWCON52      TRU64_UNIX        THIS        1    230913  OL this  0
HOST_ID=1000-0000-C920-EDEB ADAPTER_ID=1000-0000-C920-EDEB
!NEWCON53      TRU64_UNIX        OTHER       1    230913  OL other  0
HOST_ID=1000-0000-C920-EDEB ADAPTER_ID=1000-0000-C920-EDEB
!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
```

Example 6–1: Determine HSG80 Connection Names (cont.)

```
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. 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 could 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
HSG80> set !NEWCON49 operating_system = TRU64_UNIX
HSG80> restart other
HSG80> restart this
      :
HSG80> show connection
```

- 1 Set the relative offset for LUN numbering to 0. You can set the `unit_offset` to nonzero values, but use caution. Make sure you understand the impact.
- 2 Specify that the host environment connected to the Fibre Channel port is `TRU64_UNIX`. You must change each connection to `TRU64_UNIX`. This is very important. Failure to set this to `TRU64_UNIX` will prevent your system from booting correctly, recovering from run-time errors, or from booting at all. The default operating system is Windows NT, and NT uses a different SCSI dialect to talk to the HSG80 controller.
- 3 Restart both controllers to cause all changes to take effect.
- 4 Enter the `show connection` command once more and verify that all connections have the offsets set to 0 and the operating system is set to `TRU64_UNIX`.

Note

If the fibre-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 6.6.1.

6.5.3.1 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 6.5.3.1) and other configuration information, you can use the command-line interface (CLI) commands to configure the new controller and set the worldwide name. Set the worldwide name as follows:

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

6.6 Preparing to Install Tru64 UNIX and TruCluster Server on 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:

- Configure HSG80 storagesets — In this document, example storagesets are configured for both Tru64 UNIX and TruCluster Server on Fibre Channel storage. Modify the storage configuration to meet your needs (Section 6.6.1).
- Set the device unit number — The device unit number is a subset of the device name (as shown in a `show device` display). For example, in the device name `DKA100.1001.0.1.0`, the device unit number is 100 (DKA100). The Fibre Channel worldwide name (often referred to as the worldwide ID or WWID) is too long (64 bits) to be used as the device unit number. Therefore, you set a device unit number that is an alias for the Fibre Channel worldwide name (Section 6.6.2).
- 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 right disk (Section 6.6.3).

6.6.1 Configuring 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 of course any disks for applications). This document assumes 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 disk(s) are either private disk(s) on the system that will become the first cluster member, or disk(s) on a shared bus that the system can access. Whether the Tru64 UNIX disk is on a private disk or a shared disk, you should 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 will be based on the use of four 4-GB disks used to create two mirrorsets (RAID level 1) to provide reliability. The mirrorsets will be partitioned to provide partitions of appropriate sizes. Disks 30200, 30300, 40000, and 40100 will be used for the mirrorsets.

Table 6-2 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 will be used for the Tru64 UNIX and cluster member system boot disks. The other mirrorset, `CROOT-MIR`, will be 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
...
```


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

Config - Normal Termination

```
HSG80> ADD MIRRORSET BOOT-MIR DISK 30200 40000
HSG80> ADD MIRRORSET CROOT-MIR DISK 30300 40100
HSG80> INITIALIZE BOOT-MIR
HSG80> INITIALIZE CROOT-MIR
HSG80> SHOW BOOT-MIR
```

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
```

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
HSG80> CREATE_PARTITION BOOT-MIR SIZE=25
HSG80> CREATE_PARTITION BOOT-MIR SIZE=LARGEST
HSG80> CREATE_PARTITION CROOT-MIR SIZE=5
HSG80> CREATE_PARTITION CROOT-MIR SIZE=15
HSG80> CREATE_PARTITION CROOT-MIR SIZE=40
HSG80> CREATE_PARTITION CROOT-MIR SIZE=LARGEST
HSG80> SHOW BOOT-MIR
```

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	
2	2094502 (1072.38 MB)	2094507	
3	4189009 (2144.77 MB)	4189014	

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

```

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

HSG80> ADD UNIT D131 BOOT-MIR PARTITION=1 DISABLE_ACCESS_PATH=ALL 16
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 17
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 18
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

```

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

```
HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON53,!NEWCON54,!NEWCON55,!NEWCON56
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified
HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON61,!NEWCON62,!NEWCON64,!NEWCON65
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified
HSG80> set d144 ENABLE_ACCESS_PATH = !NEWCON68,!NEWCON74,!NEWCON76,!NEWCON77
Warning 1000: Other host(s) in addition to the one(s) specified can still
              access this unit.  If you wish to enable ONLY the host(s)
              specified, disable all access paths (DISABLE_ACCESS=ALL), then
              again enable the ones specified
```

```
HSG80> show d131 19
-----
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 19
-----
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 Use the `CONFIG` utility to configure the devices on the device side buses and add them to the controller configuration. The `CONFIG` utility takes about two minutes to complete. You can use the `ADD DISK` command to add disk drives to the configuration manually.
- 2 Create the `BOOT-MIR` mirrorset using disks 30200 and 30300 and the `CROOT-MIR` mirrorset using disks 40000 and 40100.
- 3 Initialize the `BOOT-MIR` and `CROOT-MIR` mirrorsets. If you want to set any initialization switches, you must do so in this step. The `BOOT-MIR` mirrorset will be used for the Tru64 UNIX and cluster member system boot disks. The `CROOT-MIR` mirrorset will be used for the cluster root (`/`), cluster `/usr` and cluster `/var` file systems, and the quorum disk.
- 4 Verify the mirrorset configuration and switches. Ensure that the mirrorsets use the correct disks.
- 5 Create appropriate sized partitions in the `BOOT-MIR` mirrorset using the percentage of the storageset that each partition will use. These partitions will be used for the two member system boot disks (25 percent or 1 GB each) and the Tru64 UNIX disk. For the last partition, the controller assigns the largest free space available to the partition (which will be close to 50 percent or 2 GB).
- 6 Create appropriate sized partitions in the `CROOT-MIR` mirrorset using the percentage of the storageset that each partition will use. These partitions will be used for the quorum disk (5 percent), cluster root partition (15 percent), `/usr` (40 percent), and `/var` file systems. For the last partition, `/var`, the controller assigns the largest free space available to the partition (which will be close to 40 percent). See the TruCluster Server *Software Installation* manual to obtain partition sizes.
- 7 Verify the `BOOT-MIR` mirrorset partitions. Ensure that the partitions are of the desired size. The partition number is in the first column, followed by the partition size and starting block.
- 8 Partition for member system 1 boot disk.
- 9 Partition for member system 2 boot disk.
- 10 Partition for Tru64 UNIX operating system disk.
- 11 Verify the `CROOT-MIR` mirrorset partitions. Ensure that the partitions are of the desired size. The partition number is in the first column, followed by the partition size and starting block.
- 12 Partition for the quorum disk.
- 13 Partition for cluster root (`/`) filesystem.

- 14 Partition for cluster `/usr` filesystem.
- 15 Partition for cluster `/var` filesystem.
- 16 Assign a unit number to each partition. When the unit is created by the `ADD UNIT` command, disable access to all hosts. This allows selective access in case there are other systems or clusters connected to the same switch as our cluster.

Record the unit name of each partition with the intended use for that partition (see Table 6–2).

- 17 Set the identifier for each storage unit. Use any number between 1 and 9999. The number you select for the storage unit shows up as the user define identifier (UDID) in the `wwidmgr -show wwid` display. It will be used by the WWID manager when setting the device unit number and `bootdef_dev` console environment variable.

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

It will also show up during the Tru64 UNIX installation to allow you to select the Tru64 UNIX installation disk.

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 (see Table 6–2).

- 18 Enable access to each unit for those hosts that you want to be able to access this unit. Because access was initially disabled to all hosts, you can ensure selective access to the units.

Use the connection name for each connection to the KGPSA host bus adapter on the host for which you want access enabled. Many of the connections used here are shown in Example 6–1.

- 19 Use the `SHOW unit` command (where `unit` is D131 through D133 and D141 through 144 in the example) to verify the identifier and that access to each unit is correct. Ensure that there is no connection to an unwanted system. Record the identifier and worldwide name for later use. Table 6–2 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

that at this point, even though the table is filled in, we do not yet know the device names or `dskn` numbers.

Table 6–2: Converting Storageset Unit Numbers to Disk Names

File System or Disk	HSG80 Unit	Worldwide Name	User Define Identifier (UDID)	Device Name	<code>dskn</code>
Member 1 boot disk	D131	6000-1FE1-0000-0D60-0009-8080-0434-002F	131	dga131.1001.0.1.0	dsk17
Member2 boot disk	D132	6000-1FE1-0000-0D60-0009-8080-0434-0030	132	dga132.1001.0.1.0	dsk16
Tru64 UNIX disk	D133	6000-1FE1-0000-0D60-0009-8080-0434-002E	133	dga133.1001.0.1.0	dsk15
Quorum disk	D141	6000-1FE1-0000-0D60-0009-8080-0434-0029	141	N/A ^a	dsk21
Cluster root (/)	D142	6000-1FE1-0000-0D60-0009-8080-0434-002A	142	N/A ^a	dsk20
/usr	D143	6000-1FE1-0000-0D60-0009-8080-0434-002B	143	N/A ^a	dsk19
/var	D144	6000-1FE1-0000-0D60-0009-8080-0434-0028	144	N/A ^a	dsk18

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

6.6.2 Setting the Device Unit Number

Set the device unit number for the Fibre Channel disks to be used as the Tru64 UNIX Version 5.0A installation disk and cluster member boot disks. Setting the device unit number allows the installation scripts to recognize a Fibre Channel disk. You have to set the device unit number because the 64-bit worldwide name is too large to be used as the device unit number. When you set the device unit number, you set an alias for the device worldwide name.

You use the WWID manager (`wwidmgr`) to define a device unit number that is an alias for Fibre Channel devices. For instance, if `DKA0` or `DKA100` are part of the device name seen in a `show dev display`, `0` or `100` is the device unit number.

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

1. Obtain the user define identifier (UDID) for the HSG80 storageset to be used as the Tru64 UNIX Version 5.0A installation disk or cluster member system boot disks. In the example in Table 6–2 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. Use the `wwidmgr -clear all` command to clear the stored Fibre Channel `wwid1`, `wwid2`, `wwid3`, `wwid4`, `N1`, `N2`, `N3`, and `N4` console environment variables. You want to start with all `wwid<n>` and `N<n>` variables clear.

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

Note

The console only creates devices for which the `wwid<n>` console environment variable has been set, and are accessible through an HSG80 N_Port as specified by the `N<n>` console environment variable also being set. These console environment variables are set with the `wwidmgr -quickset` or `wwidmgr -set wwid` commands. We use the `wwidmgr -quickset` command later.

3. Use the `wwidmgr -show wwid` command to display the UDID and worldwide names of all 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 `wwid<n>` console environment variables (`wwid0`, `wwid1`, `wwid2`, and `wwid3`) can be set. The console `show dev` command only shows those disk devices for which a `wwid<n>` console environment variable has been set using the `wwidmgr -quickset` or `wwidmgr -set` command. In this example, none of the `wwid<n>` environment variables is set.
4. Look through the `wwidmgr -show wwid display` 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. Example 6-3 shows the use of the `wwidmgr` command with the `-quickset` option to define the UDID as the device unit number as an alias for the worldwide name for each of the devices. 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).

Example 6-3 shows:

- The use of the `wwidmgr -quickset` command to set the device unit number for the Tru64 UNIX Version 5.0A installation disk to 133, the cluster member system boot disks to 131 (cluster member 1) and

132 (cluster member 2). The device unit number is an alias for the worldwide name for the storage unit.

- 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 KGPSA
 - Whether access is available through a port
 - The HSG80 port (N_Port) that will be used to access the storage unit
 - The `connected` column indicates the HSG80 controller ports that will be used to access the storage units. The HSG80 controllers are in multiple-bus failover, so only one controller is active.

Example 6–3: Using the `wwidmgr quickset` Command to Set Device Unit Number

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

Disk assignment and reachability after next initialization:

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

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

Disk assignment and reachability after next initialization:

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

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

Example 6–3: Using the wwidmgr quickset Command to Set Device Unit Number (cont.)

```

dgb131.1004.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d61      Yes

P00>>> wwidmgr -quickset -udid 132

Disk assignment and reachability after next initialization:

6000-lfe1-0000-0d60-0009-8080-0434-002e
      via adapter:      via fc nport:      connected:
dga133.1001.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d64      No
dga133.1002.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d62      Yes
dga133.1003.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d63      No
dga133.1004.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d61      Yes
dgb133.1001.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d64      No
dgb133.1002.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d62      Yes
dgb133.1003.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d63      No
dgb133.1004.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d61      Yes

6000-lfe1-0000-0d60-0009-8080-0434-002f
      via adapter:      via fc nport:      connected:
dga131.1001.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d64      No
dga131.1002.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d62      Yes
dga131.1003.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d63      No
dga131.1004.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d61      Yes
dgb131.1001.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d64      No
dgb131.1002.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d62      Yes
dgb131.1003.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d63      No
dgb131.1004.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d61      Yes

6000-lfe1-0000-0d60-0009-8080-0434-0030
      via adapter:      via fc nport:      connected:
dga132.1001.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d64      No
dga132.1002.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d62      Yes
dga132.1003.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d63      No
dga132.1004.0.1.0      pga0.0.0.1.0      5000-lfe1-0000-0d61      Yes
dgb132.1001.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d64      No
dgb132.1002.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d62      Yes
dgb132.1003.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d63      No
dgb132.1004.0.2.0      pgb0.0.0.2.0      5000-lfe1-0000-0d61      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 console is in diagnostic mode.

The disks are not reachable and you cannot boot until after the system is initialized.

Note, that in the reachability portion of the display, the storagesets are reachable from KGPSA dga through two HSG80 ports and from KGPSA dgb through two HSG80 ports. Also note, that the device unit numbers, the alias for the worldwide name of the disk device, has been set for the KGPSA for each HSG80 port. The device names have also been set for the cluster member boot disks. Record the device names.

