Bull ESCALA

SCSI RAID Adapter

Installation & Configuration Guide

ORDER REFERENCE
86 A1 44HX 02

Bull ESCALA

SCSI RAID Adapter
Installation & Configuration Guide

Hardware

April 2000

BULL ELECTRONICS ANGERS
CEDOC
34 Rue du Nid de Pie – BP 428
49004 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A1 44HX 02

Printed in France

Suggestions and criticisms concerning the form, content, and presentation of this book are invited. A form is provided at the end of this book for this purpose.

To order additional copies of this book or other Bull Technical Publications, you are invited to use the Ordering Form also provided at the end of this book.

Trademarks and Acknowledgements

We acknowledge the right of proprietors of trademarks mentioned in this book.

AIX® is a registered trademark of International Business Machines Corporation, and is being used under licence.

UNIX is a registered trademark in the United States of America and other countries licensed exclusively through the Open Group.

Year 2000

The product documented in this manual is Year 2000 Ready.

About This Book

This document explains the installation and configuration of the SCSI RAID Adapters in PCI bus computers.

Who Should Use this Book

This book addresses those who intend to install the adapter and configure the system.

Operating Systems

This Adapter operates in the AIX environment.

Document Overview

This book contains the following chapters:

Chapter 1 **SCSI RAID Adapter – Overview**
Introduces the SCSI RAID Adapters.

Chapter 2 **SCSI RAID Adapter – Installation**
Configuring the SCSI peripheral devices and cables. Then installing memory modules and the SCSI RAID Adapters.

Chapter 3 **SCSI RAID Adapter – Operating System Installation**
Installing the AIX Operating System.

Chapter 4 **SCSI RAID Adapter – Storage Devices**
Describes the RAID features, SCSI devices and accessories and gives instructions on configuring complete RAID subsystems using these products.

Chapter 5 **SCSI RAID Adapter – Storage Manager**
Installing and using the Storage Manager.

Chapter 6 **SCSI RAID Adapter – Theory of Operation**
Describes the Theory of Operation.

Appendix A **B4–4 Adapter Troubleshooting**
Troubleshooting Guide providing answers to many commonly asked questions.

Appendix A **B4–C Adapter Troubleshooting**
Troubleshooting Guide providing answers to many commonly asked questions.

Glossary Alphabetical list of terms and abbreviations used in this manual.

Index General index.

The following highlighting conventions are used in this book:

Bold Identifies commands, keywords, files, directories, or other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.

Italics Identifies parameters whose actual names or values are to be supplied by the user.

`Monospace` Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

Related Publications

Documentation for the PC–supported part of this product is provided in the form of on–line help, integrated with the software.

Other Publications

Cabling Guide for Multiple Bus Systems	86 A170JX
Adapters for Multiple Bus Systems	86A127HX
AIX and Related Products Documentation Overview	86 A2 71WE

Ordering Publications

You can order publications from your sales representative or from your point of sale.

If you received a printed copy of *Documentation Overview* with your system, use that book for information on related publications and for instructions on ordering them.

To order additional copies of this book, use order number 86 A1 44HX.

Communication Statements

The following statement applies to this product. The statement for other products intended for use with this product appears in their accompanying manuals.

Federal Communications Commission (FCC) Statement

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 particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from authorized dealers. Neither the provider nor the manufacturer are responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

European Union (EU) Statement

This product is in conformity with the protection requirements of EC Council Directives 89/336/EEC relating to electromagnetic compatibility, and with 73/23/EEC (Low Voltage Directive).

Label:

CE

Neither the provider nor the manufacturer can accept responsibility for any failure to satisfy the protection requirements resulting from a non–recommended modification of the product, including the fitting of option cards not supplied by the manufacturer.

This product has been tested and found to comply with the limits for Class B Information Technology Equipment according to CISPR 22 / European Standard EN 55022. The limits for Class B equipment were derived for typical residential environments to provide reasonable protection against interference with licensed communication devices.

International Electrotechnical Commission (IEC) Statement

This product has been designed and built to comply with IEC Standard 950.

This equipment is manufactured to the International Safety Standard EN60950 and as such is approved in the UK under the General Approval Number NS/G/1234/J/100003 for indirect connection to the public telecommunication network.

The network adapter interfaces housed within this equipment are approved separately, each one having its own independent approval number. These interface adapters, supplied by the manufacturer, do not use or contain excessive voltages. An excessive voltage is one which exceeds 70.7 V peak ac or 120 V dc. They interface with this equipment using Safe Extra Low Voltages only. In order to maintain the separate (independent) approval of the manufacturer's adapters, it is essential that other optional cards, not supplied by the manufacturer, do not use main voltages or any other excessive voltages. Seek advice from a competent engineer before installing other adapters not supplied by the manufacturer.

Avis de conformité aux normes du ministère des Commmunications du Canada

Cet appareil numérique de la classe B respecte toutes les exigences du Réglement sur le matériel brouilleur du Canada.

Canadian Department of Communications Compliance Statement

This Class B digital apparatus meets the requirements of the Canadian Interference–Causing Equipment Regulations.

この装置は，第二種情報装置（住宅地域又はその隣接した地域において使用されるべき情報装置）で住宅地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）基準に適合しております。

しかし，本装置をラジオ，テレビジョン受信機に近接してご使用になると，受信障害の原因となることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

VCCI Statement

The following is a summary of the VCCI Japanese statement in the box above.

This equipment is in the Class 2 category (information equipment to be used in a residential area or an adjacent area thereto) and conforms to the standards set by the Voluntary Control Council For Interference by Data Processing Equipment and Electronic Office Machines aimed at preventing radio interference in such residential area.

When used near a radio or TV receiver, it may become the cause of radio interference.

Read the instructions for correct handling.

Radio Protection for Germany

Dieses Gerät ist berechtigt in Übereinstimmung mit dem deutschen EMVG vom 9.Nov.92 das EG–Konformitätszeichen zu führen.

Der Aussteller der Konformitätserklärung ist die Bull Germany.

Dieses Gerät erfüllt die Bedingungen der EN 55022 Klasse B.

Table of Contents

About This Book 3

Chapter 1. SCSI RAID Adapters – Overview **1-1**
SCSI RAID Adapter – Introduction 1-1
Components 1-1
SCSI RAID Adapter Features 1-1
SCSI RAID PCI B4–4 Adapter Specifications 1-2
Ultra2 SCSI High Performance RAID B4–C Adapter Specifications 1-2
SCSI RAID Adapters 1-2
SCSI Devices 1-2
Storage Manager 1-3

Chapter 2. SCSI RAID Adapter – Installation **2-1**
Installation Overview 2-1
B4–4 Adapter Installation 2-1
B4–C Adapter Installation 2-1
Installing Additional SIMMs on a B4–4 Adapter 2-2
Installing Additional DIMMs on a B4–C Adapter 2-3
SCSI Termination 2-5
Configuring SCSI Device IDs 2-7
Hard Disk Drive SCSI-ID on a Deskside 2-7
Hard Disk Drive SCSI-ID on a Rack 2-8
Installing the Adapter 2-9
SCSI RAID Adapter LEDs 2-10
SCSI RAID Adapter LEDs on the B4–4 Adapter 2-10
SCSI RAID Adapter LEDs on the B4–C Adapter 2-11
Audible Alarm on the B4–C Adapter 2-11
Upgrading Flash Firmware 2-12
Upgrading Flash Firmware on the B4–4 Adapter 2-12
Upgrading Flash Firmware on the B4–C Adapter 2-12
NVRAM (on the B4–C Adapter only) 2-12

Chapter 3. SCSI RAID Adapter – O.S. Installation **3-1**
Installing the AIX Operating System 3-1
Packaging Installation 3-1
Adding Hard Drives 3-3
Hard Drives Recognition 3-3

Chapter 4. SCSI RAID Adapter – Storage Devices **4-1**
RAID Subsystem Configuration 4-1
Determining the RAID Level 4-1
Selecting Hard Drives 4-2
Drives LED Management 4-3
Replacing a Drive 4-3

Chapter 5. SCSI RAID Adapter – Storage Manager **5-1**
Introduction 5-1
Installing Storage Manager 5-1
Storage Manager GUI 5-1
Storage Manager Keyboard Functions 5-2
Running Storage Manager 5-2
Physical Configuration View 5-3
Logical Configuration View 5-4
Information Windows 5-5
HBA Information Windows 5-5
SCSI Device Information Windows 5-9
Array Groups 5-12
Creating an Array Group 5-13
Modifying an Array Group 5-15
Deleting an Array Group 5-15
Logical Device Address 5-16
Status 5-17
Events 5-18
Event Logging 5-18
Broadcaster 5-19
Broadcaster for AIX 5-19
Hard Drive Formatting 5-20
Formatting Hard Drives 5-20
Hard Drive Failures 5-21
Audible Alarm 5-21
Hot Swapping Drives 5-21
Hot Spares 5-21
Diagnostics for a B4–4 Adapter 5-23
SCSI Device Diagnostics 5-23
Array Diagnostics for a B4–4 Adapter 5-24
Scheduling Diagnostics for a B4–4 Adapter 5-25
Background Task Priority 5-26
I/O Monitor 5-27
Battery Backup Configuration for a B4–C Adapter 5-29
Battery Status Messages 5-29
Initial Calibration 5-30
Flash Configuration for a B4–C Adapter 5-31
Setting the Configuration 5-32

Chapter 6. SCSI RAID Adapter – Theory of Operation 6-1
Disk Caching 6-1
Software and Hardware Caching 6-1
Elevator Sorting 6-3
Write–back Algorithm 6-3
Read–ahead Algorithm 6-4
Optimizing OS Cache Size 6-6
Optimizing Adapter Cache Size 6-6
RAID 6-7
Disk Drive Striping 6-7
Definition of RAID Levels 6-9
Creating Data Redundancy 6-13
Degraded Mode 6-14
Rebuilding a Failed Hard Drive 6-15
SCSI RAID Adapter Architecture 6-16
EATA Protocol 6-16
PCI Data Transfer 6-18
PCI Bus Mastering 6-18
SCSI Interface 6-19
SCSI Device IDs 6-19
LUNs 6-19
Commands and Messages 6-19
Disconnect/Reconnect 6-19
Transfer Padding and Residue Reporting 6-19
Single–ended/Differential SCSI 6-20
Wide SCSI 6-20
SCSI Transfer Rate 6-20
Cabling Single–ended SCSI 6-20
SCSI Termination 6-22
TERMPWR 6-22
SCSI Bus ECC Protection 6-23

Appendix A. SCSI RAID B4–4 Adapter – Troubleshooting A-1
SCSI RAID B4–4 Adapter – Troubleshooting A-1

Appendix B. SCSI RAID B4–C Adapter – Troubleshooting B-1
SCSI RAID B4–C Adapter – Troubleshooting B-1

Index X-1

Glossary G-1

Chapter 1. SCSI RAID Adapters – Overview

Introduces the SCSI RAID Adapters.

SCSI RAID Adapter – Introduction

The proposed SCSI product incorporates world–class technology to deliver optimum performance for desktop systems, file servers, and multiuser hosts.

- Support for 2GB (only with B4–4 adapter), 4GB (only with B4–4 adapter), 9GB and 18GB Ultra Wide SCSI hard disk drives.
- SCSI–1, SCSI–2 SCSI–3 and LVD (only on B4–C adapter) support with active termination.
- Top bus mastering performance.
- Hardware scatter/gather support for top performance in virtual memory environments such as AIX.
- Overlapped I/O and tagged command queuing, for highest performance.
- EATA and I2O protocols that provide universal support for third–party applications and utilities.
- Storage Manager software.

Components

Type	Designation
B4-4	Ultra Wide SCSI Adapter
B4-C	ULTRA2/LVD SCSI PCI Adapter

SCSI RAID Adapter Features

SCSI RAID adapters are high–performance adapters for PCI bus computers which support SCSI storage devices in hot–pluggable carriers.

Key features include:

- Hardware RAID levels 0, 1 and 5
- Multi–level RAID and Duplexing
- Support for AIX operating systems
- Up to 64MB of cache per B4–4 adapter on adapter, up to 32MB of cache per B4–C adapter
- Hot–swappable and redundant disk drives
- Hot Spare capability
- SCSI, SCSI-2, SCSI-3, Fast Wide and Ultra Wide SCSI, and ULTRA2/LVD (only on B4–C adapter) support with active termination.
- Event logging
- Local configuration and I/O monitoring
- diagnostics (on B4–C adapters)

SCSI RAID PCI B4–4 Adapter Specifications

Maximum number	3
Bus architecture	PCI
Bus master	Yes
Data transfer rate	132 MB/second across the PCI bus
Connectors	Internal 68-position 16-bit High-Density SCSI bus External 68-position 16-bit High-Density SCSI bus

Ultra2 SCSI High Performance RAID B4–C Adapter Specifications

Maximum number	2
SCSI features	Ultra 2 SCSI–3 interface, maximum SCSI transfer rate 80MB, 16–bit SCSI bus with multi–mode SE/LVD termination, 2 SCSI busses, 30 SCSI devices, external connector, internal connector
Battery backup	72 hours with 16 MB, 2–4 3A NiMh batteries
Bus architecture	64–bit PCI 2.1 with a maximum transfer rate of 264MBs
Power Requirements	5V +/– 5%, 50Mv peak–to–peak max, 1.87 A (single channel)
Reliability	MTBF >350.000 hours
Agency Certification	FCC, CE, UL, CSA, AUSTEL, SMA
Complete Kit Content	PMK3755U2B Ultra 2 SCSI high performance RAID controller, 1 additional adapter for 1 SCSI LVD, external and internal channel, optional battery backup module adds up to 72 hours of protection for cache data in the event of a power failure

SCSI RAID Adapters

SCSI adapters are RAID–ready with integrated hardware caching.

