

HARDWARE GUIDE

MegaRAID[®] SCSI 320-2 RAID Controller

November 2002

LSI LOGIC[®]

This document contains proprietary information of LSI Logic Corporation. The information contained herein is not to be used by or disclosed to third parties without the express written permission of an officer of LSI Logic Corporation.

LSI Logic products are not intended for use in life-support appliances, devices, or systems. Use of any LSI Logic product in such applications without written consent of the appropriate LSI Logic officer is prohibited.

Document DB15-000259-00, First Edition (November 2002)

This document describes the initial release of LSI Logic Corporation's MegaRAID SCSI 320-2 Controller and will remain the official reference source for all revisions/releases of this product until rescinded by an update.

LSI Logic Corporation reserves the right to make changes to any products herein at any time without notice. LSI Logic does not assume any responsibility or liability arising out of the application or use of any product described herein, except as expressly agreed to in writing by LSI Logic; nor does the purchase or use of a product from LSI Logic convey a license under any patent rights, copyrights, trademark rights, or any other of the intellectual property rights of LSI Logic or third parties.

Copyright © 2002 by LSI Logic Corporation. All rights reserved.

TRADEMARK ACKNOWLEDGMENT

LSI Logic, the LSI Logic logo design, and MegaRAID are trademarks or registered trademarks of LSI Logic Corporation. Intel is a registered trademark of Intel Corporation. MS-DOS, Microsoft Windows and Windows NT are registered trademarks of Microsoft Corporation. UNIX is a registered trademark of the Open Group. Novell NetWare is a registered trademark of Novell Corporation. ASPI is a registered trademark of Adaptec, Inc. All other brand and product names may be trademarks of their respective companies.

CD

To receive product literature, visit us at <http://www.lsilogic.com>.

For a current list of our distributors, sales offices, and design resource centers, view our web page located at

<http://www.lsilogic.com/contacts/index.html>

FCC Regulatory Statement

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Warning: Changes or modifications to this unit not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment.

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a specific installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, try to correct the interference by one or more of the following measures:

1. Reorient or relocate the receiving antenna.
2. Increase the separation between the equipment and the receiver.
3. Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
4. Consult the dealer or an experienced radio/TV technician for help.

Shielded interface cables must be used with this product to ensure compliance with the Class B FCC limits.

Model Number: Series 518

Disclaimer – LSI Logic certifies only that this product will work correctly when this product is used with the same jumper settings, the same system configuration, the same memory module parts, and the same peripherals that were tested by LSI Logic with this product. The complete list of tested jumper settings, system Configurations, peripheral devices, and memory modules are documented in the LSI LOGIC Compatibility Report for this product. Call your LSI Logic sales representative for a copy of the Compatibility Report for this product.

Preface

This book is the primary reference and Hardware Guide for the LSI Logic MegaRAID[®] SCSI 320-2 Controller. It contains instructions for installing the MegaRAID controller and for configuring RAID arrays. It also contains background information on RAID.

The MegaRAID SCSI 320-2 controller supports single-ended and low-voltage differential (LVD) SCSI devices on two Ultra320 and Wide SCSI channels with data transfer rates up to 320 Mbytes/s.

Audience

This document is intended for people who need to install the MegaRAID SCSI 320-2 Controller in a server and then create and configure RAID arrays.

Organization

This document has the following chapters and appendixes:

- [Chapter 1, Overview](#), provides an overview of the MegaRAID SCSI 320-2 and basic SCSI features.
- [Chapter 2, Introduction to RAID](#), introduces important RAID concepts.
- [Chapter 3, RAID Levels](#), describes each supported RAID level and the factors to consider when choosing a RAID level.
- [Chapter 4, Features](#), explains the features of the MegaRAID SCSI 320-2.
- [Chapter 5, Configuring Physical Drives, Arrays, and Logical Drives](#), explains how to configure SCSI physical drives, arrays, and logical drives.

- [Chapter 6, Hardware Installation](#), explains how to install the MegaRAID SCSI 320-2 controller.
- [Chapter 7, Installing and Configuring Clusters](#), explains how to implement clustering to enable two independent servers to access the same shared data storage.
- [Chapter 8, Troubleshooting](#), provides troubleshooting information for the MegaRAID SCSI 320-2 controller.
- [Appendix A, SCSI Cables and Connectors](#), describes the SCSI cables and connectors used with the MegaRAID SCSI 320-2 controller.
- [Appendix B, Audible Warnings](#), explains the meaning of the warning tones generated by the MegaRAID SCSI 320-2 controller.
- [Appendix C, Glossary](#), defines many terms used in this manual.

Technical Support

If you need help installing, configuring, or running the MegaRAID SCSI 320-2 Controller, you may be able to find the information you need at the MegaRAID support page at <http://megaraid.lsillogic.com>

If this does not resolve your problem, you can call your LSI Logic OEM Technical Support representative at 678-728-1250. Before you call, please complete the **MegaRAID Problem Report** form.

MegaRAID Problem Report Form

Customer Information	MegaRAID Information
Name:	Today's Date:
Company:	Date of Purchase:
Address:	Invoice Number:
City/State:	Serial Number:
Country:	
Email Address:	Cache Memory:
Phone:	Firmware Version:
Fax:	BIOS Version:

MegaRAID Problem Report Form (Cont.)

System Information	
Motherboard:	BIOS manufacturer:
Operating System:	BIOS Date:
Op. Sys. Ver.:	Video Adapter:
MegaRAID Driver Ver.:	CPU Type/Speed:
Network Card:	System Memory:
Other disk controllers installed:	Other adapter cards Installed:
Description of problem:	
Steps necessary to re-create problem: 1. 2. 3. 4.	

Logical Drive Configuration

Use this form to record the configuration details for your logical drives.

Logical Drive Configuration

Logical Drive	RAID Level	Stripe Size	Logical Drive Size	Cache Policy	Read Policy	Write Policy	# of Physical Drives
LD0							
LD1							
LD2							
LD3							
LD4							
LD5							
LD6							
LD7							
LD8							
LD9							
LD10							
LD11							
LD12							
LD13							
LD14							
LD15							
LD16							
LD17							
LD18							
LD19							
LD20							
LD21							

Logical Drive Configuration (Cont.)

Logical Drive	RAID Level	Stripe Size	Logical Drive Size	Cache Policy	Read Policy	Write Policy	# of Physical Drives
LD22							
LD23							
LD24							
LD25							
LD26							
LD27							
LD28							
LD29							
LD30							
LD31							
LD32							
LD33							
LD34							
LD35							
LD36							
LD37							
LD38							
LD39							

Physical Device Layout

Use this form to record the physical device layout.

Physical Device Layout

	Channel 0	Channel 1
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		

Physical Device Layout (Cont.)

	Channel 0	Channel 1
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		

Physical Device Layout (Cont.)

	Channel 0	Channel 1
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		

Physical Device Layout (Cont.)

	Channel 0	Channel 1
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		

Contents

Chapter 1

Overview

1.1	SCSI Channels	1-1
1.2	NVRAM and Flash ROM	1-2
1.3	SCSI Connectors	1-2
1.4	Single-Ended and Differential SCSI Buses	1-2
1.5	Maximum Cable Length for SCSI Standards	1-2
1.6	SCSI Bus Widths and Maximum Throughput	1-3
1.7	Documentation	1-3
1.7.1	MegaRAID SCSI 320-2 Hardware Guide	1-4
1.7.2	MegaRAID Configuration Software Guide	1-4
1.7.3	MegaRAID Operating System Driver Installation Guide	1-4

Chapter 2

Introduction to RAID

2.1	RAID Benefits	2-1
2.1.1	Improved I/O	2-1
2.1.2	Increased Reliability	2-2
2.2	MegaRAID SCSI 320-2 – Host-Based RAID Solution	2-2
2.2.1	Host-Based RAID	2-2
2.2.2	SCSI-to-SCSI External RAID	2-3
2.3	RAID Overview	2-3
2.3.1	Physical Array	2-3
2.3.2	Logical Drive	2-3
2.3.3	Consistency Check	2-4
2.3.4	Fault Tolerance	2-4
2.3.5	Disk Striping	2-4
2.3.6	Disk Mirroring	2-6

2.3.7	Disk Spanning	2-7
2.3.8	Parity	2-8
2.3.9	Hot Spares	2-8
2.3.10	Hot Swapping	2-9
2.3.11	Disk Rebuild	2-9
2.3.12	Logical Drive States	2-10
2.3.13	SCSI Drive States	2-10
2.3.14	Disk Array Types	2-11
2.3.15	Enclosure Management	2-11

Chapter 3 RAID Levels

3.1	Selecting a RAID Level	3-1
3.2	RAID 0	3-2
3.3	RAID 1	3-3
3.4	RAID 5	3-4
3.5	RAID 10	3-6
3.6	RAID 50	3-7

Chapter 4 Features

4.1	SMART Technology	4-3
4.2	Configuration on Disk	4-3
4.3	Configuration Features	4-3
4.4	Array Performance Features	4-5
4.5	RAID Management Features	4-5
4.6	Fault Tolerance Features	4-6
4.7	Software Utilities	4-6
4.8	Operating System Software Drivers	4-7
4.9	MegaRAID SCSI 320-2 Specifications	4-7
4.10	MegaRAID Components	4-8
4.10.1	CPU	4-8
4.10.2	Cache Memory	4-9
4.10.3	MegaRAID BIOS	4-9
4.10.4	Onboard Speaker	4-9
4.10.5	Serial Port	4-9
4.10.6	SCSI Bus	4-10

4.10.7	SCSI Connectors	4-10
4.10.8	SCSI Termination	4-10
4.10.9	SCSI Firmware	4-10
4.11	RAID Management	4-11
4.11.1	MegaRAID BIOS Configuration Utility	4-11
4.11.2	WebBIOS Configuration Utility	4-11
4.11.3	Power Console Plus	4-12
4.11.4	MegaRAID Manager	4-12
4.12	Compatibility	4-12
4.12.1	Server Management	4-12
4.12.2	SCSI Device Compatibility	4-12
4.12.3	Software	4-12

Chapter 5

Configuring Physical Drives, Arrays, and Logical Drives

5.1	Configuring SCSI Physical Drives	5-1
5.1.1	Distributing Drives	5-1
5.1.2	Basic Configuration Rules	5-2
5.1.3	Current Physical Device Configuration	5-2
5.1.4	Logical Drive Configuration	5-4
5.1.5	Physical Device Layout	5-6
5.2	Configuring Arrays	5-9
5.2.1	Arranging Arrays	5-9
5.2.2	Creating Hot Spares	5-10
5.3	Creating Logical Drives	5-10
5.3.1	Configuration Strategies	5-10
5.3.2	Assigning RAID Levels	5-12
5.4	Configuring Logical Drives	5-13
5.4.1	Optimizing Data Storage	5-13
5.5	Planning the Array Configuration	5-14
5.5.1	Using the Array Configuration Planner	5-14

Chapter 6

Hardware Installation

6.1	Hardware Requirements	6-1
6.1.1	Optional Equipment	6-2
6.2	Installation Steps	6-2

6.2.1	Step 1: Unpack	6-3
6.2.2	Step 2: Power Down	6-3
6.2.3	Step 3: Install Cache Memory	6-4
6.2.4	Step 4: Check Jumper Settings	6-5
6.2.5	Step 5: Set Termination	6-8
6.2.6	Step 6 Set SCSI Terminator Power	6-13
6.2.7	Step 7: Install Battery Pack (Optional)	6-13
6.2.8	Step 8: Install MegaRAID SCSI 320-2	6-16
6.2.9	Step 9: Connect SCSI Devices	6-17
6.2.10	Step 10: Set Target IDs	6-18
6.2.11	Step 11: Power Up	6-20
6.2.12	Step 12: Run the MegaRAID BIOS Configuration Utility	6-20
6.2.13	Step 13: Install the Operating System Driver	6-21
6.3	Summary	6-21

Chapter 7

Installing and Configuring Clusters

7.1	Overview	7-1
7.2	Benefits of Clusters	7-2
7.3	Installation and Configuration	7-2
7.3.1	Driver Installation Instructions under Microsoft Windows 2000 Advanced Server	7-3
7.3.2	Network Requirements	7-4
7.3.3	Shared Disk Requirements	7-5
7.4	Cluster Installation	7-5
7.4.1	Installation Overview	7-5
7.4.2	Installing Microsoft Windows 2000	7-7
7.4.3	Setting Up Networks	7-7
7.4.4	Configuring the Cluster Node Network Adapter	7-8
7.4.5	Setting Up the First Node in Your Cluster	7-9
7.4.6	Configuring the Public Network Adapter	7-10
7.4.7	Verifying Connectivity and Name Resolution	7-11
7.4.8	Verifying Domain Membership	7-12
7.4.9	Setting Up a Cluster User Account	7-12
7.4.10	Setting Up Shared Disks	7-13
7.4.11	Configuring Shared Disks	7-14
7.4.12	Assigning Drive Letters	7-15

7.4.13	Verifying Disk Access and Functionality	7-15
7.4.14	Installing Cluster Service Software	7-16
7.4.15	Configuring Cluster Disks	7-18
7.4.16	Validating the Cluster Installation	7-25
7.4.17	Configuring the Second Node	7-25
7.4.18	Verifying Installation	7-26
7.5	Installing SCSI Drives	7-27
7.5.1	Configuring the SCSI Devices	7-28
7.5.2	Terminating the Shared SCSI Bus	7-28

Chapter 8 Troubleshooting

8.1	General Troubleshooting	8-1
8.2	BIOS Boot Error Messages	8-3
8.3	Other BIOS Error Messages	8-5
8.4	Other Potential Problems	8-5

Appendix A SCSI Cables and Connectors

A.1	68-Pin High-Density SCSI Internal Connector	A-1
A.1.1	Cable Assembly for Internal Wide SCSI Devices	A-2
A.1.2	Connecting Internal and External Wide Devices	A-3
A.1.3	Converting Internal Wide to Internal Non-Wide (Type 2)	A-4
A.1.4	Converting Internal Wide to Internal Non-Wide (Type 30)	A-5
A.1.5	Converting from Internal Wide to Internal Non-Wide (Type 3)	A-5
A.2	SCSI Cable and Connector Vendors	A-6
A.3	High-Density 68-Pin Connector Pinout for SE SCSI	A-7
A.4	68-Pin Connector Pinout for LVD SCSI	A-9

Appendix B Audible Warnings

Appendix C Glossary

Index

Customer Feedback

Chapter 1

Overview

This chapter provides an overview of the MegaRAID[®] SCSI 320-2 controller. It contains the following sections:

- [Section 1.1, “SCSI Channels”](#)
- [Section 1.2, “NVRAM and Flash ROM”](#)
- [Section 1.3, “SCSI Connectors”](#)
- [Section 1.4, “Single-Ended and Differential SCSI Buses”](#)
- [Section 1.5, “Maximum Cable Length for SCSI Standards”](#)
- [Section 1.6, “SCSI Bus Widths and Maximum Throughput”](#)
- [Section 1.7, “Documentation”](#)

The MegaRAID SCSI 320-2 controller has two SCSI channels. Throughput on each SCSI channel can be as high as 320 Mbytes/s. MegaRAID supports a low-voltage differential SCSI bus or a single-ended SCSI bus.

The MegaRAID SCSI 320-2 is a high-performance intelligent PCI-to-SCSI host adapter with RAID control capabilities. The MegaRAID SCSI 320-2 requires no special motherboard PCI expansion slot. The card includes an Intel 80303 processor. MegaRAID provides reliability, high performance, and fault-tolerant disk subsystem management.

1.1 SCSI Channels

The MegaRAID SCSI 320-2 has two Ultra320 SCSI channels. One LSI Logic dual SCSI controller supports both channels. Each SCSI channel supports up to fifteen SCSI devices.

1.2 NVRAM and Flash ROM

A 32 KB x 8 NVRAM stores RAID system configuration information. The MegaRAID SCSI 320-2 firmware is stored in flash ROM for easy upgrade.

1.3 SCSI Connectors

The MegaRAID SCSI 320-2 has two ultra high-density 68-pin external SCSI connectors and two 68-pin internal SCSI connectors for internal SCSI drives.

1.4 Single-Ended and Differential SCSI Buses

The SCSI standard defines two electrical buses:

- Single-ended bus
 - Low-voltage differential bus
-

1.5 Maximum Cable Length for SCSI Standards

[Table 1.1](#) lists the maximum SCSI cable length and number of disk drives that you can use, depending on the SCSI speeds and type of device.

Table 1.1 Maximum Cable Length for SCSI Standards

Standard	Single Ended SCSI	Low-Voltage Differential SCSI	Maximum # of Drives
Ultra SCSI	1.5 m	12 m	7
Ultra SCSI	3 m	12 m	3
Wide Ultra SCSI		12 m	15
Wide Ultra SCSI	1.5 m	12 m	7
Wide Ultra SCSI	3 m	12 m	3

Table 1.1 Maximum Cable Length for SCSI Standards (Cont.)

Standard	Single Ended SCSI	Low-Voltage Differential SCSI	Maximum # of Drives
Ultra 2 SCSI		25 m	1
Ultra 2 SCSI		12 m	7
Wide Ultra 2 SCSI		25 m	1
Wide Ultra 2 SCSI		12 m	15
Ultra320		12 m	15

1.6 SCSI Bus Widths and Maximum Throughput

Table 1.2 lists the SCSI bus widths and maximum throughput, based on the SCSI speeds.

Table 1.2 SCSI Bus Widths and Maximum Throughput

SCSI Standard	SCSI Bus Width	SCSI Throughput
Fast Wide SCSI	16 bits	20 Mbytes/s
Wide Ultra SCSI	16 bits	40 Mbytes/s
Wide Ultra 2 SCSI	16 bits	80 Mbytes/s
Wide Ultra 160 SCSI	16 bits	160 Mbytes/s
Ultra 320 SCSI	16 bits	320 Mbytes/s

1.7 Documentation

The MegaRAID SCSI 320-2 documentation set includes:

- The *MegaRAID SCSI 320-2 Hardware Guide*
- The *MegaRAID Configuration Software Guide*
- The *MegaRAID Operating System Driver Installation Guide*

1.7.1 MegaRAID SCSI 320-2 Hardware Guide

The Hardware Guide contains the RAID overview, RAID planning, and RAID system configuration information you need first. Read this document first.

1.7.2 MegaRAID Configuration Software Guide

This manual describes the software configuration utilities that you can use to configure and modify RAID systems:

- MegaRAID BIOS Configuration Utility
- WebBIOS Configuration Utility
- MegaRAID Manager
- Power Console Plus

1.7.3 MegaRAID Operating System Driver Installation Guide

This manual provides detailed information about installing the MegaRAID SCSI 320-2 operating system drivers.

Chapter 2

Introduction to RAID

This chapter introduces important RAID concepts. It contains the following sections:

- [Section 2.1, “RAID Benefits”](#)
- [Section 2.2, “MegaRAID SCSI 320-2 – Host-Based RAID Solution”](#)
- [Section 2.3, “RAID Overview”](#)

RAID (Redundant Array of Independent Disks) is a data storage method in which data, along with parity information, is distributed among two or more hard disks (called an *array*) to improve performance and reliability. The RAID array appears to the host computer as a single storage unit or as multiple logical units. I/O is expedited because several disks can be accessed simultaneously. RAID systems provide improved data storage reliability and fault tolerance compared to single-drive computers. If a disk drive in a RAID array fails, data can be reconstructed from the data and parity information on the remaining disk drives.

2.1 RAID Benefits

RAID is widely used because it improves I/O performance and increases storage subsystem reliability. RAID provides data security through fault tolerance and redundant data storage. The MegaRAID SCSI 320-2 management software configures and monitors RAID disk arrays.

2.1.1 Improved I/O

Although disk drive capabilities have improved drastically, actual performance of individual disk drives has been improved only three to four times in the last decade. RAID provides a way to achieve much better data throughput.

2.1.2 Increased Reliability

The electromechanical components of a disk subsystem operate more slowly, require more power, and generate more noise and vibration than electronic devices. These factors reduce the reliability of data stored on disks. RAID provides a way to achieve much better fault tolerance and data reliability.

2.2 MegaRAID SCSI 320-2 – Host-Based RAID Solution

RAID products are either host-based or external.

The MegaRAID SCSI 320-2 controller is a host-based RAID solution. The MegaRAID SCSI 320-2 is a PCI adapter card that is installed in any available PCI expansion slot in a host system.

2.2.1 Host-Based RAID

A host-based RAID product puts all of the RAID intelligence on an adapter card that is installed in a network server. A host-based RAID product provides the best performance. MegaRAID SCSI 320-2 is part of the file server, so it can transmit data directly across the computer's buses at data transfer speeds up to 532 Mbytes/s.

The available sequential data transfer rate is determined by the following factors:

- The sustained data transfer rate on the motherboard PCI bus
- The sustained data transfer rate on the PCI-to-PCI bridge
- The sustained data transfer rate of the SCSI controller
- The sustained data transfer rate of the SCSI devices
- The number of SCSI channels
- The number of SCSI disk drives

Host-based solutions must provide operating system-specific drivers.

2.2.2 SCSI-to-SCSI External RAID

A SCSI-to-SCSI external RAID product puts the RAID intelligence inside the RAID chassis and uses a plain SCSI host adapter installed in the network server. The data transfer rate is limited to the bandwidth of the SCSI channel. A SCSI-to-SCSI external RAID product that has two Wide SCSI channels operating at speeds up to 320 Mbytes/s must squeeze the data into a single Wide SCSI (320 Mbytes/s) channel back to the host computer.

In SCSI-to-SCSI external RAID products, the disk drive subsystem uses only a single SCSI ID, which allows you to connect multiple drive subsystems to a single SCSI controller.

2.3 RAID Overview

RAID is a collection of specifications that describes a system for ensuring the reliability and stability of data stored on large disk subsystems. A RAID system can be implemented in a number of different versions (or RAID levels). MegaRAID SCSI 320-2 supports standard RAID levels 0, 1, and 5, and RAID levels 10 and 50, special RAID versions supported by the MegaRAID controller.

2.3.1 Physical Array

A RAID array is a collection of physical disk drives governed by the RAID management software. A RAID array appears to the host computer as one or more logical drives.

2.3.2 Logical Drive

A *logical drive* is a partition in a physical array of disks that is made up of contiguous data segments on the physical disks. A logical drive can consist of any of the following:

- An entire physical array
- More than one entire physical array
- A part of an array
- Parts of more than one array

- A combination of any two of the above conditions

2.3.3 Consistency Check

A *consistency check* verifies the correctness of redundant data in a RAID array. For example, in a system with distributed parity, checking consistency means computing the parity of the data drives and comparing the results to the contents of the parity drives.

2.3.4 Fault Tolerance

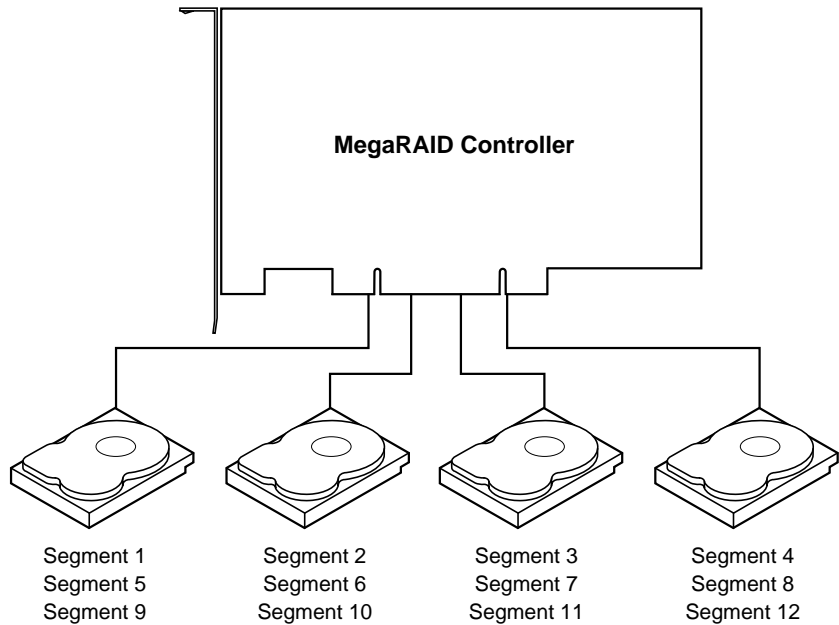
Fault tolerance is achieved through cooling fans, power supplies, and the ability to hot swap drives. MegaRAID SCSI 320-2 provides hot swapping through the hot spare feature. A hot spare drive is an unused online available drive that MegaRAID SCSI 320-2 instantly plugs into the system when an active drive fails.

