



Sun StorEdge™ T3 Disk Tray Configuration Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 U.S.A.
650-960-1300

Part No. 806-4210-12
March 2001, [Revision A](#)

[Send comments about this document to: docfeedback@sun.com](mailto:docfeedback@sun.com)

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, Solstice DiskSuite, docs.sun.com, OpenBoot, SunSolve, JumpStart, StorTools, Sun Enterprise, Sun StorEdge, Sun Ultra, Sun Fire, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, Solstice DiskSuite, docs.sun.com, SunSolve, OpenBoot, JumpStart, StorTools, Sun Enterprise, Sun StorEdge, Sun Ultra, Sun Fire, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Preface

The *Sun StorEdge T3 Disk Tray Configuration Guide* describes the recommended configurations for the Sun StorEdge™ T3 disk tray for high availability, the maximum performance, and maximum storage capability. This guide is intended for Sun™ field sales and technical support personnel.

Before You Read This Book

Read the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual* for product overview information.

How This Book Is Organized

Chapter 1 describes the connection ports and Fibre Channel loops for the Sun StorEdge T3 disk tray. It also describes basic rules and recommendations for configuring the Sun StorEdge T3 disk tray.

Chapter 2 describes how to configure the Sun StorEdge T3 disk tray's global parameters.

Chapter 3 describes how to configure Sun StorEdge T3 disk trays into partner groups to form redundant storage systems.

Chapter 4 provides reference configuration examples.

Chapter 5 describes host connections for the Sun StorEdge T3 disk tray.

Chapter 6 describes disk tray cabling for the Sun StorEdge T3 disk tray.

Using UNIX Commands

This document does not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information:

- AnswerBook2™ online documentation for the Solaris™ operating environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1 Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

TABLE P-3 Related Documentation

Application	Title	Part Number
Site Preparation	<i>Sun StorEdge T3 Disk Tray Site Preparation and Planning Guide</i>	806-4212
Installation overview	<i>Sun StorEdge T3 Disk Tray Installation Task Map</i>	806-1061
Installing, using, and servicing the disk tray	<i>Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual</i>	806-1062
Administration	<i>Sun StorEdge T3 Disk Tray Administrator's Guide</i>	806-1063
Troubleshooting	<i>Sun StorEdge T3 Disk Tray Field Service Manual</i>	806-4213
Release notes	<i>Sun StorEdge T3 Disk Tray Release Notes</i>	806-1497
Disk drive specifications	<i>18 Gbyte, 1-Inch, 10K rpm Disk Drive Specifications</i>	806-1493
	<i>36.4 Gbyte, 10K rpm Disk Drive Specifications</i>	806-1491

TABLE P-3 Related Documentation (*Continued*)

Application	Title	Part Number
	<i>73 Gbyte, 10K rpm, 1.6-Inch Disk Drive Specifications</i>	806-4800
Configuration Rules	“Configuration Rules for Mission Critical Storage” document (also known as the “Big Rules” document)	n/a
Cluster installation and service	<i>Sun Cluster 2.2 Hardware Service Manual</i>	806-4749

Accessing Sun Documentation Online

You can access a select group of Sun technical documentation on the Web. You can browse the documentation archive at:

<http://www.sun.com/products-n-solutions/hardware/docs>

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

docfeedback@sun.com

Please include the part number (806-4210-12) of your document in the subject line of your email.

Contents

Preface iii

1. Sun StorEdge T3 Disk Tray Configuration Overview 1

Product Description 1

 Connection Ports 2

 FC-AL Configurations 4

 Partner Group Configuration 4

 Standalone Configuration 5

Configuration Guidelines and Restrictions 7

Configuration Recommendations 7

Supported Platforms 8

Sun Cluster 2.2 Support 8

2. Configuring Global Parameters 9

Cache 9

 Configuring Cache for Performance and Redundancy 10

 Configuring Data Block Size 10

 Selecting a Data Block Size 11

 Enabling Mirrored Cache 12

 Configuring Cache Allocation 12

Logical Volumes	13
Guidelines for Configuring Logical Volumes	13
Determining How Many Logical Volumes You Need	14
Determining Which RAID Level You Require	14
Determining Whether You Need a Hot Spare	15
Creating and Labeling a Logical Volume	16
Setting the LUN Reconstruction Rate	16
Using RAID Levels to Configure Redundancy	17
RAID 0	17
RAID 1	18
RAID 5	18
Configuring RAID Levels	18
3. Configuring Partner Groups	19
Understanding Partner Groups	19
How Partner Groups Work	21
Creating Partner Groups	22
4. Configuration Examples	27
Direct Host Connection	28
Single Host With One Controller Unit	28
Single Host With Two Controller Units Configured as a Partner Group	29
Host Multi-Pathing Management Software	30
Single Host With Four Controller Units Configured as Two Partner Groups	32
Single Host With Eight Controller Units Configured as Four Partner Groups	34
Hub Host Connection	36
Single Host With Four Controller Units and Two Hubs Configured as Two Partner Groups	36

Single Host With Eight Controller Units and Two Hubs Configured as Four Partner Groups	38
Dual Hosts With Four Controller Units and Two Hubs	40
Dual Hosts With Eight Controller Units and Two Hubs	42
Switch Host Connection	45
Dual Hosts With Two Controller Units and Two Switches Configured as One Partner Group	45
Dual Hosts With Eight Controller Units and Two Switches Configured as Four Partner Groups	47
5. Host Connections	49
Sun Enterprise SBus+ and Graphics+ I/O Boards	50
System Requirements	50
Sun StorEdge PCI FC-100 Host Bus Adapter	51
System Requirements	51
Sun StorEdge SBus FC-100 Host Bus Adapter	52
System Requirements	52
Sun StorEdge PCI Single Fibre Channel Network Adapter	53
System Requirements	53
6. Disk Tray Cabling	55
Overview of Disk Tray Cabling	55
Data Path	55
Administration Path	56
Connecting Partner Groups	56
Single Disk Tray Configurations	58
Partner Group Configurations	59
Glossary	61

Figures

- FIGURE 1-1 Controller Card and Ports 3
- FIGURE 1-2 Interconnect Card and Ports 4
- FIGURE 1-3 Sun StorEdge T3 Disk Tray Partner Group Configuration 5
- FIGURE 1-4 Sun StorEdge T3 Disk Tray Standalone Configuration 6
- FIGURE 3-1 Sun StorEdge T3 Disk Tray Partner Group 20
- FIGURE 3-2 Connecting the Interconnect Cables 23
- FIGURE 3-3 Fully Cabled Partner Group 24
- FIGURE 4-1 Single Host Connected to One Controller Unit 28
- FIGURE 4-2 Single Host With Two Controller Units Configured as a Partner Group 29
- FIGURE 4-3 DMP or AP Failover Configuration 31
- FIGURE 4-4 Single Host With Four Controller Units Configured as Two Partner Groups 33
- FIGURE 4-5 Single Host With Eight Controller Units Configured as Four Partner Groups 35
- FIGURE 4-6 Single Host With Four Controller Units and Two Hubs Configured as Two Partner Groups 37
- FIGURE 4-7 Single Host With Eight Controller Units and Two Hubs Configured as Four Partner Groups 39
- FIGURE 4-8 Dual Hosts With Four Controller Units and Two Hubs 41
- FIGURE 4-9 Dual Hosts With Eight Controller Units and Two Hubs 44
- FIGURE 4-10 Dual Hosts With Two Controller Units and Two Switches 46
- FIGURE 4-11 Dual Hosts With Eight Controller Units and Two Switches 48
- FIGURE 5-1 Sun Enterprise 6x00/5x00/4x00/3x00 SBus+ I/O Board 50

FIGURE 5-2	Sun StorEdge PCI FC-100 Host Bus Adapter	51
FIGURE 5-3	Sun StorEdge SBus FC-100 Host Bus Adapter	52
FIGURE 5-4	Sun StorEdge PCI Single Fibre Channel Network Adapter	53
FIGURE 6-1	StorEdge Disk Tray Controller Card and Interconnect Cards	57
FIGURE 6-2	Single Disk Tray Configuration	58
FIGURE 6-3	Partner Group Tray Configuration	59

Sun StorEdge T3 Disk Tray Configuration Overview

This chapter describes the Sun StorEdge T3 disk tray and its connection ports and Fibre Channel connections. It also describes basic rules and recommendations for configuring the disk tray, and it lists supported hardware and software platforms. This chapter is organized as follows:

- “Product Description” on page 1
- “Configuration Guidelines and Restrictions” on page 7
- “Configuration Recommendations” on page 7
- “Supported Platforms” on page 8
- “Sun Cluster 2.2 Support” on page 8

Note – For installation and cabling information, refer to the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual*. For software configuration information, refer to the *Sun StorEdge T3 Disk Tray Administrator’s Guide*.

Product Description

The Sun StorEdge T3 disk tray is a high-performance storage device that provides a *redundant array of independent disks* (RAID) architecture for modular, scalable data storage. The disk tray contains the controller card and nine disk drives with Fibre Channel connectivity to the data host.

The Sun StorEdge T3 disk tray includes redundant, *hot-swappable* components and automatic notification of failed components. The disk tray can be used either as a standalone storage unit or as a building block, interconnected with other Sun

StorEdge T3 disk trays are configured in various ways to provide a storage solution optimized to the host application. The disk tray can be placed on a table top or rackmounted in a server cabinet or expansion cabinet.

The Sun StorEdge T3 disk tray *expansion unit* contains the same components as the Sun StorEdge T3 disk tray *controller unit* except for the internal RAID controller. When connected to a controller unit, the expansion unit enables you to increase your storage capacity.

See Appendix B in the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual* for an illustrated breakdown of the Sun StorEdge T3 disk tray and its component parts.

Connection Ports

The Sun StorEdge T3 disk tray has the following ports:

- On the controller card:
 - One *Fibre Channel-Arbitrated Loop* (FC-AL) port, which provides data path connectivity to the application host system.
 - One 10BASE-T Ethernet host interface port (RJ-45).

This port provides the interface between the controller card and the management host system. An unshielded twisted-pair Ethernet cable connects the controller server to the site's network hub. This interface enables the administration and management of the disk tray via the Sun StorEdge Component Manager software or the *command-line interface* (CLI).

- One RS-232 serial port (RJ-11-6).

This serial port is reserved for diagnostic procedures that can only be performed by qualified service personnel.

FIGURE 1-1 shows the controller card and its ports.

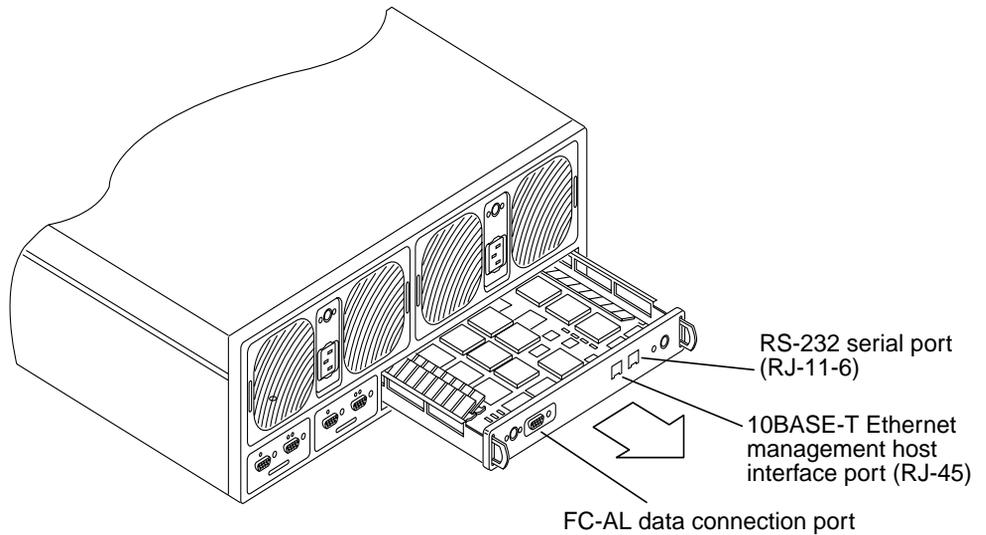


FIGURE 1-1 Controller Card and Ports

- On the *interconnect cards*:
 - Two interconnect ports: one input and one output, for interconnecting multiple Sun StorEdge T3 disk trays.

The interconnect card provides switch and failover capabilities, as well as an environmental monitor for the disk tray. Each disk tray contains two interconnect cards for redundancy (thus providing a total of four interconnect ports).

FIGURE 1-2 shows a Sun StorEdge T3 disk tray interconnect card.

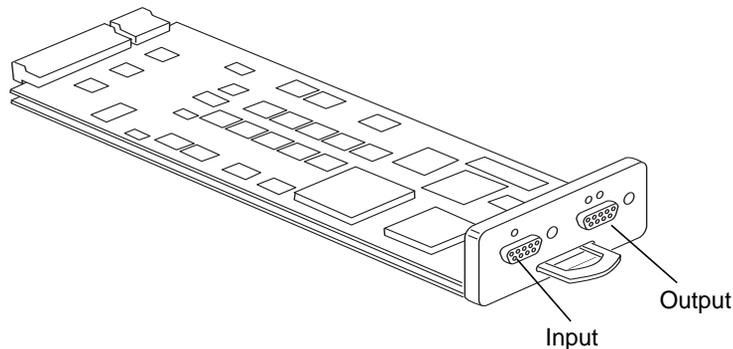


FIGURE 1-2 Interconnect Card and Ports

FC-AL Configurations

Each disk tray uses Fibre Channel-Arbitrated Loop (FC-AL) connections to connect to the application host. An FC-AL connection is a 100 Mbyte/second serial channel that allows multiple devices, such as disk drives and controllers, to be connected.

Partner Group Configuration

In a *partner group* configuration, two Sun StorEdge T3 disk trays are paired using interconnect cables. The partner group is connected to one application host through two FC-AL cables. The partner group is connected to a management host through two 10BASE-T Ethernet cables. The partner group provides *RAS*, plus failover mechanisms for controller pairs, using mirrored caches, and redundant FC-AL host channels for continuous data availability and redundant Ethernet host connections to the management host.

FIGURE 1-3 shows a Sun StorEdge T3 disk tray partner group, which includes two controller units configured as a *master controller unit* and an *alternate master controller unit*.