The `wwidmgr -quickset` command provides a reachability display (equivalent to execution of the `wwidmgr -reachability` command). The devices shown in the reachability display 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 would show that the environment variable `wwidn` is set for the three boot disks. Also, the `show n*` command shows that the units are accessible through 4 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 132 1 WWID:01000010:6000-1fe1-0000-0d60-0009-8080-0434-0030
wwid3
P00>>> show n*
N1          50001fe100000d64
N2          50001fe100000d62
N3          50001fe100000d63
N4          50001fe100000d61
```

Example 6–4 provides sample device names as displayed by the `show dev` command after using the `wwidmgr -quickset` command to set the device unit numbers.

Example 6–4: 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
dga132.1001.0.1.0      $1$DGA132      HSG80  V8.5F
dga132.1002.0.1.0      $1$DGA132      HSG80  V8.5F
dga132.1003.0.1.0      $1$DGA132      HSG80  V8.5F
dga132.1004.0.1.0      $1$DGA132      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
```

Example 6–4: Sample Fibre Channel Device Names (cont.)

dgb132.1001.0.2.0	\$1\$DGA132	HSG80	V8.5F
dgb132.1002.0.2.0	\$1\$DGA132	HSG80	V8.5F
dgb132.1003.0.2.0	\$1\$DGA132	HSG80	V8.5F
dgb132.1004.0.2.0	\$1\$DGA132	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 displayed by the console `show dev` command are those devices that have been assigned to a `wwid<n>` environment variable.

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

6.6.3 Setting the `bootdef_dev` Console Environment Variable

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.

Notes

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

After the base operating system has been installed, or after cluster software has been installed for a cluster member system, set the `bootdef_dev` console environment variable again to provide multiple boot paths (see Section 6.8).

It would do no good to set `bootdef_dev` for multiple boot paths now because the installation procedure overwrites the variable.

The following procedure is used for:

- The initial Tru64 UNIX installation before booting from CD-ROM.

- After a cluster member has been added to the cluster with `clu_add_member` (but before the member system is booted).

Note

You do not use this procedure after using `clu_create` to create the first cluster member. Before booting the first cluster member, you reset the `bootdef_dev` console environment variable to multiple boot paths.

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

1. Obtain the device name for the Fibre Channel units that you will boot from. Ensure that you choose the correct device name for the entity you are booting (Tru64 UNIX, cluster member system 2, and so on). They show up in the reachability display as shown in Example 6-3 with a **Yes** under the `connected` column. You can also use the `wwidmgr -show reachability` command to determine reachability. Example 6-4 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 6-3 and Example 6-4 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`
- `dga132.1002.0.1.0`
- `dga132.1004.0.1.0`
- `dga133.1002.0.1.0`
- `dga133.1004.0.1.0`
- `dgb131.1002.0.2.0`
- `dgb131.1004.0.2.0`
- `dgb132.1002.0.2.0`
- `dgb132.1004.0.2.0`
- `dgb133.1002.0.2.0`
- `dgb133.1004.0.2.0`

Note that 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 path(s) that show up as connected. Ensure that you set the `bootdef_dev` variable appropriately for the system and boot disk. For the example disk configuration, set `bootdef_dev` as follows:

- On the system where you are installing the Tru64 UNIX operating system (which will also be the first cluster member):

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

- For a system that has just been set up as the second (or subsequent) cluster member system:

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

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
```

or:

```
P00>>> show bootdef_dev
bootdef_dev      dga132.1002.0.1.0
```

Now you are ready to install the Tru64 UNIX operating system, or boot `genvmunix` on a cluster member system that has been added to the cluster with `clu_add_member`.

6.7 Install the Base Operating System

After reading the *TruCluster Server Software Installation* manual, and using the *Tru64 UNIX Installation Guide* as a reference, boot from the CD-ROM and perform a full installation of the Tru64 UNIX Version 5.0A operating system.

When the installation procedure displays the list of disks available for operating system installation as shown here, look for the identifier in the Location column. Verify the identifier from the table you have been preparing (see Table 6-2).

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)	dsk16	1.0	SCSI	HSG80	IDENTIFIER=132
4)	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 6-2).

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

After the installation is complete, reset the `bootdef_dev` console environment variable to provide multiple boot paths (see Section 6.8).

6.8 Resetting the `bootdef_dev` Console Environment Variable

If you set the `bootdef_dev` console environment variable to multiple paths in Section 6.6.3, the base operating system installation, `clu_create`, or `clu_add_member` 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 units that you will boot from (see Table 6-2). Ensure that you choose the correct device name for the entity you are booting (Tru64 UNIX, cluster member system 2, and so on).
2. Check the reachability display (Example 6-3) provided by the `wwidmgr -quickset` or the `wwidmgr -reachability` commands for the device names that can access the storage unit you are booting from. Check the `show dev` command output to ensure the device name may be assigned to the `bootdef_dev` console environment variable.

Notes

You should 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, you should 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 6-3, you could choose the following device names when setting the `bootdef_dev` console environment variable:

```
dga133.1001.0.1.0 1
dga133.1004.0.1.0 2
dgb133.1002.0.2.0 3
dgb133.1003.0.2.0 4
```

- 1 Path from host bus adapter A to controller B port 2
- 2 Path from host bus adapter A to controller A port 1
- 3 Path from host bus adapter B to controller A port 2
- 4 Path from host bus adapter B to controller B port 1

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

-
3. Set the `bootdef_dev` console environment variable to a comma separated list of several of the boot path(s) 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:

- For the base operating system:

```
P00>>> set bootdef_dev \
dga133.1001.0.1.0,dga133.1002.0.2.0,\
dgb133.1001.0.1.0,dgb133.1002.0.2.0
P00>>> init
      :
```

- For member system 1 boot disk:

```
P00>>> set bootdef_dev \
dga131.1001.0.1.0,dga131.1002.0.2.0,\
dgb131.1001.0.1.0,dgb131.1002.0.2.0
P00>>> init
      :
```


- For member system 2 boot disk:

```
P00>>> set bootdef_dev \
dga132.1001.0.1.0,dga132.1002.0.2.0,\
dgb132.1001.0.1.0,dgb132.1002.0.2.0
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 four is all that is guaranteed.

6.9 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 5.0A 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. Use 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
```

Note

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
```

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 (`hwmgr`), see `hwmgr(8)`.

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

6.10 Installing the TruCluster Server Software

This section covers the Fibre Channel specific procedures you need to execute before running `clu_create` to create the first cluster member or `clu_add_member` to add subsequent cluster members. It also covers the procedure you need to execute after running `clu_create` or `clu_add_member` before you boot the new cluster member into the cluster. Use the TruCluster Server *Software Installation* procedures in conjunction with this manual for the TruCluster Server software installation.

To install the TruCluster Server software, follow these steps:

1. On the system you installed the Tru64 UNIX operating system on, boot the system, and before you install the TruCluster Server software subsets, determine the `/dev/disk/dskn` values to use for cluster installation (see Section 6.9).
2. Initialize disklabels for all disks needed to create the cluster. For the example disks we are using, it is disks `dsk18 (/var)`, `dsk19 (/usr)`, `dsk20 [cluster root (/)]`, and `dsk21 (Quorum)`. For instance:

```
# disklabel -rw dsk20 HSG80
```
3. Install the TruCluster Server software subsets and run the `clu_create` command to create the first cluster member using the procedures in the TruCluster Server *Software Installation* manual. When `clu_create` terminates, do not reboot the system. Shut down the system and reset the `bootdef_dev` console environment variable to provide multiple boot paths to the member system boot disk before booting (see Section 6.8). Boot the first cluster member.
4. On the system you installed the Tru64 UNIX operating system on, run `clu_add_member` to add subsequent cluster members.

Note

The system you installed the Tru64 UNIX operating system on is already enabled to access all the member system boot disks. If you use another cluster member system, you need to use the `wwidmgr -quickset` command to set up the paths to the member system boot disk.

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. For member system 2 in the example configuration, it is the storage unit with UDID 132 (See Table 6-2):

```
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
```

- b. Set the `bootdef_dev` console environment variable to one reachable path (Yes in the connected column) to the member system boot disk (See Section 6.6.3):

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

- c. Boot `genvmunix` on the newly added cluster member system.
5. After the new kernel is built, do not reboot the new cluster member 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 (see Section 6.8).
 6. Repeat steps 4 and 5 for other cluster member systems.

6.11 Changing the HSG80 from Transparent to Multiple-Bus Failover Mode

You may be are using transparent failover mode with Tru64 UNIX Version 5.0A and TruCluster Server Version 5.0A and want to take advantage of the ability to create a no-single-point-of-failure (NSPOF) configuration, and the availability that multiple-bus failover provides over transparent failover mode.

If you are upgrading from Tru64 UNIX Version 4.0F or Version 4.0G and TruCluster Software Products Version 1.6 to Tru64 UNIX Version 5.0A and TruCluster Server Version 5.0A you may want to change from transparent failover to multiple-bus failover to take advantage of multibus support in Tru64 UNIX Version 5.0A and multiple-bus failover mode and the ability to create a NSPOF cluster.

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, it adds the connection to its table of known connections. The unit offset for the connection depends on the failover mode in effect at the time the connection is discovered. In transparent failover mode, host connections to port 1 default to an offset of 0; host connections on port 2 default to an offset of 100. Host connections on port 1 can see units 0 through 99; host connections on port 2 can see units 100 through 199.

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

If you change the failover mode from transparent failover to multiple-bus failover, the offsets in the table of known connections remain the same as if they were for transparent failover mode; the offset on port 2 remains 100. With an offset of 100 on port 2, a host cannot see units 0 through 99 on port 2. This reduces the availability. Also, if you have only a single HSG80 controller and lose the connection to port 1, you lose access to units 0 through 99.

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

Note

It would do no good to disconnect and then reconnect the cables, because once a connection is added to the table it 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.

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

1. Shut down the operating systems on all host systems that are accessing the HSG80 controllers you want to change from transparent failover to multiple-bus failover mode.

- At the HSG80, set multiple-bus failover as follows. Note that before putting the controllers in multiple-bus failover mode, you must remove any previous failover mode:

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

Note

Use the controller known to have the good configuration information.

- 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 Offset
!NEWCON49	TRU64_UNIX HOST_ID=1000-0000-C920-DA01	THIS	2	230813 ADAPTER_ID=1000-0000-C920-DA01	OL this	100
!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	100
!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
!NEWCON54	TRU64_UNIX HOST_ID=1000-0000-C920-DA01	OTHER	1	230813 ADAPTER_ID=1000-0000-C920-DA01	OL other	0
!NEWCON55	TRU64_UNIX HOST_ID=1000-0000-C920-EDEB	OTHER	2	230913 ADAPTER_ID=1000-0000-C920-EDEB	OL other	100
!NEWCON56	TRU64_UNIX HOST_ID=1000-0000-C920-DA01	OTHER	2	230813 ADAPTER_ID=1000-0000-C920-DA01	OL other	100
!NEWCON57	TRU64_UNIX HOST_ID=1000-0000-C921-09F7	THIS	2		offline	100
!NEWCON58	TRU64_UNIX HOST_ID=1000-0000-C921-09F7	OTHER	1		offline	0
!NEWCON59	TRU64_UNIX HOST_ID=1000-0000-C921-09F7	THIS	1		offline	0
!NEWCON60	TRU64_UNIX HOST_ID=1000-0000-C921-09F7	OTHER	2		offline	100
!NEWCON61	TRU64_UNIX HOST_ID=1000-0000-C921-086C	THIS	2	210513 ADAPTER_ID=1000-0000-C921-086C	OL this	100

```

!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

4. 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

```

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

- a. Use the `wwidmgr` 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 that each system will need to boot from. 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 (see Section 6.8).

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

6.12 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 5.0A and, although not needed to maintain worldwide name to target ID mappings, it may be used with TruCluster Server Version 5.0A 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.

6.12.1 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:
  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 :
    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 :
    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 :
    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 :
    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 :
    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 :
    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 :
    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) allowing the adapter to send Fibre Channel traffic into the fabric
 - Directory Server — The N_Port is the FC entity queried to determine who is present on the Fibre Channel fabric
- 3 A target ID of -1 (or -2) shows up for remote Fibre Channel devices that do not communicate using Fibre Channel protocol, the directory server, and F_Port.

Note

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

6.12.2 Using the `emxmgr` Utility Interactively

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

- Display the presence of KGPSA Fibre Channel adapters

- Display the target ID mappings for a Fibre Channel adapter
- Display the current Fibre Channel topology for a Fibre Channel adapter

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

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

Note

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

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**
#

Preparing ATM Adapters

The Compaq Tru64 UNIX operating system supports Asynchronous Transfer Mode (ATM). TruCluster Server supports the use of LAN emulation over ATM for client access.

This chapter provides an ATM overview, an example TruCluster Server cluster using ATM, an ATM adapter installation procedure, and information about verifying proper installation of fiber optic cables. See the Tru64 UNIX *Asynchronous Transfer Mode* manual for information on configuring the ATM software.

7.1 ATM Overview

In synchronous transfer methods, time-division multiplexing (TDM) techniques are used to divide the bandwidth into fixed-size channels dedicated to particular connections. If a system has nothing to transmit when its time slot comes up, that time slot is wasted. Also, if the system has lots of information to transmit, the system can only transmit when its turn comes up, even if other time slots are empty.

ATM eliminates the inefficiencies of TDM technology by sharing network bandwidth among multiple logical connections. Instead of dividing the bandwidth into fixed-size channels dedicated to particular connections, ATM uses the entire bandwidth to transmit a steady stream of fixed-size (53-byte) cells. Each cell includes a 5-byte header containing an address to identify the cell with a particular logical connection.

If a connection needs more bandwidth, it is allocated more cells. When a connection is idle, it uses no cells and consumes no bandwidth. This feature makes ATM the ideal technology for transferring voice, video, and data through private networks and across public networks.

ATM is a connection-oriented, cell-switching and multiplexing technology. Cells transit ATM networks by passing through ATM switches, which analyze information in the header to switch the cell to the output interface that connects the cell to the next appropriate switch as the cell proceeds to its destination.

The ATM switch acts as a hub in the ATM network. All devices are attached to an ATM switch, either directly or indirectly.

Most data traffic in existing customer networks is sent over Local Area Networks (LANs) such as Ethernet or Token Ring networks. The services provided by the LANs differ from those of ATM, for example:

- LAN messages are connectionless; ATM is a connection-oriented technology
- Because a LAN is based on a shared medium, it is easy to broadcast messages
- LAN addresses are based on hardware manufacturing serial numbers and are independent of the network topology

In order to use the large base of existing LAN application software, ATM defines a LAN Emulation (LANE) service that emulates services of existing LANs across an ATM network.

