TruCluster Software Products

Hardware Configuration

Part Number: AA-R88GA-TE

January 1998

Product Version:	TruCluster Production Server Software Version 1.5, TruCluster Available Server Software Version 1.5, and TruCluster MEMORY CHANNEL Software Version 1.5
Operating System and Version:	DIGITAL UNIX Version 4.0D

This manual describes how to configure the hardware for a TruCluster Software Products environment.

Digital Equipment Corporation
Maynard, Massachusetts

Contents

Figures

Tables

About This Manual

This manual describes how to set up and maintain the hardware configuration for the following TruCluster™ Software Products:

- TruCluster Production Server Software (Production Server, PS)
- TruCluster Available Server Software (Available Server, AS)
- TruCluster MEMORY CHANNEL™ (MEMORY CHANNEL, MC)

Audience

This manual is for system administrators who will set up and configure the hardware before installing one of the TruCluster Software Products kits. 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 seven chapters, a glossary, and an index. A brief description of the contents follows:

Chapter 1	Introduces the TruCluster products, describes available server environment (ASE) operation, and provides an overview of setting up TruCluster hardware.
Chapter 2	Describes the supported hardware, including some of the requirements for each hardware component.
Chapter 3	Contains information about setting up a shared SCSI bus and bus requirements.
Chapter 4	Describes how to prepare systems and storage shelves for a TruCluster configuration, and how to make shared SCSI bus connections.
Chapter 5	Describes how to set up the cluster MEMORY CHANNEL interconnect.
Chapter 6	Contains examples of Available Server and Production Server hardware configurations.

Chapter 7 Describes how to maintain the hardware and software.

Glossary Describes the TruCluster terms introduced in the documentation set.

Related Documents

See the following TruCluster manuals for assistance in cluster configuration, installation, and administration:

- TruCluster Software Products *Release Notes*—Documents known restrictions on the use of the TruCluster software products.
- TruCluster Software Products *Software Installation*—Describes how to install TruCluster Software products.
- TruCluster Software Products *Administration*—Describes administration tasks, such as those required to set up an ASE. It also shows how to configure, start, and manage distributed raw disk (DRD) services and other available services.

The following TruCluster manuals round out the TruCluster documentation:

- TruCluster Production Server Software *Application Programming Interfaces*—Describes the application programming interfaces (APIs) provided by the distributed lock manager (DLM) and cluster information services.
- TruCluster Production Server Software *MEMORY CHANNEL Application Programming Interfaces*—Describes the MEMORY CHANNEL APIs.

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.digital.com/golden-eggs`

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

In addition, you should have available the following manuals from the DIGITAL 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 Production Server:

- DECevent™
- Logical Storage Manager (LSM)
- POLYCENTER™ Advanced File System Utilities
- POLYCENTER NetWorker
- POLYCENTER Performance Manager
- POLYCENTER Console Manager

Online Documentation

Each book in the TruCluster Software documentation set is shipped as a set of Hypertext Markup Language (HTML) and graphics files in the `/TCR/doc/html` directory on the Associated Products Volume 2 CD-ROM. You can use the Netscape® Navigator™ World Wide Web browsing program to display these books.

If the DIGITAL UNIX installation program detects graphics capabilities on your system, it automatically installs Netscape Navigator. You can then invoke Netscape from an icon on the Common Desktop Environment (CDE) front panel or directly from the command line. Detailed help for Netscape is available through the help menus.

To access the TruCluster Software documentation from the viewer, click on the Open icon in the Netscape main window and enter the following file location in the Open Location: text entry field:

`file:/`*`mountpoint`*`/TCR/doc/html/BOOKSHELF.HTM`

Reader's Comments

DIGITAL welcomes any comments and suggestions you have on this and other DIGITAL 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:

  Digital Equipment Corporation
  UBPG Publications Manager
  ZKO3-3/Y32
  110 Spit Brook Road
  Nashua, NH 03062-9987

  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 DIGITAL UNIX that you are using.
- If known, the type of processor that is running the DIGITAL UNIX software.

The DIGITAL 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 DIGITAL technical support office. Information provided with the software media explains how to send problem reports to DIGITAL.

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.
`file`	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, `cat`(1) indicates that you can find information on the `cat` command in Section 1 of the reference pages.
cluster	Bold text indicates a term that is defined in the glossary.
PS	Abbreviation for the TruCluster Production Server Software.
AS	Abbreviation for the TruCluster Available Server Software.
MC	Abbreviation for the TruCluster MEMORY CHANNEL Software.

1

Introduction

This chapter provides an introduction to:

- Each of the following TruCluster Software Products:
  - TruCluster Available Server Software (Available Server, AS)
  - TruCluster Production Server Software (Production Server, PS)
  - TruCluster MEMORY CHANNEL Software (MEMORY CHANNEL, MC)
- The concept of using an **available server environment** (**ASE**) to provide highly available services and storage. An Available Server configuration consists of one ASE. A Production Server configuration must contain one ASE, and may contain up to four ASEs. A MEMORY CHANNEL configuration contains no ASEs.
- Hardware configuration concepts

The following chapters describe how to set up and maintain Production Server, Available Server, and MEMORY CHANNEL hardware configurations. See the TruCluster Software Products *Software Installation* manual for information about software installation, and the TruCluster Software Products *Administration* manual for detailed information about setting up member systems and services.

1.1 Describing the Suite of TruCluster Software Products

The DIGITAL UNIX TruCluster Software Products Version 1.5 suite of products allows multiple systems to be clustered as a single computing resource. The TruCluster product suite contains three products, as follows:

- TruCluster Available Server Software—The TruCluster Available Server Software (formerly known as DECsafe Available Server) significantly reduces down time due to hardware and software failures. Cluster activities are coordinated within an available server environment (ASE), an integrated organization of systems and external disks connected to shared SCSI buses that together provide highly available software and disk data to other systems.
- TruCluster MEMORY CHANNEL Software—The TruCluster MEMORY CHANNEL Software is an enabler for highly optimized applications that require high performance data delivery over the PCI-based MEMORY

CHANNEL interconnect. The MEMORY CHANNEL application programming interface (API) library provides APIs for MEMORY CHANNEL data channel and locking functions.

- TruCluster Production Server Software—The TruCluster Production Server Software (formerly known as TruCluster Software) combines the capabilities of the TruCluster Available Server Software and TruCluster MEMORY CHANNEL Software to provide a high-performance, highly available, and scalable operating environment.

  The TruCluster Production Server Software allows the processing components of an application to concurrently access raw devices, supporting highly parallelized database applications such as Oracle Parallel Server™ (OPS).

  The TruCluster Production Server Software **distributed lock manager** (**DLM**) synchronizes access to the resources that are shared among cooperating processes throughout the cluster. DLM provides services to applications to enforce a resource-sharing policy based on mutual exclusion or restricted sharing, plus services that notify the process that owns a resource that is blocking another resource requesting the resource. DLM allows an application to synchronize clusterwide access to shared resources, ensuring data integrity.

  The **distributed raw disk** (**DRD**) service allows a disk-based, user-level application to run within the cluster, regardless of where in the cluster the physical storage it depends upon is located. A DRD service allows applications such as distributed database systems or transaction processing systems (TPSs) parallel access to storage media from multiple cluster member systems

Each of the products in the TruCluster product suite provides a different aspect of TruCluster functionality. While the TruCluster Production Server Software offers the benefits of all three products, the TruCluster Available Server Software and the TruCluster MEMORY CHANNEL Software can also be used as standalone products in situations where only partial TruCluster functionality is required.

1.2 Overview of an Available Server Environment

Many applications, such as database applications and system services (for example, exported file systems and mail), are critical to business operations. Therefore, it is desirable to provide consistent and uninterrupted access to software or data on disks in a network environment. However, software or data on disks can become unavailable due to scheduled or unscheduled downtime for the system that is providing the resources. For example:

- System management operations such as maintenance, installations, and backups can prevent a system from providing data and applications.

- System crashes, although infrequent, interrupt network access to the disks and applications that the crashed system served.
- Network and I/O failures make data and applications unavailable.

An available server environment (ASE) is an integrated organization of systems and disks located on a **shared SCSI bus** that together provide highly available software and data to client systems. An ASE makes applications and data highly available, and can significantly reduce down time due to hardware and software failures.

In an ASE, the software provides multihost access to shared **SCSI** disks or tape and a generic failover mechanism for network-based applications and system services. In addition, the **POLYCENTER Advanced File System** (**AdvFS**), and the **Logical Storage Manager** (**LSM**) can be used within an ASE to provide fast file system recovery and high disk availability and reliability with disk mirroring.

In an ASE, you set up services for the applications or disk data that you want to make highly available. The TruCluster software failover mechanism makes applications and disk data independent of the availability of any one particular system. The applications are installed on each member system, and the disks are shared, so any member system can run an application and access data. This enables clients to have virtually uninterrupted access to resources.

When you create a service in an ASE, at a minimum, the `asemgr` utility prompts you for the following information:

- A unique name
- Either an application or storage configuration to make highly available
- Information about which member systems you want to run the service and how you want the service to behave when a failure occurs

Only certain types of applications can be made highly available with an ASE service. The application must have the following characteristics:

- The application must run on only one system at a time.
- The application must be able to be started and stopped using a set of commands that are performed in a specific order. When you set up a service, these commands are included in a set of programs called **action scripts**. Available Server software uses action scripts to fail over the services in the ASE.

The following sections descuss the ASE environment in more detail.

1.2.1 Types of ASE Services

TruCluster software provides support for various types of ASE services:

- Network File System (NFS) service—Enables you to provide highly available access to exported disk data (for example, mounted disks or system mail). When you create an NFS service, you must specify a unique Internet Protocol (IP) host name for the service name and the UNIX file systems, Advanced File System (AdvFS) filesets, or Logical Storage Manager (LSM) volumes that you want to export.

  You can also use your own action scripts in an NFS service if you want to fail over an application in addition to data. The member system that runs the service responds to the IP address that is assigned to the NFS service name and exports the service data. If the service is relocated to another member system, the new member system responds to the IP address. Clients are unaware of the change in the system that exports the data, and they experience only a temporary NFS server time out.

- Disk service—Enables you to provide highly available access to disks or a disk-based application, such as a database program. A disk service is similar to an NFS service except that no data is exported. When you create a disk service, you must specify the UNIX file systems, AdvFS filesets, or LSM volumes that you want to make highly available.

  You can also use your own action scripts in a disk service if you want to fail over an application in addition to data. Optionally, you can specify a unique IP host name for the service name, so that the member system that runs the service responds to the IP address. If the service is relocated to another member system, the new member system responds to the IP address.

- User-defined service—Enables you to provide highly available access to an application that is not disk based (for example, a login service). You must use your own action scripts in a user-defined service to fail over the application.

- Distributed raw disk (DRD) services allow a disk-based, user-level application to run within a cluster, regardless of where in the cluster the physical storage on which it depends is located. A DRD service allows an application, such as a distributed database system or transaction processing (TP) monitor, parallel access to storage media from multiple cluster members. Applications that perform I/O involving sets of large data files, random access to records within these files, and concurrent read/write data sharing can benefit from using the features of DRD. As deployed within an ASE, a DRD service can survive failures of both the server system and any mirrored disk participating in the service.

  A DRD service is applicable to the TruCluster Production Server Software product only.

- Tape service—The generic ASE tape service adds the capability to support NetWorker server failover with the TruCluster Production Server Software and TruCluster Available Server Software products. The tape service depends upon a set of one or more tape devices. There may be media changer devices and file systems associated with the tape devices. The service enables the user to configure the NetWorker server or other tape-based applications for failover.

  A user may perform existing tape applications on tape device on a shared SCSI bus. However, the user should be aware that an ASE tape service operation or system reboot of any system in the cluster on the shared SCSI bus could cause a SCSI reset, resulting in the premature termination of the user's job.

  Therefore, do not perform tape applications on tape devices on a shared SCSI bus unless you know exactly what is going on within the ASE.

1.2.2 ASE Services as They Appear to Clients

Figure 1–1 shows how the services in the ASE appear to clients.

Figure 1–1: Client View of the Available Server Environment

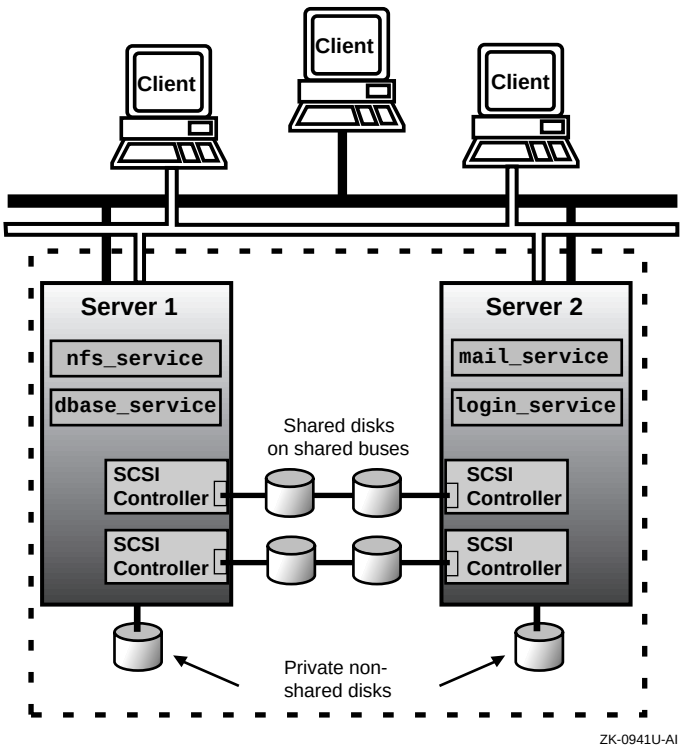

As shown in Figure 1–1, clients use the service names `nfs_service`, `dbase_service`, `mail_service`, and `login_service` to access the ASE services through a network.

An NFS service can use only disk data, or it can use applications and disk data. For example, the `nfs_service` service consists of only NFS disk data, but the `mail_service` service consists of NFS disk data and the `sendmail` command.

A disk service can also use only disk data, or it can use applications and disk data. For example, the `dbase_service` service consists of both disk data and a database application.

A user-defined service can use only applications. For example, the `login_service` service consists of only the `ifconfig` program.

A problem on Server 1 might cause the `nfs_service` and `dbase_service` to fail over to Server 2. The clients would notice a short disruption of those services, but they would once again be available when started on Server 2.

1.2.3 Comparing ASEs Between Available Server and Production Server

Both the TruCluster Available Server Software and TruCluster Production Server Software consist of the daemons, scripts, user interface, and drivers necessary to provide ASE functionality. Therefore, an ASE functions pretty much the same in both products.

With the TruCluster Available Server Software product, there can be only one ASE. All the member systems must be connected to the same shared SCSI bus (or buses).

With the TruCluster Production Server Software product, there must be at least one ASE; however, there may be multiple ASEs (a maximum of four), with between two and four systems in each ASE. There can also be up to eight systems in a cluster, but cluster members do not need to be ASE members, although they must be connected to other cluster members with the MEMORY CHANNEL interconnect.

Figure 1–2 shows a four-member configuration with one ASE. Note that all four systems are connected to both shared SCSI buses. This figure can represent either an Available Server or Production Server configuration. With a Production Server configuration, the cluster interconnect is the MEMORY CHANNEL, and with the four-member system configuration shown, there would have to be a MEMORY CHANNEL hub (which is not shown).

Figure 1–2: Hardware Configuration with One ASE

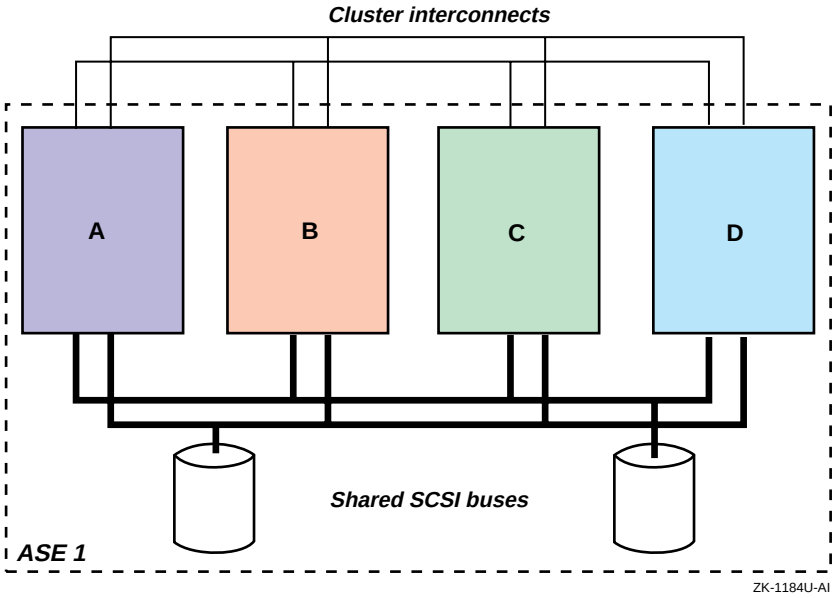

The TruCluster Production Server Software product allows multiple ASEs. Each ASE is numbered individually. All member systems in an ASE must be connected to the same shared SCSI buses. Figure 1–3 shows a four-member Production Server configuration with two ASEs. In this configuration, member systems A and B are in ASE 1 and systems C and D are in ASE 2. The cluster interconnect is the MEMORY CHANNEL and the MEMORY CHANNEL hub, which is not shown.

Figure 1–3: Hardware Configuration with Two ASEs

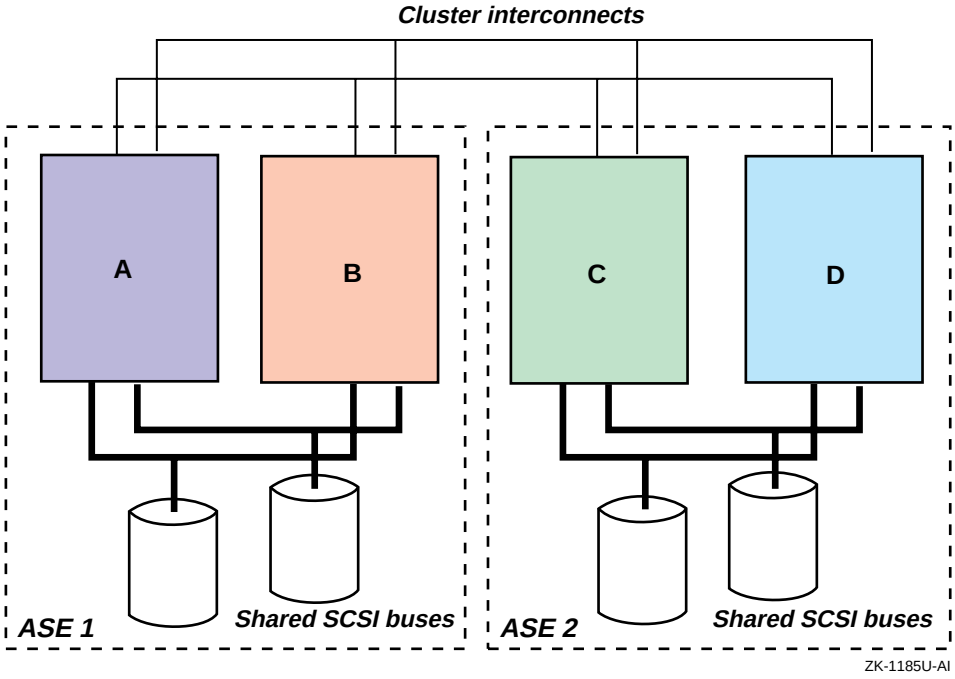

1.3 Overview of the TruCluster Available Server Hardware Configuration

Before you install the TruCluster Available Server software and set up the members and ASE services, you must set up the hardware. The TruCluster Available Server hardware configuration consists of a number of highly specific hardware components needed to support an available server environment (ASE). In addition, you must adhere to a number of configuration requirements for the SCSI bus configuration and device **firmware**. If you fail to adhere to these requirements, ASE operation will be impaired.

The hardware components needed to support TruCluster Available Server are as follows:

- Member systems—The member systems are the basic computing resources in a TruCluster Available Server configuration. The configuration can include from two to four member systems. The member systems run applications and provide access to data, and must be connected to at least one shared SCSI bus and one common network. The member systems communicate with each other and monitor the

shared devices and the network through the bus and through the network. If a hardware or software failure prevents a member system from running an application or providing data, the Available Server failover mechanism relocates resources to a viable system, which provides high availability.

- SCSI controllers—Each member system must have at least one supported SCSI controller installed in an I/O bus slot to connect the system to a shared bus. The use of multiple shared SCSI buses allows data to be mirrored across SCSI buses.
- Shared SCSI bus—Member systems must be connected to at least one shared SCSI bus, and a storage shelf is usually connected to a shared bus. Shared SCSI buses must have the same logical number on each member system, must be properly terminated, and must be within the cable length limit. You must connect devices to a shared bus in a way that allows you to disconnect them without affecting bus operation. Eight SCSI IDs are available on each bus.
- Storage configuration—Disks used in an ASE must be located in external storage expansion shelves connected to a shared SCSI bus. This enables all the member systems to access the data on the disks. Because these external units have their own power source, they are not dependent on any system's power.
- Network interconnect—Member systems must be connected to at least one common Internet Protocol (IP) subnet. Available Server uses the Internet for client access and ASE daemon communication purposes.

  For client access, Available Server allows you to configure up to four monitored network adapters (which must be on separate IP subnets). When a monitored network interface fails on a member, the status is passed to a script that you can customize. The default behavior of the script is to relocate all services from a member that experiences a failure of all monitored network interfaces. The services are relocated to a member system that has at least one functioning monitored interface.

  Additionally, the ASE daemons communicate over the primary network (set up by the running `asemgr` after the software has been installed). The primary network is monitored by default. A backup network may also be defined for ASE daemon communication. A failure of the primary network will cause ASE daemon communication to fail over to the backup network. For more information about setting up network monitoring, see the TruCluster Software Products *Administration* guide.

At a minimum, an ASE includes the following hardware:

- Two member systems
- One storage shelf

- One shared SCSI bus
- One common network

If you want to increase availability or performance, you can use additional systems, SCSI buses, or network connections. If you need more disks, you can use a **RAID** subsystem or add more SCSI buses. For example, you can use two shared SCSI buses and mirror the disks across the buses for high data reliability.

1.4 Overview of the TruCluster Production Server Hardware Configuration

As with a TruCluster Available Server Software configuration, a TruCluster Production Server hardware configuration consists of a number of highly specific hardware components. The hardware used in a TruCluster Production Server Software configuration is much the same as the hardware used in a TruCluster Available Server Software configuration, as discussed in Section 1.3, with the following exceptions:

- TruCluster Production Server Software supports from two to eight member systems.
- There must be sufficient internal and external SCSI controllers and disks to provide sufficient storage for the applications, but a shared SCSI bus is not a requirement for TruCluster Production Server Software.

  If highly available NFS, disk, tape, or user-defined services will be used, an ASE is required, and an ASE requires that the storage be on shared SCSI buses.

  Note that whereas only one ASE is supported with TruCluster Available Server Software, multiple ASEs are supported with TruCluster Production Server Software. All systems in an ASE must be on the same shared SCSI bus.
- As with TruCluster Available Server Software, at least one Internet Protocol (IP) subnet is required to allow client access to the cluster. For client access, like Available Server, Production Server allows you to configure up to four monitored network adapters (which must be on separate IP subnets). When a monitored network interface fails on a member, the status is passed to a script that you can customize. The default behavior of the script is to relocate all services from a member that experiences a failure of all monitored network interfaces. The services are relocated to a member system that has at least one functioning monitored interface.

For Production Server, daemon communication is over the MEMORY CHANNEL. If a second MEMORY CHANNEL is present, a failure of one MEMORY CHANNEL will cause daemon communication to fail over to the second MEMORY CHANNEL.

- TruCluster Production Server Software requires a least one peripheral component interconnect (PCI) MEMORY CHANNEL adapter on each system. The MEMORY CHANNEL adapters comprise the cluster interconnect for TruCluster Production Server Software, providing host-to-host communications. For a cluster with two systems, the MEMORY CHANNEL adapters are connected directly together 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 line cards. 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 Production Server Software. 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.

1.5 Overview of the TruCluster MEMORY CHANNEL Software Hardware Configuration

Like a TruCluster Production Server Software configuration, a TruCluster MEMORY CHANNEL Software hardware configuration also requires MEMORY CHANNEL adapters. However, the TruCluster MEMORY CHANNEL Software product does not support shared SCSI buses; therefore, it does not support an available server environment (ASE). You can have up to four MEMORY CHANNEL adapters with the TruCluster MEMORY CHANNEL Software product, but there is no failover.

As with TruCluster Production Server Software, the MEMORY CHANNEL adapter is required for any processor communication needed to support the MEMORY CHANNEL application programming interfaces (APIs) used.

1.6 Overview of Setting up the TruCluster Hardware Configuration

The following table provides an overview of the steps necessary to set up a TruCluster hardware configuration. This table contains entries for each of the TruCluster products: TruCluster Available Server Software (AS), TruCluster Production Server Software (PS), and TruCluster MEMORY

CHANNEL Software (MC). More specific hardware installation instructions are provided in Chapter 4.

Step	What	AS	PS	MC	Comments:
1	Plan your hardware.	X	X	X	See Chapter 4.
2	Draw a diagram of your configuration.	X	X	X[a]	Compare with examples in Chapter 4 and Chapter 6.
3	Identify all devices, cables, SCSI adapters, and so forth.	X	X	X	Use the diagram you just constructed.
4	Prepare the member systems by installing:				
	Additional Ethernet or Asynchronous Transfer Mode (ATM) network adapters.	X[b]	X		See Chapter 4.
	MEMORY CHANNEL adapters. Ensure that jumpers are set correctly.		X	X	See Chapter 5.
	SCSI bus adapters. Ensure that adapter jumpers are set correctly.	X	X		See Chapter 4.
5	Prepare the shared storage by installing disks and configuring any RAID controller subsystems.	X	X		See Chapter 4 and the documentation for the StorageWorks enclosure or RAID controller.
6	Install signal converters in StorageWorks enclosures, if applicable.	X	X		See Chapter 3 and Chapter 4.
7	Connect devices to the shared SCSI buses in each ASE. Terminate each bus. Use Y cables or trilink connectors where necessary.	X	X		See Chapter 3 and Chapter 4.
8	Connect the MEMORY CHANNEL adapters to each other or to the MEMORY CHANNEL hub as appropriate.		X	X	See Chapter 5.
9	Turn on MEMORY CHANNEL hubs and storage shelves, then turn on member systems.	X	X	X	

Step	What	AS	PS	MC	Comments:
10	Install firmware, set SCSI IDs, and enable fast bus speed as necessary.	X	X		See Chapter 4.
11	Display configuration information for each member system, and ensure that all shared disks are seen at the same device number.	X	X		See Chapter 4.

[a] Use a TruCluster Production Server Software configuration as an example, but remember, TruCluster MEMORY CHANNEL Software configurations do not use a shared SCSI bus.
[b] You may add extra network adapters for redundancy with Available Server.

2

Supported Hardware

This chapter describes the supported hardware for the TruCluster Software Products: TruCluster Available Server Software, TruCluster Production Server Software, and TruCluster MEMORY CHANNEL Software. This chapter also contains some specific requirements for some of the hardware components.

See the TruCluster Software Products *Release Notes* for the latest information about supported hardware and firmware revisions.

2.1 Member Systems

Table 2–1 lists the supported systems by product and the SCSI adapter used for shared SCSI buses on each system. Section 2.1.1 and Section 2.1.2 provide the specific requirements for member systems in a TruCluster Available Server Software (AS), TruCluster Production Server Software (PS), or TruCluster MEMORY CHANNEL Software (MC) configuration.

Table 2–1: Supported Member Systems

System	SCSI Adapter	AS	PS	MC
DEC 3000	PMAZC or KZTSA	Yes	No	No
DEC 7000 and DEC 10000	KZMSA	Yes	No	No
AlphaServer 300	KZPSA	Yes	No	No
AlphaServer 400	KZPSA	Yes	No	No
AlphaServer 800	KZPSA	Yes	Yes	Yes
AlphaServer 1000	KZPSA	Yes	No	No
AlphaServer 1000A	KZPSA	Yes	Yes	Yes
AlphaServer 1200	KZPSA	Yes	Yes	Yes
AlphaServer 2000	KZPSA	Yes	Yes	Yes
AlphaServer 2100	KZPSA	Yes	Yes	Yes
AlphaServer 2100A	KZPSA	Yes	Yes	Yes
AlphaServer 4000	KZPSA	Yes	Yes	Yes

Table 2–1: Supported Member Systems (cont.)

System	SCSI Adapter	AS	PS	MC
AlphaServer 4000A	KZPSA	Yes	Yes	Yes
AlphaServer 4100	KZPSA	Yes	Yes	Yes
AlphaServer 8200 and 8400	KZPSA	Yes	Yes	Yes

2.1.1 Available Server Member System Requirements

Member systems in a TruCluster Available Server Software environment have the following requirements:

- Each supported member system requires a minimum firmware revision. See the TruCluster Software Products *Release Notes* for information about required firmware.
- If you have a system with a PMAZC or a KZMSA installed, you can use only three member systems in the ASE.
- For TURBOchannel-based systems, the `boot_reset` console variable must be set to on. If the variable is not set to on, the TURBOchannel (TC) option self-test may fail and the system may not automatically reboot.
- A system may not boot properly if it does not have a graphics head attached. If a system does not have a graphics head, but has a terminal interface, you must use the following command:

  ```
  >>> set console serial
  ```

- You must perform the following steps to turn on and turn off a system that is connected directly to a SCSI signal converter:

  1. You must turn on a system and allow it to complete its startup diagnostics before you turn on the signal converter that is connected to the system.
  2. Before you turn off a system, you must first turn off any signal converter that is connected to the system.

- If your system is a DEC 3000 Model 300, you must follow these steps to enable the system to automatically boot as a server:

  1. Ensure that a keyboard and mouse are attached to the system with a BC13M cable, which is connected to the 15-pin D-sub in the back of the system. A graphics monitor is not required, but a keyboard and mouse must be connected to prevent errors when you turn on the system.

2. Connect a serial terminal to the 25-pin serial port in the back of the system. The serial terminal must be set to 9600 baud, 8 data bits, 1-stop bit, and no parity.
3. Attach a jumper to pins 2 and 3 (the middle pins) on the W2 jumper on the system motherboard (54-22257). The W2 jumper is located on the front of the board, to the left of the SIMM connectors.
4. Turn on the system.
5. If the system does not automatically boot, enter the following command at the console prompt:

   ```
   >>> set auto_a boot
   ```

   Turn off and turn on the system. Ensure that the BOOTDEF_DEV console variable is set to the correct boot device.

2.1.2 Production Server and MEMORY CHANNEL Member Requirements

The requirements for member systems in a TruCluster Production Server Software or TruCluster MEMORY CHANNEL Software environment are as follows:

- Each supported member system requires a minimum firmware revision. See the TruCluster Software Products *Release Notes* for information about required firmware.
- Production Server Software supports up to eight systems in a hardware configuration.
- MEMORY CHANNEL Software supports up to eight systems in a hardware configuration.
- Production Server does not support the XMI CIXCD on an AlphaServer 8200 or 8400 system.

2.2 Network Adapters

Ethernet and Asynchronous Transfer Mode (ATM) networks are supported by TruCluster Available Server Software and TruCluster Production Server Software. The following sections describe these adapters.

2.2.1 Supported Ethernet Adapters

The following Ethernet network adapters are supported:

- DE500 (PCI/Fast Ethernet)
- DEFPA (PCI/FDDI)

- DE435 (PCI/Ethernet)
- DEFEA (EISA/FDDI)
- DE422 (EISA/Lance Ethernet)
- DE425 (EISA/Ethernet)
- DEMNA (XMI/Ethernet)
- DEMFA (XMI/FDDI)
- PMAD (TURBOchannel/Ethernet)
- DEFTA (TURBOchannel/FDDI)
- DEFZA (TURBOchannel/FDDI)

2.2.2 Supported ATM Adapters

Asynchronous Transfer Mode (ATM) is a high-speed, connection-oriented, cell-switched technology. It meets the real-time networking requirements of multimedia applications, while providing increased bandwidth for current installations.

TruCluster Available Server Software supports ATM Lan Emulation (LANE) as a communication path between systems in a TruCluster environment. Both TruCluster Available Server Software and TruCluster Production Server Software products support the use of ATM Lan Emulation as a network for client access.

ATM is a cell-switching and multipexing technology where the sender negotiates a path with the network for a connection to the destination.

Information to be sent, which may be data, voice, or video, is segmented into a 53-byte, fixed-length cell (48 bytes of actual information with a 5-byte header).

An ATM network consists of the following components:

- ATM switch: A specialized system that connects one end system to another and forwards, or switches, ATM cells from one end system to another. The ATM switch contains information necessary to route the cells from one end system to another based upon cell header information.
- End system: A system physically connected to a switch. It communicates with other end systems through one or more switches.

The DGLPB-AB ATMworks 350 155 MB/sec multimode fiber optics PCI adapter is supported.

Any high-performance ATM switch with 155 MB/sec ports is supported. A minimum of one ATM switch is required for each ATM circuit.

2.3 SCSI Bus Adapter

To connect a member system to a shared SCSI bus, you must install a SCSI bus adapter in an I/O bus slot. Table 2–2 describes the supported SCSI adapters.