Note – Partner groups are *not* supported in Sun Cluster 2.2.

Note – If you plan to configure a partner group, install and set up the standalone units before you power up or configure the pair. For information about configuring partner groups, see Chapter 3 in this guide.

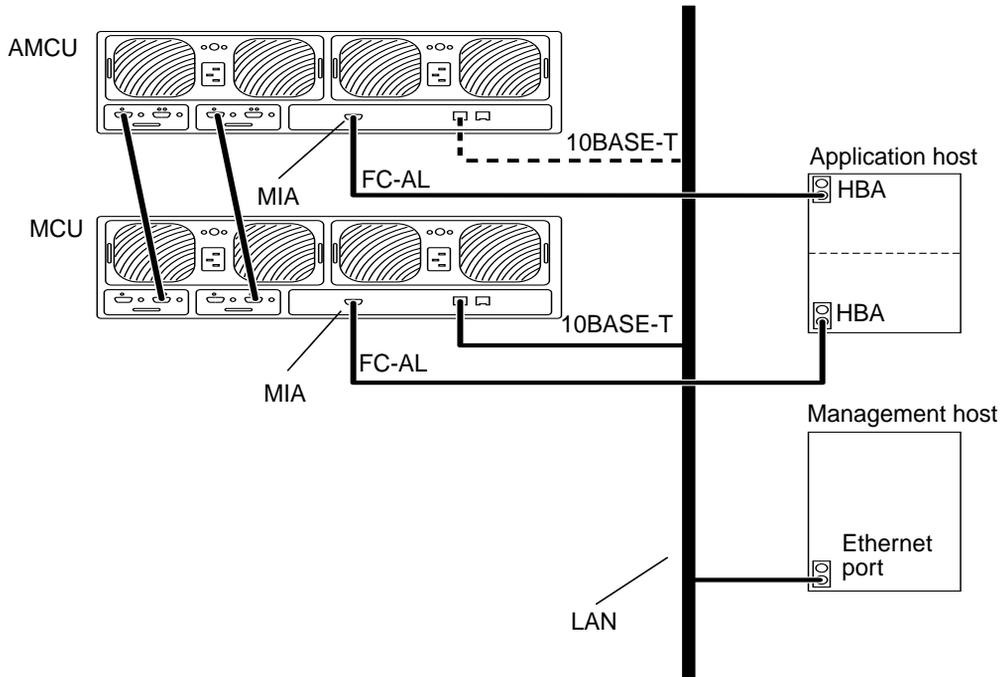


FIGURE 1-3 Sun StorEdge T3 Disk Tray Partner Group Configuration

Note – You must insert the *media interface adapter* (MIA) into the FC-AL port prior to connecting the FC-AL cable to the host bus adapter (HBA) on the application host. The 10BASE-T Ethernet connection is on the public or separate network and requires an IP address. The dashed line indicates that an IP address is not required on the alternate master. This controller will take over the IP address of the master controller unit if a failover occurs.

Standalone Configuration

In a standalone configuration, one controller unit is connected to an application host through an FC-AL cable. The controller unit is connected to a management host through an Ethernet cable. While this unit is fully populated with redundant hot-swappable components and nine disk drives, it is not a recommended configuration for high availability. In the standalone configuration, the controller unit and any device in the host-to-T3 data path (hubs, MIA, HBA, host-to-T3 data cable, etc.), is a single point of failure.

FIGURE 1-4 shows the standalone disk tray configuration with one controller.

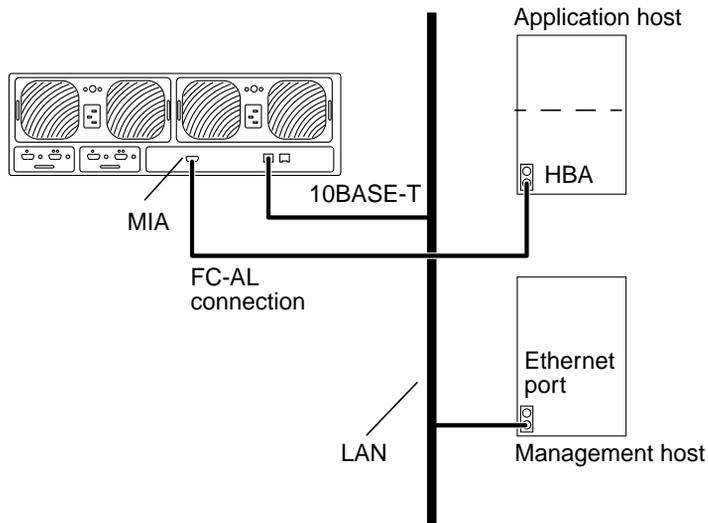


FIGURE 1-4 Sun StorEdge T3 Disk Tray Standalone Configuration

Configuration Guidelines and Restrictions

- Standalone controller unit:

The media access control (MAC) address is required to assign an IP address to the controller unit. The MAC address uniquely identifies each node of a network. The MAC address is available on the pull-out tab on the front left side of the Sun StorEdge T3 disk tray (behind the front cover).

- Partner-group configuration:

- You cannot connect a partner group to more than one host.
- You cannot use a daisy-chain configuration to link more than two controller units (in a partner group).



Caution – When configuring a partner group, be sure to use the MAC address of the master controller unit.

Note – Refer to the *Configuration Rules for Mission Critical Storage* document, also known as the “Big Rules” document, for further configuration rules.

Configuration Recommendations

- Use partner groups (partner groups are *not* supported in Sun Cluster 2.2).
- Use host based software such as VERITAS Volume Manager (VxVM) 3.0.4 or Sun Enterprise™ Server Alternate Pathing (AP) 2.3.1.
- Connect redundant paths to separate host adapters, I/O cards, and system buses.
- Configure active paths over separate system buses to maximize bandwidth.



Caution – The Sun StorEdge T3 disk tray and its global parameters must be tailored to match the I/O workload for optimum performance. Within a partner group, both units will share the same *volume* configuration, block size, and cache mode. That is, all cache parameter settings are common to both units within a partner group.

Supported Platforms

The Sun StorEdge T3 disk tray supports the following host platforms:

- Sun Ultra™ 60 and Ultra 80 workstations
- Sun Enterprise 250, 450, 3x00, 4x00, 5x00, 6x00 servers
- Sun Enterprise 10000 server

The Sun StorEdge T3 disk tray supports the following software versions:

- Solaris 2.6, Solaris 7, and Solaris 8
- VERITAS Volume Manager 3.0.4 and up with Dynamic Multi-Pathing (DMP)
- Sun Enterprise Server Alternate Pathing (AP) 2.3.1

Sun Cluster 2.2 Support

Sun Cluster 2.2 is now supported with the following restrictions:

- 1.16A level controller firmware is *required* on each Sun StorEdge T3 disk tray.
- Partner groups are *not* supported.
- Switches are *not* supported.
- Hubs *must* be used.
- The Sun StorEdge SBus FC-100 (SOC+) HBA and the onboard SOC+ interface in Sun Fire™ systems are supported.
- On Sun Enterprise 6x00/5x00/4x00/3x00 systems, a maximum of 64 Sun StorEdge T3 disk trays are supported per cluster.
- On Sun Enterprise 10000 systems, a maximum of 256 Sun StorEdge T3 disk trays are supported per cluster.
- To ensure full redundancy, host-based mirroring software such as VERITAS Volume Manager 3.0.4, Solstice DiskSuite™ (SDS) 4.2, or SDS 4.2.1 *must* be used.
- Solaris 2.6 and Solaris 8 are the only supported operating systems.

Configuring Global Parameters

When the Sun StorEdge T3 disk tray is shipped, its global parameters are set to default values. This chapter describes how to reconfigure your disk tray by changing these default values for the following parameters:

- “Cache” on page 9
- “Logical Volumes” on page 13
- “Using RAID Levels to Configure Redundancy” on page 17



Caution – If you are planning to configure a partner group from new units from the factory, be sure to install and set up the units as a partner group before you power on, and change any parameters or create/change any logical volumes. Refer to the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual* for more information.

Note – For more information on changing global parameters for the Sun StorEdge T3 disk tray, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Cache

Each controller unit has 256 Mbytes of data cache. Writing to cache improves write performance by staging data in cache, assembling the data into data stripes, and then destaging the data from cache to disk, when appropriate. This method frees the data host for other operations while cache data is being destaged, and it eliminates the read-modify-write delays seen in non-cache systems. *Read cache* improves performance by determining which data will be requested for the next read operation and prestaging this data into cache. RAID 5 performance is also improved by coalescing writes.

Configuring Cache for Performance and Redundancy

Cache mode can be set to the following values:

- **Auto.** The cache mode is determined as either write-behind or write-through, based on the I/O profile. If the disk tray has full redundancy available, then caching operates in write-behind mode. If any disk tray component is non-redundant, the caching mode is set to write-through. Read caching is always performed. Auto caching mode provides the best performance while retaining full redundancy protection.

Auto is the default cache mode for the Sun StorEdge T3 disk tray.

- **Write-behind.** All read and write operations are written to cache. An algorithm determines when the data is destaged or moved from cache to disk. Write-behind cache improves performance, because a write to a high-speed cache is faster than a write to a normal disk.

Use write-behind cache mode with a standalone disk tray configuration when you want to force write-behind caching to be used.

- **Write-through.** This cache mode forces write-through caching to be used. In write-through cache mode, data is written through cache in a serial manner and is then written to the disk. Write-through caching does not improve write performance. However, if a subsequent read operation needs the same data, the read performance is improved, because the data is already in cache.
- **None.** No reads or writes are cached.

Note – For full redundancy, set the cache mode and the mirror variable to “auto” for a partner group. This ensures that the cache is mirrored between controllers and that write-behind cache mode is in effect. If a failure occurs, the data is synchronized to disk, and then write-through mode takes effect. Once the problem has been corrected and all internal components are again optimal, the system will revert to operating in write-behind cache mode.

Configuring Data Block Size

The data block size is the amount of data written to each drive when striping data across drives. (The block size is also known as the stripe unit size.) The block size can be changed only when there are no volumes defined. The block size can be configured as 16 Kbyte, 32 Kbyte, or 64 Kbyte. The default block size is 64 Kbyte.

A cache segment is the amount of data being read into cache. A cache segment is 1/8 of a data block. Therefore, cache segments can be 2 Kbyte, 4 Kbyte or 8 Kbyte. Because the default block size is 64 Kbyte, the default cache segment size is 8 Kbyte.

Note – The disk tray data block size is independent of I/O block size. Alignment of the two is not recommended.

Selecting a Data Block Size

If the I/O initiated from the host is 4 Kbyte, a data block size of 64 Kbyte would force 8 Kbyte of internal disk I/O, wasting 4 Kbyte of the cache segment. Therefore, it would be best to configure 32 Kbyte block sizes, causing 4 Kbyte physical I/O from the disk. If sequential activity occurs, full block writes (32 Kbyte) will take place. For 8 Kbyte I/O or greater from the host, use 64 Kbyte blocks.

Applications benefit from the following data block or stripe unit sizes:

- 16-Kbyte data block size
 - Online Transaction Processing (OLTP)
 - Internet service provider (ISP)
 - Enterprise Resource Planning (ERP)
- 32-Kbyte data block size
 - NFSTM file system, version 2
 - Attribute-intensive NFS file system, version 3

- 64-Kbyte data block size
 - Data-intensive NFS file system, version 3
 - Decision Support Systems (DSS)
 - Data Warehouse (DW)
 - High Performance Computing (HPC)

Note – The data block size must be configured before any logical volumes are created on the units. Remember, this block size is used for every logical volume created on the unit. Therefore it is important to have similar application data configured per unit.

Note – For more detailed information about configuring data block or stripe unit size, refer to the “Configuration Rules for Mission Critical Storage” document (also known as the “Big Rules” document).

Data block size is universal throughout a partner group. Therefore, you cannot change it after you have created a volume. To change the data block size, you must first delete the volume(s), change the data block size, and then create new volume(s).



Caution – Unless you back up and restore the data on these volumes, it will be lost.

Enabling Mirrored Cache

By enabling mirrored cache, you can safeguard cached data if a controller fails.

Note – Mirrored cache is possible only in a redundant partner group configuration.

Configuring Cache Allocation

Cache is allocated based on the read/write mix and it is dynamically adjusted by the controller firmware, based on the I/O profile of the application. If the application profile is configured for a 100% read environment, then 100% of the cache is used for reads. If the application profile has a high number of writes, then the upper limit for writes is set to 80%.

Logical Volumes

Also called a *logical unit number* (LUN), a logical volume is one or more disk drives that are grouped together to form a single unit. Each logical volume is represented to the host as a logical unit number. Using the `format` utility on the application host, you can view the logical volumes presented by the disk tray. You can use this disk space as you would any physical disk, for example, to perform the following operations:

- Install a file system
- Use the device as a raw device (without any file system structure)
- Partition the device

Note – Individual physical disk drives are not visible from the application host. Refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide* for more information on creating logical volumes.

Guidelines for Configuring Logical Volumes

Use the following guidelines when configuring logical volumes:

- The Sun StorEdge T3 disk tray's native volume management can support a maximum of two volumes per disk tray unit.
- The minimum number of drives is based on the RAID level, as follows:
 - RAID 0 and RAID 1 require a minimum of two drives.
 - RAID 5 requires a minimum of three drives.
- Drive #9 can be designated as a *hot spare*. If designated, drive #9 will be the hot spare for all volumes in the disk tray.
- A partial drive configuration is not allowed.
- Volumes cannot span disk tray units.

Consider the following questions when configuring logical volumes:

- How many logical volumes do you need? (1 or 2)
- What RAID level do you require?
- Do you need a hot spare?

Determining How Many Logical Volumes You Need

You can configure a volume into seven partitions (also known as *slices*) using the `format(1M)` utility. Alternatively, you can configure virtually a large number of partitions (also known as *subdisks*) using VERITAS Volume Manager. Therefore, Sun StorEdge T3 disk trays are best configured as one large volume.

Applications benefit from the following logical volume or LUN configurations:

- Two LUNs per disk tray
 - OLTP
 - ISP
 - ERP
 - NFS, version 2
 - Attribute-intensive NFS, version 3
- One LUN per disk tray
 - Data-intensive NFS, version 3
 - DSS
 - DW
 - HPC

Note – If you are creating new volumes or changing the volume configuration, you must first manually rewrite the label of the previous volume using the `autoconfigure` option of the `format(1M)` UNIX host command. For more information on this procedure, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.