The LAN emulation environment groups hosts into an emulated LAN (ELAN) which has the following characteristics:

- Identifies hosts through their 48-bit media access control (MAC) number
- Supports multicast and broadcast services through point-to-multipoint connections or through a multicast server
- Supports any protocol that uses an IEEE broadcast LAN
- Provides the appearance of a connectionless service to participating end systems

One or more emulated LANs can run on the same ATM network. Each ELAN is independent of the others and users cannot communicate directly across emulated LAN boundaries. Communication between ELANs is possible only through routers or bridges.


Each ELAN is composed of:

- A set of LAN emulation clients (LECs): An LEC resides in each end system and performs data forwarding, address resolution, and control functions that provide a MAC-level emulated Ethernet interface to higher-level software and other entities within the emulated LAN.
- A LAN emulation service, which normally resides on an ATM switch and consists of:
 - LAN Emulation Configuration Server (LECS): An LECS implements the assignment of individual LAN emulation clients to different emulated LANs. It provides the client with the ATM address of the LAN emulation server.
 - LAN Emulation Server (LES): An LES implements the control coordination function for the emulated LAN by registering and resolving MAC addresses and route descriptors to ATM addresses.

- Broadcast and Unknown Server (BUS): A BUS handles broadcast data sent by a LAN emulation client, all multicast data, and data sent by a LAN emulation client before the ATM address has been resolved.

Figure 7–1 shows an ATM network with two emulated LANs. Hosts A and B are LECs on ELAN1. Hosts C, D, and E are LECs on ELAN2. The LECS, the LES, and the BUS are server functions resident on the ATM switch (even though they are shown separately).

Figure 7–1: Emulated LAN Over an ATM Network



Use LAN emulation over ATM in a TruCluster Server cluster for client system access (Memory Channel is the cluster interconnect).

7.2 Installing ATM Adapters

Warning

Some fiber optic equipment can emit laser light that can injure your eyes. Never look into an optical fiber or connector port. Always assume the cable is connected to a light source.

Note

Do not touch the ends of the fiber optics cable. The oils from your skin can cause an optical power loss.

Use the following steps to install an ATMworks adapter. See the *ATMworks 350 Adapter Installation and Service* guide for more information. Be sure to use the antistatic ground strap.

1. Remove the adapter extender bracket if the ATMworks 350 is to be installed in an AlphaServer 2100 system.
2. Remove the option slot cover from the appropriate PCI or TURBOchannel slot.
3. Install the adapter module.
4. Install the multimode fiber optics (SC connectors) cables as follows:
 - Remove the optical dust caps.
 - Line up the transmit cable connector with the transmit port and the receive cable connector with the receive port and insert the SC connectors. The ATMworks transmit port is identified by an arrow exiting a circle. The receive port is identified by an arrow entering a circle.

Listen for the click indicating that the connector is properly seated.

Note

Ensure that the bend radius of any fiber optic cable exceeds 2.5 cm (1 inch) to prevent breaking the glass.

When removing an SC connector, do not pull on the cable. Pull on the cable connector only.

To verify that the cables are connected correctly, see Section 7.3.

7.3 Verifying ATM Fiber Optic Cable Connectivity

The fiber optic cables from some suppliers are not labeled or color coded, and as the system and ATM switch may be separated by a great distance, verifying that the cables are connected correctly may be difficult.

The ATMworks adapters start sending idle cells when the ATM driver is enabled. The adapter sends idle cells even when no data is being sent. ATM switches provide an indication that they are receiving the idle cells.

To verify that the fiber optic cables are properly connected, follow these steps:

1. Verify that both the transmit and receive connectors are seated properly at both the ATM adapter and the ATM switch.
2. Verify that the following ATM subsets have been installed with this command:

```
# /usr/sbin/setld -i | grep ATM
```

- OSFATMBASE: ATM Commands
- OSFATMBIN: ATM Kernel Modules

Additionally, after the ATM subsets have been installed, verify that a new kernel has been built with the following kernel options selected (/sbin/sysconfig -q atm):

- Asynchronous Transfer Mode (ATM)
- ATM UNI 3.0/3.1 Signalling for SVCs
- LAN Emulation over ATM (LANE)

3. Enable the ATM driver with the following command:

```
# /usr/sbin/atmconfig up driver=driver_name
```

In the command, `driver_name` is `lta#` for the ATMworks 350. The number sign (#) is the adapter number.

To enable `lta0` to initiate contact with the network, enter the following command:

```
# /usr/sbin/atmconfig up driver=lta0
```

4. Check the ATM switch for an indication that it is receiving idle cells. The following table provides the indication for a few ATM switches. If you do not have one of these switches, check the documentation for your switch to determine how the switch indicates that it is cabled correctly.

ATM Switch	Indicator	Comments
Compaq GIGAswitch	PHY	Illuminated green LED indicates that the switch is receiving idle cells from the ATM adapter.
Bay Networks® Centillion100	En	Illuminated green LED indicates that the switch is receiving idle cells from the ATM adapter.
SynOptics® LattisCell 10114	Link	Illuminated green LED indicates that the switch is receiving idle cells from the ATM adapter.
CISCO® Systems LightStream 1010	TX	The switch starts transmitting data as soon as it receives idle cells. The green TX LED will flash on and off.
FORESystems® ForeRunner A S X200	TX	The Yellow TX LED will be on steady.

5. If you do not have an indication that confirms a correct cable connection, swap the transmit and receive connectors on one end of the cable and recheck the indicators.
6. If you still do not have a correct cable connection, you probably have a bad cable.

7.4 ATMworks Adapter LEDs

The ATMworks adapter has two LEDs that indicate the status of the adapter and its connections to the network, the Network LED, and the Module LED. The Network LED is labeled with a number sign (#) under the LED. The Module LED is labeled with an incomplete circle under the LED. The meaning of the LEDs is shown in Table 7-1.

Table 7-1: ATMworks Adapter LEDs

Network LED	Module LED	Description
Off	Off	PCI slot is not receiving power, or the ATMworks driver has not been loaded.
Off/Amber/Green	Green	ATMworks driver is loaded and the module is OK.
Amber	Amber	ATMworks adapter is in reset mode.
Off	Amber	The adapter diagnostics failed.
Green	Green/Off	A physical link connection has been made.
Amber	Green	There is no physical link connection.

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.

Notes

Section 8.6 and Section 8.7 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 8.11 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) 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 shown 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




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
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

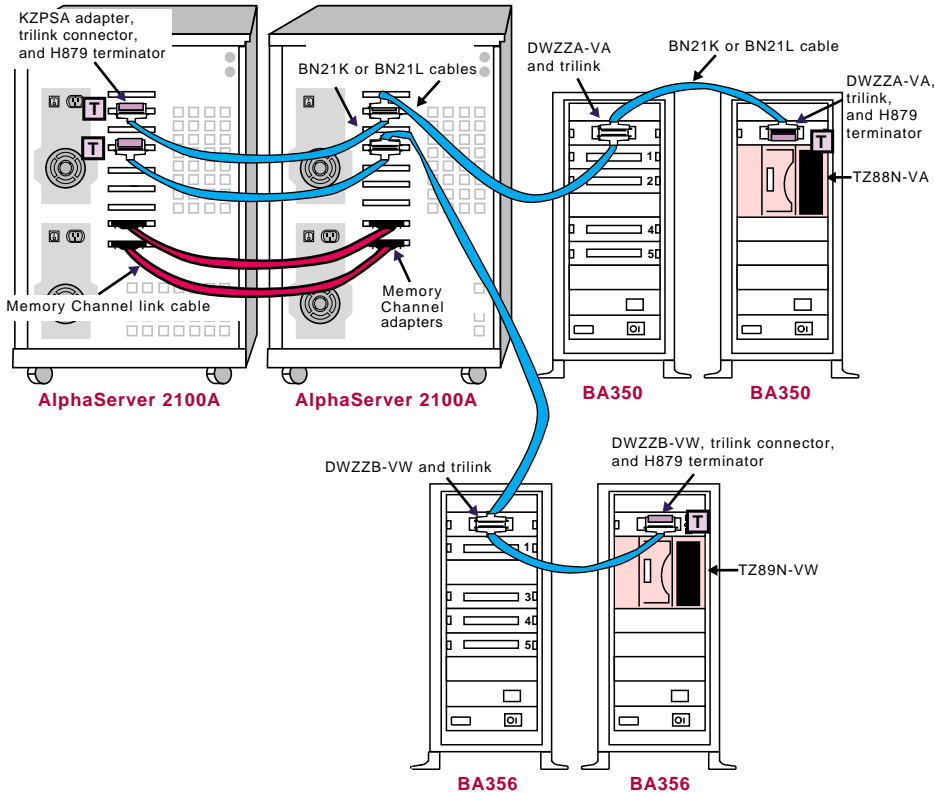
There are no special cabling restrictions specific to the TZ88N-VA; it is installed in a BA350 StorageWorks enclosure. A DWZZA-VA installed in slot 0 of the BA350 provides the connection to the shared SCSI bus. The tape drive takes up three slots, so two SCSI IDs are unavailable for disks in this StorageWorks enclosure. Another BA350 may be daisy chained to allow the use of the SCSI IDs unavailable in the first StorageWorks enclosure due to the TZ88 tape drive.

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 two shared SCSI buses. The top shared bus has a BA350 with disks at SCSI IDs 1, 2, 4, and 5. The other BA350 contains a TZ88N-VA at SCSI ID 3.

Figure 8–2: Shared SCSI Buses with SBB Tape Drives



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 8–6 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 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 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) 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 8–2 shows the switch settings for the DS-TZ89N-VW.

Table 8–2: 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

Table 8–2: DS-TZ89N-VW Switch Settings (cont.)

SCSI ID	SCSI ID Selection Switches							
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
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 as 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 trilink 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 trilink for that BA356 to provide termination for the shared SCSI bus.

Figure 8–2 shows a TruCluster Server cluster with two shared SCSI buses. The bottom shared bus has a BA356 with disks at SCSI IDs 1, 3, 4, and 5. The other BA356 contains a DS-TZ89N-VW at SCSI ID 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 8–6 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 or BN21L 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 or BN21L 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 or BN21L 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 capable of holding up to 40 GB of data per Compactape IV cartridge using 2:1 compression. It is capable of storing/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* (185292-002)
- *DLT Tape Drive User Guide Supplement* (340949-002)

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 to 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 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 maximum length (cable part number 328215-004 is a 20-meter

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 cables) as long as the SCSI bus maximum length of 3 meters (fast SCSI) is not exceeded. Ensure that the tape drive on the end of the bus is terminated with terminator part number 341102-001.

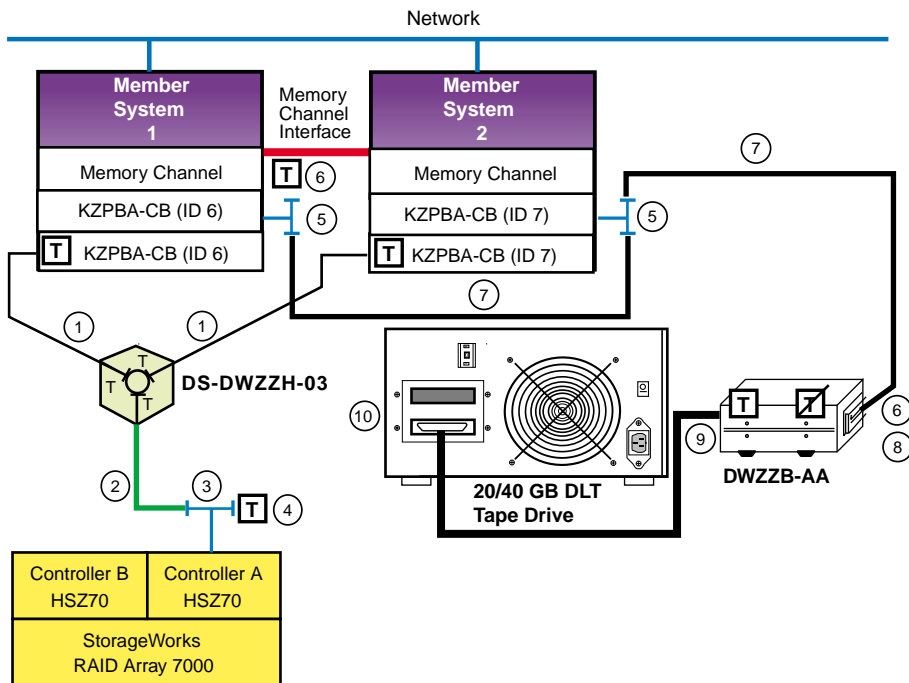
You can add additional shared SCSI buses with Compaq 20/40 GB DLT Tape Drives by adding additional DWZZB-AA/Compaq 20/40 GB DLT Tape Drive combinations.

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



NOTE: This drawing is not to scale.

ZK-1604U-AI

Table 8–3 shows the components used to create the cluster shown in Figure 8–5.

Table 8–3: 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 8–3: Hardware Components Used to Create the Configuration Shown in Figure 8 — 5 (cont.)

Callout Number	Description
9	199629-002 or 189636-002 (1.8-meter cable)
10	341102-001 terminator

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

^b The maximum length of the BN37A cable must not exceed 25 meters.

^c The maximum combined length of these cables must not exceed 25 meters.

8.4 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.

8.4.1 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 five 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.

8.4.2 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 8–6 shows a configuration of a TZ885

for use on a shared SCSI bus. The TZ885 in this figure has had the SCSI ID set to 0 (zero). 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 trilink 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 trilink or Y cable to the differential shared SCSI bus with BN21K or BN21L cables. Ensure that the trilink 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, system, and disk SCSI IDs.

Figure 8–6: Cabling a Shared SCSI Bus with a TZ885



8.5 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/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.

8.5.1 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 8–7). 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 8–7: TZ887 DLT MiniLibrary Rear Panel



ZK-1461U-AI

8.5.2 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 8–8 shows a configuration with a TZ887 for use on a shared SCSI bus. The TZ887 in this figure would have the SCSI ID set to 0. The member systems use SCSI IDs 6 and 7, and the disks are located in the BA356 slots at SCSI IDs 1-5.