Table 2–2: Supported SCSI Adapters

SCSI Bus Adapter	Transmission Method
KZPSA[a]	Differential
KZTSA[b]	Differential
KZMSA[b]	Single-ended
PMAZC[b]	Single-ended

[a]Used for both Available Server and Production Server
[b]Used for Available Server only

The following sections describe the SCSI adapters in more detail.

2.3.1 KZPSA SCSI Adapter Requirements

KZPSA SCSI adapters have the following requirements:

- If you have a KZPSA adapter installed in an AlphaServer 1000, 1000A, 2000, 2100, or 2100A system, you must set the `bus_probe_algorithm` console variable to `new` by entering the following command:

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

- On AlphaServer 1000A and 2100A systems, updating the firmware on the KZPSA SCSI adapter is not supported when the adapter is behind the PCI-to-PCI bridge.
- Before you install a KZPSA adapter in a PCI slot, you must remove the termination from the adapter.

2.3.2 KZTSA SCSI Adapter Requirements

KZTSA SCSI adapters have the following requirements:

- If you are using both a KZTSA adapter and a KZMSA adapter on a shared bus in the available server environment (ASE), you must disable the `wide` parameter on the system in which the KZTSA adapter is installed. However, you must first obtain the SCSI ID for the KZMSA adapter. To display the target ID of the KZTSA adapter and then disable the parameter, use the following commands on the system in which the KZTSA is installed:

```
>>> t tc#
>>> set target_id wide disable
```

The # specifies the TC slot in which the KZTSA adapter is installed. The *target_id* specifies the SCSI ID of the KZMSA adapter.

2.3.3 KZMSA SCSI Adapter Requirements

KZMSA SCSI adapters have the following requirements:

- If your system has a KZMSA adapter installed, you can use only three member systems in the ASE.
- You must attach a DWZZA-AA signal converter to each KZMSA adapter that you connect to the shared bus.
- The KZMSA hardware revision level is contained in its boot ROM. You must contact your field service representative to update the KZMSA boot ROM to Revision F03, if necessary. You can determine the hardware revision level by using console commands, or by examining the 23-class part number printed on the EPROM located at module position E7. The part numbers and associated hardware revisions are shown in the following table:.

Part number	Hardware Revision
23-368E9-01	Revision F01
23-386E9-01	Revision F02
23-419E9-01	Revision F03

- KZMSA adapters must have one of the following revision part numbers:
  - 609-3400546
  - 609-3400563

  If the two NCR 53C710 chips on the KZMSA adapter are not marked with either of these part numbers, return the adapter to DIGITAL for replacement.
- You must enable the Disable Reset configuration option for each KZMSA channel that you use for a shared bus. Keep the option disabled if the channel will not be used for a bus.
- Available Server supports up to 16 KZMSA adapters installed in XMI slots.

2.3.4 PMAZC TURBOchannel SCSI Controller Requirements

PMAZC TURBOchannel SCSI controllers have the following requirements:

- If your system has a PMAZC controller installed, you can have at most three member systems in the ASE.
- If you update the firmware for a PMAZC controller, the update procedure may change the PMAZC option settings to the default settings. For example, the SCSI ID may be reset to 7 and the bus speed may be reset to slow.
- If you want to run a PMAZC controller in Fast SCSI mode, DIGITAL recommends that you use DWZZA signal converters and a differential bus in your hardware configuration.
- When a system with a PMAZC controller installed is turned on, it may hang on the PMAZC self-tests if the module is not terminated either internally or externally.

2.4 MEMORY CHANNEL Hardware

Table 2–3 describes the supported MEMORY CHANNEL hardware.

Table 2–3: Supported MEMORY CHANNEL Hardware

MEMORY CHANNEL Hardware Component	Designation	Description
MEMORY CHANNEL PCI adapter	CCMAA-AA or CCMAA-BA	MEMORY CHANNEL PCI adapter used as the cluster interconnect.
MEMORY CHANNEL hub	CCMHA-AA	PC-class enclosure that is populated with line cards and used to connect MEMORY CHANNEL adapters. A hub is required if you have more than two member systems.
Line card	CCMLA-AA	Installed in a MEMORY CHANNEL hub to allow connection of more than two systems.
Link cable	BC12N-10 (10 meters)	Connects a MEMORY CHANNEL adapter to a hub or one MEMORY CHANNEL adapter to another MEMORY CHANNEL adapter (virtual hub).

The MEMORY CHANNEL hardware requirements are as follows:

- 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).
- The maximum length of a MEMORY CHANNEL link cable is 3 meters (10 feet).
- 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.

2.5 Disk Devices

Table 2–4 lists the supported disk devices for shared SCSI buses and the data paths available for each disk type.

Table 2–4: Supported Disk Devices

Disk	Data Path
RZ26	Narrow
RZ26L	Narrow and wide
RZ26N	Narrow and wide
RZ28	Narrow and wide
RZ28B	Narrow
RZ28D	Narrow and wide
RZ28L	Narrow and wide
RZ28M	Narrow and wide
RZ29	Narrow and wide
RZ29B	Narrow and wide
RZ29L	Narrow and wide
RZ40	Narrow and wide
RZ40L	Narrow and wide
RZ1BB	Narrow and wide
RZ1CB	Narrow and wide
RZ1DB	Narrow and wide

The requirements for disk devices are as follows:

- Supported disks require a minimum firmware revision. See the TruCluster Software Products *Release Notes* for information about the required firmware.

- Disks on shared SCSI buses must be installed in external storage shelves.
- TruCluster does not support Prestoserve on any shared disk.

2.6 Supported RAID Controllers

RAID controllers provide high performance, high availability, and high connectivity access to SCSI devices through a shared SCSI bus. Table 2–5 lists the RAID controllers supported in the Available Server (AS) and Production Server (PS) environments.

Table 2–5: Supported RAID Controllers

RAID Controller	AS	PS
HSZ10-A*x*	Yes	No
SWXRA-Z1 (HSZ20)	Yes	Yes
HSZ40-A*x*	Yes	Yes
HSZ40-B*x*	Yes	Yes
HSZ40-C*x*	Yes	Yes
HSZ50-A*x*	Yes	Yes
HSZ70	Yes	Yes

RAID controllers have the following requirements:

- The HSZ10 controller must be used in an available server environment (ASE) that includes a PMAZC controller exclusively.
- The HSZ10 controller uses only one SCSI ID.
- The HSZ40 controller can be configured with one to four SCSI IDs.
- The HSZ40 controller can be configured with one to eight **logical unit numbers (LUNs)** for each SCSI ID. For any number of disks, having more SCSI IDs with fewer LUNs for each SCSI ID provides better performance than a configuration with fewer SCSI IDs and more LUNs.
- Each RAID controller requires a specific minimum Hierarchical Storage Operating Firmware (HSOF) revision. See the TruCluster Software Products *Release Notes* for information about required HSOF revisions.
- Production Server supports single-controller, single-bus HSZ40 array configurations and dual-controller, single-bus HSZ40 configurations.

2.7 Disk Storage Shelves

Table 2–6 lists the supported disk storage shelves.

Table 2–6: Supported Disk Storage Shelves

Storage Shelf	Transmission Mode	Data Path	Internal SCSI Bus Length
BA350	Single-ended	Narrow	0.9 Meter
BA353	Single-ended	Narrow	0.9 Meter
BA356	Single-ended	Wide	1.0 Meter
UltraSCSI BA356	UltraSCSI Single-ended	Wide	1.0 Meter

Disk storage shelves have the following requirements:

- Storage shelves must be external with an independent power supply.
- You must connect a SCSI signal converter to a storage shelf that has a single-ended SCSI interface, such as a BA350, BA353, or non-UltraSCSI BA356 storage shelf.
- A shared differential UltraSCSI bus is connected to the DS-BA35X-DA personality module in an UltraSCSI BA356.

The following sections discuss these storage shelves in more detail.

2.7.1 BA350 Storage Shelf

Up to seven narrow (8-bit) StorageWorks building blocks (SBB) 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.

The BA350 storage shelf contains internal SCSI bus termination and a SCSI bus jumper. There are occasions when the termination must be removed from the BA350 (for example, when daisy chaining two BA350s together). 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 2–1 shows the relative locations of the BA350 SCSI bus terminator and SCSI bus jumper. They are accessed from the rear of the box. For TruCluster operations, you must install the J jumper.

Figure 2–1: BA350 Internal SCSI Bus

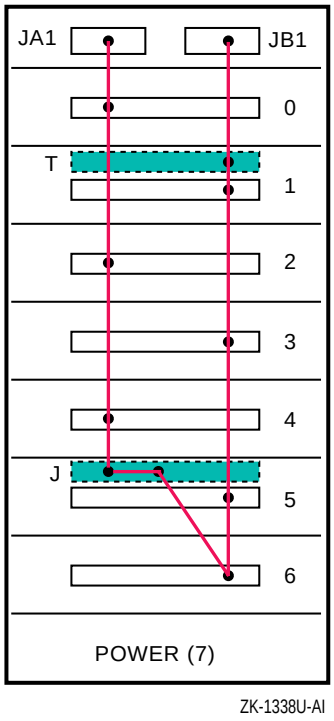

2.7.2 BA353 Storage Shelf

The BA353 is probably of little use in TruCluster configurations; it only has three slots for shared disks.

The SCSI ID for disks installed in a BA353 is defined by device address switches on the back of the BA353. The switches are located to the left of the SCSI input and SCSI output connectors, as shown in Figure 2–2.

The switches are marked as Left (Slot 1), Center (Slot 2), and Right (Slot 3). Slot 1 is the leftmost slot when the BA353 is viewed from the front.

The On position of a switch generates a logic 1 in the device address, and switch one is the least significant bit (LSB) in the device address. The SCSI IDs shown in Figure 2–2 would be 0, 1, and 2, left, center, and right.

Figure 2–2: BA353 Device Address Switches and SCSI Connectors

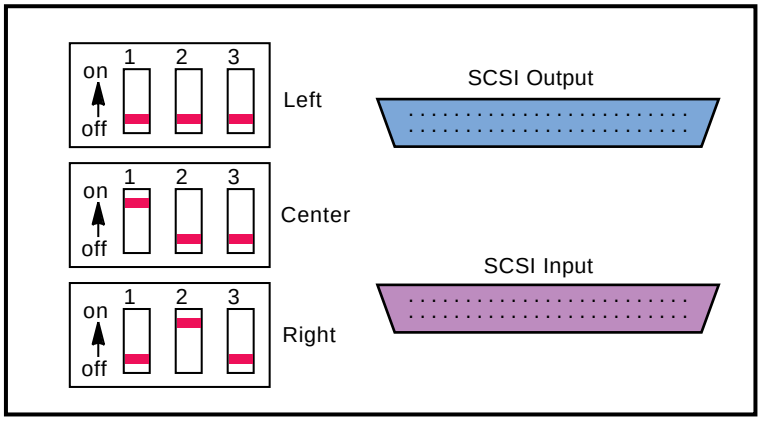

2.7.3 BA356 Storage Shelf

There are two variations of the BA356 used in TruCluster configurations: the UltraSCSI BA356 and the non-UltraSCSI BA356.

An example of the non-UltraSCSI BA356 is the BA356-KC, which has a wide, single-ended internal SCSI bus. It uses the BA35X-MH 16-bit personality module and has a 150-watt power supply.

The 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 and the shared, differential UltraSCSI bus. The UltraSCSI BA356 uses an 180-watt power supply.

2.7.3.1 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. Also, like the BA350, the SBB SCSI IDs are based upon the slot they are installed in, but the switches on the personality module (BA35X-MH) have to be set to off, the default switch positions. Verify that the personality module switches are all off.

Figure 2–3 shows the relative location of the BA356 SCSI bus jumper, BA35X-MF. The jumper is accessed from the rear of the box. For TruCluster operations, 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 two BA356s together. In this case, when the cable is connected to JB1, the personality module terminator is disabled.

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 environment.

You can use the position behind slot 1 to store the SCSI bus terminator or jumper.

Figure 2–3 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 TruCluster operations, you must install the J jumper.

Figure 2–3: BA356 Internal SCSI Bus

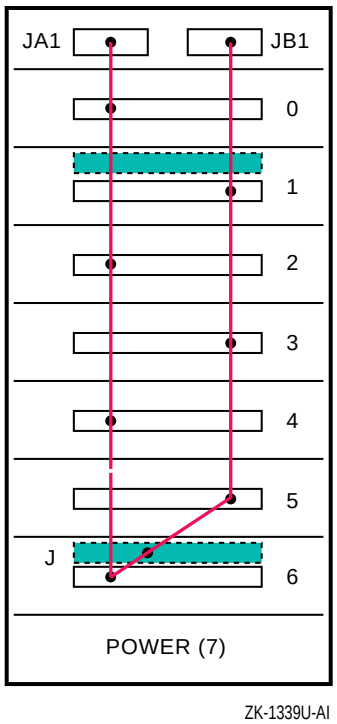

Note that JA1 and JB1 are located on the personality module (in the top of the box when it is standing vertically). JB1, on the front of the module, is

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 2–4):

- 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 2–4: BA356 Jumper and Terminator Module Identification Pins

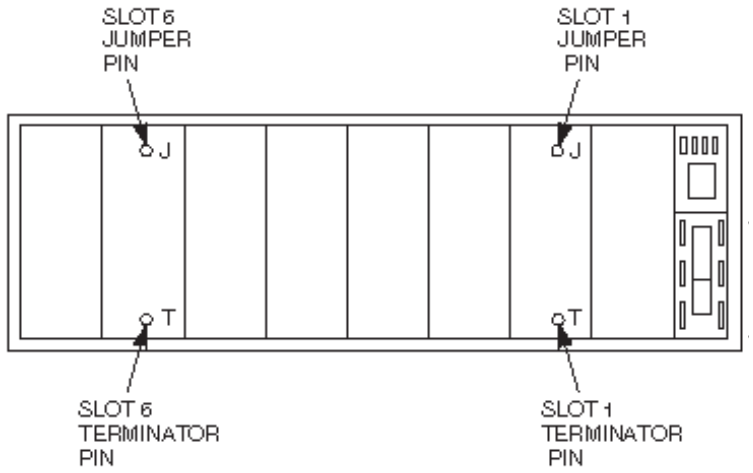

2.7.3.2 UltraSCSI BA356 Storage Shelf

The UltraSCSI BA356 can also hold up to seven StorageWorks building blocks (SBBs). These SBBs are UltraSCSI single-ended wide devices. The SBB SCSI IDs are based upon the slot they are installed in, but the switches on the personality module (BA35X-DA) have to be set to off, the default switch positions. Verify that the personality module switches are all off.

The jumper module is positioned behind slot 6 as with the non-UltraSCSI BA356 shown in Figure 2–3. For TruCluster operations, 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 2–4.

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 switch pack S4. For normal cluster operations, S4-1 and S4-2 should normally be on. Section 3.6.2.2 discusses these switches in more detail.

2.8 SCSI Signal Converters

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

Table 2–7 lists the supported SCSI signal converters.

Table 2–7: Supported SCSI Signal Converters

Device	Description
DWZZA-AA	Standalone unit Converts single-ended, narrow SCSI to differential, wide SCSI. Use with a BA350 or BA353 storage shelf.
DWZZA-VA	SBB Converts single-ended, narrow SCSI to differential, wide SCSI. Install in a BA350 or BA353 storage shelf.
DWZZB-AA	Standalone unit Converts single-ended, wide SCSI to differential, wide SCSI. Use with a BA356 storage shelf.
DWZZB-VW	SBB Converts single-ended, wide SCSI to differential, wide SCSI. Install in a BA356 storage shelf.

The requirements for SCSI signal converters are as follows:

- For each type of SCSI signal converter, a TruCluster environment requires a specific, minimum hardware revision. See the TruCluster Software Products *Release Notes* for information about required hardware revisions.
- 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

2.9 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 Very High Density Cable Interconnect (VHDCI) 68-pins (UltraSCSI). 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 2–8 describes each supported cable and the context in which you would use the cable.

Table 2–8: Supported SCSI Cables

Cable	Connector Density	Pins	Configuration Use
BN21V-0B	One high, two low	50-pin	This Y cable attaches to a PMAZC and can be terminated if necessary.
BN21W-0B	Three high	68-pin	A Y cable that attaches to a KZTSA, KZPSA, HSZ10, HSZ40, HSZ50, or the differential side of a signal converter. It can be terminated if necessary.

Table 2–8: Supported SCSI Cables (cont.)

Cable	Connector Density	Pins	Configuration Use
BN21R or BN23G	One high, one low	50-pin	Connects narrow devices such as the PMAZC, KZMSA, BA350, or BA353 to the single-ended side of a DWZZA signal converter. (Replaces the BC09D cable.)
BN21H or BN21J	Two high	50-pin	Connects narrow devices, such as a BA350 or a BA353.
BC19J or BC06P[a]	Two low	50-pin	Connects BN21V-0B Y cables to each other or the single-ended end of a DWZZA to a TZ885 .
BN21K or BN21L	Two HD68	68-pin	Connects BN21W Y cables or wide devices. For example, connects KZTSAs, KZPSAs, HSZ10s, HSZ40s, the differential sides of two SCSI signal converters, or a DWZZB-AA to a BA356.
BN38C or BN38D	One HD68, one VHDCI	VHDCI to HD68	Connect a high-density (HD68) Y cable (BN21W-0B) to a VHDCI trilink.
BN37A	Two VHDCI	VHDCI to VHDCI	Connect two VHDCI trilinks to each other.

[a]The BC06P is no longer available for purchase.

The requirement for SCSI cables is as follows:

- 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 2–9 describes the supported SCSI terminators and the context in which you would use them.

Table 2–9: Supported SCSI Terminators

Terminator	Density	Pins	Configuration Use
H8574-A or H8860-AA	Low	50-pin	Terminates a BN21V-0B Y cable. Used with Available Server configurations only.
H879-AA	High	68-pin	Terminates an H885 trilink connector or BN21W-0B Y cable.

Table 2–9: Supported SCSI Terminators (cont.)

Terminator	Density	Pins	Configuration Use
12-37004-04	High	50-pin	Terminates a BA353 input connector.
12-41667-01	High	50-pin	Terminates an unused PMAZC port. Used with Available Server configurations only.
H8863-AA	VHDCI	68-pin	Terminate a VHDCI trilink connector.

Table 2–10 describes the supported trilink connectors and the context in which you would use them.

Table 2–10: Supported Trilink Connectors

Trilink Connector	Number of Connectors	Pins	Density	Configuration Use
H885	Three	68-pin	High	Attaches to high-density, 68-pin cables or devices, such as a KZTSA, KZPSA, HSZ10, HSZ40, HSZ50, or the differential side of a SCSI signal converter. Can be terminated with an H879-AA terminator if necessary.
H8861-AA	Three	68-pin	VHDCI	Attaches to VHDCI 68-pin cables, UltraSCSI BA356 JA1, HSZ70 RAID controllers. Can be terminated with an H8863-AA terminator if necessary.

The requirement for trilink connectors is 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 HD68 trilink if installing it will block an adjacent peripheral component interconnect (PCI) port. Use a BN21W-0B Y cable instead.

2.11 Supported Tape Devices (AS, PS)

The following tape devices are supported by the TruCluster Available Server Software and TruCluster Production Server Software products:

- TZ88—The TZ88 is a streaming DIGITAL Linear Tape (DLT) cartridge tape device capable of holding up to 40 GB data per CompacTape IV cartridge when using 2:1 compression. It is capable of storing/retrieving data at a rate of up to 10.8 GB per hour.

Two TZ88 models are supported: The TZ88N-TA, a table-top model, and the TZ88N-VA, a StorageWorks building block (SBB) 5.25-inch carrier. The SBB version takes up three slots in a BA350 StorageWorks enclosure.

The TZ88 uses CompacTape III, CompacTape IIIXT, or CompacTape IV media.

Other TZ88 models (for instance, the TZ88N-AV fast-wide, differential drive) are supported as the tape devices for tape loaders. These TZ88 tape drives cannot be used as standalone drives.

- TZ89—Another streaming DLT cartridge tape drive, the TZ89 has a capacity of up to 70 GB per cartridge (CompacTape IV) when using 2:1 compression. In compressed mode it can back up 36 GB of data per hour.

  Two TZ89 models are supported: The DS-TZ89N-TA, a table-top model, and the DS-TZ89N-VW, a StorageWorks building block (SBB) 5.25-inch carrier. The SBB version takes up three slots in a BA356 StorageWorks enclosure.

  The TZ89 also uses CompacTape III, CompacTape IIIXT, or CompacTape IV media.

  Other TZ89 models (for instance, the TZ89N-AV fast-wide differential drive) are supported as the tape devices for tape loaders. These TZ89 tape drives cannot be used as standalone drives.

- TZ885—The TZ885 is a DLT magazine tape subsystem combining a cartridge tape drive and an automatic cartridge loader. It uses a five-cartridge (CompacTape IV) removable magazine to provide a mini-tape library with a total capacity of 200 GB of compressed data. The TZ885 uses the TZ88N-AX tape drive, a single-ended drive. The SCSI bus connector is 50-pin low-density, single-ended.

3

Shared SCSI Bus Requirements and Configuration

TruCluster Production Server Software allows you to configure one or more available server environments (ASEs), depending on the number of systems in the cluster. TruCluster Available Server Software allows you to configure one ASE. An ASE 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 and describes requirements for the shared SCSI bus. In addition to using only the supported hardware described in the TruCluster Software Products *Release Notes*, adhering to the requirements described in this chapter will ensure that your ASE operates correctly.

This chapter discusses the following topics:

- Shared SCSI bus configuration requirements (Section 3.1)
- Numbering SCSI devices (Section 3.2)
- SCSI bus performance (Section 3.3)
- SCSI bus device identification numbers (Section 3.4)
- SCSI bus length (Section 3.5)
- Using SCSI signal converters (Section 3.6)
- SCSI bus termination (Section 3.7)
- Connecting devices to a shared SCSI bus (Section 3.8)

In some sections of this manual, a distinction is made between the different implementations of the Small Computer System Interconnect (SCSI), SCSI, Fast 10 (F10), and UltraSCSI. Most of the manual uses SCSI as the generic term, which could be any of the SCSI implementations.

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.
- The logical bus numbers you assign to the adapters for shared buses must be higher than the bus numbers for local buses. The starting number for shared SCSI buses must be high enough to allow for expansion of local buses without requiring the shared buses to be renumbered.
- SCSI bus specifications limit the number of devices on an 8-bit (narrow) SCSI bus to 8. The limit is 16 devices on a 16-bit SCSI bus (wide). However, although Available Server and Production Server supports wide, differential and UltraSCSI devices, DIGITAL UNIX Version 4.0 supports only 8 devices on a SCSI bus. See Section 3.4 for more information.
- The length of each physical bus is strictly limited. See Section 3.5 for more information.
- For each physical bus or UltraSCSI bus segment, you can have only two terminators, one at each end.

  Use trilink connectors and Y cables to connect devices to a shared bus, so you can disconnect the devices without affecting bus termination. See Section 3.7 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 or data paths. See Section 3.6 for information about SCSI signal converters.
- You can use up to 30 shared buses in an ASE, and a shared SCSI bus can be in only one ASE.

  A member system can be in only one ASE, but you can set up multiple ASEs in a TruCluster Production Server Software environment, if you have four or more member systems.
- 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 ASE, you must be able to isolate the device from the shared bus without affecting bus termination.
- All DIGITAL UltraSCSI host adapters support UltraSCSI disks at UltraSCSI speeds in UltraSCSI BA356 shelves or behind the HSZ70 RAID controller. 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.
- Differential UltraSCSI adapters may be connected to both a non-UltraSCSI BA356 shelf (via a DWZZB-VW) and the new 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.
- The HSZ70 UltraSCSI RAID controller has a wide differential UltraSCSI host bus with a Very High Density Cable Interconnect (VHDCI) connector. HSZ70 controllers will work with fast, wide differential SCSI adapters (for example, KZPSA) at fast SCSI speeds. A converter cable is necessary.
- 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.
- An older BA356 that has been retrofitted with a BA35X-HH 180-watt power supply and BA35X-DA personality module is not the same as an UltraSCSI BA356.

3.2 Numbering SCSI Buses

All member systems must recognize the disks on a shared bus at the same device number. A device number is obtained from the logical bus number, which is defined in the system configuration file. If you connect a shared bus to SCSI controllers that have the same logical bus number on each system, the shared disks will have the same device number on each system.

Bus numbers are assigned to SCSI controllers during the kernel configuration process, and they are specified in the kernel configuration file. When you configure a kernel by running the `doconfig` program using the generic kernel, an algorithm is used to probe the SCSI controllers installed in the system. As the probe algorithm encounters the adapters, it assigns logical bus numbers to the SCSI controllers in sequence, starting with 0.

Because systems have different numbers of internal buses, in an available server environment (ASE), the ase_fix_config script is used to manually assign logical bus numbers and override the probe algorithm. If you choose to have the software installation procedure rebuild your kernel, it automatically runs the script and prompts you for information about bus numbers. You can also manually invoke the script with the following command:

```
# /var/ase/sbin/ase_fix_config –c  CONFIG-FILE
```

CONFIG-FILE is the name of the system's configuration file, which is stored in the /sys/conf directory.

Before you install your SCSI controllers, plan your bus configuration. DIGITAL recommends that you install SCSI controllers for the local buses in the lower bus slots in your system, leave some empty slots, and then install the SCSI controllers for the shared buses. If you use this method, you can install additional SCSI controllers for both the local buses and the shared buses without disrupting the shared bus numbering scheme. For example, if a maximum of eight local controllers will be installed in a system, install the first controller for a shared bus in slot 8. Then, install another controller for a shared bus in the next highest slot (7), and so on.

See the TruCluster Software Products *Software Installation* manual for more information about the ase_fix_config script. In addition, some SCSI controllers have two ports (or channels) so you can connect two shared SCSI buses to each controller. For example, a PMAZC TURBOchannel SCSI controller has ports A and B, and a KZMSA XMI-to-SCSI adapter has channels 0 and 1. Both ports do not have to be used, but any unused ports must be terminated.

If you are using dual-port SCSI controllers, a shared bus must be connected to the same port on each system. For example, a shared bus that is connected to port A (or channel 0) on one SCSI controller must be attached to port A (or channel 0) on all the other SCSI controllers.

3.3 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.

3.3.1 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:
  - PMAZC (Available Server only)
  - KZMSA (Available Server only)
  - BA350, BA353, and 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, enables you to use longer cables, and uses 68-pin, high-density or 68-pin VHDCI connectors (UltraSCSI). Devices with differential SCSI interfaces include the following:
  - KZPSA
  - KZTSA (Available Server only)
  - HSZ10 (Available Server only), HSZ40, HSZ50, and HSZ70 controllers
  - Differential side of a SCSI signal converter or personality module

You cannot use the two transmission methods in the same physical bus. 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. See Section 3.6 for information about using SCSI signal converters.

UltraSCSI does not support the use of a signal converter such as the DWZZA or DWZZB for Available Server or Production Server Version 1.5. The DS-BA35X-DA personality module functions as a signal converter for the UltraSCSI BA356. It is the interface between the shared differential UltraSCSI bus and the BA356 internal single-ended SCSI bus.

3.3.2 Data Path

There are two possible 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 and a 32-bit data path for UltraSCSI. This mode increases the amount of data that is transferred in parallel on the bus.

Usually, but not always, single-ended devices are narrow, and differential devices are wide. The BA356 storage shelves are examples of single-ended, wide devices. If you want to connect devices that have different data paths, use a SCSI signal converter between the devices.

3.3.3 Bus Speed

Bus speeds vary depending upon the type of SCSI bus: SCSI, SCSI-2, fast, wide SCSI-2, or UltraSCSI and the data path width. The following sections discuss these bus speeds in more detail.

3.3.3.1 SCSI, SCSI-2, and Fast Wide SCSI-2

SCSI and SCSI-2 controllers operate in two data transfer speeds: standard or slow mode and fast bus speed. In slow mode, the narrow 8-bit SCSI bus achieves up to 5 MB/sec. Fast SCSI-2 bus speed (the fast synchronous transfer option) enables I/O devices to attain high peak-rate transfers in synchronous mode. Production Server supports only fast bus mode.

To set the bus speed on a SCSI controller, use either console commands or the Loadable Firmware Update (LFU) utility, depending on the type of SCSI controller.

Although fast SCSI bus speed (sometimes referred to as Fast 10) doubles the transmission rate to 10 MB/sec, it reduces the maximum cable length for each single-ended bus from 6 meters to 3 meters. Fast bus speed on a wide (16-bit) SCSI bus increases the maximum bus bandwidth to 20 MB/sec (referred to as Fast 20). See Section 3.5 for more information about bus length.

3.3.3.2 UltraSCSI Bus Speed

Fast bus speed in an UltraSCSI bus increases the maximim bus bandwidth to 40 MB/sec.

3.4 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 7). 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.