Caution – Removing and reconfiguring the volume will destroy all data previously stored there.

Determining Which RAID Level You Require

For a new installation of a Sun StorEdge T3 disk tray, the default configuration is 8+1 RAID 5, without a hot spare.

In general, RAID 5 is efficiently managed by the RAID controller hardware. This efficiency is apparent when compared to RAID 5 software solutions such as VERITAS Volume Manager.

The following applications benefit most from the RAID controller hardware of the Sun StorEdge T3 disk tray:

- Data-intensive NFS file system, version 3
- DSS
- DW
- HPC

Note – For more information about RAID levels, see “Using RAID Levels to Configure Redundancy” later in this chapter. For more detailed information about configuring RAID storage, refer to the “Configuration Rules for Mission Critical Storage” document (also known as the “Big Rules” document).

Determining Whether You Need a Hot Spare

If you choose to include a hot-spare disk drive in your configuration, you must specify it when you create the first volume in the disk tray. If you want to add a hot spare at a later date, you must remove the existing volume(s) and recreate the configuration.

Note – Only one hot spare is allowed per disk tray and it is only usable for the disk tray in which it is configured. The hot spare must be configured as drive #9.

Drive #9 will be the hot spare in the unit. So, for example, should a drive failure occur on drive #7, drive #9 is synchronized automatically with the entire LUN to reflect the data on drive #7. Once the failed drive (#7) is replaced, the controller unit will automatically copy the data from drive #9 to the new drive, and drive #9 will become a hot spare again.

Tip – Although they are not required, hot spares are always recommended for mission-critical configurations because they allow the controller unit to reconstruct the data from the RAID group and only take a performance hit while the reconstruction is taking place. If a hot spare is not used, the controller unit remains in write-through cache mode until the failed drive is replaced and reconstruction is complete (which could take an extended period of time). During this time, the disk tray is operating in degraded mode.

If there is no hot spare, the reconstruction of the data will begin when the failed drive is replaced, provided RAID 1 or RAID 5 is used.

For suggested hot-spare configurations for the Sun StorEdge T3 disk tray, based on RAID level (RAID 1 versus RAID 5) and application size (small/medium versus medium/large) refer to the “Configuration Rules for Mission Critical Storage” Creating and Labeling a Logical Volume.

Creating and Labeling a Logical Volume

You must set the RAID level and the hot-spare disk when creating a logical volume. For the Solaris operating system to recognize a volume, it must be labeled with the `format` or `fmthard` command.



Caution – Removing and reconfiguring a logical volume will destroy all data previously stored there.

Setting the LUN Reconstruction Rate

Note – When a failed drive is disabled, the volume is operating without further redundancy protection, so the failed drive needs to be replaced as soon as possible.

If the volume has a hot spare configured and that drive is available, the data on the disabled drive is reconstructed on the hot-spare drive. When this operation is complete, the volume is operating with full redundancy protection, so another drive in the volume may fail without loss of data.

After a drive has been replaced, the original data is automatically reconstructed on the new drive. If no hot spare was used, the data is regenerated using the RAID redundancy data in the volume. If the failed drive data has been reconstructed onto a hot spare, once the reconstruction has completed, a copy-back operation begins where the hot spare data is copied to the newly replaced drive.

You can also configure the rate at which data is reconstructed, so as not to interfere with application performance. Reconstruction rate values are low, medium, and high as follows:

- low is the slowest and has the lowest impact on performance
- medium is the default
- high is the fastest and has the highest impact on performance

Note – Reconstruction rates can be changed while a reconstruction operation is in process. However, the changes don't take effect until the current reconstruction has completed.

Using RAID Levels to Configure Redundancy

The RAID level determines how the controller reads and writes data and *parity* on the drives. The Sun StorEdge T3 disk tray can be configured with RAID level 0, RAID level 1 (1+0) or RAID level 5. The factory-configured LUN is a RAID 5 LUN.

Note – The default RAID level (5) can result in very large volumes; for example, 128 Gbytes in a configuration of single 7+1 RAID 5 LUN plus hot spare, with 18 Gbyte drives. Some applications cannot use such large volumes effectively. The following two solutions can be used separately or in combination:

- First, use the partitioning utility available on the data host's operating system. In the Solaris environment, use the `format` utility, which can create up to seven distinct partitions per volume. Note that in the case of the configuration described above, if each partition is equal in size, this will result in 18-Gbyte partitions, which still may be too large to be used efficiently by legacy applications.
- Second, you can use third-party software on the host system to create as many partitions as desired from a given volume. In the Solaris environment, you can use VERITAS Volume Manager or Solaris Logical Volume Management (SLVM) formerly known as Solstice DiskSuite (SDS) for this purpose.

Note – For information on using the `format` utility, refer to the `format (1M)` man page. For more information on third-party software or VERITAS Volume Manager, refer to the documentation for that product.

RAID 0

Data blocks in a RAID 0 volume are striped across all the drives in the volume in order. There is no parity data, so RAID 0 uses the full capacity of the drives. There is, however, no redundancy. If a single drive fails, all data on the volume is lost.

RAID 1

Each data block in a RAID 1 volume is mirrored on two drives. If one of the mirrored pair fails, the data from the other drive is used. Because the data is mirrored in a RAID 1 configuration, the volume has only half the capacity of the assigned drives. For example, if you create a 4-drive RAID 1+0 volume with 18-Gbyte drives, the resulting data capacity is $4 \times 18 / 2 = 36$ Gbytes.

RAID 5

In a RAID 5 configuration, data is striped across the drives in the volumes in segments, with parity information being striped across the drives, as well. Because of this parity, if a single drive fails, data can be recovered from the remaining drives. Two drive failures cause all data to be lost. A RAID 5 volume has the data capacity of all the drives in the logical unit, less one. For example, a 5-drive RAID 5 volume with 18-Gbyte drives has a capacity of $(5 - 1) \times 18 = 72$ Gbytes.

Configuring RAID Levels

The Sun StorEdge T3 disk tray is preconfigured at the factory with a single LUN, RAID level 5 redundancy and no hot spare. Once a volume has been configured, you cannot reconfigure it to change its size, RAID level, or hot spare configuration. You must first delete the volume and create a new one with the configuration values you want.

Configuring Partner Groups

Sun StorEdge T3 disk trays can be interconnected in *partner groups* to form a redundant and larger storage system.

Note – Partner groups are *not* supported in Sun Cluster 2.2.

This chapter describes how to configure Sun StorEdge T3 disk tray partner groups, and it includes the following sections:

- “Understanding Partner Groups” on page 19
- “Creating Partner Groups” on page 22

Understanding Partner Groups

A partner group consists of a “master” controller unit and a “slave” or alternate master controller unit. These controllers are connected with interconnect cables, as shown in FIGURE 3-1.

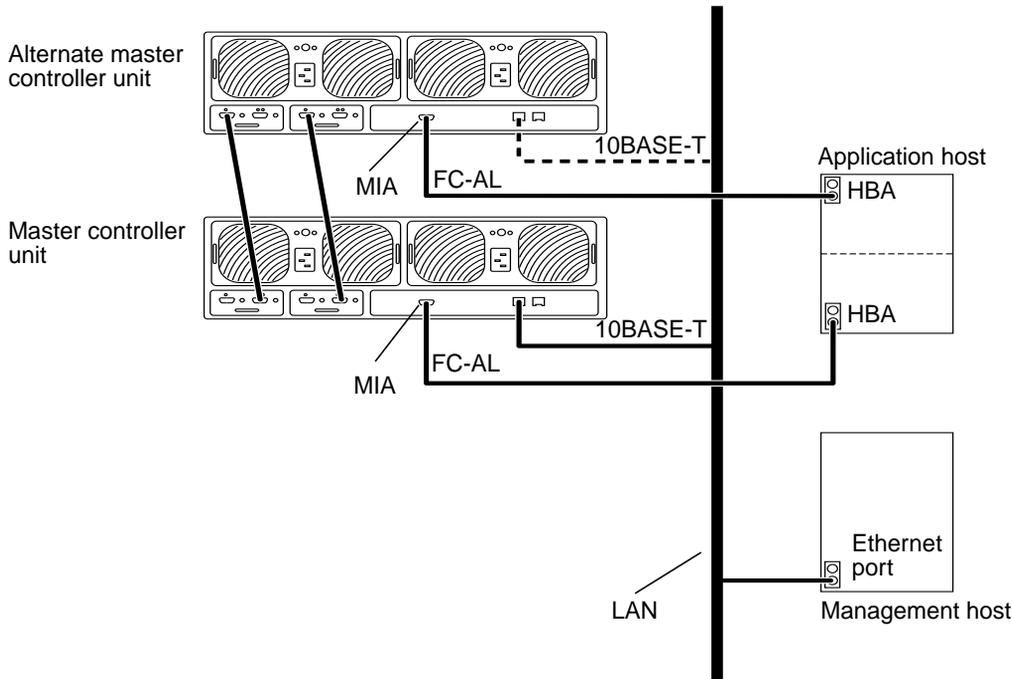


FIGURE 3-1 Sun StorEdge T3 Disk Tray Partner Group

Note – You must insert the *media interface adapter* (MIA) into the FC-AL port prior to connecting the FC-AL cable to the host bus adapter (HBA) on the application host. The 10BASE-T Ethernet connection is on the public or private network and requires an IP address. The dashed line indicates that an IP address is not required on the alternate master unit. This controller will take over the IP address of the master controller unit if a failover occurs.

When two units are connected together, they form a redundant partner group. This group provides controller redundancy. Because the controller is a single point of failure in a standalone configuration, this redundancy allows an application host to access data even if a controller fails. This configuration offers multipath and LUN failover features.

The partner group connection also allows for a single point of control. The bottom unit will assume the role of the master, and from its Ethernet connections, it will be used to monitor and administer the unit installed above it.

The master unit will set the global variables within this storage system, including cache block size, cache mode, and cache mirroring.

Note – For information about setting or changing these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Any controller unit will boot from the master unit's drives. All configuration data, including `syslog` information, is located on the master unit's drives.

How Partner Groups Work

If the master unit fails and the “heartbeat” between it and the alternate master stops, this failure causes a controller failover, where the alternate master assumes the role of the master unit. The new master (formerly the alternate master) takes the IP address and the MAC address from the old master and begins to function as the administrator of the storage system. It will also be able to access the former master unit's drives. The former master unit's drives will still be used to store `syslog` information, system configuration information, and bootcode. Should it become necessary to reboot the storage system while the master unit is inactive, the alternate master will use the former master unit's drives to boot from.

Note – After the failed master controller is back online, it remains the alternate master controller and, as a result, the original configuration has been modified from its original state.

In a redundant partner group configuration, the units can be set to do a path failover operation. Normally the volumes or LUNs that are controlled by one unit are not accessible to the controller of the other. The units can be set so that if a failure in one controller occurs, the remaining one will accept I/O for the devices that were running on the failed controller.

Note – In order for a feature such as VERITAS DMP to access a LUN through both controllers in a redundant partner group, the `mp_support` parameter must be set to `rw` to enable this feature.

Creating Partner Groups

Partner groups can be created in two ways:

- From new units
- From existing standalone units

▼ To Create a Partner Group From New Units

To set up a partner group from two new Sun StorEdge T3 disk trays, follow these steps before powering on the disk trays.



Caution – If there is any customer data that needs to be preserved, ensure that the data is backed up to tape before starting this procedure.

Note – This procedure assumes that you have already connected the fiber-optic cable, media interface unit, AC power cords, and 10BASE-T interface cables. For specific instructions on connecting these cables, refer to Chapter 2 of the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual*.

1. Use the interconnect cables to interconnect units, as shown in figure FIGURE 3-2.

Make sure you connect the cables to the correct interconnect (loop) card connectors exactly as shown in the FIGURE 3-2. (Port 1 on the alternate master is always connected to port 2 on the master unit.) Tighten the retaining screws.

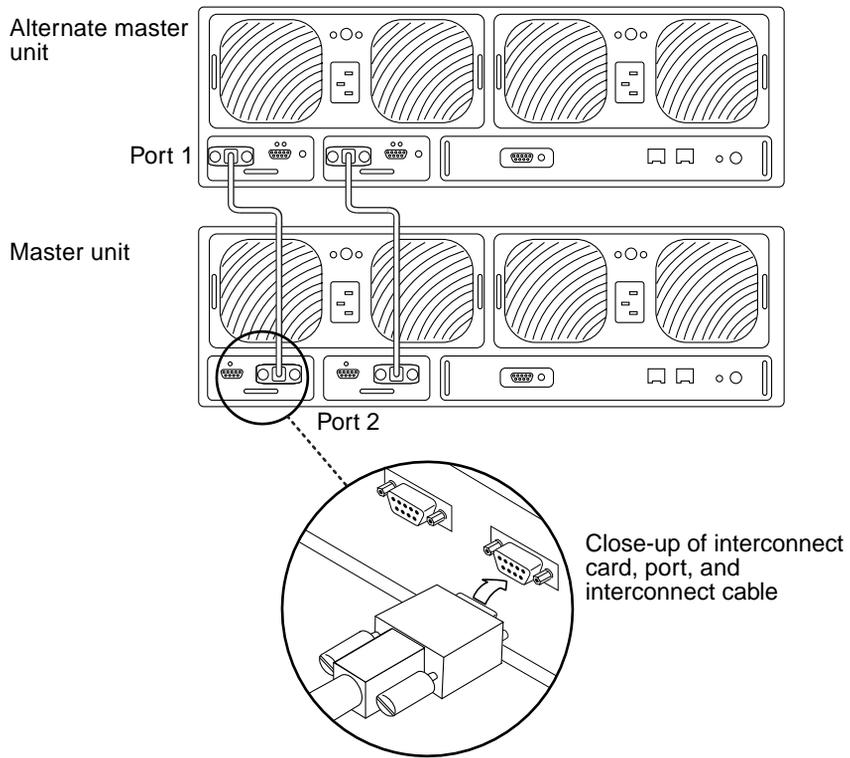


FIGURE 3-2 Connecting the Interconnect Cables

A fully cabled partner group is shown in FIGURE 3-3.