After the hot spare is automatically moved into the RAID subsystem, the failed drive is automatically rebuilt. The RAID disk array continues to handle request while the rebuild occurs.

2.3.5 Disk Striping

Disk striping writes data across multiple disk drives instead of just one disk drive, as shown in [Figure 2.1](#).

Figure 2.1 Disk Striping



Disk striping involves partitioning each disk drive's storage space into stripes that can vary in size from 2 to 128 Kbytes. These stripes are interleaved in a repeated, sequential manner. The combined storage space is composed of stripes from each drive. MegaRAID SCSI 320-2 supports stripe sizes of 2, 4, 8, 16, 32, 64, or 128 Kbytes.

For example, in a four-disk system using only disk striping (as in RAID level 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple drives are accessed simultaneously; but disk striping does not provide data redundancy.

2.3.5.1 Stripe Width

Stripe width is a measure of the number of disks involved in an array where striping is implemented. For example, a four-disk array with disk striping has a stripe width of four.


2.3.5.2 Stripe Size

The *stripe size* is the length of the interleaved data segments that MegaRAID SCSI 320-2 writes across multiple drives. MegaRAID SCSI 320-2 supports stripe sizes of 2, 4, 8, 16, 32, 64, or 128 Kbytes.

2.3.6 Disk Mirroring

With *disk mirroring* (used in RAID 1), data written to one disk drive is simultaneously written to another disk drive, as shown in [Figure 2.2](#).

Figure 2.2 Disk Mirroring



If one disk drive fails, the contents of the other disk drive can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100% data redundancy. Since the contents of the disk drive are completely written to a second drive, it does not matter if one of the drives fails. Both drives contain the same data at all times. Either drive can act as the operational drive.


Although disk mirroring provides 100% redundancy, it is expensive because each drive in the system must be duplicated.

2.3.7 Disk Spanning

Disk spanning allows multiple disk drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 60 Gbyte disk drives can be combined to appear to the operating system as one single 240 Gbyte drive.

Disk spanning alone does not provide reliability or performance enhancements. Spanned logical drives must have the same stripe size and must be contiguous. In [Figure 2.3](#), two RAID 1 arrays are turned into a RAID 10 array.

Figure 2.3 Disk Spanning



The controller supports a span depth of eight. This means that eight RAID 1, 3, or 5 arrays can be spanned to create one logical drive.

Table 2.1 describes how disk spanning is used for RAID 10 and RAID 50.

Table 2.1 Spanning for RAID 10 and RAID 50

Level	Description
10	Configure RAID 10 by spanning two contiguous RAID 1 logical drives. The RAID 1 logical drives must have the same stripe size.
50	Configure RAID 50 by spanning two contiguous RAID 5 logical drives. The RAID 5 logical drives must have the same stripe size.

Note: Spanning two contiguous RAID 0 logical drives does not produce a new RAID level or add fault tolerance. It does increase the size of the logical volume and improves performance by doubling the number of spindles.

2.3.8 Parity

Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets. Parity data does not fully duplicate the parent data sets. In RAID, this method is applied to entire drives (dedicated parity) or to stripes across all disk drives in an array (distributed parity).

RAID 5 combines distributed parity with disk striping. If a single disk drive fails, it can be rebuilt from the parity and the data on the remaining drives. Parity provides redundancy for one drive failure without duplicating the contents of entire disk drives, but parity generation can slow the write process.

2.3.9 Hot Spares

A *hot spare* is an extra, unused disk drive that is part of the disk subsystem. It is usually in standby mode, ready for service if a drive fails. Hot spares permit you to replace failed drives without system shutdown or user intervention.

MegaRAID SCSI 320-2 implements automatic and transparent rebuilds using hot spare drives, providing a high degree of fault tolerance and zero downtime. The MegaRAID SCSI 320-2 RAID Management software allows you to specify physical drives as hot spares. When a hot spare is needed, the MegaRAID SCSI 320-2 controller assigns the hot spare that

has a capacity closest to and at least as great as that of the failed drive to take the place of the failed drive.

Note: Hot spares are used only in arrays with redundancy—for example, RAID levels 1, 5, 10, and 50. A hot spare connected to a specific MegaRAID SCSI 320-2 controller can be used only to rebuild a drive that is connected to the same controller.

2.3.10 Hot Swapping

Hot swapping is the manual replacement of a defective physical disk unit while the computer is still running. When a new drive has been installed, you must issue a command to rebuild the drive. The MegaRAID controller can be configured to detect the new disks and to rebuild the contents of the disk drive automatically.

2.3.11 Disk Rebuild

You rebuild a disk drive by recreating the data that had been stored on the drive before the drive failed. Rebuilding can be done only in arrays with data redundancy such as RAID level 1, 5, 10, and 50.

Standby (warm spare) rebuild is employed in a mirrored (RAID 1) system. If a disk drive fails, an identical drive is immediately available. The primary data source disk drive is the original disk drive.

A hot spare can be used to rebuild disk drives in RAID 1, 5, 10, or 50 systems. If a hot spare is not available, the failed disk drive must be replaced with a new disk drive so the data on the failed drive can be rebuilt.

The MegaRAID SCSI 320-2 controller automatically and transparently rebuilds failed drives with user-definable rebuild rates. If a hot spare is available, the rebuild starts automatically when a drive fails. MegaRAID SCSI 320-2 automatically restarts the system and the rebuild if the system goes down during a rebuild.

2.3.11.1 Rebuild Rate

The *rebuild rate* is the fraction of the compute cycles dedicated to rebuilding failed drives. A rebuild rate of 100% means the system is

totally dedicated to rebuilding the failed drive. The MegaRAID SCSI 320-2 rebuild rate can be configured between 0% and 100%. At 0%, the rebuild is only done if the system is not doing anything else. At 100%, the rebuild has a higher priority than any other system activity.

2.3.12 Logical Drive States

[Table 2.2](#) describes the logical drive states.

Table 2.2 Logical Drive States

State	Description
Optimal	Drive operating condition is good. All configured drives are online.
Degraded	Drive operating condition is not optimal. A configured drive has failed or is offline.
Failed	Drive has failed.
Offline	Drive is not available to MegaRAID SCSI 320-2.

2.3.13 SCSI Drive States

A SCSI disk drive can be in one of the states described in [Table 2.3](#).

Table 2.3 SCSI Drive States

State	Description
Online (ONLIN)	The drive is functioning normally and is a part of a configured logical drive.
Ready (READY)	The drive is functioning normally but is not part of a configured logical drive and is not designated as a hot spare.
Hot Spare (HOTSP)	The drive is powered up and ready for use as a spare in case an online drive fails.
Fail (FAIL)	A fault has occurred in the drive, placing it out of service.
Rebuild (REB)	The drive is being rebuilt with data from a failed drive.

2.3.14 Disk Array Types

Table 2.4 describes the RAID disk array types.

Table 2.4 Disk Array Types

Type	Description
Software-Based	The array is managed by software running in a host computer using the host CPU bandwidth. The disadvantages associated with this method are the load on the host CPU and the need for different software for each operating system.
SCSI to SCSI	The array controller resides outside of the host computer and communicates with the host via a SCSI adapter in the host. Array management software runs in the controller. It is transparent to the host and independent of the host operating system. The disadvantage is the limited data transfer rate of the SCSI channel between the SCSI adapter and the array controller.
Bus-Based	The array controller resides on the bus (for example, a PCI or EISA bus) in the host computer and has its own CPU to generate parity and handle other RAID functions. A bus-based controller can transfer data at the speed of the host bus but is limited to the bus it is designed for. MegaRAID SCSI 320-2 resides on a PCI bus, which can handle data transfer at up to 132 Mbytes/s. With MegaRAID SCSI 320-2, the channel can handle data transfer rates up to 320 Mbytes/s per SCSI channel.

2.3.15 Enclosure Management

Enclosure management is the intelligent monitoring of the disk subsystem by software and/or hardware.

The disk subsystem can be part of the host computer or can be separate from it. Enclosure management helps you stay informed of events in the disk subsystem, such as a drive or power supply failure. Enclosure management increases the fault tolerance of the disk subsystem.

Chapter 3

RAID Levels

This chapter describes each supported RAID level and the factors to consider when choosing a RAID level. It contains the following sections:

- [Section 3.1, “Selecting a RAID Level”](#)
 - [Section 3.2, “RAID 0”](#)
 - [Section 3.3, “RAID 1”](#)
 - [Section 3.4, “RAID 5”](#)
 - [Section 3.5, “RAID 10”](#)
 - [Section 3.6, “RAID 50”](#)
-

3.1 Selecting a RAID Level

To ensure the best performance, you should select the optimal RAID level when you create a system drive. The optimal RAID level for your disk array depends on a number of factors:

- The number of drives in the disk array
- The capacity of the drives in the array
- The need for data redundancy
- The disk performance requirements


3.2 RAID 0

RAID 0 provides disk striping across all drives in the RAID subsystem. RAID 0 does not provide any data redundancy, but does offer the best performance of any RAID level. RAID 0 breaks up data into smaller blocks and then writes a block to each drive in the array. The size of each block is determined by the stripe size parameter, set during the creation of the RAID set. RAID 0 offers high bandwidth. By breaking up a large file into smaller blocks, MegaRAID SCSI 320-2 can use several drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance.

Uses	RAID 0 provides high data throughput, especially for large files. Suitable for any environment that does not require fault tolerance.
Strong Points	Provides increased data throughput for large files. No capacity loss penalty for parity.
Weak Points	Does not provide fault tolerance. All data lost if any drive fails.
Drives	1 to 30 The initiator takes one ID per channel. This leaves 15 IDs available for each channel.

Figure 3.1 shows a RAID 0 array with four disk drives.

Figure 3.1 RAID 0 Array



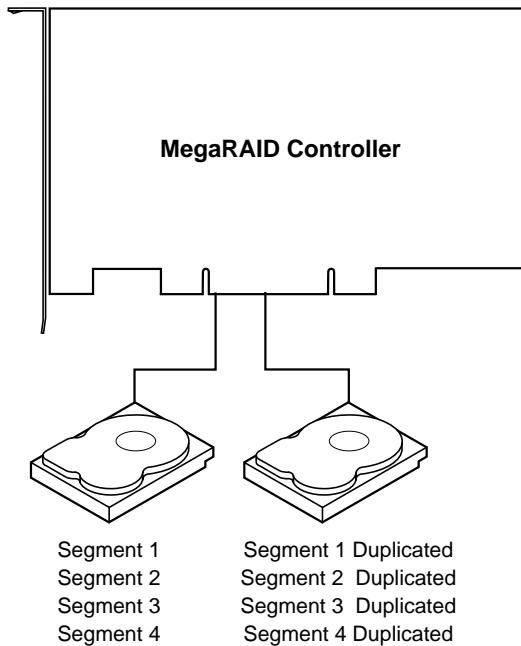
3.3 RAID 1

In RAID 1, the MegaRAID SCSI 320-2 duplicates all data from one drive to a second drive. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity.

Uses	Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.
Strong Points	Provides complete data redundancy. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
Weak Points	Requires twice as many disk drives. Performance is impaired during drive rebuilds.
Drives	2

Figure 3.2 shows a RAID 1 array.

Figure 3.2 RAID 1 Array



3.4 RAID 5

RAID 5 includes disk striping at the byte level and parity. In RAID 5, the parity information is written to several drives. RAID 5 is best suited for networks that perform many small I/O transactions simultaneously.

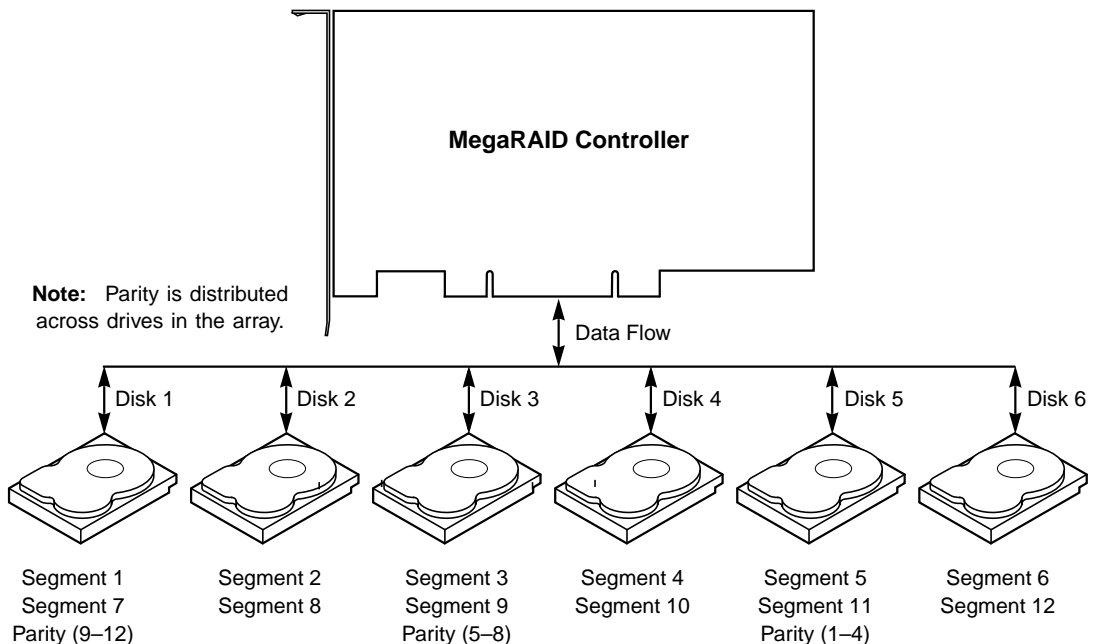
RAID 5 addresses the bottleneck issue for random I/O operations. Since each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based

exclusive-or assist make RAID 5 performance exceptional in many different environments.

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications, because each drive can read and write independently. If a drive fails, the MegaRAID SCSI 320-2 uses the distributed parity data to recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy and good performance in most environments
Weak Points	Disk drive performance is reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 to 30

Figure 3.3 shows a RAID 5 array with six disk drives.

Figure 3.3 RAID 5 Array




3.5 RAID 10

RAID 10 is a combination of RAID 0 and RAID 1. RAID 10 has mirrored drives. RAID 10 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 1 RAID set. Each RAID 1 RAID set then duplicates its data to its other drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. RAID 10 can sustain one to four drive failures while maintaining data integrity, if each failed disk is in a different RAID 1 array.

Uses	Works best for data storage that must have 100% redundancy of mirrored arrays and that also needs the enhanced I/O performance of RAID 0 (striped arrays). RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except RAID 1.
Drives	4 to 30 (must be a multiple of 2)

Figure 3.4 shows a RAID 10 array with four disk drives.

Figure 3.4 RAID 10 Array



3.6 RAID 50

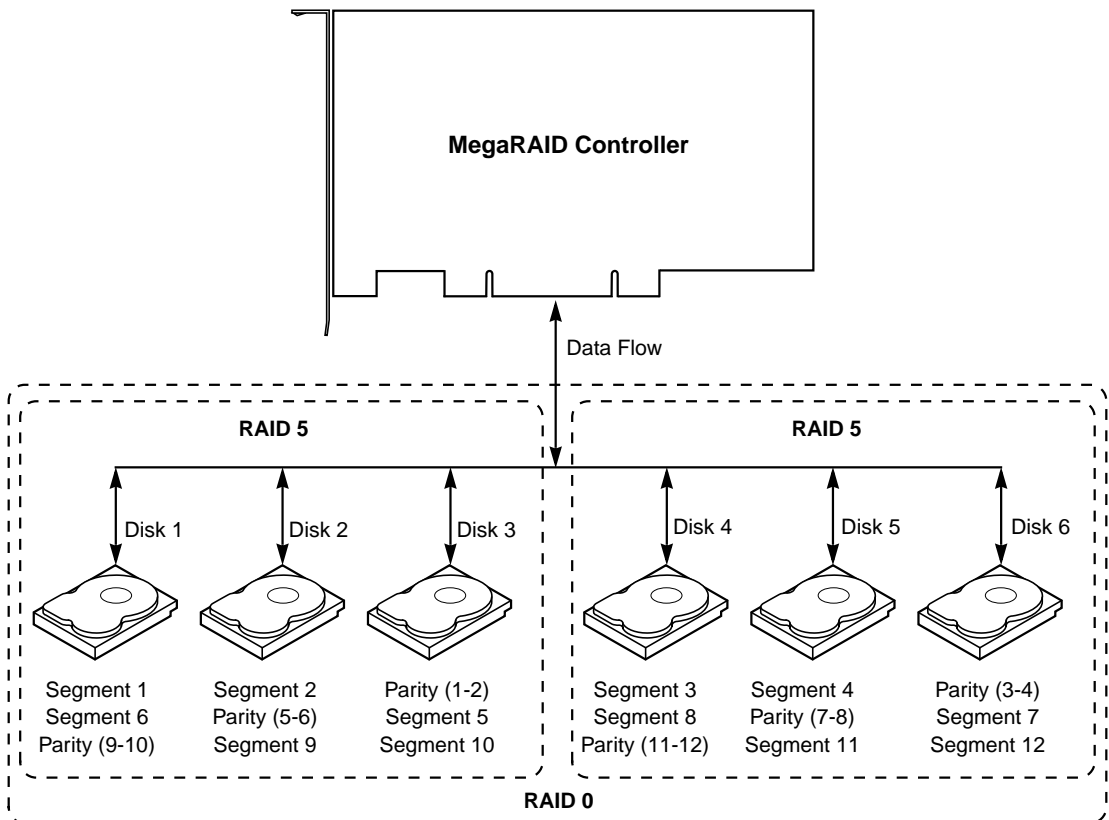
RAID 50 provides the features of both RAID 0 and RAID 5, including both parity and disk striping across multiple drives. RAID 50 is best implemented on two RAID 5 disk arrays with data striped across both disk arrays. RAID 50 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 5 RAID set. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, and then writes the blocks of data and parity to each drive in the array. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID 50 can sustain one to four drive failures while maintaining data integrity, if each failed disk is in a different RAID 5 array.

Uses	Works best when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.
Strong Points	Provides high data throughput, data redundancy, and very good performance.
Weak Points	Requires 2 to 4 times as many parity drives as RAID 5.
Drives	6 to 30 The initiator takes one ID per channel. This leaves 15 IDs available for each channel.

Figure 3.5 shows a RAID 50 array with six disk drives.

Figure 3.5 RAID 50 Array



Chapter 4

Features

This chapter explains the features of the MegaRAID SCSI 320-2. It contains the following sections:

- [Section 4.1, “SMART Technology”](#)
- [Section 4.2, “Configuration on Disk”](#)
- [Section 4.3, “Configuration Features”](#)
- [Section 4.4, “Array Performance Features”](#)
- [Section 4.5, “RAID Management Features”](#)
- [Section 4.6, “Fault Tolerance Features”](#)
- [Section 4.7, “Software Utilities”](#)
- [Section 4.8, “Operating System Software Drivers”](#)
- [Section 4.9, “MegaRAID SCSI 320-2 Specifications”](#)
- [Section 4.10, “MegaRAID Components”](#)
- [Section 4.11, “RAID Management”](#)
- [Section 4.12, “Compatibility”](#)

MegaRAID is a family of high performance intelligent PCI-to-SCSI host adapters with RAID control capabilities. The MegaRAID SCSI 320-2 has two SCSI channels that support Ultra320 and Wide SCSI with data transfer rates of up to 320 Mbytes/s. Each SCSI channel supports up to 15 Wide devices and up to seven non-Wide devices.

MegaRAID SCSI 320-2 features include:

- Remote configuration and array management through MegaRAID WebBIOS
- High-performance I/O migration path while preserving existing PCI-SCSI software

- SCSI data transfers up to 320 Mbytes/s
- Synchronous operation on a wide LVD SCSI bus
- Up to 15 LVD SCSI devices on each of the Wide buses
- Up to 256 Mbytes of 3.3 V PC100 (or faster) SDRAM cache memory in one single-sided or double-sided DIMM socket (Cache memory is used for read and write-back caching and for RAID 5 parity generation.)
- Non-volatile random access memory (NVRAM) storage for RAID configuration data
- Audible alarm
- Direct memory access (DMA) chaining support
- Separate DRAM bus
- Support for differential or single-ended SCSI with active termination
- Up to 12 MegaRAID 320-2 cards per system
- Support for up to 15 SCSI devices per channel
- Support for RAID levels 0, 1, 5, 10, and 50
- Support for scatter/gather and tagged command queuing
- Ability to multithread up to 256 commands simultaneously
- Support for multiple rebuilds and consistency checks with transparent user-definable priority setting
- Support for variable stripe sizes for all logical drives
- Automatic detection of failed drives
- Automatic and transparent rebuild of hot spare drives
- Hot swap of new drives without taking the system down
- Optional battery backup for up to 72 hours of data retention
- Server clustering support
- Optional firmware provides multi-initiator support
- Server failover
- Software drivers for major operating systems

4.1 SMART Technology

The MegaRAID Self Monitoring Analysis and Reporting Technology (SMART) feature detects up to 70% of all predictable drive failures. SMART monitors the internal performance of all motors, heads, and drive electronics. You can recover from drive failures through online physical drive migration.

4.2 Configuration on Disk

Configuration on Disk (drive roaming) saves configuration information both in nonvolatile random access memory (NVRAM) on the MegaRAID SCSI 320-2, and on the disk drives it controls. If the MegaRAID SCSI 320-2 is replaced, the new MegaRAID controller can detect the actual RAID configuration, maintaining the integrity of the data on each drive, even if the drives have changed channel and/or target ID.

4.3 Configuration Features

[Table 4.1](#) contains the configuration features for the MegaRAID 320-2.

Table 4.1 Configuration Features

Specification	Feature
RAID levels	0, 1, 5, 10, and 50
SCSI channels	2
Maximum number of drives per channel	15
Array interface to host	PCI 2.2
PCI bus master	Supports write invalidate
Drive interface	Wide Ultra320
Upgradable cache memory sizes	32, 64, 128, or 256 Mbytes
Cache function	Write-back, Write-through, Adaptive Read Ahead, Non Read Ahead, Read Ahead

Table 4.1 Configuration Features (Cont.)

Specification	Feature
Multiple logical drives/arrays per controller	Up to 40 logical drives per controller
Maximum number of MegaRAID controllers per system	12
Online capacity expansion	Yes
Hot spare support	Yes
Flashable firmware	Yes
Hot swap devices supported	Yes
Non-disk devices supported	Yes
Mixed capacity hard drives	Yes
Number of internal SCSI connectors	2
Number of external SCSI connectors	2
Support for hard drives with capacities of more than 8 Gbytes	Yes
Clustering support (Failover control)	Yes
Online RAID level migration	Yes
No reboot necessary after expansion	Yes
More than 200 Qtags per physical drive	Yes
Hardware clustering support on the board	Yes
User-specified rebuild rate	Yes

4.4 Array Performance Features

Table 4.2 lists the array performance features.