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
```

This order specifies that 7 is the highest priority, and 0 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.

The SCSI ID for a disk in a BA350 or BA356 storage shelf corresponds to its slot location. In addition, you can set the SCSI IDs for some types of disks in a BA350 or BA356 storage shelf by using switches in the rear of some StorageWorks building blocks (SBBs). The SCSI ID for a disk in a BA353 storage shelf is set by the device address switches on the rear of the box.

The HSZ10 controller uses only one SCSI ID. You set the SCSI ID for an HSZ10 controller using switches on the unit. You can configure the HSZ40 controller with one to four SCSI target addresses.

3.5 SCSI Bus Length

There is a limit to the length of the cables in a shared SCSI bus. The total cable length for a physical bus or UltraSCSI 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 physical bus. If you have devices with different transmission methods, you will have both single-ended and differential physical buses, each of which must be terminated only at both ends and must adhere to the rules on bus length.

Note

In an Available Server or Production Server configuration you always have single-ended SCSI bus segments since all of the storage shelves use a single-ended bus.

Table 3–1 describes the maximum cable length for a physical SCSI bus.

Table 3–1: SCSI Bus Length

SCSI Bus	Bus Speed	Maximum Cable Length
Narrow, single-ended	5 MB/second	6 meters
Narrow, single-ended fast	10 MB/second	3 meters
Wide differential, fast	20 MB/second	25 meters
Differential UltraSCSI	40 MB/second	24 meters [a]

[a]The length of the differential UltraSCSI cables between the terminators. You do not need to compute the length of the single-ended bus segments in storage shelves, trilink connectors, or the SCSI adapters.

The total length of a physical bus must include the amount of cable that is located inside each system and disk storage shelf. This length varies, depending on the device. For example, the length of cable inside a BA350, BA353, or BA356 storage shelf is approximately 1.0 meter.

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.

3.6 Using Signal Converters

A SCSI 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 bus to isolate bus segments for maintenance purposes.

Each SCSI signal converter has a single-ended side with either a narrow or a wide data path, and a differential side with a wide data path.

Note

Some UltraSCSI documentation uses the UltraSCSI "bus expander" term when referring to the DWZZA, DWZZB, and

UltraSCSI signal converters. Other UltraSCSI documentation refers to some UltraSCSI products as bus extender/converters.

For Available Server and Production Server Version 1.5 there are no supported 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, except the DS-BA35X-DA personality module. However, because the DS-BA35X-DA personality module does convert a differential bus to a single-ended bus, it is discussed with the signal converters.

A SCSI signal converter is required in the following cases:

- You want to connect devices with different transmission modes or data path widths.
- You are using a DEC 7000 or DEC 10000 system in an available server environment (ASE). The KZMSA SCSI adapter does not have removable termination resistors.
- You want to run a PMAZC TURBOchannel SCSI controller using fast SCSI bus speed.

3.6.1 Types of SCSI Bus Signal Converters

Signal converters can be standalone units, StorageWorks building blocks (SBBs) that are installed in a storage shelf disk slot, or as with the UltraSCSI BA356, the SBB Shelf Differential I/O Module (pesonality module DS-BA35X-DA). You must use the signal converter or personality module that is appropriate for your hardware configuration.

For example, use a DWZZA-VA signal converter to connect a BA350 or BA353 (single-ended and narrow) storage shelf to a differential device, but use a DWZZB-VW signal converter to connect a non-UltraSCSI BA356 (single-ended and wide) storage shelf to a differential adapter. The DS-BA35X-DA personality module is used in an UltraSCSI BA356 to connect a differential adapter to the single-ended disks in the BA356.

Table 2–7 shows the types of supported SCSI signal converters.

3.6.2 SCSI Signal Converter Removable Termination

The DWZZA and DWZZB signal converters have removable termination. The DS-BA35X-DA personality module has switch selectable termination. The following sections describe these signal converters in more detail.

3.6.2.1 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 trilink connector to this side. To remove the differential termination, remove the five 14-pin SIP resistors (located near the differential end of the signal converter). You can terminate the trilink connector to terminate the differential bus. If you detach the trilink connector from the signal converter, the shared SCSI bus is still terminated.

In most cases, you must keep the termination in the single-ended side to terminate the single-ended bus connecting the signal converter and the single-ended device. Verify that the termination is active. A DWZZA should have jumper J2 installed. Jumpers W1 and W2 should be installed in a DWZZB.

However, if you have a DWZZA-VA signal converter installed in a BA353 storage shelf, you must remove the termination from the single-ended side. To do this, remove jumper J2.

Figure 3–1 shows the status of internal termination for a standalone SCSI signal converter that has a trilink connector attached to the differential side.

Figure 3–1: Standalone SCSI Signal Converter

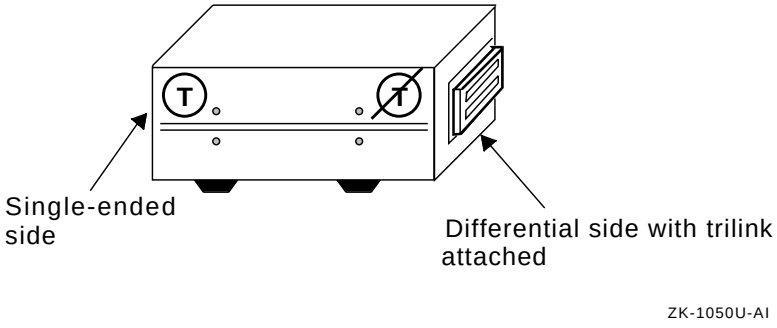

Figure 3–2 shows the status of internal termination for an SBB SCSI signal converter that has a trilink connector attached to the differential side.

Figure 3–2: SBB SCSI Signal Converter

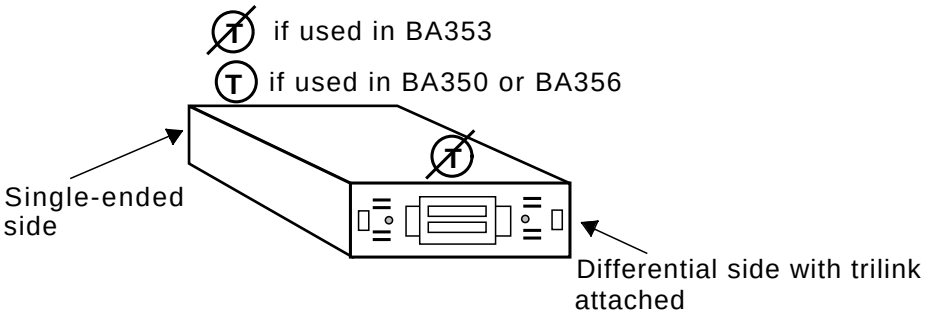

3.6.2.2 DS-BA35X-DA Termination

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 BA356.

The personality module controls termination for the external differential UltraSCSI bus segment, and for both ends of the internal UltraSCSI 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, which allows the use of the BA356 in the middle of the bus or external termination for a 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.

Note

S4-3 and S4-4 have no function on the DS-BA35X-DA personality module.

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

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

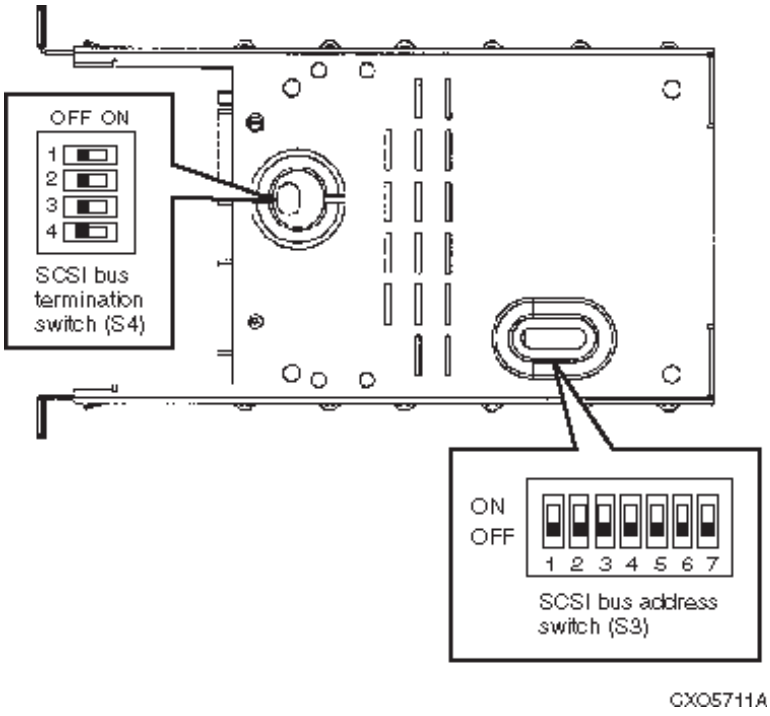

3.7 SCSI Bus Termination

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:

- Threr are only two terminators for each bus segment.
- If it is possible, bus termination must be external.

Note

Generally, when using tape loaders on a shared SCSI bus, cluster SCSI bus termination principles will not be adhered to. The SCSI bus is not terminated externally to the tape loader. Therefore, the tape loader must be on the end of the shared SCSI bus.

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 or some other method of termination. For example, PMAZC, KZTSA, and KZPSA adapters, BA350 and BA356 storage shelves, SCSI signal converters, and the UltraSCSI BA356 personality module (DS-BA35X-DA) have internal termination, but BA353 and HSZ40 controllers have automatic termination. Depending on how you set up a shared bus, you may have to enable or disable device termination.

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, you must use external termination, enabling you to detach the device without affecting the bus termination.

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 a terminator to the Y cable or trilink connector to terminate the bus. 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 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 3–4 shows a BN21V-0B Y cable, which you attach to a PMAZC TURBOchannel SCSI controller that has had its onboard termination removed.

Figure 3–4: BN21V-0B Y Cable

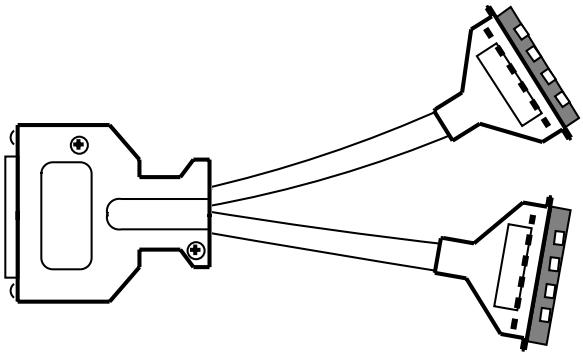

Figure 3–5 shows a BN21W-0B Y cable, which you may attach to a KZTSA or KZPSA SCSI adapter that has had its onboard termination removed. You can use the BN21W-0B Y cable with an HSZ10, HSZ40, or HSZ50 controller. Use a trilink connector with a KZTSA instead of the Y cable.

Figure 3–5: BN21W-0B Y Cable

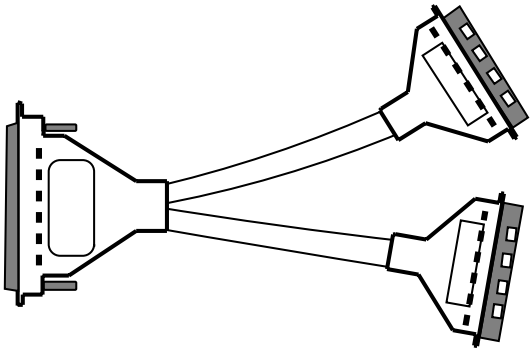

Figure 3–6 shows an HD68 trilink connector (H885), which you may attach to a KZTSA or KZPSA adapter that has its onboard termination removed, an HSZ10, HSZ40, or HSZ50 controller, or the unterminated differential side of a SCSI signal converter.

Figure 3–6: HD68 Trilink Connector (H885)

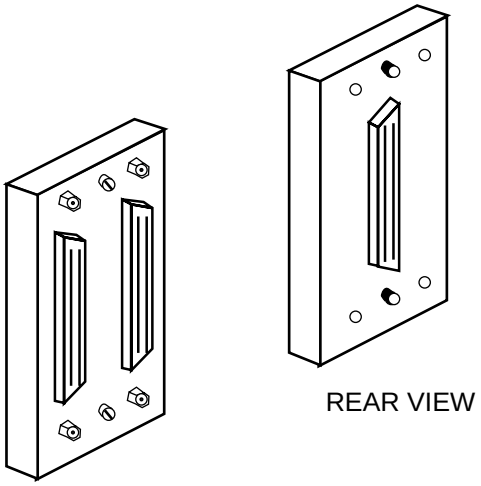

Note

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 KZPSA SCSI adapters on some AlphaServer systems.

Figure 3–7 shows a VHDCI trilink connector (UltraSCSI), which you may attach to an HSZ70 or UltraSCSI BA356 personality module that has the external SCSI bus termination disabled.

Figure 3–7: VHDCI Trilink Connector (H8861-AA)

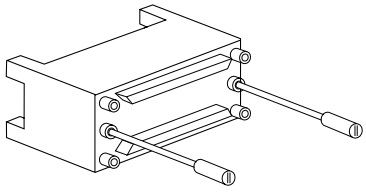

Figure 3–8 shows an Available Server shared SCSI bus that includes two DEC 3000 Model 500 systems with PMAZC TURBOchannel SCSI controllers installed, a BA350 storage shelf, and one single-ended bus. A BN21V-0B Y cable is attached to the unterminated A port in each PMAZC controller. (The unused B port is terminated with terminator part number 12-41667-01.) The BA350 storage shelf is in the middle of the bus, so its internal termination is removed. The shared bus is terminated by H8574-A or H8860-AA terminators attached to each BN21V-0B cable.

Figure 3–8: Terminating a Single-Ended Bus with Y Cables

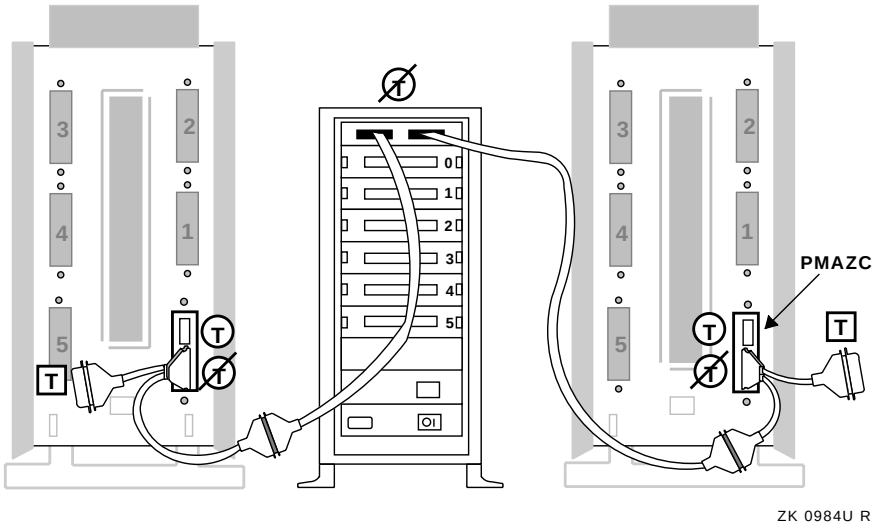

If this configuration is used and a Y cable is disconnected from a PMAZC controller, that system is not available. However, the shared SCSI bus is still operable because the bus termination is maintained, as shown in Figure 3–9.

Figure 3–9: Disconnecting a Y Cable

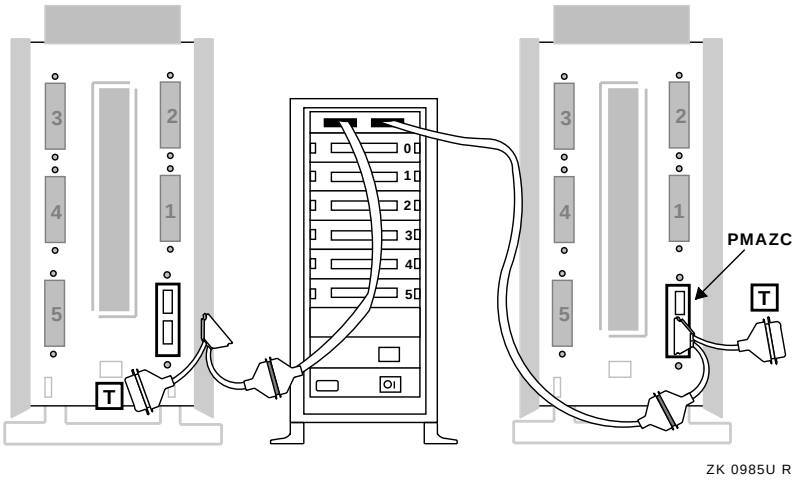

If you have the same hardware as described in Figure 3–9, you can create an alternative hardware configuration by placing one of the systems in the middle of the shared bus and the storage shelf at the end of the bus. In this case, the BA350 storage shelf internal termination is used to terminate the bus, as shown in Figure 3–10.

Figure 3–10: Terminating a Single-Ended Bus at the Storage Shelf

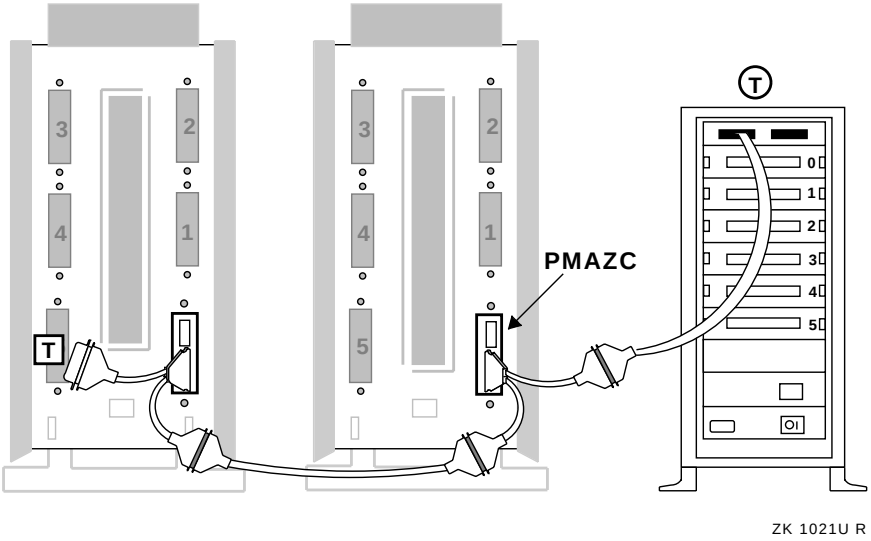

However, if you have the configuration shown in Figure 3–10 and you disconnect the cable from the BA350 storage shelf, the single-ended shared bus is not terminated and ASE operation is impaired. To be able to isolate

a single-ended storage shelf from a shared SCSI bus, you must connect the shelf to the single-ended side of a SCSI signal converter and attach a trilink connector to the differential side. Using this configuration, you can disconnect both the storage shelf and the signal converter from the shared bus without affecting ASE operation.

Figure 3–11 shows a hardware configuration that uses signal converters. The entire shared bus consists of three single-ended buses and one differential bus.

Figure 3–11: Using SCSI Signal Converters on a Bus

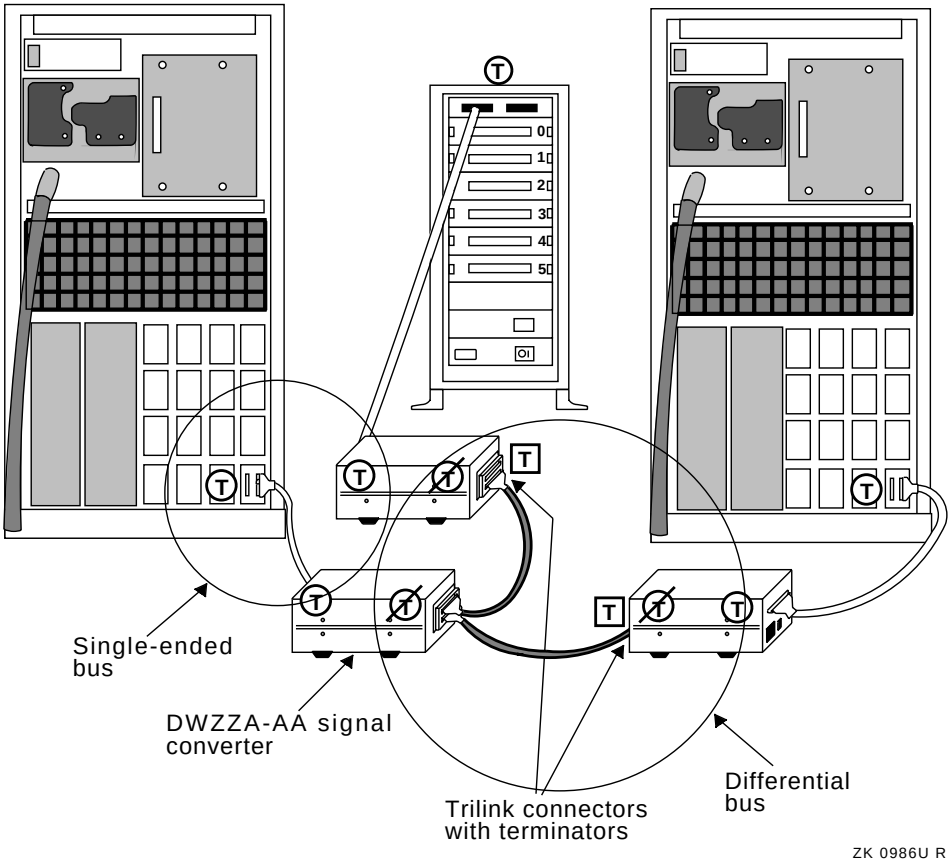

In Figure 3–11, there is a single-ended bus between each single-ended device (KZMSA adapter or BA350 storage shelf) and the single-ended side of a DWZZA-AA signal converter. The single-ended bus is terminated by the internal termination in the device and in the signal converter, as shown in Figure 3–12. If a cable is disconnected from the single-ended device, the

single-ended bus is not terminated; however, the differential bus is not affected if you have turned off the SCSI signal converter.

Figure 3–12: Terminating a Single-Ended Bus

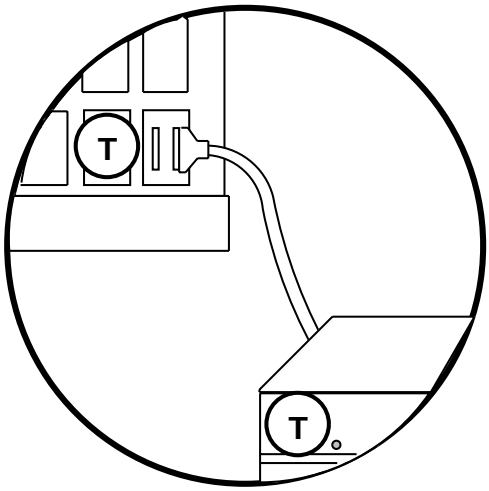

ZK-0989U-AI

In Figure 3–11, the differential bus is connected to the trilink connectors attached to the unterminated differential side of each DWZZA-AA signal converter. The differential bus is terminated by the terminators attached to the trilink connectors at the ends of the bus, as shown in Figure 3–13. If a trilink connector is disconnected from a DWZZA-AA, the differential bus is still terminated.

Figure 3–13: Terminating a Differential Bus

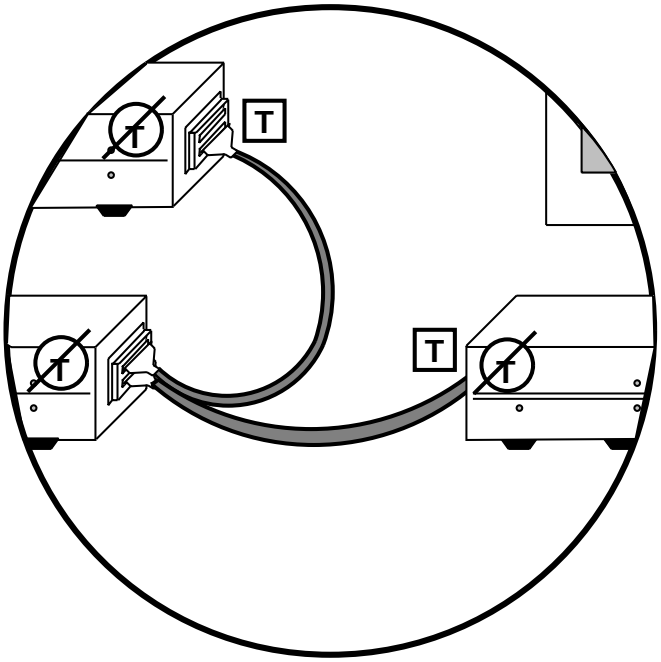

ZK-0988U-AI

3.8 Connecting Devices on a Shared SCSI Bus

There are a number of requirements for connecting devices on a SCSI bus. The requirements are:

- Use SCSI signal converters to connect devices with different transmission modes or data paths.
- Use trilink connectors or Y cables to connect devices to the shared bus so that they can be disconnected without affecting bus termination.
- Terminate each bus segment only at the ends.
- Adhere to the restrictions on bus length.

All the member systems and at least one storage shelf are connected to each shared SCSI bus. You can directly connect two devices only if they have the same transmission method (single-ended or differential) and data path (narrow or wide). The SCSI cable used to connect devices depends on the transmission mode and data path of the devices.

The following sections describe how to connect devices on a shared SCSI bus. After you understand the SCSI bus connections that you need for your configuration, you can prepare the systems and storage shelves.

3.8.1 Connecting Single-Ended Devices

Single-ended devices usually have a narrow data path. However, the BA356 storage shelf is an example of a single-ended device that has a wide data path. You cannot directly connect devices with different data paths; instead, you must use two SCSI signal converters to connect the devices. See Section 3.8.3 for information about connecting narrow and wide devices.

Figure 3–14 shows the cable required for a bus connection between two single-ended, narrow devices, which include:

- PMAZC TURBOchannel SCSI controller and KZMSA XMI to SCSI adapter
- BA350 and BA353 storage shelf
- Single-ended side of a DWZZA signal converter

Figure 3–14: Bus Connection Between Single-Ended, Narrow Devices

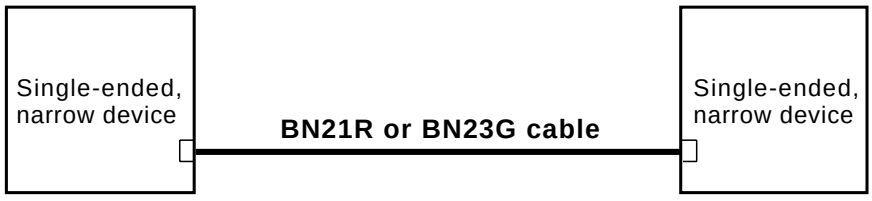

Figure 3–15 shows the cable required for a bus connection between two single-ended, wide devices, which include:

- BA356 storage shelf
- Single-ended side of a DWZZB signal converter

Figure 3–15: Bus Connection Between Single-Ended, Wide Devices

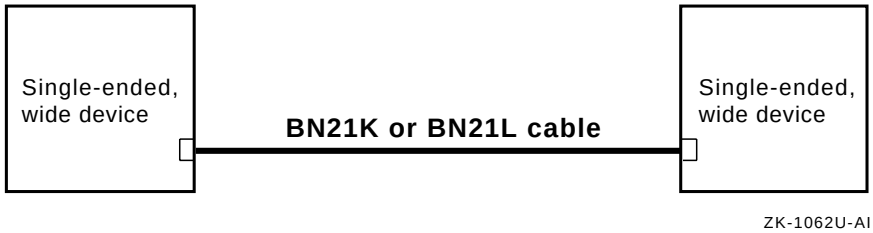

3.8.2 Connecting Differential Devices

Differential devices always have a wide data path. Figure 3–16 shows the cable required for a bus connection between two differential, wide devices, which include:

- KZTSA and KZPSA adapters
- HSZ10 and HSZ40 controllers
- Differential side of a SCSI signal converter

Figure 3–16: Bus Connection Between Differential Devices

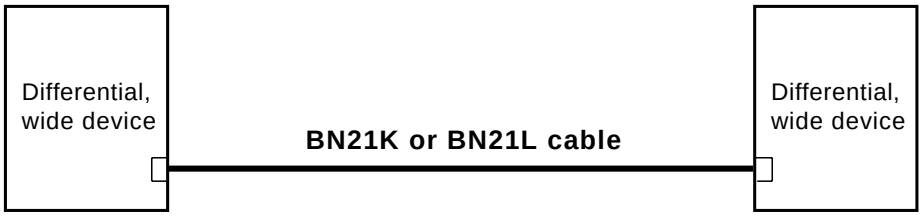

3.8.3 Connecting Single-Ended Devices to Differential Devices

Use a SCSI signal converter to connect devices with different transmission modes or data paths. See Section 3.6 for detailed information about using SCSI signal converters.

Figure 3–17 shows how to connect a device with a single-ended and narrow SCSI interface (for example, a BA350 or BA353 storage shelf) to a device with a differential and wide SCSI interface (for example, a KZPSA adapter) by using a DWZZA-VA signal converter installed in a disk slot.

Figure 3–17: Bus Connection Between a Single-Ended, Narrow Device and a Differential Device

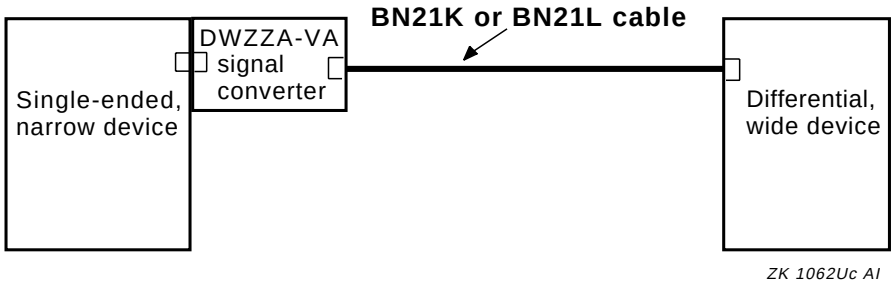

If instead of using a DWZZA-VA you use a standalone DWZZA-AA signal converter, connect the DWZZA-AA to the single-ended, narrow device using a BN21R or BN23G cable.

Figure 3–18 shows how to connect a device with a single-ended and wide SCSI interface (for example, a BA356 storage shelf) to a device with a differential and wide SCSI interface by using a DWZZB-VW signal converter, which is installed in a disk slot.

Figure 3–18: Bus Connection Between a Single-Ended, Wide Device and a Differential Device

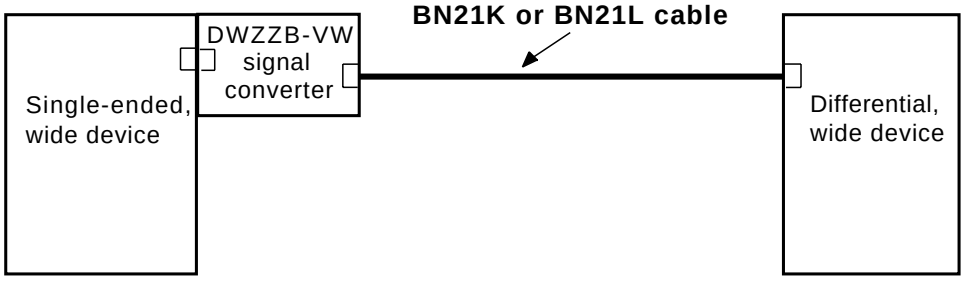

If instead of using a DWZZB-VW you use a standalone DWZZB-AA signal converter, connect the DWZZB-AA to the single-ended, wide device using a BN21K or BN21L cable.

To connect two devices with single-ended SCSI interfaces, connect each device to a SCSI signal converter, and then connect the differential sides of the signal converters by using a BN21K or BN21L cable.

4

Configuring TruCluster Hardware

This chapter describes how to prepare the systems and storage shelves for TruCluster products, including how to connect devices to a shared SCSI bus for the TruCluster Available Server Software and TruCluster Production Server Software products. 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 product. Therefore, you must have the documentation that describes how to install the individual pieces of hardware. This documentation should arrive with the hardware.

For the TruCluster MEMORY CHANNEL Software product, you only need to install the MEMORY CHANNEL hardware. An available server environment (ASE) is not supported for use with TruCluster MEMORY CHANNEL Software so a shared SCSI bus is not needed.

For the Available Server product you will install at least one shared SCSI bus to form an ASE. You can have up to 30 shared SCSI buses, but there can be only one ASE (all member systems are connected to the same shared SCSI bus(es) and the same network(s)).

For a Production Server configuration, if you have four or more systems, you can have multiple ASEs. Each member system in the ASE must be connected to the same shared SCSI buses. All systems in the cluster must be connected via the cluster interconnect (MEMORY CHANNEL).

Note

For Production Server, a cluster member does not have to be a member of an ASE, but there must be at least one ASE in the cluster.

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. Some tape loaders can only be installed at the end of a shared SCSI bus.

- 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 and connecting Y cables, trilink connectors, or 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 4.3.1.2 or Section 4.4.3 as appropriate). You can then install the Available Server or Production Server software as described in TruCluster Software Products *Administration* manual.

Read Section 4.1 and Section 4.2 then proceed to the designated section as follows:

- If you are installing hardware for the TruCluster Available Server Software or TruCluster Production Server Software product using a peripheral component interconnect (PCI) SCSI adapter, see Section 4.3 and Table 4–2.
- If you are installing hardware for the TruCluster Available Server Software product using a PMAZC, KZTSA, or KZMSA SCSI controller, see Section 4.4 and Table 4–4.
- If you are installing hardware for use with TruCluster MEMORY CHANNEL Software, see Chapter 5 for information on installing MEMORY CHANNEL hardware.

4.1 Planning Your TruCluster Hardware Configuration

Before you set up a TruCluster 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 for Production Server, the number of ASEs.*

  You can use two to four member systems for TruCluster Available Server Software and two to eight member systems for TruCluster Production Server Software. 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.

  For Production Server, if you have four or more systems, you can set up two or more ASEs. All systems in the ASE must be connected to the same shared SCSI bus(es).

Each member system must have a supported SCSI adapter for each shared SCSI bus connection. Available Server configurations must have at least one network adapter for cluster communications. For Production Server, there must be enough PCI slots for the cluster interconnect(s) (MEMORY CHANNEL) 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 Production Server, the cluster interconnect is the MEMORY CHANNEL; for Available Server the interconnect is the network. However, you can use redundant cluster interconnects to protect against an interconnect failure and for easier hardware maintenance.

  For Production Server, if you have more than two member systems, you must have one MEMORY CHANNEL hub for each interconnect.

  If you have multiple network adapters for Available Server, a failure of the primary network adapter will cause communications to fail over to another network adapter. You can have up to four network adapters monitored for failover by TruCluster Available Server Software. They are set up by the `asemgr`. See the TruCluster Software Products *Administration* manual for more information.

- *Number of shared SCSI buses in each ASE and the storage on each shared bus*.

  You can set up only one ASE in Available Server or a Production Server hardware configuration that has two or three member systems. You can set up two or more ASEs if your Production Server configuration has four or more member systems. Multiple SCSI buses increase your storage capacity in an ASE. You can connect up to 30 shared buses in each ASE.

  In addition, DEC RAID array controllers allow you to increase your storage capacity and protect against disk and controller failures. Mirroring data across shared buses by using the Logical Storage Manager (LSM) provides you with more reliable and available data.

- *Presence of highly available tape devices*.

  If you will be using highly available tape devices with either the TruCluster Available Server Software or TruCluster Production Server Software products, you must use a KZPSA PCI-to-SCSI adapter for the shared SCSI bus. The use of highly available tape devices is not supported on a shared SCSI bus with a PMAZC, KZTSA, or KZMSA SCSI controller.

- TruCluster Available Server Software and TruCluster Production Server Software do not support the ML200-*xx* series of Prestoserve NVRAM hardware.

  You cannot use Prestoserve in a TruCluster Available Server Software or TruCluster Production Server Software environment to cache I/O operations for shared SCSI disks associated with AS services (such as distributed raw disk (DRD)). Because data in the Prestoserve buffer cache of one member is not accessible to other member systems, TruCluster software cannot provide correct failover when Prestoserve is being used.

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

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.
	Set up two ASEs (Production Server only).
Member system availability	Increase the number of member systems.
Cluster interconnect availability	Use redundant cluster interconnects.
Disk availability	Mirror disks across shared buses.
	Use a DEC RAID array controller.
Shared storage capacity	Increase the number of shared buses.
	Use a DEC 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 an Available Server or Production Server installation, so you may need the firmware release notes. Obtain the firmware release notes from 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:

1. At the console prompt, or using the system startup log if the DIGITAL UNIX operating system is running, determine the drive number of the CD-ROM.
2. Boot the DIGITAL 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 DIGITAL UNIX version installed (or to be installed) into the drive.
5. Mount the CD-ROM as follows (`/dev/rz4c` is used as an example CD-ROM drive):

   ```
   # mount -rt cdfs -o noversion /dev/rz4c/mnt
   ```

6. Copy the appropriate release notes to your system disk. In this example, obtain the firmware release notes for the AlphaServer 4000/4100 from the Alpha Firmware Update 3.9 CD-ROM:

   ```
   # cp /mnt/doc/alpha4100_v48_fw_relnote.txt as4100-rel-notes
   ```

7. Unmount the CD-ROM drive.

   ```
   # umount /mnt
   ```

8. Print the release notes.

4.3 TruCluster Hardware Installation Using a PCI SCSI Adapter

Both Available Server and Production Server 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. See the TruCluster Software Products *Release Notes* and check for the revisions supported by the TruCluster products.

Follow the steps in Table 4–2 to start the procedure for TruCluster hardware installation. For Available Server or Production Server Version 1.5, the table only includes the KZPSA PCI-to-SCSI adapter. Note that you would install only MEMORY CHANNEL for TruCluster MEMORY CHANNEL

Software. For TruCluster Available Server Software, skip the first step as you do not use MEMORY CHANNEL. Also, you may save time by installing the MEMORY CHANNEL adapters, redundant network adapters (if applicable), and KZPSA SCSI adapters all at the same time.

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

The last step of Table 4–2 directs you to the procedures for highly available tape device installation. In general, tape devices on the shared SCSI must be on the end of the shared SCSI bus. Also, highly available tape devices are only supported on a shared SCSI bus attached to a KZPSA PCI-to-SCSI adapter.

Caution

Static electricity can damage modules and electronic components. DIGITAL recommends using a gounded antistatic wrist strap and a grounded work surface when handling modules.

Table 4–2: Configuring TruCluster 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 b c]
2	Install network adapters if required to provide network failover for TruCluster Available Server Software.	
	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.	Section 4.7 and *ATMworks 350 Adapter Installation and Service*
3	Install a PCI SCSI adapter for each shared SCSI bus in each member system.[d]	
	Install a KZPSA PCI-to-SCSI adapter.	*KZPSA PCI-to-SCSI Storage Adapter Installation and User's Guide*, Section 4.3.1, and Table 4–3
4	Attach a BN21W-0B Y cable to each KZPSA SCSI adapter.	

Table 4–2: Configuring TruCluster Hardware for Use With a PCI SCSI Adapter (cont.)

Step	Action	Refer to:
5	Install the remaining SCSI bus hardware (DWZZA(B), storage shelves, cables, and terminators).	Section 4.5
	Preparing a BA350 storage shelf.	Section 4.5.1
	Preparing a BA353 storage shelf.	Section 4.5.2
	Preparing a non-Ultra BA356 storage shelf.	Section 4.5.3
	Preparing an Ultra BA356 storage shelf.	Section 4.5.4
	Preparing a DEC RAID controller on a shared SCSI bus.	Section 4.5.5
6	Install the highly available tape device hardware and cables on the shared SCSI bus as follows:	Section 4.6
	TZ88	Section 4.6.1
	TZ89	Section 4.6.2
	TZ885	Section 4.6.3

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.

Install the cables for the shared SCSI bus and storage, then install the highly available tape devices and cables.

[a]This is the only step required for hardware options with the TruCluster MEMORY CHANNEL Software product.
[b]Skip this step for the TruCluster Available Server Software product.
[c]If you install additional KZPSA SCSI adapters or an extra network adapter at this time, delay testing the MEMORY CHANNEL until you have installed all hardware.
[d]For Production Server, a cluster member does not have to be a member of an ASE, but there must be at least one ASE in the cluster.

4.3.1 Installing a KZPSA PCI-to-SCSI Adapter

The KZPSA PCI-to-SCSI bus adapter is installed in a PCI slot of the supported member system.

The KZPSA is a fast, wide differential adapter with only a single port, so only one differential shared SCSI bus can be connected to a KZPSA adapter.

The KZPSA 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. For the TruCluster Production Server Software product, the KZPSA must be set to fast speed.

Use the steps in Table 4–3 to set up the KZPSA for an Available Server or Production Server configuration.

Table 4–3: Installing the KZPSA

Step	Action	Refer to:
1	Remove the KZPSA internal termination resistors, Z1, Z2, Z3, Z4, and Z5.	*KZPSA PCI-to-SCSI Storage Adapter Installation and User's Guide* and Figure 4–1
2	Install a KZPSA PCI-to-SCSI bus adapter in the PCI slot corresponding to the logical bus to be used for the shared SCSI bus.	*KZPSA PCI-to-SCSI Storage Adapter Installation and User's Guide*
3	Power up the system and use the show console commands to display the installed devices and information about the KZPSAs on the AlphaServer systems.	Section 4.3.1.2, Example 4–1, Example 4–2, Example 4–3, and Example 4–4
4	If necessary, update the KZPSA firmware to the appropriate revision.	Firmware release notes for the system (see Section 4.2) and Section 4.3.1.3
5	Set the KZPSA bus speed and SCSI bus ID as appropriate.	Section 4.3.1.4
6	Use the show console commands, and ensure that all SCSI adpaters for a shared SCSI bus are at different SCSI IDs by checking the SCSI ID on each member system. **Note** You will have problems if you have two or more SCSI adapters at the same SCSI ID.	Section 4.3.1.2
7	Return to step 4 of Table 4–2.	

The following sections describe how to install a KZPSA adapter in more detail.

4.3.1.1 KZPSA Termination Resistors

The KZPSA internal termination is disabled by removing termination resistors Z1 through Z5, as shown in Figure 4–1.

Figure 4–1: KZPSA Termination Resistors

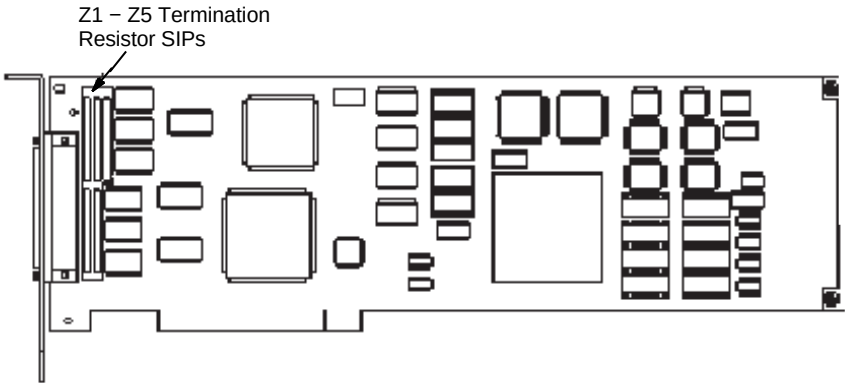

4.3.1.2 Displaying Hardware Configuration Information

Use the console commands to display information about your system configuration. For example, use the `show config` and `show device` console commands to display some or all the following information:

- Devices installed in each I/O bus slot
- Firmware and hardware revision, SCSI ID, and bus speed for each installed SCSI bus adapter
- Disk firmware revisions

Example 4–1 shows the output from the `show config` console command entered on an AlphaServer 8200 system.

Example 4–1: Displaying an AlphaServer 8200 Configuration