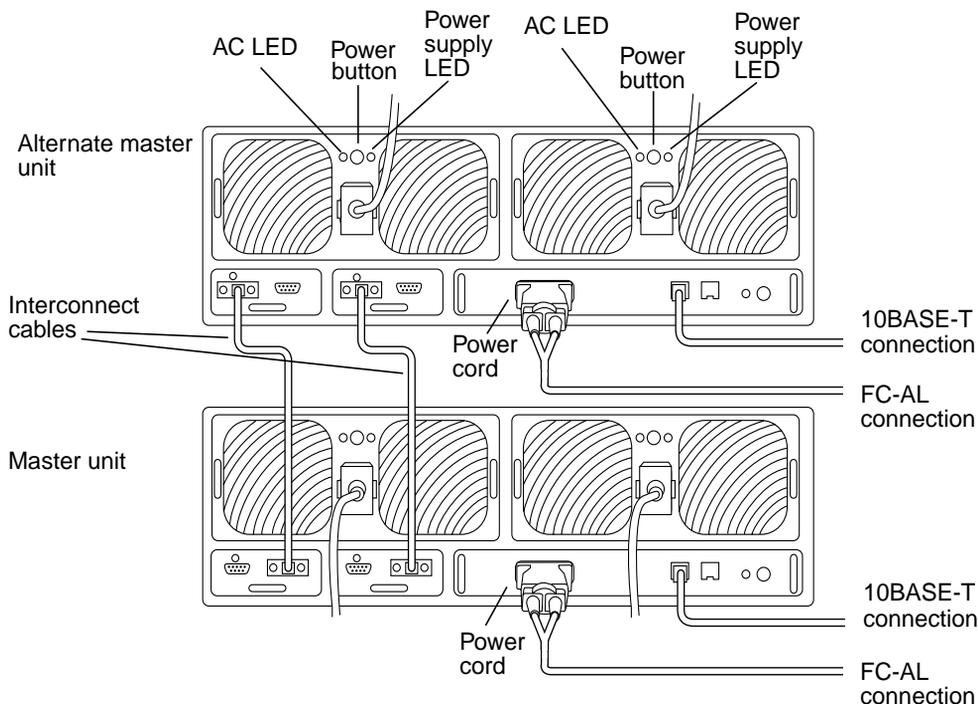


FIGURE 3-3 Fully Cabled Partner Group

2. Power on and verify the hardware configuration.

a. Plug the other end of the power cords from the disk trays into AC outlets.

Use alternate power sources to make sure that *power and cooling unit* features are redundant.

b. Verify that AC power is present on each power and cooling unit.

The AC LED on each power and cooling unit will be solid amber and the fans will turn at low speed.

c. Press the power button on the power and cooling units to power on the disk trays.

The AC and power supply (PS) LEDs on the power and cooling units will be green.

d. Check the LEDs at the front and back of the unit to ensure that all components are receiving power and are functional.

While the drives are spinning up, the LEDs will blink. The disk tray boot time will take up to several minutes, after which all LEDs should be solid green, indicating that the unit is receiving power and there is no drive activity.

For more information on LEDs and how to interpret them, refer to the *Sun StorEdge T3 Disk Tray Installation, Operation, and Service Manual*.

Note – The master unit is defined automatically upon power up. Allow enough time for the alternate master to be booted up by the master.

3. Check that the following global parameters are set, as follows:

- `mp_support = rw`
- `cache = auto`
- `mirror = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

4. Create the data volumes.

Refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

▼ To Create a Partner Group From Existing Standalone Units

Refer to Section 10.1 “Connecting Single Controller Units to Form a Partner Group” of the *Sun StorEdge T3 Disk Tray Field Service Manual* for detailed instructions on performing this procedure.

Configuration Examples

This chapter includes the following sample reference configurations for the Sun StorEdge T3 disk tray:

Direct Host Connection

- “Single Host With One Controller Unit” on page 28
- “Single Host With Two Controller Units Configured as a Partner Group” on page 29
- “Single Host With Four Controller Units Configured as Two Partner Groups” on page 32
- “Single Host With Eight Controller Units Configured as Four Partner Groups” on page 34

Hub Host Connection

- “Single Host With Four Controller Units and Two Hubs Configured as Two Partner Groups” on page 36
- “Single Host With Eight Controller Units and Two Hubs Configured as Four Partner Groups” on page 38
- “Dual Hosts With Four Controller Units and Two Hubs” on page 40
- “Dual Hosts With Eight Controller Units and Two Hubs” on page 42

Switch Host Connection

- “Dual Hosts With Two Controller Units and Two Switches Configured as One Partner Group” on page 45
- “Dual Hosts With Eight Controller Units and Two Switches Configured as Four Partner Groups” on page 47

Although there are many supported configurations, these reference configurations provide the best solution for many installations.

Direct Host Connection

Single Host With One Controller Unit

FIGURE 4-1 shows one application host connected through an FC-AL cable to one Sun StorEdge T3 disk tray controller unit. The 10BASE-T connection connects the controller to a management host via a LAN on a public or separate network, and requires an IP address. You must insert the media interface adapter (MIA) into the FC-AL connection port prior to connecting the FC-AL cable to the HBA on the application host.

Note – This configuration is not recommended for RAS functionality because the controller is a single point of failure.

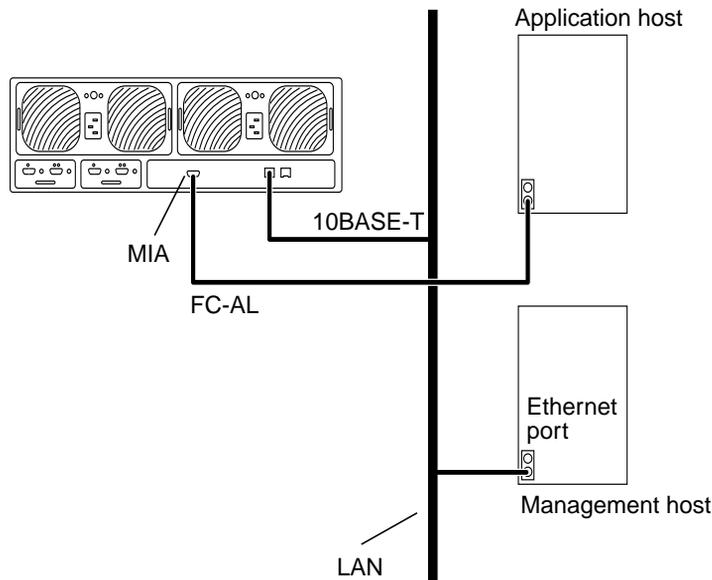


FIGURE 4-1 Single Host Connected to One Controller Unit

Single Host With Two Controller Units Configured as a Partner Group

FIGURE 4-2 shows one application host connected through FC-AL cables to one Sun StorEdge T3 disk tray partner group, which consists of two Sun StorEdge T3 disk trays. The 10BASE-T connection on the master controller is on a public or separate network and requires an IP address for the partner group. In the event of a failover, the alternate master controller will use the master controller's IP address and MAC address.

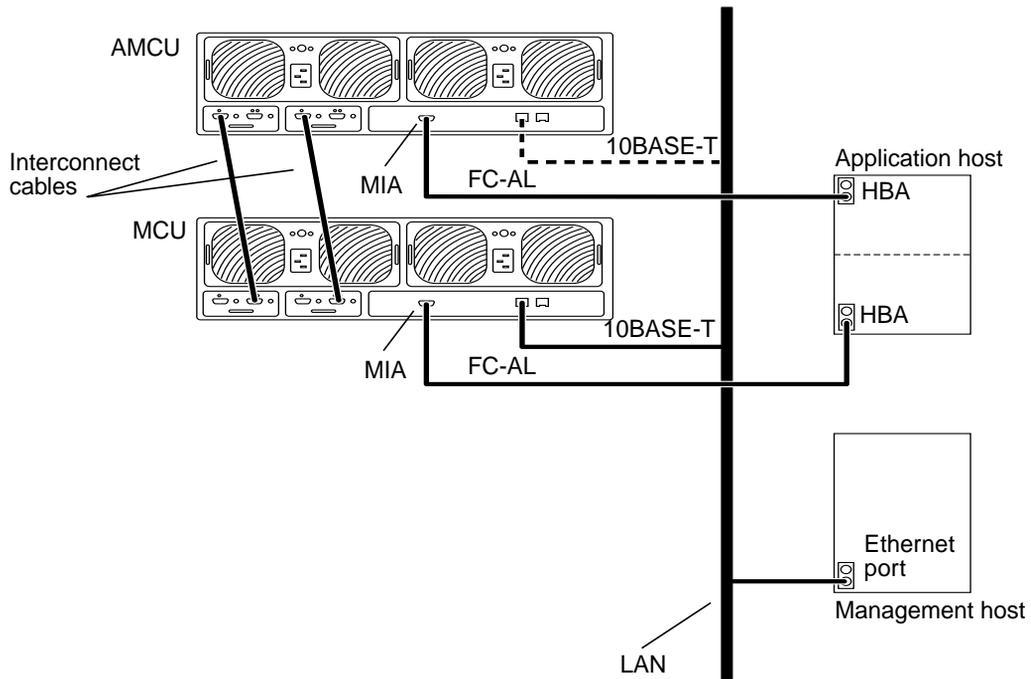


FIGURE 4-2 Single Host With Two Controller Units Configured as a Partner Group

Note – The dashed line indicates that an IP address is not necessary. This controller will take over the IP address and the MAC address of the MCU if a failover occurs.

Note – This configuration is a recommended enterprise configuration for RAS functionality because there is no single point of failure. This configuration supports Dynamic Multi-Pathing (DMP) or Alternate Pathing (AP) for failover only.

The following three global parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Host Multi-Pathing Management Software

While the Sun StorEdge T3 disk tray is a redundant device that automatically reconfigures itself whenever a failure occurs on any of its internal components, a host-based solution is needed for a redundant data path. Dynamic Multi-Pathing (DMP) is provided by VERITAS Volume Manager, version 3.0.4 and above. Alternate Pathing (AP) is provided by Sun Enterprise Server Alternate Pathing (AP) 2.3.1.

During normal operation, I/O moves on the host channel connected to the controller that “owns” the disks. This path is an “active path.” During failover operation, DMP or AP directs all I/O to the alternate channel’s controller. This path is the “failover path.”

When the master controller fails, the alternate master controller becomes the master. When the failed controller is repaired, the new controller immediately boots, goes online and becomes the alternate master controller. The former alternate master controller remains as the current master controller.

Note – DMP, which is provided by VERITAS Volume Manager, or AP must be installed on the application host to achieve a fully redundant configuration.

FIGURE 4-3 shows a DMP or AP failover configuration.

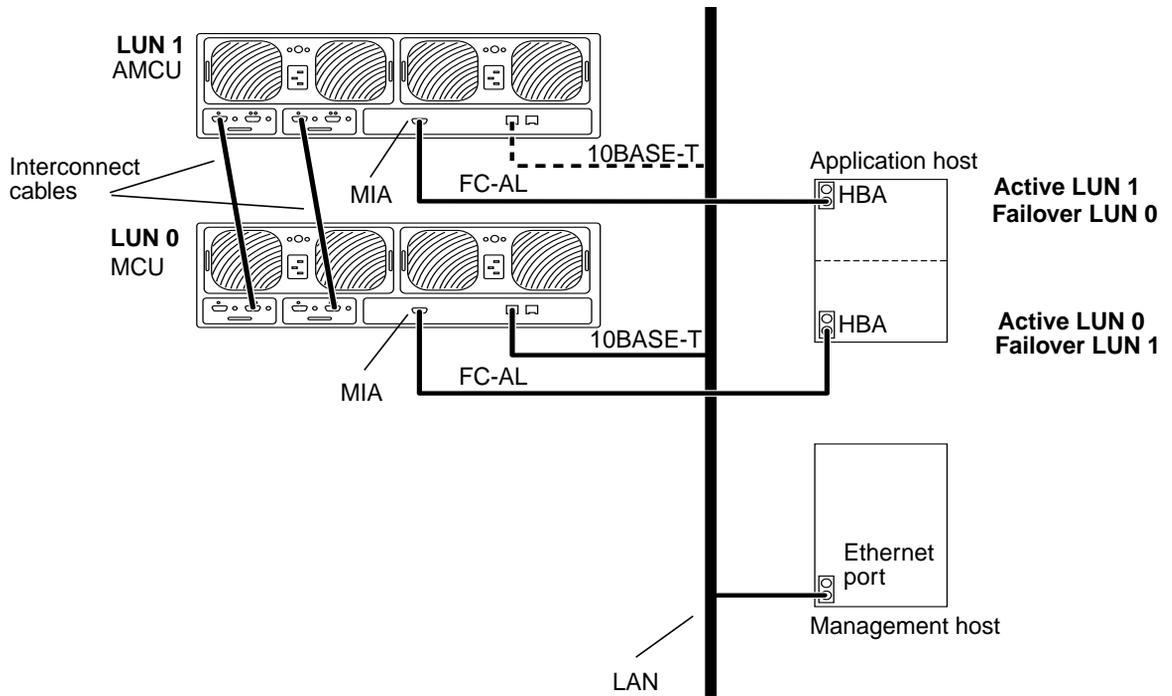


FIGURE 4-3 DMP or AP Failover Configuration

Single Host With Four Controller Units Configured as Two Partner Groups

FIGURE 4-4 shows one application host connected through FC-AL cables to four Sun StorEdge T3 disk trays configured as two separate partner groups. This configuration can be used for capacity and I/O throughput requirements.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

The following three parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Host-based alternate pathing, such as VERITAS DMP or Sun Enterprise Server AP, must also be installed on the host.