SCSI RAID adapters provide Wide internal or external SCSI connectors and cables.

SCSI RAID adapter with four SIMM or DIMM sockets accept up to 64MB of cache (32MB on B4–C), providing maximum performance and data integrity.

Environmental sensing capabilities are standard on SCSI RAID adapters. On–board temperature and voltage sensors for the server, as well as storage cabinet voltage and fan speed sensors, pass critical server management data to Storage Manager for logging and notification to the system administrator.

SCSI Devices

SCSI devices include hard drives in hot–pluggable carriers designed for insertion in any of the 5 1/4” device bays of the computer.

Hard Drives

2GB (B4–4 only), 4GB (B4–4 only), 9GB or 18GB Ultra Wide SCSI hard drives, enclosed in hot–pluggable carriers, snap into any device bay in the computer.

A fault LED on each hard drive indicates a failed drive. Failed hard drives may be replaced without interruption of normal system operation by simply pulling out the drive carrier. When a replacement drive is inserted, SCSI RAID adapter automatically spins–up the drive and starts the RAID rebuild process. It is not necessary to set switches or run special software.

For immediate automatic replacement of failed drives, unused drives can be designated as Hot Spares. Upon failure of any drive in a RAID–1 or RAID–5 array, the Hot Spare will automatically replace the failed drive and the missing data will be rebuilt onto the Hot Spare without user intervention.

Storage Manager

Every SCSI RAID adapter is shipped with a Storage Manager, an easy–to–use graphical application which performs on–line control, monitoring and analysis. Real–time I/O load analysis, on–line diagnostics, event logging, and device inventory management are available. Storage Manager gives administrators unequalled flexibility and dramatically simplifies and reduces the system administration cost.

- **Hardware Inventory Management** – Storage Manager automatically discovers and inventories all SCSI devices and hardware in your system. Complete information, including each device's type, make, model, operating status, transfer speed, capacity, and configured options, can be viewed on the screen or printed to create a complete record of your installation.

Note: This applies only to SCSI RAID Adapters and attached devices.

- **RAID Configuration** – Storage Manager makes RAID setup and configuration a snap. By following a simple graphical procedure, Storage Manager automatically constructs the RAID 0, 1 or 5 configuration that best meets your requirements.
- **On–line Scheduled Diagnostics** – Storage Manager's diagnostics for hard drives and disk arrays may be executed without taking drives or arrays off–line or interrupting normal system operation. Diagnostics may be scheduled to occur during off–hours so as not to impact system performance during peak–load times.
- **Real–time I/O Analysis** – SCSI RAID adapter products are instrumented to collect I/O statistics. Storage Manager presents this data in real–time in an easy–to–understand format. By analyzing the relative number of disk reads and writes, the distribution of record sizes, and the number of cache hits and misses, network administrators can determine when they need to scale the I/O processing capabilities of their storage subsystem.
- **Event Logging**– SCSI RAID adapter products monitor and record the operating condition of your storage subsystem and maintain a complete log of key events. These event logs may be viewed through a scrolling window and filtered to show only events of interest. Storage Manager's event logging capability shortens diagnosis time and allows the network administrator to take preventive action before users and data are affected.
- **Event Notification** – Storage Manager allows you to designate who will be notified when a failure or other key event occurs. Designated individuals are automatically notified of events through e–mail or network broadcast, allowing for timely response.

Chapter 2. SCSI RAID Adapter – Installation

Installation Overview

B4–4 Adapter Installation

The PCI SCSI RAID B4–4 adapter consists of 3 parts:

1. The SCSI Adapter
2. Additional Single SCSI Bus Ultra Wide
3. SIMM Memory Modules.

These parts are factory–assembled, with, at least, a 16MB memory module, and look like a long format PCI board, see Figure 1, on page 2-2.

The PCI SCSI RAID adapter installation is summarized as follows:

1. Install one or more SIMMs in the SCSI RAID adapter, see page 2-2.
2. Configure SCSI IDs, cables and termination for SCSI devices in the Server, see page 2-5.
3. Install the adapter and SCSI devices in the Server, see page 2-9.
4. Configure RAID subsystems, see page 4-1.

B4–C Adapter Installation

The PCI SCSI RAID B4–4 adapter consists of 4 parts:

1. The SCSI Adapter
2. Additional Single SCSI Bus Ultra Wide/LVD
3. DIMM Memory Modules.
4. Battery Backup Module

These parts are factory–assembled, with a 32MB memory module, and look like a long format PCI board, see Figure 2, on page 2-3.

The PCI SCSI RAID adapter installation is summarized as follows:

1. Install two DIMMs in the SCSI RAID adapter, see page 2-3.
2. Configure SCSI IDs, cables and termination for SCSI devices in the Server, see page 2-5.
3. Install the adapter and SCSI devices in the Server, see page 2-9.
4. Configure RAID subsystems, see page 4-1.

Installing Additional SIMMs on a B4–4 Adapter

SCSI RAID adapters support up to 64MB of cache using four on–board SIMM sockets. One, two, three or four SIMMs may be installed in the adapter. **At least one SIMM must be fitted.**

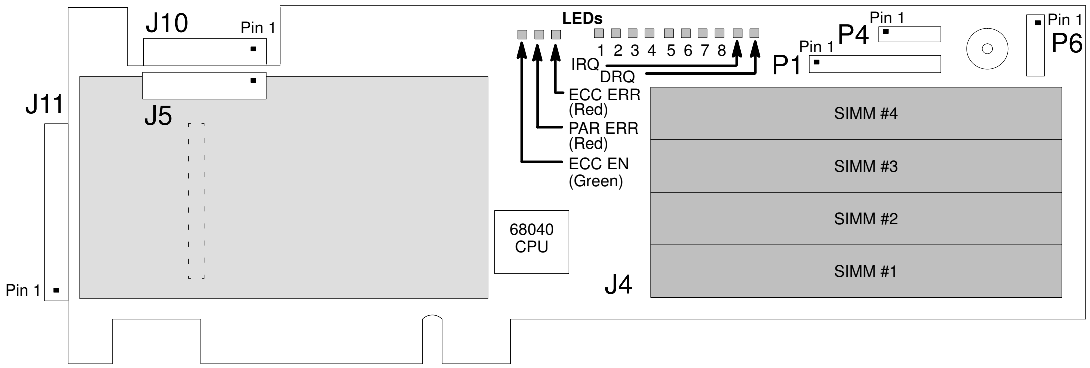

Legend:

- P1 Firmware Flash Clear Jumper, pins 5 & 6
- P4 NVRAM Clear Jumper, pins 1 & 2 / Remote Alarm Connector
- P6 Disk Activity LED Connector
- J10 Internal Wide SCSI Connector (Bus 0)
- J11 External Wide SCSI Connector (Bus 0)
- J5 Internal Wide SCSI Connector (Bus 1)

Figure 1. SIMM Installation

1. Install the SIMMs in the SIMM sockets. Standard 36–bit 4MB and 16MB low–profile single–sided SIMMs may be used.

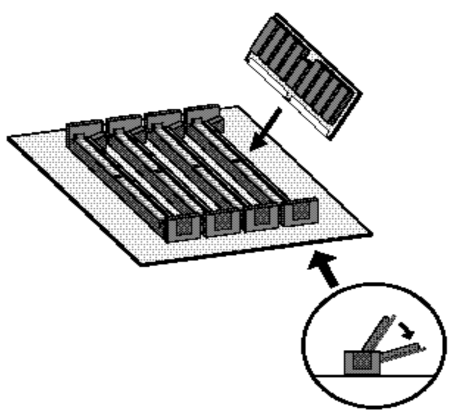

Note: Different SIMM sizes may be mixed, however, Socket #1 must always contain the largest SIMM.

2. To confirm that the SIMMs have been properly installed, the amount of cache memory present will be displayed by Storage Manager when the Host Bus Adapter (HBA) (SCSI adapter) icon is double–clicked.

Note: LED functions are described in SCSI RAID Adapter LEDs, on page 2-10.

Installing Additional DIMMs on a B4–C Adapter

SCSI RAID adapters support up to 64MB of cache using two on–board DIMM sockets. **The two DIMM must be fitted.**

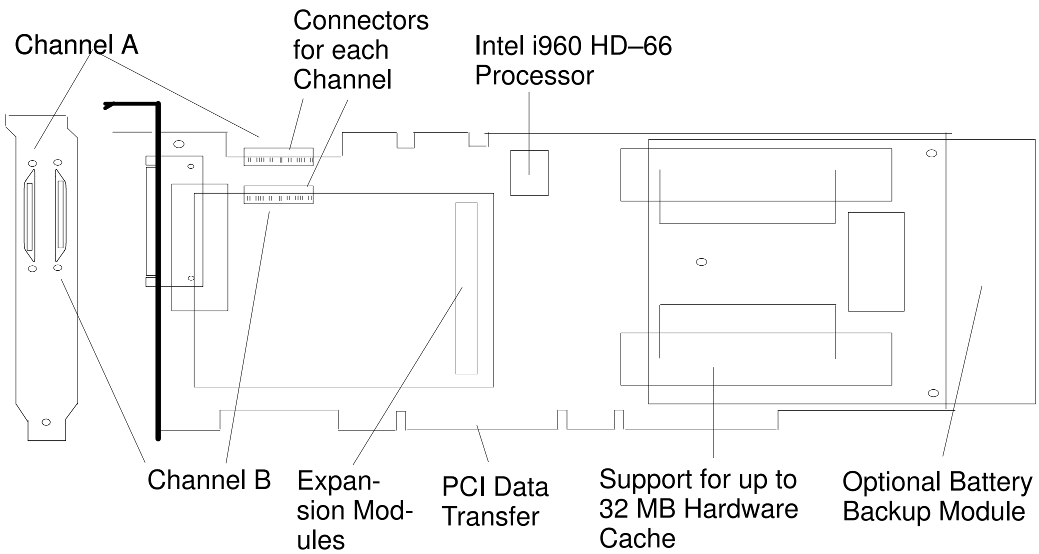

Figure 2. DIMM Installation

1. Unplug the Battery Backup module to have access to DIMMs sockets (3 screws to unfit).

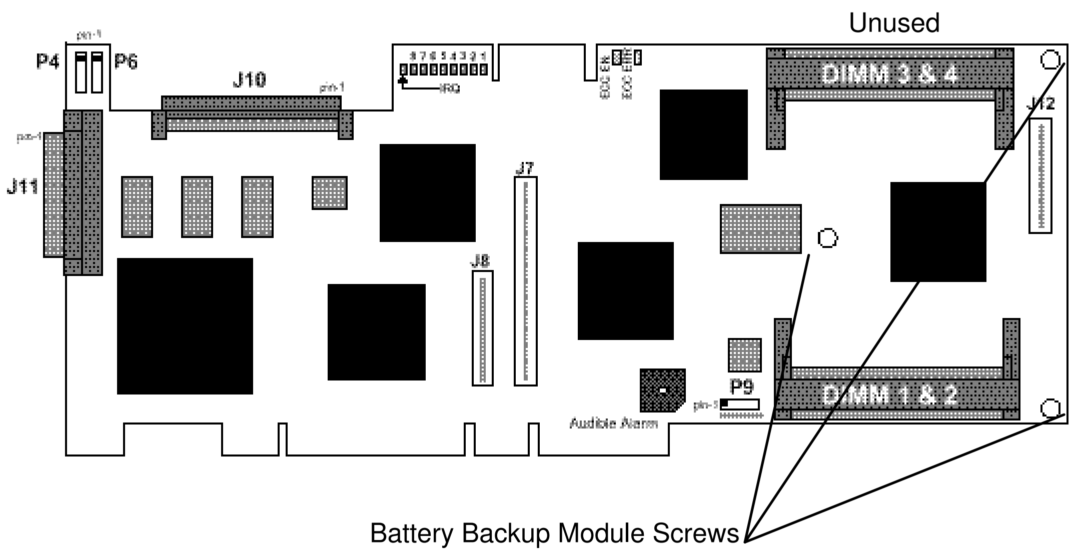

P4 Pins 1 & 2 NVRAM Clear Jumper
Pins 3 & 4 (reserved, do not use)
P6 Disk Activity LED Connector
P9 Pins 1 & 2 LOAD
Pins 3 & 4 RUN

J7 32–bit Bus Expansion Module Connector
J8 64–bit Bus Expansion Module Connector
J10 Internal Wide Ultra2 SCSI Connector (Bus 0)
J11 External Wide Ultra2 SCSI Connector (Bus 0)
J12 Battery Backup Module Connector

2. Install the 16MB DIMMs in the DIMM sockets.

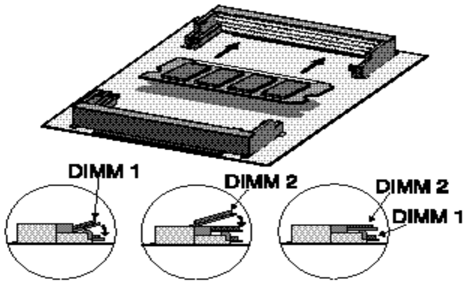

3. Reassemble the adapter and plug it into the system.
4. To confirm that the DIMMs have been properly installed, the amount of cache memory present will be displayed by Storage Manager when the Host Bus Adapter (HBA) (SCSI adapter) icon is double–clicked.