```
>>> 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
```

Example 4–1: Displaying an AlphaServer 8200 Configuration (cont.)

```
 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+  DECpci MC         181011    000B     mc0
⋮
```

Example 4–2 shows the output from the `show device` console command entered on an AlphaServer 8200 system.

Example 4–2: 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              HSZ50-AX X29Z
dkf1.0.0.1.1      DKf1              HSZ50-AX X29Z
dkf2.0.0.1.1      DKf2              HSZ50-AX X29Z

polling for units on kzpsa0, slot 2, bus 0, hose1...
kzpsa0.4.0.2.1   dke     TPwr 1 Fast 1 Bus ID 7   F01 A10
dkg100.1.0.2.1   DKg100                 RZ26N   0586
dkg200.2.0.2.1   DKg200                 RZ26    392A
dkg300.3.0.2.1   DKg300                 RZ26N   0586
⋮
```

Example 4–2 shows that the system has a KZPSA PCI-to-SCSI adapter installed, with fast SCSI bus speed enabled, a SCSI ID of 7, hardware revision level F01, and firmware revision A10 (which needs to be updated). The example also shows three disks on the SCSI bus.

The indication that TPwr equals 1 indicates that the KZPSA is providing termination power to the SCSI bus, which is the correct configuration. All

SCSI bus adapters in a TruCluster configuration should be generating termination power.

Example 4–3 shows the output from the `show device` console command entered on an AlphaServer 2000.

Example 4–3: Displaying Devices on an AlphaServer 2000

```
P00 >>>show device
dka0.0.0.1.0          DKA0        RZ28              D41C
dka100.1.0.1.0        DKA100      RZ28              D41C
dka200.2.0.1.0        DKA200      RZ28              D41C
dka300.3.0.1.0        DKA300      RZ28              D41C
dka600.6.0.1.0        DKA600      RRD43             1084
dkb0.0.0.6.0          DKB0        HSZ40-Bx          V21Z
dkb100.1.0.6.0        DKB100      RZ28              D41C
dkb300.3.0.6.0        DKB300      RZ28B             0006
dva0.0.0.1000.0       DVA0        RX26
ewa0.0.0.0.0          EWA         08-00-2B-E2-7C-81
pka0.7.0.1.0          PKA0        SCSI Bus ID 7
pkb0.6.0.6.0          PKB0        SCSI Bus ID 4     F01  A10
P00>>>
```

Example 4–3 shows that the firmware for KZPSA `pkb` is also at revision A10 and needs to be updated. Also, the SCSI bus speed is not shown in the `show device` output for the AlphaServer 2000. To obtain the bus speed for this system, use the `show pkb*` console command as shown in Example 4–4.

Note

We are not interested in `pka` as that is the internal SCSI bus.

Example 4–4: Displaying KZPSA Bus Speed With the show pkb* Command

```
P00 >>>show pkb*
pkb0_fast       0
pkb0_host_id    4
pkb0_termpwr    1
P00>>>
```

Example 4–4 shows that the bus speed is slow (`pkb0_fast` is equal to 0). The `pkb0_termpwr` variable being equal to 1 indicates that the KZPSA is generating termination power to the SCSI bus (as it should be).

4.3.1.3 Updating the KZPSA Adapter Firmware

You must check, and update as necessary, the system and SCSI adapter firmware. The firmware may be out of date. Read the release notes for the applicable system/SCSI adapter.

If the Standard Reference Manual console (SRM), Advanced RISC Computing console (ARC), or KZPSA firmware is not current, as the KZPSA in the previous examples, read the firmware release notes, then 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 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 `pkb0`:

```
UPD>  update pkb0
```

After updating the firmware and verifying this with the `verify` command, reset the system by cycling the power.

4.3.1.4 Setting KZPSA SCSI Bus ID and Speed

If the SCSI ID is not correct, or if it was reset to 7 by the firmware update utility, or you need to change the KZPSA speed, use the `set` console command for most AlphaServers. Use the LFU utility for the AlphaServer 8200/8400s. Select the `modify` command (for instance `modify kzpsa2`).

Set the SCSI bus ID with the `set` command as shown in the following example:

```
>>> set pkn_0_host_id #
```

The *n* specifies the KZPSA ID, which you obtain from the `show device` console command. 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.

Example 4–5 shows how to determine the present SCSI ID and bus speed, and then set the KZPSA SCSI ID to 6 and bus speed to fast for `pkb0`.

Example 4–5: Setting KZPSA SCSI Bus ID and Speed

```
P00>>> show pkb0_host_id
4
P00>>> show pkb0_fast
0
P00>>> set pkb0_host_id 6
P00>>> set pkb0_fast 1
P00>>> show pkb0_host_id
6
P00>>> show pkb0_fast
1
```

4.4 Available Server Hardware Installation Using a PMAZC, KZTSA, or KZMSA SCSI Controller

This section describes how to set up the TruCluster Available Server Software hardware for installations that use the PMAZC, KZTSA, or KZMSA SCSI controllers.

Note

The use of a highly available tape device is not supported with PMAZC, KZTSA, or KZMSA SCSI controllers.

To begin your configuration, find the SCSI controller in the left column of Table 4–4 and refer to information shown in the right column. Read the preparatory text and follow the steps in the table to generate the specified hardware configuration.

You will install and set up the SCSI controller, and then go to a section that covers the type of storage you are using.

If you have a configuration with different types of SCSI controllers, you will have to use a different procedure for each controller.

Also, if you are using multiple types of storage, you will have to address each type individually.

Table 4–4: Configuring Hardware for Available Server Configurations Using PMAZC, KZTSA, or KZMSA SCSI Controllers

If your controller is:	Refer to:
PMAZC	Section 4.4.1 and Table 4–5
KZTSA	Section 4.4.4 and Table 4–6
KZMSA	Section 4.4.5 and Table 4–8

4.4.1 Configuring an Available Server With PMAZCs

Note

If you have a TruCluster Available Server Software configuration that includes a PMAZC SCSI controller, you can have only three systems in the configuration.

Table 4–5 covers the steps necessary to configure PMAZC SCSI controllers in an Available Server configuration.

Table 4–5: Configuring an Available Server Using the PMAZC TURBOchannel SCSI Controller

Step	Action	Refer to:
1	For each system using a PMAZC on the shared bus, shut down the system and install the PMAZC.	*Dual SCSI Module (PMAZC-AA)*
	If necessary, install jumper W1 to enable the `setld` console utility to set the PMAZC SCSI ID, bus speed, or to update the firmware.	Figure 4–2
2	Turn on the system power and set the PMAZC SCSI ID and speed as necessary.	Section 4.4.3, Example 4–6, Example 4–7, and Example 4–8
3	If the firmware has to be updated, boot from the Alpha Systems Firmware CD-ROM.	Firmware release notes for the system (see Section 4.2)
4	Turn off the system power and remove PMAZC jumper W1. Store it on an empty jumper rest.	Figure 4–2
	If your shared SCSI bus will be single-ended (no DWZZAs), disable the PMAZC internal termination by removing the jumper for the appropriate port (W2 (port A) and W3 (port B)).	
	If the shared SCSI bus will be differential, ensure that the appropriate PMAZC termination jumpers (W2 (port A) or W3 (port B)) are installed to provide termination for one end of each single-ended SCSI bus.	
5	Install any network adapters, if required, to provide network failover.	
	Install the Ethernet or FDDI network adapters.	User's guide for the applicable Ethernet or FDDI adapter, and the user's guide for the applicable system
6	Install the cables, DWZZAs, DWZZBs, trilinks (or Y cables) to create the shared SCSI bus with the storage shelves as follows:	
	Preparing a BA350 storage unit.	Section 4.5.1
	Preparing a BA353 storage unit.	Section 4.5.2
	Preparing a Non-Ultra BA356 storage unit.	Section 4.5.3

Table 4–5: Configuring an Available Server Using the PMAZC TURBOchannel SCSI Controller (cont.)

Step	Action	Refer to:
	Preparing an Ultra BA356 storage unit.	Section 4.5.4
	Preparing a DEC RAID controller on a shared SCSI bus.	Section 4.5.5

4.4.2 PMAZC Dual SCSI Module Jumpers

Figure 4–2 shows the jumpers on a PMAZC SCSI controller.

Figure 4–2: PMAZC Dual SCSI Module Jumpers

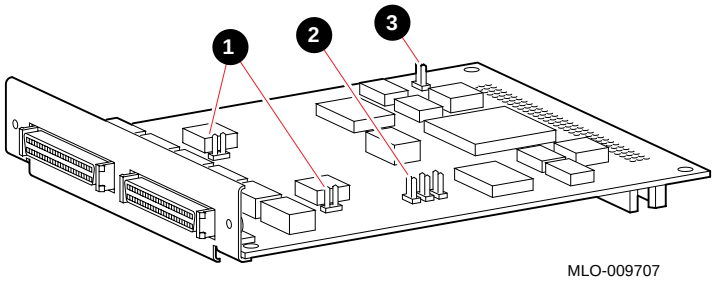

The following list explains the callouts in Figure 4–2:

1. W2 and W3 terminator jumpers. When installed, these jumpers provide the required termination to one end of the two SCSI buses. W2 is for port A, and is the leftmost jumper in the figure. W3 is for port B. They are shown as being removed.
2. Jumper rests used to store jumpers that have been removed.
3. W1 is the flash memory write jumper. Do not install it except to update the ROM code or when using the `setid` utility to change the SCSI ID or bus speed.

4.4.3 Verifying and Setting PMAZC and KZTSA SCSI ID and Bus Speed

This section provides examples to display and change the SCSI ID or bus speed for a PMAZC or KZTSA TURBOchannel SCSI controller.

To display the SCSI ID and bus speed for a PMAZC or KZTSA SCSI ID, shut down the system. Use the console `show config` command to

determine the PMAZC or KZTSA configurations. Example 4–6 shows that the DEC 3000 Model 500 has PMAZC-AA SCSI controllers in two TURBOchannel slots, TC0 and TC1, and a KZTSA in TURBOchannel slot TC3.

Example 4–6: Displaying a DEC 3000 Configuration

```
>>>show config
                    DEC 3000  -  M500
              Digital Equipment Corporation
     VPP PAL X5.48-82000101/OSF PAL X1.35-82000201 -
                    Build on 20-JUL-1994 11:07:03.31
TCINFO      DEVNAM       DEVSTAT
------      --------     -------
                 CPU     OK KN15-AA  -V5.1-5748-t19D-sV;.?
                                     -DECchip 21064 P2.1
                 OSC     150 MHZ
                ASIC     OK
                 MEM     OK
8
                 CXT     OK
7
                 NVR     OK
                 SCC     OK
                  NI     OK
                ISDN     OK
6
                SCSI     OK
1-PMAZC-AA      TC1
0-PMAZC-AA      TC0
3-KZTSA-AA      TC3
```

To display the SCSI ID or bus speed for a specific PMAZC or KZTSA, use the `t tc# cnfg` console command shown in Example 4–7 and Example 4–9. In this command, the number sign (#) specifies the TURBOchannel slot number. Example 4–7 shows that the PMAZC-AAs in TURBOchannel slot 0 and slot 1 both have SCSI IDs of 7 and are set to slow speed.

Example 4–7: Displaying the PMAZC SCSI ID and Bus Speed

```
>>> t tc0 cnfg
DEC PMAZC-AA V2.0  Port A Slow  Port B Slow (Dual SCSI [53CF96])
 BOOTDEV   ADDR   DEVTYPE   NUMBYTES   RM/FV   WP   DEVNAM   REV
 -------   ----   -------   --------   -----   --   ------   ---
..HostID.. A/7    INITR
..HostID.. B/7    INITR
```

Example 4–7: Displaying the PMAZC SCSI ID and Bus Speed (cont.)

```
>>> t tc1 cnfg
DEC PMAZC-AA V2.0  Port A Slow  Port B Slow (Dual SCSI [53CF96])
 BOOTDEV   ADDR   DEVTYPE   NUMBYTES   RM/FX   WP   DEVNAM   REV
 -------   ----   -------   --------   -----   --   ------   ---
..HostID.. A/7    INITR
 DKB000    B/0/0   DISK          1GB      FX        RZ26    T386
 DKB100    B/1/0   DISK          1GB      FX        RZ26    392A
..HostID.. B/7    INITR
```

To set the SCSI ID or bus speed for both PMAZC ports, use the following `t` command; for example:

```
>>> t tc# setid x y
```

The number sign (#) is the TURBOchannel slot, `x` is the SCSI ID or speed (s = slow and f = fast) for port A and `y` is the SCSI ID or speed for port B.

Example 4–8 shows the commands to set the SCSI ID to 6 for both ports and to set the speed to fast for the PMAZC in TURBOchannel slot 1, and to then verify the changes.

Note

If the PMAZC W1 jumper is not on, you will receive the "`Put JMPR ON...`" console message.

Example 4–8: Setting the PMAZC SCSI ID and Bus Speed

```
>>> t tc1 setid 6 6
Precharging
.........................................................
Erasing
.........................................................
Programming
.........................................................
Checksum GOOD
>>> t tc1 setid f f
Precharging
.........................................................
Erasing
.........................................................
Programming
.........................................................
Checksum GOOD
>>> t tc1 cnfg
DEC PMAZC-AA V2.0  Port A Fast Port B Fast (Dual SCSI [53CF96])
 BOOTDEV   ADDR   DEVTYPE   NUMBYTES   RM/FV  WP   DEVNAM   REV
```

Example 4–8: Setting the PMAZC SCSI ID and Bus Speed (cont.)

```
 -------   ----   -------   --------   -----  --   ------   ---
..HostID.. A/6    INITR
..HostID.. B/6    INITR
```

Example 4–9 shows how to display the SCSI ID for the KZTSA in TURBOchannel slot TC3 (from Example 4–6). Note that the KZTSA only has one port.

Example 4–9: Displaying the KZTSA SCSI ID and Bus Speed

```
>>> t tc3 cnfg
DEC KZTSA-AA A09           (SCSI = 7, Slow)

-------------------------------------------------
DEV      PID              VID      REV     SCSI DEV
======= ================ ======== ======= ========
dka0000 HSZ40-Bx (C) DEC DEC      V21Z    DIR
dka0100 RZ28     (C) DEC DEC      442D    DIR
dka0300 RZ28B    (C) DEC DEC      0006    DIR
>>>
```

Set the KZTSA SCSI ID or bus speed as you did with a PMAZC. Note however, that the KZTSA has only one port. Example 4–10 shows how to set the SCSI ID to 5 and set the bus speed to fast for the KZTSA in TURBOchannel slot 1.

Note in Example 4–10, that after you change the KZTSA SCSI ID, you must reset the SCSI bus (`init`) to effect the ID change. A SCSI bus reset is not needed to change the speed.

Example 4–10: Setting the KZTSA SCSI ID and Bus Speed

```
>>> t tc3 setid 5
>>> t tc3 setid f
>>> t tc3 cnfg
DEC KZTSA-AA A09           (SCSI = 7, Fast)

-------------------------------------------------
DEV      PID              VID      REV     SCSI DEV
======= ================ ======== ======= ========
dka0000 HSZ40-Bx (C) DEC DEC      V21Z    DIR
dka0100 RZ28     (C) DEC DEC      D41C    DIR
dka0300 RZ28B    (C) DEC DEC      0006    DIR
>>> init
```

Example 4–10: Setting the KZTSA SCSI ID and Bus Speed (cont.)

```
>>> t tc3 cnfg
DEC KZTSA-AA A09           (SCSI = 5, Fast)

-----------------------------------------------
DEV      PID              VID      REV     SCSI DEV
======= ================ ======== ======= ========
dka0000 HSZ40-Bx (C) DEC DEC      V21Z    DIR
dka0100 RZ28     (C) DEC DEC      D41C    DIR
dka0300 RZ28B    (C) DEC DEC      0006    DIR
>>>
```

4.4.4 Configuring an Available Server With KZTSA SCSI Adapters

This section is specific to a Available Server configuration using KZTSA TURBOchannel-to-SCSI adapters. The KZTSA is a differential single-channel SCSI adapter. Using a KZTSA in a DEC 3000 system simplifies hardware configuration and reduces the total number of required DWZZAs. When used with a DEC RAID subsystem, you do not have to use a DWZZA in the Available Server configuration.

Use Table 4–6 to set up a Available Server configuration using KZTSA TURBOchannel-to-SCSI adapters.

Table 4–6: Configuring an Available Server With KZTSA TURBOchannel SCSI Adapters

Step	Action	Refer to:
1	For each DEC 3000 system that will have a KZTSA on the shared SCSI bus, shut down the system and install the KZTSA.	*KZTSA SCSI Storage Adapter Installation and User's Guide*
	Disable the KZTSA internal SCSI termination by removing the J1, J2, J3, J6, and J7 terminator packs.	Figure 4–3
2	Install any network adapters, if required, to provide network failover.	
	Install the Ethernet or FDDI network adapters.	User's guide for the applicable Ethernet or FDDI adapter, and the user's guide for the applicable system
3	The default SCSI ID for a KZTSA is 7. Turn on the system power and set the KZTSA SCSI ID if necessary.	Example 4–6, Example 4–9, and Example 4–10

Table 4–6: Configuring an Available Server With KZTSA TURBOchannel SCSI Adapters (cont.)

Step	Action	Refer to:
4	If the firmware has to be updated, boot from the Alpha Systems Firmware Update CD-ROM and update the firmware.	Firmware release notes (see Section 4.2)
5	Install a BN21W-0B Y cable or H885-AA trilink connector on the external SCSI connector of each KZTSA.	
6	Install the cables, DWZZAs, DWZZBs, trilinks (or Y cables) to create the shared SCSI bus with the storage shelves as follows:	
	Preparing a BA350 storage unit.	Section 4.5.1
	Preparing a BA353 storage unit.	Section 4.5.2
	Preparing a Non-Ultra BA356 storage unit.	Section 4.5.3
	Preparing an Ultra BA356 storage unit.	Section 4.5.4
	Preparing a DEC RAID controller on a shared SCSI bus.	Section 4.5.5

Figure 6–5 shows a configuration that is similar to a configuration with two DEC 3000/500 systems using KZTSA TURBOchannel SCSI adapters. Replace the AlphaServers with two DEC 3000/500 systems with installed KZPSAs and connect the trilinks to the KZTSAs. All cabling is the same.

Figure 6–6 shows an Available Server configuration with two AlphaServer 2100 systems with KZPSA PCI SCSI adapters on a shared bus with a BA350 storage shelf. The only difference for a DEC 3000/500 with a KZTSA TURBOchannel SCSI adapter is to swap the AlphaServer 2100 systems for the DEC 3000 systems and connect the trilink connectors to the KZTSAs.

The KZTSA jumpers and LEDs are shown in Figure 4–3.

Figure 4–3: KZTSA Jumpers and Termination

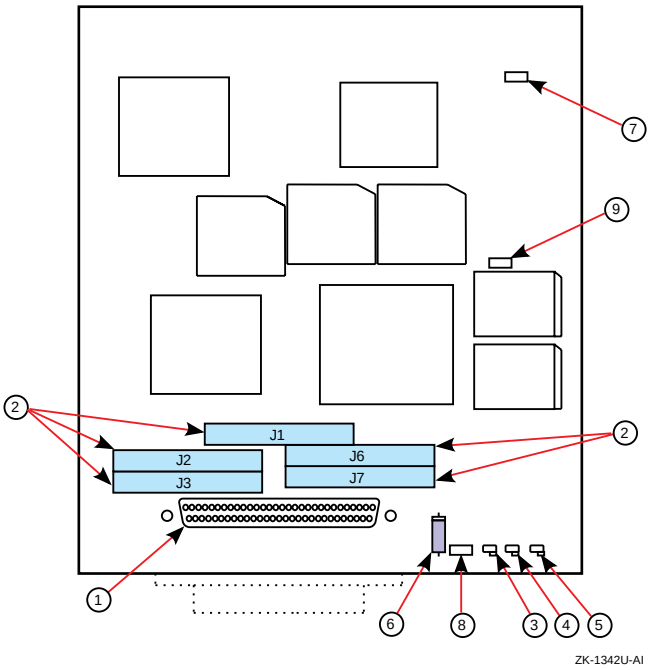

The following list describes the callouts in Figure 4–3:

1. Internal SCSI bus P-connector
2. Near-end SCSI bus terminator packs
3. Yellow LED—Power-on self-test passed
4. Red LED—Power-on self-test failed
5. Green LED—SCSI bus terminator power is functional
6. Jumper W1—Installed: Inline fuse that protects the onboard SCSI bus terminator power supply
7. Jumper W2—Not installed: Manufacturing use only
8. Jumper W3—Installed: Enables terminator power onto the SCSI bus
9. Jumper W4—Not installed: Manufacturing use only

4.4.5 Configuring an Available Server With KZMSA XMI-to-SCSI Adapters

The KZMSA is an XMI-to-SCSI adapter used in DEC 7000 or DEC 10000 systems. It is a dual channel, single-ended SCSI controller. You cannot remove the KZMSA internal termination. It must be used with a

DWZZA-AA signal converter to provide the proper SCSI bus termination, and to allow a KZMSA and its associated system to be isolated for maintenance purposes. When using a KZMSA for a shared SCSI bus in an Available Server configuration, make sure that you are connecting the bus to the same KZMSA channel as on other KZMSA or PMAZC SCSI controllers.

Each KZMSA used for a shared SCSI bus in an Available Server configuration must have the revision F03 boot ROM. If necessary, a revision F01 or F02 boot ROM must be replaced with a revision F03 boot ROM. Table 4–7 shows the part numbers for the various revisions of KZMSA boot ROMs.

Table 4–7: KZMSA Boot ROM Part Numbers

Part Number	Revision
23-368E9-01	F01
23-386E9-01	F02
23-419E9-01	F03

You can determine the KZMSA hardware revision by booting the LFU utility and using the console commands, or by examining the 23-class part number printed on the boot ROM located at module position E7. The LFU utility is covered in Section 4.4.6.

You can only use KZMSAs with Revision D NCR 53C710 chips in an Available Server configuration. The chip must have part numbers 609-3400546 or 609-3400563.

Follow the steps in Table 4–8 to set up an Available Server configuration using KZMSA XMI-to-SCSI adapters.

Note

You can have only three systems in an Available Server configuration that include a KZMSA XMI-to-SCSI adapter.

Table 4–8: Configuring an Available Server using a KZMSA XMI-to-SCSI Adapter

Step	Action	Refer to:
1	For each DEC 7000 or DEC 10000 system using a KZMSA on the shared bus, shut down the system and install the KZMSA in an XMI slot.	*KZMSA Adapter Installation Guide*
2	Install the Ethernet or FDDI network adapters, as required, to provide network failover.	User's guide for the applicable Ethernet or FDDI adapter and the user's guide for the applicable system
3	Boot the Loadable Firmware Update (LFU) utility to configure the KZMSA hardware.	Firmware release notes for the system (see Section 4.2), Section 4.4.6, and Example 4–11
	Update the KZMSA firmware if necessary.	Example 4–12 and Example 4–13
	Set the SCSI IDs for the KZMSA.	Example 4–12 and Example 4–14
	Enable the `Disable Reset` configuration option for any KZMSA channel that will be used for a shared SCSI bus, and disable the option for any channel not used on a shared SCSI bus.	Example 4–12 and Example 4–14
	Enable (disable) fast SCSI speed for the KZMSA.	Example 4–12 and Example 4–14
4	Because the KZMSA internal termination cannot be removed, you will need a DWZZA-AA for each KZMSA XMI-to-SCSI adapter in the Available Server configuration.	
	For the DWZZA-AA to be connected to the KZMSA, ensure that the single-ended SCSI jumper, J2, is installed and remove the five differential terminator resistor SIPs.	
5	For each KZMSA used in the Available Server configuration, install a BN21R or BN23G cable between the KZMSA connector for the appropriate channel and the DWZZA-AA single-ended connector.	
6	Install the cables, DWZZAs, DWZZBs, trilinks (or Y cables) to create the shared SCSI bus with the storage shelves as follows:	
	Preparing a BA350 storage unit.	Section 4.5.1

Table 4–8: Configuring an Available Server using a KZMSA XMI-to-SCSI Adapter (cont.)

Step	Action	Refer to:
	Preparing a BA353 storage unit.	Section 4.5.2
	Preparing a Non-Ultra BA356 storage unit.	Section 4.5.3
	Preparing an Ultra BA356 storage unit.	Section 4.5.4
	Preparing a DEC RAID controller on a shared SCSI bus.	Section 4.5.5

Figure 6–4 shows an Available Server configuration with a DEC 7000 system with a KZMSA XMI-to-SCSI adapter on a shared bus with a DEC 3000/500 with PMAZC TURBOchannel SCSI controller and a BA350 storage shelf.

Figure 6–7 shows an Available Server with two DEC 7000 systems with KZMSA XMI SCSI adapters on a shared bus with an HSZ40 unit.

4.4.6 Using the LFU with a DEC 7000 or DEC 10000 System

If you are using a DEC 7000 or DEC 10000 system with a KZMSA in an Available Server configuration, you may have to update the KZMSA firmware, change the SCSI ID or bus speed, or enable or disable the `Disable Reset` option.

For the DEC 7000 and DEC 10000 systems, use the Loadable Firmware Update (LFU) utility to perform these hardware tasks. Shut down the system then load the LFU, as shown in Example 4–11.

Example 4–11: Booting the LFU Utility

```
>>> show device kzmsa0    1
polling for units on kzmsa0, slot2, xmi0...
dka100.1.0.2.0 dka100 RRD42
>>> boot -flag 0,80 DKA100    2
>>> kzmsa_lfu.exe    3
Booting...

****** Loadable Firmware Update Utility ******

------------------------------------------------------------
Function    Description
------------------------------------------------------------

Display     Displays the system's configuration table.
Exit        Returns to loadable offline operating environment.
```

Example 4–11: Booting the LFU Utility (cont.)

```
List        Lists the device types and firmware revisions
            supported by this revision of LFU.
Modify      Modifies port parameters and device attributes.
Show        Displays device mnemonic, hardware and firmware
            revisions.
Update      Replaces current firmware with loadable data image.
Verify      Compares loadable and device images.
? or Help   Scrolls the function table.
------------------------------------------------------------

Function?     4
```

The following list describers the callouts in Example 4–11:

1. At the console prompt, use the `show device kzmsa0` command to determine the name of the RRD42 drive.

   Load the CD-ROM into an RRD42 caddy and insert the caddy into the RRD42 drive. The CD-ROM that includes both the LFU utility and the KZMSA revision 5.6 firmware has the label `Alpha Systems Firmware Update 4.0`.

2. Boot the LFU utility.

3. When prompted, specify the name of the secondary bootstrap file, `kzmsa_lfu.exe`.

4. At the LFU `Function?` prompt, enter the command for the task you want to perform.

You can display information about the hardware configuration with the LFU utility using the `display` command, as shown in Example 4–12.

Example 4–12: Using the LFU Utility to Display Hardware Configuration

```
Function? display    1
    Name      Type     Rev      Mnemonic FW Rev   HW Rev
LSB
0+  KN7AA     (8001)   0000     kn7aa0   1.0      E04
5+  MS7AA     (4000)   0000     ms7aa0   N/A      A01
7+  MS7AA     (4000)   0000     ms7aa1   N/A      A01
8+  IOP       (2000)   0001     iop0     N/A      A

C0 XMI               xmi0
8+  DWLMA     (102A)   A5A6     dwlma0   N/A      A
B+  KZMSA     (0C36)   5143     kzmsa0   4.3      F01    2
C+  KZMSA     (0C36)   5143     kzmsa1   4.3      F01    2
E+  DEMNA     (0C03)   060B     demna0   6.8

C1 XMI
1+  KZMSA     (0C36)   5343     kzmsa2   4.3      F03    3
```

Example 4–12: Using the LFU Utility to Display Hardware Configuration (cont.)

```
2+  KZMSA     (0C36)   5343     kzmsa3   4.3      F03     3
8+  DWLMA     (102A)   A5A6     dwlma1   N/A      A
Function?
```

The following list describes the callouts in Example 4–12:

1. Enter the `display` command to display the configuration.
2. `KZMSA0` and `KZMSA1` have the revision 4.3 firmware and the revision F01 hardware.
3. `KZMSA2` and `KZMSA3` have the revision 4.3 firmware and the revision F03 hardware.