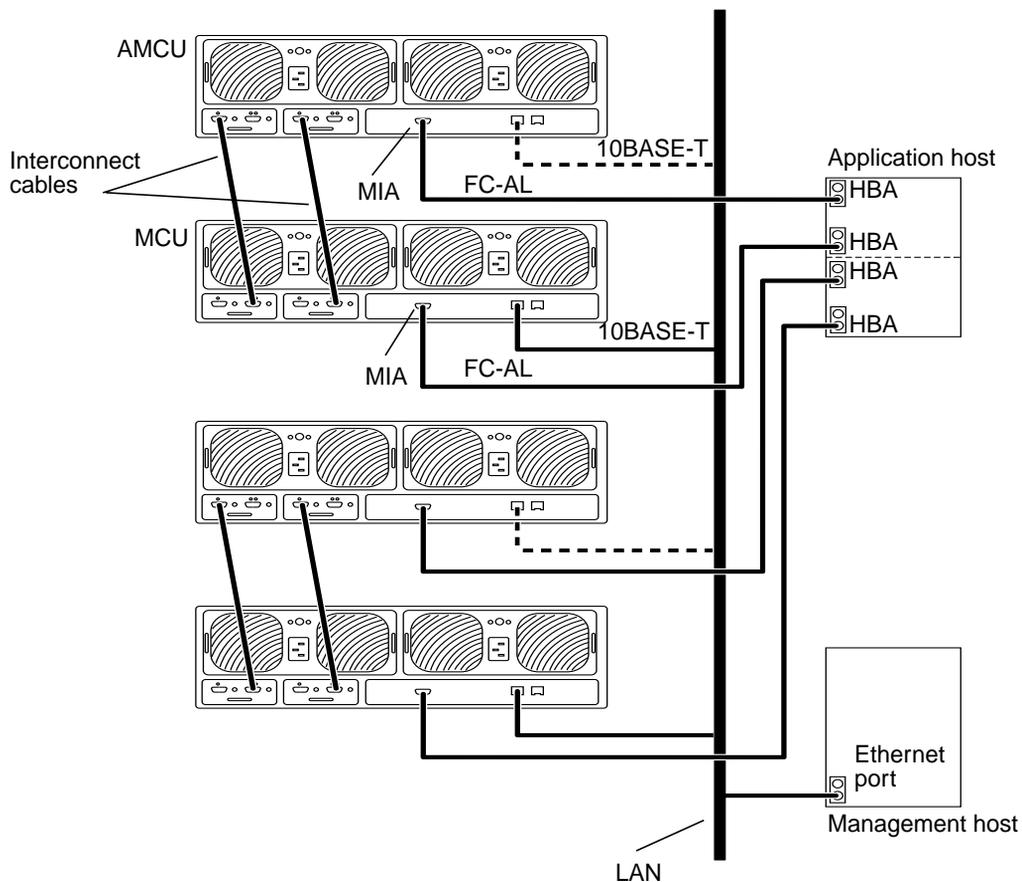


FIGURE 4-4 Single Host With Four Controller Units Configured as Two Partner Groups

Note – You must insert the media interface adapter (MIA) into the FC-AL port prior to connecting the FC-AL cable to the HBA on the application host. The 10BASE-T connection is on a public or separate network and requires an IP address. The dashed line indicates that an IP address is not necessary. This controller will take over the IP address and the MAC address of the master unit if a failover occurs.

Single Host With Eight Controller Units Configured as Four Partner Groups

FIGURE 4-5 shows one application host connected through FC-AL cables to eight Sun StorEdge T3 disk trays, forming four partner groups. This configuration is the maximum allowed in a 72-inch cabinet. This configuration can be used for footprint and I/O throughput.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

The following three parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Host-based alternate pathing, such as VERITAS DMP or Sun Enterprise Server AP, must also be installed on the host.

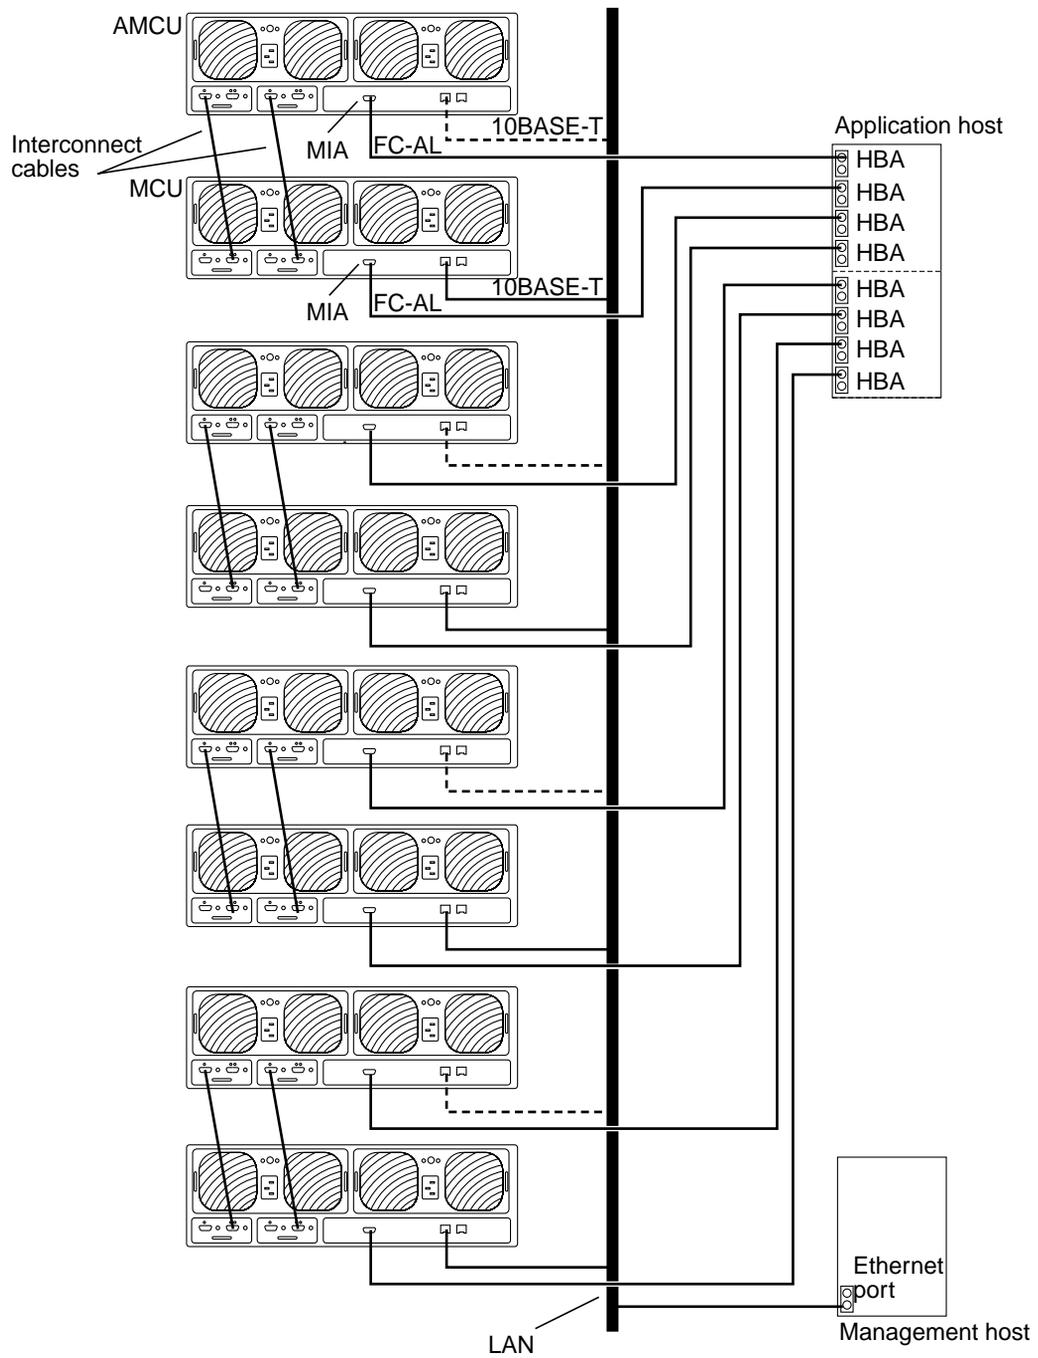


FIGURE 4-5 Single Host With Eight Controller Units Configured as Four Partner Groups

Hub Host Connection

Single Host With Four Controller Units and Two Hubs Configured as Two Partner Groups

FIGURE 4-6 shows one application host connected through FC-AL cables to two hubs and two Sun StorEdge T3 disk tray partner groups, which consist of four Sun StorEdge T3 disk trays. The 10BASE-T connection on the master controller is on a public or separate network and requires an IP address for the partner group. In the event of a failover, the alternate master controller will use the master controller's IP address and *MAC address*.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting a Sun StorEdge T3 disk tray to a hub. A Sun StorEdge T3 disk tray can be connected to any available port on the hub.

Each T3 disk tray needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the disk tray target addresses are set starting with target address 1 for the bottom disk tray and continuing to the top disk tray. Use the `port list` command to verify that all T3 disk trays have a unique target address. Refer to Appendix A.3.12 of the *Sun StorEdge T3 Disk Tray Administrator's Guide* for further details.

The following three parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the Sun StorEdge T3 Disk Tray Administrator's Guide.

Host-based alternate pathing, such as VERITAS DMP or Sun Enterprise Server AP, must also be installed on the host.

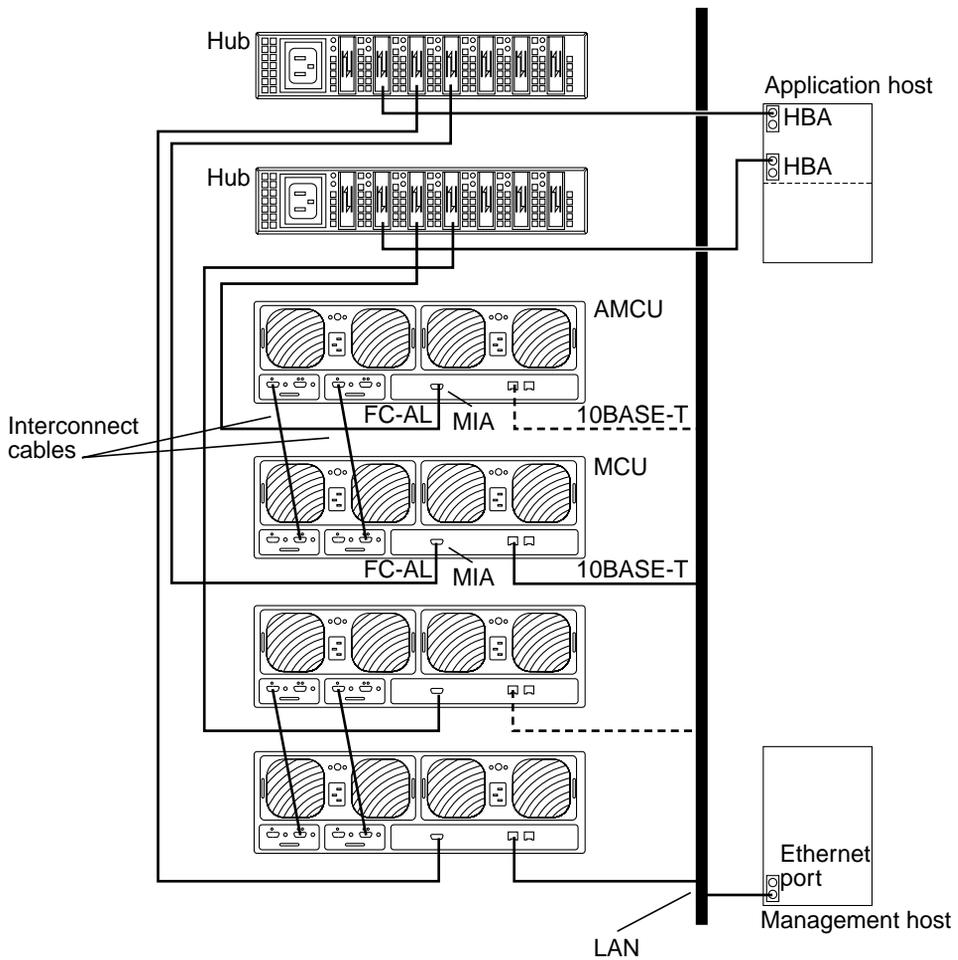


FIGURE 4-6 Single Host With Four Controller Units and Two Hubs Configured as Two Partner Groups

Single Host With Eight Controller Units and Two Hubs Configured as Four Partner Groups

FIGURE 4-7 shows one application host connected through FC-AL cables to two hubs and to eight Sun StorEdge T3 disk trays, forming four partner groups. This configuration is the maximum allowed in a 72-inch cabinet. This configuration can be used for footprint and I/O throughput.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting a Sun StorEdge T3 disk tray to a hub. A Sun StorEdge T3 disk tray can be connected to any available port on the hub.

Each T3 disk tray needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the disk tray target addresses are set starting with target address 1 for the bottom disk tray and continuing to the top disk tray. Use the `port list` command to verify that all T3 disk trays have a unique target address. Refer to Appendix A.3.12 of the *Sun StorEdge T3 Disk Tray Administrator's Guide* for further details.

The following three parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Host-based alternate pathing, such as VERITAS DMP or Sun Enterprise Server AP, must also be installed on the host.