Table 4.2 Array Performance Features

Specification	Feature
Host data transfer rate	532 Mbytes/s
Drive data transfer rate	320 Mbytes/s
Stripe sizes	2, 4, 8, 16, 32, 64, or 128 Kbytes

4.5 RAID Management Features

Table 4.3 lists the RAID management features.

Table 4.3 RAID Management Features

Specification	Feature
Support for SNMP	Yes
Performance Monitor provided	Yes
Remote control and monitoring	Yes
Drive roaming	Yes
Support for concurrent multiple stripe sizes	Yes
Windows NT, 2000, XP, and .NET server support using a GUI client utility	Yes

4.6 Fault Tolerance Features

Table 4.4 lists the fault tolerance features.

Table 4.4 Fault Tolerance Features

Specification	Feature
Support for SMART	Yes
Optional battery backup for cache memory	Standard. Provided on the MegaRAID Controller. Up to 72 hours data retention
Enclosure management	SCSI-accessed fault-tolerant enclosure (SAF-TE) compliant
Drive failure detection	Automatic
Drive rebuild using hot spares	Automatic and transparent
Parity generation and checking	Software and hardware

4.7 Software Utilities

Table 4.5 lists the software utility features.

Table 4.5 Software Utilities

Specification	Feature
FlexRAID reconfiguration on the fly	Yes
FlexRAID RAID level migration on the fly	Yes
FlexRAID online capacity expansion	Yes
Remote configuration and management over the Internet	Yes
Graphical user interface	Yes
Diagnostic utility	Yes
Management utility	Yes
Bootup configuration using MegaRAID Manager	Yes
Online read, write, and cache policy switching	Yes

4.8 Operating System Software Drivers

MegaRAID SCSI 320-2 includes a DOS software configuration utility, and drivers for:

- Windows NT 4.0
- Windows 2000
- Windows .NET
- Windows XP
- Novell NetWare 5.1, 6.0
- Red Hat Linux 7.2, 7.3
- DOS version 6.xx or later

The DOS drivers for MegaRAID are contained in the firmware on the MegaRAID controller, except for the DOS ASPI and CD drivers. Call your LSI OEM support representative or access the web site at www.lsilogic.com for information about drivers for other operating systems.

4.9 MegaRAID SCSI 320-2 Specifications

Table 4.6 lists the specifications for the SCSI 320-2.

Table 4.6 MegaRAID SCSI 320-2 Specifications

Parameter	Specification
Card size	6.875" x 4.2" (half length PCI)
Processor	Intel 80303 @ 100 MHz
SCSI processor	One LSI Logic 53C1030 SCSI controller
Bus type	PCI 2.2
Bus data transfer rate	Up to 266 Mbytes/s
Cache configuration	64 or 128Mbyte DIMM
Firmware	1 MB 8 flash ROM

Table 4.6 MegaRAID SCSI 320-2 Specifications (Cont.)

Parameter	Specification
Nonvolatile RAM	32 KB 8 for storing RAID configuration
Memory type	One 72-bit 168-pin SDRAM DIMM socket provides write-through or write-back caching on a logical drive basis. It also provides adaptive read-ahead.
Operating voltage	5.00 V 0.25 V and 3.3 V +/- 0.3 V
SCSI controller	2 SCSI controllers for Ultra320 and Wide support
SCSI data transfer rate	Up to 320 Mbytes/s
SCSI bus	Low voltage differential or SE
SCSI termination	Active, low-voltage differential, or single-ended
Termination disable	Automatic through cable detection
Devices per SCSI channel	Up to 15 wide or seven non-wide SCSI devices. Up to 6 non-disk SCSI drives per MegaRAID controller.
SCSI device types supported	Synchronous or Asynchronous. Disk and non-disk.
RAID levels supported	0, 1, 5, 10, and 50
SCSI connectors	Two 68-pin internal high-density connectors for 16-bit SCSI devices. Two ultra-high density 68-pin external connectors
SCSI cables	Up to 25 m if using low voltage differential
Serial port	3-pin RS232C-compatible connector, used for test purposes only

4.10 MegaRAID Components

4.10.1 CPU

The MegaRAID SCSI 320-2 controller uses the 64-bit Intel 80303 Intelligent I/O processor with an embedded 32-bit 80960 Jx RISC processor that runs at 100 MHz. This processor directs all functions of the controller including command processing, PCI and SCSI bus

transfers, RAID processing, drive rebuilding, cache management, and error recovery.

4.10.2 Cache Memory

Cache memory resides in a single 72-bit DIMM socket that requires one unbuffered 3.3 V SDRAM single-sided or double-sided DIMM. Possible configurations are 32, 64, 128, or 256 Mbytes.

MegaRAID supports write-through or write-back caching, which can be selected for each logical drive. To improve performance in sequential disk accesses, MegaRAID does not use read-ahead caching for the current logical drive. The default setting for the read policy is *Normal*, meaning no read-ahead caching. You can disable read-ahead caching.

Warning: Write caching is not recommended for the physical drives. When write cache is enabled, loss of data can occur when power is interrupted.

4.10.3 MegaRAID BIOS

The BIOS resides on a 1 MB or 2 MB 8 flash ROM for easy upgrade. The MegaRAID BIOS supports INT 13h calls to boot DOS without special software or device drivers. The MegaRAID BIOS provides an extensive setup utility that can be accessed by pressing <Ctrl> <M> at BIOS initialization. The MegaRAID Configuration Utility is described in the *MegaRAID Configuration Software Guide*.

4.10.4 Onboard Speaker

The MegaRAID SCSI 320-2 has an onboard tone generator for audible warnings when system errors occur. Audible warnings can be generated through this speaker. The audible warnings are listed in Appendix B.

4.10.5 Serial Port

The MegaRAID SCSI 320-2 has a 3-pin RS232C-compatible serial port connector, which is used for test purposes only.

4.10.6 SCSI Bus

The MegaRAID SCSI 320-2 controller has two Ultra320 Wide SCSI channels that support low voltage differential SCSI devices with active termination. Both synchronous and asynchronous devices are supported. The MegaRAID controller provides automatic termination disable via cable detection. Each channel supports up to 15 wide or seven non-wide SCSI devices at speeds up to 320 Mbytes/s per SCSI channel. The MegaRAID controller supports up to six non-disk devices per controller.

The SCSI bus mode defaults to LVD for each SCSI channel. If a single-ended device is attached to a SCSI channel, the MegaRAID controller automatically switches to SE mode for that SCSI channel.

4.10.7 SCSI Connectors

The MegaRAID SCSI 320-2 has two types of SCSI connectors:

- Two 68-pin high density internal connectors
- Two 68-pin very-high-density external connectors

4.10.8 SCSI Termination

The MegaRAID SCSI 320-2 uses active termination on the SCSI bus, conforming to Alternative 2 of the SCSI-2 specifications. Termination enable/disable is automatic through cable detection.

4.10.9 SCSI Firmware

The MegaRAID SCSI 320-2 firmware handles all RAID and SCSI command processing and also supports the features listed in [Table 4.7](#).

Table 4.7 SCSI Firmware

Feature	Description
Disconnect/reconnect	Optimizes SCSI bus seek
Tagged command queuing	Multiple tags to improve random access
Scatter/gather	Multiple address/count pairs

Table 4.7 SCSI Firmware (Cont.)

Feature	Description
Multi-threading	Up to 255 simultaneous commands with elevator sorting and concatenation of requests per SCSI channel
Stripe size	Variable for all logical drives: 2, 4, 8, 16, 32, 64, or 128 Kbytes
Rebuild	Multiple rebuilds and consistency checks, with user-definable priority

4.11 RAID Management

RAID management is provided by software utilities that manage and configure the RAID system and MegaRAID SCSI 320-2, create and manage multiple disk arrays, control and monitor multiple RAID servers, provide error statistics logging, and provide online maintenance. They include:

- MegaRAID BIOS Configuration Utility
- WebBIOS Configuration Utility
- Power Console Plus
- MegaRAID Manager

4.11.1 MegaRAID BIOS Configuration Utility

The BIOS Configuration Utility (<Ctrl><M>) is used to configure and maintain RAID arrays, format hard drives, and manage the RAID system. It is independent of any operating system. See the *MegaRAID Configuration Software Guide* for additional information.

4.11.2 WebBIOS Configuration Utility

The WebBIOS Configuration Utility is an HTML-based utility used to configure and maintain RAID arrays, format hard drives, and manage the RAID system. It is independent of any operating system. See the *MegaRAID Configuration Software Guide* for additional information.

4.11.3 Power Console Plus

Power Console Plus runs in Windows NT, 2000, XP, and .NET. It configures, monitors, and maintains multiple RAID servers from any network node or a remote location. See the *MegaRAID Configuration Software Guide* for additional information.

4.11.4 MegaRAID Manager

MegaRAID Manager is a character-based, non-GUI utility for Linux and Novell NetWare that changes policies and parameters, and monitors RAID systems. See the *MegaRAID Configuration Software Guide* for additional information.

4.12 Compatibility

MegaRAID SCSI 320-2 compatibility issues include:

- Server management
- SCSI device compatibility
- Software compatibility

4.12.1 Server Management

As a simple network management protocol (SNMP) agent, MegaRAID SCSI 320-2 supports all SNMP managers.

4.12.2 SCSI Device Compatibility

The MegaRAID SCSI 320-2 supports SCSI hard drives, CD-ROM drives, and tape drives.

4.12.3 Software

All SCSI backup and utility software should work with the MegaRAID SCSI 320-2. This software is not provided with the MegaRAID controller.

Chapter 5

Configuring Physical Drives, Arrays, and Logical Drives

This chapter explains how to configure SCSI physical drives, arrays, and logical drives connected to the MegaRAID SCSI 320-2 controller. It contains the following sections:

- [Section 5.1, “Configuring SCSI Physical Drives”](#)
 - [Section 5.2, “Configuring Arrays”](#)
 - [Section 5.3, “Creating Logical Drives”](#)
 - [Section 5.4, “Configuring Logical Drives”](#)
 - [Section 5.5, “Planning the Array Configuration”](#)
-

5.1 Configuring SCSI Physical Drives

SCSI physical drives must be organized into logical drives. The arrays and logical drives that you construct must be able to support the RAID level that you select. The MegaRAID SCSI 320-2 controller has two SCSI channels.

5.1.1 Distributing Drives

Distribute the disk drives across all channels for optimal performance. It is best to stripe across channels instead of down channels. Performance is most affected for sequential reads and writes. The MegaRAID SCSI 320-2 controller supports SCSI CD-ROM drives, SCSI tape drives, and other SCSI devices as well as SCSI hard disk drives. For optimal performance, all non-disk SCSI devices should be attached to one SCSI channel.

5.1.2 Basic Configuration Rules

You should observe the following guidelines when connecting and configuring SCSI devices in a RAID array:

- Attach non-disk SCSI devices to a single SCSI channel that does not have any disk drives.
- Distribute the SCSI hard disk drives equally among all available SCSI channels except any SCSI channel that is being reserved for non-disk drives.
- You can place up to 30 physical disk drives in a logical array, depending on the RAID level.
- An array can contain SCSI devices that reside on an array on any channel.
- Include all drives that have the same capacity in the same array.
- Make sure any hot spare has a capacity that is at least as large as the largest drive that may be replaced by the hot spare.
- When replacing a failed drive, make sure that the replacement drive has a capacity that is at least as large as the drive being replaced.

Note: Be sure to back up your data regularly, even when using RAID.

5.1.3 Current Physical Device Configuration

Use [Table 5.1](#) to record the current configuration for your physical devices on SCSI channels 0 and 1.

Table 5.1 Physical Device Configuration

SCSI ID	Device Description	Termination?
SCSI Channel 0		
0		
1		
2		
3		

Table 5.1 Physical Device Configuration (Cont.)

SCSI ID	Device Description	Termination?
4		
5		
6		
8		
9		
10		
11		
12		
13		
14		
15		
SCSI Channel 1		
0		
1		
2		
3		
4		
5		
6		
8		
9		
10		
11		
12		
13		

Table 5.1 Physical Device Configuration (Cont.)

SCSI ID	Device Description	Termination?
14		
15		

5.1.4 Logical Drive Configuration

Use [Table 5.2](#) to record the configuration for your logical drives.

Table 5.2 Logical Drive Configuration

Logical Drive	RAID Level	Stripe Size	Logical Drive Size	Cache Policy	Read Policy	Write Policy	# of Physical Drives
LD0							
LD1							
LD2							
LD3							
LD4							
LD5							
LD6							
LD7							
LD8							
LD9							
LD10							
LD11							
LD12							
LD13							
LD14							
LD15							
LD16							

Table 5.2 Logical Drive Configuration (Cont.)

Logical Drive	RAID Level	Stripe Size	Logical Drive Size	Cache Policy	Read Policy	Write Policy	# of Physical Drives
LD17							
LD18							
LD19							
LD20							
LD21							
LD22							
LD23							
LD24							
LD25							
LD26							
LD27							
LD28							
LD29							
LD30							
LD31							
LD32							
LD33							
LD34							
LD35							
LD36							
LD37							
LD38							
LD39							

5.1.5 Physical Device Layout

Use [Table 5.3](#) to record the physical device layout.

Table 5.3 Physical Device Layout

	Channel 0	Channel 1
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		

Table 5.3 Physical Device Layout (Cont.)

	Channel 0	Channel 1
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		

Table 5.3 Physical Device Layout (Cont.)

	Channel 0	Channel 1
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		

Table 5.3 Physical Device Layout (Cont.)

	Channel 0	Channel 1
Target ID		
Device type		
Logical drive number/Drive number		
Manufacturer/Model number		
Firmware level		

5.2 Configuring Arrays

Connect the physical drives to the MegaRAID SCSI 320-2, configure the drives, then initialize them. The number of physical disk drives that an array can support depends on the firmware version.

For the MegaRAID SCSI 320-2, an array can consist of up to 30 physical disk drives, depending on the RAID level (see Chapter 3 for more information). This controller supports up to 40 logical drives per controller. The number of drives in an array determines the RAID levels that can be supported.

5.2.1 Arranging Arrays

You must arrange the arrays to provide additional organization for the drive array. You must arrange arrays so that you can create system drives that can function as boot devices.

You can sequentially arrange arrays with an identical number of drives so that the drives in the group are spanned. Spanned drives can be treated as one large drive. Data can be striped across multiple arrays as one logical drive.

You can create spanned drives by using the MegaRAID Configuration Utility or the MegaRAID Manager. See the *MegaRAID Configuration Software Guide* for additional information.

5.2.2 Creating Hot Spares

Any drive that is present, formatted, and initialized, but is not included in an array or logical drive is automatically designated as a hot spare.

You can also designate drives as hot spares using the MegaRAID BIOS Configuration Utility, the MegaRAID Manager, or Power Console Plus. See the *MegaRAID Configuration Software Guide* for additional information.

5.3 Creating Logical Drives

Logical drives are arrays or spanned arrays that are presented to the operating system. You must create one or more logical drives.

The logical drive capacity can include all or any portion of an array. The logical drive capacity can also be larger than an array by using spanning. The MegaRAID SCSI 320-2 supports up to 40 logical drives.

5.3.1 Configuration Strategies

The most important factors in RAID array configuration are: drive capacity, drive availability (fault tolerance), and drive performance. You cannot configure a logical drive that optimizes all three factors, but it is easy to choose a logical drive configuration that maximizes one factor at the expense of the other two factors, although needs are seldom that simple.

5.3.1.1 Maximize Capacity

RAID 0 achieves maximum drive capacity, but does not provide data redundancy. Maximum drive capacity for each RAID level is shown below. OEM-level firmware that can span up to 4 logical drives is

assumed. [Table 5.4](#) describes the RAID levels, including the number of drives required, and the capacity.

Table 5.4 Capacity for RAID Levels

RAID Level	Description	Drives Required	Capacity
0	Striping without parity	1 – 30	(Number of disks) X capacity of smallest disk
1	Mirroring	2	(Capacity of smallest disk) X (1)
5	Striping with floating parity drive	3 – 30	(Number of disks) X (capacity of smallest disk) - (capacity of 1 disk)
10	Mirroring and Striping	4 – 30 (Must be a multiple of 2.)	(Number of disks) X (capacity of smallest disk) / (2)
50	RAID 5 and Striping	6 – 30 (Must be a multiple of the # of arrays.)	(Number of disks) X (capacity of smallest disk) – (capacity of 1 disk X number of arrays)

5.3.1.2 Maximizing Drive Availability

You can maximize the availability of data on the physical disk drive in the logical array by maximizing the level of fault tolerance. [Table 5.5](#) describes the levels of fault tolerance for the RAID levels.

Table 5.5 Fault Tolerance for RAID Levels

RAID Level	Fault Tolerance Protection
0	No fault tolerance.
1	100% protection through data mirroring.
5	100% protection through striping and parity. The data is striped and parity data is written across a number of physical disk drives.
10	100% protection through data mirroring.
50	100% protection through data striping and parity. All data is striped and parity data is written across all drives in two or more arrays.

5.3.1.3 Maximizing Drive Performance

You can configure an array for optimal performance. But optimal drive configuration for one type of application will probably not be optimal for any other application. [Table 5.6](#) lists basic guidelines for the performance characteristics for RAID drive arrays at each RAID level.

Table 5.6 Performance Characteristics for RAID Levels

RAID Level	Performance Characteristics
0	Excellent for all types of I/O activity, but provides no data security.
1	Provides data redundancy and good performance.
5	Provides data redundancy and good performance in most environments.
10	Provides data redundancy and excellent performance.
50	Provides data redundancy and very good performance.

5.3.2 Assigning RAID Levels

Only one RAID level can be assigned to each logical drive. [Table 5.7](#) shows the drives required per RAID level.

Table 5.7 Number of Physical Drives per RAID Level

RAID Level	Minimum # of Physical Drives	Maximum # of Physical Drives
0	1	30
1	2	2
5	3	30
10	4	30
50	6	30

5.4 Configuring Logical Drives

After you have installed the MegaRAID SCSI 320-2 controller in the server and have attached all physical disk drives, perform the following actions to prepare a RAID array:

1. Optimize the MegaRAID SCSI 320-2 controller options for your system. See Chapter 3 for additional information.
2. Press <Ctrl><M> to run the MegaRAID Manager.
3. If necessary, perform a low-level format of the SCSI drives that will be included in the array and the drives to be used for hot spares.
4. Customize the RAID array and define and configure one or more logical drives by selecting Easy Configuration or New Configuration.
5. Create and configure one or more system drives (logical drives) by selecting the RAID level, cache policy, read policy, and write policy.
6. Save the configuration.
7. Initialize the system drives.

After initialization, you can install the operating system.

5.4.1 Optimizing Data Storage

5.4.1.1 Data Access Requirements

Each type of data stored in the disk subsystem has a different frequency of read and write activity. If you know the data access requirements, you can more successfully determine a strategy for optimizing the disk subsystem capacity, availability, and performance.

Servers that support Video on Demand typically read the data often, but write data infrequently. Both the read and write operations tend to be long. Data stored on a general-purpose file server involves relatively short read and write operations with relatively small files.

5.4.1.2 Array Functions

You must first define the major purpose of the disk array. Will this disk array increase the system storage capacity for general-purpose file and print servers? Does this disk array support any software system that

must be available 24 hours per day? Will the information stored in this disk array contain large audio or video files that must be available on demand? Will this disk array contain data from an imaging system? You must identify the purpose of the data to be stored in the disk subsystem before you can confidently choose a RAID level and a RAID configuration.

5.5 Planning the Array Configuration

Fill out [Table 5.8](#) to help plan this array.

Table 5.8 Factors for Planning the Array Configuration

Question	Answer
Number of MegaRAID SCSI channels	
Number of physical disk drives in the array	
<hr/>	
Purpose of this array. Rank the following factors:	
Maximize drive capacity	
Maximize the safety of the data (fault tolerance)	
Maximize hard drive performance and throughput	
<hr/>	
Number of hot spares	
Amount of cache memory installed on MegaRAID SCSI 320-1	
Are all of the disk drives and the server protected by an uninterruptible power supply?	

5.5.1 Using the Array Configuration Planner

[Table 5.9](#) lists the possible RAID levels, fault tolerance, and effective capacity for all possible drive configurations for an array consisting of one to eight drives. This table does not take into account any hot spare (standby) drives. You should always have a hot spare drive in case of drive failure. RAID 1 requires two drives, RAID 10 at least four, and RAID 50 at least six.

Use [Table 5.9](#) to plan the array configuration.

Table 5.9 Array Configuration Planner

# of Drives	Possible RAID Levels	Relative Performance	Fault Tolerance	Effective Capacity
1	RAID 0	Excellent	No	100%
2	RAID 0	Excellent	No	100%
2	RAID 1	Good	Yes	50%
3	RAID 0	Excellent	No	100%
3	RAID 5	Good	Yes	67%
4	RAID 0	Excellent	No	100%
4	RAID 5	Good	Yes	75%
4	RAID 10	Excellent	Yes	50%
5	RAID 0	Excellent	No	100%
5	RAID 5	Good	Yes	80%
6	RAID 0	Excellent	No	100%
6	RAID 5	Good	Yes	83%
6	RAID 10	Excellent	Yes	50%
6	RAID 50	Good	Yes	67%
7	RAID 0	Excellent	No	100%
7	RAID 5	Good	Yes	86%
8	RAID 0	Excellent	No	100%
8	RAID 5	Good	Yes	87%
8	RAID 10	Excellent	Yes	50%
8	RAID 50	Good	Yes	75%

Chapter 6

Hardware Installation

This chapter explains how to install the MegaRAID SCSI 320-2 controller. It contains the following sections:

- [Section 6.1, "Hardware Requirements"](#)
 - [Section 6.2, "Installation Steps"](#)
 - [Section 6.3, "Summary"](#)
-

6.1 Hardware Requirements

You must have the following in order to install the MegaRAID SCSI 320-1 controller and create arrays:

- A host computer with the following:
 - A motherboard with 5 V/3.3 V PCI expansion slots that has an available expansion slot
 - Support for PCI version 2.2 or later
 - Intel Pentium, Pentium Pro, or more powerful CPU
 - Floppy drive, color monitor, VGA adapter card, mouse, and keyboard
- The MegaRAID SCSI 320-2 Installation CD
- The necessary internal and/or external SCSI cables and terminators (this depends on the number and type of SCSI devices to be attached)
- An Uninterruptible Power Supply (UPS) for the entire system
- Ultra320 SCSI disk drives and other SCSI devices, as desired

Important: The MegaRAID SCSI 320-2 controller must be installed in a PCI expansion slot.

6.1.1 Optional Equipment

You may also want to install SCSI cables that interconnect the MegaRAID SCSI 320-2 to external SCSI devices.

6.2 Installation Steps

The MegaRAID SCSI 320-2 provides extensive customization options. If you need only basic MegaRAID SCSI 320-2 features and your computer does not use other adapter cards with resource settings that may conflict with MegaRAID SCSI 320-2 settings, even custom installation can be quick and easy.

[Table 6.1](#) lists the hardware installation steps. Each step is described in detail in the following pages.

Table 6.1 Hardware Installation Steps

Step	Action	Additional Information
1	Unpack the MegaRAID controller and inspect for damage. Make sure all items are in the package.	If damaged, call your LSI Logic OEM support representative.
2	Turn the computer off and remove the power cord, and remove the cover.	
3	Install cache memory on the MegaRAID card.	32 Mbytes minimum cache memory is required.
4	Check the jumper settings on the MegaRAID SCSI 320-2 controller.	See page 6-5 for the MegaRAID SCSI 320-2 jumper settings.
5	Set SCSI termination.	
6	Set SCSI terminator power (TermPWR).	
7	Connect the battery harness.	Optional
8	Install the MegaRAID SCSI 320-2 card.	
9	Connect the SCSI cables to SCSI devices.	
10	Set the target IDs for the SCSI devices.	