Note: LED functions are described in SCSI RAID Adapter LEDs, on page 2-11.

SCSI Termination

The devices on each physical end of a SCSI cable must be terminated. This may include the SCSI adapter and/or the SCSI devices depending upon the cabling method chosen. SCSI termination for SCSI RAID adapters is configured through the Configuration Utility or from the Configure Host Bus Adapter window accessible from the adapter's Information Window in Storage Manager. Adapter termination may be set to On or Off, causing termination for all SCSI signals to be enabled or disabled respectively. (The default setting is "Auto" that allows the adapter to automatically set/unset termination by sensing physical lines). An additional setting of High Only enables termination for signals which are present on Wide SCSI devices but not on 8–bit SCSI devices. This allows Wide cables and 8-bit cables to be simultaneously attached to the adapter.

If only one SCSI cable is attached to the adapter then the SCSI device on the opposite end of the cable must be terminated and adapter termination must be set to On as shown below:

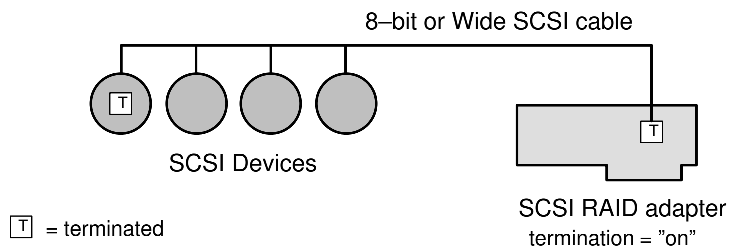

Figure 3. SCSI Single Cable End Termination

This is the only configuration (with wide SCSI cable) allowed. The other following configuration are for information only.

In some cases, it may be desirable to attach the adapter in the middle of an internal SCSI cable as shown below:

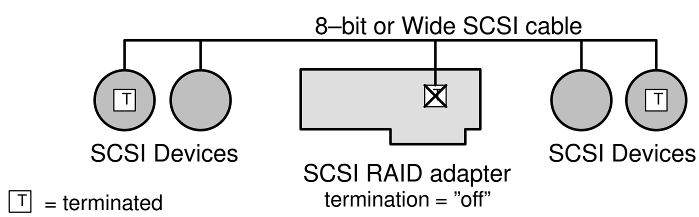

Figure 4. SCSI Single Cable Center Termination

No other cables may be attached to the adapter. The SCSI devices on each end of the cable must be terminated and adapter termination must be set to Off.

If two cables are used and SCSI devices on both cables are either all 8-bit or all Wide, then the devices on the end of each cable must be terminated and adapter termination must be set to Off as shown below:

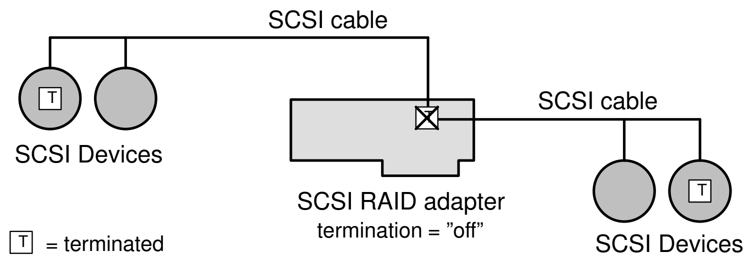

Figure 5. SCSI Two-Cable Termination

If one cable contains only 8-bit devices (either on an 8-bit cable or by using SCSI connector convertors) and the other cable (either internal or external) is Wide, then the SCSI devices on the end of each cable must be terminated and the adapter termination must be set to High Only as shown below:

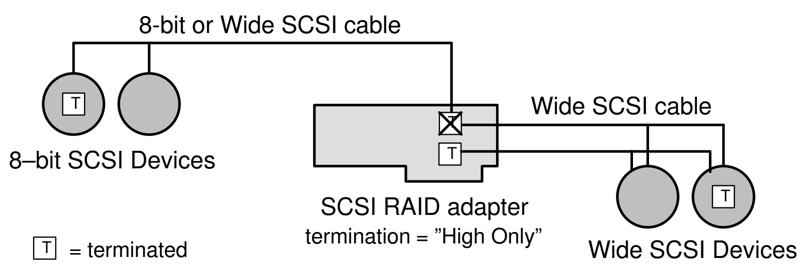

Figure 6. SCSI 8-bit/Wide Cables Termination

Configuring SCSI Device IDs

Hard Disk Drive SCSI-ID on a Deskside

Each SCSI device must be identified on the SCSI bus by a logical address, called SCSI-ID.

Hard disk drives are identified by a SCSI-ID which is automatically defined as the hard disk drive is installed in a given position. No manual intervention is required.

The following table details the correspondence between physical position and SCSI-ID for hard disk drives installed on your system.
It details also the suggested installation sequence.

The suggested sequence fills in:

- First Areas 5 and 4, which can be managed by the Bus 0 of the SCSI controller, and
- Then Areas 3 and 2, which must be managed by the bus 1 of the SCSI controller.

The installation sequence is not directly linked to the SCSI-ID. Even if the installation sequence is not the one reported here, the correspondence between physical position and SCSI-ID remains unchanged.

HARD DISK DRIVE CONFIGURATION TABLE

Device Area	Physical Position	SCSI-ID	Bus
1	-	-	
	-	-	
2	D10	8	1
	D11*	9	1
	D12*	10	1
3	D7	0	1
	D8	1	1
	D9	2	1
4	D4	8	0
	D5	9	0
	D6	10	0
5	D1 (system disk)	0	0
	D2	1	0
	D3	2	0

Legend:

D*n* = D hard disk drive
n number giving the installation sequence of devices inside the system unit.

* Revert the installation sequence of D11 and D12, in case of 18Gbyte disk.

Hard Disk Drive SCSI-ID on a Rack

Each SCSI device must be identified on the SCSI bus by a logical address, called SCSI-ID.

Hard disk drives are identified by a SCSI-ID which is automatically defined as the hard disk drive is installed in a given position. No manual intervention is required.

Note: If you purchased the hard disk drive from a supplier other than your system supplier, be sure that the device itself has no SCSI-ID set but has the motor delay jumper set. Refer to the documentation delivered with the device.

The following tables detail the correspondence between physical position and SCSI-ID for hard disk drives installed in the drawer(s).

Notice that there is not a direct link between the installation sequence and the SCSI-ID. Even if the installation sequence is not the one reported here, the correspondence between physical position and SCSI-ID remains unchanged.

- **HARD DISK DRIVE CONFIGURATION TABLE (Drawer Rear View)**

Device Area	**3**			**2**			**1**	
Physical Position	D3	D2	D1	D6	D5	D4	-	-
SCSI-ID	2	1	0	10	9	8	-	-
Bus-ID	1	1	1	0	0	0	-	-

Legend:
D*n* = D hard disk drive
n number giving the installation sequence of devices inside the drawer.

For details on SCSI IDs, saa SCSI Device IDs, on page 6-19

Devices under RAID Configuration

- Connection of disk devices on two RAID busses.

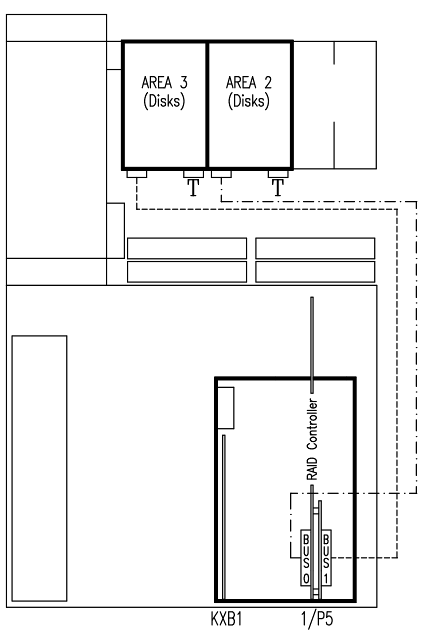

Legend:
T Terminating plug

Installing the Adapter

1. Install the PCI SCSI RAID Adapter in an empty PCI slot of the system. The slot must be a 64–bit PCI slot when a B4–C adapter is installed.
2. Connect the SCSI cables to the adapter.

SCSI RAID Adapter LEDs

SCSI RAID Adapter LEDs on the B4–4 Adapter

Ten activity LEDs on SCSI RAID adapters provide a means to visually monitor adapter activity. Several different adapter states are indicated by the LED as outlined in the following sections. Three additional LEDs indicate the status of the cache memory on the adapter. See Figure 1, on page 2-2.

LED Display During Adapter Idle

When no commands are in progress and all SCSI activity has ceased, the adapter enters the Idle state. This is indicated by a rotating bit pattern in LEDs 1 through 8.

LED Display During Adapter Active

By viewing these LED's, the user can determine information concerning cache hits and misses, disk reads and writes, and computer bus reads and writes, see Figure 7. In addition several LED's are provided for troubleshooting purposes. See **Troubleshooting**, on page A-1, for more information on interpreting these LED patterns.

LED	Function
1	Busy
2	Computer Bus Transfer to Adapter
3	Computer Bus Transfer from Adapter
4	Cache Hit
5	Disk Read–Ahead Active
6	Disk Read
7	Disk Write
8	Adapter Reset
9	Interrupt Pending to Computer
10	DRQ Asserted to Computer

Figure 7. LED Display States

LED Display During Power–up

During the power–up sequence, SCSI RAID adapter's LEDs will display the amount of expansion cache present. The expansion cache size in megabytes is displayed as a binary pattern, with lit LEDs representing binary one. LED 1, the leftmost LED, is the least significant bit. For example, if 4MB of cache memory is installed, LED 3 will light. If 12MB is installed, LEDs 3 and 4 will light.

SCSI RAID adapter then enter a fifteen–second waiting period in order to give all SCSI peripherals time to power–up. During this period, LEDs 1, 2, 5, 6 and LEDs 3, 4, 7, 8 alternately flash.

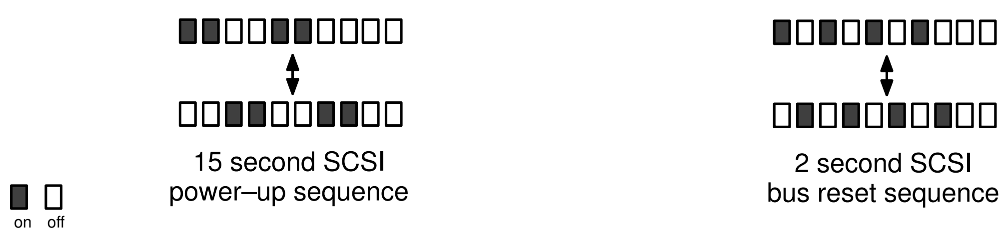

Figure 8. Power–up Sequence LED Display

RAID then resets the SCSI bus(ses) and waits for 2 seconds while LEDs 1, 3, 5, 7 and LEDs 2, 4, 6, 8 alternately flash. Starting with SCSI ID 0, the adapter then attempts to establish contact with each SCSI device. As devices 0–7 are scanned, LEDs 1–8 light in sequence from left to right. As devices 8–15 are scanned, LEDs 8–1 light in sequence from right to left. This sequence is repeated for each SCSI bus on the adapter.

During the scan, any device which reports that it is still spinning up will cause the adapter to wait for up to 30 seconds while LEDs 1, 3, 5, 7 and LEDs 2, 4, 6, 8 alternately flash. The adapter will issue a spin–up command to any device which reports that it is not spinning up.

Once contact has been established with the hard drives, the adapter will enter the idle state and display a slower, rotating bit pattern which repeats approximately once per second.

Cache Status LEDs

Three LEDs on the SCSI RAID adapter (refer to Appendix D) indicate the status of the on–board cache RAM.