If the KZMSA firmware is not up to the correct revision, use the LFU utility `update` command to update it. Note that the CD-ROM containing the firmware must be installed in the RRD42. The `update` command has the following format:

`update kzmsa #`

The number sign (#) indicates the number of the KZMSA, which is to have the firmware updated.

Example 4–13 shows how to update the firmware for `KZMSA2` to Version 5.6.

Example 4–13: Using the LFU Utility to Update KZMSA Firmware

```
Function? update kzmsa2    1

Update kzmsa2? [Y/(N)]  Return

WARNING: updates may take several minutes to
         complete for each device.

DO NOT ABORT!
kzmsa2 Updating to 5.6... Reading Device...
       Verifying 5.6... PASSED.

Function?  display    2
    Name      Type     Rev      Mnemonic FW Rev   HW Rev
LSB
0+  KN7AA     (8001)   0000     kn7aa0   1.0      E04
5+  MS7AA     (4000)   0000     ms7aa0   N/A      A01
7+  MS7AA     (4000)   0000     ms7aa1   N/A      A01
8+  IOP       (2000)   0001     iop0     N/A      A
```

Example 4–13: Using the LFU Utility to Update KZMSA Firmware (cont.)

```
C0 XMI                xmi0
8+  DWLMA     (102A)   A5A6     dwlma0   N/A      A
B+  KZMSA     (0C36)   5143     kzmsa0[3] 4.3     F01
C+  KZMSA     (0C36)   5143     kzmsa1[3] 4.3     F01
E+  DEMNA     (0C03)   060B     demna0    6.8

C1 XMI
1+  KZMSA     (0C36)   5356     kzmsa2[4] 5.6     F03
2+  KZMSA     (0C36)   5343     kzmsa3[5] 4.3     F03
8+  DWLMA     (102A)   A5A6     dwlma1    N/A     A
Function?
```

The following list describes the callouts in Example 4–13:

1. Update the firmware for `kzmsa2`.
2. Display the configuration to verify that the firmware has been updated.
3. `kzmsa0` and `kzmsa1` are still at firmware revision 4.3.
4. `kzmsa2` is now at firmware revision 5.6.
5. `kzmsa3` is still at firmware revision 4.3.

Use the LFU utility `modify kzmsa #` command to display detailed information about a specific KZMSA and to:

- Change the SCSI ID
- Enable or disable the fast SCSI option for a particular channel
- Enable or disable the `Disable Reset` option

Example 4–14 shows how to use the LFU utility `modify` command to display detailed information, set the SCSI ID, enable fast SCSI bus speed, and enable the `Disable Reset` option for `kzmsa2`.

Example 4–14: Using the LFU Utility to Modify KZMSA Options

```
Function?  modify kzmsa2     [1]
kzmsa2
Local Console: ENABLED
Log Selftest Errors: ENABLED
Log NRC 53C710 RBD Errors: ENABLED
Log XMI RBD Errors: ENABLED
Log XZA RBD Errors: ENABLED
RBD Error Logging: DISABLED
RBD Error Frame Overflow: DISABLED Read Only
Hard Error Frame Overflow: DISABLED Read Only
Soft Error Frame Overflow: DISABLED Read Only
```

Example 4–14: Using the LFU Utility to Modify KZMSA Options (cont.)

```
FW Update Error Frame Overflow: DISABLED Read Only
Disable Reset Channel 0: DISABLED  2
Disable Reset Channel 1: DISABLED  2
Chnl 0 Fast SCSI: DISABLED  3
Chnl 1 Fast SCSI: DISABLED  3
Channel_0 ID: 07  4
Channel_1 ID: 07  4
Module Serial Numbers: *SG90XXX455*
Do you wish to modify any of these parameters? [y/(n)]  Return
Local Console: ENABLED Change? [y/(n)]  Return
Log Selftest Errors: ENABLED Change? [y/(n)]  Return
.
.
.
Disable Reset Channel 0: DISABLED Change? [y/(n)] y    5
Disable Reset Channel 1: DISABLED Change? [y/(n)] y    5
Chnl 0 Fast SCSI: DISABLED Change? [y/(n)] y    6
Chnl 1 Fast SCSI: DISABLED Change? [y/(n)] y    6
Channel_0 ID: 07 Change? [y/(n)] y    7
Valid ID is a value from 0 to 7.
Enter new Channel ID: 6    7
Channel_1 ID: 07 Change? [y/(n)] y 7
Valid ID is a value from 0 to 7.
Enter new Channel ID: 6    7
Module Serial Numbers: *SG90XXX455* Change? [y/(n)] n
Local Console: ENABLED
Log Selftest Errors: ENABLED
.
.
.
Log NRC 53C710 RBD Errors: ENABLED
Log XMI RBD Errors: ENABLED
Log XZA RBD Errors: ENABLED
RBD Error Logging: DISABLED
RBD Error Frame Overflow: DISABLED Read Only
Hard Error Frame Overflow: DISABLED Read Only
Soft Error Frame Overflow: DISABLED Read Only
FW Update Error Frame Overflow: DISABLED Read Only
Disable Reset Channel 0: ENABLED 8
Disable Reset Channel 1: ENABLED 8
Chnl 0 Fast SCSI: ENABLED  9
Chnl 1 Fast SCSI: ENABLED  9
Channel_0 ID: 06 10
Channel_1 ID: 06 10
Module Serial Numbers: *SG909T1455*
Modify kzmsa2 with these parameter values? [y/(n)]  y    11
Function?  exit
```

The following list describes the callouts in Example 4–14:

1. Execute the LFU `modify` command to modify the options for `kzmsa2`. The present options are displayed first.
2. The `Disable Reset` option for both channels is disabled.
3. The fast SCSI option is disabled for both channels.
4. The ID for both channels is 7.
5. Enable the `Disable Reset` option for channels 0 and 1.
6. Enable the `Fast SCSI` option for channels 0 and 1.
7. Change the SCSI ID for channels 0 and 1 to 6.
8. The LFU utility is set up to enable the `Disable Reset` option.
9. The LFU utility is set up to enable the `Fast SCSI` option.
10. The LFU utility is set up to set the SCSI ID for both channels to 6.
11. Entering `y` causes the options to be changed to the requested values.

4.5 Preparing the Storage Configuration

An ASE provides you with high data availability through service failover, disk mirroring, and fast file system recovery. TruCluster supports both the Logical Storage Manager (LSM) and RAID technology. You must determine the storage configuration that will meet your needs. Mirroring disks across two shared buses provides the most highly available data.

Chapter 2 describes the supported storage shelves, disk devices, and controllers for DEC RAID subsystems. See the TruCluster Software Products *Release Notes* to determine the latest revisions of hardware and firmware. To determine if your disk has the correct firmware, use console commands or the `scu` utility, or examine the messages that are displayed when you turn on the system.

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. For detailed information about installation and configuration, see your storage shelf (or DEC RAID subsystem) documentation.

The following sections describe how to prepare storage shelves for a shared SCSI bus.

4.5.1 Preparing a BA350 Storage Shelf

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

single-ended devices to form a single-ended shared SCSI bus or it can be used with a DWZZA and connected to a differential shared SCSI bus.

4.5.1.1 Connecting a BA350 Storage Shelf to a Single-Ended Shared SCSI Bus

The BA350 storage shelf is a single-ended, narrow device, and may be connected to one of the following single-ended, narrow devices:

- PMAZC TURBOChannel SCSI controller
- KZMSA XMI-to-SCSI adapter
- BA350 or BA353 storage shelf
- Single-ended end of a DWZZA

To connect a BA350 storage shelf directly to another single-ended, narrow device, see Figure 6–2 and follow these steps:

1. If the BA350 storage shelf is at the end of the single-ended bus, ensure that the BA350 SCSI terminator jumper (see Figure 2–1) is installed. If the BA350 storage shelf is in the middle of the single-ended bus, remove the BA350 internal termination.
2. Connect the BA350 to a single-ended, narrow device with a BN21R or BN23G cable. One cable will be connected to the BA350 (JA1) if it is on the end of the bus, two if it is in the middle of the bus.

Note

The adapter(s) not on the end of the shared SCSI bus must have their internal termination removed.

If the BA350 is connected directly to a single-ended device without the use of Y cables, the SCSI bus adapter termination is used to terminate the shared SCSI bus. Devices will not be able to be removed from the shared bus for maintenance purposes.

3. Install the disks in the BA350 storage shelf.

Figure 6–2 shows a sample Available Server configuration with two systems and a BA350 on a single-ended shared SCSI bus. Figure 6–3 shows an Available Server configuration with two BA350 storage shelves.

4.5.1.2 Connecting a BA350 Storage Shelf to a Differential Shared SCSI Bus

To prepare a BA350 storage shelf for connection to a differential shared SCSI bus, follow these steps:

1. Ensure that the BA350 storage shelf's internal termination and jumper is installed (see Figure 2–1).
2. Obtain the SCSI signal converter that you want to use, either a DWZZA-AA or a DWZZA-VA.
3. Ensure that the DWZZA single-ended termination jumper, J2, is installed. Remove the termination from the differential end by removing the five 14-pin differential terminator resistor SIPs.
4. Attach a trilink connector to the differential side of the DWZZA-AA or DWZZA-VA signal converter.
5. Attach the SCSI signal converter to the BA350 storage shelf as follows:

   If you are using a DWZZA-VA signal converter, install it in slot 0 (see Figure 4–4).

   If you are using a DWZZA-AA signal converter, use a BN21R or BN23G cable to connect the BA350 shelf's input connector, JA1, to the single-ended side of the DWZZA-AA. This connection forms a single-ended bus segment that must be terminated at both ends and must adhere to SCSI bus length restrictions. See Table 3–1 for the bus length restrictions.
6. Connect the trilink connector or Y cable to another trilink connector or Y cable with BN21K or BN21L cables, forming the differential shared SCSI bus segment.
7. If the BA350 is on the end of the shared SCSI bus, terminate the bus by installing an H879-AA on one side of the trilink connector or Y cable.
8. Install the disks in the BA350 storage shelf. The SCSI IDs for disks in a BA350 storage shelf correspond to their slot location.

Figure 4–4 shows a BA350 storage shelf with a DWZZA-VA signal converter installed in disk slot 0. A trilink connector is attached to the signal converter.

Figure 4–4: BA350 Storage Shelf, DWZZA-VA, and Trilink Connector

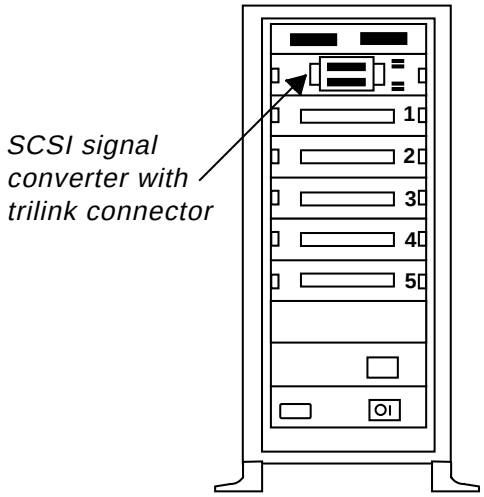

Figure 6–4 shows a configuration with a DEC 3000/500 with PMAZC TURBOchannel SCSI controller and a DEC 7000 in a differential Available Server configuration with a BA350 storage shelf.

Figure 6–6 shows an Available Server configuration with two systems and a BA350 with installed DWZZA-VA.

4.5.2 Preparing a BA353 Storage Shelf

A BA353 storage shelf provides access to SCSI devices through an 8-bit, single-ended, narrow SCSI-2 interface. You can use it to house disks for use on a single-ended shared SCSI bus. With a DWZZA, it may be used along with a differential shared SCSI bus. The BA353 has limited use as it has only three slots, one of which may contain a DWZZA-VA.

4.5.2.1 Connecting a BA353 Storage Shelf to a Single-Ended Shared SCSI Bus

The BA353 storage shelf is a single-ended, narrow device, so may be connected directly to one of the following single-ended, narrow devices:

- PMAZC TURBOChannel adapter
- KZMSA XMI-to-SCSI controller
- BA350 or BA353 storage shelf

To connect a BA353 storage shelf in a configuration with a single-ended, narrow shared SCSI bus, follow these steps:

1. If the BA353 storage shelf is at the end of the shared bus, leave the SCSI output connector open to terminate the bus (see Figure 2–2).
2. Use a BN21R or BN23G cable to connect the BA353 storage shelf input connector to a single-ended, narrow device.
3. Install disks in the BA353, ensuring that each disk has a unique SCSI ID on the shared bus. The SCSI IDs for disks in a BA353 storage shelf are set by the device address switches on the rear of the shelf (see Figure 2–2).

Note

The adapter(s) not on the end of the shared SCSI bus must have their internal termination removed.

If the BA353 is connected directly to a single-ended device without the use of Y cables (and external termination), the SCSI bus adapter termination is used to terminate the shared SCSI bus. You will not be able to be disconnect devices from the shared bus for maintenance purposes.

The configuration shown in Figure 6–2 is similar to a single-ended shared SCSI bus using a BA353, just picture a BA353 instead of a BA350.

4.5.2.2 Connecting a BA353 Storage Shelf to a Differential Shared SCSI Bus

You must use a DWZZA-AA or DWZZA-VA to connect a BA353 storage shelf to a differential SCSI bus segment.

To prepare a BA353 storage shelf for connection to a differential shared SCSI bus, follow these steps:

1. Obtain a DWZZA-AA or a DWZZA-VA SCSI signal converter.

   If you will use a DWZZA-VA:

   a. Remove the DWZZA-VA differential termination by removing the five 14-pin differential terminator resistor SIPs. Remove the termination from the single-ended side by removing jumper J2.

   b. Install the DWZZA-VA signal converter in any disk slot in the BA353 storage shelf.

   c. Terminate the BA353 SCSI input connector using a 50-pin, high-density, single-ended terminator (12-37004-04).

If you are using a DWZZA-AA signal converter:

a. Ensure that jumper J2 is installed in the DWZZA-AA signal converter to enable the single-ended termination. Remove the differential termination by removing the five 14-pin differential terminator resistor SIPs.

b. Connect the DWZZA-AA single-ended side to the BA353 input connector with a BN21R or BN23G cable. This connection forms a single-ended bus segment that is terminated at one end by the DWZZA single-ended termination and on the other end by the BA353 output termination. The single-ended bus must adhere to SCSI bus length restrictions. See Table 3–1 for the length restrictions.

2. Ensure that the BA353 storage shelf's SCSI output connector is open. If two BA353's are daisy chained, ensure that the output connector on the last storage shelf is open.
3. Install disks in the BA353, ensuring that each disk has a unique SCSI ID on the shared bus. The SCSI IDs for disks in a BA353 storage shelf are set by the device address switches on the rear of the shelf (see Figure 2–2).
4. Attach a trilink connector or Y cable to the differential side of the DWZZA-AA or DWZZA-VA signal converter.
5. Connect the differential devices (using trilink connectors or Y cables) together with BN21K or BN21L cables to form the differential bus.
6. Attach H879-AA terminators to the trilink connectors or Y cables at the end of the differential bus.

4.5.3 Preparing a Non-Ultra BA356 Storage Shelf

A non-Ultra BA356 storage shelf provides access to SCSI devices through a 16-bit, single-ended, wide SCSI-2 interface. In a cluster configuration, you would not connect a non-Ultra BA356 directly to any single-ended device other than another non-Ultra BA356 or the single-ended side of a DWZZB.

To prepare a non-Ultra BA356 storage shelf for a differential shared SCSI bus configuration, follow these steps:

1. Ensure that the BA356 storage shelf's internal termination is enabled. (It is enabled as long as there is not a cable attached to the JB1 connector on the personality module.)
2. Install the disks in the shelf. Ensure that the switches on the personality module are all set off to allow the slot position to

determine the SCSI bus ID. If you are daisy chaining two BA356 storage shelves, switches 1, 2, and 3 must be set to on in the second BA356. See the BA356 *User's Guide* for information about disk device addresses in BA356 configurations. Note that only eight devices are supported on a shared SCSI bus.

3. Obtain the SCSI signal converter that you want to use, either a DWZZB-AA or a DWZZB-VW.

   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.

4. If you are using a DWZZB-VW signal converter, install it in slot 0.

   If you are using a DWZZB-AA signal converter, connect the single-ended side of the DWZZB-AA to the BA356 input connector, JA1, on the personality module with a BN21K or BN21L cable. 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 single-ended bus must adhere to SCSI bus length restrictions, including the length of the SCSI bus internal to the BA356. See Table 3–1 for the length restrictions.

5. Attach an H885 trilink connector or BN21W-0B Y cable to the differential side of the DWZZB-AA or DWZZB-VW signal converter.

6. Connect all trilink connectors or Y cables that will be used to form the differential shared SCSI bus to each other with BN21K or BN21L cables.

7. Attach an H879-AA to the two trilink connectors or Y cables on the end of the bus.

The configuration shown in Figure 6–6 is very similar to a configuration using a BA356 storage enclosure. Just visualize replacing the BA350 with a BA356, and replacing the DWZZA-VA with a DWZZB-VW.

4.5.4 Preparing an Ultra BA356 Storage Shelf

An Ultra BA356 storage shelf is connected to a shared wide UltraSCSI bus, and provides access to UltraSCSI devices on the internal, single-ended wide UltraSCSI bus. The interface between the buses is the DS-BA35X-DA personality module, which is installed in the Ultra BA356.

To prepare an Ultra BA356 storage shelf for a differential shared SCSI bus configuration, follow these steps:

1. Disable the Ultra BA356 differential termination. Ensure that personality module (DS-BA35X-DA) switch pack 4 switches S4-1 and S4-2 are off (see Figure 3–3).

Note

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

2. Enable the device position in the Ultra BA356 to determine the SCSI ID. Ensure that all personality module switch pack 3 switches are off (see Figure 3–3).
3. Ensure that the BA35X-MJ jumper module is installed behind slot 6 (see Section 2.7.3.2 and Figure 2–3).
4. Install an H8861-AA VHDCI trilink on personality module connector JA1.
5. Cable and terminate the UltraSCSI bus as follows:
   - The Ultra BA356 is at the end of the shared SCSI bus:
     – Install an H8863-AA terminator on one side of the H8861-AA trilink.
     – Install a BN38C cable (an HD68 to a VHDCI) between the closest BN21W-0B Y cable and the H8861-AA trilink.
     – Connect two BN21W-0B Y cables from different systems with a BN21K or BN21L cable.
     – Terminate the other end of the shared SCSI bus by installing an H879-AA terminator on the BN21W-0B with an open connector.
   - The Ultra BA356 is in the middle of the shared SCSI bus:
     – Install BN38C cables (an HD68 to a VHDCI) between the H8861-AA trilink and adjacent BN21W-0B Y cables.
     – Connect the adjacent BN21W-0B Y cables with a BN21K or BN21L cable.
     – Terminate both ends of the shared SCSI bus by installing an H879-AA terminator on the open connector on the BN21W-0B Y cables at both ends of the bus.

4.5.5 Placing a DEC RAID Controller on a Shared SCSI Bus

A DEC RAID array controller provides high performance, high availability, and high connectivity access to SCSI devices through the shared SCSI buses.

Before you connect a DEC 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. For information about setting up a DEC 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*
- *Working in DIGITAL UNIX Systems DIGITAL StorageWorks HSZ50 Array Controller HSOF 5.01 Configuration Manual*
- *HSZ70 Array Controller HSOF Version 7.0 Configuration Manual*

The HSZ10 controller uses only one SCSI ID and can be used in an ASE that uses only PMAZC TURBOChannel SCSI controllers. You can configure the HSZ20, HSZ40, and HSZ50 DEC RAID array controllers with one to four SCSI IDs. You can configure the HSZ70 DEC RAID array controller with one to seven SCSI IDs.

Because the HSZ10, HSZ20, HSZ40, and HSZ50 have a wide differential connection on the host side, you connect them to one of the following differential devices:

- KZPSA or KZTSA adapter
- Another HSZ10, HSZ20, HSZ40, or HSZ50
- Differential side of a DWZZA or DWZZB SCSI signal converter (which can be connected to a KZMSA or PMAZC single-ended SCSI controller)

The HSZ70 is an UltraSCSI device, so it must be connected using Very High Density Cable Interconnect (VHDCI) components. You can use an HSZ70 on a shared SCSI bus with a KZPSA host adapter.

To connect an HSZ10, HSZ40, or HSZ50 controller to a shared SCSI bus, follow these steps:

1. Attach an H885 trilink connector to the DEC RAID controller port.

2. Use a BN21K or BN21L cable to connect the trilink connector to a BN21W-0B Y cable attached to a differential SCSI controller or the differential end of a signal converter.
3. Terminate the differential bus by attaching an H879-AA terminator to the H885 trilink connector or BN21W-0B Y cables at the end of the 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.

Figure 6–7 shows two DEC 7000 systems with KZMSA XMI-to-SCSI adapters on a shared bus with an HSZ40.

To connect an HSZ70 DEC RAID controller to a shared UltraSCSI bus, follow these steps:

1. Attach an H8861-AA VHDCI trilink connector to each HSZ70 controller port.
2. Conect a BN37A cable between the H8861-AA trilinks on the HSZ70 controller ports.
3. Cable and terminate the Ultra SCSI bus as follows:
   - The HSZ70 is on the end of the shared SCSI bus:
     - Attach an H8863-AA terminator to the open connector of the H8861-AA trilink at the end of the shared SCSI bus.
     - Connect the closest BN21W-0B Y cable to the open connector on the H8861-AA trilink with a BN38C (HD68 to VHDCI) cable. If there is only one HSZ70 controller, the H8861-AA trilink also has an H8863-AA terminator.
     - Connect all pairs of BN21W-0B Y cables from different systems with BN21K or BN21L cables.
     - Terminate the other end of the shared SCSI bus by attaching an H879-AA terminator to the BN21W-0B Y cable with an open connector.
   - If the HSZ70 is in the middle of the shared SCSI bus:
     - Install a BN38C (HD68 to VHDCI) cable between the H8861-AA trilink open connectors and a BN21W-0B Y cable.
     - If there are more than two systems, connect the BN21W-0B Y cables to each other with BN21K or BN21L cables.
     - Attach H879-AA terminators to the open connectors of the BN21W-0B Y cables on each end of the shared SCSI bus.

Figure 6–8 shows two AlphaServer 8200 systems in an Available Server configuration with two HSZ70 DEC RAID controllers. Note that the SCSI bus adapters are KZPSA PCI-to-SCSI adapters. In this figure, the HSZ70s are on the end of the shared SCSI bus.

Figure 6–9 shows a figure similar to Figure 6–8 with two AlphaServer 8200 systems in an Available Server configuration with two HSZ70 DEC RAID controllers. As with Figure 6–8, the SCSI bus adapters are KZPSA PCI-to-SCSI adapters. In this figure, the HSZ70s are in the middle of the shared SCSI bus so the cables required are different.

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

1. Referring to the *RAID Array 310 Deskside Subsystem (SWXRA-ZX) Hardware User's Guide*, open the SWXRA-Z1 cabinet, locate the SCSI bus converter board, and:
   - Remove the five differential terminator resistor SIPs.
   - Ensure that the W1 and W2 jumpers are installed to enable the single-ended termination on one end of the bus.

   Note

   The RAID Array 310 SCSI bus converter board is the same logic board used in the DWZZB signal converter.

2. Attach an H885 trilink connector to the SCSI input connector (on the back of the cabinet).
3. Use a BN21K or BN21L cable to connect the trilink connector to a trilink 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 trilink 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.

4.6 Preparing the Tape Devices

The topics in this section provide information on preparing the various tape devices for use on a shared SCSI bus with the TruCluster Production Server Software and TruCluster Available Server Software products.

Tape devices are only supported on shared SCSI buses with KZPSA PCI SCSI adapters.

Note

To achieve system performance capabilities, DIGITAL recommends placing no more than one TZ89 drive on a SCSI bus.

4.6.1 Preparing the TZ88 for Shared Bus Usage

Two versions of the TZ88 are supported, the TZ88N-TA table-top 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.

4.6.1.1 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 4–9. Figure 4–5 shows the TZ88N-VA with the backplane interface connector and SCSI ID switch pack.

Figure 4–5: TZ88N-VA SCSI ID Switches

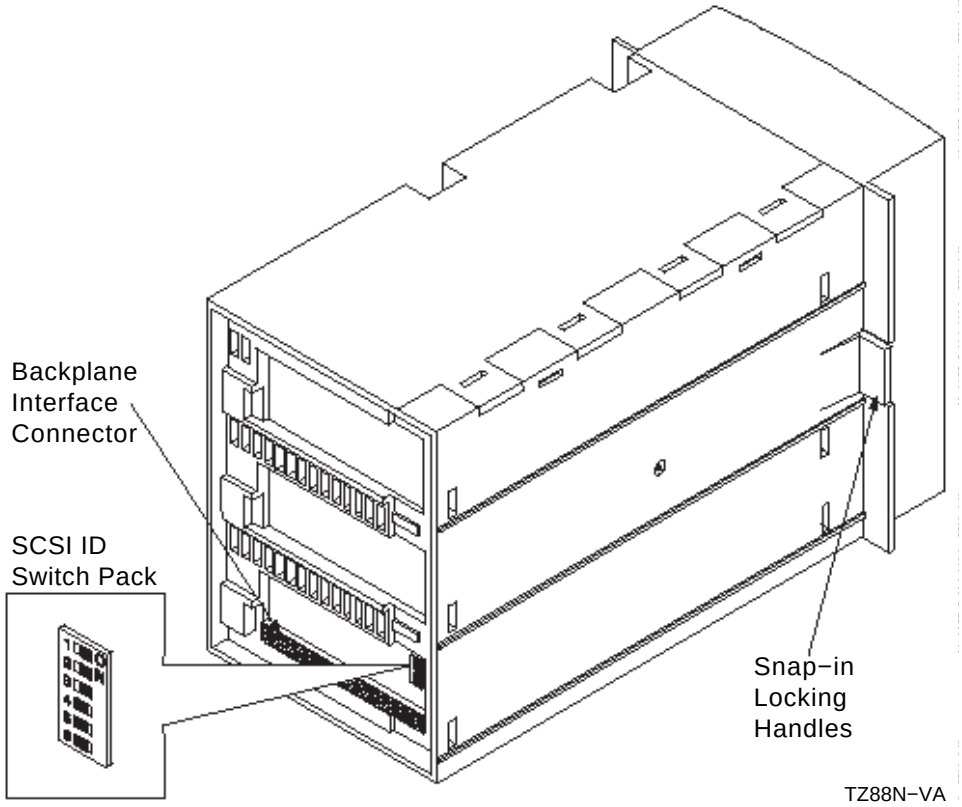

Table 4–9: TZ88N-VA Switch Settings

SCSI ID	SCSI ID Selection Switches					
	1	2	3	4	5	6
Automatic[a]	Off	Off	Off	On	On	On
0	Off	Off	Off	Off	Off	Off
1	On	Off	Off	Off	Off	Off
2	Off	On	Off	Off	Off	Off
3	On	On	Off	Off	Off	Off
4	Off	Off	On	Off	Off	Off
5	On	Off	On	Off	Off	Off

Table 4–9: TZ88N-VA Switch Settings (cont.)

SCSI ID	SCSI ID Selection Switches					
6	Off	On	On	Off	Off	Off
7	On	On	On	Off	Off	Off

[a]SBB tape drive SCSI ID is determined by the SBB physical slot.

4.6.1.2 Cabling the TZ88N-VA

There are no special cabling requirements 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 trilink connector on the DWZZA-VA differential end allows connection to the shared bus. An H879-AA terminator is installed on the trilink for the BA350 on the end of the bus to provide shared SCSI bus termination.

Figure 4–6 shows a Production Server configuration 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 4–6: Shared SCSI Buses with SBB Tape Drives

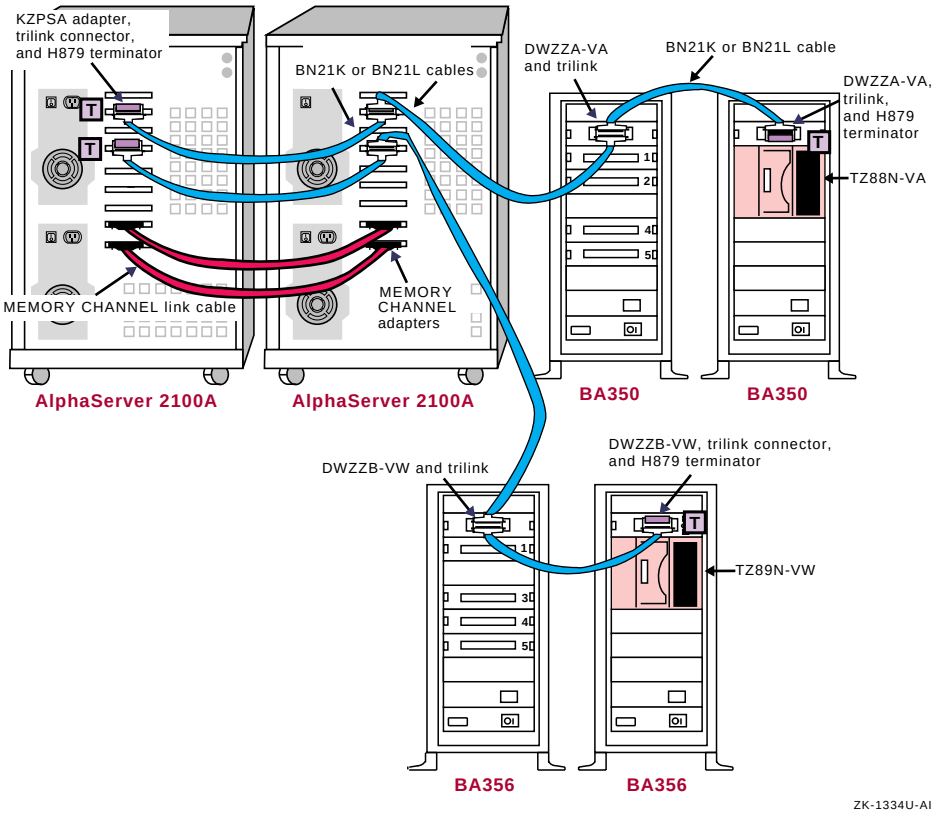

4.6.1.3 Setting the TZ88N-TA SCSI ID

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.

4.6.1.4 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. Figure 4–8 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 for each TZ88N-TA.

Ensure that the DWZZA jumper J2 is 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 or Y cable to the differential end of the DWZZA-AA.
3. Connect the single-ended end of the DWZZA-AA to the TZ88N-TA with a BC19J cable.
4. Install a H8574-A or H8890-AA terminator on the other TZ88N-TA 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 TZ88N-TA tape drives to the differential shared SCSI bus by adding additional DWZZA-AA/TZ88N-TA combinations.

Note

Ensure that there is no conflict with tape drive, system, and disk SCSI IDs, and that you keep the number of SCSI devices to a maximum of eight.

4.6.2 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.

4.6.2.1 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 4–7 shows a view of the DS-TZ89N-VW showing the backplane interface connector and SCSI ID switch pack.

Figure 4–7: DS-TZ89N-VW SCSI ID Switches

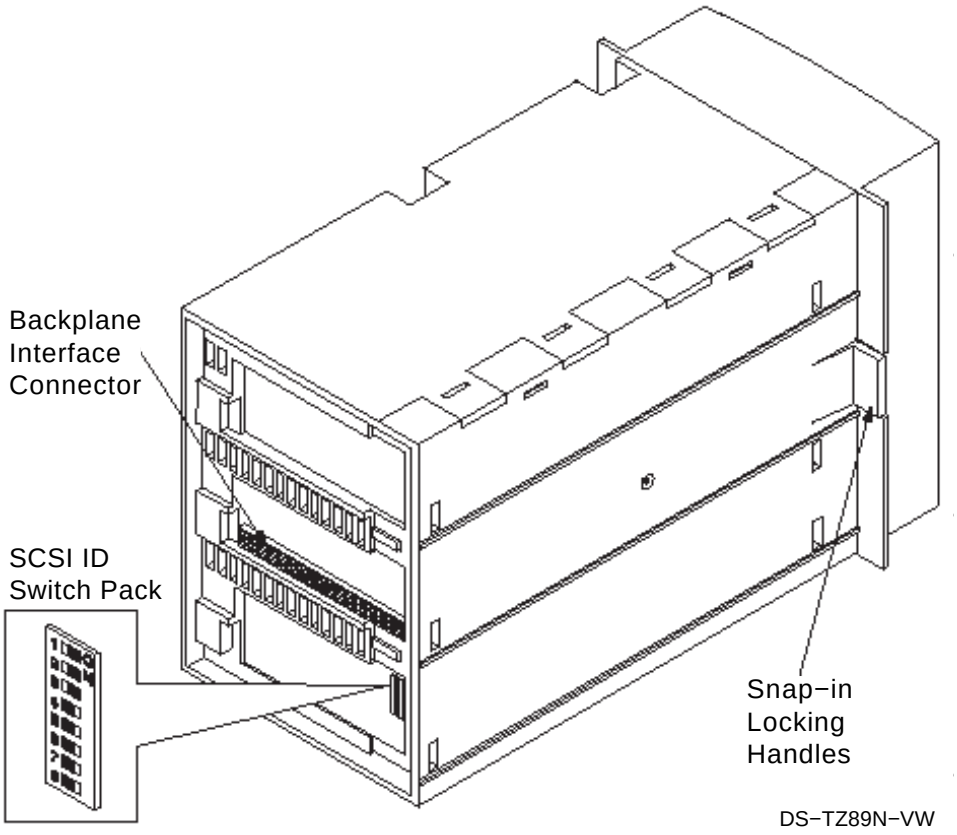

The SCSI ID is selected by switch positions, which must be selected before the tape drive is installed in the BA356. Table 4–10 shows the switch settings for the DS-TZ89N-VW.