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, host bus adapter, and disk SCSI IDs.

Figure 8–8: Cabling a Shared SCSI Bus with a TZ887



8.6 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.

There are six 68-pin, high-density SCSI connectors located on the back of the MiniLibrary; two SCSI connectors for each drive and two for the robotics controller. The TL891 uses a 0.3-meter 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 0.3-meter 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.

8.6.1 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 to 9 representing the cartridge slots. Rectangles present on this line indicate cartridges 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, an indication 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 once again.

2. Depress 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:) for which you wish to change the SCSI bus ID. 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 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.

8.6.2 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 0.3-meter 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 0.3-meter 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 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 0.3-meter 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 0.3-meter 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 0.3-meter 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 0.3-meter 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 8–9 shows an example of a TruCluster Server cluster with a TL892 connected to two shared SCSI buses.

Figure 8–9: TruCluster Server Cluster with a TL892 on Two Shared SCSI Buses



8.7 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.

8.7.1 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 seven inches (the height of each module) with a DS-TL800-AA pass-through mechanism extension. A seven-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.

8.7.2 Preparing the DLT MiniLibraries for Shared SCSI Bus Usage

The following sections describe how to prepare the DLT MiniLibraries in more detail. It is assumed that the expansion unit, base modules, and pass-through and motor mechanisms have been installed.

8.7.2.1 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 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 8-10 as 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 0.3-meter 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.

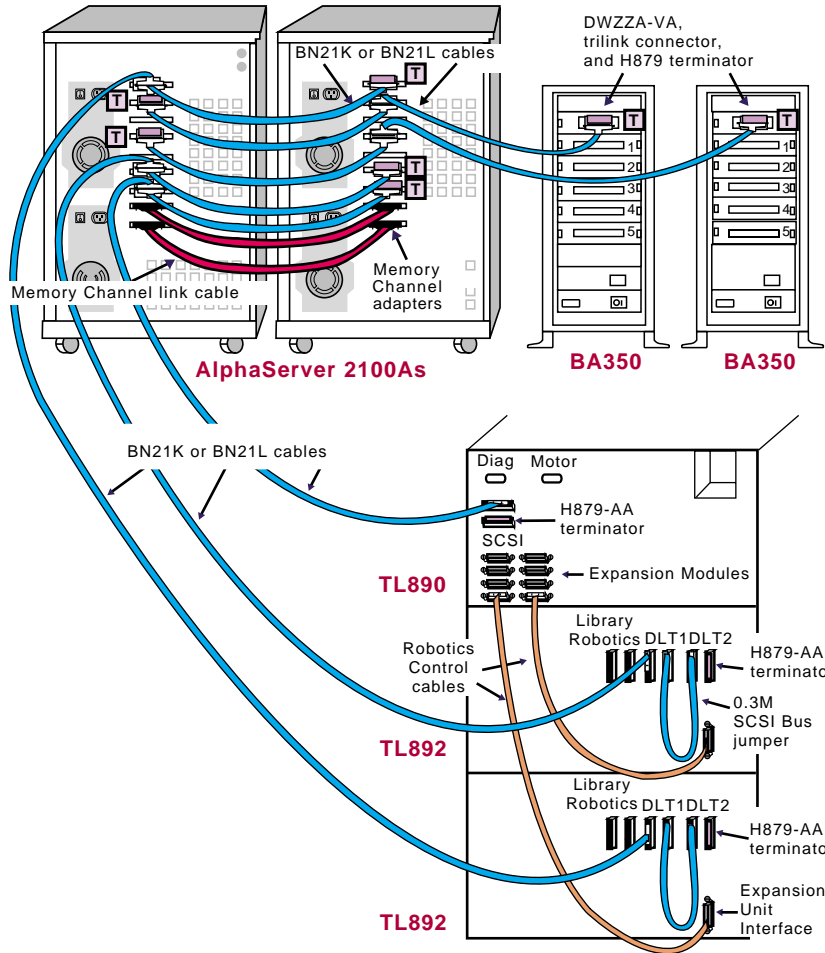
Notes

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 8–10 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. Note that the pass-through mechanism and cable to the library robotics motor is not shown in this figure.

Figure 8–10: TL890 and TL892 DLT MiniLibraries on Shared SCSI Buses



NOTE: This drawing is not to scale.

ZK-1398U-AI

8.7.2.2 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. Depress 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 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.

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.