- The green **ECC Enabled** LED, when lit, indicates that all installed SIMMs are ECC SIMMs, and thus the adapter cache and internal data paths are ECC protected. In that case, hard drives formatted for 528–byte sectors will support SCSI bus ECC protection.
- The red **Parity Error** LED, when lit, indicates that a parity error has been detected in one of the parity SIMMs. This LED remains lit until the adapter is powered–down. Cache failure information is recorded in the adapter error log and can be viewed through Storage Manager.
- The red **ECC Error** LED, when lit, indicates that a correctable or non–correctable error has been detected in one of the ECC SIMMs. The LED will stay lit after the error has been corrected until the adapter is powered–down. Cache failure information is recorded in the adapter error log and can be viewed through Storage Manager.

SCSI RAID Adapter LEDs on the B4–C Adapter

LED Indicators

The B4–C Adapter has two LED indicators labeled CHARGE and TRICKLE. These LEDS indicate the following conditions:

Status	TRICKLE LED	CHARGE LED
Initial Powerup (1)	Momentary flash	On
Fast Charge	Off	On
Trickle Charging	On	Off
Charge inhibit	Off	On
Discharge	Flashing	On
Battery not installed	On	On

(1) If battery voltage and the ambient temperature are acceptable, the TRICKLE LED will not remain lit.

Use the Battery Configuration dialog in Storage Manager to view the current status of the module.

Audible Alarm on the B4–C Adapter

The failure of a drive which is a member of an array attached to the adapter causes the audible alarm to sound. The alarm stops automatically (after the initial systen scan) when you start Storage Manager.

Upgrading Flash Firmware

Upgrading Flash Firmware on the B4–4 Adapter

A flash firmware upgrade is made automatically at boot time. The flash firmware file is: /etc/microcode/firm3334.fwi

If this upgrade is unsuccessful, causing the adapter to hang, the new firmware can be temporarily disabled and the upgrade attempted again using the following procedure.

1. Power-off the system.
2. Install a shorting jumper across pins 5 and 6 of P1 on the adapter, as shown below:

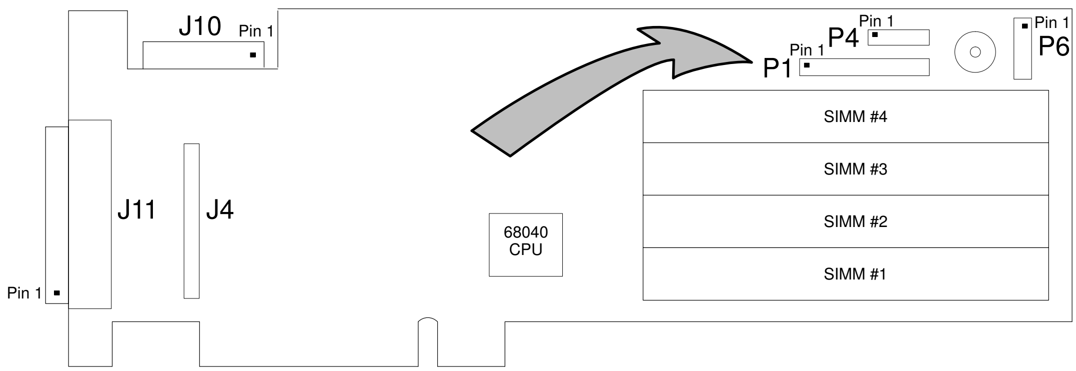

Figure 9. Installing Jumper on P1

3. Power-up the system. Boot AIX and attempt the flash update again.
4. Power-off the system and remove the jumper.

Upgrading Flash Firmware on the B4–C Adapter

A flash firmware upgrade is made automatically at boot time. The flash firmware file is: /etc/microcode/firm375.fwi

NVRAM (on the B4–C Adapter only)

SCSI RAID adapters retain their setup parameters even when powered off. These parameters are stored on the adapter in a small area of electrically erasable memory. There is a slight possibility that through improper configuration, the adapter can be put into a state in which it hangs the system during boot. If this happens, the parameters in the Non–Volatile Random Access Memory (NVRAM) can be restored to their default settings by following the steps below:

1. Power off the system.
2. Place a shorting jumper across pins 1 and 2 of P4 on the adapter (See page 2-3).
3. Power up the system and wait until the LEDs on the adapter begin flashing.
4. Power off the system and remove the jumper.

The adapter may then be re-configured using the Configuration Utility or Storage Manager.

Chapter 3. SCSI RAID Adapter – O.S. Installation

Installing the AIX Operating System.

Installing the AIX Operating System

SCSI RAID adapter associated software is located in the **AIX CD–ROM**. The package contains the AIX driver and associated Storage Manager.

Packaging Installation

To install SCSI RAID adapter software on AIX, first configure the new adapter according to the instructions given in **Configuring SCSI Termination**, on page 2-5. AIX will automatically detect the new adapter, but the driver is not installed yet.

Perform the following steps:

1. Insert the **AIX CD–ROM** in the drive,
2. Using the ”root” login type: `smit cfgmgr`,
3. Select the */dev/cd0* and type **Enter**.

Then all the software is automatically installed and configured by default.

The packages can be installed as a device driver located in CD–ROMs by following the normal AIX operation and selecting the devices named:

devices.bullraid (common)

devices.pci.441000a4 (for B4–4)

devices.pci.441000a5 (for B4–C)

AIX Installation using the CD-ROM

After booting from the AIX CD-ROM the following menu appears:

```
                 Welcome to Base Operating System
                   Installation and Maintenance

Type the number of your choice and press Enter. Choice is
indicated by >>>.


    1 Start Install Now with Default Settings
    2 Change/Show Installation Settings and Install
    3 Start Maintenance Mode for System Recovery
    >>> 4 Bull PCI Raid Configuration

    88 Help ?
    99 Previous Menu

>>> Choice [4]:
```

When AIX is to be installed on a RAID drive, the array can be created using Storage Manager (4th Menu item) before any installation operation:

- Using ESC-c command Create a RAID drive, then with Arrows and Tab keys select which RAID type.

- Using Tab key to change and Space key to select the drives on which the RAID is based.
- ESC-i and ESC-d to Include the physical drives and complete RAID definition.
- ESC-o to confirm the creation.
- ESC-f to access the main menu, then the ’s’ key to ”Set System Configuration”, to start physical RAID building.
- if the RAID is type 1 the Storage Manager proposes two modes: Copy or Clear on the physical drives.

Note the SCSI ID of the created RAID array for further identification of the disk on which to install AIX.

- ESC-f to access the main menu, then the ’x’ key to exit the Storage Manager.

for a B4–4 adapter, before installing AIX on the new RAID drive, wait for completion of the RAID array build (Amber LED in front of the disk goes OFF).

for a B4–C adapter, AIX may be installed on the new RAID drive, while the RAID array is building (Amber LED in front of the disk is Flashing).

It is possible to watch the RAID build progress re-using the Storage Manager.

Then select the drives on which AIX is to be installed (1st Menu item), and follow the standard installation operation.

To be sure of correct operation, note the SCSI– ID of the created RAID drives while in Storage Manager and then find the associated drive in the proposed list.

Migrate from a Stand-alone System Disk Drive to RAID 1 System Drive

Note: You cannot expand a single system disk to a RAID 1 system disk if you are booted on the system disk itself.

Boot AIX using AIX CD-ROM (previous paragraph)

Enter in the BULL PCI Raid Configuration.

Create the RAID 1 with both the current system drive and the other drive selected.

After ”Set System Configuration” a dialog appears asking for the RAID 1 mode.

- Select Copy mode.

A new dialog appears.

- Choose the copy direction. (direction in which to copy the System Drive to the other drive)

Wait for the end if RAID 1 building.

Then Reboot; the system is now RAID 1 drive array.

Storage Manager on Graphical ASCII Emulation Mode

Start ASCII Emulation with the command:

```
/usr/lpp/dpt/dptmgrac –simple
```

This emulation mode does not support the underline attribute to show a selected item. There is no indication on the screen to identify a selected item; all selections, made using the ”space-bar”, must be remembered or noted manually.

Commands and actions are performed using ESC- key sequences, where the key is usually the first letter of the command or action and identified as underlined or reverse video.

’Caps Lock’ must be unlocked to have ESC- key sequences available.

Adding Hard Drives

To add SCSI hard drives to an AIX system, follow the steps below:

1. Insert the new SCSI device into the guide slots, tilt the latches until they are anchored to the disk cage and lock carrier by closing the latches.
2. If the new drives are to be members of a new or existing Array Groups, then run Storage Manager to create or modify these groups. Each array Group will appear to the operating system as a single hard drive. Select each new hard drive and Array Group LSU icon and note the HBA and ID number.

The HBA represents the occurrence of the SCSI RAID adapter on the machine and identifies the SCSI ID on the SCSI bus.

Hard Drives Recognition

To access or create an AIX Logical Volume on a physical or logical defined disk drive trough SCSI RAID adapter an explanation is necessary. AIX knows the disk drives as **hdisk** objects and their location code (given with lsdev –Ccdisk command) and SCSI RAID adapter as LSU (given by the Storage Manager).

Location code on AIX is like aa–bb–cc–dd,0 where:

aa identifies PCI bus where adapter is located

bb identifies PCI slot where adapter is located

cc identifies the SCSI bus

dd identifies drive SCSI ID

LSU on the adapter is like A,B,C,0 where:

A identifies HBA

B identifies SCSI bus on the adapter

C identifies drives SCSI ID

HBAs are numbered according to their AIX logical name number. B4–4 first (ex.: sra0, sra1, srai0, srai1,...)

Example: 2 adapters: 1 B4–4 and 1 B4–C

sra0 –> HBA0

srai0 –> HBA1

Further association between AIX Location Code and LSU numbering is:

CC <–> B

DD <–> C

The application "`/usr/lpp/diagnostics/pci/dpt_who`" can also be used to identify AIX disks and LSUs.

Chapter 4. SCSI RAID Adapter – Storage Devices

Describes the RAID subsystems, SCSI devices and accessories and gives instructions on configuring complete RAID subsystems using these products.

RAID Subsystem Configuration

RAID subsystems may be configured a variety of ways, depending upon the storage capacity, performance, degree of fault–tolerance and types of SCSI devices desired. Determining the proper configuration for your application consists of the following steps:

1. Determine RAID level(s) to be utilized. This should be based on the performance and capacity requirements of the system, and whether or not fault–tolerance is desired.
2. Determine the number and capacity of hard drives including Hot–Spares. Depending upon the RAID level, different numbers of drives are required to provide the same usable storage capacity.
3. Select the amount of hardware cache memory to be installed. At least 4MB for a B4–4 adapter and 32MB for a B4–C adapter is required and delivered for each adapter. B4–4 adapter can be configured with up to 64MB of cache. Adding more cache typically increases performance.

A detailed description of each step above is given in the next section of this chapter.

Determining the RAID Level

Use the following descriptions as guidelines for determining which RAID level is best for your system. Multiple arrays of different levels may be configured on a single adapter. For details, see **RAID**, on page 6-7. Any number of arrays of the same RAID level may be combined into a single array group which appears to your computer as one hard drive. Array groups can be spanned over the 2 SCSI RAID adapter channels. The number of drives which may be included in an array group depends upon the RAID level as summarized below:

RAID–0	**Data striping across multiple drives**
Benefits:	Maximized storage capacity
	Best performance
	Requires as few as two drives
Drawbacks:	No fault tolerance
Description:	A RAID–0 array is simply a group of hard drives which have been combined into a single logical storage unit using striping. This logical storage unit looks like a single large disk drive to the computer, but the storage space is interleaved so that disk accesses will be evenly distributed across every drive in the group. A RAID–0 group must contain at least two drives, but there is no limit to the maximum number of drives which may be included in the group. The drives may be connected to any combination of adapter channels.
Capacity:	Since no drive capacity is used for data redundancy, the capacity of a RAID–0 array will be equal to the combined capacity of all drives in the array.
Performance:	Adding more drives to a RAID–0 array increases performance in multiuser systems since more user data requests can be serviced simultaneously.
Drive faults:	RAID–0 arrays are not fault–tolerant. If any drive in the array fails, the entire logical storage unit will fail.

RAID–1 **Mirrored data is written to pairs of drives**

Benefits: Fastest fault–tolerant architecture

Least performance degradation with failed drive

Fastest rebuild

Requires as few as two drives

Drawbacks: Least efficient storage capacity

Description: A RAID–1 array is a mirrored pair of drives which contain identical data. The RAID–1 pair looks like a single drive to the computer and has the same capacity as a single drive but can perform two read operations at the same time. Although both drives in a RAID–1 array must be connected to the same adapter, an unlimited number of RAID–1 arrays connected to any combination of adapters may be combined through striping into a single RAID–1 ”array group”.

Capacity: Since one half of all drives in the array group contain redundant data, the usable storage capacity of the group is half the combined capacity of all drives.

Performance: Both read and write operations are fast. Write operations are considerably faster than RAID–5 arrays since much less overhead is required. Adding more drive pairs to a RAID–1 array group increases performance in multiuser systems since more user data requests can be serviced simultaneously.

Drive Faults: Unless both drives of the same mirrored pair fail, the RAID–1 group will not fail and no data will be lost. When a drive fails in a RAID–1 array group, the array is said to be functioning in degraded mode. Degraded mode in a RAID–1 array group is only slightly slower than normal operation, since all read operations which would normally be performed by the failed drive are re–routed to its mirrored drive. Write operations are not impacted. Array rebuilds occur quickly with minimal performance impact since the contents of the replaced drive are simply copied from its mirrored drive.