Table 4–10: 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 4–10: 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.

4.6.2.2 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 4–6 shows a Production Server configuration 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.

4.6.2.3 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 (maximum is 15); push the button below the switch to decrease the SCSI ID.

4.6.2.4 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 4–8 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. Configure the shared SCSI bus for use with a DS-TZ89N-TA as follows:

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 or Y cable 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 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 wide, single-ended SCSI bus may be daisy chained from one single-ended tape drive to another with BN21N 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, and that you keep the number of SCSI devices on the shared SCSI bus to a maximum of eight.

4.6.3 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.

4.6.3.1 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.

4.6.3.2 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. Figure 4–8 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 zero. To configure the shared SCSI bus for use with a TZ885, follow these steps:

1. You will need one DWZZA-AA for each TZ885 tape drive.

   Ensure that the DWZZA jumper J2 is 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 or Y cable to the differential end of the DWZZA-AA.
3. Connect the single-ended end of the DWZZA-AA to the TZ885 with a BC19J 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-AA/TZ885 combinations.

Note

Ensure that there is no conflict with tape drive, system, and disk SCSI IDs, and that you keep the number of SCSI devices to a maximum of eight.

Figure 4–8: Cabling a Shared SCSI Bus With a TZ885

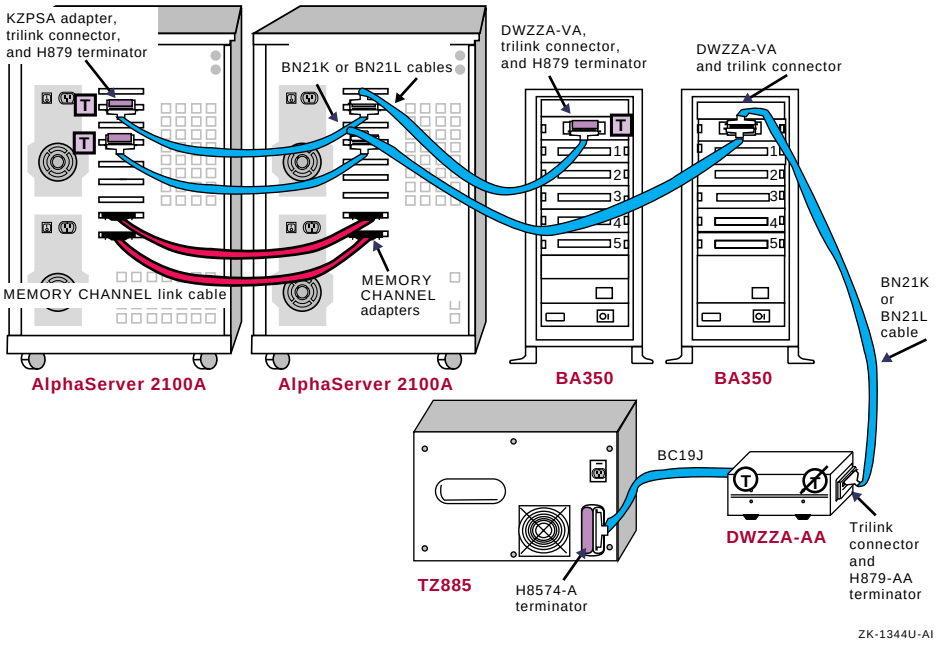

4.7 Preparing ATM Adapters

The DIGITAL UNIX base operating system supports Asynchronous Transfer Mode (ATM). TruCluster Available Server Software supports the use of LAN emulation (LANE) over ATM for client access and as the network for cluster communications. TruCluster Production Server Software supports the use of LAN emulation over ATM for client access.

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

4.7.1 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.

Asynchronous Transfer Mode (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 boundries. Communication between ELANs is possible only through routers or bridges.

Each emulated LAN 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 4–9 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 Broadcast and Unknown Server are server functions resident on the ATM switch (even though they are shown separately).

Figure 4–9: Emulated LAN Over an ATM Network

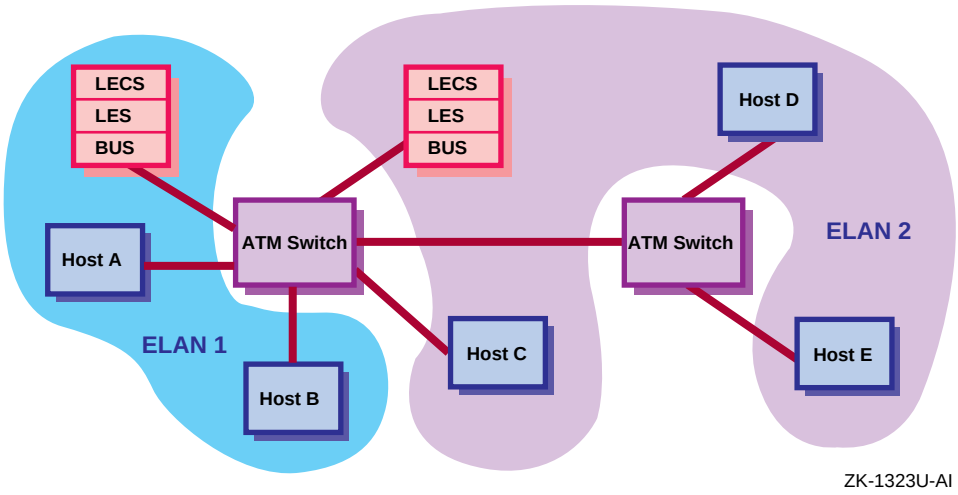

TruCluster Available Server Software and TruCluster Production Server Software use LAN emulation over ATM as follows:

- TruCluster Available Server Software: LAN emulation over ATM can be used as the primary or backup interconnect for communication between ASE members. LAN emulation over ATM may also be used for client system access.
- TruCluster Production Server Software: LAN Emulation over ATM may be used for client system access (MEMORY CHANNEL is the cluster interconnect).

4.7.2 Example TruCluster Available Server Software Configuration Using ATM

A typical TruCluster Available Server Software configuration using LAN emulation over ATM for client access is shown in Figure 4–10. In this configuration, each ASE server and all the clients on GIGAswitch/ATM A are on one emulated LAN. Each ASE server and all clients on GIGAswitch/ATM B are on another emulated LAN. Clients on one ELAN can communicate with clients on the other ELAN through the router.

Figure 4–10: Example TruCluster Available Server Software Configuration using LAN Emulation Over ATM

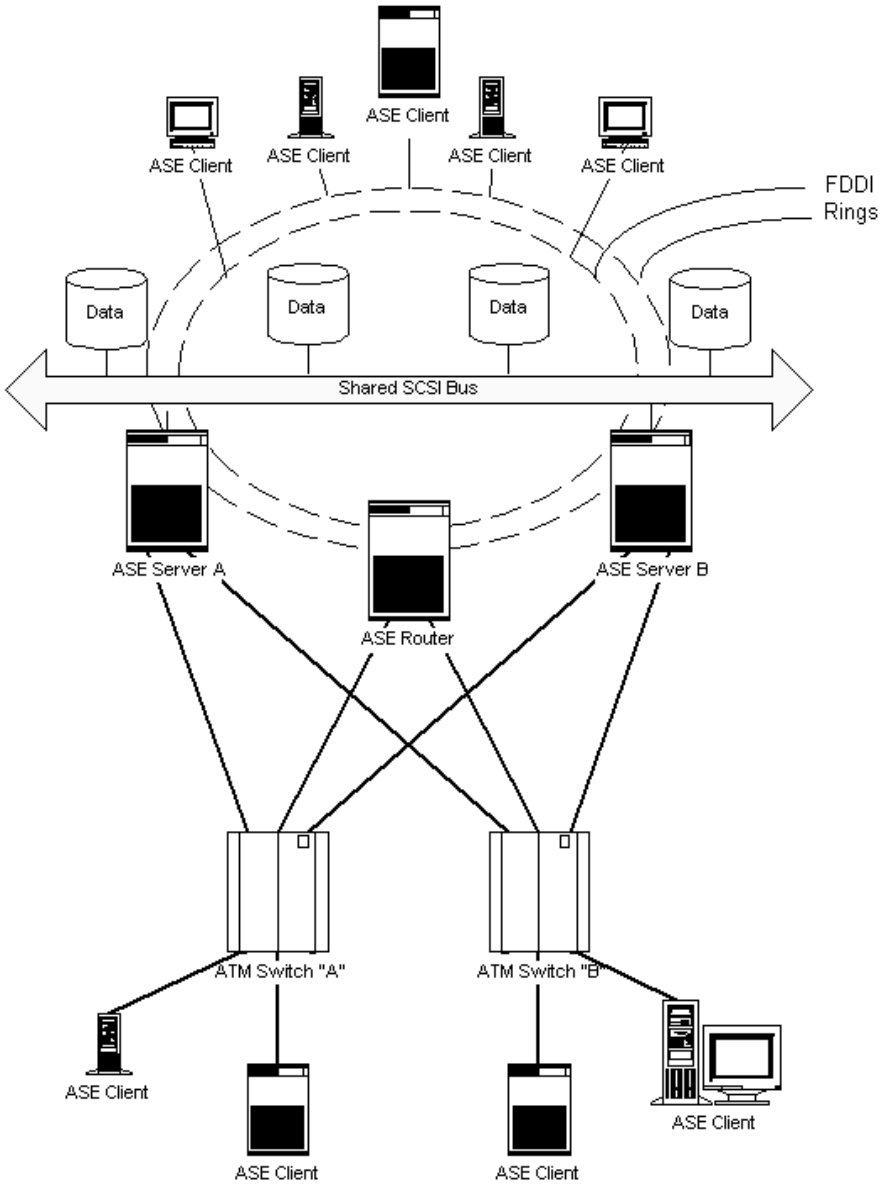

4.7.3 Installing ATM Adapters

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 arror 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 4.7.4.

4.7.4 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
DIGITAL 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.

4.7.5 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 4–11.

Table 4–11: 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.

Table 4–11: ATMworks Adapter LEDs (cont.)

Network LED	Module LED	Description
Green	Green/Off	A physical link connection has been made.
Amber	Green	There is no physical link connection.

5

Setting Up the MEMORY CHANNEL Cluster Interconnect

This chapter describes MEMORY CHANNEL configuration requirements, and describes how to set up the MEMORY CHANNEL cluster interconnect, including setting up a MEMORY CHANNEL hub and connecting link cables.

If you are setting up a TruCluster Available Server Software configuration, you can ignore this chapter as MEMORY CHANNEL hardware is not used.

Section 2.4 lists the supported MEMORY CHANNEL hardware. See the MEMORY CHANNEL *User's Guide* for illustrations and detailed information about installing jumpers, MEMORY CHANNEL adapters, and hubs. See Chapter 7 for information about changing and maintaining the MEMORY CHANNEL hardware.

TruCluster Production Server Software and TruCluster MEMORY CHANNEL Software currently support only MEMORY CHANNEL for cluster interconnects. You can have two MEMORY CHANNEL adapters for TruCluster Production Server Software to protect against an interconnect failure and for easier hardware maintenance. If the active MEMORY CHANNEL adapter fails, cluster communications fails over to the inactive MEMORY CHANNEL adapter.

You can have up to four MEMORY CHANNEL adapters for TruCluster MEMORY CHANNEL Software but there is no failover between adapters. For each interconnect, you must install one MEMORY CHANNEL adapter in a peripheral component interconnect (PCI) slot in each member system.

See Section 5.1 for a discussion on MEMORY CHANNEL requirements.

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 5.2).
2. Install the MEMORY CHANNEL adapter into a PCI slot on each system (Section 5.3).
3. If you have more than two systems in the cluster, install a MEMORY CHANNEL hub (Section 5.4).

4. Connect the MEMORY CHANNEL cables (Section 5.5).
5. After you complete steps 1 through 4 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 MEMORY CHANNEL Requirements

Ensure that you abide by the following MEMORY CHANNEL requirements:

- MEMORY CHANNEL adapters require a specific, minimum firmware revision. See the TruCluster Software Products *Release Notes* for information about required firmware revisions.
- 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 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 8200/8400 systems, the MEMORY CHANNEL adapter must be installed in slots 0-7 of a DWLPA-CA PCI; there are no restrictions for a DWLPB-CA PCI bus.

- 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 MEMORY CHANNEL hub).
- If you use multiple MEMORY CHANNEL interconnects, each interconnect must use the same mode, either virtual hub mode or standard mode.
- If you are using virtual hub mode, you must connect the first MEMORY CHANNEL adapter installed in one system to the first adapter installed in the other system, and you must connect the second adapter installed in one system to the second adapter installed in the other system.

  If a Production Server or MEMORY CHANNEL 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 in each system must be connected to one MEMORY CHANNEL hub. The second MEMORY CHANNEL adapter in each system must be connected to a second MEMORY CHANNEL hub.
- The maximum length of a MEMORY CHANNEL link cable is 3 meters (10 feet).

  Always check a 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.

5.2 Setting the MEMORY CHANNEL Adapter Jumpers

The MEMORY CHANNEL module has 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 (jumpering middle and left pins). Confirm that the jumpers are set properly. The jumpers are right next to the factory/maintenance cable connector and are shown in Table 5–1.

Table 5–1: MEMORY CHANNEL Jumper Configuration

If mode is:	Jumper:	Example:
Standard	Middle and left pins	
Virtual: VH0	Middle and right pins	
Virtual: VH1	None needed; store the jumper on the right post	

If you are upgrading from virtual hub mode to standard mode (or from standard hub mode to virtual hub mode), be sure to change the jumpers on all MEMORY CHANNEL adapters.

5.3 Installing the MEMORY CHANNEL Adapter

Install the MEMORY CHANNEL adapter in an appropriate peripheral component interconnect (PCI) slot. Secure the module at the backplane. Ensure that the screw is tight to maintain proper grounding, but do not over torque the screw. An over-torqued screw can pull the module partially out of the backplane connector.

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, or for an AlphaServer 8200/8400, 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.

5.4 Installing the MEMORY CHANNEL Hub

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

- You must install the hub within 3 meters (10 feet) of each of the systems.
- If the hub has a voltage selection switch on the back of the hub, select the correct voltage for your location.
- Ensure that the hub contains a CCMLA line card for each system in the cluster (the hub comes with four line cards). If you have a four-node cluster, you may want to install an extra line card for troubleshooting use.

5.5 Installing the MEMORY CHANNEL Cables

To set up a MEMORY CHANNEL interconnect, use the BC12N-10 MEMORY CHANNEL link cables to connect MEMORY CHANNEL adapters and, optionally, MEMORY CHANNEL hubs. The maximum length of the link cable is 3 meters (10 feet).

5.5.1 Connecting MEMORY CHANNEL Link Cables in Virtual Hub Mode

For a MEMORY CHANNEL virtual hub configuration (two nodes in the cluster), connect the BC12N-10 link cables between the MEMORY CHANNEL adapter installed in each of the systems.

Caution

Be very careful when installing the link cables. Insert the cables straight in. When removing the cables, rock them vertically, not horizontally.

Gently push the cables' connector into the receptacle, 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:

- Both adapters in a system must have the same jumper setting, either VH0 or VH1.
- You must connect the first MEMORY CHANNEL adapter installed in one system to the first adapter installed in the other system, and connect the second adapter in one system to the second adapter in the other system.

5.5.2 Connecting MEMORY CHANNEL Link Cables in Standard Hub Mode

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

If you are setting up redundant interconnects, the following rules 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 line cards that are installed in the same slot position in each hub. For example, if you connect one adapter to a line card installed in slot one in one hub, you must connect the other adapter in that system to a line card installed in slot one of the second hub.

Figure 5–1 shows MEMORY CHANNEL adapters connected to line cards that are in the same slot position in the MEMORY CHANNEL hubs.

Figure 5–1: Connecting MEMORY CHANNEL Adapters to Hubs

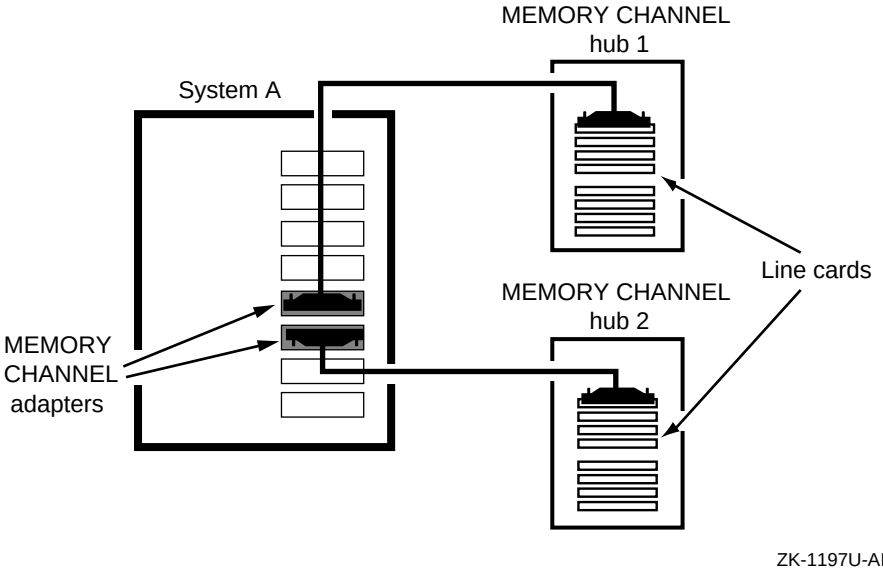

5.6 Running MEMORY CHANNEL Diagnostics

After the MEMORY CHANNEL adapters, hubs, and link 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.
  - It is run as part of the initialization sequence when the system is powered up.
  - Can be run 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).
  - Is designed to isolate problems to the MEMORY CHANNEL adapter, BC12N-10 link cables, hub line cards, 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.

Caution

Do not run the `mc_cable` MEMORY CHANNEL diagnostic on a system that is cabled to a live cluster. The diagnostic transmissions on the link are seen by other systems in the cluster and may cause problems.

Example 5–1 shows a sample output from one node 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 one reports that `mca0` is on line but has not communicated with the MEMORY CHANNEL adapter on node 0.
3. Node one 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 one report off line.
7. A Ctrl/C on node one terminates the test.

6

Sample Hardware Configurations

This chapter describes some common Available Server and Production Server hardware configurations. For the Available Server configurations, figures containing generic devices are shown first, and then figures with actual devices are shown. For the Production Server configurations, most of the figures contain generic devices. See Chapter 4 for information about connecting devices to a shared SCSI bus and preparing systems and storage shelves.

6.1 Sample Available Server Hardware Configurations

This section describes some common Available Server hardware configurations.

6.1.1 Available Server Configurations that Use Two Systems and One Shared Bus

Figure 6–1 shows a minimum Available Server hardware configuration consisting of two systems, one storage shelf, and one shared SCSI bus.

Figure 6–1: Two Systems and One Shared Bus

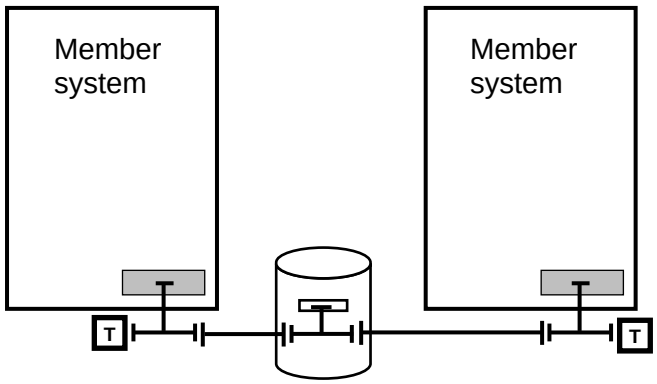

ZK-1116U-AI

The following sections provide detailed examples of two-system, one-bus configurations.

6.1.1.1 Single-Ended Device Configurations

In the configuration shown in Figure 6–2, the SCSI controllers and the storage shelf are single-ended, narrow devices. There is one single-ended bus between the two terminated Y cables.

Figure 6–2: Two DEC 3000/500 Systems With PMAZCs and a BA350 Storage Shelf

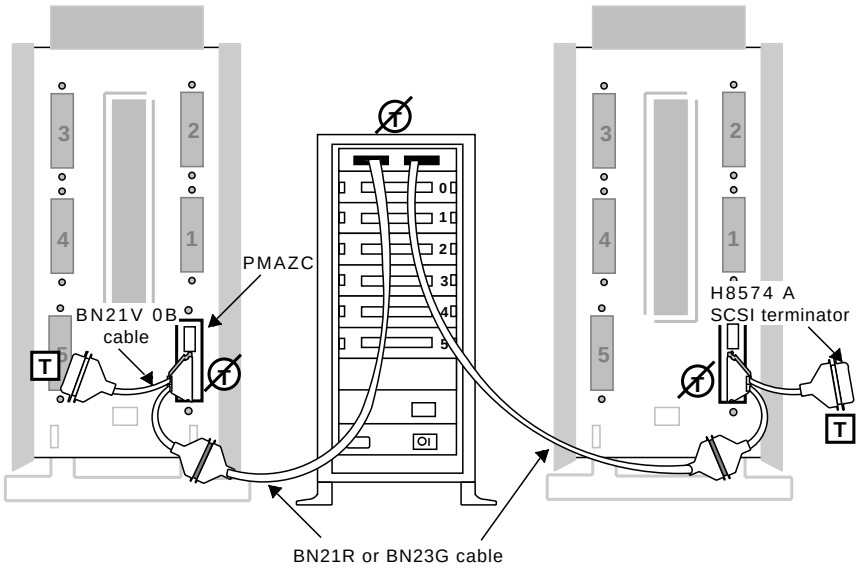

To create the configuration in Figure 6–2, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two DEC 3000 Model 500s	Set the PMAZC SCSI IDs to 6 and 7.
SCSI Controllers	Two PMAZC	Remove port A's termination; keep port B terminated.

Hardware	Description	Configuration Information
Storage	BA350 with six disks	Remove the termination and install disks in slots 0 to 5.
Cables	Two BN21V-0Bs (Y)	Attach to the PMAZCs.
	Two BN21Rs or BN23Gs	Connect the Y cables and BA350.
Terminators	Two H8574-As	Attach to the Y cables.

If you have the same devices as in Figure 6–2, you can create an alternative configuration by putting the BA350 storage shelf at the end of the shared bus. If you use this configuration, keep the internal termination in the BA350 storage shelf to terminate the bus, and use a BC19J cable to connect the two systems.

If you want two BA350 storage shelves in your configuration, use a BN21H or BN21J cable to connect the shelves, as shown in Figure 6–3.

Figure 6–3: Two DEC 3000/500 Systems With PMAZCs and Two BA350 Shelves

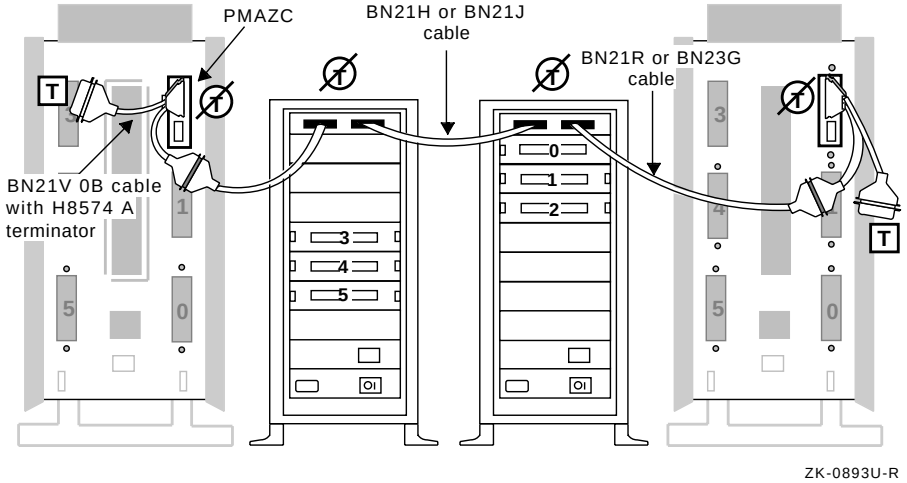

Figure 6–4 shows a configuration consisting of two types of single-ended SCSI controllers and one single-ended storage shelf. This configuration uses a DEC 7000 system, which uses KZMSA XMI-to-SCSI bus adapters. You must use a DWZZA-AA with the KZMSA because the KZMSA does not have removable termination. Therefore, to provide a differential bus for all devices, you must connect SCSI signal converters to each of the single-ended devices.

Figure 6–4: A DEC 7000 System, a DEC 3000/500 System With a PMAZC, and a BA350 Storage Shelf

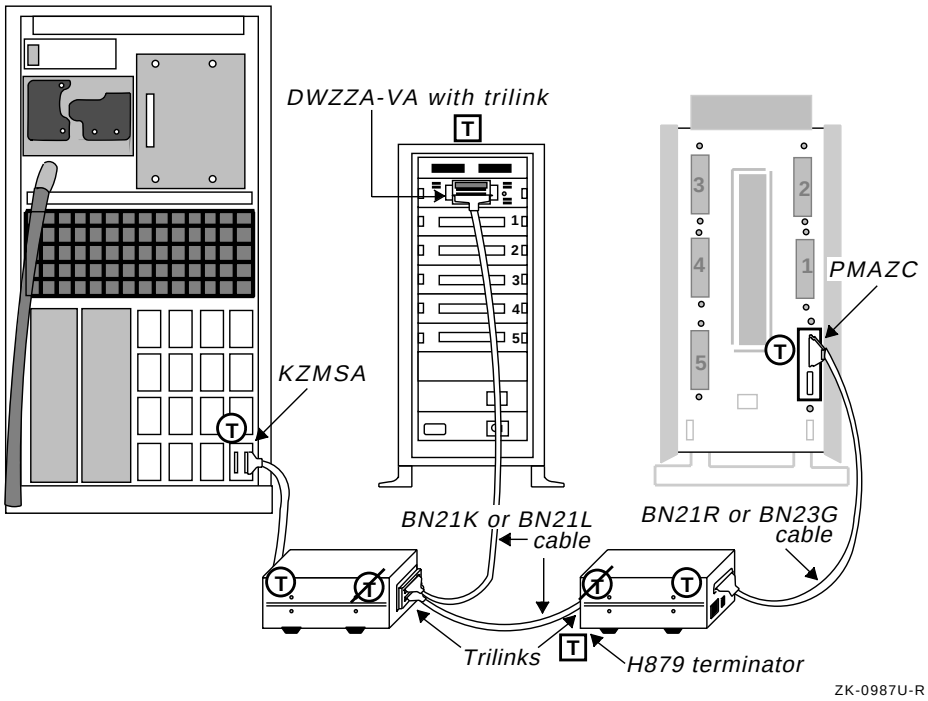

To create the configuration in Figure 6–4, you need the following hardware:

Hardware	Description	Configuration Information
Systems	One DEC 3000 Model 500	Set the PMAZC SCSI ID to 7.
	One DEC 7000	Set the KZMSA SCSI ID to 6.
SCSI Controllers	One PMAZC and one KZMSA	Terminate both PMAZC ports.
Storage	BA350 with five disks	Terminate the storage shelf and install disks in slots 1 to 5.
Signal Converters	Two DWZZA-AAs and one DWZZA-VA	Keep the termination in the single-ended side, and remove the termination from the differential side.
Connectors	Three H885 trilinks	Attach to the differential side of the DWZZAs.

Hardware	Description	Configuration Information
Cables	Two BN21Rs or BN23Gs	Connect the PMAZC and KZMSA to the DWZZAs.
	Two BN21Ks or BN21Ls	Connect the trilinks.
Terminators	Two H879-AAs	Attach to the trilinks at the ends of the bus.

6.1.1.2 Differential Device Configuration

Figure 6–5 shows a configuration consisting of two systems and one HSZ40 DEC RAID controller, all differential devices. There is one shared differential SCSI bus between the terminated H885 trilink connector and terminated BN21W-0B Y cable.

Figure 6–5: Two AlphaServer 8200 Systems and an HSZ40 Unit

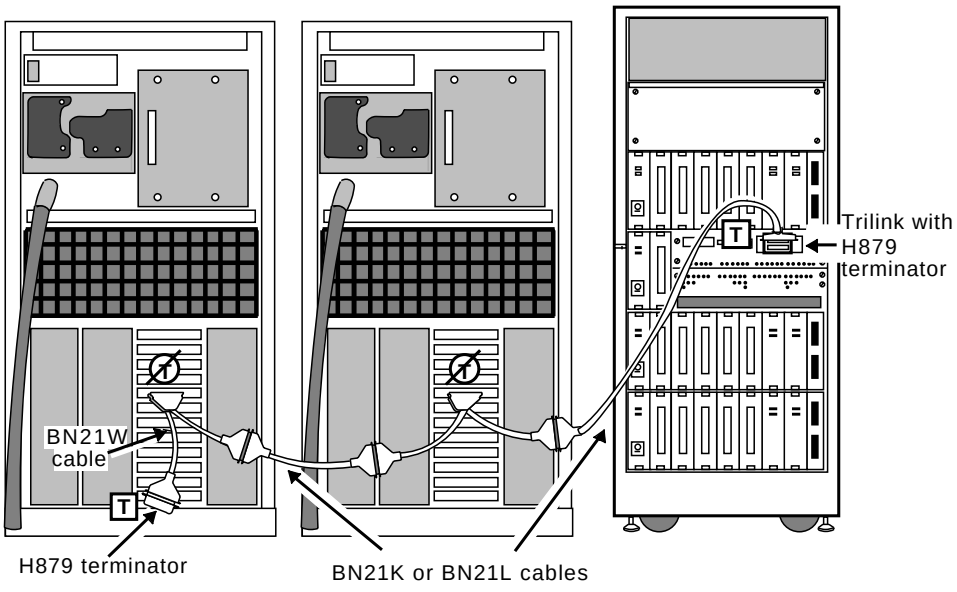

To create the configuration in Figure 6–5, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two AlphaServer 8200s	Set the KZPSAs SCSI IDs to 6 and 7.
SCSI Controllers	Two KZPSAs	Remove the termination.
Storage	HSZ40	N/A

Hardware	Description	Configuration Information
Connectors	One H885 trilink	Attach to the HSZ40.
Cables	Two BN21W-0Bs (Y)	Attach to the KZPSAs.
	Two BN21Ks or BN21Ls	Connect the Y cables and trilink.
Terminators	Two H879-AAs	Attach to the Y cable and trilink at the ends of the bus.

6.1.1.3 Single-Ended Device and Differential Device Configurations

Figure 6–6 shows a configuration with two systems with KZPSA differential SCSI adapters and one single-ended storage shelf with an installed DWZZA-VA SCSI signal converter.

Figure 6–6: Two AlphaServer 2100 Systems and a BA350 Storage Shelf

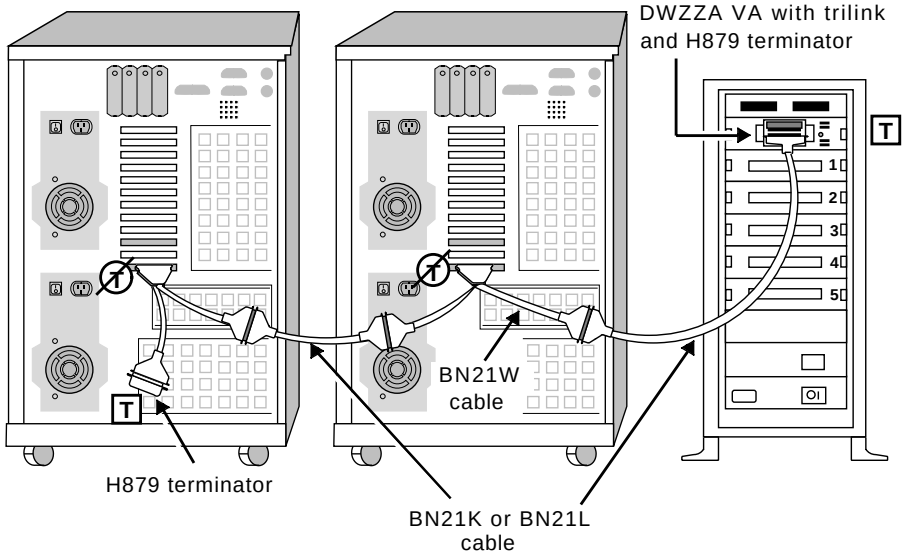

To create the configuration in Figure 6–6, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two AlphaServer 2100s	Set the KZPSAs SCSI IDs to 6 and 7.
SCSI Controllers	Two KZPSAs	Remove the termination.
Storage	BA350 with five disks	Terminate the storage shelf and install the disks in slots 1 to 5.
Signal Converters	One DWZZA-VA	Terminate the single-ended side, and remove the termination from the differential side.
Connectors	One H885 trilink	Attach to the DWZZA-VA.
Cables	Two BN21W-0Bs (Y)	Attach to the KZPSAs.
	Two BN21Ks or BN21Ls	Connect the Y cables and trilink.
Terminators	Two H879-AAs	Attach to the trilink and Y cable at the ends of the bus.