8.7.2.3 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 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
```

8.7.2.4 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 wish 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 wish 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 you wish 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 you wish to configure.
11. If there are no more items to be configured, press the Escape button until the Default window is displayed.

8.8 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.

8.8.1 TL894 Robotic Controller Required Firmware

Robotic firmware Version S2.20 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 *TL81X/TL894 Automated Tape Library for DLT Cartridges Diagnostic Software User's Manual*.

8.8.2 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 8–4 lists the default SCSI IDs.

Table 8–4: TL894 Default SCSI ID Settings

SCSI Device	SCSI Address
Robotics Controller	0
Tape Drive 0	2
Tape Drive 1	3

Table 8–4: TL894 Default SCSI ID Settings (cont.)

SCSI Device	SCSI Address
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 8-4, 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.

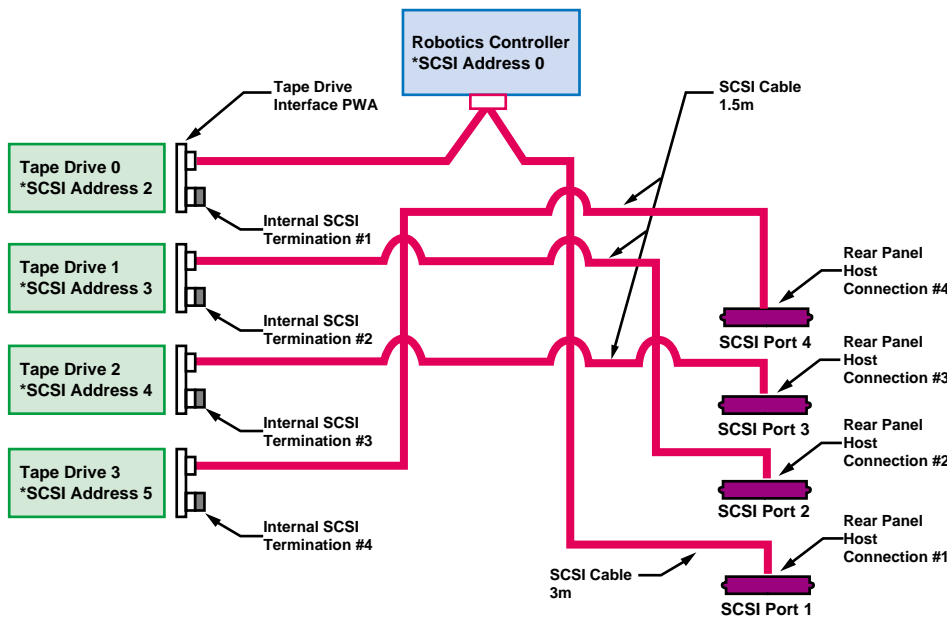
- Press and release the up or down button and the SELECT button simultaneously and verify that System On-line or System Off-line is displayed in the SDA.

8.8.3 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, called the four-bus configuration, is shown in Figure 8–11. 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 8–11: TL894 Tape Library Four-Bus Configuration



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) supplied in the accessories kit 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 do not recommend that you 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.

This is not wrong, but by reconfiguring in this manner, the length of the SCSI bus is increased by 1.5 meters, and may cause problems if SCSI bus length is of concern.

In a future revision of the previously mentioned guide, the bus configuration figures will be modified to show all SCSI buses terminated at the tape drives.

8.8.4 Connecting the TL894 Tape Library to the Shared SCSI Bus


The TL894 tape libraries have up to 3 meters of internal SCSI cabling per SCSI bus. Because of the internal SCSI cable lengths, it is not possible to 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, 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 trilink connectors or Y cables. However, the ASE must be shut down to remove a tape loader from the shared bus.

Figure 8–12 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 8.8.3 and Figure 8–11). 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 8–12, 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 8–12: Shared SCSI Buses with TL894 in Two-Bus Mode



NOTE: This drawing is not to scale.

ZK-1625U-AI

8.9 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 either 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* (EK-TL895-IG)
- *TL895 DLT Library Operator's Guide* (EK-TL895-OG)
- *TL895 DLT Tape Library Diagnostic Software User's Manual* (EK-TL895-UM)

For more information on upgrading from five to six or seven tape drives, see the *TL895 Drive Upgrade Instructions* manual.

Note

There are rotary switches on the library printed circuit board 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 ID 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.

8.9.1 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*.

8.9.2 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 8–5 lists the TL895 default SCSI IDs.

Table 8–5: 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 you wish 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:
 - Press the lock icon on the vertical bar of the control panel.
 - On the Password screen, press the User button.
 - A screen appears informing you that the new security level has been set. Press the Okay 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 affect.

8.9.3 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 8–13 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 8–13: TL895 Tape Library Internal Cabling



ZK-1397U-AI

8.9.4 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 8.9.2).

The actual upgrade is beyond the scope of this manual. See the *TL895 Drive Upgrade Instructions* manual for upgrade instructions.

8.9.5 Connecting the TL895 Tape Library to the Shared SCSI Bus

The TL895 tape library has up to 3 meters of internal SCSI cabling per SCSI bus. Because of the internal SCSI cable lengths, it is not possible to 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.

8.10 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.

8.10.1 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.

8.10.2 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 8–6.

Table 8–6: 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
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.

8.10.3 Setting the MUC SCSI ID

The multi-unit controller (MUC) SCSI ID is set with switch 1, 2, and 3, as shown in Table 8–7. 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 8–7: 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.

8.10.4 Tape Drive SCSI IDs

Each tape library arrives with default SCSI ID selections. The TL893 is shown in Table 8–8. The TL896 is shown in Table 8–9.

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 8–8: 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 8–9: 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

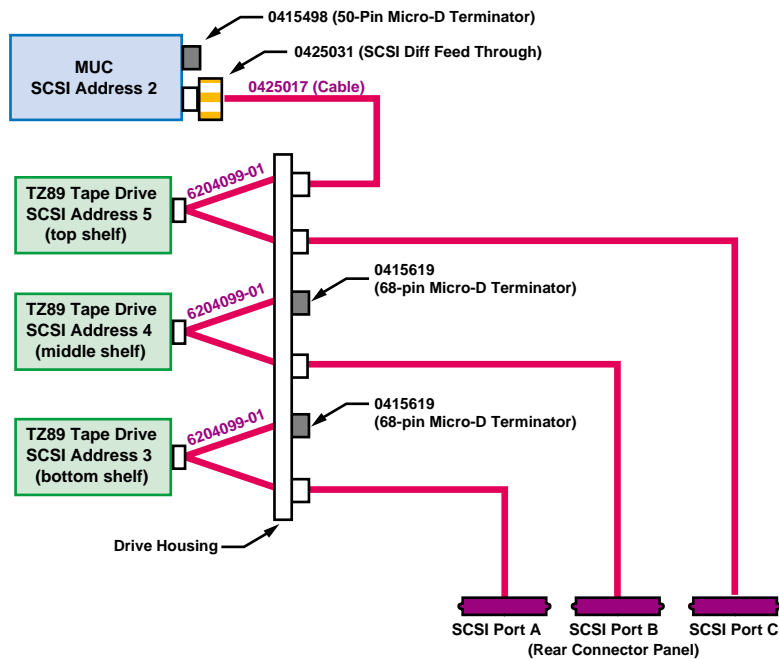
8.10.5 TL893 and TL896 Automated Tape Library Internal Cabling

The default internal cabling configurations for the TL893 and TL896 Automated Tape Libraries (ATLs) is 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 8–14 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 8–14 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 8–14: 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 8–15 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 8–15 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 8-15: TL896 Six-Bus Configuration



8.10.6 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 of internal SCSI cabling on each SCSI bus. Because of the internal SCSI cable lengths, it is not possible to 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 not place more than two drives on any one SCSI bus in these tape libraries.

Figure 8–16 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 8–16, 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 8–16: Shared SCSI Buses with TL896 in Three-Bus Mode



NOTE: This drawing is not to scale.

ZK-1626U-AI

8.11 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.

8.11.1 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.

8.11.1.1 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.

8.11.1.2 TL881 and TL891 MiniLibrary Rackmount Components

A TL881 or TL891 base unit (which contains the tape drive(s)) 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. Once 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 as 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 seven inches (the height of each module). A seven-inch gap may be left between modules (providing there is sufficient space), but additional pass-through mechanism extensions must be used.

8.11.1.3 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 8–10 shows the capacity and performance of a

TL881 or TL891 MiniLibrary in configurations set up for either maximum performance or maximum capacity.

Table 8–10: 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/hour)	1.32 TB (66 cartridges)	50 MB/sec (180 GB/hour)	2.31 TB (66 cartridges)
Capacity	1	4	3 MB/sec (10.8 GB/hour)	1.8 TB (90 cartridges)	10 MB/sec (36 GB/hour)	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 could be up to 40 GB/cartridge compressed.

^f Up to 5 MB/sec per drive.

^g Based on 35 GB/cartridge uncompressed. It could 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.

8.11.1.4 DLT MiniLibrary Part Numbers

Table 8–11 shows the part numbers for the TL881 and TL891 DLT MiniLibrary systems. Part numbers are only shown for the TL881 fast, wide differential components.

Table 8–11: 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

Table 8–11: DLT MiniLibrary Part Numbers (cont.)

DLT Library Component	Number of Tape Drives	Tabletop/Rackmount	Part Number
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.

8.11.2 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.

8.11.2.1 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*.

8.11.2.1.1 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 to 9 representing the cartridge slots. Rectangles present on this line indicate cartridges 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 once again.

2. Depress 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:) for which you wish to change the SCSI bus ID. 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 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.

8.11.2.1.2 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 8–17). A 0.3-meter 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 0.3-meter 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 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 trilink connector on the bus to the leftmost connector (as viewed from the rear) of the MiniLibrary. The 328215-004 is a 20-meter cable.
2. Install a 0.3-meter 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 0.3-meter 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 0.3-meter 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 0.3-meter 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 8–17 shows an example of a TruCluster configuration with a TL891 standalone MiniLibrary connected to two shared SCSI buses.

Figure 8–17: TL891 Standalone Cluster Configuration



ZK-1627U-AI

Table 8–12 shows the components used to create the cluster shown in Figure 8–17.

Table 8–12: 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 8–12: 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.

^b The maximum length of the BN37A cable must not exceed 25 meters.

^c The maximum combined length of these cables must not exceed 25 meters.

8.11.2.2 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 unit(s), 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.

8.11.2.2.1 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 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 8–18 as 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 a base unit or data unit 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 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 0.3-meter 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 8–18 shows a TL891 DLT MiniLibrary configuration with an expansion unit, a base units, 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 as it does not contain any tape drives but tape cartridges only. Note that the pass-through mechanism and cable to the library robotics motor is not shown in this figure.

For more information on cabling the units, see Section 8.11.2.1.2. 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 8–18: TL881 DLT MiniLibrary Rackmount Configuration



ZK-1628U-AI

Table 8–13 shows the components used to create the cluster shown in Figure 8–18.

Table 8–13: 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.

^b The maximum length of the BN37A cable must not exceed 25 meters.

^c The maximum combined length of these cables must not exceed 25 meters.

8.11.2.2.2 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. Depress 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 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.

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.

8.11.2.2.3 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 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.

8.11.2.2.4 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 wish 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 wish 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 you wish 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 you wish 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.

8.12 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.

8.12.1 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 and Installation Guide* (146585-001)
- *Reference Guide* (146583-001)
- *Maintenance and Service Guide* (155898-001)
- *Tape Drive Upgrade Guide* (146582-001)

Note

These tape devices have been qualified for use on shared SCSI buses with both the KZPSA-BB and KZPBA-CB host bus adapters.

8.12.2 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 six different part numbers, based on the number of tape drives:

Order Number	Number of Tape Drives
146205-B23	6
146205-B24	8
146205-B25	10
146205-B26	12
146205-B27	14
146205-B28	16

A tape library with a capacity for additional tape drives may be upgraded with part number 146209-B21, which adds a 35/70 DLT tape drive. See the Compaq StorageWorks ESL9000 Series Tape Library *Tape Drive Upgrade Guide* (146582-001) for more information.

8.12.3 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 Digital Linear Tape (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 are short on 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 recommend that storage not be placed on shared SCSI buses that have tape drives.

The following sections describe how to prepare the ESL9326D Enterprise Library in more detail.

8.12.3.1 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 V80 or later firmware.

8.12.3.2 Library Electronics and Tape Drive SCSI IDs


The default robotics and tape drive SCSI IDs are shown in Figure 8–19. 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* (146583-001).

8.12.3.3 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 8–19 shows the default cabling for an ESL9326D Enterprise Library with 16 tape drives. Note that 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 not present are not installed on the SCSI bulkhead.

Figure 8–19: ESL9326D Internal Cabling



ZK-1705U-AI

Note

Each internal cable is up to 2.5 meters 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, and the internal tape drive SCSI bus length is 5 meters. Therefore, you must limit the external SCSI bus cables to 20 meters maximum.

8.12.3.4 Connecting the ESL9326D Enterprise Library to the Shared SCSI Bus

The ESL9326D Enterprise Library has 5 meters of internal SCSI bus cabling for each pair of tape drives. Because of the internal SCSI bus lengths, it is not possible to use a tralink 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. Each SCSI bus must be terminated at the end of the SCSI bus by installing a terminator on the SCSI bulkhead SCSI connector. Therefore, TruCluster Server configurations using the ESL9326D Enterprise Library must ensure that the tape library is on the end of the shared SCSI bus.

Note

We recommend that disk storage devices be placed on separate shared SCSI buses.

Use 328215-001 (5-meter), 328215-002 (10-meter), 328215-003 (15-meter), 328215-004 (20-meter), or BN21K (BN21L) cables of the appropriate length to connect the ESL9326D Enterprise Library to a shared SCSI bus. Do not use a cable longer than 20 meters. 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 shown in Table 8–14 to form shared SCSI buses.

Table 8–14: 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 .3-meter 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 the customer is short on host bus adapters and wants to place more than two tape drives on a single SCSI bus (a configuration that we do not recommend).

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:

- Using SCSI bus signal converters (Section 9.1)
- SCSI bus termination in externally terminated TruCluster Server configurations (Section 9.2)
- Overview of the BA350, BA356, and UltraSCSI BA356 disk storage shelves (Section 9.3)
- Preparing the storage configuration for external termination using Y cables and trilinks (Section 9.4)
 - Preparing the storage shelves for an externally terminated TruCluster Server configuration (Section 9.4.1)
 - Connecting multiple storage shelves, for instance a BA350 and a BA356, two BA356s, or two UltraSCSI BA356s (Section 9.4.2)
 - Using the HSZ20, HSZ40, or HSZ50 RAID array controllers (Section 9.4.3)
- Radial configurations using the HSZ40 or HSZ50 RAID array controllers (Section 9.4.4)

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.

9.1 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 "bus expander" term 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.

9.1.1 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 could install a DWZZB-VW in an UltraSCSI BA356,