RAID–5 **Data striping with parity**

Benefits: Fault tolerant

Most efficient storage capacity

Drawbacks: Worst performance degradation with failed drive

Slowest rebuild

Requires at least three drives

Description: Each RAID–5 array consists of three or more drives which must be attached to the same adapter. However, an unlimited number of RAID–5 arrays connected to any combination of adapters may be combined through striping into a single array group.

Capacity: Since the equivalent capacity of one drive in each RAID–5 array contains redundant data, the usable storage capacity of the combined array group equals the combined capacity of all drives minus one drive for each array.

Performance: Read operations are fast but write operations require much more overhead than in RAID–1 arrays and thus are slower. Adding more drives to a RAID–5 array group increases performance in multiuser systems since more user data requests can simultaneously be serviced by the array.

Drive faults: Unless two drives within the same RAID–5 array fail, the RAID–5 group will not fail and no data will be lost. When a drive fails in a RAID–5 array group, the array is said to be functioning in degraded mode. Degraded mode and array rebuilds in a RAID–5 array group are considerably slower than RAID–1 arrays since missing data must be synthesized by XORing data read from the remaining drives in the array.

Selecting Hard Drives

Determine the number of hard drives you need to yield the total usable capacity required by your system. Refer to the preceding section to calculate the capacity depending upon the RAID level selected. SCSI RAID adapter hard drives are available in 2GB, 4GB and 9GB capacities.

Since adding more drives to an array group typically increases performance, you may wish to configure your system using a larger number of smaller capacity drives. Although it is supported, it is not recommended that you mix drive capacities within an array group.

Optionally, RAID–1 and RAID–5 arrays can be protected by a Hot Spare drive, providing immediate, automatic array rebuild when a drive fails. Any hard drive model above may be configured as a Hot Spare as long as it is the same or larger capacity as the drives it protects. A single Hot Spare will protect all RAID–1 and RAID–5 arrays attached to the same adapter as the Hot Spare. Therefore, if your arrays are to span multiple adapters, one Hot Spare will be required per adapter.

Drives LED Management

Each hard disk drive, connected to the SCSI RAID adapter, has two LEDs that display the drive status. The green/blue LED is the drive activity LED and is on whenever the drive is powered on and spinning. The amber LED is the drive fault LED and indicates a failed drive.

Green or Blue (LVD mode) LED ON = drive motor ON.

Amber LED ON (solid, not flashing) = drive fault, indicates a failed drive.

The amber LED flashes to indicate several different types of operations:

- Amber LEDs on all drives flash during power–up.
- Amber LEDs on all drives, in an array, flash during the initial array build operation.
- When a failed drive is replaced, the new drive's amber LED flashes during the array rebuild operation.
- A drive's amber LED flashes during a low–level format operation.

Replacing a Drive

CAUTION:
Do not remove a drive from its location if either the green/blue LED is lit or the amber LED is flashing.

SCSI RAID Adapter drives may be replaced without powering down the cabinet or interrupting normal operation. Any failed drive may be removed by simply pulling out the drive carrier. When a replacement drive is inserted, SCSI RAID adapter automatically spins–up the drive and starts the rebuild process.

It is not necessary to set switches or run special software.

For immediate automatic replacement of failed drives, extra drives can be designated as Hot Spares. Upon failure of any drive in a RAID–1 or RAID–5 array, the Hot Spare will automatically replace the failed drive and the missing data will be rebuilt onto the Hot Spare as a background activity without user intervention and without interruption of normal system operation.

Before replacing a SCSI device, assure that the device is not active (as indicated by the green or Blue LED OFF) and the amber fault LED is not flashing.

Use the following procedure to install, remove or replace a drive from the cabinet:

1. Open the front door of the cabinet.
2. To remove an existing SCSI device, open the two latches on each side of the device together and slide the disk carrier out of the cage.
3. Insert the replacement SCSI device into the guide slots and tilt the latches until they are anchored to the disk cage and lock carrier by closing the latches.
4. When powered–up, the green device LED will turn ON, indicating that the drive is ready. The amber fault LED will either be off or will flash if an array rebuild operation is building data onto that drive.

Chapter 5. SCSI RAID Adapter – Storage Manager

Installing and using the Storage Manager.

Introduction

The Storage Manager utility performs several functions:

- Allows user to create, modify and delete disk arrays.
- Provides on–line functions for the SCSI storage subsystem such as event logging and notification, diagnostics and I/O monitoring.

Installing Storage Manager

When setting up a system, once the SCSI hardware has been configured and installed, the next step is to run Storage Manager's install process. This will allow you to verify that your SCSI hardware has been configured correctly and to create disk arrays.

1. At this point, you can optionally inspect your SCSI hardware configuration as displayed by Storage Manager:
   a. Verify that all SCSI devices and adapters are shown. If any devices are missing from the display then exit Storage Manager and check your hardware.
   b. Check all devices marked with a yellow warning flag. This indicates that Storage Manager has detected a potential problem with that device such as a SCSI ID conflict. Click on the device icon to view information about the problem.
   c. View the information window for each adapter to verify that all plug–on modules and SIMMs are shown.
2. At this point you should create any desired disk arrays as outlined later in this chapter. Array Groups may be created or modified at any time after system installation.
3. Exit Storage Manager when finished. You will be prompted to save the configuration changes. If you elect to save the configuration then arrays which were created or modified will be built at this time. For large arrays, this process may take several hours. You may exit Storage Manager and perform other activity on the system while the build operation continues. However, the Array Groups being built may not be accessed until the build has completed. If you wish to monitor the progress of the build for any Array Group, Storage Manager may be re–entered and the information window for the array may be viewed.

Storage Manager GUI

Storage Manager's Graphical User Interface (GUI) works optimally with a Video Graphics Adapter (VGA) or better monitor. In text mode, Storage Manager's functionality is unchanged but instead of graphical icons, SCSI devices are represented by shaded squares with identifying text.

Storage Manager Keyboard Functions

Before running Storage Manager, be sure that the mouse driver has been installed according to the manufacturer's directions. Storage Manager may also be controlled via the arrow and tab keys. The table below lists some Storage Manager mouse functions and their equivalent keyboard key sequence:

Storage Manager function	Key sequence	UNIX key sequence
Select pull–down menu or button	Alt+Underlined letter	Esc, then Underlined letter
Exit pull–down menu	Esc	Esc, then Esc
Move to next icon	Tab	Tab
Move to previous icon	Shift+Tab	n/a
View icon information window	Enter	Enter
Scroll up or down	Up or Down arrows	Up or Down arrows
Scroll up or down one page	Page Up or Down	n/a
Move to top or bottom	Home or End	n/a
Select drive icon	Space	Space

Running Storage Manager

After initial installation, AIX Storage Manager is started by typing:

```
/usr/lpp/dpt/dptmgram
```

Storage Manager will initially scan for hardware installed on the computer on which it is run. If SCSI RAID adapters are found, the SCSI hardware configuration will be displayed.

Physical Configuration View

The first screen to be displayed by Storage Manager is always the ”Physical Configuration View”. This screen displays each host adapter in the system along with the attached SCSI busses and devices. Icons representing hard drives, are displayed in order of Host Bus adapter (HBA) number and SCSI ID. (A complete list of possible icons may be viewed by selecting ”Legend of Icons” from the ”Help” menu.)

Hard drives which are part of an Array Group are represented with the word ”RAID” overlaid on the drive icon. Hard drives which have been assigned as Hot Spares are represented with a white cross on the drive icon.

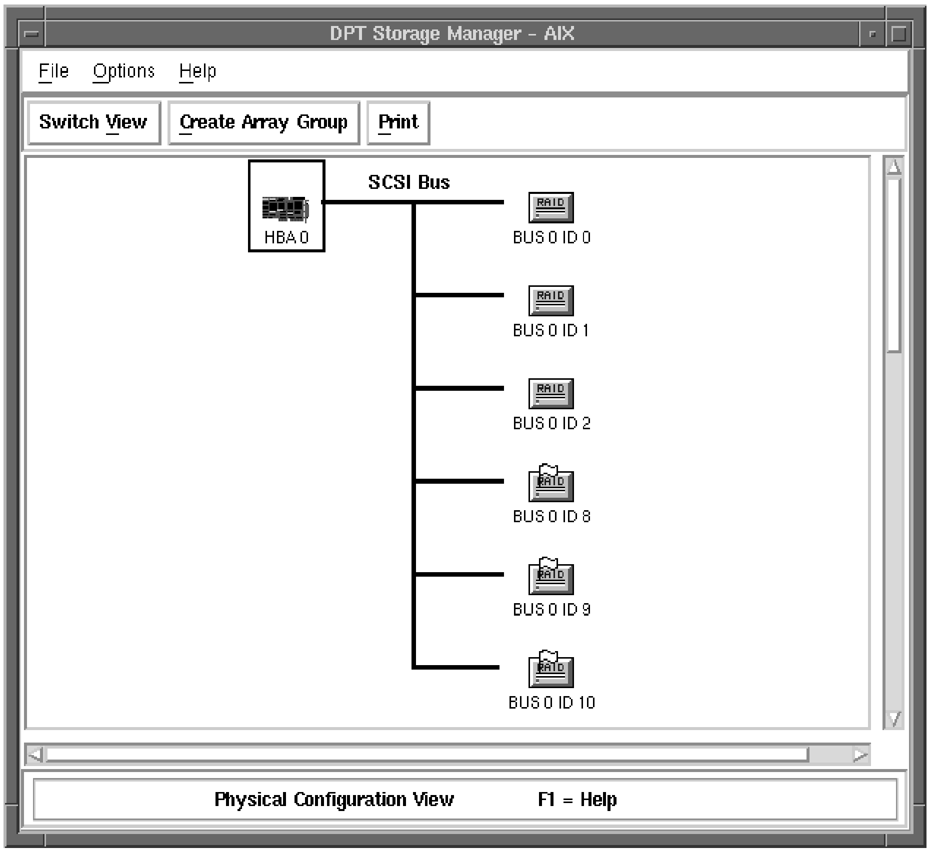

Figure 10. Physical Configuration View Screen

Logical Configuration View

The ”Switch View” button toggles between the Physical Configuration View and the Logical Configuration View screens. On the right side of the Logical Configuration View screen are all physical SCSI devices which are attached to BULL SCSI RAID adapters. On the left side of the screen are the associated logical devices as seen by the computer.

Non–disk devices have the same logical and physical icons. Hard disks can be seen by the computer either as individual drives or as members of Array Groups. In either case, the drive or Array Group is represented on the left side of the screen as an LSU (Logical Storage Unit). Hardware Arrays which make up that Array Group are displayed as RAID–1 or RAID–5 icons between the LSU on the left and the drives on the right.

Devices are displayed in order of device type, with all non–disk devices displayed first, followed by all un–arrayed hard disks, hot spares, and finally, all Array Groups by RAID level.

Figure 11 shows a Logical Configuration View screen with one un–arrayed hard disk and a RAID–5 Array Group composed of four hard disks.

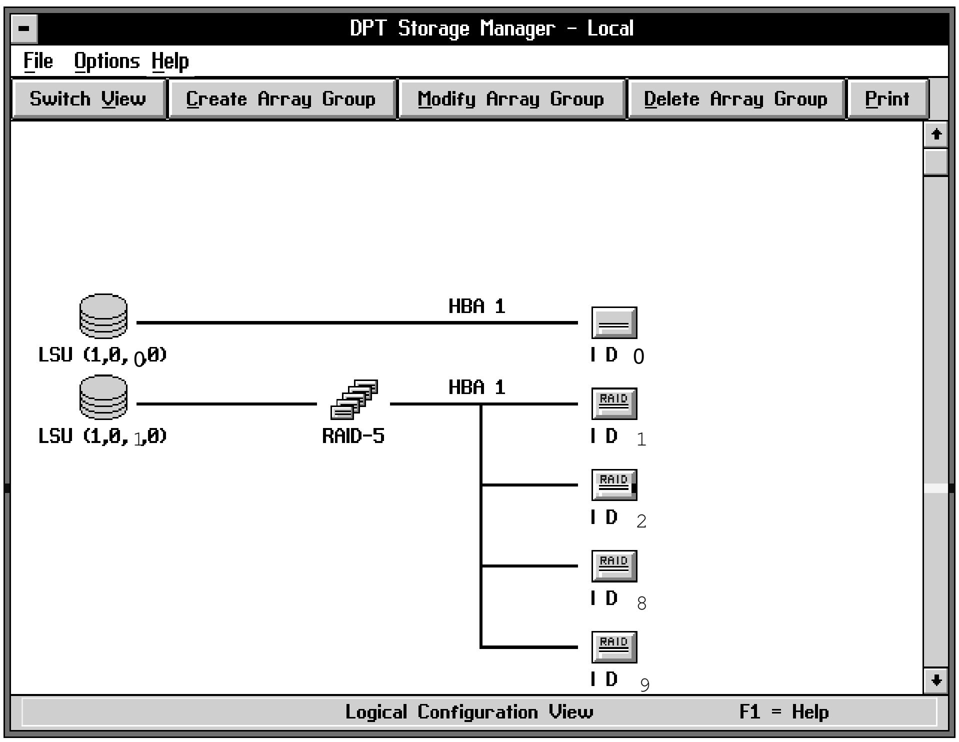

Figure 11. Logical Configuration View Screen

Information Windows

Double–clicking on an HBA (Host Bus Adapter) or SCSI device icon will cause an ”Information Window” to be displayed with information and buttons for that HBA or device. (Alternately, **Tab** and **Shift–Tab** may be used to highlight the text under the icon. Pressing **Enter** will then display the window.)