Table 6.1 Hardware Installation Steps (Cont.)

Step	Action	Additional Information
11	Replace the computer cover and turn the power on.	Be sure the SCSI devices are powered up before, or at the same time as, the host computer.
12	Run the MegaRAID BIOS Configuration Utility.	Optional.
13	Install software drivers for the desired operating systems.	

6.2.1 Step 1: Unpack

Unpack and install the hardware in a static-free environment. Remove the MegaRAID SCSI 320-2 controller card from the anti-static bag and inspect it for damage. If the card appears damaged, or if any item listed below is missing, contact LSI Logic or your MegaRAID OEM support representative. The MegaRAID SCSI 320-2 controller is shipped with the following:

- The Driver and Documentation CD, which contains these items:
 - The *MegaRAID Configuration Software Guide*
 - The *MegaRAID Operating System Driver Installation Guide*
 - The *MegaRAID SCSI 320-2 Hardware Guide*
 - The software license agreement
 - The MegaRAID configuration utilities for DOS
- The warranty registration card

6.2.2 Step 2: Power Down


Turn off the computer, remove the power cord, then remove the cover. Make sure the computer is turned off and disconnected from any networks before installing the controller card.

6.2.3 Step 3: Install Cache Memory

Important: A minimum of 32 Mbytes of cache memory is required. The cache memory must be installed before the MegaRAID controller is operational.

Install cache memory DIMMs on the MegaRAID controller card in the cache memory socket. Use a 72-bit 3.3 V single-sided or double-sided 168-pin unbuffered DIMM. Lay the controller card component-side up on a clean static-free surface. The memory socket is mounted flush with the MegaRAID card, so the DIMM is parallel to the MegaRAID card when properly installed. The DIMM clicks into place, indicating proper seating in the socket. The MegaRAID card is shown lying on a flat surface below.

Figure 6.1 Installing DIMM Memory



6.2.3.1 Installing or Changing Memory

Important: The battery pack harness or cable must be disconnected from J10 on the MegaRAID SCSI 320-2 Ultra320 card before you add or remove memory.

Perform the following steps to install or change memory:

1. Bring down the operating system properly. Make sure that cache memory has been flushed.

You must perform a system reset if operating under DOS. When the computer reboots, the MegaRAID controller will flush cache memory.

2. Turn the computer power off. Disconnect the power cables from the computer.
3. Remove the computer cover.
4. Disconnect the battery pack cable from the MegaRAID controller.
5. Remove the MegaRAID controller.

6. You can now add or remove DRAM modules from the MegaRAID controller, as described above.
7. Reattach the battery pack harness to J10 on the MegaRAID controller.
8. Reinstall the MegaRAID controller in the computer. Follow the instructions in this chapter.
9. Replace the computer cover and turn the computer power on.

6.2.3.2 Recommended Memory Vendors

Call LSI Logic Technical Support at 678-728-1250 for a current list of recommended memory vendors.

6.2.4 Step 4: Check Jumper Settings

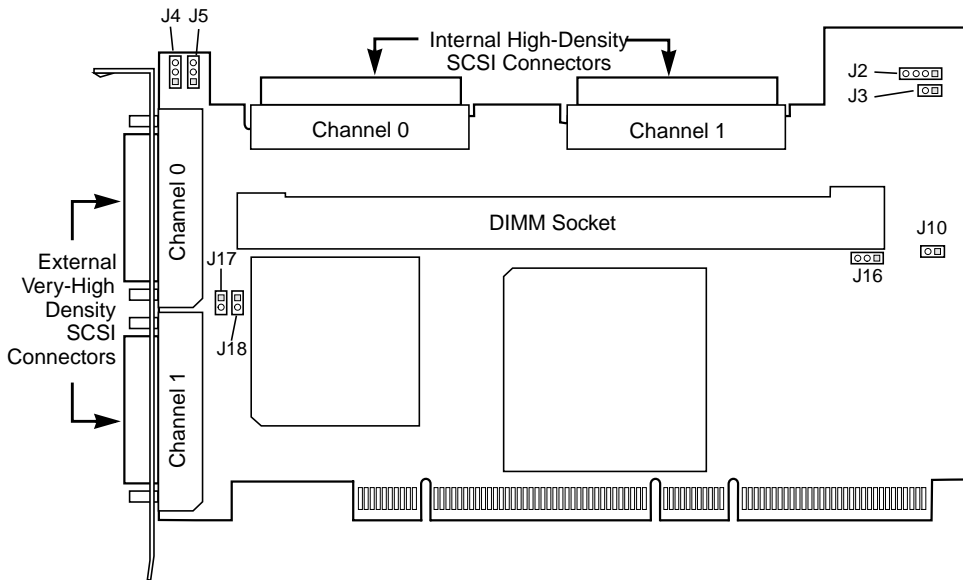
Make sure the jumper settings on the MegaRAID SCSI 320-2 card are correct. The jumpers are set at the factory and you probably do not need to change them. [Table 6.2](#) lists the jumpers that you should check.

Table 6.2 Jumpers for the MegaRAID SCSI 320-2

Item	Description	Type
J2	SCSI activity LED connector	4-pin header
J3	Write pending indicator	2-pin header
J4	SCSI Termination Enable Channel 0	3-pin header
J5	SCSI Termination Enable Channel 1	3-pin header
J10	Battery connector	3-pin header
J16	BIOS Enable	2-pin header
J17	Termination Power Channel 0	2-pin header
J18	Termination Power Channel 1	2-pin header

[Figure 6.2](#) shows the location of these items on the MegaRAID SCSI 320-2 controller.

Figure 6.2 MegaRAID SCSI 320-2 Controller Layout



6.2.4.1 J2 SCSI Activity LED

J2 is a four-pin connector that attaches to a cable that connects to the hard disk LED mounted on the computer enclosure. The LED indicates data transfers (SCSI bus activity.)

6.2.4.2 J3 Dirty Cache LED

J3 is a two-pin header for the dirty cache LED. This can be connected to an LED on the computer enclosure. The LED will be lit when data in the cache has not yet been written to the storage device.

6.2.4.3 J4/J5 Termination Enable

J4 and J5 are three-pin headers that specify the type of SCSI termination for channel 0 and channel 1, respectively. [Table 6.3](#) shows the jumper

settings. Leave at the default setting (jumper on pins 1 and 2) to allow the MegaRAID controller to automatically set its own SCSI termination.

Table 6.3 Pinout for J4/J5 Termination Enable

Type of SCSI Termination	J4/J5 Setting
Software control of SCSI termination using drive detection (default).	Short pins 1-2
Permanently disable all onboard SCSI termination.	Short pins 2-3
Permanently enable all onboard SCSI termination.	OPEN

6.2.4.4 J10 Battery Pack Connector

J10 is a 3-pin connector that attaches to the optional battery pack. [Table 6.4](#) displays the J10 pinout.

Table 6.4 J10 External Battery Pinout

Pin	Signal Description
3	+BATT Terminal (red wire)
2	Thermistor (white wire)
1	-BATT Terminal (black wire)

6.2.4.5 J16 Onboard BIOS Enable

J16 is a 2-pin header that enables or disables the MegaRAID onboard BIOS. The onboard BIOS should be enabled (J16 unjumpered) for normal board operation. [Table 6.5](#) shows the jumper settings.

Table 6.5 Pinout for J16 BIOS Enable

J16 Setting	Onboard BIOS Status
Unjumpered	Enabled
Jumpered	Disabled

6.2.4.6 J17/J18 SCSI Bus Termination Power

J17 and J18 are 2-pin jumpers that control the termination power setting for channel 0 and channel 1, respectively. Leave both jumpers at the default setting (jumper installed on pins 1 and 2) to allow the PCI bus to provide termination power. (When the jumpers are removed, the SCSI bus provides termination power.)

6.2.5 Step 5: Set Termination

Each MegaRAID SCSI channel can be individually configured for termination enable mode by setting the J4 and J5 jumpers. You must terminate the SCSI bus properly. Set termination at both ends of the SCSI cable. The SCSI bus is an electrical transmission line and must be terminated properly to minimize reflections and losses.

For a disk array, set SCSI bus termination so that removing or adding a SCSI device does not disturb termination. An easy way to do this is to connect the card to one end of the SCSI cable and to connect a terminator module at the other end of the cable. The connectors between the two ends can connect SCSI devices. Disable termination on the SCSI devices. See the manual for each SCSI device to disable termination.

6.2.5.1 SCSI Termination

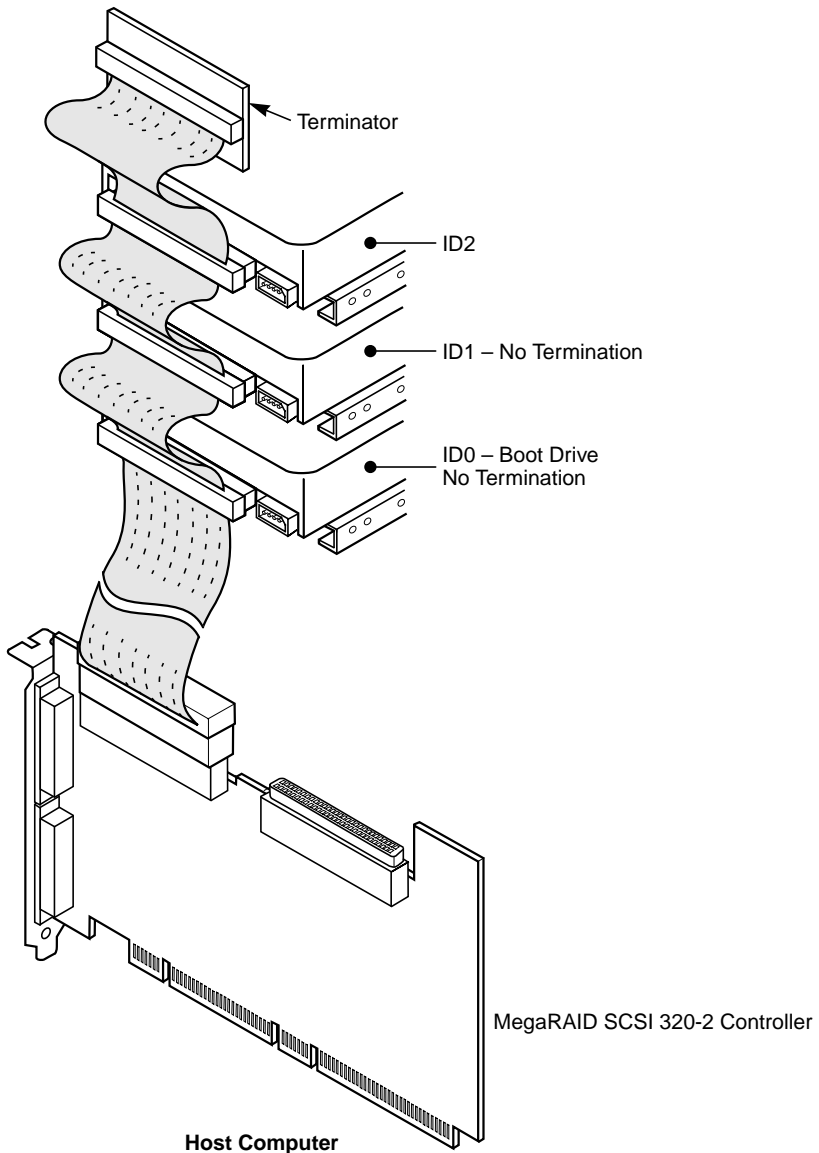
You can let the card automatically provide SCSI termination at one end of the SCSI bus. You can terminate the other end of the SCSI bus by attaching an external SCSI terminator module to the end of the cable or by attaching a SCSI device that internally terminates the SCSI bus at the end of the SCSI channel.

Use standard external SCSI terminators on a SCSI channel operating at 10 Mbytes/s or higher synchronous data transfer.

6.2.5.2 Terminating Internal SCSI Disk Arrays

Set the termination so that SCSI termination and termination power are intact when any hard drive is removed from a SCSI channel, as shown in [Figure 6.3](#). Also, jumpers J4 and J5 should always be set to allow the termination to be controlled by software. (See [Section 6.2.4.3, "J4/J5 Termination Enable,"](#) page 6-6.)

Figure 6.3 Termination of Internal SCSI Disk Arrays




6.2.5.3 Terminating External Disk Arrays

In most array enclosures, the end of the SCSI cable has an independent SCSI terminator module that is not part of any SCSI drive. In this way,

SCSI termination is not disturbed when any drive is removed, as shown in [Figure 6.4](#):


Figure 6.4 Terminating External Disk Arrays



6.2.5.4 Terminating Internal and External Disk Arrays

You can use both internal and external drives with the MegaRAID SCSI 320-2. You still must make sure that the proper SCSI termination and termination power is preserved, as shown in [Figure 6.5](#):

Figure 6.5 Terminating Internal and External Disk Arrays



Note: *Termination enabled from last SCSI drive

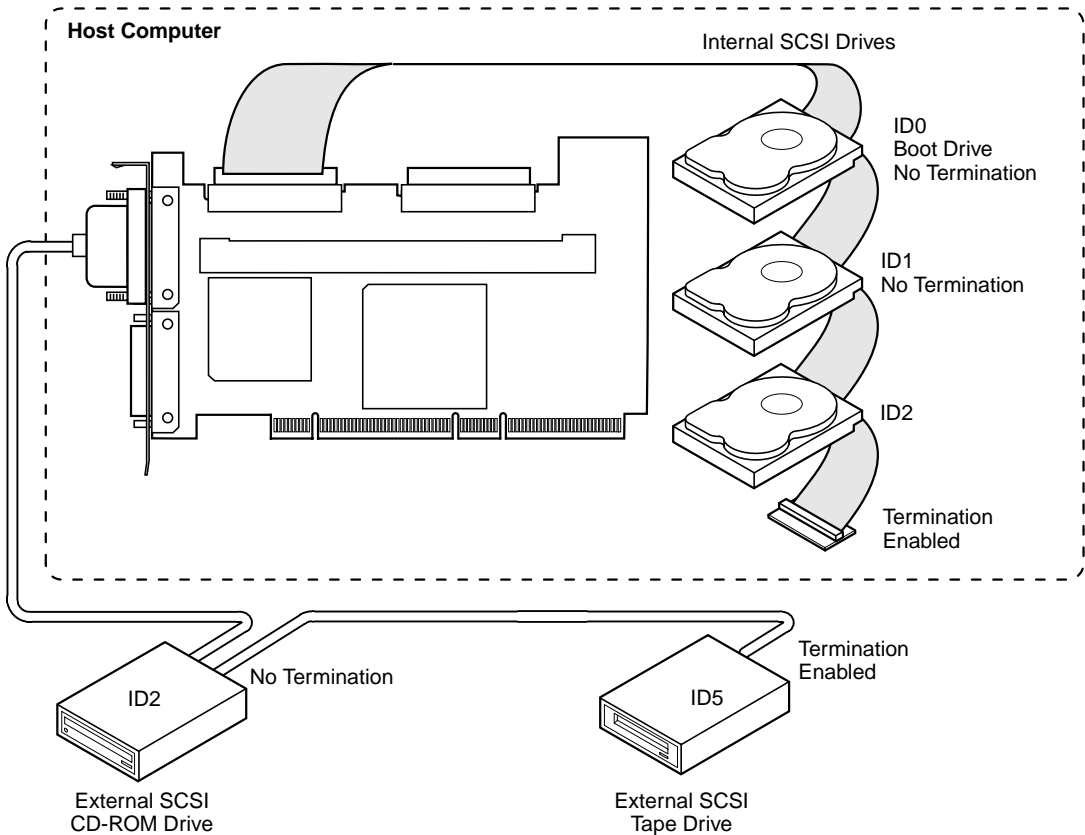
6.2.5.5 Connecting Non-Disk SCSI Devices

SCSI tape drives and SCSI CD-ROM drives must each have a unique SCSI ID regardless of the SCSI channel they are attached to. The general rule for Unix systems is:

- Tape drive set to SCSI ID 2
- CD-ROM drive set to SCSI ID 5

Make sure that no hard drives are attached to the same SCSI Channel as the non-disk SCSI devices. Drive performance will be significantly degraded if SCSI hard disk drives are attached to this channel.

Figure 6.6 Connecting Non-Disk SCSI Devices



6.2.6 Step 6 Set SCSI Terminator Power

J17 and J18 control the termination power setting for the MegaRAID 320-2 SCSI channels, as explained in [Section 6.2.4.6, "J17/J18 SCSI Bus Termination Power," page 6-8](#). By default (jumper installed on pins 1 and 2 of J17 and J18), the PCI bus supplies termination power.

See the documentation for each SCSI device for information about enabling TermPWR.

Important: The SCSI channels need Termination power to operate. If a channel is not being used and no auxiliary power source is connected, change the jumper setting for that channel to supply TermPWR from the PCI bus.


6.2.7 Step 7: Install Battery Pack (Optional)

You can install a battery pack on the RAID controller using the J10 connector. The battery pack provides emergency power to the MegaRAID 320-2 SCSI controller in the event of power failure.

- **LSI Logic Part Number:** BAT-NIMH-3.6-02
- **Description:** Battery, NIMH, 3.6 V, 600 mA onboard battery pack with mounting brackets

You install the battery pack by attaching it to the controller board through the backside of the board, using the three holes in the board. Then connect the three wires from the battery pack to J10, the external battery connector (the black wire connects to pin 1). [Figure 6.7](#) shows an installed battery pack (upper right corner).

Figure 6.7 MegaRAID Controller with Backup Battery Module



6.2.7.1 Configuring the Battery Pack

After you install the MegaRAID controller you must configure the battery pack in the BIOS Configuration Utility. To do this, you select the Objects menu, then select Battery Backup. [Table 6.6](#) explains the battery backup menu options.

Table 6.6 Backup Battery Menu Options

Menu Item	Explanation
Battery Pack	PRESENT appears if the battery pack is properly installed; ABSENT if it is not.
Temperature	GOOD appears if the temperature is within the normal range. HIGH appears if the module is too hot.
Voltage	GOOD appears if the voltage is within the normal range. BAD appears if the voltage is out of range.
Fast Charging	COMPLETED appears if the fast charge cycle is done. CHARGING appears if the battery pack is charging.
No. of Cycles	Must be configured when you first install a battery pack. When you select this option, a small pop-up window appears asking if you want to reset the charge cycles. Select YES to reset the number of cycles to zero. After 1100 charge cycles, the life of the battery pack is assumed to be over and you must replace it.

6.2.7.2 Charging the Battery Pack

The battery pack is shipped uncharged, and you must charge it for 6 hours before you can use it. The battery pack will not supply power for the full data retention time until it is fully charged.

It is a good idea to set the MegaRAID controller cache write policy option to *Write-Through* during the battery pack charging period. After the battery pack is fully charged, you can change the cache write policy to *Write-Back*.

6.2.7.3 Changing the Battery Pack

The MegaRAID configuration software warns you when the battery pack must be replaced. A new battery pack should be installed every 1 to 5 years. The procedure for changing the battery pack is as follows:

1. Bring down the operating system properly. Make sure that cache memory has been flushed. You must perform a system reset if operating under DOS.

When the computer reboots, the MegaRAID SCSI 320-2 Ultra320 controller flushes cache memory.

2. Turn the computer power off and remove the computer cover.
3. Remove the MegaRAID controller.
4. Disconnect the battery pack cable or harness from J10 on the MegaRAID SCSI 320-2 Ultra320 controller.
5. Install a new battery pack and connect the new battery pack to J10.
6. Disable write-back caching using MegaRAID Manager or Power Console Plus.

6.2.7.4 Disposing of a Battery Pack

Warning: Do not dispose of the MegaRAID battery pack by fire. Do not mutilate the battery pack. Do not damage it in any way. Toxic chemicals can be released if it is damaged. Do not short-circuit the battery pack.

The material in the battery pack contains heavy metals that can contaminate the environment. Federal, state, and local laws prohibit disposal of some rechargeable batteries in public landfills. These

batteries must be sent to a specific location for proper disposal. Call the Rechargeable Battery Recycling Corporation at 352-376-6693 (FAX: 352-376-6658) for an authorized battery disposal site near you. For a list of battery disposal sites, write to:

Rechargeable Battery Recycling Corporation
2293 NW 41st Street
Gainesville FL 32606
Voice: 352-376-6693
FAX: 352-376-6658


Important: Most used Nickel-Metal Hydride batteries are not classified as hazardous waste under the federal RCRA (Resource Conservation and Recovery Act). Although Minnesota law requires that Nickel-Metal Hydride batteries be labeled “easily removable” from consumer products, and that Nickel-Metal Hydride batteries must be collected by manufacturers, the Minnesota Pollution Control Agency (MPCA) has granted a temporary exemption from these requirements.

LSI Logic reminds you that you must comply with all applicable battery disposal and hazardous material handling laws and regulations in the country or other jurisdiction where you are using an optional battery pack on the MegaRAID SCSI 320-2 controller.

6.2.8 Step 8: Install MegaRAID SCSI 320-2

Select a 3.3 V or 5 V PCI slot and align the MegaRAID SCSI 320-1 controller bus connector with the slot. Insert the MegaRAID SCSI 320-2 card in the PCI slot, as shown in [Figure 6.8](#). Press down gently but firmly to make sure that the card is properly seated in the slot. The bottom edge of the controller card should be flush with the slot. Attach the bracket to the computer frame with the bracket screw.

Figure 6.8 Installing the MegaRAID SCSI 320-2 Controller



6.2.9 Step 9: Connect SCSI Devices

Use SCSI cables to connect SCSI devices to the MegaRAID SCSI 320-1. The MegaRAID SCSI 320-2 has the following connectors:

- Two internal high-density 68-pin SCSI connectors: J7 is for SCSI channel 0, J8 is for SCSI channel 1.
- Two external very high-density 68-pin SCSI connector: J9 is for SCSI channel 0, J19 is for SCSI channel 1.

Use this procedure to connect SCSI devices:

1. Disable termination on any SCSI device that does *not* sit at the end of the SCSI bus.
2. Configure all SCSI devices to supply TermPWR.
3. Set proper target IDs (TIDs) for each SCSI device.
4. Distribute SCSI devices evenly across the SCSI channels for optimum performance.
5. Do not exceed the maximum cable length for the number and type of SCSI devices you are using. (See [Table 1.1](#) for details.)
6. Try to connect all non-disk SCSI devices to a SCSI channel that has no SCSI disk drives connected to it.

6.2.9.1 Cable Suggestions

System throughput problems can occur if SCSI cable use is not maximized. Here are some cabling guidelines:

- Use the shortest SCSI cables.
- LVD mode cable lengths should be no more than 25 meters with two devices and no more than 12 meters with eight devices.
- Use active termination.
- Avoid clustering the stubs.
- Cable stub length should be no more than 0.1 meter (4 inches).
- Route SCSI cables carefully.
- Use high-impedance cables.
- Do not mix cable types (choose flat cable for inside the enclosure, and round shielded cables for outside the enclosure.)
- Ribbon cables have fairly good cross-talk rejection characteristics.