but you would waste a disk slot and it would 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.

9.1.2 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.


9.1.2.1 DWZZA and DWZZB Signal Converter Termination

Both the single-ended side and the differential side of each DWZZA and DWZZB signal converter has 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 9–1 shows the status of internal termination for a standalone SCSI signal converter that has a tralink connector attached to the differential side.


Figure 9–1: Standalone SCSI Signal Converter



ZK-1050U-AI

Figure 9–2 shows the status of internal termination for an SBB SCSI signal converter that has a trilink connector attached to the differential side.

Figure 9–2: SBB SCSI Signal Converter



ZK-1576U-AI

9.1.2.2 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 9.3.2.2 for information on how to select the device SCSI IDs in an UltraSCSI BA356.

Figure 9-3 shows the relative positions of the two DS-BA35X-DA switch packs.

Figure 9-3: DS-BA35X-DA Personality Module Switches



ZK-1411U-AI

9.2 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. Note that you may use external termination with an UltraSCSI hub, but it is not the recommended way.

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 9.4.3.

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 trilink connectors to connect a device to a shared SCSI bus.

By attaching a Y cable or trilink 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 trilink connector to terminate the bus. For UltraSCSI devices, attach an H8863-AA terminator to the H8861 trilink connector. If you disconnect the Y cable or trilink 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 trilink connector to a properly terminated shared bus without connecting the Y cable or trilink connector to a device. If you do this, you can connect a device to the Y cable or trilink

connector at a later time without affecting bus termination. This allows you to expand your configuration without shutting down the cluster.

Figure 9-4 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 trilink, and a trilink connector elsewhere.

Figure 9-4: BN21W-0B Y Cable





Figure 9-5 shows an HD68 trilink 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 9–5: HD68 Trilink Connector (H885-AA)



ZK-1140U-AI

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.

9.3 Overview of Disk Storage Shelves

The following sections provide an introduction to the BA350, BA356, and UltraSCSI BA356 disk storage shelves.

9.3.1 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 9-6 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 9–6: BA350 Internal SCSI Bus



ZK-1338U-AI

9.3.2 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.

9.3.2.1 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 9–7 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. Note that 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 9–7 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 9–7: BA356 Internal SCSI Bus




ZK-1339U-AI

Note that 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 9–8):

- 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 9–8: BA356 Jumper and Terminator Module Identification Pins



9.3.2.2 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 9–7. 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 the same as for the non-UltraSCSI

BA356, as shown in Figure 9–8. 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 9.1.2.2.

9.4 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 (DRD), service failover through the cluster application availability (CAA) subsystem, disk mirroring, and fast file system recovery. TruCluster Server supports mirroring of the clusterwide root (/) file system, the member-specific boot disks, and the cluster quorum disk through hardware RAID only. You can mirror the clusterwide /usr and /var file systems and the data disks 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.

After completing the following sections and setting up your RAID storagesets, you should be ready to cable your host bus adapters to storage when they have been installed (see Chapter 10).

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

If you need to use a BA350 or non-UltraSCSI BA356 with an UltraSCSI BA356 storage shelf, extrapolate the needed information from Section 9.4.1 and Section 9.4.2.

Later sections describe how to install cables to configure an HSZ20, HSZ40, or HSZ50 in a TruCluster Server configuration with two member systems.

9.4.1 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 would 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 would 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.

9.4.1.1 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 9.3.1 and Figure 9-6).

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.

9.4.1.2 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) or BN21L-01 (1-meter) 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 cable keeps the single-ended

SCSI bus (cable and BA356) under the 3-meter limit to still allow high speed operation.

If you are using a DWZZB-VW, install it in slot 0 of the BA356.

9.4.1.3 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 9.3.2.1, Figure 9–7, and Figure 9–8).
2. Set the SCSI bus ID switches on the UltraSCSI BA356 personality module (DS-BA35X-DA, Figure 9–3) 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 9–3).

Note

S4-3 and S4-4 are not used on the DS-BA35X-DA.

9.4.2 Connecting Storage Shelves Together

Section 9.4.1 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 9.4.2.1)
- Two BA356s (Section 9.4.2.2)
- Two UltraSCSI BA356s (Section 9.4.2.3)

9.4.2.1 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 9–9), follow these steps:

1. Complete the steps in Section 9.4.1.1 and Section 9.4.1.2 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 9–9 shows a two-member TruCluster Server configuration using a BA350 and a BA356 for storage.

Figure 9–9: BA350 and BA356 Cabled for Shared SCSI Bus Usage



ZK-1595U-AI

Table 9–1 shows the components used to create the cluster shown in Figure 9–9 and Figure 9–10.

Table 9–1: Hardware Components Used for Configuration Shown in Figure 8–9 and Figure 8–10

Callout Number	Description
1	BN21W-0B Y cable
2	H879-AA terminator
3	BN21K (or BN21L) cable ^a
4	H885-AA trilink connector

^a The maximum combined length of the BN21K (or BN21L) cables must not exceed 25 meters.

9.4.2.2 Connecting Two BA356s for Shared SCSI Bus Usage

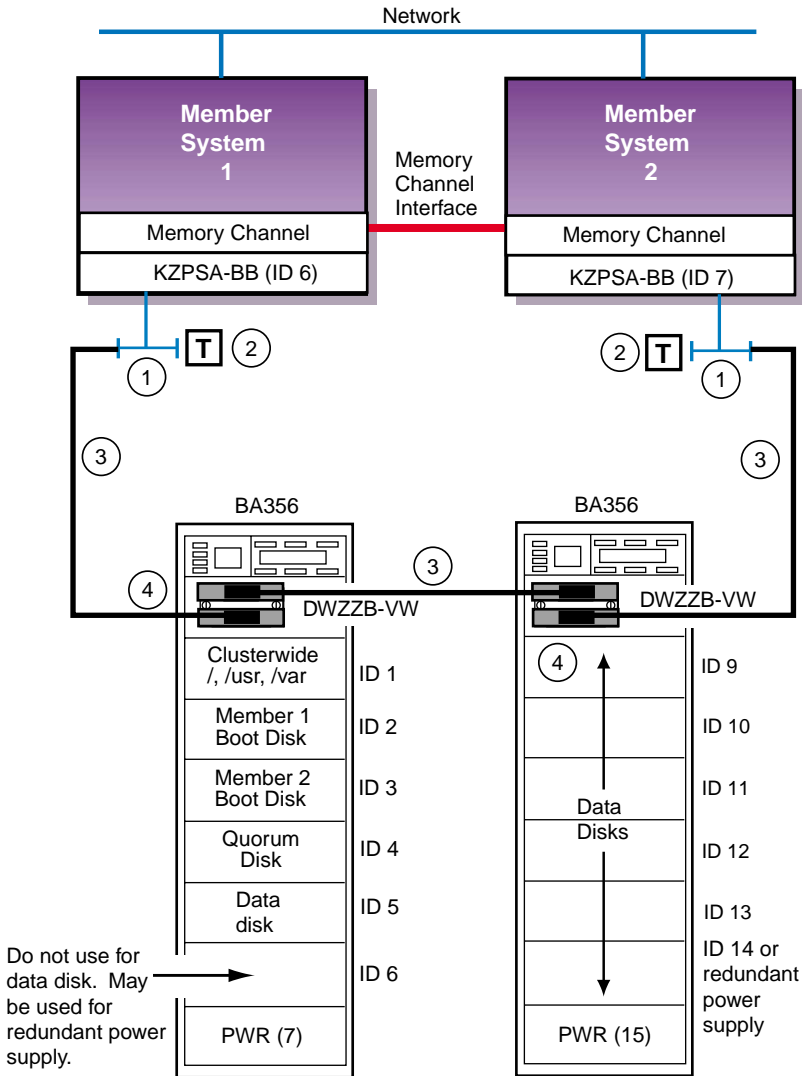
When you use two BA356 storage shelves on a shared SCSI bus in a TruCluster configuration, one BA356 must be configured for SCSI IDs 0 through 6 and the other configured for SCSI IDs 8 through 14.

To prepare two BA356 storage shelves for shared SCSI bus usage (see Figure 9–10), follow these steps:

1. Complete the steps of Section 9.4.1.2 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 9–10 shows a two member TruCluster Server configuration using two BA356s for storage.

Figure 9–10: Two BA356s Cabled for Shared SCSI Bus Usage



ZK-1592U-AI

Table 9–1 shows the components used to create the cluster shown in Figure 9–10.

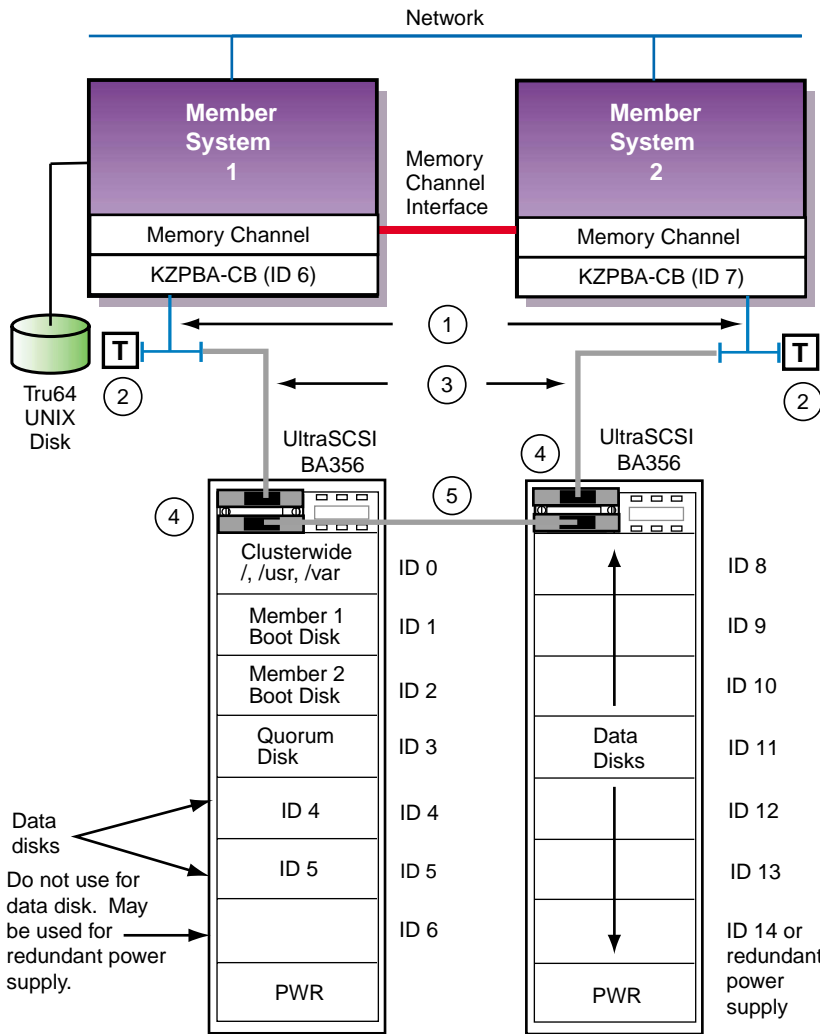
9.4.2.3 Connecting Two UltraSCSI BA356s for Shared SCSI Bus Usage

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 9–11) follow these steps:

1. Complete the steps of Section 9.4.1.3 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 trilink 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 trilink connectors. Install the trilink 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 trilink 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 trilink 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 trilink 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 trilink connector on the DS-BA35X-DA personality modules.

Figure 9–11: Two UltraSCSI BA356s Cabled for Shared SCSI Bus Usage



ZK-1598U-AI

Table 9–2 shows the components used to create the cluster shown in Figure 9–11.

Table 9–2: Hardware Components Used for Configuration Shown in Figure 9–11

Callout Number	Description
1	BN21W-0B Y cable
2	H879-AA HD68 terminator
3	BN38C (or BN38D) cable ^a
4	H8861-AA VHDCI tralink connector
5	BN37A cable ^a

^a The maximum combined length of the BN38C (or BN38D) and BN37A cables on one SCSI bus segment must not exceed 25 meters.

9.4.3 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.


9.4.3.1 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 trilink connector.
2. Attach an H885-AA trilink connector to each RAID controller port. Attach the H885-AA trilink 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 0.15-meter cable) between the H885-AA trilink 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 controller(s).
 - 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 9–12 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 9–12: Externally Terminated Shared SCSI Bus with Mid-Bus HSZ50 RAID Array Controllers

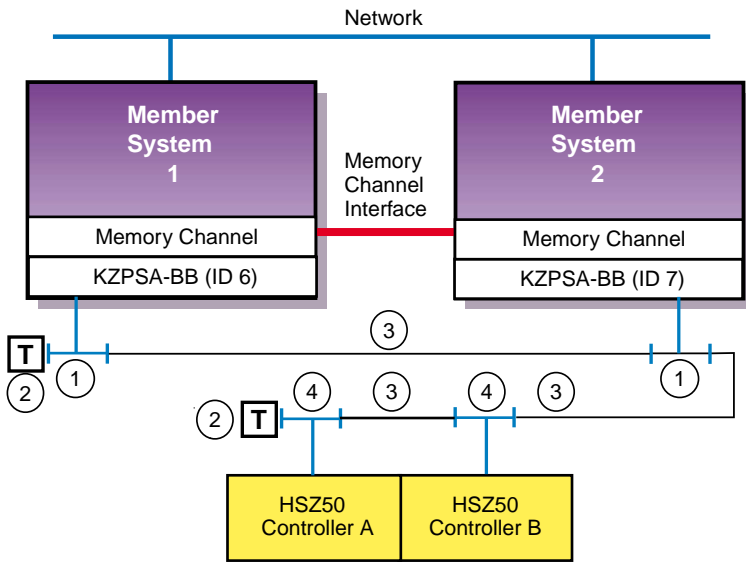


ZK-1596U-AI

Table 9–3 shows the components used to create the cluster shown in Figure 9–12 and Figure 9–13.

Figure 9–13 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 9–12, the SCSI bus adapters are KZPSA-BB PCI-to-SCSI adapters. They could be KZPBA-CB host bus adapters without changing any cables.

Figure 9–13: Externally Terminated Shared SCSI Bus with HSZ50 RAID Array Controllers at Bus End



ZK-1597U-AI

Table 9–3 shows the components used to create the cluster shown in Figure 9–12 and Figure 9–13.

Table 9–3: Hardware Components Used for Configuration Shown in Figure 8–12 and Figure 8–13

Callout Number	Description
1	BN21W-0B Y cable
2	H879-AA terminator
3	BN21K (or BN21L) cable ^{a b}
4	H885-AA tralink connector

^a The maximum combined length of the BN21K (or BN21L) cables must not exceed 25 meters.

^b The cable between the H885-AA tralink connectors on the HSZ50s could be a BN21L-0B, a 0.15-meter cable.

9.4.3.2 Cabling an HSZ20 in a Cluster using External Termination

To connect a SWXRA-Z1 (HSZ20 controller) to a shared SCSI bus, follow these steps:

- 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.

9.4.4 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 wish 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 0.15-meter 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 which 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.

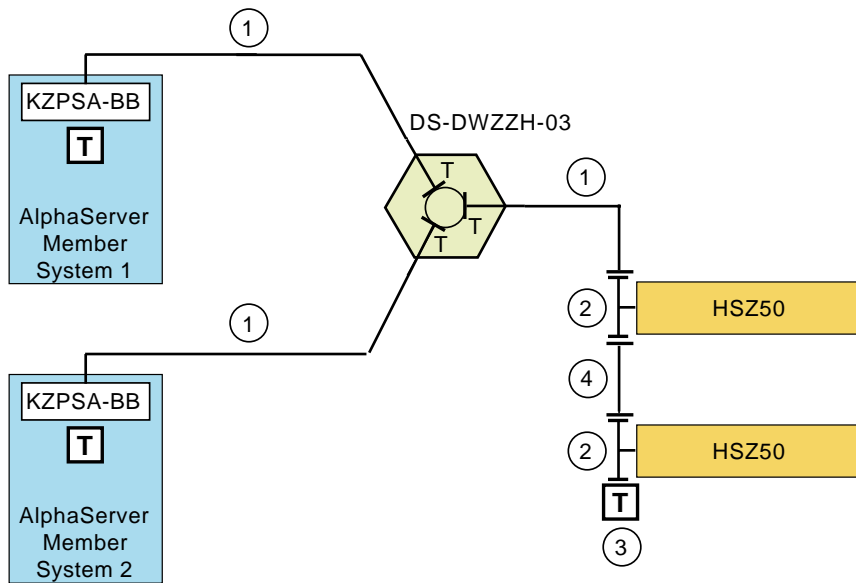
Note

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 9–14 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 9–14: TruCluster Server Cluster Using DS-DWZZH-03, SCSI Adapter with Terminators Installed, and HSZ50



ZK-1415U-AI

Table 9–4 shows the components used to create the cluster shown in Figure 9–14.

Table 9–4: Hardware Components Used in Configuration Shown in Figure 9–14

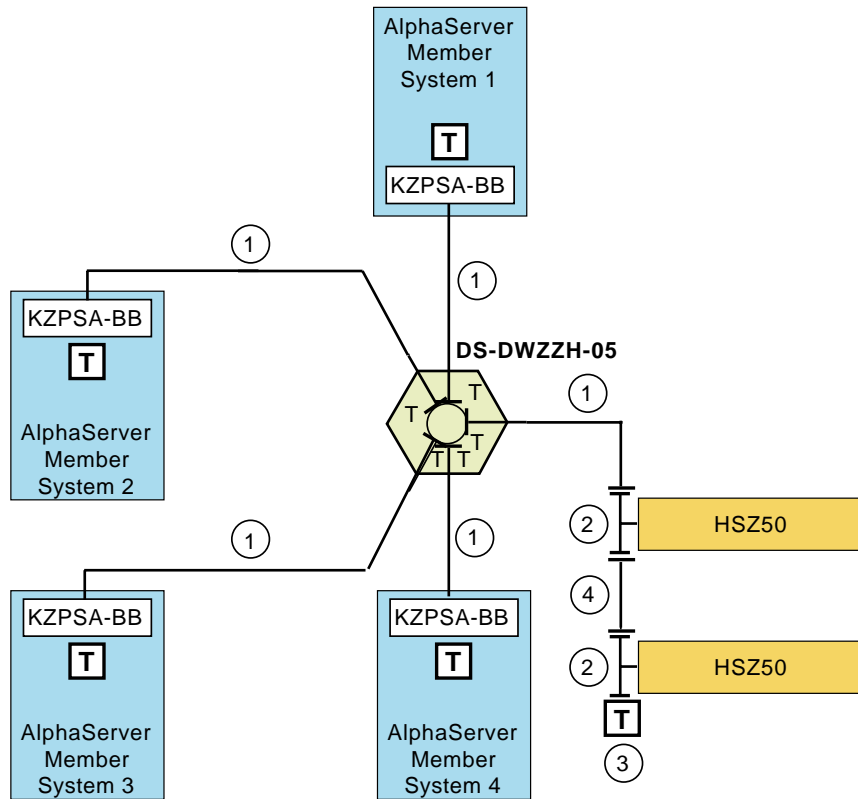
Callout Number	Description
1	BN38C cable ^{a b}
2	H885-AA HD68 tralink connector
3	H879-AA HD68 terminator
4	BN21K or BN21L cable ^b

^a The maximum length of the BN38C cable on one SCSI bus segment must not exceed 25 meters.

^b The maximum combined length of the BN38C and BN21K (or BN21L) cables on the storage SCSI bus segment must not exceed 25 meters.

Figure 9–15 shows a sample configuration that uses KZPSA-BB SCSI adapters, a DS-DWZZH-05 UltraSCSI hub, and an HSZ50 RAID array controller.

Figure 9–15: TruCluster Server Cluster Using KZPSA-BB SCSI Adapters, a DS-DWZZH-05 UltraSCSI Hub, and an HSZ50 RAID Array Controller



ZK-1449U-AI

Note

The systems shown in Figure 9–15 use KZPSA-BB SCSI adapters. They could be KZPBA-CB UltraSCSI adapters without changing any cables in the configuration.

Table 9–4 shows the components used to create the cluster shown in Figure 9–15.

10

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 for external termination or radial connection to non-UltraSCSI RAID array controllers (HSZ40 and HSZ50). 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. This documentation should arrive 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.

10.1 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.

It is assumed that you have already configured and cabled your storage subsystems as described in Chapter 9. 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 10–1 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.

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 10–1: Configuring TruCluster Server Hardware for Use with a PCI SCSI Adapter

Step	Action	Refer to:
1	Install the Memory Channel module(s), cables, and hub(s) (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 Chapter 7 and <i>ATMworks 350 Adapter Installation and Service</i>
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: Externally terminated host bus adapter:	Section 10.1.1 and Table 10–2 Section 10.1.2 Table 10–3

^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.

10.1.1 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 (t_{empwr}) 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 t_{empwr} 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.

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 2.4.1).
- 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 2.4.1.

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 10–2 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 10–2: 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 10.1.4.4, Figure 10–1, 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>

Table 10–2: Installing the KZPSA-BB or KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub (cont.)

Step	Action	Refer to:
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 are 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 cable between the KZPBA-BB or KZPBA-CB UltraSCSI host adapter and a DWZZH port.	
Notes		
The maximum length of a SCSI bus segment is 25 meters, 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 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 10.1.4.5
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.		
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 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 10.1.3 and Example 10–1 through Example 10–4

Table 10–2: Installing the KZPSA-BB or KZPBA-CB for Radial Connection to a DWZZH UltraSCSI Hub (cont.)

Step	Action	Refer to:
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 10.1.4.1 through Section 10.1.4.3 and Example 10–6 through Example 10–9
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. 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 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 9.4.4

10.1.2 Installing a KZPSA-BB or KZPBA-CB Using External Termination

Use the steps in Table 10–3 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 10–3: 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 10.1.4.4, Figure 10–1, 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 are 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 10.1.4.5
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 <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 display</code> and <code>isp</code> in the <code>show device display</code> to determine which devices are KZPBA-CBs.	Section 10.1.3 and Example 10–1 through Example 10–4

Table 10–3: 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 10.1.4.1 through Section 10.1.4.3 and Example 10–6 through Example 10–9
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 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 9.4
	BA350 storage shelf.	Section 9.3.1, Section 9.4.1.1, and Section 9.4.2.1
	Non-UltraSCSI BA356 storage shelf.	Section 9.3.2.1, Section 9.4.1.2, and Section 9.4.2.2
	Ultra BA356 storage shelf.	Section 9.3.2.2, Section 9.4.1.3, and Section 9.4.2.3
	HSZ40 or HSZ50 RAID array controller	Section 9.4.3
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 8.4
	TZ887	Section 8.5
	TL891/TL892 MiniLibrary	Section 8.6

Table 10–3: Installing a KZPSA-BB or KZPBA-CB for use with External Termination (cont.)

Step	Action	Refer to:
	TL890 with TL891/TL892	Section 8.7
	TL894	Section 8.8
	TL895	Section 8.9
	TL893/TL896	Section 8.10
	TL881/TL891 DLT MiniLibraries	Section 8.11
	Compaq ESL9326D Enterprise Library	Section 8.12

Notes

If you install tape devices on the shared SCSI buses, ensure that you understand how the particular tape device(s) 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.

10.1.3 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 10–1 shows the output from the `show config` console command on an AlphaServer 4100 system.

Example 10–1: 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
Memory 64 MB SYNC                     0         0000    mem2
Memory 64 MB SYNC                     0         0000    mem3
```