Figure 6–7 shows a configuration consisting of two systems with single-ended SCSI controllers connected to SCSI signal converters and one differential storage shelf.

Figure 6–7: Two DEC 7000 Systems and an HSZ40 Unit

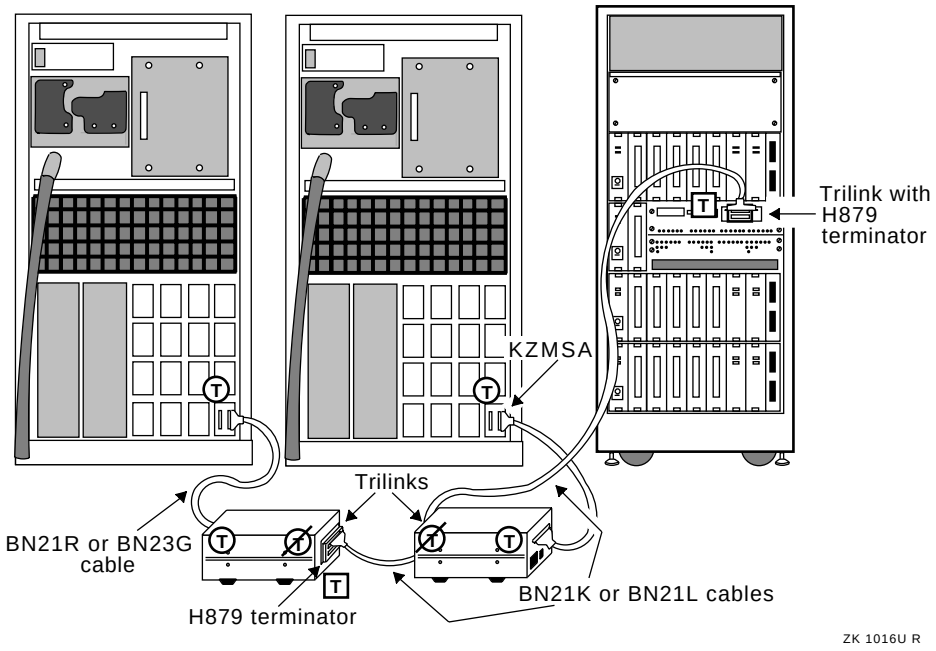

To create the configuration in Figure 6–7, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two DEC 7000s	Set the KZMSAs SCSI IDs to 6 and 7.
SCSI controllers	Two KZMSAs	N/A
Storage	HSZ40 RAID subsystem	N/A
Signal Converters	Two DWZZA-AAs	Terminate the single-ended side, and remove the termination from the differential side.
Connectors	Three H885 trilinks	Attach to the differential side of the DWZZA-AAs and to the HSZ40.
Cables	Two BN21Rs or BN23Gs	Connect the KZMSAs to the single-ended side of the DWZZA-AAs.

Hardware	Description	Configuration Information
	Two BN21Ks or BN21Ls	Connect the trilinks.
Terminators	Two H879-AAs	Attach to the trilinks at the ends of the bus.

6.1.1.4 UltraSCSI Configurations

Figure 6–8 shows two AlphaServer 8200 systems in an Available Server configuration with two HSZ70 DEC RAID controllers. Note that the SCSI bus adapters are KZPSA PCI-to-SCSI adapters. In this figure, the HSZ70s are on the end of the shared SCSI bus.

Figure 6–8: UltraSCSI Shared Bus With an HSZ70 on the Bus End

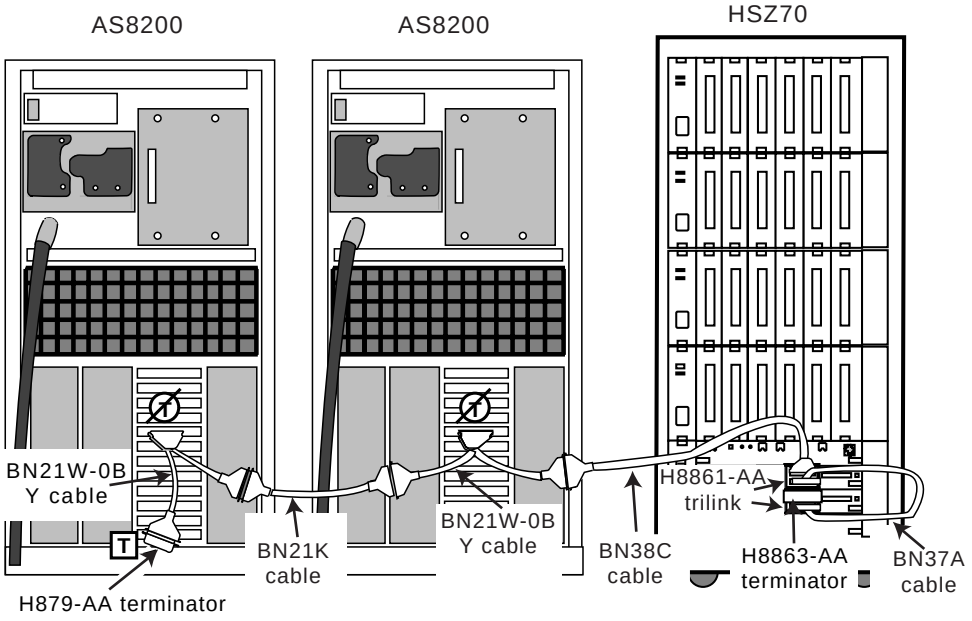

To create the configuration in Figure 6–8, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two DEC 8200s	Set the KZPSAs SCSI IDs to 6 and 7.
SCSI controllers	Two, which may be KZPSAs	Remove the KZPSA onboard SCSI termination resistor SIPs.

Hardware	Description	Configuration Information
Storage	HSZ70 DEC RAID subsystem with two controllers	N/A
Connectors	Two H8861-AA VHDCI trilink connectors	Attach to the HSZ70 controller port connectors.
Cables	Two BN21W-0B Y cables	Install on KZPSA SCSI controller.
	One BN38C	Install between one of the BN21W-0B Y cables and an H8861-AA VHDCI trilink connector.
	One BN21K or BN21L	Install between BN21W-0B Y cables.
Cables	One BN37A	Install between the H8861-AA trilinks.
	One H8863-AA VHDCI differential terminator	Attach to the H8861-AA open connector.
Terminators	One H879-AA	Attach to the BN21W-0B Y cable with a free connector.

Figure 6–9 shows a figure with two AlphaServer 8200 systems in an Available Server configuration with two HSZ70 DEC RAID controllers in the middle of a shared SCSI bus.

Figure 6–9: HSZ70 in the Middle of an UltraSCSI Shared Bus

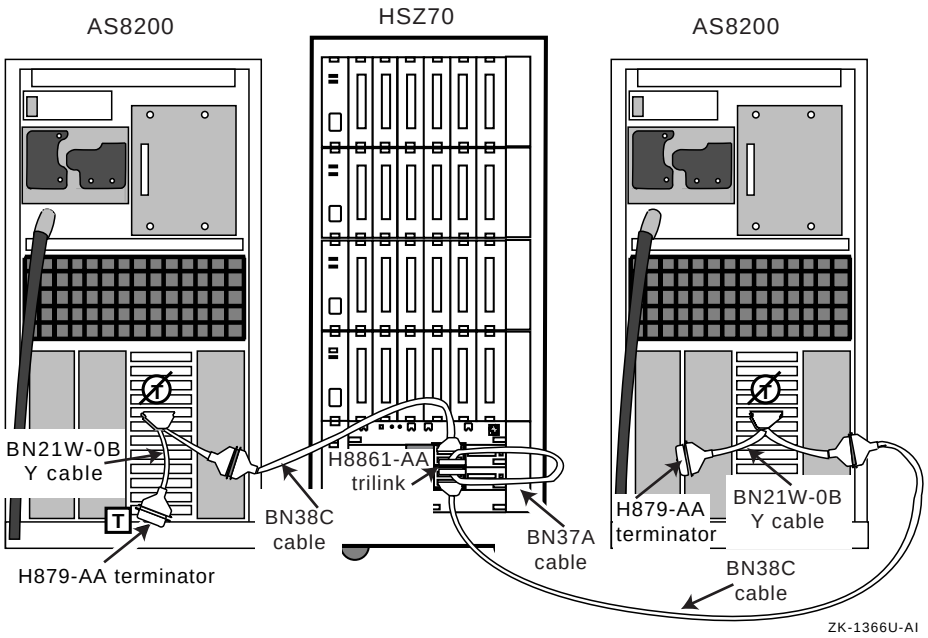

To create the configuration in Figure 6–9, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two DEC 8200s	Set the KZPSAs SCSI IDs to 6 and 7.
SCSI controllers	KZPSAs	Remove the KZPSA onboard SCSI termination resistor SIPs.
Storage	HSZ70 DEC RAID subsystem with two controllers	N/A
Connectors	Two H8861-AA VHDCI trilink connectors	Attach to the HSZ70 contoller port connectors.
Cables	Two BN21W-0B Y cables	Install on KZPSA SCSI controller.
	Two BN38C	Install between BN21W-0B Y cables and an H8861-AA VHDCI trilink connector.

Hardware	Description	Configuration Information
	One BN37A	Install between the H8861-AA trilinks.
Terminators	Two H879-AAs	Attach to the BN21W-0B Y cables at the ends of the bus.

6.1.2 Available Server Configurations That Use Two Systems and Two Shared Buses

More complicated Available Server hardware configurations involve two shared buses. DIGITAL recommends that you set up one bus at a time to avoid confusion.

Figure 6–10 shows an ASE with two systems, two shared buses, and a storage shelf on each bus.

Figure 6–10: Two Systems and Two Shared Buses

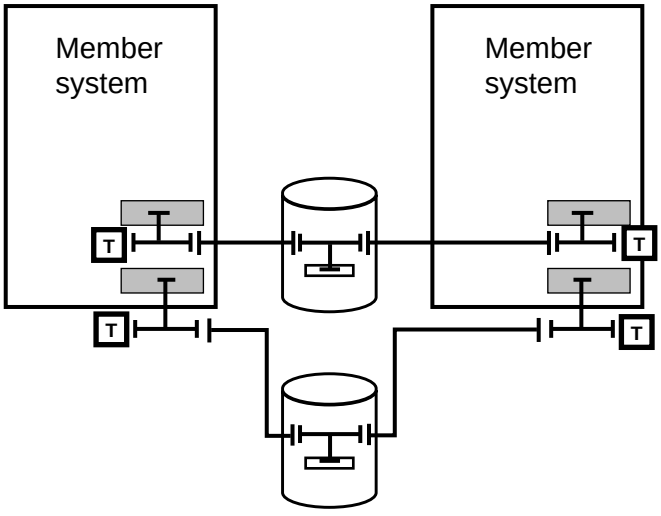

ZK-1113U-AI

Figure 6–11 shows a detailed two-bus configuration that includes two single-ended systems and storage shelves, which are connected to two single-ended buses. The bus formed from the PMAZC modules' port B is shaded to differentiate it from the bus formed from port A.

Figure 6–11: Two DEC 3000/500 Systems With PMAZCs, Two BA350 Storage Shelves, and Two Shared Buses

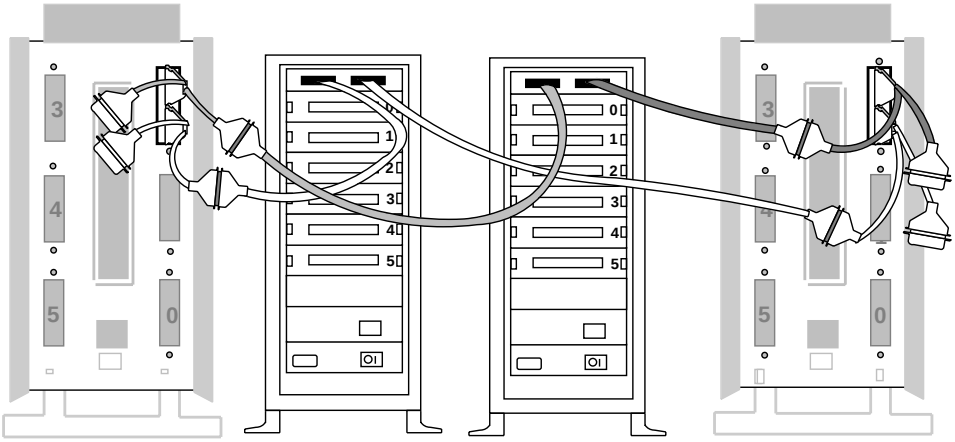

To create the configuration shown in Figure 6–11, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two DEC 3000 Model 500s	Set the PMAZCs SCSI IDs to 6 and 7.
SCSI Controllers	Two PMAZCs	Remove PMAZC internal termination.
Storage	Two BA350s with six disks each	Remove the termination and install disks in slots 0 to 5.
Cables	Four BN21V-0B Y cables	Attach to PMAZCs.
	Four BN21Rs or BN23Gs	Connect the Y cables and BA350s.
Terminators	Four H8574-As	Attach to the Y cables.

Figure 6–12 shows a more complicated two-bus configuration. This configuration includes two DEC 7000 systems, two BA350 storage shelves, one at the end of each bus, and six DWZZA-AA signal converters, connected to two shared buses. Attach terminators to the trilink connectors at each end of each differential bus to terminate the differential shared SCSI buses. The shared bus connected to KZMSA channel 1 is shaded to differentiate it from the bus connected to KZMSA channel 0.

Figure 6–12: Two DEC 7000 Systems, Two BA350 Storage Shelves, and Two Shared Buses

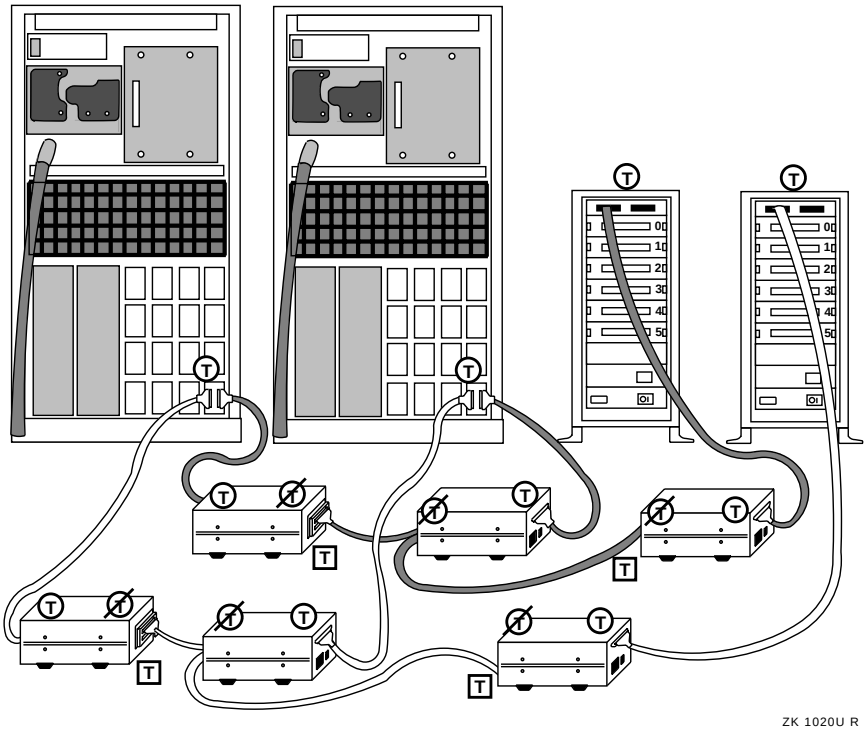

To create the configuration shown in Figure 6–12, you need the following hardware:

Hardware	Description	Configuration Information
Systems	Two DEC 7000s	Set the KZMSAs SCSI IDs to 6 and 7.
SCSI Controllers	Two KZMSAs	N/A
Storage	Two BA350s with six disks each	Terminate the storage shelves and install disks in slots 0 to 5.
Signal Converters	Six DWZZAs	Terminate the single-ended side, and remove the termination from the differential side.
Connectors	Six H885 trilinks	Attach to the DWZZAs.

Hardware	Description	Configuration Information
Cables	Six BN21Rs or BN23Gs	Connect the KZMSAs and BA350s to the single-ended side of the DWZZAs.
	Four BN21Ks or BN21Ls	Connect the trilinks.
Terminators	Four H879-AAs	Attach to the trilinks on the DWZZAs at the ends of the differential bus.

6.1.3 Available Server Configurations That Use Three Systems and Two Shared Buses

Figure 6–13 shows an Available Server configuration with three systems, two shared buses, and storage on each bus.

Figure 6–13: Three Systems and Two Shared Buses

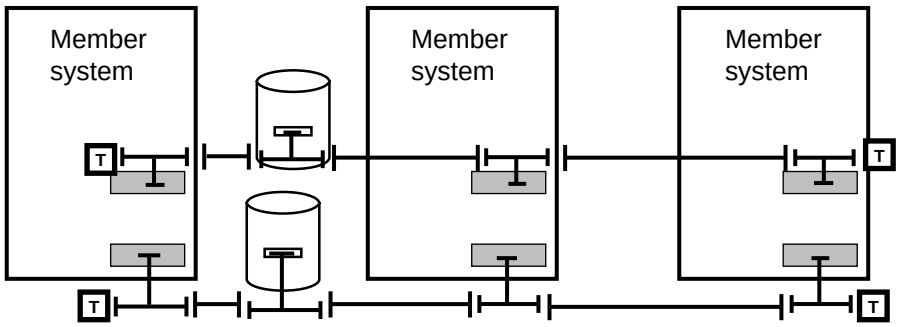

ZK-1114U-AI

6.1.4 Available Server Configurations That Use Four Systems and Three Shared Buses

Figure 6–14 shows an Available Server configuration with four systems, three shared buses, and a storage shelf on each bus.

Figure 6–14: Four Systems and Three Shared Buses

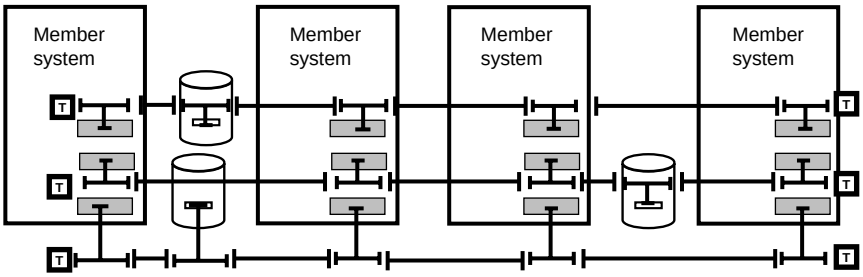

6.2 Sample Production Server Hardware Configurations

This section describes some common Production Server hardware configurations. It does not describe all the possible configurations. You can set up any configuration that adheres to the hardware listed in Chapter 2 and the configuration requirements described in Chapter 4.

6.2.1 Understanding the Configuration Diagrams

Most of the figures in this section use generic devices and other components, such as cables and terminators. To set up your configuration, replace a generic device or component with one listed in Chapter 2.

For example, you can replace a member system with any supported system that has the necessary number of peripheral component interconnect (PCI) slots. Because member systems have different numbers of PCI slots, not all systems can be used in every configuration. You can replace a shared storage shelf with an HSZ40 controller or with a BA350 storage shelf that is connected to a SCSI signal converter.

Except for Figure 6–15, the sample configurations in this chapter show redundant MEMORY CHANNEL interconnects and two shared SCSI buses in each available server environment (ASE).

Figure 6–15 shows the parts of a Production Server hardware configuration.

Figure 6–15: Hardware Configuration Components

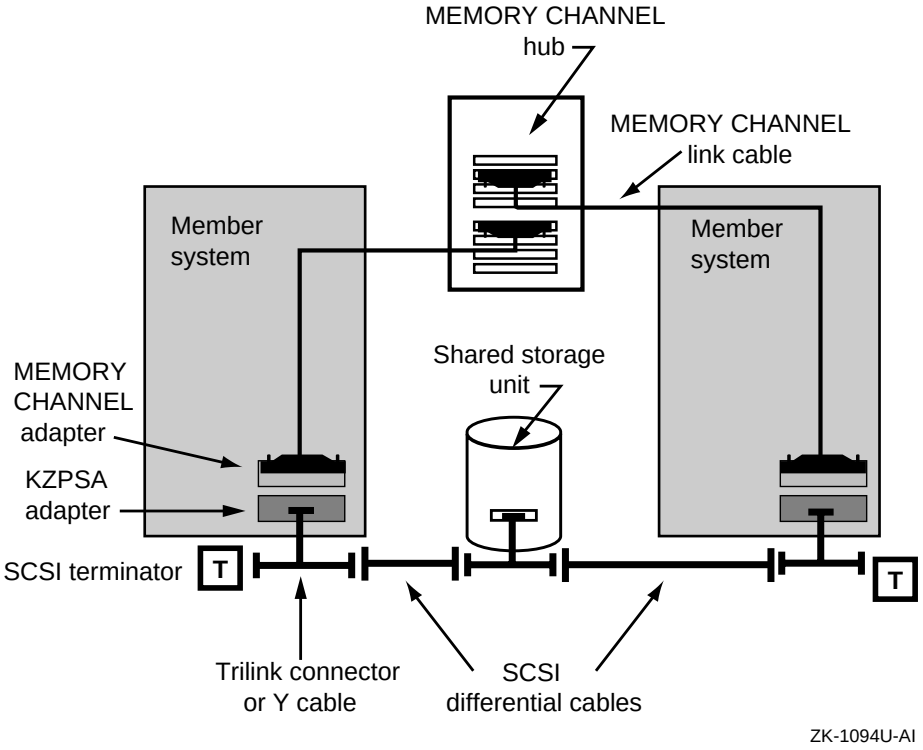

Figure 6–15 shows the following:

- Member systems—You can use any combination of member systems listed in Chapter 2. Member systems have different numbers of PCI slots.
- KZPSA adapters—Install KZPSA PCI-to-SCSI adapters in PCI slots to provide member system connections to shared SCSI buses.
- Shared storage shelves—Use any combination of storage shelves listed in Chapter 2. If you have a single-ended storage shelf, such as a BA350, you must connect it to a SCSI signal converter.
- Trilink connectors or Y cables—Use trilink connectors and Y cables to connect KZPSA SCSI adapters, DEC RAID controllers, or SCSI signal converters installed in single-ended storage shelves to shared SCSI buses.
- SCSI differential cables—Use the appropriate cables to connect the trilink connectors and Y cables attached to member systems and storage shelves.
- SCSI terminators—Terminate both ends of each shared SCSI bus.

- MEMORY CHANNEL adapters—Install MEMORY CHANNEL adapters in PCI slots to connect systems to the cluster interconnects.
- MEMORY CHANNEL hub—A MEMORY CHANNEL hub is required if you have more than two member systems. You can install up to eight line cards in a hub.
- MEMORY CHANNEL link cables—Use a MEMORY CHANNEL link cable to connect a MEMORY CHANNEL adapter to another adapter, or to connect an adapter to a line card installed in a hub.

The following sections show some common Production Server hardware configurations.

6.2.2 Production Server Two-Member System Configuration

Figure 6–16 shows a Production Server hardware configuration that includes two member systems, two terminated shared SCSI buses, and redundant MEMORY CHANNEL interconnects in a virtual hub configuration (there is no MEMORY CHANNEL hub).

Figure 6–16: Two-System Configuration

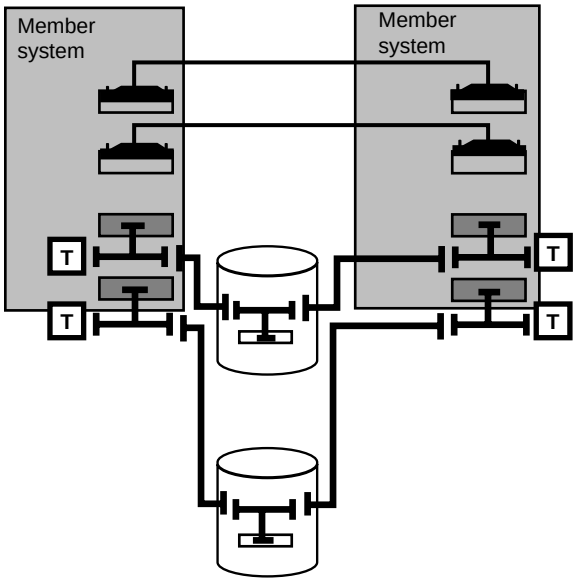

ZK-1183U-AI

Figure 6–17 shows an actual example of the previous configuration using the following hardware:

- Two AlphaServer 2100A systems with KZPSA adapters and H885 trilink connectors

- Two BA350 storage shelves with installed DWZZA-VA signal converters and trilink connectors attached to the signal converters
- Two shared SCSI buses consisting of BN21K or BN21L cables that are connected to the trilink connectors and terminated with H879-AA terminators
- Redundant MEMORY CHANNEL adapters connected by MEMORY CHANNEL link cables to form the MEMORY CHANNEL interconnects

Figure 6–17: Actual Two-System Configuration

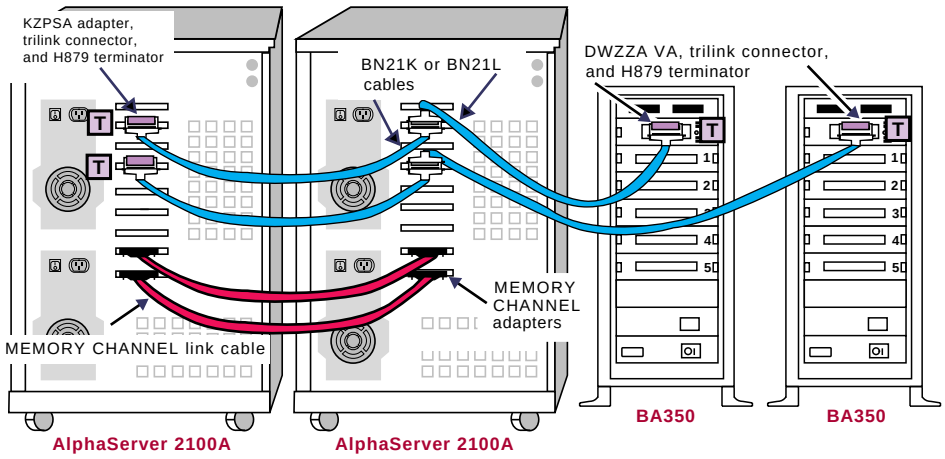

6.2.3 Production Server Three-Member System Configuration

Figure 6–18 shows a Production Server hardware configuration that includes three member systems, two terminated shared SCSI buses, and redundant MEMORY CHANNEL interconnects in a standard hub configuration (it uses MEMORY CHANNEL hubs). Because this configuration requires four PCI slots, it is not supported on AlphaServer 2000 or 2100 systems.

Figure 6–18: Three-System Configuration

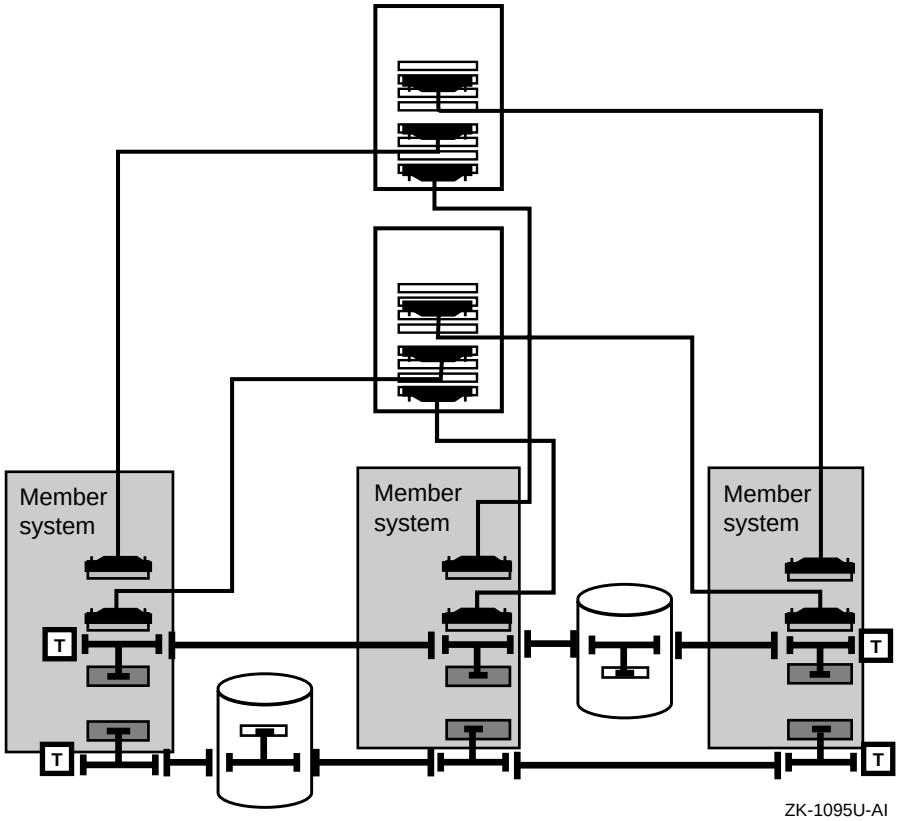

6.2.4 Production Server Four-Member System Configurations

A four-system hardware configuration can include either one or two ASEs. Chapter 1 describes the performance and availability benefits of one- and two-ASE configurations.

Figure 6–19 shows a four-member Production Server configuration that consists of one ASE with two shared SCSI buses, and redundant MEMORY CHANNEL interconnects with MEMORY CHANNEL hubs. Because this configuration requires four PCI slots, it is not supported on AlphaServer 2000 or 2100 systems.

Figure 6–19: Four-System, One-ASE Configuration

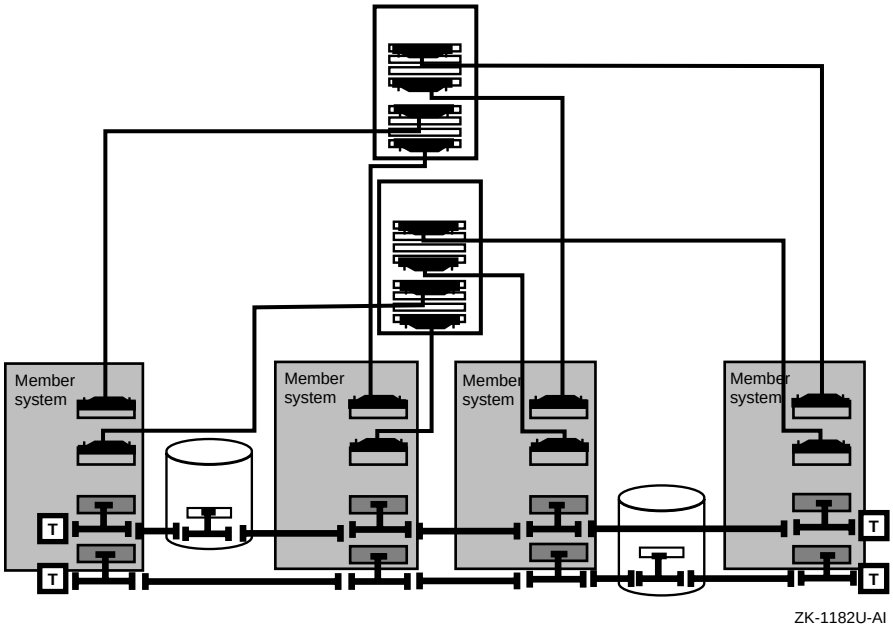

Figure 6–20 shows a Production Server hardware configuration with two ASEs. Each ASE consists of two member systems, each connected to two shared SCSI buses. Each member system also has redundant MEMORY CHANNEL interconnects in a standard hub configuration (including MEMORY CHANNEL hubs). Because this configuration requires four PCI slots, it is not supported on AlphaServer 2000 or 2100 systems.

Figure 6–20: Four-System, Two-ASE Configuration

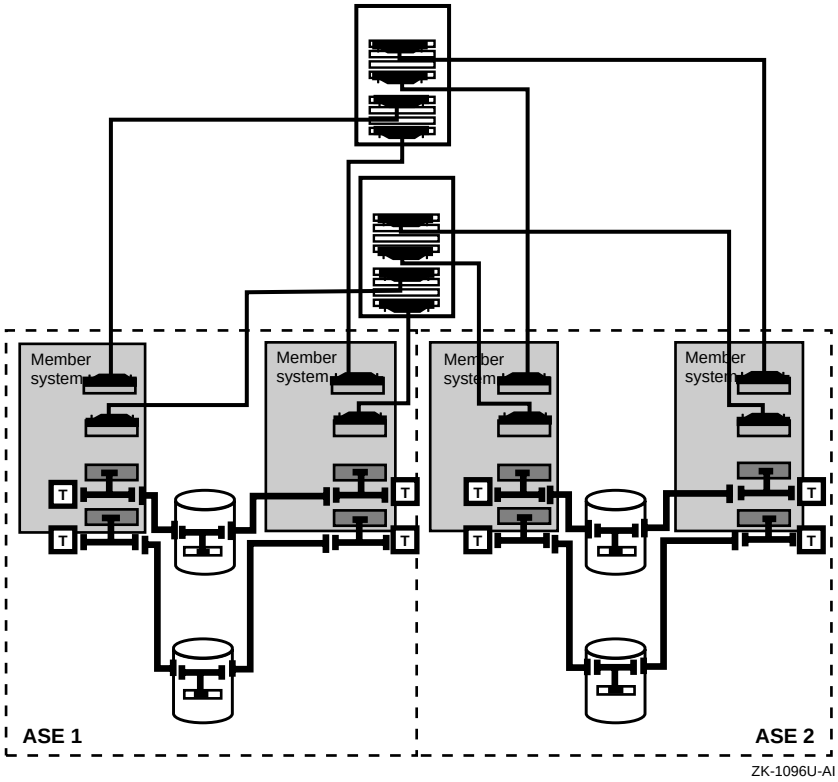

7

Maintaining the Hardware Configuration

This chapter describes how to maintain an Available Server or Production Server hardware configuration. It discusses the following topics:

- Preparing to change the hardware configuration (Section 7.1)
- Stopping available server environment (ASE) activity (Section 7.2)
- Shutting down and starting up the cluster in a Production Server environment (Section 7.3)
- Maintaining member systems (Section 7.4)
- Adding and removing storage shelves (Section 7.5)
- Maintaining disks in an ASE (Section 7.6)
- Adding and removing shared buses (Section 7.7)
- Disconnecting and connecting SCSI signal converters (Section 7.8)
- Maintaining MEMORY CHANNEL interconnects (Section 7.9)

7.1 Preparing to Change the Hardware Configuration

If you want to change your hardware configuration while maintaining available server environment (ASE) operation, make sure that any shared SCSI buses remain terminated. If you use trilink connectors and Y cables to connect devices to the shared SCSI buses, you can disconnect the devices without affecting the bus termination. In addition, if you connect an extra trilink connector or Y cable to a shared bus, you can attach a device to it and expand your configuration without affecting the bus termination. See Chapter 3 for information about maintaining bus termination.