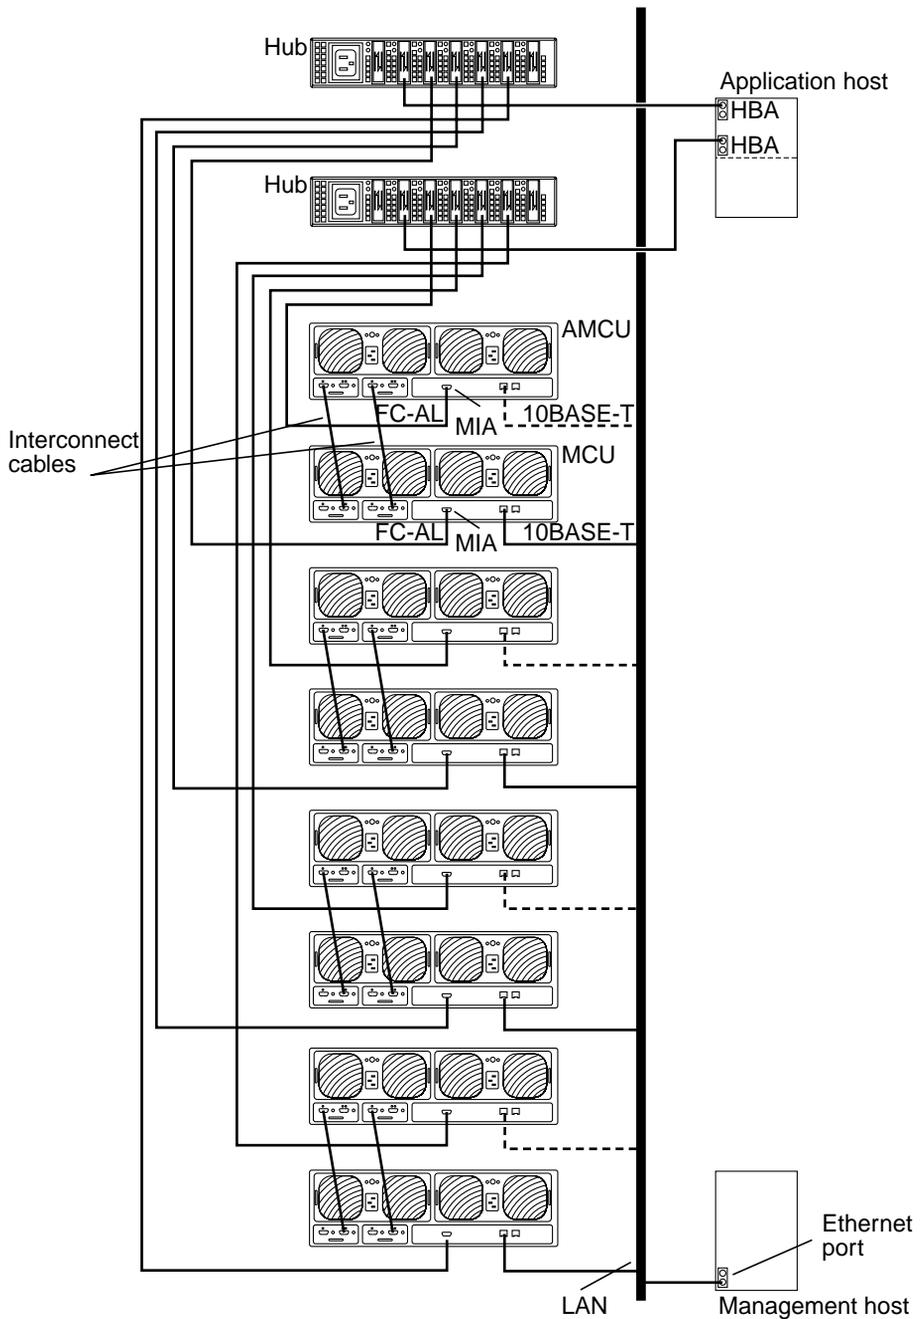


FIGURE 4-7 Single Host With Eight Controller Units and Two Hubs Configured as Four Partner Groups

Dual Hosts With Four Controller Units and Two Hubs

FIGURE 4-8 shows two application hosts connected through FC-AL cables to two hubs and four Sun StorEdge T3 disk trays. This configuration, also known as a Multi-Initiator configuration, can be used for footprint and I/O throughput. The following limitations should be evaluated when proceeding with this configuration

- Multi-Initiator configurations are only supported with single tray T3 configurations (no partner groups).
- Because only single trays are supported in Multi-Initiator configurations, to avoid the risk caused by any T3 disk tray or data path single point of failure, Sun Microsystems, Inc. recommends that host-based mirroring software such as VERITAS Volume Manager or Solaris Logical Volume Manager be used for host-based mirroring of T3 volumes on the disk trays.
- When configuring more than one partner group or single T3 disk tray to share a single FC-AL loop, as with a hub, T3 disk tray target addresses need to be set to unique values.
- With 1.16 level controller firmware, Sun Microsystems, Inc. recommends that RAID 0 T3 volumes *not be used*. This is a release note bug (refer to bug 4374280 listed in the *Sun StorEdge T3 Disk Tray Release Notes*).

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting a Sun StorEdge T3 disk tray to a hub. A Sun StorEdge T3 disk tray can be connected to any available port on the hub.

Each T3 disk tray needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the disk tray target addresses are set starting with target address 1 for the bottom disk tray and continuing to the top disk tray. Use the `port list` command to verify that all T3 disk trays have a unique target address. Refer to Appendix A.3.12 of the *Sun StorEdge T3 Disk Tray Administrator's Guide* for further details.

The following two parameters must be set on the master controller, as follows:

- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

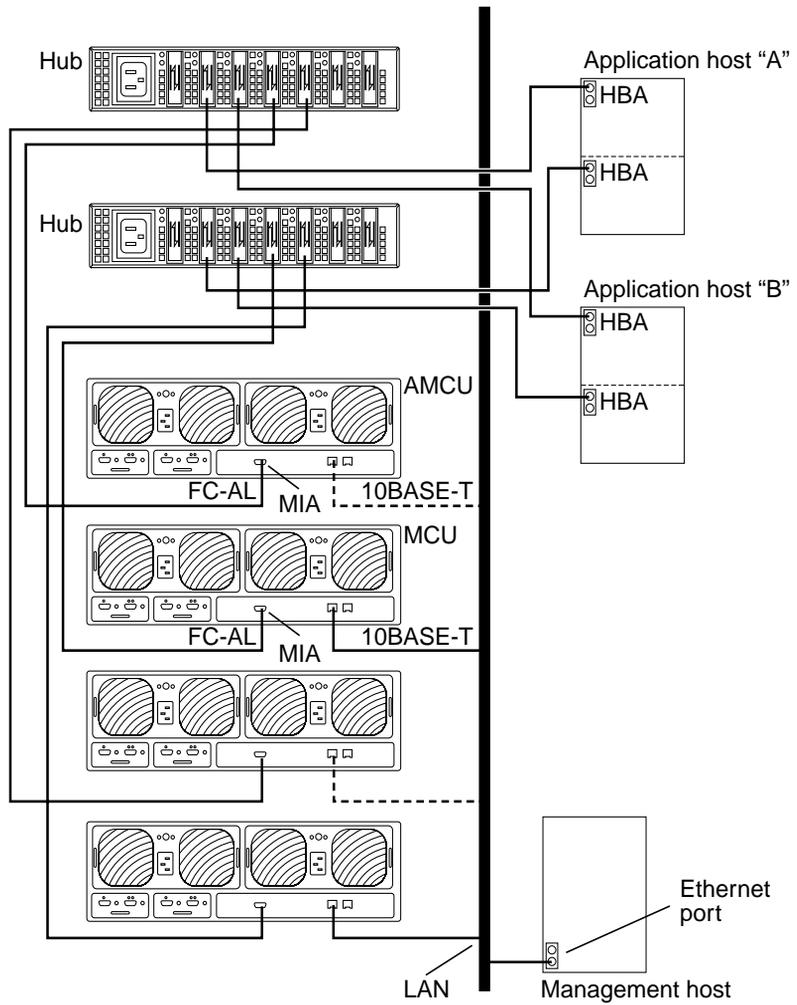


FIGURE 4-8 Dual Hosts With Four Controller Units and Two Hubs

Dual Hosts With Eight Controller Units and Two Hubs

FIGURE 4-9 shows two application hosts connected through FC-AL cables to two hubs and eight Sun StorEdge T3 disk trays. This configuration, also known as a Multi-Initiator configuration, can be used for footprint and I/O throughput. The following limitations should be evaluated when proceeding with this configuration

- Multi-Initiator configurations are only supported with single tray T3 configurations (no partner groups).
- Because only single trays are supported in Multi-Initiator configurations, to avoid the risk caused by any T3 disk tray or data path single point of failure, Sun Microsystems, Inc. recommends that host-based mirroring software such as VERITAS Volume Manager or Solaris Logical Volume Manager be used for host-based mirroring of T3 volumes on the disk trays.

Note – This configuration, running host-based mirroring features from VERITAS Volume Manager or Solaris Logical Volume Manager, represents four trays of data mirrored to the other four trays using host-based mirroring.

- When configuring more than one partner group or single T3 disk tray to share a single FC-AL loop, as with a hub, T3 disk tray target addresses need to be set to unique values.
- With 1.16 level controller firmware, Sun Microsystems, Inc. recommends that RAID 0 T3 volumes *not be used*. This is a release note bug (refer to bug 4374280 listed in the *Sun StorEdge T3 Disk Tray Release Notes*).

This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting a Sun StorEdge T3 disk tray to a hub. A Sun StorEdge T3 disk tray can be connected to any available port on the hub.

Each T3 disk tray needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the disk tray target addresses are set starting with target address 1 for the bottom disk tray and continuing to the top disk tray. Use the `port list` command to verify that all T3 disk trays have a unique target address. Refer to Appendix A.3.12 of the *Sun StorEdge T3 Disk Tray Administrator's Guide* for further details.

The following two parameters must be set on the master controller, as follows:

- cache mode = auto
- cache mirroring = auto

Note – For information on setting these parameters, refer to the Sun StorEdge T3 Disk Tray Administrator's Guide.

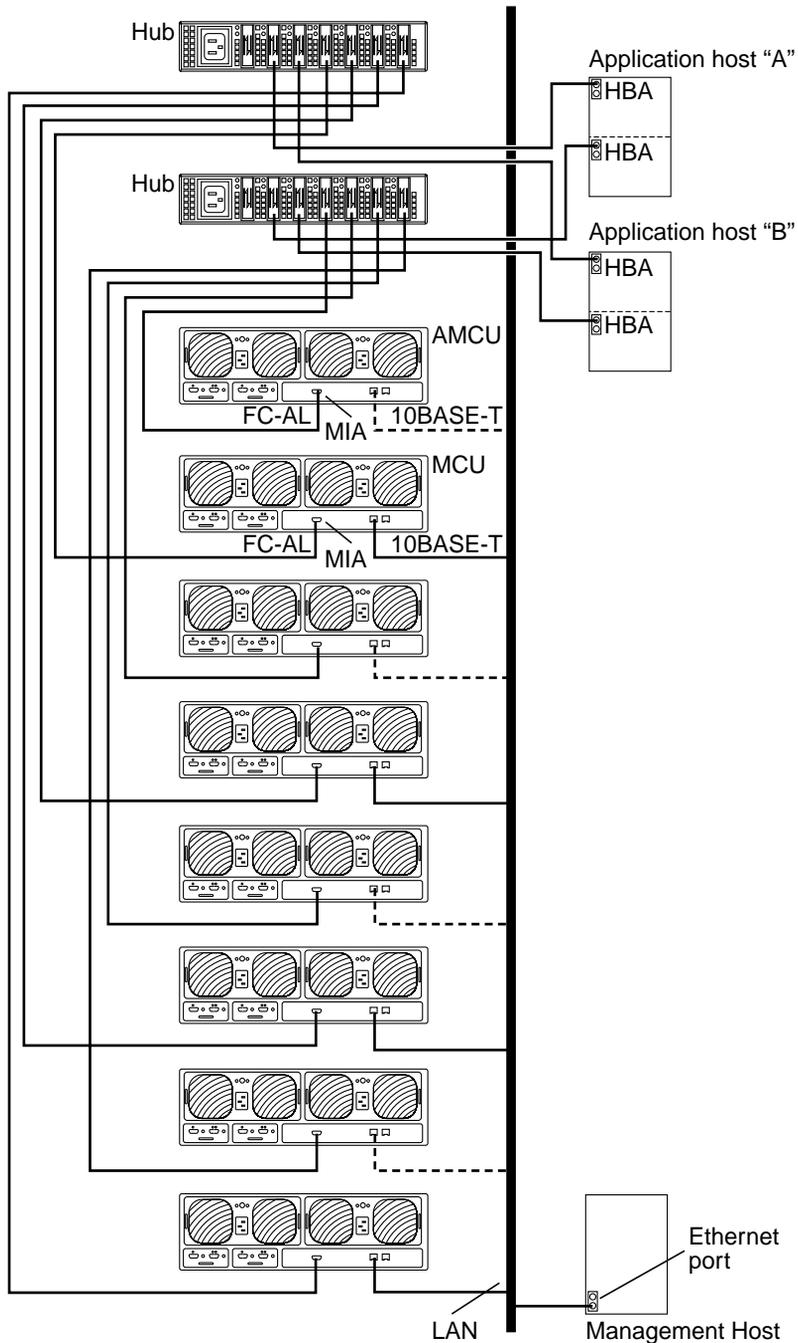


FIGURE 4-9 Dual Hosts With Eight Controller Units and Two Hubs

Switch Host Connection

Dual Hosts With Two Controller Units and Two Switches Configured as One Partner Group

FIGURE 4-10 shows two application hosts connected through FC-AL cables to two switches and one Sun StorEdge T3 disk tray partner group, which consists of two Sun StorEdge T3 disk trays. The 10BASE-T connection on the master controller is on a public or separate network and requires an IP address for the partner group. In the event of a failover, the alternate master controller will use the master controller's IP address and *MAC address*.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Each T3 disk tray needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the disk tray target addresses are set starting with target address 1 for the bottom disk tray and continuing to the top disk tray. Use the `port list` command to verify that all T3 disk trays have a unique target address. Refer to Appendix A.3.12 of the *Sun StorEdge T3 Disk Tray Administrator's Guide* for further details.

The following three parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Host-based alternate pathing, such as VERITAS DMP or Sun Enterprise Server AP, must also be installed on the host.