Example 10–1: Displaying Configuration on an AlphaServer 4100 (cont.)

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 10–2 shows the output from the `show device console` command entered on an AlphaServer 4100 system.

Example 10–2: 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 ispl (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
dkd2.0.0.4.1 DKd2 HSZ50-AX X29Z
```

Example 10–2: Displaying Devices on an AlphaServer 4100 (cont.)

```
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 10–3 shows the output from the `show config` console command entered on an AlphaServer 8200 system.

Example 10–3: 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 10–4 shows the output from the `show device console` command entered on an AlphaServer 8200 system.

Example 10–4: 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
```

10.1.4 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.

10.1.4.1 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 10–5 shows the `pk` console environment variables for an AlphaServer 4100.

Example 10–5: 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
```

Example 10–5: Displaying the pk* Console Environment Variables on an AlphaServer 4100 System (cont.)

```
pke0_termpr      1

pkf0_fast        1
pkf0_host_id     7
pkf0_termpr      1
```

Compare the `show pk*` command display in Example 10–5 with the `show config` command in Example 10–1 and the `show dev` command in Example 10–2. Note that there are no `pk*` devices in either display. Example 10–2 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 10–5 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 10–6 shows the use of the `show isp` console command to display the console environment variables for KZPBA-CBs on an AlphaServer 8x00.

Example 10–6: 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 10–3 and Example 10–4 show five `isp` devices; `isp0`, `isp1`, `isp2`, `isp3`, and `isp4`. In Example 10–6, 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 10–3 and Example 10–4 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 10–7 shows the use of the `show pk` console command to display the console environment variables for KZPSA-BBs on an AlphaServer 8x00.

Example 10–7: 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 10–7: Displaying Console Variables for a KZPSA-BB on an AlphaServer 8x00 System (cont.)

```
pkc0_fast      1
pkc0_host_id   7
pkc0_termprw  on
```

10.1.4.2 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 10–8 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 10–5.

Example 10–8: 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
```

10.1.4.3 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 10–9 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 10–9: 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
```

10.1.4.4 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 10-1.

Figure 10-1: KZPSA-BB Termination Resistors



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

10.1.4.5 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 Standard 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 wish 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.

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	User Define Iden- tifier (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

A

- ACS V8.5, 2-4
- arbitrated loop, 6-7
- ATL
 - TL893, 8-40, 8-41
 - TL896, 8-40, 8-41
- ATM
 - atmconfig command, 7-4
 - connecting cables, 7-4
 - installation, 7-1
 - LANE, 7-1
- atmconfig command, 7-4
- availability
 - increasing, 4-3

B

- BA350, 9-9
 - preparing, 9-15
 - preparing for shared SCSI usage, 9-15
 - termination, 9-3, 9-15
- BA356, 9-9
 - DS-DWZZH-03 installed in, 2-9, 3-9, 4-7, 9-29, 10-3
 - DS-DWZZH-05 installed in, 3-10
 - jumper, 9-9, 9-12
 - personality module address switches, 9-11
 - preparing, 9-15, 9-17
 - preparing for shared SCSI usage, 9-16
 - SCSI ID selection, 9-16
 - selecting SCSI IDs, 9-11
 - termination, 9-3, 9-9, 9-12

BA370

- DS-DWZZH-03 installed in, 2-9, 3-9
- bootdef_dev, 6-46, 6-48, 6-50, 6-51, 6-54
 - setting, 6-46, 6-48, 6-49, 6-51, 6-54
- bus hung message, 2-8
- bus_probe_algorithm, 2-6
- buses
 - data paths, 3-5
 - extending differential, 9-2
 - narrow data path, 3-5
 - speed, 3-5
 - terminating, 3-7, 9-5, 9-8
 - wide data path, 3-5

C

- cables
 - BC12N-10, 5-7
 - BN39B-04, 5-7
 - BN39B-10, 5-7
 - ESL9326D, 8-68
 - supported, 2-10
- cabling
 - Compaq 20/40 GB DLT Tape Drive, 8-10
 - DS-TZ89N-TA, 8-8
 - DS-TZ89N-VW, 8-7
 - ESL9326D, 8-65, 8-68
 - TL881/891 DLT MiniLibrary, 8-54, 8-58
 - TL890, 8-24
 - TL891, 8-20, 8-24

- TL892, 8-20, 8-24
- TL893, 8-47
- TL894, 8-34
- TL895, 8-40
- TL896, 8-47
- TZ885, 8-13
- TZ887, 8-16
- TZ88N-TA, 8-4
- TZ88N-VA, 8-3
- changing HSG80 failover modes, 6-55
- cluster
 - expanding, 3-7, 9-6
 - increasing availability, 4-3
 - planning, 4-2
- cluster interconnects
 - increasing availability, 4-2
- clusterwide file systems
 - allocating a disk for, 1-4
- command
 - atmconfig, 7-4
 - CONFIGURATION RESTORE, 6-33
 - emxmgr, 6-59
 - emxmgr -d, 6-59
 - emxmgr -m, 6-59
 - emxmgr -t, 6-60
 - init, 6-24, 6-45, 6-48, 6-50
 - SAVE_CONFIGURATION, 6-33
 - set bootdef_dev, 6-48, 6-50
 - SET FAILOVER COPY = THIS_CONTROLLER, 1-14
 - SET MULTIBUS_FAILOVER COPY = THIS_CONTROLLER, 3-18
 - show config, 4-9t, 10-5t, 10-7t
 - show device, 4-9t, 10-5t, 10-7t
 - SHOW THIS_CONTROLLER, 6-31
 - wwidmgr, 6-40
 - wwidmgr -clear, 6-41
 - wwidmgr -quickset, 6-42
 - wwidmgr -show, 6-24, 6-42, 6-45
- Compaq 20/40 GB DLT Tape Drive, 8-9

- cabling, 8-10
- capacity, 8-9
- cartridges, 8-9
- connectors, 8-9
 - setting SCSI ID, 8-9
- configuration
 - restrictions, 2-5
- CONFIGURATION RESTORE
 - command, 6-33
- configuring base unit as slave, 8-26, 8-61
- connections to HSG80, 6-55
- connectors
 - supported, 2-11
- console variable
 - bus_probe_algorithm, 2-6

D

- data path
 - for buses, 3-5
- default SCSI IDs
 - ESL9326D, 8-66
 - TL881/TL891, 8-53
 - TL890, 8-29
 - TL891, 8-29
 - TL892, 8-29
 - TL893, 8-43
 - TL894, 8-30
 - TL895, 8-37
 - TL896, 8-44
- device name, 6-40
- device unit number, 6-40
 - setting, 6-41
- diagnostics
 - Memory Channel, 5-12
- differential SCSI buses
 - description of, 3-4
- differential transmission
 - definition, 3-4
- disk devices
 - restrictions, 2-7
 - setting up, 3-16, 9-14
- disk placement

- clusterwide /usr, 1–10
- clusterwide /var, 1–10
- clusterwide root, 1–10
- member boot, 1–10
- quorum, 1–10
- disklabel, 6–53
- displaying device information
 - KZPBA-CB, 4–9t, 10–5t, 10–7t
 - KZPSA-BB, 10–5t, 10–7t
- DLT
 - Compaq 20/40 GB DLT Tape Drive, 8–9
 - TZ885, 8–13
 - TZ887, 8–15
- DLT MiniLibrary
 - Configuring TL881/TL891 as slave, 8–61
 - Configuring TL891 as slave, 8–26
 - TL881, 8–48
 - TL891, 8–48
- DS-BA356
 - DS-DWZZH-03 installed in, 2–9, 3–9, 4–7, 9–29, 10–3
 - DS-DWZZH-05 installed in, 3–10
- DS-BA35X-DA personality module, 3–3, 3–5, 4–8, 9–2, 9–3
- DS-DWZZH-03, 3–9
 - bus connectors, 3–9
 - bus isolation, 2–9
 - description, 2–9
 - installed in, 2–9, 3–9, 4–7, 9–29, 10–3
 - internal termination, 3–9
 - radial disconnect, 2–9
 - SBB, 3–9
 - SCSI ID, 3–9
 - support on, 3–9
 - termpwr, 3–9
 - transfer rate, 2–9
- DS-DWZZH-05, 3–9
 - bus connectors, 3–10
 - bus isolation, 2–9

- configurations, 3–15
- description, 2–9
- fair arbitration, 3–10
- installed in, 3–10, 3–11
- internal termination, 3–9
- radial disconnect, 2–9
- SBB, 3–10
- SCSI ID, 3–10
- termpwr, 3–9
- transfer rate, 2–9
- DS-TZ89N-TA
 - cabling, 8–8
 - setting SCSI ID, 8–8
- DS-TZ89N-VW
 - cabling, 8–7
 - setting SCSI ID, 8–5
- dual-redundant controllers, 1–14
- DWZZA
 - incorrect hardware revision, 2–8
 - termination, 9–3, 9–16
 - upgrade, 2–8
- DWZZB
 - termination, 9–3, 9–16
- DWZZH-03
 - (See DS-DWZZH-03)

E

- emxmgr, 6–59
 - displaying adapters, 6–59
 - displaying target ID mapping, 6–59
 - displaying topology, 6–60
 - use, 6–59, 6–61
 - using interactively, 6–62
- enterprise library
 - (See ESL9326D)
- environment variable
 - bootdef_dev, 6–46, 6–48, 6–49, 6–51, 6–54
 - N, 6–41
 - wwid, 6–41
- ESA12000

- configuring, 2–5
- port configuration, 2–5
- replacing controllers of, 6–32
- transparent failover mode, 2–5
- unit configuration, 2–5

ESL9000 series tape library
(See ESL9326D)

ESL9326D

- cables, 8–68
- cabling, 8–65, 8–68
- capacity, 8–65
- firmware, 8–66
- internal cabling, 8–67
- number of drives, 8–65
- part numbers, 8–65
- SCSI connectors, 8–68
- setting SCSI IDs, 8–66
- tape cartridges, 8–65
- tape drives, 8–65
- termination, 8–68
- upgrading, 8–65

F

- F_Port, 6–5
- fabric, 6–5
- failover mode
 - changing, 6–55
 - multiple-bus, 6–55
 - set nofailover, 6–56
 - transparent, 6–55
- Fibre Channel
 - arbitrated loop, 6–7
 - data rates, 6–4
 - distance, 6–4
 - F_Port, 6–5
 - fabric, 6–6
 - FL_Port, 6–5
 - frame, 6–4
 - N_Port, 6–5
 - NL_Port, 6–5
 - point-to-point, 6–6
 - restrictions, 2–3

- supported configurations, 6–8
- switch installation, 6–15
- terminology, 6–4
- topology, 6–5, 6–61

file

- /var/adm/messages, 6–26

firmware

- ESL9326D, 8–66
- obtaining release notes, 4–4
- reset system for update, 10–19
- update CD-ROM, 4–4
- updating, 10–18
- updating KZPSA, 10–18

FL_Port, 6–5

G

- GBIC, 6–16

H

- hardware components
 - Fibre Channel, 2–3
 - SCSI adapters, 2–6
 - SCSI cables, 2–10
 - SCSI signal converters, 2–8
 - storage shelves, 9–9
 - terminators, 2–11
 - trilink connectors, 2–11
- hardware configuration
 - bus termination, 3–7, 9–5
 - disk devices, 3–16, 9–14
 - hardware requirements for, 2–1
 - hardware restrictions for, 2–1
 - requirements, 3–1, 9–1
 - SCSI bus adapters, 2–6
 - SCSI bus speed, 3–5
 - SCSI cables, 2–10
 - SCSI signal converters, 9–2
 - storage shelves, 3–16, 9–14
 - supported cables, 2–1
 - supported terminators, 2–1
 - supported trilinks, 2–1

- supported Y cables, 2-1
- terminators, 2-11
- trilink connectors, 2-11
- host bus adapters
 - (*See* KGPSA, KZPBA-CB, KZPSA-BB)
- HSG80 controller
 - ACS, 2-4
 - changing failover modes, 6-55
 - configuring, 2-5, 6-26
 - multiple-bus failover, 6-28
 - obtaining the worldwide name of, 6-31
 - port configuration, 2-5
 - port_n_topology, 6-29
 - replacing, 6-32
 - resetting offsets, 6-55
 - setting controller values, 6-26, 6-28
 - transparent failover mode, 2-5
 - unit configuration, 2-5
- HSZ failover
 - multiple-bus, 1-15
 - transparent, 1-14
- HSZ20 controller
 - and shared SCSI bus, 9-24
- HSZ40 controller, 1-14
 - and shared SCSI bus, 9-24
- HSZ50 controller, 1-14
 - and shared SCSI bus, 9-24
- HSZ70 controller, 1-14
 - and fast wide differential SCSI, 3-3
- HSZ80 controller, 1-14
- hwmgr, 6-52

I

- I/O buses
 - number of, 2-6
- initialize, 6-24, 6-45, 6-48, 6-50
- installation, 3-16

- (*See also* hardware configuration)
- ATM adapter, 7-1
- KGPSA, 6-22
- KZPSA, 10-3
- MC2, 5-10
- MC2 cables, 5-9
- Memory Channel, 5-5
- Memory Channel cables, 5-7
- Memory Channel hub, 5-6
- optical converter, 5-6
- optical converter cables, 5-10
- switch, 6-15
- internal cabling
 - ESL9326D, 8-67
 - TL893, 8-44
 - TL896, 8-44

J

- jumpers
 - MC1 and MC1.5 (CCMAA), 5-2
 - MC2 (CCMAB), 5-3
 - MC2 (CCMLB), 5-5
 - Memory Channel, 5-2

K

- KGPSA
 - GLM, 6-23
 - installing, 6-22
 - mounting bracket, 6-22
 - obtaining the worldwide name of, 6-26
- KZPBA-CB
 - displaying device information, 4-9t, 10-5t, 10-7t
 - restrictions, 2-6
 - termination resistors, 4-9t, 10-4t, 10-7t
 - use in cluster, 4-6, 10-2

KZPSA-BB
displaying device information,
10-5t, 10-7t
installation, 10-3
restrictions, 2-6
setting bus speed, 10-17
setting SCSI ID, 10-17
setting termination power, 10-17
termination resistors, 10-4t, 10-7t
updating firmware, 10-18
use in cluster, 10-2