6.2.10 Step 10: Set Target IDs

Set target identifiers (TIDs) on the SCSI devices. Each device in a specific SCSI channel must have a unique TID in that channel. Non-disk devices (CD-ROM or tapes) should have unique SCSI IDs *regardless of the channel where they are connected*. See the documentation for each SCSI device to set the TIDs. The MegaRAID SCSI 320-2 controller

automatically occupies TID 7 on each SCSI channel. Eight-bit SCSI devices can only use the TIDs from 0 to 6. 16-bit devices can use the TIDs from 0 to 15. The arbitration priority for a SCSI device depends on its TID. [Table 6.7](#) shows the relative priority of each Target ID:

Table 6.7 Priority of Target IDs

Priority	Highest						Lowest					
TID	7	6	5	...	2	1	0	15	14	...	9	8

Important: Non-disk devices (CD-ROM or tape drive) should have unique SCSI IDs regardless of the channel they are connected to. ID 0 cannot be used for non-disk devices because they are limited to IDs 1 through 6. There is a limit of six IDs per controller for non-disk devices.

6.2.10.1 Example of MegaRAID SCSI 320-2 ID Mapping

[Table 6.8](#) provides an example of ID mapping for the SCSI 320-2.

Table 6.8 Example of Mapping for SCSI 320-2

ID	Channel 0	Channel 1
0	A1-1	A1-2
1	A2-1	CD
2	CD	A2-3
3	A2-5	A2-6
4	CD	A3-1
5	A4-1	Tape
6	Optical	A5-1
7	Reserved	Reserved
8	A5-2	A5-3
9	A5-6	A5-7
10	A6-1	A6-2
11	A6-4	A6-5

Table 6.8 Example of Mapping for SCSI 320-2 (Cont.)

ID	Channel 0	Channel 1
12	A6-7	A6-8
13	A7-2	A7-3
14	A7-5	A7-6
15	A7-8	A8-1

6.2.11 Step 11: Power Up

Replace the computer cover and reconnect the AC power cords. Turn power on to the host computer. Set up the power supplies so that the SCSI devices are powered up at the same time as or before the host computer. If the computer is powered up before a SCSI device, the device might not be recognized.

During boot, the MegaRAID SCSI 320-2 BIOS message appears:

```
MegaRAID SCSI 320-2 Disk Array Adapter BIOS Version x.xx  
date  
Copyright (c) LSI Logic Corporation  
Firmware Initializing... [ Scanning SCSI Device ..(etc).. ]
```

The firmware takes several seconds to initialize. During this time the adapter scans the SCSI channel(s). When ready, the following appears:

```
Host Adapter-1 Firmware Version x.xx DRAM Size 16 MB  
0 Logical Drives found on the Host Adapter  
0 Logical Drives handled by BIOS  
Press <Ctrl><M> to run MegaRAID BIOS Configuration Utility
```

The <Ctrl><M> utility prompt times out after several seconds. The MegaRAID SCSI 320-2 host adapter (controller) number, firmware version, and cache DRAM size are displayed in the second portion of the BIOS message. The numbering of the controllers follows the PCI slot scanning order used by the host motherboard.

6.2.12 Step 12: Run the MegaRAID BIOS Configuration Utility

Press <Ctrl><M> to run the MegaRAID BIOS Configuration Utility. See the *MegaRAID Configuration Software Guide* for more information.

6.2.13 Step 13: Install the Operating System Driver

MegaRAID can operate under MS-DOS or any DOS-compatible operating system using the standard AT BIOS INT 13h Hard Disk Drive interface. To operate with other operating systems, you must install software drivers. MegaRAID provides software drivers on the Driver and Documentation CD for the following operating systems:

- MS-DOS version 6.xx or later
- Microsoft Windows NT 4.0, Windows 2000, Windows XP, Windows .NET
- Novell NetWare 5.1, 6.0
- Red Hat Linux 7.2, 7.3

Note: Refer to the *MegaRAID Driver Installation Guide* for the procedures used to install operating system drivers.

6.3 Summary

This chapter discussed hardware installation. Configure the RAID system using software configuration utilities. See the *MegaRAID Configuration Software Guide* for all information about MegaRAID SCSI 320-2 software utilities. [Table 6.9](#) lists these utility programs.

Table 6.9 Configuration Utilities and Operating Systems

Configuration Utility	Operating System
MegaRAID BIOS Configuration Utility	Independent of the operating system
WebBIOS Configuration Utility	Independent of the operating system
MegaRAID Manager	DOS Red Hat Linux 7.2, 7.3 Novell NetWare 5.1, 6.0
Power Console Plus	Microsoft Windows NT Windows 2000 Windows XP Windows .NET

Chapter 7

Installing and Configuring Clusters

This chapter explains how clusters work and how to install and configure them. It has the following sections:

- [Section 7.1, “Overview”](#)
 - [Section 7.2, “Benefits of Clusters”](#)
 - [Section 7.3, “Installation and Configuration”](#)
 - [Section 7.4, “Cluster Installation”](#)
 - [Section 7.5, “Installing SCSI Drives”](#)
-

7.1 Overview

Physically, a cluster is a grouping of two independent servers that can access the same shared data storage and provide services to a common set of clients. With current technology, this usually means servers connected to common I/O buses and a common network for client access.

Logically, a cluster is a single management unit. Any server can provide any available service to any authorized client. The servers must have access to the same shared data and must share a common security model. Again, with current technology, this generally means that the servers in a cluster will have the same architecture and run the same version of the same operating system.

7.2 Benefits of Clusters

Clusters provide three basic benefits:

- Improved application and data availability
- Scalability of hardware resources
- Simplified management of large or rapidly growing systems

7.3 Installation and Configuration

Use the following procedure to install and configure your system as part of a cluster.

1. Unpack the controller, following the instructions in Chapter 6.
2. Set the hardware termination for the controller as “always on”. See [Section 6.2.4.3, “J4/J5 Termination Enable,” page 6-6](#), for more information.
3. Configure the IDs for the drives in the enclosure.
4. Install one controller at a time. Press <Ctrl> <M> at BIOS initialization to configure the options in [step 5 – step 11](#). Do not attach the disks yet.
5. Set the controller to Cluster Mode in the Objects > Adapter > Cluster Mode menu.
6. Disable the BIOS in the Objects > Adapter > Enable/Disable BIOS menu.
7. Change the initiator ID in the Objects > Adapter > Initiator ID menu.
8. Power down the first server.
9. Attach the controller to the shared array.
10. Configure the first controller to the desired arrays using the Configure > New Configuration menu.

Note: Use the whole array size of any created array. Do not create partitions of different sizes on the RAID arrays from the BIOS Configuration Utility (<Ctrl><M>), because they

cannot be failed over individually when assigned driver letters in Windows 2000.

11. Follow the on-screen instructions to create arrays and save the configuration.
12. Repeat [step 4](#) – [step 7](#) for the second controller.
13. Power down the second server.
14. Attach the cables for the second controller to the shared enclosure and power up the second server.
15. If a configuration mismatch occurs, enter the <Ctrl> <M> utility, then go to the Configure-> View/Add Configuration > View Disk menu to view the disk configuration.
16. Save the configuration.
17. Proceed to the driver installation for a Microsoft cluster environment.

7.3.1 Driver Installation Instructions under Microsoft Windows 2000 Advanced Server

After the hardware is set up for the MS cluster configuration, perform the following procedure to configure the driver under Microsoft Windows 2000 Advanced Server

When the controller is added after a Windows 2000 Advanced Server installation, the operating system detects the controller.

1. When the Found New Hardware Wizard screen displays the detected hardware device, click on **Next**.
2. When the next screen appears, select **Search for a suitable driver....** and click on **Next**.

The Locate Driver Files screen appears.

3. Insert the floppy diskette with the appropriate driver disk for Windows 2000, then select **Floppy disk drives** on the screen and click on **Next**.

The Wizard detects the device driver on the diskette and the "Completing the upgrade device driver" wizard displays the name of the controller.

4. Click on **Finish** to complete the installation.

5. Repeat [step 1](#) – [step 4](#) to install the device driver on the second system.
After the cluster is installed, and both nodes are in the Microsoft Windows 2000 Advanced Server, installation will detect a SCSI processor device.
6. On the Found New Hardware Wizard prompt, choose to display a list of the known drivers, so that you can select a specific driver.
7. Click on **Next**.
8. Select the driver that you want to install for the device. If you have a disk with the driver you want to install, click on **Have Disk**.
9. Select **Other devices** from the list of hardware types, then click on **Next**.
10. Insert the disk containing the driver into the selected drive and click on **OK**.
11. Select the processor device and click on **Next**.
12. On the final screen, click on **Finish** to complete the installation.
13. Repeat the process on the peer system.

7.3.2 Network Requirements

The network requirements for clustering are:

- A unique NetBIOS cluster name
- Five unique, static IP addresses:
 - Two addresses are for the network adapters on the internal network
 - Two addresses are for the network adapters on the external network
 - One address is for the cluster itself
- A domain user account for Cluster Service (all nodes must be part of the same domain.)
- Two network adapters for each node—one for connection to the external network and the other for the node-to-node internal cluster network. If you do not use two network adapters for each node, your

configuration is unsupported. HCL certification requires a separate private network adapter.

7.3.3 Shared Disk Requirements

Disks can be shared by the nodes. The requirements for sharing disks are as follows:

- Physically attach all shared disks, including the quorum disk, to the shared bus.
- Make sure that all disks attached to the shared bus are seen from all nodes. You can check this at the setup level in <Ctrl> <M> (the BIOS configuration utility.) See [Section 7.5, “Installing SCSI Drives,” page 7-27](#), for installation information.
- Assign unique SCSI identification numbers to the SCSI devices and terminate the devices properly. Refer to the storage enclosure manual for information about installing and terminating SCSI devices.
- Configure all shared disks as basic (not dynamic).
- Format all partitions on the disks as NTFS.

It is best to use fault-tolerant RAID configurations for all disks. This includes RAID levels 1, 5, 10, or 50.

7.4 Cluster Installation

7.4.1 Installation Overview

During installation, some nodes are shut down, and other nodes are rebooted. This is necessary to ensure uncorrupted data on disks attached to the shared storage bus. Data corruption can occur when multiple nodes try to write simultaneously to the same disk, if that disk is not yet protected by the cluster software.

Table 7.1 shows which nodes and storage devices should be powered on during each step.

Table 7.1 Nodes and Storage Devices

Step	Node 1	Node 2	Storage	Comments
Set Up Networks	On	On	Off	Make sure that power to all storage devices on the shared bus is turned off. Power on all nodes.
Set up Shared Disks	On	Off	On	Power down all nodes. Next, power on the shared storage, then power on the first node.
Verify Disk Configuration	Off	On	On	Shut down the first node. Power on the second node.
Configure the First Node	On	Off	On	Shut down all nodes. Power on the first node.
Configure the Second Node	On	On	On	Power on the second node after the first node was successfully configured.
Post-installation	On	On	On	All nodes should be active.

Before installing the Cluster Service software you must follow the steps below:

- Install Windows 2000 Advanced Server or Windows 2000 Datacenter Server on each node
- Set up networks
- Set up disks

Note: These steps must be completed on every cluster node before proceeding with the installation of Cluster Service on the first node.

To configure the Cluster Service on a Windows 2000-based server, you must be able to log on as administrator or have administrative permissions on each node. Each node must be a member server, or must be a domain controller inside the same domain. A mix of domain controllers and member servers in a cluster is not acceptable.

7.4.2 Installing Microsoft Windows 2000

Install Microsoft Windows 2000 on each node. See your Windows 2000 manual for information.

Log on as administrator before you install the Cluster Services.

7.4.3 Setting Up Networks

Note: Do not allow both nodes to access the shared storage device before the Cluster Service is installed. In order to prevent this, power down any shared storage devices and then power up nodes one at a time. Install the Clustering Service on at least one node and make sure it is online before you power up the second node.

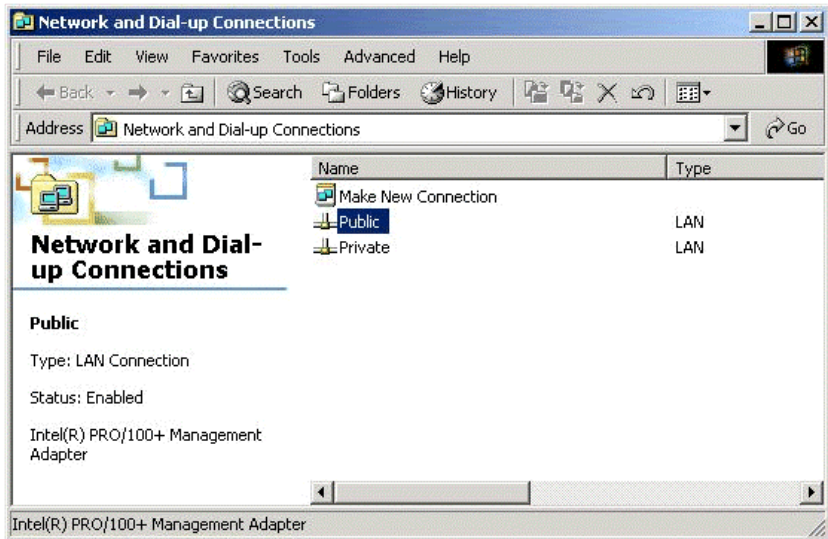
Install at least two network card adapters per each cluster node. One network card adapter card is used to access the public network. The second network card adapter is used to access the cluster nodes.

The network card adapter that is used to access the cluster nodes establishes the following:

- Node to node communications
- Cluster status signals
- Cluster Management

Check to make sure that all the network connections are correct. Network cards that access the public network must be connected to the public network. Network cards that access the cluster nodes must connect to each other.

Verify that all network connections are correct, with private network adapters connected to other private network adapters only, and public network adapters connected to the public network. View the Network and Dial-up Connections screen to check the connections.



Note: Use crossover cables for the network card adapters that access the cluster nodes. If you do not use the crossover cables properly, the system will not detect the network card adapter that accesses the cluster nodes. If the network card adapter is not detected, then you cannot configure the network adapters during the Cluster Service installation. However, if you install Cluster Service on both nodes, and both nodes are powered on, you can add the adapter as a cluster resource and configure it properly for the cluster node network in Cluster Administrator.

7.4.4 Configuring the Cluster Node Network Adapter

Note: The wiring determines which network adapter is private and which is public. For the purposes of this chapter, the first network adapter (Local Area Connection) is connected to the public network, and the second network adapter (Local Area Connection 2) is connected to the private cluster network. This may not be the case in your network.

7.4.4.1 Renaming the Local Area Connections

In order to make the network connection more clear, you can change the name of the Local Area Connection (2). Renaming it will help you identify

the connection and correctly assign it. Follow these steps to change the name:

1. Right-click on the **Local Area Connection 2** icon.
2. Click on **Rename**.
3. In the text box, type
Private Cluster Connection
and then press **Enter**.
4. Repeat steps 1-3 to change the name of the public LAN network adapter to Public Cluster Connection.
5. The renamed icons should look like those in the picture above. Close the Networking and Dial-up Connections window. The new connection names automatically replicate to other cluster servers as the servers are brought online.

7.4.5 Setting Up the First Node in Your Cluster

Follow these steps to set up the first node in your cluster.

1. Right-click on **My Network Places**, then click on **Properties**.
2. Right-click the **Private Connection** icon.
3. Click on **Status**. The Private Connection Status window shows the connection status, as well as the speed of connection.

If the window shows that the network is disconnected, examine cables and connections to resolve the problem before proceeding.
4. Click on **Close**
5. Right-click **Private Connection** again
6. Click on **Properties**.
7. Click on **Configure**.
8. Click on **Advanced**. The network card adapter properties window displays.
9. You should set network adapters on the private network to the actual speed of the network, rather than the default automated speed selection.

Select the network speed from the drop-down list. Do not use “Auto-select” as the setting for speed. Some adapters can drop packets while determining the speed.

Set the network adapter speed by clicking the appropriate option, such as Media Type or Speed.

10. Configure identically all network adapters in the cluster that are attached to the same network, so they use the same Duplex Mode, Flow Control, Media Type, and so on. These settings should stay the same even if the hardware is different.
11. Click on **Transmission Control Protocol/Internet Protocol (TCP/IP)**.
12. Click on **Properties**.
13. Click on the radio-button for **Use the following IP address**.
14. Enter the IP addresses you want to use for the private network.
15. Type in the subnet mask for the network.
16. Click the Advanced radio button, then select the WINS tab.
17. Select **Disable NetBIOS over TCP/IP**.
18. Click **OK** to return to the previous menu. *Perform this step for the private network adapter only.*

7.4.6 Configuring the Public Network Adapter

Note: It is strongly recommended that you use static IP addresses for all network adapters in the cluster. This includes both the network adapter used to access the cluster nodes and the network adapter used to access the LAN (Local Area Network). If you must use a dynamic IP address through DHCP, access to the cluster could be terminated and become unavailable if the DHCP server goes down or goes offline.

The use of long lease periods is recommended to assure that a dynamically assigned IP address remains valid in the event that the DHCP server is temporarily lost. In all cases, set static IP addresses for the private network connector. Note that Cluster Service will recognize only one network interface per subnet.

7.4.7 Verifying Connectivity and Name Resolution

Perform the following steps to verify that the network adapters are working properly:

Note: Before proceeding, you must know the IP address for each network card adapter in the cluster. You can obtain it by using the IPCONFIG command on each node.

1. Click on **Start**.
2. Click on **Run**.
3. Type `cmd` in the text box.
4. Click on **OK**.
5. Type `ipconfig /all` and press **Enter**. IP information displays for all network adapters in the machine.
6. If you do not already have the command prompt on your screen, click on **Start**.
7. Click on **Run**.
8. In the text box, type:

```
cmd
```

9. Click on **OK**.

10. Type

```
ping ipaddress
```

where *ipaddress* is the IP address for the corresponding network adapter in the other node. For example, assume that the IP addresses are set as follows:

Node	Network Name	Network Adapter IP Address
1	Public Cluster Connection	192.168.0.171
1	Private Cluster Connection	10.1.1.1
2	Public Cluster Connection	192.168.0.172
2	Private Cluster Connection	10.1.1.2

In this example, you would type:

```
Ping 192.168.0.172
```

and

```
Ping 10.1.1.1
```

from Node 1.

Then you would type:

```
Ping 192.168.0.172
```

and

```
10.1.1.1
```

from Node 2.

To confirm name resolution, ping each node from a client using the node's machine name instead of its IP number.

7.4.8 Verifying Domain Membership

All nodes in the cluster have to be members of the same domain and capable of accessing a domain controller and a DNS Server. You can configure them as either member servers or domain controllers. If you plan to configure one node as a domain controller, you should configure all other nodes as domain controllers in the same domain as well.

7.4.9 Setting Up a Cluster User Account

The Cluster Service requires a domain user account that the Cluster Service can run under. You must create the user account before installing the Cluster Service. The reason for this is that setup requires a user name and password. This user account should not belong to a user on the domain. Follow these steps to set up a cluster user account.

1. Click on **Start**.
2. Point to Programs, then point to Administrative Tools.
3. Click on **Active Directory Users and Computers**.
4. Click the plus sign (+) to expand the domain name (if it is not already expanded.)
5. Click on **Users**.

6. Right-click on **Users**.
7. Point to New and click on **User**.
8. Type in the cluster name and click on **Next**.
9. Set the password settings to User Cannot Change Password and Password Never Expires.
10. Click on **Next**, then click on **Finish to create this user**.

Note: If your company's security policy does not allow the use of passwords that never expire, you must renew the password on each node before password expiration. You must also update the Cluster Service configuration.

11. Right-click on **Cluster** in the left pane of the Active Directory Users and Computers snap-in.
12. Select **Properties** from the context menu.
13. Click on **Add Members to a Group**.
14. Click on **Administrators** and click on **OK**. This gives the new user account administrative privileges on this computer.
15. Close the Active Directory Users and Computers snap-in.

7.4.10 Setting Up Shared Disks

Warning: Make sure that Windows 2000 Advanced Server or Windows 2000 Datacenter Server and the Cluster Service are installed and running on one node before you start an operating system on another node. If the operating system is started on other nodes before you install and configure Cluster Service and run it on at least one node, the cluster disks will have a high chance of becoming corrupted.

To continue, power off all nodes. Power up the shared storage devices. Once the shared storage device is powered up, power up node one.

7.4.10.1 Quorum Disk

The quorum disk stores cluster configuration database checkpoints and log files that help manage the cluster. Windows 2000 makes the following quorum disk recommendations:

- Create a small partition [Use a minimum of 50 Mbytes as a quorum disk. Windows 2000 generally recommends a quorum disk to be 500 Mbytes.]
- Dedicate a separate disk for a quorum resource. The failure of the quorum disk would cause the entire cluster to fail; therefore, Windows 2000 strongly recommends that you use a volume on a RAID disk array.

During the Cluster Service installation, you have to provide the drive letter for the quorum disk.

Note: For our example, we use the letter **E** for the quorum disk drive letter.

7.4.11 Configuring Shared Disks

Perform the following steps to configure the shared disks:

1. Right-click on **My Computer**.
2. Click on **Manage**, then click on **Storage**.
3. Double-click on **Disk Management**.
4. Make sure that all shared disks are formatted as NTFS and are designated as Basic. If you connect a new drive, the Write Signature and Upgrade Disk Wizard starts automatically.

If this occurs, click on **Next** to go through the wizard. The wizard sets the disk to dynamic, but you can uncheck it at this point to set it to basic.

To reset the disk to Basic, right-click on **Disk #** (where # identifies the disk that you are working with) and click on **Revert to Basic Disk**.

5. Right-click on unallocated disk space.
6. Click on **Create Partition...**
The Create Partition Wizard begins.
7. Click on **Next** twice.
8. Enter the desired partition size in Mbytes and click on **Next**.
9. Accept the default drive letter assignment by clicking on **Next**.
10. Click on **Next** to format and create a partition.

7.4.12 Assigning Drive Letters

After you have configured the bus, disks, and partitions, you must assign drive letters to each partition on each clustered disk. Follow these steps to assign drive letters.

Note: Mountpoints is a feature of the file system that lets you mount a file system using an existing directory without assigning a drive letter. Mountpoints is not supported on clusters. Any external disk that is used as a cluster resource must be partitioned using NTFS partitions and must have a drive letter assigned to it.

1. Right-click on the desired partition and select **Change Drive Letter and Path**.
2. Select a new drive letter.
3. Repeat steps 1 and 2 for each shared disk.
4. Close the Computer Management window.

7.4.13 Verifying Disk Access and Functionality

Perform the following steps to verify disk access and functionality:

1. Click on **Start**.
2. Click on **Programs**. Click on **Accessories**, then click on **Notepad**.
3. Type some words into Notepad and use the File/Save As command to save it as a test file called test.txt. Close Notepad.
4. Double-click on the **My Documents** icon.
5. Right-click on test.txt and click on **Copy**.
6. Close the window.
7. Double-click on **My Computer**.
8. Double-click on a shared drive partition.
9. Click on **Edit** and click on **Paste**.
10. A copy of the file should now exist on the shared disk.
11. Double-click on **test.txt** to open it on the shared disk.
12. Close the file.

13. Highlight the file and press the **Del** key to delete it from the clustered disk.
14. Repeat the process for all clustered disks to make sure they can be accessed from the first node.

After you complete the procedure, shut down the first node, power on the second node and repeat the procedure above. Repeat again for any additional nodes. After you have verified that all nodes can read and write from the disks, turn off all nodes except the first, and continue with this guide.

7.4.14 Installing Cluster Service Software

Before you begin the Cluster Service Software installation on the first node, make sure that all other nodes are either powered down or stopped and that all shared storage devices are powered on.

To create the cluster, you must provide the cluster information. The Cluster Configuration Wizard allows you to input this information. Follow these steps to use the Wizard:

1. Click on **Start**.
2. Click on **Settings**, then click on **Control Panel**.
3. Double-click on **Add/Remove Programs**.
4. Double-click on **Add/Remove Windows Components**.
5. Select **Cluster Service**, then click on Next.