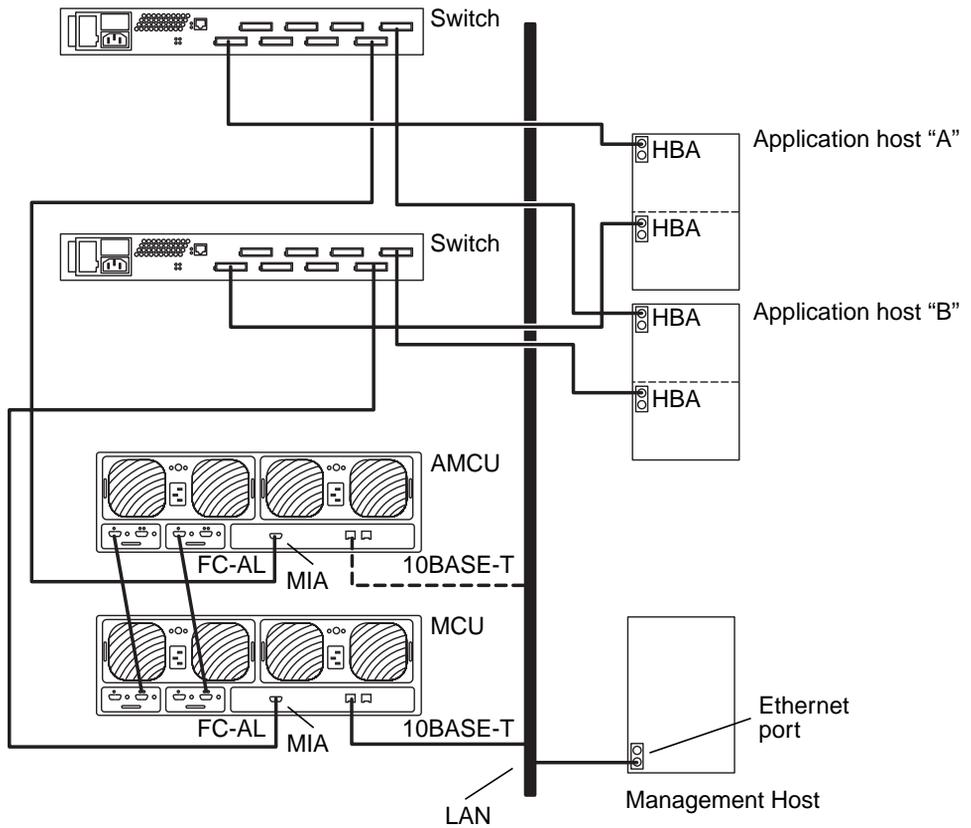


FIGURE 4-10 Dual Hosts With Two Controller Units and Two Switches

Dual Hosts With Eight Controller Units and Two Switches Configured as Four Partner Groups

FIGURE 4-11 shows two application hosts connected through FC-AL cables to two switches and four Sun StorEdge T3 disk tray partner groups, which consist of eight Sun StorEdge T3 disk trays. The 10BASE-T connection on the master controller is on a public or separate network and requires an IP address for the partner group. In the event of a failover, the alternate master controller will use the master controller's IP address and *MAC address*.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Each T3 disk tray needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the disk tray target addresses are set starting with target address 1 for the bottom disk tray and continuing to the top disk tray. Use the `port list` command to verify that all T3 disk trays have a unique target address. Refer to Appendix A.3.12 of the *Sun StorEdge T3 Disk Tray Administrator's Guide* for further details.

The following three parameters must be set on the master controller, as follows:

- `mp_support = rw`
- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3 Disk Tray Administrator's Guide*.

Host-based alternate pathing, such as VERITAS DMP or Sun Enterprise Server AP, must also be installed on the host.

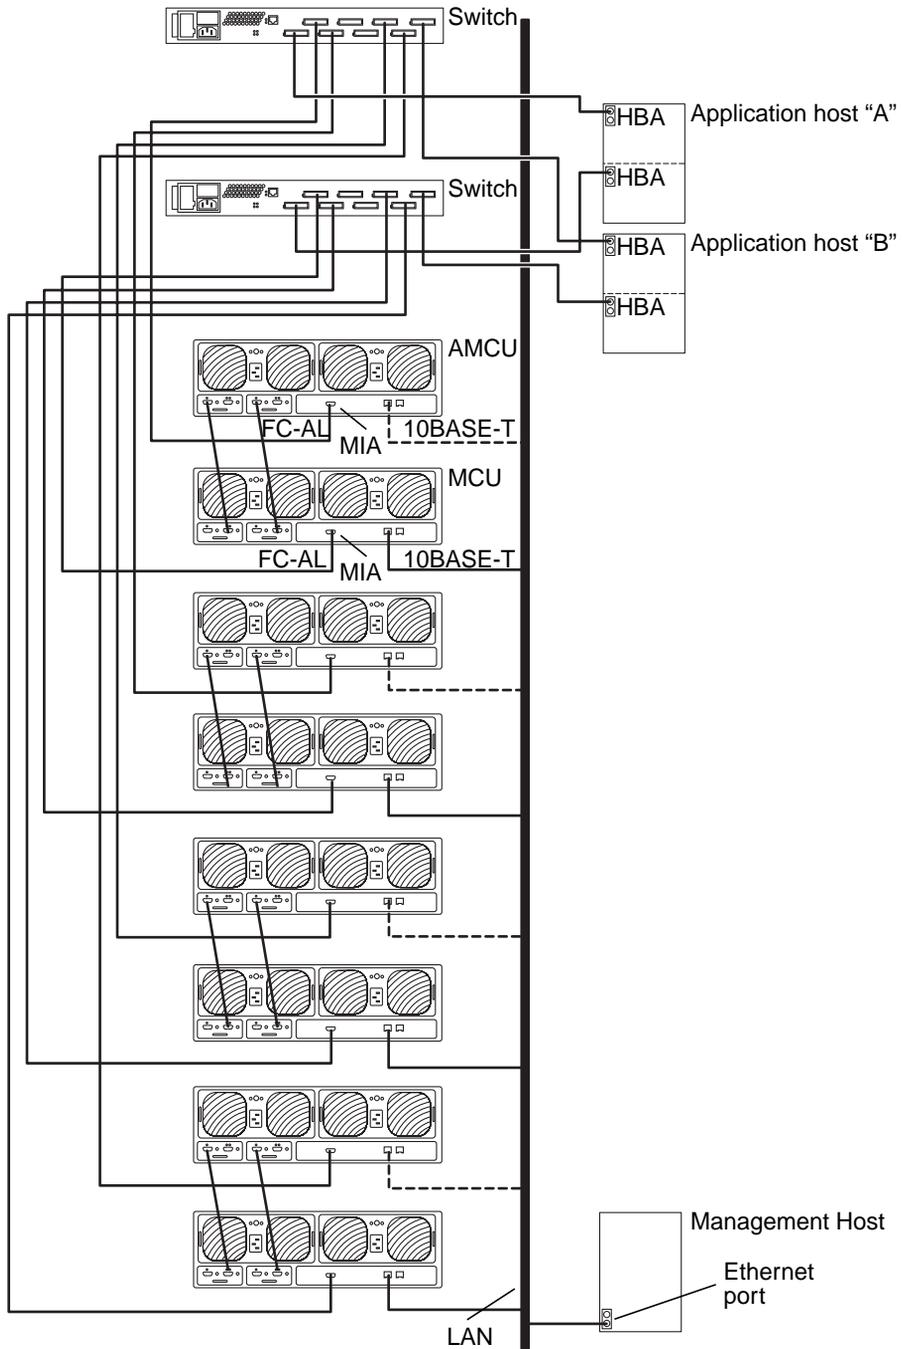


FIGURE 4-11 Dual Hosts With Eight Controller Units and Two Switches

Host Connections

This chapter describes how to connect the Sun StorEdge T3 disk tray to the following devices:

- “Sun Enterprise SBus+ and Graphics+ I/O Boards” on page 50
- “Sun StorEdge PCI FC-100 Host Bus Adapter” on page 51
- “Sun StorEdge SBus FC-100 Host Bus Adapter” on page 52
- “Sun StorEdge PCI Single Fibre Channel Network Adapter” on page 53

Sun Enterprise SBus+ and Graphics+ I/O Boards

The SBus+ and Graphics+ I/O boards each provide mounting for two *gigabit interface converters* (GBICs). For more detailed information about these I/O boards, refer to the *Sun Enterprise 6x00/5x00/4x00/3x00 Systems SBus+ and Graphics+ I/O Boards Installation Guide*, part number 805-2704. FIGURE 5-1 shows an Enterprise 6x00/5x00/4x00/3x00 SBus+ I/O board.

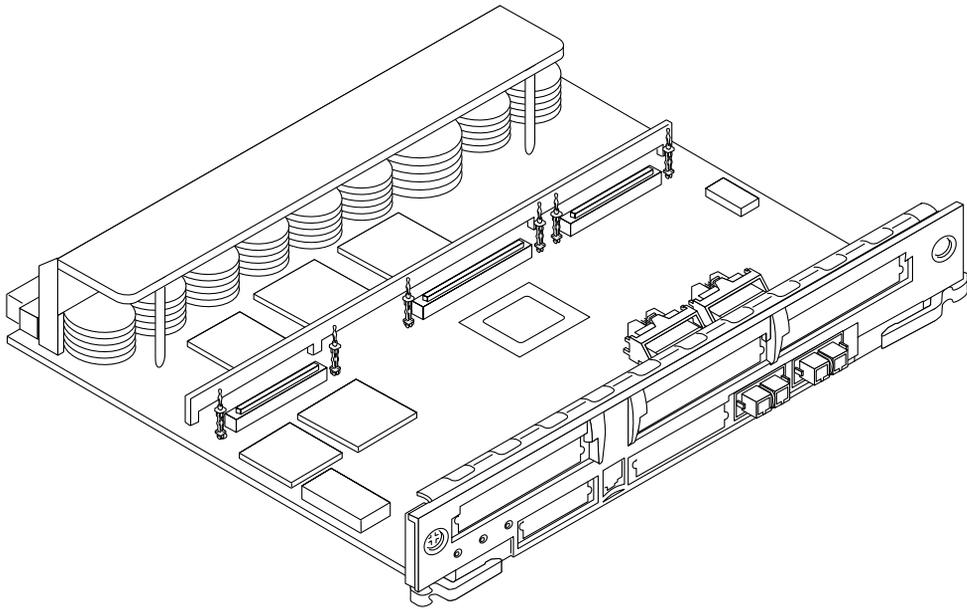


FIGURE 5-1 Sun Enterprise 6x00/5x00/4x00/3x00 SBus+ I/O Board

System Requirements

Your system must meet the following hardware and software requirements:

- Sun Enterprise 6x00/5x00/4x00/3x00 system
- An available I/O board slot
- OpenBoot™ PROM, version 3.2.10 or later
- Solaris 2.6 or higher version of the Solaris operating system

Sun StorEdge PCI FC-100 Host Bus Adapter

The Sun StorEdge PCI FC-100 host bus adapter is a 33-MHz, 100 Mbytes/second, single-loop Fibre Channel PCI host bus adapter with an onboard GBIC. This host bus adapter is PCI Version 2.1-compliant. For more detailed information about this product, refer to the *Sun StorEdge PCI FC-100 Host Adapter Installation Manual*, part number 805-3682. FIGURE 5-2 shows a Sun StorEdge PCI FC-100 host bus adapter.

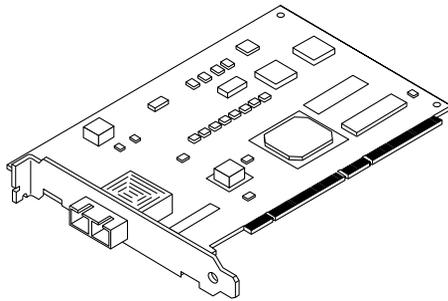


FIGURE 5-2 Sun StorEdge PCI FC-100 Host Bus Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available PCI port
- Solaris 2.6 or higher version of the Solaris operating system

Sun StorEdge SBus FC-100 Host Bus Adapter

The Sun StorEdge SBus FC-100 host bus adapter is a single-width Fibre Channel SBus card with a Sun Serial Optical Channel (SOC+) ASIC (application-specific integrated circuit). You can connect up to two loops to each card, using hot-pluggable GBICs. For more detailed information about this product, refer to the *Sun StorEdge SBus FC-100 Host Adapter Installation and Service Manual*, part number 802-7572. FIGURE 5-3 shows a Sun StorEdge SBus FC-100 host bus adapter.

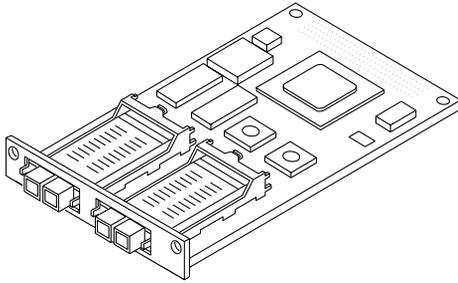


FIGURE 5-3 Sun StorEdge SBus FC-100 Host Bus Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available SBus port
- Solaris 2.6 or higher version of the Solaris operating system

Sun StorEdge PCI Single Fibre Channel Network Adapter

The Sun StorEdge PCI Single Fibre Channel network adapter is a Fibre Channel PCI card with one onboard optical receiver. This network adapter is PCI Version 2.1-compliant. For more detailed information about this product, refer to the *Sun StorEdge PCI Single Fibre Channel Network Adapter Installation Guide*, part number 806-7532-10. FIGURE 5-4 shows a Sun StorEdge PCI Single Fibre Channel network adapter.

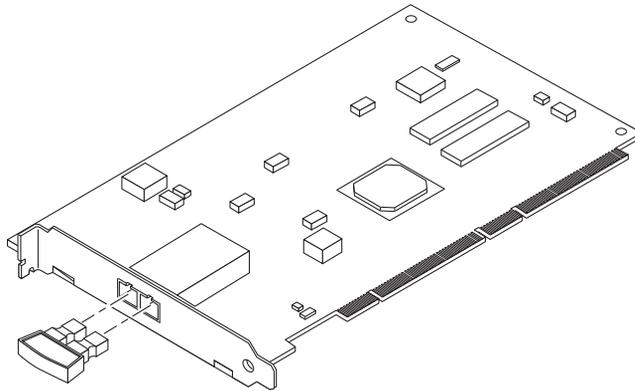


FIGURE 5-4 Sun StorEdge PCI Single Fibre Channel Network Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available PCI port
- Solaris 7 or higher version of the Solaris operating system

Disk Tray Cabling

This chapter describes the disk tray configurations supported by the Sun StorEdge T3 disk tray, and it includes the following sections:

- “Overview of Disk Tray Cabling” on page 55
- “Single Disk Tray Configurations” on page 58
- “Partner Group Configurations” on page 59

Overview of Disk Tray Cabling

The Sun StorEdge T3 disk tray has the following connections:

- One FC-AL interface to the application host
- One 10BASE-T Ethernet interface to the management host (via a LAN) for administration purposes
- One serial interface to be used for service tasks by qualified service personnel only
- Interconnect ports for configuring Sun StorEdge T3 disk trays into partner groups

Data Path

For the data path (FC-AL) connection, the only way to attach a controller unit to an application host is in “direct attached mode.” The controller unit attaches directly to the host FC connector.

Administration Path

For the administration path, each controller unit has a 10BASE-T Ethernet connector. For each installed controller, an Ethernet connection and IP address are required. The administration server uses this link to set up and manage the Sun StorEdge T3 disk trays, using the Sun StorEdge Component Manager facility.

Note – In a partner group, only one of the two Ethernet connections is active at any time. The second Ethernet connection is used for redundancy.