HBA Information Windows

Information windows for HBAs display configuration information reported by that SCSI adapter, Figure 12. At the top left of the window in the ”Controller” section are the HBA ”Model”(The F model suffix indicating floppy support is not detectable by Storage Manager and will not be displayed in the Model field), ”Firmware” revision and amount of installed ”Cache”. ”ECC”, if checked, indicates that only ECC SIMMs are installed and therefore the adapter internal data paths are protected by ECC. At the top right is a display showing ”Attached Modules”, including SIMMs. At the bottom of the window are the SCSI bus and computer bus types supported. Buttons available include ”Configure”, ”Event Log”, ”I/O Monitor” and ”Print”. ”Test Alarm” tests the audible alarm on the adapter.

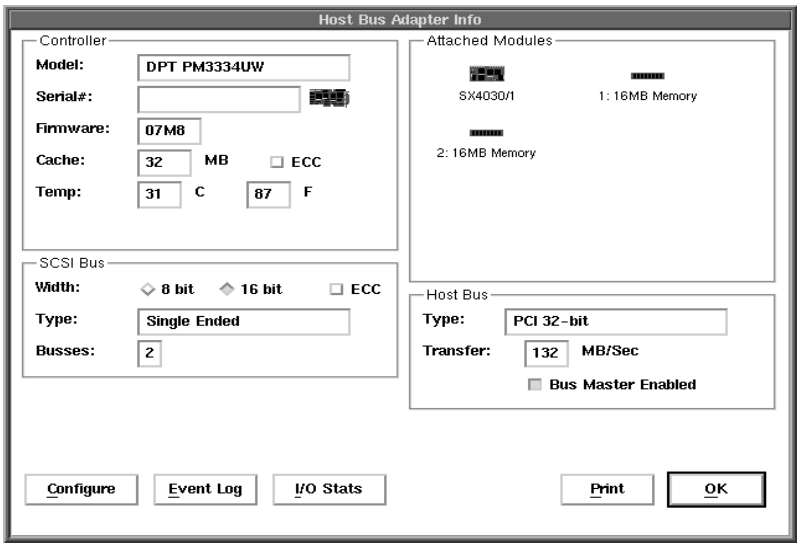

Figure 12. HBA Information Window for a B4–4 adapter

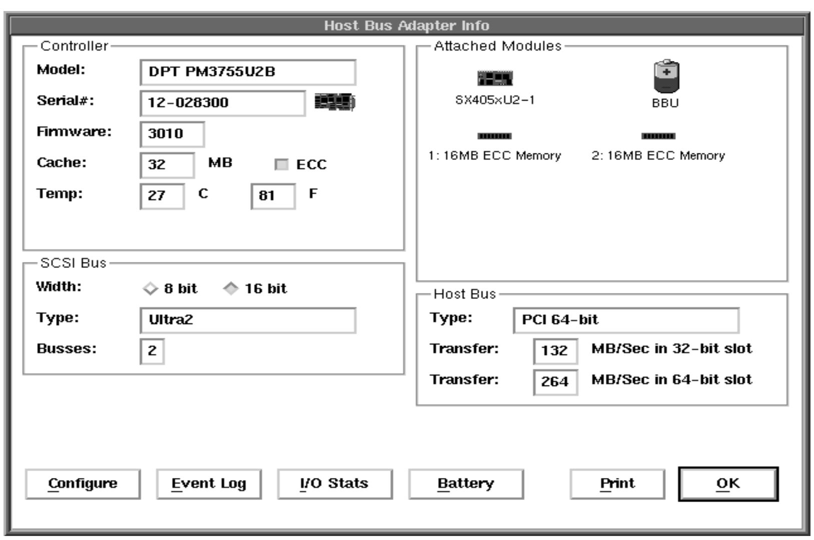

Figure 13. HBA Information Window for a B4–C adapter

HBA Configuration

To modify configurable hardware parameters for the SCSI RAID adapter click on the "Configure" button in the adapter's information window. This will cause the "Configure Host Bus Adapter" window to appear.

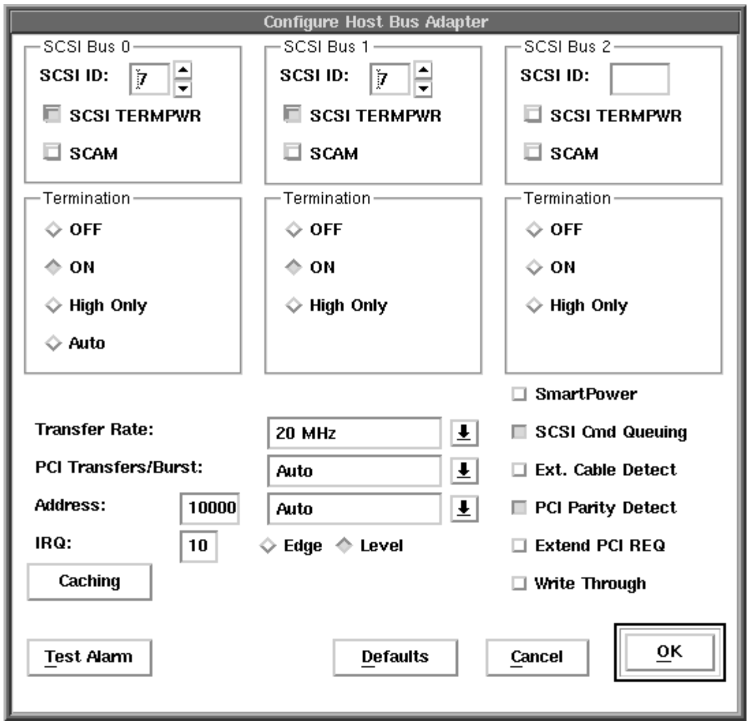

Figure 14. HBA Configuration Window for a B4–4 adapter

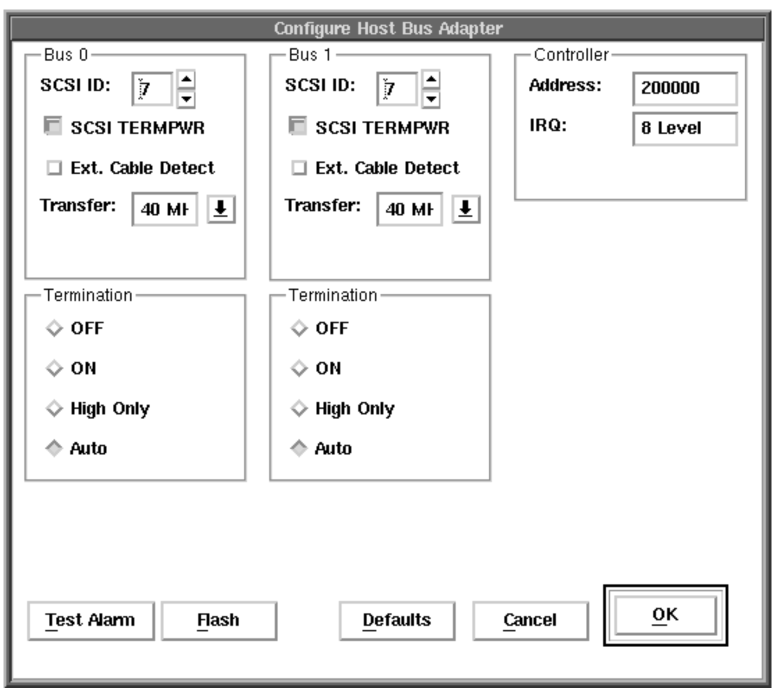

Figure 15. HBA Configuration Window for a B4–C adapter

Selecting the "Caching" button causes the "HBA Caching Configuration" window to be displayed (see next section). Selecting "Defaults" causes the parameters to return to their default settings. The window may be exited without saving changes by selecting "Cancel" or may be exited with changes by selecting "OK".

- **SCSI ID:** SCSI RAID adapters are configured by default at SCSI ID 7. This value normally need not be changed but, if necessary, for more information on selecting SCSI IDs, see **Configuring SCSI Termination**, on page 2-5.
- **SCSI TERMPWR:** SCSI RAID adapters, by default supply termination power for other SCSI devices via the TERMPWR line on the SCSI cable. If you wish to disable this parameter, see **TERMPWR**, on page 6-22.
- **SCAM:** SCSI RAID adapters do not support SCAM automatic configuration of SCSI IDs.
- **Termination:** This section controls the adapter SCSI termination. The default of ”On” should not be changed if only one SCSI cable is attached to the adapter and the adapter is the last device on that cable. Otherwise, see **Configuring SCSI Termination**, on page 2-5, for instructions on configuring adapter termination.
- **SCSI Transfer Rate:** SCSI RAID adapters automatically negotiate with each SCSI device at power–up to determine the optimum SCSI transfer rate. This section sets the maximum rate for SCSI transfers and should be left at the default except during troubleshooting of SCSI bus errors. (If setting the maximum rate to a lower rate eliminates SCSI bus data errors, this is usually an indication of improper SCSI cable length or termination.)
- **Ext. Cable Detect:** SCSI RAID adapters can detect the presence of an external SCSI cable and automatically lower the maximum SCSI bus transfer rate to 5MHz (for both internal and external cables) when an external cable is present. Since many external SCSI cabling methods are not capable of handling maximum transfer rates without causing data errors, enabling this feature can protect against data integrity problems if external SCSI devices (such as tape and CD–ROMs) with long cables are occasionally attached.
- **I/O Address:** With SCSI RAID adapters, this parameter may be used to override the I/O address automatically assigned by the plug–and–play BIOS.
- **IRQ:** This read–only parameter displays the adapter's interrupt vector setting. With SCSI RAID adapters, IRQ is automatically assigned by the plug–and–play BIOS.
- **PCI Transfers/Burst:** This appears only for PCI adapters and specifies the minimum number of 4–byte words transferred by the adapter per bus master burst. The adapter waits until it has accumulated enough data to transfer the specified number of words before requesting the bus. If PCI Transfers/Burst is left at the default of Auto the adapter will choose the best value (typically 32). Setting this to a higher value may increase throughput in some systems.
- **SmartPower:** SCSI RAID adapters do not support automatic green mode during periods of inactivity.
- **PCI Parity Detect:** This parameter appears only for PCI adapters. SCSI RAID PCI adapters check PCI bus parity for motherboards that support parity. This parameter must be manually disabled for some PCI motherboards that incorrectly report parity support, to disable the adapter's parity checking.
- **Extend PCI REQ:** This parameter appears only for SCSI RAID PCI adapters and controls the duration of the PCI Bus Request signal. The default setting of disabled forces the adapter to release the bus more frequently. Enabling this option causes the adapter to assert the bus request signal for the entire duration of the data transfer, thus keeping the bus until the transfer has completed or the adapter is forced to release the bus by another bus master or the latency timer. With some PCI motherboards, enabling this option may provide higher throughput.
- **SCSI Cmd Queuing:** SCSI RAID adapters support SCSI Tagged Command Queuing, allowing the adapter to transfer up to 64 commands to a SCSI device. For details, see **Command Queuing**, on page 6-16. Disabling this option should not typically be necessary since the adapter will interrogate each SCSI device to determine if it is capable of command queuing.
- **Write–Through:**In Write Through mode, data is written to the final destination before a write operation is reported as complete.

HBA Caching Configuration for B4–4 Adapter only

Selecting the ”Caching” button from the ”Configure Host Bus Adapter” window causes the ”HBA Caching Configuration” window to be displayed. From this window, the parameters described below may be modified. Selecting ”Defaults” causes the parameters to return to their default settings. The window may be exited without saving changes by selecting ”Cancel” or may be exited with changes by selecting ”OK”.

Note: Additional caching parameters which may be set separately for each SCSI hard drive are modified from the ”Device Caching Configuration” windows, accessible from each hard drive’s information window.