Cluster Service files are located on the Windows 2000 Advanced Server or Windows 2000 Datacenter Server CD-ROM.

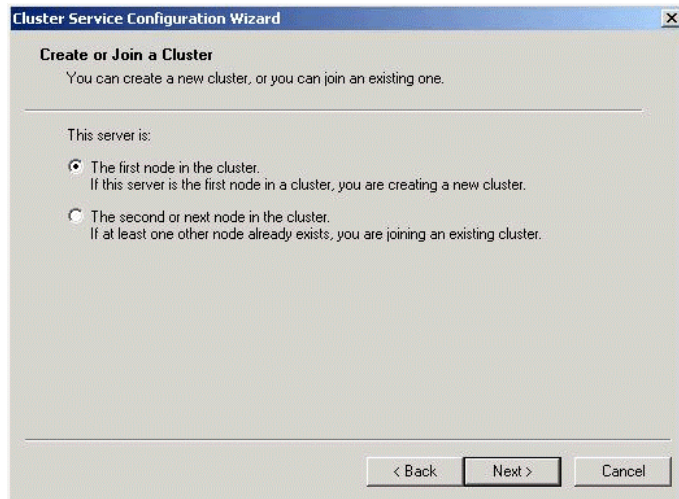
6. Enter `x:\i386` (where *x* is the drive letter of your CD-ROM). If you installed Windows 2000 from a network, enter the appropriate network path instead. (If the Windows 2000 Setup flash screen displays, close it.)
7. Click on **OK**. The Cluster Service Configuration Window displays.
8. Click on **Next**.

The Hardware Configuration Certification window appears.

9. Click on **I Understand** to accept the condition that Cluster Service is supported only on hardware listed on the Hardware Compatibility List.

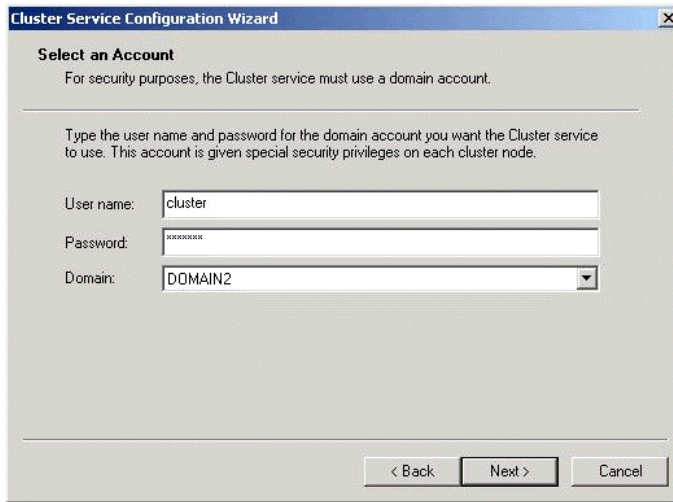
This is the first node in the cluster; therefore, you must create the cluster itself.

10. Select the first node in the cluster, as shown below and then click on **Next**.



11. Enter a name for the cluster (up to 15 characters), and click on **Next**. (In our example, the cluster is named ClusterOne.)
12. Type the user name of the Cluster Service account that you created during the pre-installation. (In our example, the user name is cluster.) Do not enter a password.
13. Type the domain name, then click on **Next**.

At this point the Cluster Service Configuration Wizard validates the user account and password.

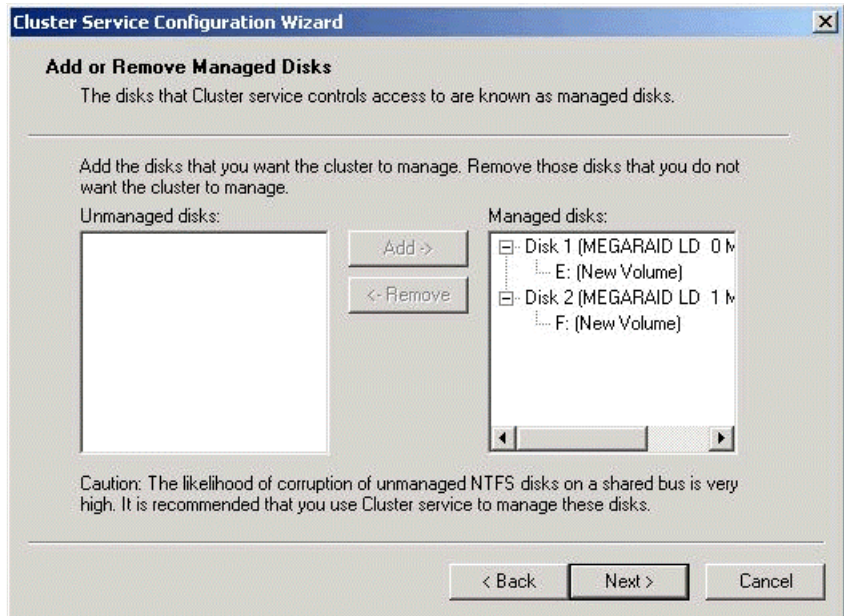


14. Click on **Next**.

The Add or Remove Managed Disks screen displays next. This screen is in the following section about configuring cluster disks.

7.4.15 Configuring Cluster Disks

Windows 2000 Managed Disks displays all SCSI disks, as shown on the screen below. It displays SCSI disks that do not reside on the same bus as the system disk. Because of this, a node that has multiple SCSI buses will list SCSI disks that are not to be used as shared storage. You must remove any SCSI disks that are internal to the node and not to be shared storage.



In production clustering scenarios, you need to use more than one private network for cluster communication to avoid having a single point of failure. Cluster Service can use private networks for cluster status signals and cluster management. This provides more security than using a public network for these roles. In addition, you can use a public network for cluster management, or you can use a mixed network for both private and public communications.

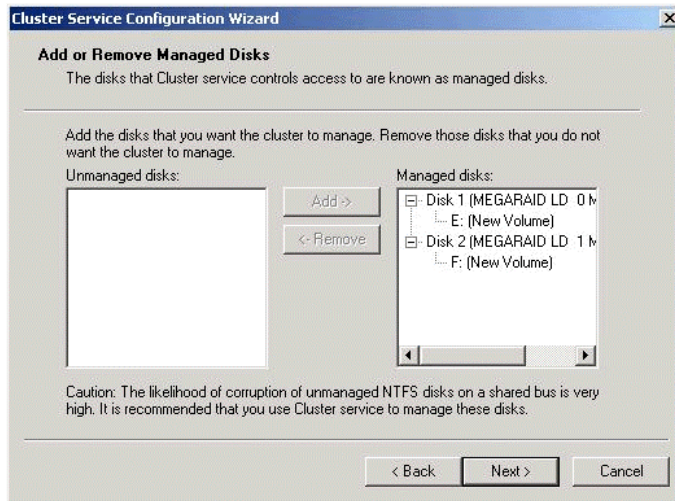
In any case, verify that at least two networks are used for cluster communication; using a single network for node-to-node communication creates a potential single point of failure. We recommend that you use multiple networks, with at least one network configured as a private link between nodes and other connections through a public network. If you use more than one private network, make sure that each uses a different subnet, as Cluster Service recognizes only one network interface per subnet.

This document assumes that only two networks are in use. It describes how you can configure these networks as one mixed and one private network.

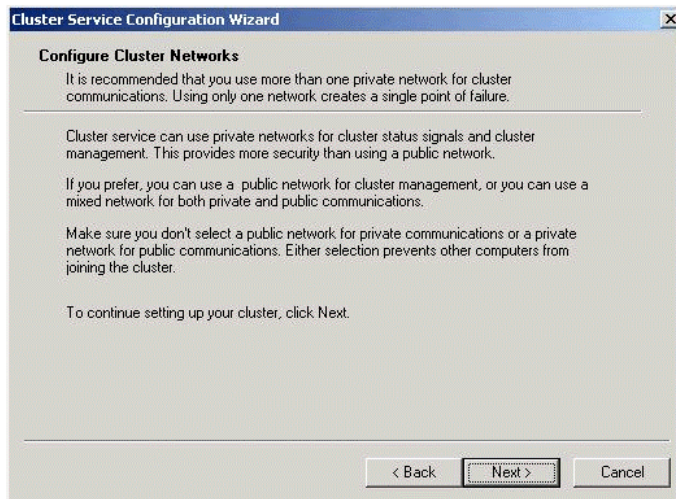
The order in which the Cluster Service Configuration Wizard presents these networks can vary. In this example, the public network is presented first.

Follow these steps to configure the clustered disks:

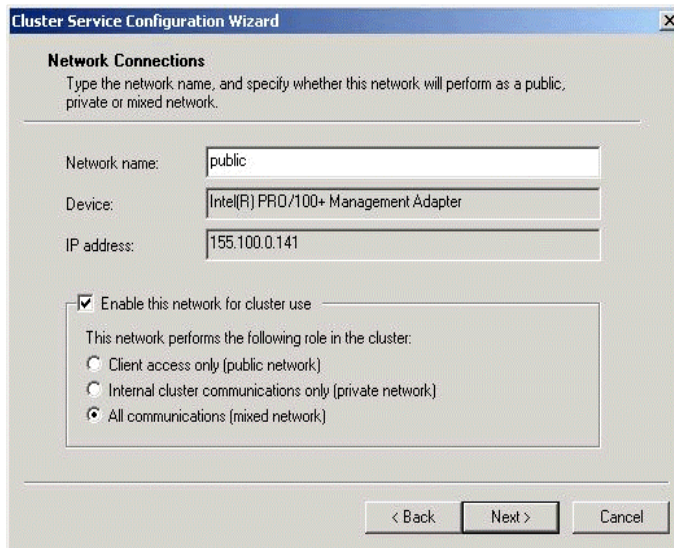
1. The Add or Remove Managed Disks dialog box specifies disks on the shared SCSI bus that will be used by Cluster Service. Add or remove disks as necessary, then click on **Next**.



The following screen displays.

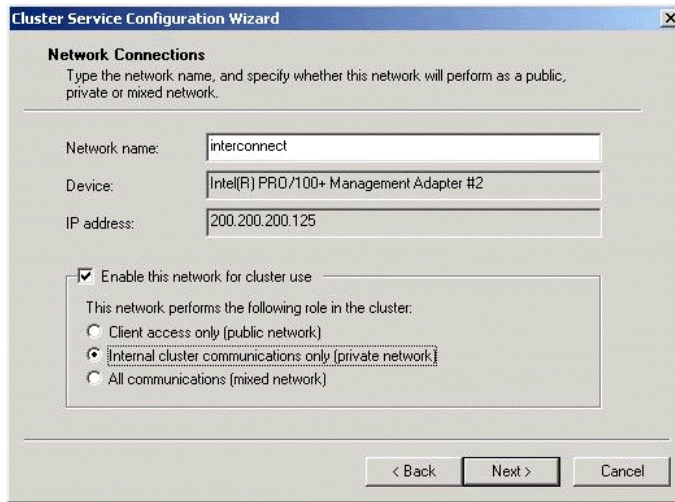


2. Click on **Next** in the Configure Cluster Networks dialog box.
3. Verify that the network name and IP address correspond to the network interface for the *public* network.
4. Check the box **Enable this network for cluster use**.
5. Select the option **All communications (mixed network)**, as shown below, and click on **Next**.



The next dialog box configures the private network. Make sure that the network name and IP address correspond to the network interface used for the *private* network.

6. Check the box **Enable this network for cluster use**.
7. Select the option **Internal cluster communications only**, then click on **Next**.

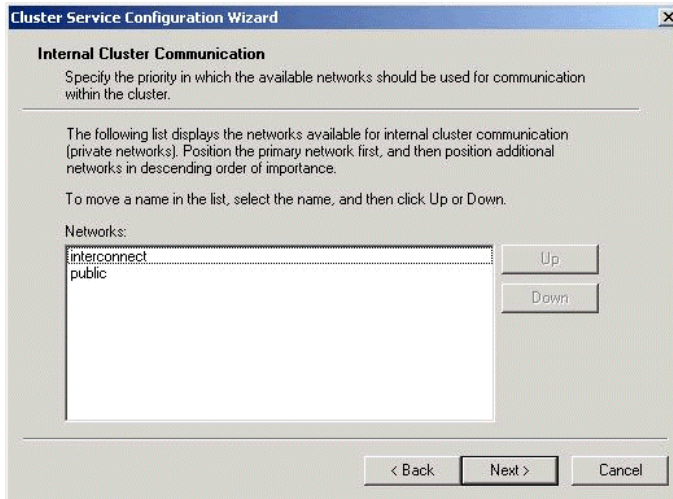


In this example, both networks are configured so that both can be used for internal cluster communication. The next dialog window offers an option to modify the order in which the networks are used. Because Private Cluster Connection represents a direct connection between nodes, it remains at the top of the list.

In normal operation, this connection is used for cluster communication. In case of the Private Cluster Connection failure, Cluster Service automatically switches to the next network on the list—in this case Public Cluster Connection.

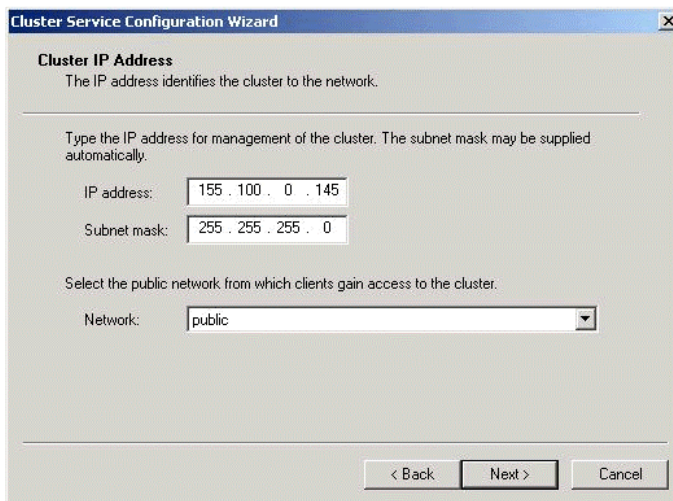
8. Verify that the first connection in the list is the Private Cluster Connection, then click on **Next**.

Note: Always set the order of the connections so that the Private Cluster Connection is first in the list.



9. Enter the unique cluster IP address and Subnet mask for your network, then click on **Next**.

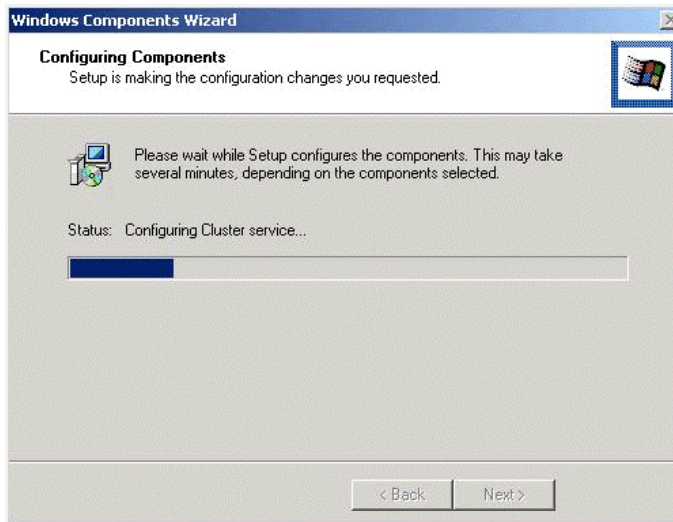
The Cluster Service Configuration Wizard shown below automatically associates the cluster IP address with one of the public or mixed networks. It uses the subnet mask to select the correct network.



10. Click **Finish** to complete the cluster configuration on the first node.

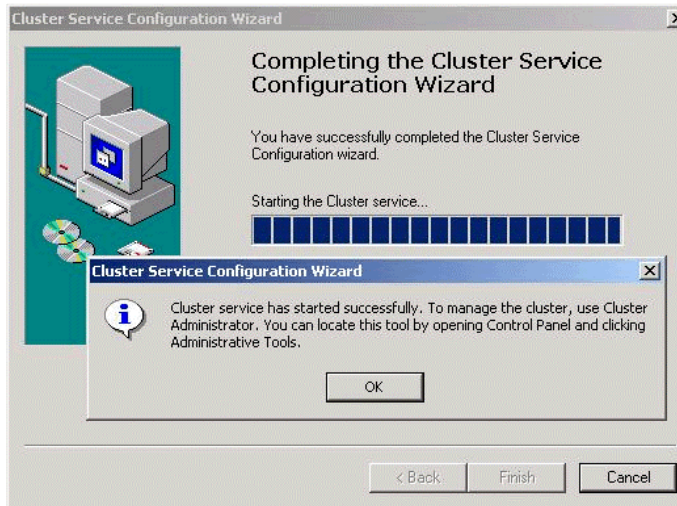


The Cluster Service Setup Wizard completes the setup process for the first node by copying the files needed to complete the installation of Cluster Service.



After the files are copied, the Cluster Service registry entries are created, the log files on the quorum resource are created, and the Cluster Service is started on the first node.

11. When a dialog box appears telling you that Cluster Service has started successfully. Click on **OK**.



12. Close the Add/Remove Programs window.

7.4.16 Validating the Cluster Installation

Use the Cluster Administrator snap-in to validate the Cluster Service installation on the first node. Follow these steps to validate the cluster installation.

1. Click on **Start**.
2. Click on **Programs**.
3. Click on **Administrative Tools**.
4. Click on **Cluster Administrator**.

The Cluster Administrator screen displays. If your snap-in window is similar to the one shown in the screen, your Cluster Service was successfully installed on the first node. You are now ready to install Cluster Service on the second node.

7.4.17 Configuring the Second Node

Note: For this procedure, have node one and all shared disks powered on, then power up the second node.

Installation of Cluster Service on the second node takes less time than on the first node. Setup configures the Cluster Service network settings on the second node based on the configuration of the first node.

Installation of Cluster Service on the second node begins the same way as installation on the first node. The first node must be running during installation of the second node.

Follow the same procedures used to install Cluster Service on the first node, with the following differences:

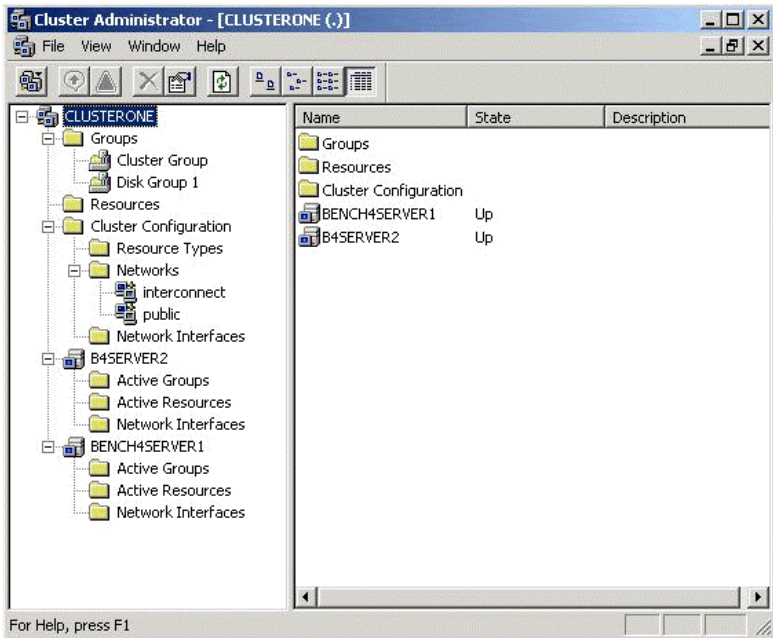
1. In the Create or Join a Cluster dialog box, select **The second or next node in the cluster**, then click **Next**.
2. Enter the cluster name that was previously created (it is MyCluster in this example), and click **Next**.
3. Leave **Connect to cluster as** unchecked. The Cluster Service Configuration Wizard automatically supplies the name of the user account selected when you installed the first node. Always use the same account you used when you set up the first cluster node.
4. Enter the password for the account (if there is one), then click **Next**.
5. At the next dialog box, click **Finish** to complete configuration.
The Cluster Service will start.
6. Click **OK**.
7. Close Add/Remove Programs.
8. If you install additional nodes, repeat the preceding steps to install Cluster Service on all other nodes.

7.4.18 Verifying Installation

There are several ways to verify that Cluster Service was successfully installed. Here is a simple one:

1. Click **Start**, click **Programs**, click **Administrative Tools**, then click **Cluster Administrator**.

The presence of two nodes (pictured below) shows that a cluster exists and is in operation.



2. Right-click the group Disk Group 1 and select the option **Move**. This option moves the group and all its resources to another node. After a short period of time, the Disk F: G: will be brought online on the second node. If you watch the screen, you will see this shift. Close the Cluster Administrator snap-in.

Congratulations! You have completed installing Cluster Service on all nodes. The server cluster is fully operational. Now, you are ready to install cluster resources, such as file shares, printer spoolers, cluster aware services like IIS, Message Queuing, Distributed Transaction Coordinator, DHCP, WINS, or cluster aware applications like Exchange or SQL Server.

7.5 Installing SCSI Drives

This information is provided as a generic instruction set for SCSI drive installations. If the SCSI hard disk vendor's instructions conflict with the instructions in this section, always use the instructions supplied by the vendor.

The SCSI bus listed in the hardware requirements must be configured prior to installation of Cluster Services. This includes:

- Configuring the SCSI devices.
- Configuring the SCSI controllers and hard disks to work properly on a shared SCSI bus.
- Properly terminating the bus. The shared SCSI bus must have a terminator at each end of the bus. It is possible to have multiple shared SCSI buses between the nodes of a cluster.

In addition to the information on the next page, refer to the documentation from the SCSI device manufacturer or the SCSI specifications, which can be ordered from the American National Standards Institute (ANSI). The ANSI web site contains a catalog that you can search for the SCSI specifications.

7.5.1 Configuring the SCSI Devices

Each device on the shared SCSI bus must have a unique SCSI ID. Since most SCSI controllers default to SCSI ID 7, part of configuring the shared SCSI bus will be to change the SCSI ID on one controller to a different SCSI ID, such as SCSI ID 6. If there is more than one disk that will be on the shared SCSI bus, each disk must also have a unique SCSI ID.

Some SCSI controllers reset the SCSI bus when they initialize at boot time. If this occurs, the bus reset can interrupt any data transfers between the other node and disks on the shared SCSI bus. Therefore, SCSI bus resets should be disabled if possible.

7.5.2 Terminating the Shared SCSI Bus

You can connect **Y** cables to devices if the device is at the end of the SCSI bus. You can then attach a terminator to one branch of the **Y** cable to terminate the SCSI bus. This method of termination requires either disabling or removing any internal terminators the device has.

Note: Any devices that are not at the end of the shared bus must have their internal termination disabled.

Chapter 8

Troubleshooting

This chapter provides troubleshooting information for the MegaRAID SCSI 320-2 controller. It contains the following sections:

- [Section 8.1, “General Troubleshooting”](#)
- [Section 8.2, “BIOS Boot Error Messages”](#)
- [Section 8.3, “Other BIOS Error Messages”](#)
- [Section 8.4, “Other Potential Problems”](#)

8.1 General Troubleshooting

[Table 8.1](#) lists the general problems that can occur, along with suggested solutions.

Table 8.1 General Problems and Suggested Solutions

Problem	Suggested Solution
Some operating systems do not load in a computer with a MegaRAID SCSI 320-2 controller.	Check the system BIOS configuration for PCI interrupt assignments. Make sure some Interrupts are assigned for PCI. Initialize the logical drive before installing the operating system.
One of the hard drives in the array fails often.	Check the drive error counts using Power Console Plus. (See the <i>MegaRAID Configuration Software Guide</i> for more information.) Format the drive. Rebuild the drive If the drive continues to fail, replace it with another drive with the same capacity.