Connecting Partner Groups

The Sun StorEdge T3 disk tray also has two interconnect cards that connect the disk trays in a partner group. These interconnect cards have two ports (in and out).

Note – Partner groups are *not* supported in Sun Cluster 2.2.

Note – In a standalone configuration, these interconnect boards cannot be used to connect to the administrative console or to the application host. These interconnect cards are used solely for ensuring redundancy and failover mechanisms in partner groups.

FIGURE 6-1 shows a Sun StorEdge T3 disk tray with its controller card and interconnect cards.

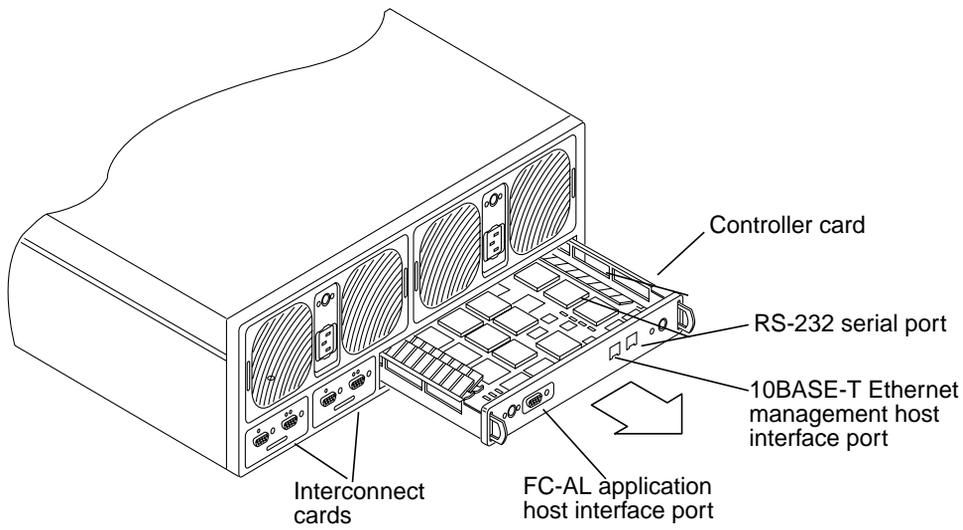


FIGURE 6-1 StorEdge Disk Tray Controller Card and Interconnect Cards

Single Disk Tray Configurations

All single disk tray configurations can be derived from FIGURE 6-2:

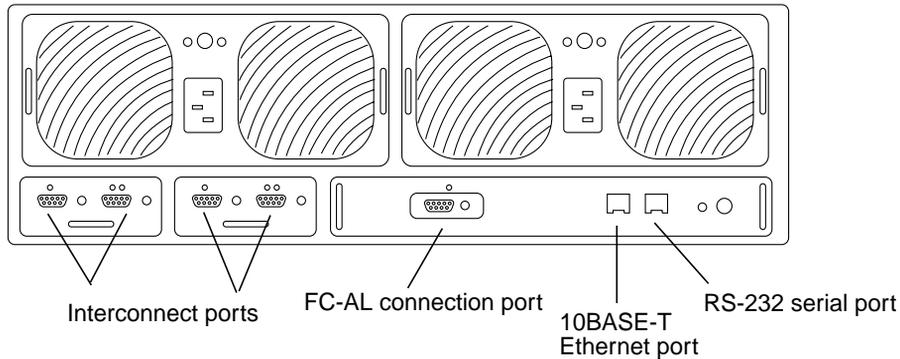


FIGURE 6-2 Single Disk Tray Configuration

The following configuration rules apply to single disk tray configurations:

- The interconnect ports, which are used only in partner group configurations, cannot be used for host connections.
- The FC-AL connection provides a data path to the application host.
- The 10BASE-T Ethernet unshielded twisted-pair connection provides a link to the management host.
- The RS-232 serial port is used solely for diagnostics and service by qualified service personnel only.
- Fiber-optic cable lengths between 2 and 500 meters, using short-wave laser and 50-micron fiber-optic cable, are supported.

Partner Group Configurations

All partner group disk tray configurations can be derived from FIGURE 6-3:

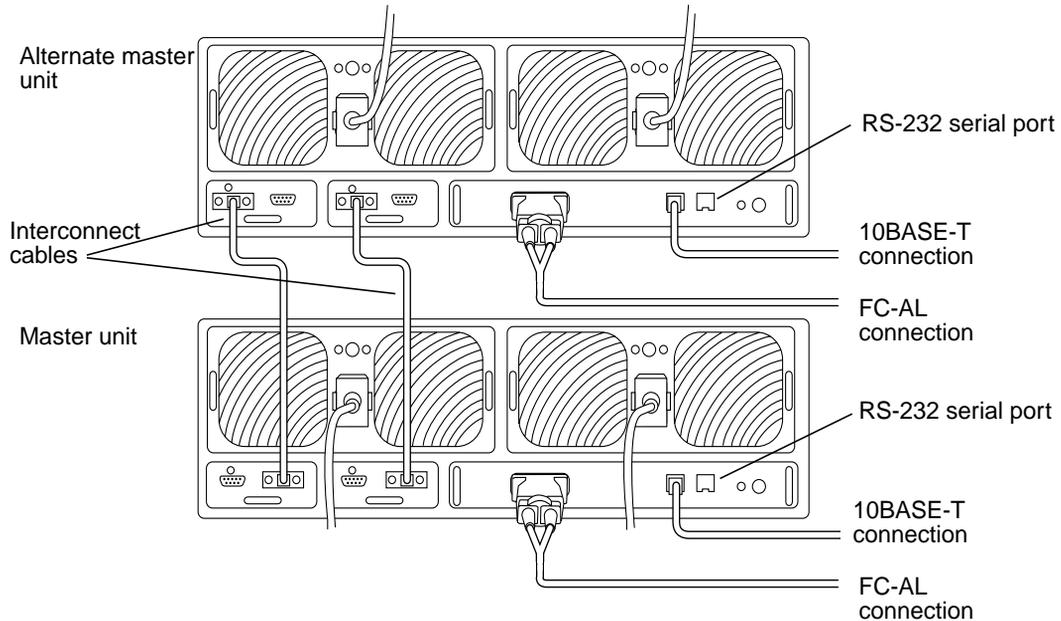


FIGURE 6-3 Partner Group Tray Configuration

The following rules configuration rules apply to partner group configurations:

- The interconnect ports, which are used only in partner group configurations, cannot be used for host connections.
- The FC-AL connection provides a data path to the application host.
- The 10BASE-T Ethernet unshielded twisted-pair connection provides a link to the management host.
- The RS-232 serial port is used solely for diagnostics and service by qualified service personnel only.
- Fiber-optic cable lengths between 2 and 500 meters, using short-wave laser and 50-micron fiber-optic cable, are supported.

This configuration is optimal because it provides full redundancy to the application hosts. Failover mechanisms are provided within the trays, but the application host has to provide data-path failover mechanisms, such as Dynamic Multi-Pathing from VERITAS Volume Manager or Alternate Pathing from Sun Enterprise Server Alternate Pathing.

Glossary

A

**alternate master
controller unit**

The secondary disk tray unit in a partner group that provides failover capability from the master unit.

auto cache mode

The default cache mode for the Sun StorEdge T3 disk tray. In a fully redundant configuration, cache is set to write-behind mode. In a nonredundant configuration, cache is set to write-through. Read caching is always performed.

auto disable

The Sun StorEdge T3 disk tray default that automatically disables a disk drive that has failed.

C

**command line interface
(CLI)**

The interface between the Sun StorEdge T3 disk tray's pSOS operating system and the user in which the user types commands to administer the disk tray.

controller unit

A disk tray that includes a controller card. It can be used as a standalone unit, or configured with other Sun StorEdge T3 disk trays.

F

**Fibre Channel-
Arbitrated Loop
(FC-AL)**

A 100-Mbyte/second serial channel that allows connection of multiple devices (disk drives and controllers).

**field-replaceable unit
(FRU)**

A component that can be removed and replaced by a field service engineer or a system administrator.

G

**gigabit interface
converter (GBIC)**

An adapter used on an SBus card to convert fiber-optic signal to copper.

**gigabyte (GB or
Gbyte)**

One gigabyte is equal to one billion bytes.

H

hot spare

A drive in a RAID 1 or RAID 5 configuration that contains no data and acts as a standby in case another drive fails.

hot-swappable

The ability to remove and replace a field-replaceable unit (FRU) while the system remains powered on and operational.

I

interconnect card

A disk tray component that contains the interface circuitry and two connectors for interconnecting multiple Sun StorEdge T3 disk tray units.

interconnect cable

An FC-AL cable with a unique switched loop architecture that is used to interconnect multiple Sun StorEdge T3 disk trays.

L

**light-emitting diode
(LED)**

A device that converts electrical energy into light that is used to display activity.

**logical unit number
(LUN)**

One or more drives that can be grouped into a unit; also called a “volume.”

M

master controller unit

The main controller unit in a partner group configuration.

**media access control
(MAC) address**

A unique address that identifies a device on a network.

**media interface adapter
(MIA)**

An adapter that converts fiber-optic light signals to copper.

**megabyte (MB or
Mbyte)**

One megabyte is equal to one million bytes.

**megabytes per second
(MB/s)**

A performance measurement of the sustained data transfer rate.

P

parity

Additional information stored with data on a disk that enables the controller to rebuild data after a drive failure.

partner group

A pair of interconnected controller units.

**power and cooling
unit**

A FRU component in the Sun StorEdge T3 disk tray. It contains a power supply, cooling fans, and an integrated UPS battery. There are two power and cooling units in a Sun StorEdge T3 disk tray.

R

read caching Data for future retrieval, to reduce disk I/O as much as possible.

**redundant array of
independent disks
(RAID)**

A configuration in which multiple drives are combined into a single virtual drive, to improve performance and reliability.

**reliability, availability,
serviceability (RAS)**

A term to describe product features that include high availability, easily serviced components, and very dependable.

S

system area Located on the disk drive label, the space that contains configuration data, boot firmware, and file system information.

U

**uninterruptable power
source (UPS)**

A component within the power and cooling unit. It supplies power from a battery in the case of an AC power failure.

V

volume Also called a LUN, a volume is one or more drives grouped into a unit for data storage.

W

write caching Data used to build up stripes of data, eliminating the read-modify-write overhead. Write caching improves performance for applications that are writing to disk.

Index

NUMERICS

- 10BASE-T Ethernet
 - administration path, 56
 - connections, 4 to 5

A

- administration path, 56
- alternate master controller unit (AMCU)
 - definition, 4
- Alternate Pathing (AP)
 - in configuration recommendations, 7
 - in partner group configuration, 30 to 31
- auto cache mode, 10

C

- cabling
 - partner group configuration, 59
 - single disk tray configuration, 58
- cache
 - allocation, configuring, 12
 - for improving performance, 9
 - mirrored, enabling, 12
 - setting cache modes, 10
- cache segment, 11
- cluster support
 - See SunCluster 2.2 support
- configuration, restrictions and recommendations, 7

- configurations
 - direct host
 - single host with eight controllers, 34 to 35
 - single host with four controllers, 32 to 33
 - single host with one controller, 28
 - single host with two controllers, 29 to 30
 - hot spare, 15
 - hub host
 - dual hosts with eight controllers, 42 to 44
 - dual hosts with four controllers, 40 to 41
 - single host with eight controllers, 38 to 39
 - single host with four controllers, 36 to 37
 - partner group, 4 to 5
 - partner group cabling, 59
 - single disk tray cabling, 58
 - standalone, 5 to 6
 - switch host
 - dual hosts with eight controllers, 47 to 48
 - dual hosts with two controllers, 45 to 46
- connections
 - 10BASE-T Ethernet, 4 to 5
 - FC-AL, 4
 - RS-232 serial, 2

D

- data block size
 - definition, 10
- data path, 55
- Dynamic Multi-Pathing (DMP)
 - in configuration recommendations, 7
 - in partner group configuration, 30 to 31

F

FC-AL

- connections, 4
- data path, 55

Fibre Channel-Arbitrated Loop (FC-AL)

- See FC-AL

H

HBA

- SOC+, 52
- Sun StorEdge PCI FC-100, 51
- Sun StorEdge PCI Single Fibre Channel network adapter, 53
- Sun StorEdge SBus FC-100, 52

hot spare

- default value, 18
- determining whether to use, 15

I

I/O boards

- Sun Enterprise SBus+ and Graphics+, 50

interconnect cards

- description, 3 to 4
- in partner groups, 56

IP addresses, for AMCU and MCU, 5

L

logical unit (LUN)

- See LUNs

LUNs

- and applications, 14
- creating and labeling, 16
- default value, 18
- definition, 13
- determining how many are needed, 14
- guidelines for configuring, 13
- reconstruction rate, setting, 16
- viewing, 13

M

MAC address, 7

master controller unit (MCU)

- in a partner group, 4, 21
- parameters controlled by, 8

MCU

- See master controller unit

media access control (MAC) address

- See MAC address

media interface adapter (MIA), 5, 20

P

parameters, tailored to I/O workload, 8

partner groups

- configuration, 4 to 5
- creating
 - from new units, 22 to 25
- definition, 19
- direct host
 - single host with eight controllers, 34 to 35
 - single host with four controllers, 32 to 33
 - single host with two controllers, 29 to 30
- how they work, 21
- hub host
 - dual hosts with eight controllers, 42 to 44
 - dual hosts with four controllers, 40 to 41
 - single host with eight controllers, 38 to 39
 - single host with four controllers, 36 to 37
- setting up before changing parameters, 9
- sharing parameter settings, 8
- switch host
 - dual hosts with eight controllers, 47 to 48
 - dual hosts with two controllers, 45 to 46
- using AP, 30 to 31
- using DMP, 30 to 31

R

RAID

- and applications, 15
- configuring for redundancy, 17
- default level, 18
- determining level required, 14
- levels, defined, 17

S

SOC+ HBA, 52

stripe unit size

See data block size

Sun Cluster 2.2 support, 8

Sun Enterprise SBus+ and Graphics+ I/O boards

See I/O boards

Sun StorEdge PCI FC-100 HBA, 51

Sun StorEdge PCI Single Fibre Channel network
adapter, 53

Sun StorEdge SBus FC-100 HBA, 52

W

write-behind cache mode, 10

write-through cache mode, 10