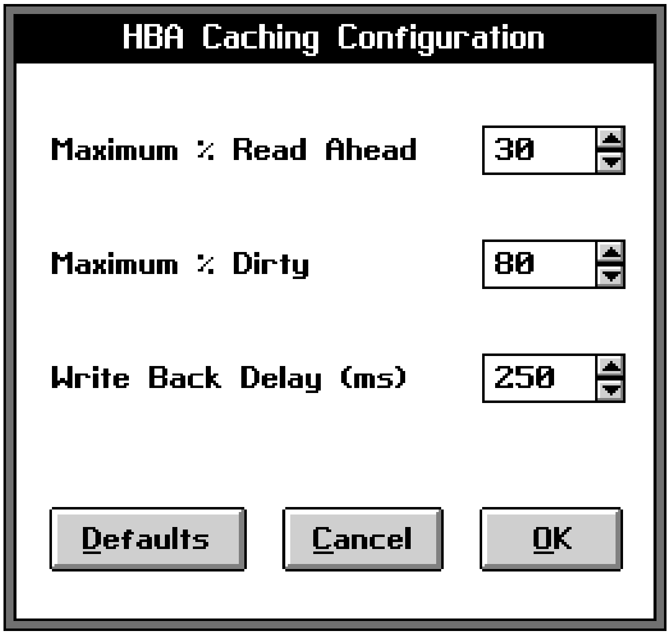

Figure 16. HBA Caching Configuration Window for B4–4 adapter

- **Maximum % Read Ahead:** If the SCSI RAID adapter determines that a read command from the computer is part of a pattern of sequential reads, additional sequential sectors will be read into cache so that future read commands will result in cache hits. The maximum percentage of adapter cache which is allowed to hold read–ahead data may be modified by this parameter.
- **Maximum % Dirty:** The maximum percentage of adapter cache which is allowed to hold dirty data may be modified by this parameter. Flush operations are automatically initiated whenever the percentage of cache which contains dirty data exceeds this value. The flush will continue and the execution of further write commands from the computer will be delayed until the percentage of cache which is dirty falls under this threshold.
- **Write Back Delay:** Whenever a SCSI RAID adapter has not received an I/O request for a hard drive for this time period, it will begin to flush dirty pages from that drive to disk. (Flush operations copy data from cache to disk but do not remove the data from cache, thereby allowing possible future read hits.)

SCSI Device Information Windows

At the top of the information window, Figure 17, for SCSI devices is a ”Description” field showing the manufacturer and model along with the ”Revision” number of the device. Following this are the ”Address” and the ”Capacity” of the device in Megabytes. For removable media disk devices, the capacity is that reported by the currently inserted media. Disk devices also display the number of sectors on the media along with the physical and logical sector size. Other information including ”Transfer” rate, and whether the device supports removable media and Error Correcting Code (ECC) protection over the SCSI bus is also reported.

Below this, the current device ”Status” is displayed. (Some status conditions are also indicated by colored flags on the device icon.) ”SCSI Capabilities” reported by the device include support for ”Soft Reset”, ”Command Queuing”, ”Linked Commands”, ”Synchronous” SCSI, 16–bit or 32–bit ”Wide” SCSI, ”Relative Addressing”, ”S.M.A.R.T.” and whether the device conforms to the SCSI–2 and SCSI–3 specs.

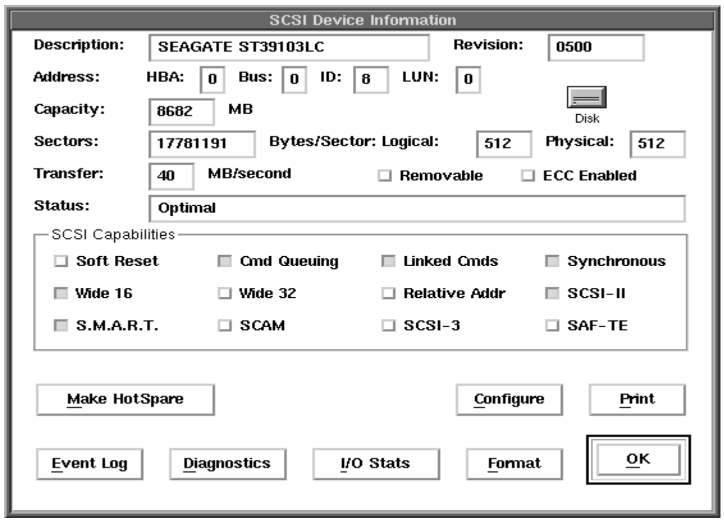

Figure 17. SCSI Device Information Window for a B4–4 adapter

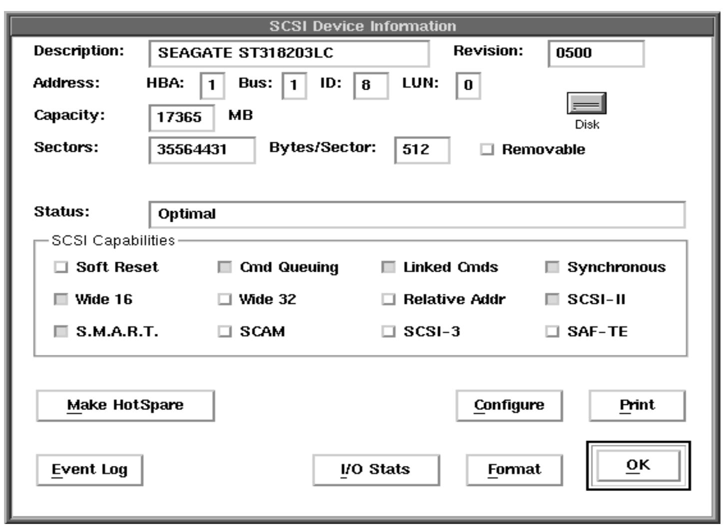

Figure 18. SCSI Device Information Window for a B4–C adapter

Various buttons are available depending upon the device type. Functions always include ”Print”. Hard disk devices also offer ”Event Log”, ”Diagnostics”, ”I/O Monitor” and ”Caching”. Hard disks which are not members of Array Groups also include ”Make HotSpare” and ”Format” buttons. For hard disks which are members of Array Groups, the name and RAID level of the array to which they belong, along with the ”Stripe Size” for that array, are reported.

Device Caching Configuration for B4–4 Adapter

Some caching parameters may be set individually for each hard drive. Selecting the ”Caching” button from a hard drive’s information window causes the ”Device Caching Configuration” window, see figure 19, to be displayed for that drive. From this window, the parameters described below may be modified. Selecting ”Defaults” causes the parameters for that drive to return to their default settings. The window may be exited without saving changes by selecting ”Cancel” or may be exited with changes by selecting ”OK”.

Note: Additional caching parameters which affect all hard drives on the adapter are modified from the ”HBA Caching Configuration” window, accessible from the adapter’s information window.

By default, all the drives are in ”Write Through” cache mode parameter. This mode is the most secured but the less efficient.

This parameter must be changed only when the drive is in a stand alone mode (not part of a RAID).

To change this Write Cache mode when part of a RAID:

- backup the RAID data on a tape,
- delete the RAID,
- change the drives Write Cache mode,
- recreate the RAID,
- restore data from the tape.

Do not mix Write Cache modes on the same RAID.

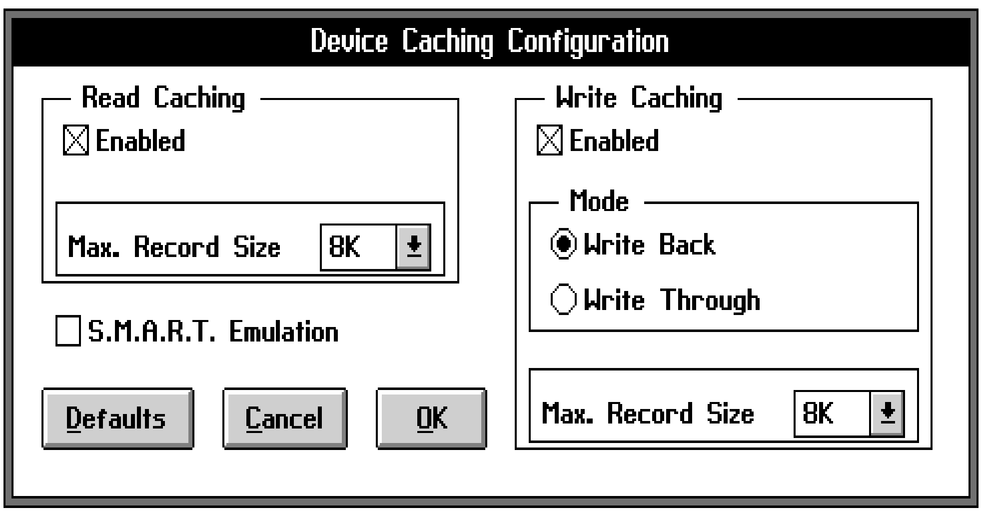

Figure 19. Device Caching Configuration Window

- **Read Caching Enabled:** If this box is not checked, no data will be cached during disk read operations.
- **Maximum Record Size:** When reading data from disk, SCSI RAID adapters normally limit the largest record size that will be cached to 8 Kbytes. This prevents long I/O operations such as large file copies from displacing large amounts of data from cache. If the computer issues a read request larger than 8KB, the data will not be cached. The maximum record size cached may be set to a different limit for each hard drive.
- **Write Caching Enabled:** If this box is not checked, no data will be cached during disk write operations.
- **Write Caching Mode:** In ”Write Through” mode, all data is written to disk for each write command before command complete status is returned to the computer (even though the

data may also be cached for later use). In ”Write Back” mode, writes to disk may be deferred until after command completion.

- **Maximum Record Size:** This parameter operates in the same manner as the Maximum Record Size parameter for read caching.
- **S.M.A.R.T. Emulation:** SCSI RAID adapters support drives which provide S.M.A.R.T. predictive failure reporting based on head flying height, read currents, etc. For drives that do not provide S.M.A.R.T. failure prediction, this parameter enables the adapter to emulate S.M.A.R.T. failure prediction based on drive spin–up and seek times.

Array Group Caching Configuration for B4–C Adapter

Selecting the “Caching” button from the “Array Group Information” window causes the “Array Group Caching Configuration” window to be displayed.

From this window, the parameters described below can be modified.

Selecting “Defaults” causes the parameters to return to their default settings.

The window may be exited without saving changes by selecting “Cancel” or may be exited with changes by seleting “OK”.

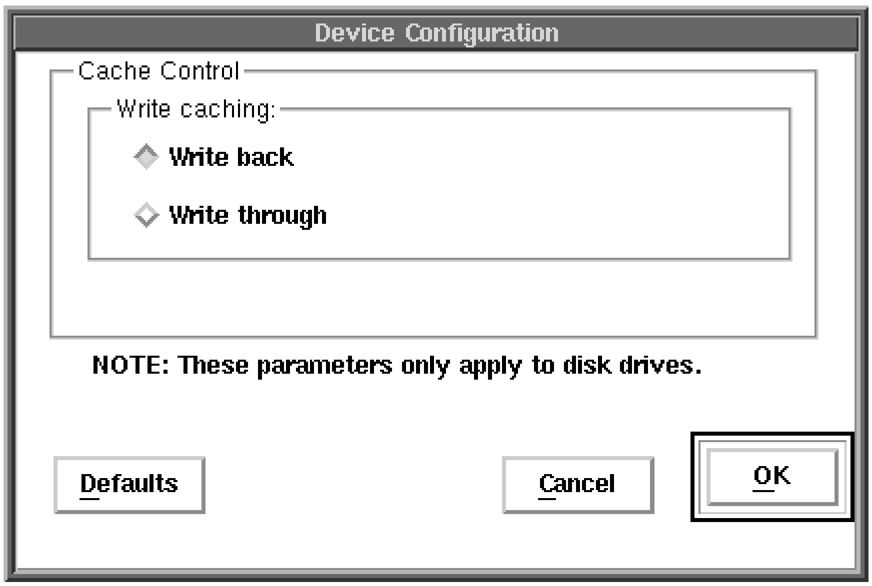

Figure 20. Array Group Caching Configuration Window

Write caching: In “Write through” mode, all data is written to disk for each wtrite command before command complete status is returned to the computer (even though the data can also be cached for later use). In “Write back” mode, writes to disk can be deferred until after command completion. “Write back” mode provides better performance.

On the opposite of the B4–4 adapter, where an external reliable power supply (full system on Battery Backup power supply) is needed before turning on the write–back mode to avoid data loss problems, the B4–C adapter has an embedded Battery Backup module (with 72 hours capacity when fully loaded), so that the write–back feature can be safety used.

Array Groups

In the tool bar at the top of the Logical Configuration View screen are the "Create Array Group", "Modify Array Group" and "Delete Array Group" buttons. These buttons allow drives to be added or removed from RAID 0, 1 or 5 Array Groups.

SCSI RAID adapters can implement any combination of RAID 0, 1 or 5 disk arrays in adapter hardware. RAID–1 Hardware Arrays are always composed of two drives. RAID–5 Hardware Arrays must contain at least three drives. All drives in a Hardware Array must be attached to the same adapter. However, any number of Hardware Arrays (one or more) of the same RAID level may be combined into an "Array Group". (RAID–0 Array Groups are composed of any combination of individual drives.) All the drives in an Array Group will then appear as a single LSU (Logical Storage Unit) to the computer.

Figure 21 shows two RAID–1 Array Group composed of RAID–1 Hardware Arrays and a RAID–5 Array Group composed of a single RAID–5 Hardware Array.