Table 8.1 General Problems and Suggested Solutions (Cont.)

Problem	Suggested Solution
Pressed <Ctrl><M>. Ran Megaconf.exe and tried to make a new configuration. The system hangs when scanning devices.	Check the drives IDs on each channel to make sure each device has a different ID. Check the termination. The device at the end of the channel must be terminated. Replace the drive cable.
Multiple drives connected to the MegaRAID SCSI 320-2 use the same power supply. There is a problem spinning the drives all at once.	Set the drives to spin on command. This will allow MegaRAID SCSI 320-2 to spin two devices simultaneously.
Pressing <Ctrl><M> or running megaconf.exe does not display the Management Menu.	These utilities require a color monitor.
At system power-up with the MegaRAID installed, the screen display is garbled.	At least 32 Mbytes of memory must be installed before power-up.
Cannot flash or update the EEPROM.	You may need a new EEPROM.
Firmware Initializing... appears and remains on the screen.	Make sure that TERMPWR is being properly provided to each peripheral device populated channel. Make sure that each end of the channel chain is properly terminated using the recommended terminator type for the peripheral device. The channel is automatically terminated at the MegaRAID SCSI 320-2 card if only one cable is connected to a channel. Make sure that memory modules are PC100 or faster. Make sure that the MegaRAID SCSI 320-2 controller is properly seated in the PCI slot.
What is the maximum number of MegaRAID adapters per computer?	Currently, all the utilities and drivers support up to 12 MegaRAID adapters per system.
What SCSI IDs can a non-hard disk device have and what is maximum number allowed per adapter?	Non-hard disk devices can accommodate only SCSI IDs 1, 2, 3, 4, 5, or 6, regardless of the channel used. A maximum of six non-hard disk devices are supported per MegaRAID SCSI 320-2 adapter.
Why does a failed logical array still get a drive assignment?	To maintain the DOS Path statement integrity.

8.2 BIOS Boot Error Messages

Table 8.2 describes BIOS error messages that can display when you boot the system, and suggested solutions.

Table 8.2 BIOS Boot Error Messages

Message	Problem	Suggested Solution
Adapter BIOS Disabled. No Logical Drives Handled by BIOS	The MegaRAID BIOS is disabled. Sometimes the BIOS is disabled to prevent booting from the BIOS.	Enable the BIOS using the MegaRAID BIOS Configuration Utility.
Host Adapter at Baseport xxxx Not Responding	The BIOS cannot communicate with the adapter firmware.	Make sure the MegaRAID SCSI 320-2 is properly installed. Try moving the MegaRAID controller to another slot. Run the MegaRAID Manager Diagnostics to verify that the MegaRAID controller is functioning properly
No MegaRAID SCSI 320-2 Adapter	The BIOS cannot communicate with the adapter firmware.	Make sure the MegaRAID SCSI 320-2 is properly installed. Try moving the MegaRAID card to another PCI slot. Run the MegaRAID Manager Diagnostics to verify that MegaRAID is functioning properly.
Configuration of nonvolatile RAM (NVRAM) and drives mismatch. Run View/Add Configuration option of Configuration Utility. Press any key to run the Configuration Utility.	The configuration stored in the MegaRAID SCSI 320-2 adapter does not match the configuration stored in the drives.	<ol style="list-style-type: none">1. Press a key to run MegaRAID Manager.2. Select View/Add Configuration from the Configure menu.3. Use View/Add Configuration to examine both the configuration in NVRAM and the configuration stored on the disk drives.4. Resolve the problem by selecting one of the configurations.
1 Logical Drive Failed	A logical drive failed to sign on.	Make sure all physical drives are properly connected and are powered on. Run MegaRAID Manager to find out if any physical drives are not responding. Reconnect, replace, or rebuild any drive that is not responding.

Table 8.2 BIOS Boot Error Messages (Cont.)

Message	Problem	Suggested Solution
X Logical Drives Degraded	x number of logical drives signed on in a degraded state.	Make sure all physical drives are properly connected and are powered on. Run MegaRAID Manager to find if any physical drives are not responding. Reconnect, replace, or rebuild any drive that is not responding.
1 Logical Drive Degraded	A logical drive signed on in a degraded state.	Make sure all physical drives are properly connected and are powered on. Run MegaRAID Manager to find out if any physical drives are not responding. Reconnect, replace, or rebuild any drive that is not responding.
Insufficient memory to run BIOS. Press any key to continue...	Not enough MegaRAID memory to run MegaRAID BIOS.	Make sure MegaRAID memory has been properly installed.
Insufficient Memory	Not enough memory on the MegaRAID adapter to support the current configuration.	Make sure MegaRAID memory has been properly installed.
The following SCSI IDs are not responding: Channel x:a.b.c	The physical drives with SCSI IDs a, b, and c are not responding on SCSI channel x.	Make sure the physical drives are properly connected and are powered on.

8.3 Other BIOS Error Messages

Table 8.3 describes other BIOS error messages, their meaning, and suggested solutions.

Table 8.3 Other BIOS Error Messages

Message	Problem	Suggested Solution
Following SCSI disk not found and no empty slot available for mapping it	The physical disk roaming feature did not find the physical disk with the displayed SCSI ID. No slot is available to map the physical drive. MegaRAID cannot resolve the physical drives into the current configuration.	Reconfigure the array.
Following SCSI IDs have the same data y, z Channel x: a, b, c	The physical drive roaming feature found the same data on two or more physical drive on channel x with SCSI IDs a, b, and c. MegaRAID cannot determine the drive that has the duplicate information.	Remove the drive or drives that should not be used.
Unresolved configuration mismatch between disks and NVRAM on the adapter	The configuration stored in the MegaRAID NVRAM does not match the configuration stored on the drives.	<ol style="list-style-type: none">1. Press a key to run MegaRAID Manager.2. Select View/Add Configuration from the Configure menu.3. Use View/Add Configuration to examine both the configuration in NVRAM and the configuration stored on the disk drives.4. Resolve the problem by selecting one of the configurations.

8.4 Other Potential Problems

Table 8.4 describes other potential problems.

Table 8.4 Other Potential Problems

Topic	Information
DOS ASPI	MEGASPI.SYS, the MegaRAID DOS ASPI manager, uses 6 Kbytes of system memory once it is loaded.

Table 8.4 Other Potential Problems (Cont.)

Topic	Information
CD drives under DOS	At this time, copied CDs are not accessible from DOS even after loading MEGASPI.SYS and MEGACDR.SYS.
Physical drive errors	To display the MegaRAID Manager Media Error and Other Error options, select the Objects menu, then Physical Drive. Select a physical drive and press <F2>. The windows displays the number of errors. A Media Error is an error that occurred while actually transferring data. An Other Error is an error that occurs at the hardware level because of a device failure, poor cabling, bad termination, signal loss, etc.
Virtual sizing	The FlexRAID virtual sizing option enables RAID expansion. FlexRAID virtual sizing must be enabled to increase the size of a logical drive or add a physical drive to an existing logical drive. Run MegaRAID Manager by pressing <Ctrl> <M> to enable FlexRAID virtual sizing. Select the Objects menu, then select the Logical Drive menu. Select View/Update Parameters. Set FlexRAID Virtual Sizing to Enabled.
BSD Unix	We do not provide a driver for BSDI Unix. MegaRAID SCSI 320-2 does not support BSDI Unix.
Multiple LUNs	MegaRAID SCSI 320-2 supports one logical unit number (LUN) per each target ID. No multiple LUN devices are supported.
MegaRAID power requirements	The maximum MegaRAID SCSI 320-2 power requirements are 15 W at 5 V and 3 A.
SCSI bus requirements	The ANSI specification dictates the following: <ul style="list-style-type: none"> • The maximum signal path length between terminators is 3 meters when using up to 4 maximum capacitance (25 pF) devices and 1.5 meters when using more than 4 devices. • SCSI devices should be uniformly spaced between terminators, with the end devices located as close as possible to the terminators. • The characteristic impedance of the cable should be 90 +/- 6 ohms for the /REQ and /ACK signals and 90 +/- 10 ohms for all other signals. • The stub length (the distance from the controller's external connector to the mainline SCSI bus) shall not exceed 0.1m (approximately 4 inches). • The spacing of devices on the mainline SCSI bus should be at least three times the stub length. • All signal lines shall be terminated once at both ends of the bus powered by the TERMPWR line.

Table 8.4 Other Potential Problems (Cont.)

Topic	Information
Windows NT Installation	<p>When Windows NT is installed using a bootable CD, the devices on the MegaRAID SCSI 320-2 are not recognized until after the initial reboot. The Microsoft documented workaround is in SETUP.TXT, which is on the CD. Perform the following steps to install drivers when Setup recognizes a supported SCSI host adapter but doesn't make the devices attached to it available for use:</p> <ol style="list-style-type: none">1. Restart Windows NT Setup.2. Press <F6> to prevent Windows NT Setup from performing disk controller detection when Windows NT Setup displays the following: <code>Setup is inspecting your computer's hardware configuration...</code> This allows you to install the driver from the drivers disk you created. All SCSI adapters must be installed manually.3. Press <S> to display a list of supported SCSI host adapters when Windows NT Setup displays the following: <code>Setup could not determine the type of one or more mass storage devices installed in your system, or you have chosen to manually specify an adapter,</code>4. Select Other from the bottom of the list.5. Insert the drivers disk you made when prompted to do so and select MegaRAID from this list. In some cases, Windows NT Setup repeatedly prompts to swap disks. Windows NT will recognize any devices attached to this adapter.6. Repeat this step for each host adapter not already recognized by Windows NT Setup.

Appendix A

SCSI Cables and Connectors

The MegaRAID SCSI 320-2 provides several different types of SCSI connectors. The connectors are:

- Two 68-pin high density internal connectors
 - Two 68-pin very high density external connectors
-

A.1 68-Pin High-Density SCSI Internal Connector

Each SCSI channel on the MegaRAID SCSI 320-2 controller has a 68-pin high density 0.050 inch pitch unshielded connector.


These connectors provide all signals needed to connect the MegaRAID SCSI 320-2 to Wide SCSI devices. The connector pinouts are for a single-ended primary bus (P-CABLE) as specified in SCSI-3 Parallel Interface X3T9.2, Project 885-D, revision 12b, dated July 2, 1993.

The cable assemblies that interface with the 68-pin connector are:

- Flat ribbon or twisted pair cable for connecting internal Wide SCSI devices
- Round shielded cable for connecting external Wide SCSI devices

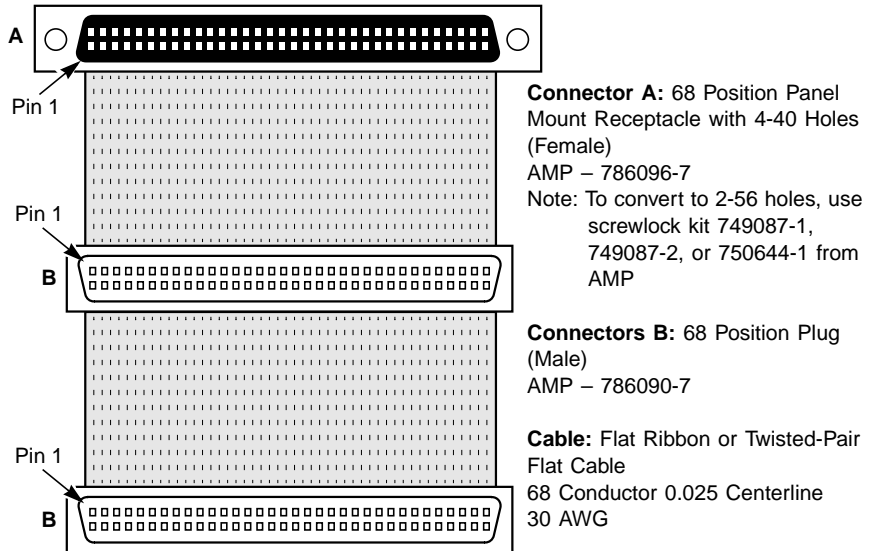
A.1.1 Cable Assembly for Internal Wide SCSI Devices

The cable assembly for connecting internal Wide SCSI devices is shown below.



A.1.2 Connecting Internal and External Wide Devices

The cable assembly for connecting internal Wide and external Wide SCSI devices is shown below.




A.1.3 Converting Internal Wide to Internal Non-Wide (Type 2)

The cable assembly for converting internal Wide SCSI connectors to internal non-Wide (Type 2) SCSI connectors is shown below.

68 Position Connector Contact Number		50 Position Connector Contact Number
6	● ————— ●	1
40	● ————— ●	2
7	● ————— ●	3
41	● ————— ●	4
	⋮	
49	● ————— ●	20
16	● ————— ●	21
50	● ————— ●	22
17	Open	23
51	Open	24
18	Open	25
52	● ————— ●	26
19	● ————— ●	27
	⋮	
29	● ————— ●	47
63	● ————— ●	48
30	● ————— ●	49
64	● ————— ●	50

Table 1: Connector Contact Connection for Wide to Non-Wide Conversion




Connector A: 68 Position Plug (Male)
AMP- 749925-5

Connector B: 50 Position IDC Receptacle (Female)
AMP - 499252-4, 1-746285-0, 1-746288-0

Wire: Twisted-Pair Flat Cable or Laminated Discrete Wire Cable
25 Pair 0.050 Centerline
28 AWG

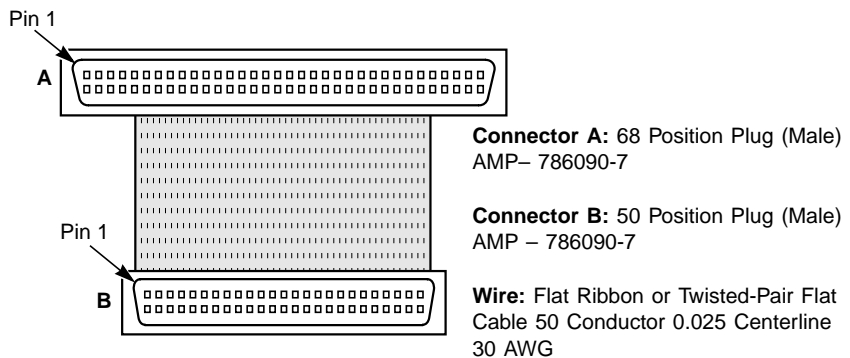
A.1.4 Converting Internal Wide to Internal Non-Wide (Type 30)

The cable assembly for converting internal Wide SCSI connectors to internal non-Wide (Type 30) SCSI connectors is shown below.



A.1.5 Converting from Internal Wide to Internal Non-Wide (Type 3)

The cable assembly for converting internal Wide SCSI connectors to internal non-Wide (Type 3) SCSI connectors is shown below.



A.2 SCSI Cable and Connector Vendors

Table A.1 lists SCSI cable vendors, and contact information.

Table A.1 SCSI Cable Vendors

Manufacturer	Telephone Number
Cables To Go	Voice: 800-826-7904 Fax: 800-331-2841
System Connection	Voice: 800-877-1985
Technical Cable Concepts	Voice: 714-835-1081
GWC	Voice: 800-659-1599

Table A.2 table lists SCSI connector vendors.

Table A.2 SCSI Connector Vendors

Manufacturer	Connector Part Number	Back Shell Part Number
AMP	749111-4	749193-1
Fujitsu	FCN-237R050-G/F	FCN-230C050-D/E
Honda	PCS-XE50MA	PCS-E50LA

A.3 High-Density 68-Pin Connector Pinout for SE SCSI

Table A.3 lists the pinout for the high-density 68-pin connectors for single-ended SCSI.

Table A.3 High-Density 68-Pin Connector Pinout (SE SCSI)

Signal	Connector Pin	Cable Pin	Cable Pin	Connector Pin	Signal
Ground	1	1	2	35	-DB(12)
Ground	2	3	4	36	-DB(13)
Ground	3	5	6	37	-DB(14)
Ground	4	7	8	38	-DB(15)
Ground	5	9	10	39	-DB(P1)
Ground	6	11	12	40	-DB(0)
Ground	7	13	14	41	-DB(1)
Ground	8	15	16	42	-DB(2)
Ground	9	17	18	43	-DB(3)
Ground	10	19	20	44	-DB(4)
Ground	11	21	22	45	-DB(5)
Ground	12	23	24	46	-DB(6)
Ground	13	25	26	47	-DB(7)
Ground	14	27	28	48	-DB(P)
Ground	15	29	30	49	SWAP L
Ground	16	31	32	50	SHELF_OK
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
Reserved	19	37	38	53	Reserved
Ground	20	39	40	54	FAULT_CLK H
Ground	21	41	42	55	-ATN

Table A.3 High-Density 68-Pin Connector Pinout (SE SCSI)

Signal	Connector Pin	Cable Pin	Cable Pin	Connector Pin	Signal
Ground	22	43	44	56	FAULT_DATA H
Ground	23	45	46	57	-BSY
Ground	24	47	48	58	-ACK
Ground	25	49	50	59	-RST
Ground	26	51	52	60	-MSG
Ground	27	53	54	61	-SEL
Ground	28	55	56	62	-C/D
Ground	29	57	58	63	-REQ
Ground	30	59	60	64	-I/O
Ground	31	61	62	65	-DB(8)
Ground	32	63	64	66	-DB(9)
Ground	33	65	66	67	-DB(10)
Ground	34	67	68	68	-DB(11)

The following applies to the high-density SCSI connector described in [Table A.3](#):

- A hyphen before a signal name indicates that signal is active low.
- The connector pin refers to the conductor position when using 0.025 inch centerline flat ribbon cable with a high-density connector (AMPLIMITE.050 Series connectors).
- Eight-bit devices connected to the P-Cable must leave the following signals open: -DB (8), -DB (9), -DB (10), -DB (11), -DB(12), -DB (13), -DB (14), -DB 15), and -DB (P1).
- All other signals should be connected as defined.

Caution: Lines labeled RESERVED should be connected to Ground in the bus terminator assemblies or in the end devices on the SCSI cable. RESERVED lines should be open in the other SCSI devices, but can be connected to Ground.

A.4 68-Pin Connector Pinout for LVD SCSI

Table A.4 lists the pinout for the 68-pin connector for LVD SCSI.

Table A.4 68-Pin Connector Pinout for LVD SCSI

Signal	Connector Pin	Cable Pin	Cable Pin	Connector Pin	Signal
+DB(12)	1	1	2	35	-DB(12)
+DB(13)	2	3	4	36	-DB(13)
+DB(14)	3	5	6	37	-DB(14)
+DB(15)	4	7	8	38	-DB(15)
+DB(P1)	5	9	10	39	-DB(P1)
+DB(0)	6	11	12	40	-DB(0)
+DB(1)	7	13	14	41	-DB(1)
+DB(2)	8	15	16	42	-DB(2)
+DB(3)	9	17	18	43	-DB(3)
+DB(4)	10	19	20	44	-DB(4)
+DB(5)	11	21	22	45	-DB(5)
+DB(6)	12	23	24	46	-DB(6)
+DB(7)	13	25	26	47	-DB(7)
+DB(P)	14	27	28	48	-DB(P)
Ground	15	29	30	49	Ground
DIFFSENS	16	31	32	50	Ground
TERMPWR	17	33	34	51	TERMPWR
TERMPWR	18	35	36	52	TERMPWR
Reserved	19	37	38	53	Reserved
Ground	20	39	40	54	Ground
+ATN	21	41	42	55	-ATN
Ground	22	43	44	56	Ground

Table A.4 68-Pin Connector Pinout for LVD SCSI (Cont.)

Signal	Connector Pin	Cable Pin	Cable Pin	Connector Pin	Signal
+BSY	23	45	46	57	-BSY
+ACK	24	47	48	58	-ACK
+RST	25	49	50	59	-RST
+MSG	26	51	52	60	-MSG
+SEL	27	53	54	61	-SEL
+C/D	28	55	56	62	-C/D
+REQ	29	57	58	63	-REQ
+I/O	30	59	60	64	-I/O
+DB(8)	31	61	62	65	-DB(8)
+DB(9)	32	63	64	66	-DB(9)
+DB(10)	33	65	66	67	-DB(10)
+DB(11)	34	67	68	68	-DB(11)

Note: The “connector pin” in the table refers to the conductor position when using flat-ribbon cable.

Appendix B

Audible Warnings

The MegaRAID SCSI 320-1 RAID controller has an onboard tone generator that indicates events and errors.

Note: This is available only if the optional series 502 Battery Backup Unit (BBU) is installed.

Table B.1 Audible Warnings and Descriptions

Tone Pattern	Meaning	Examples
Three seconds on and one second off	A logical drive is offline.	One or more drives in a RAID 0 configuration failed. Two or more drives in a RAID 1, or 5 configuration failed.
One second on and one second off	A logical drive is running in degraded mode.	One drive in a RAID 5 configuration failed.
One second on and three seconds off	An automatically initiated rebuild has been completed.	While you were away from the system, a disk drive in a RAID 1, or 5 configuration failed and was rebuilt.

Appendix C

Glossary

Array	A grouping of individual disk drives that combines the storage space on the disk drives into a single segment of contiguous storage space. MegaRAID can group disk drives on one or more SCSI channels into an array.
Array Management Software	Software that provides common control and management for a disk array. Array management software most often executes in a disk controller or intelligent host bus adapter, but it can also execute in a host computer. When array management software executes in a disk controller or adapter, it is often called <i>firmware</i> .
Array Spanning	Array spanning by a logical drive combines storage space in two arrays of disk drives into a single, contiguous storage space in a logical drive. MegaRAID logical drives can span consecutively numbered arrays that have the same number of disk drives. Array spanning promotes RAID levels 1 and 5 to RAID levels 10 and 50, respectively. See also <i>Disk Spanning</i> .
Asynchronous Operations	Operations that bear no relationship to each other in time and can overlap. The concept of asynchronous I/O operations is central to independent access arrays in throughput-intensive applications.
Cache I/O	A small amount of fast memory that holds recently accessed data. Caching speeds subsequent access to the same data. It is most often applied to processor-memory access, but it can also be used to store a copy of data accessible over a network. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), the data is read from cache memory immediately and the main memory read is aborted (or not started.) If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.

Channel	An electrical path for the transfer of data and control information between a disk and a disk controller.
Consistency Check	An examination of the disk system to determine whether all conditions are valid for the specified configuration (such as parity.)
Cold Swap	A cold swap requires that you turn the power off before replacing a defective disk drive in a disk subsystem.
Data Transfer Capacity	The amount of data per unit time moved through a channel. For disk I/O, data transfer capacity (bandwidth) is expressed in megabytes per second (Mbytes/s).
Degraded	Used to describe a disk drive that has become non-functional or has decreased in performance.
Disk	A nonvolatile, randomly addressable, rewritable mass storage device, including both rotating magnetic and optical disks and solid-state disks, or nonvolatile electronic storage elements. It does not include specialized devices such as write-once-read-many (WORM) optical disks, nor does it include so-called RAM disks implemented using software to control a dedicated portion of a host computer volatile random access memory.
Disk Array	A collection of disks from one or more disk subsystems combined with array management software. The software controls the disks and presents them to the array operating environment as one or more virtual disks.
Disk Duplexing	A variation on disk mirroring in which a second disk adapter or host adapter and redundant disk drives are present.
Disk Mirroring	Writing duplicate data to more than one (usually two) disk drives to protect against data loss in the event of device failure. Disk mirroring is a common feature of RAID systems.
Disk Spanning	Disk spanning allows multiple disk drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 36 Gbyte disk drives can be combined to appear to the operating system as one single 144 Gbyte drive. See also <i>Array Spanning</i> and <i>Spanning</i> .
Disk Striping	A type of disk array mapping. Consecutive stripes of data are mapped round-robin to consecutive array members. A striped array (RAID 0)

provides high I/O performance at low cost, but provides lower data reliability than any of its member disks.