L

LAN emulation
(*See* LANE)
LANE, 7-1
LFU, 10-18
booting, 10-18
starting, 10-18
updating firmware, 10-18
link cable
installation, 5-7
Loadable Firmware Update utility
(*See* LFU)
Logical Storage Manager
(*See* LSM)
LSM
mirroring across SCSI buses, 1-11
mirroring clusterwide /usr, 1-12
mirroring clusterwide /var, 1-12
mirroring clusterwide data disks,
1-12

M

mc_cable, 5-12
mc_diag, 5-12
member systems
improving performance, 4-2
increasing availability, 4-2
requirements, 2-1
Memory Channel

diagnostics, 5-12
installation, 5-2, 5-5
interconnect, 2-3
jumpers, 5-2
versions, 2-2
Memory Channel diagnostics
mc_cable, 5-12
mc_diag, 5-12
Memory Channel hub
installation, 5-6
Memory Channel interconnects
restrictions, 2-2
setting up, 5-1
message
bus hung, 2-8
MiniLibrary
TL881, 8-48
TL891, 8-48
minimum cluster configuration, 1-5
MUC, 8-42
setting SCSI ID, 8-43
MUC switch functions
TL893, 8-42
TL896, 8-42
multi-unit controller
(*See* MUC)
multimode fibre, 6-16, 7-1
multiple-bus failover, 1-15, 3-18,
3-22, 6-28
changing from transparent failover,
6-55
NSPOF, 3-18
setting, 6-28, 6-56

N

N_Port, 6-5
NL_Port, 6-5
node name, 6-31
non-Ultra BA356 storage shelf
preparing, 9-15
NSPOF, 1-13, 3-18

O

- optical cable, 6–16
- optical converter
 - cable connection, 5–6
 - installation, 5–6

P

- part numbers
 - ESL9326D, 8–65
- partitioned storagesets, 3–18
- performance
 - improving, 4–2
- personality module, 3–3, 9–13
 - (*See also* signal converters)
- planning the hardware configuration, 4–2
- point-to-point, 6–6
- port name, 6–31
- powering up
 - TL881/891 DLT MiniLibrary, 8–62
- preparing storage shelves
 - BA350, 9–15
 - BA350 and BA356, 9–18
 - BA356, 9–16, 9–20
 - UltraSCSI BA356, 9–17, 9–21
- Prestoserve
 - cannot be used in a cluster, 4–3

Q

- quorum disk
 - and LSM, 1–5
 - configuring, 1–5
 - number of votes, 1–5

R

- RA8000
 - configuring, 2–5
 - port configuration, 2–5

- replacing controllers of, 6–32
- transparent failover mode, 2–5
- unit configuration, 2–5
- radial connection
 - bus termination, 3–8
 - UltraSCSI hub, 3–8
- RAID, 1–13
- RAID array controllers
 - advantages of use, 3–17
 - and shared SCSI bus, 9–24
 - preparing, 9–24
 - using in ASE, 9–24
- replacing
 - HSG80 controller, 6–32
- requirements
 - SCSI bus, 3–1, 9–1
- reset, 6–24, 6–45
- resetting offsets, 6–55
- restrictions, 2–5
 - disk devices, 2–7
 - KZPBA-CB adapters, 2–6
 - KZPSA adapters, 2–6
 - Memory Channel interconnects, 2–2
 - SCSI bus adapters, 2–6

S

- SAVE_CONFIGURATION
 - command, 6–33
- SC connector
 - (*See* subscriber connector)
- SCSI
 - number of devices supported, 3–2
 - SCSI bus with BA350 and BA356, 9–18
 - SCSI bus with Two BA356s, 9–20
 - SCSI bus with two UltraSCSI BA356s, 9–21
 - SCSI buses
 - (*See* shared SCSI buses)
 - SCSI cables

- (*See cables*)
- requirement, 2–10
- SCSI controllers
 - bus speed for, 3–5
- SCSI ID selection, 9–17
 - BA356, 9–16
- SCSI IDs
 - BA350, 9–9
 - BA350 storage shelves, 9–15
 - BA356, 9–11, 9–16
 - HSZ20 controller, 9–24
 - HSZ40 controller, 9–24
 - HSZ50 controller, 9–24
 - in BA356, 9–11
 - in UltraSCSI BA356, 9–13
 - RAID subsystem controllers, 9–24
 - requirement, 3–5
 - UltraSCSI BA356, 9–13, 9–17
- SCSI targets
 - number of, 2–5
- SCSI terminators
 - supported, 2–11
- SCSI-2 bus, 3–5
- SCSI-3, 6–4
- selecting BA356 disk SCSI IDs, 9–11
- selecting UltraSCSI BA356 disk SCSI IDs, 9–13
- set bootdef_dev, 6–48, 6–50, 6–54
- setting bus speed
 - KZPSA, 10–17
- setting SCSI ID
 - Compaq 20/40 GB DLT Tape Drive, 8–9
 - DS-TZ89N-TA, 8–8
 - DS-TZ89N-VW, 8–5
 - KZPSA, 10–17
 - MUC, 8–43
 - TL881/891 DLT MiniLibrary, 8–63
 - TL891, 8–18
 - TL892, 8–18
 - TL893, 8–43
 - TL894, 8–30
 - TL896, 8–43
 - TZ885, 8–13
 - TZ887, 8–15
 - TZ88N-TA, 8–4
 - TZ88N-VA, 8–2
- setting SCSI IDs
 - ESL9326D, 8–66
- setting the SCSI ID
 - TL881/891 DLT MiniLibrary, 8–53
- shared SCSI buses, 4–3
 - adding devices, 9–6
 - assigning SCSI IDs, 3–6
 - cable length restrictions, 3–6
 - connecting devices, 3–7, 9–6
 - device addresses, 3–5
 - differential, 3–4
 - number of, 2–6, 4–3
 - requirements, 3–2
 - single-ended, 3–4
 - using triline connectors, 9–6
 - using Y cables, 9–6
- shared storage
 - BA350 storage shelf, 9–15
 - increasing capacity, 4–2, 4–3
 - non-UltraSCSI BA356 storage shelf, 9–15
 - RAID subsystem array controller, 9–24
 - UltraSCSI BA356 storage shelf, 9–15, 9–17
- shortwave, 6–23
- SHOW THIS_CONTROLLER
 - command, 6–31
- signal converters, 9–2
 - creating differential bus, 9–2
 - differential I/O module, 9–2
 - differential termination, 9–3
 - DS-BA35X-DA personality module, 3–5, 9–4
 - extending differential bus length, 9–2
 - fast SCSI bus speed, 9–2
 - overview, 9–2
 - requirement, 9–2
 - restrictions, 2–8

- SBB, 9-2
- single-ended termination, 9-3
- standalone, 9-2
- terminating, 9-2
- termination, 9-3
- single-ended SCSI buses
 - description of, 3-4
- single-ended transmission
 - definition, 3-4
- storage shelves, 9-8, 9-9, 9-13
 - attaching to shared SCSI bus, 9-8, 9-13
 - BA350, 9-9
 - BA356, 9-9
 - overview, 9-8, 9-13
 - setting up, 3-16, 9-14
- subscriber connector, 6-16
- switch
 - 10Base-T Ethernet connection, 6-15
 - changing password, 6-21
 - changing user names, 6-21
 - front panel, 6-15, 6-18
 - GBIC, 6-16
 - installing, 6-15
 - interface module, 6-16
 - overview, 6-15
 - setting Ethernet IP address, 6-18
 - setting switch name, 6-21
 - telnet session, 6-21
- system reset, 6-24, 6-45

T

- table of connections, 6-55
- termination, 9-13
 - BA356, 9-11
 - DWZZA, 9-16
 - DWZZB, 9-16
 - ESL9326D, 8-68

- terminating the shared bus, 3-7, 9-5
 - UltraSCSI BA356, 9-14
- termination resistors
 - KZPBA-CB, 4-9t, 10-4t, 10-7t
 - KZPSA, 10-4t, 10-7t
 - KZPSA-BB, 10-7t
- terminators
 - supported, 2-11
- TL881, 8-48
- TL881/891 DLT MiniLibrary
 - cabling, 8-54, 8-58
 - capacity, 8-49, 8-51
 - components, 8-49
 - configuring base unit as slave, 8-61
 - models, 8-49
 - performance, 8-51
 - powering up, 8-62
 - setting the SCSI ID, 8-53, 8-63
- TL890
 - cabling, 8-24
 - default SCSI IDs, 8-29
 - powering up, 8-28
 - setting SCSI ID, 8-29
- TL891, 8-18, 8-48
 - cabling, 8-20, 8-24
 - configuring as slave, 8-26
 - default SCSI IDs, 8-20, 8-29
 - setting SCSI ID, 8-18, 8-19, 8-29
 - shared SCSI usage, 8-18
- TL892, 8-18
 - cabling, 8-20, 8-24
 - configuring as slave, 8-26
 - default SCSI IDs, 8-20, 8-29
 - setting SCSI ID, 8-18, 8-19, 8-29
 - shared SCSI usage, 8-18
- TL893, 8-40, 8-41
 - cabling, 8-44, 8-47
 - MUC switch functions, 8-42
 - setting SCSI ID, 8-43
- TL894
 - cabling, 8-34

- setting SCSI ID, 8–30
- TL895
 - cabling, 8–40
- TL896, 8–40, 8–41
 - cabling, 8–44, 8–47
 - MUC switch functions, 8–42
 - setting SCSI ID, 8–43
- transparent failover, 1–14, 3–17
 - changing to multiple-bus failover, 6–55
- trilink connectors
 - connecting devices with, 9–6
 - requirement, 2–11
 - supported, 2–11
- TZ88, 8–1
 - versions, 8–1
- TZ885, 8–13
 - cabling, 8–13
 - setting SCSI ID, 8–13
- TZ887, 8–15
 - cabling, 8–16
 - setting SCSI ID, 8–15
- TZ88N-TA, 8–1
 - cabling, 8–4
 - setting SCSI ID, 8–4
- TZ88N-VA, 8–1
 - cabling, 8–3
 - setting SCSI ID, 8–2
- TZ89, 8–5

U

- UltraSCSI BA356
 - disable termination, 9–17
 - DS-BA35X-DA personality module, 3–3
 - fast narrow SCSI drives, 3–3
 - fast wide SCSI drives, 3–3
 - jumper, 9–14
 - personality module address
 - switches, 9–13
 - power supply, 3–3
 - preparing, 9–15, 9–17

- preparing for shared SCSI usage, 9–17
 - SCSI ID selection, 9–17
 - selecting SCSI IDs, 9–13
 - termination, 9–14
- UltraSCSI host adapter
 - host input connector, 3–3
 - with non-UltraSCSI BA356, 3–3
 - with UltraSCSI BA356, 3–3
- UltraSCSI hubs, 3–8
- unshielded twisted pair
 - (*See* UTP)
- upgrade
 - DWZZA, 2–8
- upgrading
 - ESL9326D, 8–65
- utility
 - hwmgr, 6–52
 - wwidmgr, 6–45, 6–49
- UTP, 7–1

V

- /var/adm/messages, 6–26
- variable
 - (*See* environment variable)
- Very High Density Cable Interconnect
 - (*See* VHDCI)
- VHDCI, 3–3
 - acronym defined, 3–3
 - HSZ70 host connector, 3–3

W

- WorldWide ID Manager
 - (*See* wwidmgr)
- worldwide name
 - description, 6–25
- wwidmgr, 6–40
 - clear, 6–41
 - quickset, 6–42
 - show, 6–24, 6–42, 6–45

Y

Y cables

connecting devices with, 9-6
supported, 2-10

How to Order Tru64 UNIX Documentation

You can order documentation for the Tru64 UNIX operating system and related products at the following Web site:

<http://www.businesslink.digital.com/>

If you need help deciding which documentation best meets your needs, see the Tru64 UNIX *Documentation Overview* or call **800-344-4825** in the United States and Canada. In Puerto Rico, call **787-781-0505**. In other countries, contact your local Compaq subsidiary.

If you have access to Compaq's intranet, you can place an order at the following Web site:

<http://asmorder.nqo.dec.com/>

The following table provides the order numbers for the Tru64 UNIX operating system documentation kits. For additional information about ordering this and related documentation, see the *Documentation Overview* or contact Compaq.

Name	Order Number
Tru64 UNIX Documentation CD-ROM	QA-6ADAA-G8
Tru64 UNIX Documentation Kit	QA-6ADAA-GZ
End User Documentation Kit	QA-6ADAB-GZ
Startup Documentation Kit	QA-6ADAC-GZ
General User Documentation Kit	QA-6ADAD-GZ
System and Network Management Documentation Kit	QA-6ADAE-GZ
Developer's Documentation Kit	QA-6ADAG-GZ
Reference Pages Documentation Kit	QA-6ADAF-GZ

Reader's Comments

TruCluster Server

Hardware Configuration
AA-RHGWB-TE

Compaq welcomes your comments and suggestions on this manual. Your input will help us to write documentation that meets your needs. Please send your suggestions using one of the following methods:

- This postage-paid form
- Internet electronic mail: readers_comment@zk3.dec.com
- Fax: (603) 884-0120, Attn: UBPG Publications, ZKO3-3/Y32

If you are not using this form, please be sure you include the name of the document, the page number, and the product name and version.

Please rate this manual:

	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clarity (easy to understand)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization (structure of subject matter)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Figures (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Index (ability to find topic)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Usability (ability to access information quickly)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please list errors you have found in this manual:

Page	Description
_____	_____
_____	_____
_____	_____
_____	_____


Additional comments or suggestions to improve this manual:

What version of the software described by this manual are you using? _____

Name, title, department _____
Mailing address _____
Electronic mail _____
Telephone _____
Date _____

----- Do Not Cut or Tear - Fold Here and Tape -----

COMPAQ




NO POSTAGE
NECESSARY IF
MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 33 MAYNARD MA

POSTAGE WILL BE PAID BY ADDRESSEE

COMPAQ COMPUTER CORPORATION
UBPG PUBLICATIONS MANAGER
ZKO3-3/Y32
110 SPIT BROOK RD
NASHUA NH 03062-2698



----- Do Not Cut or Tear - Fold Here -----

Cut on This Line