If you are unable to maintain a terminated shared bus, you must shut down the cluster and then change the hardware configuration. Section 7.2 describes how to shut down the cluster.

For the TruCluster Production Server Software product, because the MEMORY CHANNEL is the cluster interconnect, using redundant MEMORY CHANNEL interconnects and MEMORY CHANNEL hubs allow you to easily change your configuration without shutting down the cluster.

Some maintenance tasks require you to use the asemgr utility. See asemgr(8) and the TruCluster Software Products *Administration* manual for information about the utility.

7.2 Stopping ASE Activity

If you cannot isolate a device and maintain a terminated shared bus, you must stop all available server environment (ASE) activity before you can perform maintenance on that device.

Before you stop ASE activity, if you have not already done so, use the asemgr utility to obtain information about each of your ASE services. You should obtain information such as:

- Service name
- Placement policy
- Storage configuration information:
  - Exports list
  - Mount table
  - AdvFS configuration
  - LSM configuration

To stop all ASE activity, follow these steps:

1. Use the asemgr utility to put each ASE service off line. This stops the services.
2. Invoke the /sbin/init.d/asemember stop command on all the member systems to stop the ASE daemons.

After you stop ASE activity, you can perform the desired maintenance.

To restart ASE activity, follow these steps:

1. Invoke the /sbin/init.d/asemember start command on all the member systems to restart the ASE daemons.
2. Use the asemgr utility to put the ASE services on line.

7.3 Shutting Down and Starting Up a Production Server Cluster

To shut down all activity in the cluster, stop the cluster daemons and then stop all shared bus activity.

To stop the cluster daemons, enter the following command on each member system:

```
# /sbin/init.d/clumember stop
```

To start the cluster daemons, enter the following command on each member system:

```
# /sbin/init.d/clumember start
```

To stop and then start the cluster daemons, enter the following command on each member system:

```
# /sbin/init.d/clumember restart
```

To stop or restart ASE activity on a shared SCSI bus in the Production Server environment, follow the steps in Section 7.2.

7.4 Maintaining Member Systems

Occasionally, a member system will need maintenance. For example, you may need to disconnect a member system from the shared SCSI buses to install new hardware. You may want to replace a member system with a newer model or add a member system to your configuration.

Depending on how you set up the shared SCSI buses, you may be able to perform system maintenance without shutting down the available server environment (ASE). The following sections describe how to perform some common system maintenance tasks.

7.4.1 Shutting Down a Member System

To shut down a member system, use the `asemgr` utility to delete the member system from the ASE. This causes any ASE services running on the member system to relocate to another member system. You can also manually relocate the services running on the member system, and then shut down the system in the usual way.

Note

You cannot delete a member if it is included in the list of members that are favored to run the service, according to the service's Automatic Service Placement (ASP) policy. See the TruCluster Software Products *Administration* manual for information about deleting members and ASP policies.

If the system is connected to a SCSI signal converter, you must first turn off the signal converter that is connected to the system and then turn off the system.

To turn on a system that is connected to a SCSI signal converter, you must turn on the system and allow it to complete its startup diagnostics before you turn on the signal converter. Then, invoke the `asemgr` utility on a member system to add the system to the ASE.

7.4.2 Adding a Member System to the Configuration

To add a member system to your configuration, you must install a SCSI bus adapter for each shared SCSI bus the system will be attached to. Then, you must connect the system to all the shared SCSI buses in its ASE (see Section 7.7). Depending on your hardware configuration, you may be able to add a system without shutting down the cluster.

If you have an extra trilink connector or Y cable already connected to all the shared SCSI buses, you can add a member system to your hardware configuration without shutting down. In this case, you can connect the member system to the shared buses without affecting the bus termination and cluster operation. Otherwise, you must shut down the cluster as described in Section 7.2 and Section 7.3.

Additionally, for a Production Server configuration, you may have to install one or two MEMORY CHANNEL adapters. You must shut down the cluster to add a member system in the following cases:

- You are changing from virtual hub mode to standard mode and you must add a MEMORY CHANNEL hub to your configuration. You must shut down the cluster because you must change the jumpers in the MEMORY CHANNEL adapters of all current systems.
- You have a single MEMORY CHANNEL interconnect that includes a hub, but there is no line card available for the new system. You must shut down the cluster because you must power down a hub in order to install the line card.

See Section 7.9 for more information on adding MEMORY CHANNEL interconnects.

7.4.3 Removing a Member System from the Configuration

When you remove a system from an ASE, you do not have to shut down all ASE activity if disconnecting the system from the shared bus does not cause the bus to be unterminated. You can delete the member system from the ASE, shut down and turn off the system, as described in Section 7.2 and Section 7.3, and then disconnect the system from the shared bus.

If you used trilink connectors or Y cables to connect the system to the shared SCSI buses, you can remove a system as follows:

1. If you will replace the system, use the `asemgr` utility to relocate any ASE services running on the member system. If you will not replace the system, use the `asemgr` utility to delete the system from the ASE.
2. Disconnect the system from the shared SCSI buses and the cluster interconnects.

7.4.4 Performing CPU Maintenance

Sometimes you must disconnect a member system to perform maintenance. If the system can be isolated from the shared bus without affecting the bus termination, you can perform the maintenance and the availability of the ASE services is not affected.

If the member system cannot be isolated from the shared bus without affecting the bus termination, you must shut down all ASE activity to perform the maintenance, as described in Section 7.2. Your ASE services are unavailable while you perform the maintenance.

If you can isolate the member system from the shared bus, you can perform hardware maintenance on the system's CPU as follows:

1. Use the `asemgr` utility to relocate the services running on the member system.
2. Delete the member system from the ASE by using the `asemgr` utility.

Note

You cannot delete a member if it is included in the list of members that are favored to run the service, according to the service's Automatic Service Placement (ASP) policy. See the TruCluster Software Products *Administration* manual for information on ASP policies.

3. Shut down the system.
4. Disconnect the member from the shared bus. Make sure that the bus is still terminated so that it functions correctly.
5. Perform the CPU maintenance.
6. Connect the member to the shared bus.
7. Turn on the system.

8. Add the member system to the ASE with the `asemgr` utility.

7.4.5 Adding and Removing Network Interfaces

To add a network interface to a member system in an existing ASE, follow these steps:

1. Delete the member system from the ASE.

Note

You cannot delete a member if it is included in the list of members that are favored to run the service, according to the service's Automatic Service Placement (ASP) policy. See the TruCluster Software Products *Administration* manual for information about deleting members and ASP policies.

2. Turn off the system.
3. Install the network interface.
4. Turn on and reboot the system.
5. Configure the new network interface.
6. Run the `asemgr` utility on an existing member system and add the system to the ASE.
7. Run the `asemgr` utility on the system to specify the new network interface.

To remove a network interface from a member system, follow these steps:

1. Run the `asemgr` utility on the member system to delete the network interface.
2. Delete the member system from the ASE.
3. Turn off the system.
4. Remove the network interface.
5. Turn on and reboot the system.
6. Run the `asemgr` utility on an existing member system and add the system to the ASE.

7.5 Adding and Removing Shared Storage Shelves

If you want to connect another storage shelf to a shared bus without shutting down the ASE, you must have an extra trilink connector or Y

cable already connected to the shared SCSI bus. If your configuration meets this requirement, you can connect the storage shelf to the shared bus without affecting the bus termination and cluster operation. Otherwise, you must shut down the ASE as described in Section 7.2.

You can disconnect a storage shelf from a shared SCSI bus without shutting down the ASE, if you used a trilink connector or Y cable to connect the shelf to the bus.

In addition, if you disconnect a storage shelf from a shared SCSI bus (without affecting the bus termination) or remove a disk from a slot, any service that uses the disks is stopped, unless the disks are part of a mirrored Logical Storage Manager (LSM) volume or are contained in a RAID set.

If you want to connect or disconnect a storage shelf with a single-ended SCSI interface, see Section 7.8 for information about connecting and disconnecting SCSI signal converters.

7.6 Maintaining Disks in an ASE

Most basic system management tasks on the shared disks in the available server environemnt (ASE) are the same as in a noncluster environment. However, you must be careful when performing maintenance on any disk on a shared SCSI bus because of the constant activity on the bus. To perform some types of maintenance, such as upgrading disk firmware, you must either isolate the device from the shared bus or shut down the cluster.

The following sections describe how to maintain the disks in the cluster.

7.6.1 Setting a Disk On Line and Off Line

If you want to set a disk that is used in a service off line, you must ensure that a running service is not using the disk, unless the disk is part of a Logical Storage Manager (LSM) mirrored volume or a mirrored RAID device.

If a disk is being used by a service, you can temporarily stop the service by using the `asemgr` utility's interactive facility or command-line interface to set the service off line. For example, you can use the following command syntax:

asemgr -x [*service*]

After you set the service off line, use the `scu` utility to set the disk off line. Setting a disk off line spins down the disk and allows you to remove it from the storage shelf. For example, to set the `/dev/rz28c` disk off line, enter the following command:

#**scu -f /dev/rrz28c stop**

After you perform the maintenance on a disk, you can set the disk on line. For example:

scu -f /dev/rrz28c start

After you set the disk on line, you can use the `asemgr` utility's interactive facility or command-line interface to set the service that uses the disk on line. For example, you can use the following command syntax:

asemgr -s [*service*]

7.6.2 Adding and Removing Disks

When you add a disk to your hardware configuration, you install it in the storage shelf. The disk must have a unique SCSI ID. In addition, you may have to update the system configuration files to ensure that the systems recognize the new disk. See Section 4.3.1.2 for information about recognizing shared disks in the cluster.

When you remove a disk from a storage shelf, you must ensure that a running service is not using the disk, unless the disk is part of a LSM mirrored logical volume or a mirrored RAID device.

If a disk that you want to remove is being used by a service and the disk will be replaced, you can temporarily stop the service by using the `asemgr` utility to put the service off line. You can then replace the disk and use the `asemgr` utility to put the service that uses the disk on line. You may have to back up the disk before you remove it and then restore the information to the new disk.

If a disk that you want to remove is being used by a service and the disk will not be replaced, use the `asemgr` utility to modify the service and remove the disk from the service. You can then remove the disk from the storage shelf.

To physically remove a disk from a storage shelf, partially pull out the disk from its slot (about 3 to 5 centimeters), wait for the disk to spin down, then completely remove the disk from the slot.

Caution

If you remove the disk from the storage shelf without waiting enough time to allow the disk to spin down, the torque induced by the gyroscopic effect may cause you to drop the disk.

7.6.3 Backing Up and Restoring Disks

Disks that are local to the member (that is, internal disks not shared disks) are not affected by Available Server and can be backed up and restored with the usual methods. Disks that are on the ASE shared bus need special consideration. You do not have to shut down your system to single-user mode to perform safe backups.

There are three ways to back up a disk used in an ASE:

- Use the `asemgr` utility to relocate the service to a specific member so the service will not move. You must back up or restore the disks from this member.
- Use the `asemgr` utility to put the service that uses the disks off line, stopping the service. Back up the disks from any member system. Advanced File System (AdvFS) or LSM disks must be configured on the system from which you are performing the backup.
- Use POLYCENTER NetWorker Save and Restore to back up the disks in the ASE services. See the TruCluster Software Products *Administration* manual for information about ASE services. See the NetWorker Version 3.2 documentation for information about using NetWorker to back up an ASE service's storage.

For UNIX file systems, back up the disk using the `dump` command and the raw device file `/dev/rrz`*nn*. Use the `restore` command to restore a disk. For AdvFS filesets, from the member that is running the service, you can use the `clonefset` command to clone a fileset, and then use the `vdump` and `vrestore` commands to back up and restore the cloned fileset.

7.6.4 Handling Disk Failures

The way failed disks are handled in an ASE depends on whether you are using LSM or RAID. If a failure occurs in a disk that is not part of an LSM or RAID mirrored volume, the service stops. After the disk has been replaced and any data restored, you can restart the service.

If a disk that is used in an LSM volume fails, see the DIGITAL UNIX *Logical Storage Manager* manual for information about replacing failed LSM disks.

If a failure occurs in a disk that is part of an LSM or RAID mirrored volume, you can replace the disk while the service is running. If a disk that is mirrored with RAID fails, see the RAID documentation for information about how to handle this situation.

After a failed or previously unavailable part of an LSM mirrored volume becomes available again, you can reincorporate the device into the service

by resynchronizing the mirrored volume outside of the cluster and then rereserving the devices. You rereserve devices by using the `asemgr` utility and choosing the Advanced Utilities menu item. This method will not interrupt the service.

See the TruCluster Software Products *Administration* manual for more information on handling disk failures with AdvFS and LSM.

7.7 Adding and Removing Shared Buses

If you want to add a shared SCSI bus to your hardware configuration, you must shut down the cluster, as described in Section 7.2, and prepare the systems and storage shelves for the new shared bus connection.

You can remove a shared SCSI bus without shutting down the cluster if you used trilink connectors or Y cables to connect the member systems and storage shelves to the shared SCSI bus. If your configuration meets this requirement, you can disconnect all the devices from the bus.

7.8 Disconnecting and Connecting SCSI Signal Converters

If you are using a storage shelf with a single-ended SCSI interface in your hardware configuration, it must be connected to a SCSI signal converter.

If you want to disconnect a SCSI signal converter (and the single-ended storage shelf) from a shared bus, you must turn off the SCSI signal converter before disconnecting the cables. To reconnect it to the shared bus, connect the cables before turning on the SCSI signal converter.

Use the power switch to turn off a standalone SCSI signal converter (DWZZA-AA or DWZZB-AA). To turn off a StorageWorks building block (SBB) SCSI signal converter (DWZZA-VA or DWZZB-VW), pull it from its disk slot.

7.9 Maintaining MEMORY CHANNEL Interconnects in a Production Server Environment

The following sections contains information about maintaining MEMORY CHANNEL interconnects. See the MEMORY CHANNEL *User's Guide* for detailed information about maintaining the MEMORY CHANNEL hardware.

If you are adding a new system to your configuration, see Section 7.4.2 for information about connecting the system to the MEMORY CHANNEL interconnects.

7.9.1 Determining the Primary MEMORY CHANNEL Interconnect

Some MEMORY CHANNEL interconnect maintenance tasks require you to determine which interconnect is the primary (active) interconnect and which interconnect is the secondary (inactive) interconnect. Examine the startup message for the last boot and note which peripheral component interconnect (PCI) slot contains the adapter for the primary interconnect.

The following is an example of the startup messages:

```
mchan0: Module revision = 11
mchan0: jumpered as HUB configuration
mchan0 at pci0 slot 7
mchan1: Module revision = 11
mchan1: jumpered as HUB configuration
mchan1 at pci0 slot 8
```

7.9.2 Adding or Removing MEMORY CHANNEL Interconnects

If you want to change from a single MEMORY CHANNEL interconnect to redundant MEMORY CHANNEL interconnects without shutting down the cluster, follow these steps to add an interconnect:

1. Use the `asemgr` utility to relocate the available server environment (ASE) services running on a member system. If you are using virtual hub mode (you are not using a MEMORY CHANNEL hub), do this on the VH0 system.
2. Shut down the system.
3. Install another MEMORY CHANNEL adapter in the system, setting the jumpers for the mode of the existing interconnect, in either standard mode (if you are using a hub) or virtual hub mode.
4. Connect a MEMORY CHANNEL link cable to the newly installed MEMORY CHANNEL adapter.
5. If you are using a MEMORY CHANNEL hub, connect the link cable to the line card which occupies the same hub slot position as the line card to which the first adapter is connected. See Chapter 5 for information about connecting adapters to line cards in hubs.
6. Turn on the system.
7. Repeat these steps for each system. If you are using a MEMORY CHANNEL hub, turn on the hub after the last system is connected. If you are not using a MEMORY CHANNEL hub, connect the adapter in the first system to the adapter in the second system with the link cable.

To change from redundant MEMORY CHANNEL interconnects to a single MEMORY CHANNEL interconnect, you must remove an interconnect. To do this, follow these steps:

1. Turn off the hub on the interconnect that you want to remove.
2. Use the `asemgr` utility to relocate the ASE services running on a member system. If you are using virtual hub mode (you are not using a MEMORY CHANNEL hub), do this on the VH0 system.
3. Shut down the system.
4. Deinstall the MEMORY CHANNEL adapter that is connected to the interconnect you want to remove.
5. Reboot the system.
6. Perform steps 2 through 5 on all the systems.

7.9.3 Adding a MEMORY CHANNEL Hub

If you want to change from virtual hub mode, which does not require a MEMORY CHANNEL hub, to standard mode, you must add a hub to your configuration. To do this, you must shut down the cluster because you must change the jumpers on the MEMORY CHANNEL adapters.

See the MEMORY CHANNEL *User's Guide* for information about adapter jumpers.

7.9.4 Replacing a MEMORY CHANNEL Hub

If you need to replace a MEMORY CHANNEL hub and you have only one MEMORY CHANNEL interconnect, you must shut down the cluster as described in Section 7.2.

If you have redundant MEMORY CHANNEL interconnects, you can turn off the hub, replace the hub, and then reconnect the cables as described in Chapter 5. Then, you must reboot the member systems, one at a time.

7.9.5 Connecting a MEMORY CHANNEL Adapter to a Line Card

If a system is connected to a line card that fails and you have an extra line card available in each MEMORY CHANNEL hub, you can connect a MEMORY CHANNEL adapter to a line card without shutting down the cluster. However, you must reboot the system after you connect the system to the line card.

If you have redundant interconnects, make sure that you connect the system to line cards that are in the same slot position in the hubs.

7.9.6 Disconnecting and Connecting a Link Cable

If you have a single MEMORY CHANNEL interconnect and you disconnect a MEMORY CHANNEL link cable, the member system that was connected to the link cable will crash. You must then reconnect the link cable and reboot the system.

If you have redundant interconnects and you disconnect a MEMORY CHANNEL link cable that is part of the secondary interconnect (the inactive interconnect), you can reconnect the cable.

However, if you have redundant interconnects and you disconnect a MEMORY CHANNEL link cable that is part of the primary interconnect (the active interconnect), you must reconnect the cable and then reboot the member system.

Glossary

The terms in this glossary are commonly used in a TruCluster Software environment.

action script

Scripts that are used to make an application or data highly available by configuring an application or data on a member system. Action scripts break down a procedure (for example, starting an application or exporting data) into a series of steps, which are performed in order when executing that procedure. There are five types of action scripts: add, delete, start, stop, and check action scripts, and there are two versions of each type: internal action scripts, which cannot be modified manually, and user-defined action scripts, which allow you to customize the behavior of the service.

adapter

A device that converts the protocol and hardware interface of one bus type into that of another bus.

address switches

Electrical switches on the side or rear of some disk drives that determine the SCSI address setting for the drive.

advanced RISC computing

External interface to console firmware for operating systems that expect firmware compliance with the Advanced RISC Computing Standard Specification.

ARC

See advanced RISC computing.

available server environment

A set of systems, disks, shared SCSI buses and software that allows you to configure applications and disks so that they are highly available to client systems.

ASE

See available server environment.

ASE ID

A number from 0 to 63 that identifies an ASE within a cluster and allows the asemgr utility to generate unique clusterwide names for DRD special

files. Each ASE in a cluster has its own distinct ASE ID. All cluster members in the same ASE use the same ASE ID.

ASE service

A service that an administrator sets up in an ASE by using the asemgr utility. TruCluster software uses a service to maintain the availability of applications or data. A service consists of a unique name, an ASP policy, an application or disk specification, and action scripts that contain the commands to start and stop the application or to fail over the disk data. The action scripts implement the status changes for the service by performing necessary configuration changes and starting and stopping processes.

A member system in an ASE runs a service until a hardware or software failure or an explicit action by an administrator causes the service to run on another member system in the ASE.

Automatic service placement policy

Enables you to control which member systems are allowed to run a service. You must specify an ASP policy when you add a service. For example, you can allow any member system to run a service, or you can restrict a service to a specific member system or systems.

ASP Policy

See automatic service placement policy.

availability

The amount of time that hardware or software is available during the time it is scheduled to be available. For the TruCluster software, the ability to function despite a specific hardware or software failure. See also **highly available.** To make an ASE service available despite a particular failure, it is necessary to make the hardware and software it depends on capable of operating despite that failure. For example, a DRD service can be made available despite an MEMORY CHANNEL interconnect failure by configuring a redundant MEMORY CHANNEL interconnect so that if the primary MEMORY CHANNEL interconnect fails, the DRD service will use the other MEMORY CHANNEL interconnect.

bus

Flat or twisted-wire cable or a blackplane composed of individual identical circuits. A bus interconnects computer system components to provide communications paths for addresses, data, and control information.

client

A computer system that uses resources provided by another computer, called a **server**.

cluster

A loosely coupled collection of servers that share storage and other resources and make applications and data highly available. A cluster consists of communications media, member systems, peripheral devices, and applications. The systems communicate over a high-performance interconnect.

cluster configuration map

A file (`/etc/CCM`) that statically records the hardware configuration of a cluster for display by the Cluster Monitor utility. You use the cluster_map_create utility to generate a cluster configuration map when you first configure a cluster and, subsequently, each time you add or remove hardware.

cluster interconnect

Private physical bus employed by cluster members for intracluster communications.

Cluster Monitor

Cluster software component that provides a graphical view of the cluster configuration. You can use the Cluster Monitor utility to monitor the availability of services and the connectivity among member systems in the cluster. You can also use it to manage services and to start disk management applications.

cold swap

The ability to turn off power to a device, replace it, and then turn on power to the device.

connection manager

Cluster software component that coordinates participation of systems in the cluster, and maintains cluster integrity when computers join or leave the cluster.

differential SCSI bus

A SCSI bus where the signal's level is determined by the potential difference between two wires.

distributed lock manager

Cluster software component that synchronizes access to shared resources among cooperating processes throughout the cluster.

DLM

See distributed lock manager.

distributed raw disk

A storage technology that uses an ASE service to provide clusterwide access to a disk. The service exports a raw disk to all member systems. The

raw disk must be on a shared SCSI bus. If the member system running the DRD service fails, the service can fail over to another member system on the same shared SCSI bus.

DRD

See distributed raw disk.

failover

A transfer of the responsibility to provide an ASE service. A failover occurs when a hardware or software failure causes a service to restart on a viable member system.

Fast SCSI

An optional mode of SCSI-2 that allows transmission rates of up to 10 MB per second.

fast bus speed

A bus speed that uses the fast synchronous transfer option, enabling I/O devices to attain high peak-rate transfers (10 MB per second) in synchronous mode.

firmware

Software code stored in hardware.

highly available

In the TruCluster software, the ability to survive any single hardware or software failure.

A cluster can be considered highly available if the hardware and software provides protection against any single failure, such as a system or disk failure or a SCSI cable disconnection.

An ASE service can be considered highly available if the hardware it depends on provides protection against any single failure, and the service is configured to fail over in case of a failure.

hot swap

The ability to replace a device on a shared bus while the bus is active.

hot standby

A member system that is available to run an ASE service if the primary member system running the service fails.

local bus

See private SCSI bus.

lock file

A file that indicates that operations on one or more other files are restricted or prohibited. The presence of the lock file can be used as the

indication, or the lock file can contain information describing the nature of the restrictions.

Logical Storage Manager

A disk storage management tool that protects against data loss, improves disk I/O performance, and customizes the disk configuration.

System administrators use LSM to perform disk management functions without disrupting users or applications accessing data on those disks.

In an ASE, you can use LSM to mirror disks across shared SCSI buses. This results in greater data reliability and integrity. You can use a DRD service to make an LSM volume accessible clusterwide.

LSM

See Logical Storage Manager.

LSM disk group

A group of LSM disks that share a common configuration. The configuration information for an LSM disk group consists of a set of records describing objects including LSM disks, LSM volumes, LSM plexes, and LSM subdisks that are associated with the LSM disk group. Each LSM disk group has an administrator-assigned name that can be used to reference that LSM disk group.

LSM volume

An LSM volume is a DIGITAL UNIX special device that contains data used by a UNIX file system, a database, or other applications. LSM transparently places an LSM volume between applications and a physical disk. Applications then operate on the LSM volume rather than on the physical disk. For example, a file system is created on an LSM volume rather than on a physical disk.

An LSM volume presents block and raw interfaces that are compatible in their use with disk partition special devices. Because an LSM volume is a virtual device, it can be mirrored, spanned across disk drives, moved to use different storage, and striped using administrative commands. The configuration of an LSM volume can be changed using LSM utilities without disrupting applications or file systems that are using the LSM volume.

LSM plex

An LSM plex is a copy of an LSM volume's logical data address space, sometimes known as a mirror. An LSM volume can have up to eight LSM plexes associated with it. A read can be satisfied from any LSM plex, while a write is directed to all LSM plexes.

logical unit number
A physical or virtual peripheral device addressable through a target. LUNs use their target's bus connection to communicate on a SCSI bus.

LUN
See logical unit number.

member system
The basic computing resource in a cluster. A member system must be physically connected to a cluster interconnect and at least one shared SCSI bus. The connection manager dynamically determines cluster membership based on communications among the cluster members.

MEMORY CHANNEL
A PCI-based cluster interconnect that promotes fast and reliable communications between cluster members.

MEMORY CHANNEL interconnect
MEMORY CHANNEL interconnect. A type of cluster interconnect that consists of a MEMORY CHANNEL adapter installed in a PCI slot in each member system, one or more MEMORY CHANNEL link cables to connect the adapters, and an optional MEMORY CHANNEL hub.

mount point
A directory file that is the name of a mounted file system.

network
Two or more computing systems that are linked for the purpose of exchanging information and sharing resources.

network interface
The network adapter and the software that allows a system to communicate over a network.

partition
An abnormal condition in which nodes in an existing TruCluster software configuration divide into two independent clusters.

Peripheral component interconnect
An industry-standard expansion I/O bus that is a synchronous, asymmetrical I/O channel.

PCI
See peripheral component interconnect.

private SCSI bus
A SCSI bus that connects private storage to the local system.

private storage

A storage device on a private SCSI bus. Storage devices include hard disk, floppy disk, and compact disk drives, tape drives, and other devices.

redundant array of inexpensive disks

A technique that organizes disk data to improve performance and reliability. RAID has three attributes:

- It is a set of physical disks viewed by the user as a single logical device or multiple logical devices.
- Disk data is distributed across the physical set of drives in a defined manner.
- Redundant disk capacity is added so data can be recovered if a drive fails.

RAID

See redundant array of inexpensive disks.

redundant

Describes duplicate hardware that provides spare capacity that can be used when a component fails.

relocate a service

To stop an ASE service on one member system and restart it on another member system.

relocation policy

See ASP policy.

script

A program to be interpreted and executed by the shell.

SCSI

See Small Computer System Interface.

SCSI-2

An extension to the original SCSI standard featuring multiple systems on the same bus and hot swap. Hot swap is the ability to replace a device on a shared bus while the bus is active. The SCSI-2 standard is ANSI standard X3.T9.2/86-109.

SCSI adapter

A storage adapter that provides a connection between an I/O bus and a SCSI bus.

SCSI bus

A bus that supports the transmission and signalling requirements of a SCSI protocol. See shared SCSI bus and private SCSI bus.

SCSI bus speed
The data transfer speed for a SCSI bus. SCSI bus speed can be either slow, up to 5 million bytes per second, or fast, up to 10 million bytes per second.

SCSI controller
An adapter or module that is installed in a member system's I/O bus slot that provides a connection to a shared SCSI bus.

SCSI device
A SCSI controller, peripheral controller, or intelligent peripheral that can be attached to a SCSI bus.

SCSI ID
Unique address that identifies a device on a SCSI bus.

server
A computing system that provides a specific set of applications or data to clients. For a service in an ASE, the server is the member system that is currently running the service.

service
See ASE service.

shared SCSI bus
A SCSI bus that is connected to more than one member system and, optionally, one or more storage devices.

shared storage
Disks that are connected to a shared SCSI bus.

signal converter
Converts signals between a single-ended SCSI bus and a differential SCSI bus.

single-ended SCSI bus
A signal path in which one data lead and one ground lead are utilized to make a device connection. This transmission method is economical, but is more susceptible to noise than a differential SCSI bus.

Small Computer System Interface
An American National Standards Institute (ANSI) standard interface for connecting disks and other peripheral devices to a computer system. SCSI-based devices can be configured in a series, with multiple devices on the same bus. In this manual, SCSI refers to SCSI-2. SCSI is pronounced *skuh-zee*.

SRM
External interface to console firmware for operating systems that expect firmware compliance with the Alpha System Reference Manual (SRM).

standard mode
A MEMORY CHANNEL interconnect configuration that uses an MEMORY CHANNEL hub to connect MEMORY CHANNEL adapters. To set up an MEMORY CHANNEL interconnect in standard mode, use a link cable to connect each MEMORY CHANNEL adapter to a line card installed in an MEMORY CHANNEL hub.

StorageWorks
DIGITAL's modular storage subsystem (MSS), which consists of a family of mass storage products that can be configured to meet current and future storage needs.

subset
An installable software module that is compatible with the DIGITAL UNIX setld software installation utility.

system bus
The private (nonshared) interconnect used on the CPU subsystem. This bus connects the processor module, the memory module, and the I/O module.

target
A device that can be addressed by a SCSI ID on a SCSI bus.

terminator
Resistor array device used for terminating a SCSI bus. A SCSI bus must be terminated at its two physical ends.

tie-breaker disk
One to three disks used by the connection manager to prevent cluster partitions in a two-member cluster that does not use a hub.

trilink connector
A connector that joins two cables to a single device.

virtual hub mode
A MEMORY CHANNEL interconnect configuration that does not use an MEMORY CHANNEL hub to connect MEMORY CHANNEL adapters. Virtual hub mode is supported only for clusters that have two member systems. To set up an MEMORY CHANNEL interconnect in virtual hub mode, use an MEMORY CHANNEL link cable to connect the MEMORY CHANNEL adapter in one member system to the corresponding MEMORY CHANNEL adapter in the other member system.

warm swap
To replace a device on a shared bus while the bus is not active.

Y cable
A cable that joins two cables to a single device.

Index

A

B

C

D

F

H

I

J

K

L

M

N

O

P

R

S

T

U

V

Y

How to Order Additional Documentation

Technical Support

If you need help deciding which documentation best meets your needs, call 800-DIGITAL (800-344-4825) before placing your electronic, telephone, or direct mail order.

Electronic Orders

To place an order at the Electronic Store, dial 800-234-1998 using a modem from anywhere in the USA, Canada, or Puerto Rico. If you need assistance using the Electronic Store, call 800-DIGITAL (800-344-4825).

Telephone and Direct Mail Orders

Your Location	Call	Contact
Continental USA, Alaska, or Hawaii	800-DIGITAL	Digital Equipment Corporation P.O. Box CS2008 Nashua, New Hampshire 03061
Puerto Rico	809-754-7575	Local Digital subsidiary
Canada	800-267-6215	Digital Equipment of Canada Attn: DECdirect Operations KAO2/2 P.O. Box 13000 100 Herzberg Road Kanata, Ontario, Canada K2K 2A6
International	—	Local Digital subsidiary or approved distributor
Internal (submit an Internal Software Order Form, EN-01740-07)	—	SSB Order Processing – NQO/V19 *or* U.S. Software Supply Business Digital Equipment Corporation 10 Cotton Road Nashua, NH 03063-1260

Reader's Comments

TruCluster Software Products
Hardware Configuration
AA-R88GA-TE

Digital 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)	☐	☐	☐	☐
Clarity (easy to understand)	☐	☐	☐	☐
Organization (structure of subject matter)	☐	☐	☐	☐
Figures (useful)	☐	☐	☐	☐
Examples (useful)	☐	☐	☐	☐
Index (ability to find topic)	☐	☐	☐	☐
Usability (ability to access information quickly)	☐	☐	☐	☐

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

digital™

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

DIGITAL EQUIPMENT CORPORATION
UEG PUBLICATIONS MANAGER
ZKO3−3/Y32
110 SPIT BROOK RD
NASHUA NH 03062−9987

Do Not Cut or Tear − Fold Here

Cut on
Dotted
Line