Disk Subsystem	A collection of disks and the hardware that connects them to one or more host computers. The hardware can include an intelligent controller, or the disks can attach directly to a host computer I/O bus adapter.
Double Buffering	A technique that achieves maximum data transfer bandwidth by constantly keeping two I/O requests for adjacent data outstanding. A software component begins a double-buffered I/O stream by issuing two requests in rapid sequence. Thereafter, each time an I/O request completes, another is immediately issued. If the disk subsystem is capable of processing requests fast enough, double buffering allows data to be transferred at the full-volume transfer rate.
Failed Drive	A drive that has ceased to function or consistently functions improperly.
Fast SCSI	A variant on the SCSI-2 standard. It uses the same 8-bit bus as the original SCSI-1, but runs at up to 10 Mbytes (double the speed of SCSI-1.)
Firmware	Software stored in read-only memory (ROM) or Programmable ROM (PROM). Firmware often controls the behavior of a system when it is first turned on. A typical example would be a monitor program in a computer that loads the full operating system from disk or from a network and then passes control to the operating system.
FlexRAID Power Fail Option	The FlexRAID Power Fail option allows a reconstruction to restart if a power failure occurs. This is the advantage of this option. The disadvantage is, once the reconstruction is active, the performance is slower because an additional activity is added.
Format	The process of writing zeros to all data fields in a physical drive (disk drive) to map out unreadable or bad sectors. Because most disk drives are factory formatted, formatting is usually only done if a hard disk generates many media errors.
GByte	Gigabyte. Shorthand for 1,000,000,000 (10 to the ninth power) bytes. One Gbyte is equivalent to 1,000 Mbytes.
Host-based Array	A disk array with an array management software in its host computer rather than in a disk subsystem.

Host Computer	Any computer to which disks are directly attached. Mainframes, servers, workstations, and personal computers can all be considered host computers.
Hot Spare	A stand-by disk drive ready for use if a drive in an array fails. A hot spare does not contain any user data. Up to eight disk drives can be assigned as hot spares for an adapter. A hot spare can be dedicated to a single redundant array, or it can be part of the global hot-spare pool for all arrays controlled by the adapter.
Hot Swapping	The substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running. Hot swaps are done manually.
I/O Driver	A host computer software component (usually part of the operating system) that controls the operation of peripheral controllers or adapters attached to the host computer. I/O drivers communicate between applications and I/O devices, and in some cases they participate in data transfer.
Initialization	The process of writing zeros to the data fields of a logical drive and generating the corresponding parity to put the logical drive in a Ready state. Initializing erases previous data and generates parity so that the logical drive will pass a consistency check. Arrays can work without initialization, but they can fail a consistency check because the parity fields have not been generated.
Logical Disk	A set of contiguous chunks on a physical disk. Logical disks are used in array implementations as constituents of logical volumes or partitions. Logical disks are normally transparent to the host environment, except when the array containing them is being configured.
Logical Drive	A virtual drive within an array that can consist of more than one physical drive. Logical drives divide the contiguous storage space of an array of disk drives or a spanned group of arrays of drives. The storage space in a logical drive is spread across all the physical drives in the array or spanned arrays. Configure at least one logical drive for each array.
Mapping	The conversion between multiple data addressing schemes, especially conversions between member disk block addresses and block addresses of the virtual disks presented to the operating environment by array management software.

Mbyte	(Megabyte) An abbreviation for 1,000,000 (10 to the sixth power) bytes. One Mbyte equals 1,000 Kbytes (kilobytes).
Multi-threaded	Having multiple concurrent or pseudo-concurrent execution sequences. Used to describe processes in computer systems. Multi-threaded processes allow throughput-intensive applications to efficiently use a disk array to increase I/O performance.
Operating Environment	The operating environment includes the host computer where the array is attached, any I/O buses and adapters, the host operating system, and any additional software required to operate the array. For host-based arrays, the operating environment includes I/O driver software for the member disks, but does not include array management software, which is regarded as part of the array itself.
Parity	Parity is an extra bit added to a byte or word to reveal errors in storage (in RAM or disk) or transmission. Parity is used to generate a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets, although it does not fully duplicate the parent data sets. In RAID, this method is applied to entire drives or stripes across all disk drives in an array. Parity data is distributed among all the drives in the system. If a single disk drive fails, the drive can be rebuilt from the parity of the respective data on the remaining drives.
Partition	An array virtual disk made up of logical disks rather than physical ones. Also known as logical volume.
Physical Disk	A hard disk drive that stores data. A hard disk drive consists of one or more rigid magnetic discs rotating about a central axle with associated read/write heads and electronics.
Physical Disk Roaming	The ability of some adapters to detect when disk drives have been moved to a different slot in the computer—for example, after a hot swap.
Protocol	A set of formal rules describing how to transmit data, especially across a network. Low level protocols define the electrical and physical standards to be observed, bit- and byte- ordering, and the transmission and error detection and correction of the bit stream. High level protocols deal with the data formatting, including the message syntax, the terminal-to-computer dialogue, character sets, and sequencing of messages.

RAID	Redundant Array of Independent Disks. A data storage method in which data, along with parity information, is distributed among two or more hard disks (called an <i>array</i>) to improve performance and reliability. A RAID disk subsystem improves I/O performance on a server using only a single drive. The RAID array appears to the host server as a single storage unit. I/O is expedited because several disks can be accessed simultaneously.
RAID Levels	A style of redundancy applied to a logical drive. It can increase the performance of the logical drive and can decrease usable capacity. Each logical drive must have a RAID level assigned to it. The RAID level drive requirements are: RAID 0 requires one or more physical drives, RAID 1 requires exactly two physical drives, RAID 5 requires at least three physical drives. RAID levels 10 and 50 result when logical drives span arrays. RAID 10 results when a RAID 1 logical drive spans arrays. RAID 50 results when a RAID 5 logical drive spans arrays.
RAID Migration	RAID migration is used to move between optimal RAID levels or to change from a degraded redundant logical drive to an optimal RAID 0. In Novell, the utility used for RAID migration is MEGAMGR; in Windows NT, it is Power Console Plus. If a RAID 1 array is being converted to a RAID 0 array, instead of performing RAID migration, one drive can be removed and the other reconfigured on the controller as a RAID 0. This is due to the same data being written to each drive.
Read-Ahead	A memory caching capability in some adapters that allows them to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the additional data will be needed soon. Read-Ahead supplies sequential data faster, but it is not as effective when accessing random data.
Ready State	A condition in which a workable disk drive is neither online nor a hot spare and is available to add to an array or to designate as a hot spare.
Rebuild	The regeneration of all data from a failed disk in a RAID level 1, or 5 array to a replacement disk. A disk rebuild normally occurs without interruption of application access to data stored on the array virtual disk.
Rebuild Rate	The percentage of CPU resources devoted to rebuilding.
Reconstruct	The act of remaking a logical drive after changing RAID levels or adding a physical drive to an existing array.

Redundancy	The provision of multiple interchangeable components to perform a single function to cope with failures or errors. Redundancy normally applies to hardware; disk mirroring is a common form of hardware redundancy.
Replacement Disk	A disk available to replace a failed member disk in a RAID array.
Replacement Unit	A component or collection of components in a disk subsystem that is always replaced as a unit when any part of the collection fails. Typical replacement units in a disk subsystem includes disks, controller logic boards, power supplies, and cables. Also called a hot spare.
SAF-TE	SCSI-accessed fault-tolerant enclosure. An industry protocol for managing RAID enclosures and reporting enclosure environmental information.
SCSI	(Small computer system interface) A processor-independent standard for system-level interfacing between a computer and intelligent devices, including hard disks, floppy disks, CD-ROM, printers, scanners, etc. Multiple SCSI devices can be connected to a single host adapter on the computer's bus. SCSI transfers eight or 16 bits in parallel and can operate in either asynchronous or synchronous modes. The synchronous transfer rate is up to 320 Mbytes/s (for Ultra320 SCSI). SCSI connections normally use single-ended drivers, as opposed to differential drivers.
SCSI Channel	MegaRAID controls the disk drives through SCSI-2 buses (channels) over which the system transfers data in either Fast and Wide or Ultra SCSI mode. Each adapter can control up to three SCSI channels. Internal and external disk drives can be mixed on channels 0 and 1, but not on channel 2.
SCSI Drive State	A SCSI physical drive can be in one of these states: <ul style="list-style-type: none"> • Online: Powered-on and operational. • Ready: Functioning normally, but not part of a configured logical drive and not designated as a hot spare. • Hot Spare: Powered-on stand-by disk drive, ready for use if an online disk fails. • Fail: Out of service, due to a fault occurring on the drive. • Rebuild: Currently being rebuilt with data from a failed drive.

Service Provider	The Service Provider (SP) is a program that resides in the desktop system or server and is responsible for all DMI activities. This layer collects management information from products (whether system hardware, peripherals, or software), stores that information in the DMI database, and passes it to management applications as requested.
SNMP	Simple Network Management Protocol. The most widely used protocol for communicating management information between the managed elements of a network and a network manager. SNMP focuses primarily on the network backbone. The Internet standard protocol that manages nodes on an Internet Protocol (IP) network.
Spanning	Array spanning by a logical drive combines storage space in two arrays of disk drives into a single, contiguous storage space in a logical drive. MegaRAID logical drives can span consecutively numbered arrays that each consist of the same number of disk drives. Array spanning promotes RAID levels 1 and 5 to RAID levels 10 and 50, respectively. See also <i>Disk Spanning</i> .
Spare	A disk drive available to back up the data of other drives.
Stripe Size	The amount of data contiguously written to each disk. You can specify stripe sizes of 4, 8, 16, 32, 64, and 128 Kbytes for each logical drive. For best performance, choose a stripe size equal to or smaller than the block size used by the host computer.
Stripe Width	The number of disk drives across which the data is striped.
Striping	Segmentation of logically sequential data, such as a single file, so that segments can be written to multiple physical devices in a round-robin fashion. This technique is useful if the processor can read or write data faster than a single disk can supply or accept it. While data is being transferred from the first disk, the second disk can locate the next segment. Data striping is used in some modern databases and in certain RAID devices.
Terminator	A resistor connected to a signal wire in a bus or network for impedance matching to prevent reflections—for example, a 50 ohm resistor connected across the end of an Ethernet cable. SCSI buses and some LocalTalk wiring schemes also require terminators.

- Ultra320** A subset of Ultra3 SCSI that allows a maximum throughput of 320 Mbytes/s, which is twice as fast as Wide Ultra2 SCSI. Ultra320 SCSI provides 320 Mbytes/s on a 16-bit connection.
- Virtual Sizing** FlexRAID virtual sizing is used to create a logical drive up to 80 Gbytes. A maximum of 40 logical drives can be configured on a RAID controller, and RAID migration is possible for all logical drives except the fortieth. Because it is not possible to do migration on the last logical drive, the maximum space available for RAID migration is 560 Gbytes.
- Wide SCSI** A variant on the SCSI-2 interface. Wide SCSI uses a 16-bit bus, double the width of the original SCSI-1. Wide SCSI devices cannot be connected to a SCSI-1 bus. Like Fast SCSI, Wide SCSI supports transfer rates up to 20 Mbytes/s.

Index

Numerics

- 0 DIMM socket [6-4](#)
- 160M and Wide SCSI [4-1](#)
- 68-pin connector pinout for LVD SCSI [A-9](#)
- 68-Pin High Density Connectors [A-1](#)

A

- AMI Part Number
 - Battery [6-13](#)
- AMPLIMITE .050 Series connectors [A-8](#)
- Array [C-1](#)
- Array Configuration Planner [5-15](#)
- Array management software [C-1](#)
- Array performance features [4-5](#)
- Array spanning [C-1](#)
- Assigning Drive Letters [7-15](#)
- Assigning RAID levels [5-12](#)
- Asynchronous operations [C-1](#)
- Audible Warnings [B-1](#)
- Automatic failed drive detection and rebuild [4-6](#)

B

- Battery Pack [6-13](#)
- BIOS Boot Error Messages [8-3](#)
- BIOS Configuration Utility [6-20](#)
- BIOS message [6-20](#)
- Bus Data Transfer Rate [4-7](#)
- Bus Type [4-7](#)
- Bus-based [2-11](#)

C

- Cable assembly for internal Wide SCSI devices [A-2](#)
- Cables To Go [A-6](#)
- Cache Configuration [4-7](#)
- Cache I/O [C-1](#)
- Cache Memory [4-9](#)
 - Installing [6-4](#)
- Card Size [4-7](#)
- Changing DRAM Modules [6-15](#)
- Changing the Battery Pack [6-15](#)
- Channel [C-2](#)
- Charging the Battery Pack [6-15](#)
- Cluster Configuration Wizard [7-16](#)
- Cluster Disks
 - Configuration [7-18](#)
- Cluster Installation [7-5](#)
 - Overview [7-5](#)

- Validation [7-25](#)
- Cluster Node Network Adapter
 - Configuration [7-8](#)
- Cluster Service
 - Assigning Drive Letters [7-15](#)
 - Cluster Node Network Adapter [7-8](#)
 - Cluster User Account [7-12](#)
 - Configuring Cluster Disks [7-18](#)
 - Connectivity and Name Resolution [7-11](#)
 - Disk Access and Functionality [7-15](#)
 - Domain Membership [7-12](#)
 - SCSI Drive Installations [7-27](#)
 - Setting Up Networks [7-7](#)
 - Shared Disks Configuration [7-14](#)
 - Shared Disks Setup [7-13](#)
 - Software Installation [7-16](#)
 - Validating the Cluster Installation [7-25](#)
- Cluster User Account
 - Setup [7-12](#)
- Clustering
 - Network Requirements [7-4](#)
 - Shared Disk Requirements [7-5](#)
- Cold swap [C-2](#)
- Compatibility [4-12](#)
- Components [4-8](#)
- Configuration on Disk [4-3](#)
- Configuration Strategies [5-10](#)
- Configuring Arrays [5-9](#)
- Configuring Logical Drives [5-13](#)
- Configuring SCSI Physical Drives
 - Distributing Drives [5-1](#)
- Configuring SCSI physical drives [5-1](#)
- Connecting internal and external Wide devices [A-3](#)
- Consistency check [2-4](#), [C-2](#)
- Converting from internal Wide to internal non-Wide (Type 3) [A-5](#)
- Converting internal Wide to internal non-Wide [A-4](#)
- Converting internal Wide to internal non-Wide (Type 30) [A-5](#)
- CPU [4-8](#)
- Creating hot spares [5-10](#)
- Creating logical drives [5-10](#)

D

- Data redundancy
 - Using mirroring [2-6](#)
- Data transfer capacity [C-2](#)
- Dedicated parity [2-8](#)
- Degraded [2-10](#), [C-2](#)
- Devices per SCSI Channel [4-8](#)
- DIMM socket [6-4](#)
- DIMMs [6-4](#)

- Disconnect/reconnect [4-10](#)
- Disk [C-2](#)
- Disk Access and Functionality [7-15](#)
- Disk array [C-2](#)
- Disk Array Types [2-11](#)
- Disk duplexing [C-2](#)
- Disk mirroring [2-6](#), [C-2](#)
- Disk rebuild [2-9](#)
- Disk spanning [2-7](#), [C-2](#)
- Disk striping [2-4](#), [C-2](#)
- Disk subsystem [C-3](#)
- Disposing of a Battery Pack [6-15](#)
- Documentation [1-3](#)
- DOS [4-7](#)
- Double buffering [C-3](#)
- Drive roaming [4-3](#)
- Drivers [6-21](#)

E

- Enclosure management [2-11](#)

F

- Fail [2-10](#)
- Failed [2-10](#)
- Failed drive [C-3](#)
- Fast SCSI [C-3](#)
- Fault tolerance [2-4](#)
- Fault tolerance features [4-6](#)
- Features [4-1](#)
- Firmware [4-7](#), [C-3](#)
- Flash ROM [1-2](#)
- FlexRAID Power Fail option [C-3](#)
- Format [C-3](#)

G

- GB [C-3](#)
- Glossary [C-1](#)
- GWC [A-6](#)

H

- Hardware Installation [6-1](#)
- High-density 68-pin SCSI connector pinout [A-7](#)
- High-density single-ended connector [A-8](#)
- Host computer [C-4](#)
- Host-based array [C-3](#)
- Host-based RAID solution [2-2](#)
- Hot spare [2-8](#), [2-10](#), [C-4](#)
 - Using during disk rebuild [2-9](#)
- Hot swap [2-9](#), [C-4](#)

I

- I/O driver [C-4](#)
- Initialization [C-4](#)
- Install Cache Memory [6-4](#)
- Install Drivers [6-21](#)
- Installation steps
 - Custom [6-2](#)

J

- J10 External Battery [6-7](#)
- Jumpers
 - Setting [6-5](#)

L

- Linux
 - Red Hat [4-7](#)
- Logical disk [C-4](#)
- Logical drive [2-3](#), [C-4](#)
- Logical drive states [2-10](#)

M

- Mapping [C-4](#)
- MB [C-5](#)
- MegaRAID BIOS [4-9](#)
- MegaRAID BIOS Configuration Utility [4-11](#)
- MegaRAID Manager [4-12](#), [6-15](#)
- MegaRAID SCSI 320-1 card
 - Installing [6-16](#)
- MegaRAID SCSI 320-1 card layout [6-5](#)
- MegaRAID Specifications
 - Bus Data Transfer Rate [4-7](#)
 - Bus Type [4-7](#)
 - Cache Configuration [4-7](#)
 - Card Size [4-7](#)
 - Devices per SCSI Channel [4-8](#)
 - Firmware [4-7](#)
 - Nonvolatile RAM [4-8](#)
 - Operating Voltage [4-8](#)
 - Processor [4-7](#)
 - RAID Levels Supported [4-8](#)
 - SCSI Bus [4-8](#)
 - SCSI cables [4-8](#)
 - SCSI Connectors [4-8](#)
 - SCSI Controller [4-8](#)
 - SCSI Data Transfer Rate [4-8](#)
 - SCSI Device Types Supported [4-8](#)
 - Serial Port [4-8](#)
 - Termination Disable [4-8](#)
- Mirroring [2-6](#)
- MS-DOS [6-21](#)
- Multi-threaded [C-5](#)
- Multi-threading [4-11](#)

N

- Nonvolatile RAM [4-8](#)
- Novell NetWare [4-7](#), [6-21](#)
- NVRAM [1-2](#)

O

- Offline [2-10](#)
- Onboard Speaker [4-9](#)
- Online
 - Drive state [2-10](#)
- Operating environment [C-5](#)
- Operating system software drivers [4-7](#)
- Operating Voltage [4-8](#)
- Optimal [2-10](#)
- Optimizing Data Storage [5-13](#)

OS/2 2.x 4-12
Other BIOS error messages 8-5

P

Parity 2-8, C-5
Partition C-5
Physical array 2-3
Physical disk C-5
Physical disk roaming C-5
Physical drive 2-3
Power Console Plus 4-12, 6-15
Processor 4-7
Product specifications 4-7
Protocol C-5

R

RAID C-6
 Benefits 2-1
 Introduction to 2-1
RAID 0 3-2
RAID 1 3-3
 Spanning to configure RAID 10 2-8
RAID 10 3-6
 Configuring 2-8
RAID 5 3-4
 Spanning to make RAID 50 2-8
RAID 50 3-7
 Configuring 2-8
RAID benefits
 Improved I/O 2-1
 Increased reliability 2-2
RAID levels 1-v, 3-1, C-6
RAID Levels Supported 4-8
RAID management 4-11
RAID management features 4-5
RAID migration C-6
RAID overview 2-3
Read-ahead C-6
Ready 2-10
Ready state C-6
Rebuild 2-10, 4-11
Rebuild rate 2-10, C-6
Rebuilding a disk 2-9
Reconnect 4-10
Reconstruct C-6
Reconstruction C-6
Red Hat Linux 6-21
RedAlert 4-12
Redundancy C-7
Replacement disk C-7
Replacement unit C-7

S

SAF-TE C-7
Scatter/gather 4-10
SCO Unix 4-12
SCSI C-7
SCSI backup and utility software 4-12
SCSI Bus 4-8, 4-10
SCSI bus 4-10
SCSI bus widths and maximum throughput 1-3
SCSI buses 1-2

SCSI Cable Vendors A-6
SCSI cables 4-8
 Connecting 6-17
SCSI cables and connectors A-1
SCSI channel C-7
SCSI Connector Vendors A-6
SCSI Connectors 4-8
SCSI connectors 4-10
SCSI Controller 4-8
SCSI Data Transfer Rate 4-8
SCSI Device Compatibility 4-12
SCSI Device Types Supported 4-8
SCSI Devices
 Configuration 7-28
SCSI Drive Installations 7-27
SCSI Drive State C-7
SCSI firmware 4-10
SCSI Termination 4-8
SCSI termination 4-10, 6-8
 Set 6-8
SCSI terminator power (TermPWR
 Setting) 6-13
SCSI to SCSI 2-11
SCSI-to-SCSI RAID product 2-3
Serial Port 4-8
Serial port 4-9
Server management 4-12
Service provider C-8
Setting SCSI termination 6-8
Shared Disks
 Configuration 7-14
 Setup 7-13
Shared SCSI Bus
 Termination 7-28
SMART Technology 4-3
SNMP C-8
SNMP agent 4-12
SNMP managers 4-12
Software utilities 4-6
Software-based 2-11
Spanning 2-7, C-8
Spare C-8
Standby rebuild 2-9
Stripe size 2-6, 4-11, C-8
Stripe width 2-5, C-8
Striping C-8
System Connection A-6

T

Tagged command queuing 4-10
Target IDs
 Setting 6-18
Technical Cable Concepts A-6
Technical Support 1-vi
Termination Disable 4-8
Terminator C-8
Troubleshooting 8-1

U

Ultra3-SCSI (320M) C-9
UnixWare 4-12
Unpack 6-3

V

Virtual sizing [C-9](#)

W

WebBIOS Configuration Utility [4-11](#)

Wide SCSI [C-9](#)

Windows .NET [6-21](#)

Windows 2000 [6-21](#)

Windows NT [4-7](#), [6-21](#)

Windows XP [6-21](#)

Write-back caching [6-15](#)

Customer Feedback

We would appreciate your feedback on this document. Please copy the following page, add your comments, and fax it to us at the number shown.

If appropriate, please also fax copies of any marked-up pages from this document.

Important: Please include your name, phone number, fax number, and company address so that we may contact you directly for clarification or additional information.

Thank you for your help in improving the quality of our documents.

Reader's Comments

Fax your comments to: LSI Logic Corporation
Technical Publications
M/S E-198
Fax: 408.433.4333

Please tell us how you rate this document: *MegaRAID SCSI 320-2 Hardware Guide*. Place a check mark in the appropriate blank for each category.

	Excellent	Good	Average	Fair	Poor
Completeness of information	___	___	___	___	___
Clarity of information	___	___	___	___	___
Ease of finding information	___	___	___	___	___
Technical content	___	___	___	___	___
Usefulness of examples and illustrations	___	___	___	___	___
Overall manual	___	___	___	___	___

What could we do to improve this document?

If you found errors in this document, please specify the error and page number. If appropriate, please fax a marked-up copy of the page(s).

Please complete the information below so that we may contact you directly for clarification or additional information.

Name _____ Date _____

Telephone _____ Fax _____

Title _____

Department _____ Mail Stop _____

Company Name _____

Street _____

City, State, Zip _____

Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>