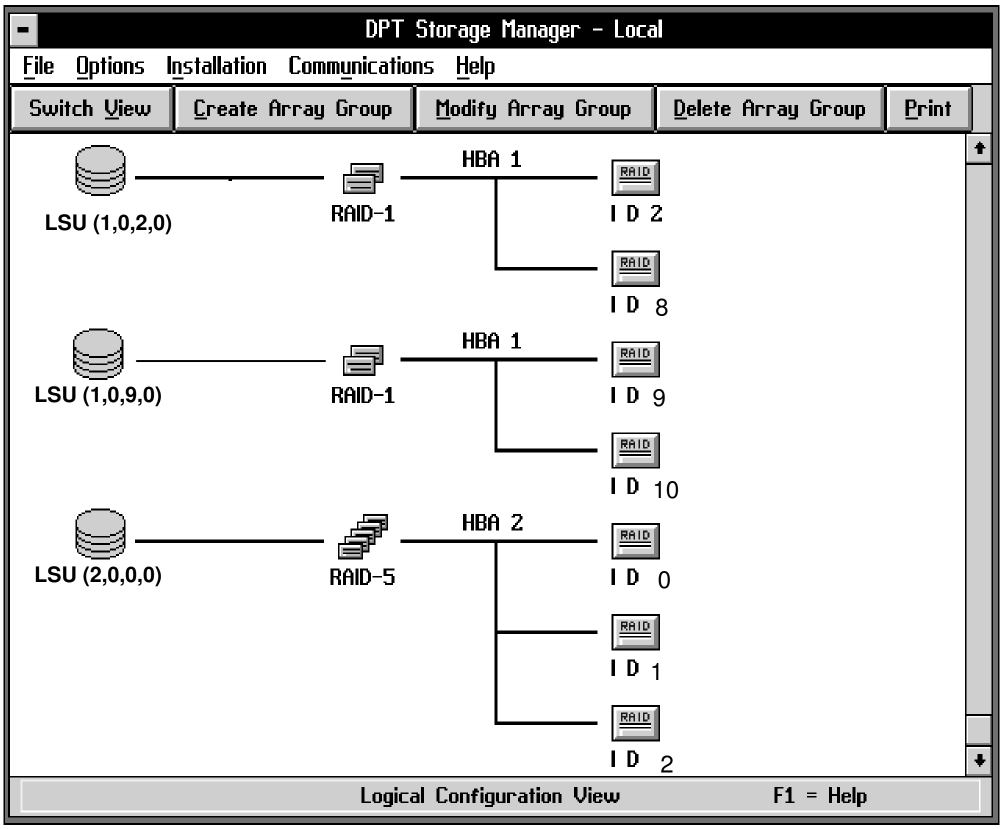

Figure 21. RAID–1 Array Group Window

In the Logical Configuration View screen, the information window for any Array Group may be viewed by clicking on that Array Group's icon. Likewise, information windows for the Hardware Arrays which make up that Array Group may be viewed by clicking on the icons for those arrays.

At the top of the information window, Figure 22, is the array "Name", RAID level and icon, and the Logical Device "Address". Following the Address, the estimated total usable "Capacity" and Array "Status" is displayed. If the information window is for a Hardware Array, a scrolling window displays all "Hot Spares" which protect that array. The "Components" scrolling window displays address, model and stripe size for each hard drive member of the Hardware Array. If the information window is for an Array Group, the "Components" scrolling window displays address or name and stripe size for each Hardware Array which is a member of the group. For RAID–0 Array Groups, member hard drives are displayed.

Buttons include ”Print”, ”Event Log”, ”Diagnostics”, ”I/O Monitor” and ”Name”. In addition, a ”Rebuild” button will appear for redundant arrays which have a failed drive.

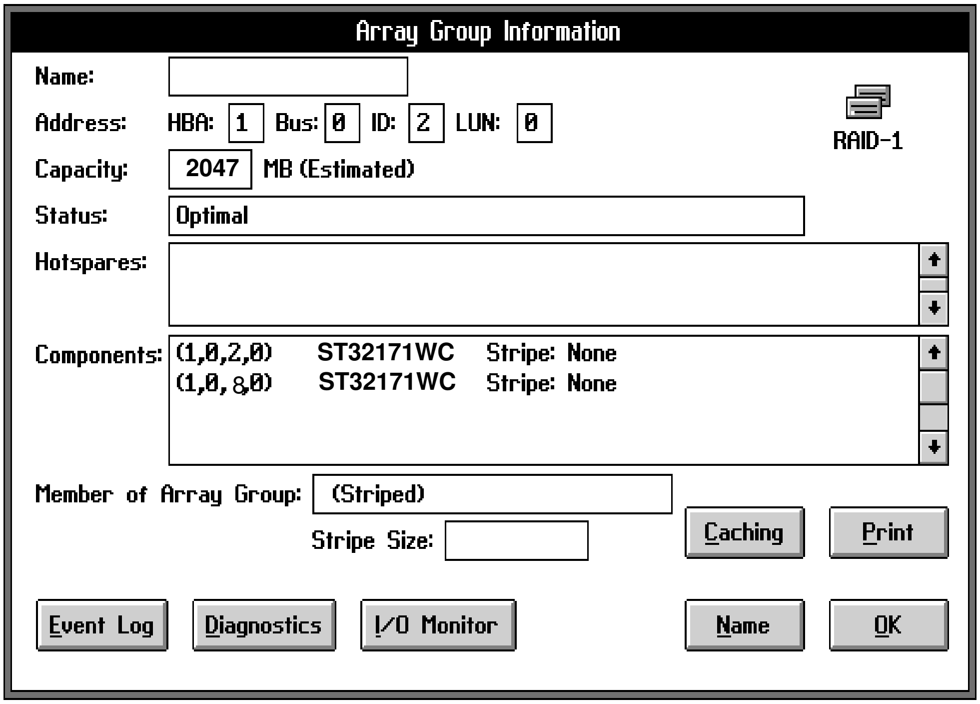

Figure 22. Array Group Information Window for a B4–4 Adapter

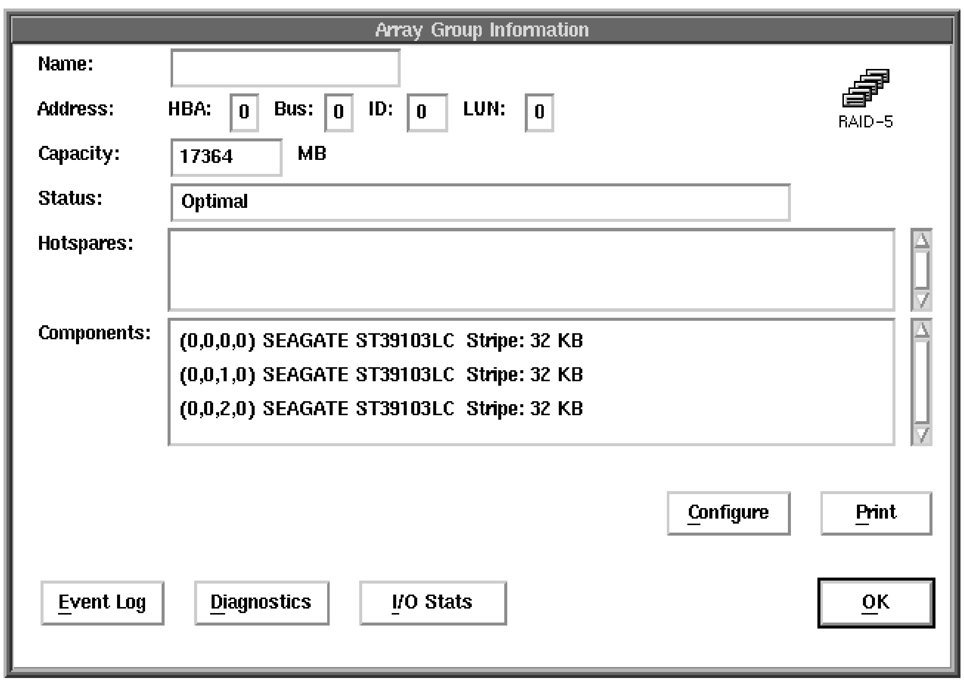

Figure 23. Array Group Information Window for a B4–C Adapter

Creating an Array Group

To create an Array Group, follow the steps below:

1. Select the ”Create Array Group” button.
2. When the ”Select Array Type” window appears, Figure , specify whether or not drive fault tolerance is desired, and whether the Array Group is to be optimized for capacity or performance. As you make your selections, the ”Chosen Array Parameters” will

automatically change to indicate which RAID level and stripe size best suit your requirements. The RAID level and stripe size may be manually specified by selecting the ”Override” button. When you are ready to proceed, select ”Continue”.

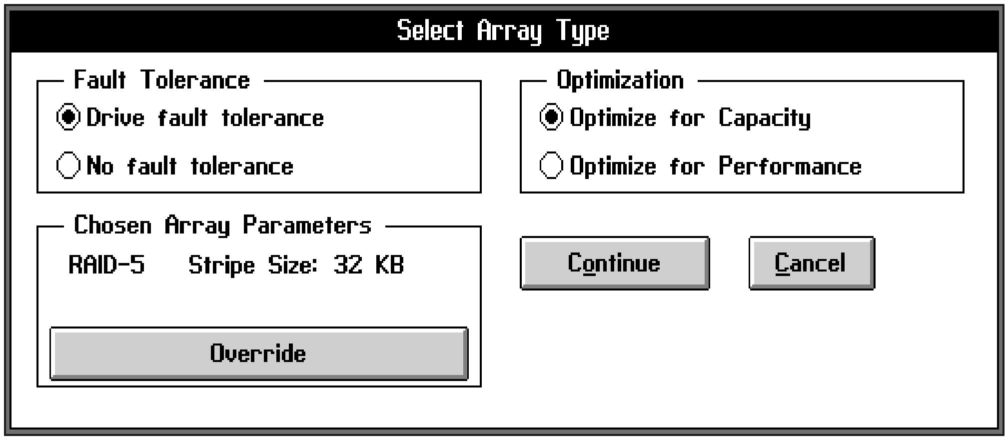

Figure 24. Select Array Type Window

3. The Logical Configuration View screen will appear with the caption: ”Choosing drives for Array (RAID–n)” where n is the RAID level chosen. At the screen top are three buttons labeled ”Include Drive”, ”Remove Drive” and ”Done”. At this point you must select the drives to be included in the Array Group:

   d. To add drives to the Array Group, mark each drive to be added by clicking on it. A green check mark indicates that a drive is selected. Selecting the ”Include Drive” button will cause the marked drive icons to move down to join the new Array Group at the bottom of the screen. You may need to scroll the display down to view the Array Group.

   e. To remove drives from the Array Group, click on the drive icons to be removed and then select the ”Remove Drive” button. The drive icons will move back up to join the other hard drive icons. You may need to scroll the display up to view the drive icons.

   f. During this process some drives may be displayed in a pale blue ”ghosted” color. This indicates that these drives cannot be included in the Array Group unless the present configuration is changed. Either more drives must be added to the Array Group or these drives must be removed from the Array Group. (See **Array Groups**, on page 5-12 for rules regarding the number of drives which may be included in Array Groups.)

4. When you have finished selecting the drives to be included in the new Array Group, select the ”Done” button. You may need to scroll the screen to view the new array. The icon for the Array Group will appear with a black ”missing” flag until the build process has been initiated when you save changes upon exiting Storage Manager.

5. When finished, exit Storage Manager. You will be prompted to save the configuration changes. If you elect to save the configuration, any arrays which were created will be built at this time. For large redundant arrays, this process may take several hours to complete. You may exit Storage Manager and perform other activity on the system while the build continues. However, the Array Groups being built may not be accessed until the build operation has completed. (You may start the build without exiting Storage Manager by selecting ”Set System Configuration” from the ”File” menu.)

   If you wish to monitor the progress of the build operation, Storage Manager may be re–entered and the information window for the Array Group may be viewed. Information windows for Hardware Arrays which are members of the Array Group may also be viewed to monitor the progress for that component of the Array Group.

Modifying an Array Group

The number of drives included in an existing Array Group may be changed by the following procedure.

1. If any data has been written to the Array Group, backup that data onto a storage media such as tape. All data on the Array Group will be destroyed when the array is modified.
2. From the Logical Configuration View screen, select the LSU icon of the Array Group you wish to modify. Then select the ”Modify Array Group” button.
3. Drives may now be added or removed from the Array Group as described in **Creating an Array Group**, on page 5-13.
4. When you are finished, select the ”Done” button. The icon belonging to the modified Array Group will appear with a black ”missing” flag until the array has been built.
5. When finished, exit Storage Manager. You will be prompted to save the configuration changes. If you elect to save the configuration, any arrays which were modified will be built at this time. For large redundant arrays, this process may take several hours to complete. You may exit Storage Manager and perform other activity on the system while the build continues. However, the Array Groups being built may not be accessed until the build operation has completed. (You may start the build without exiting Storage Manager by selecting ”Set System Configuration” from the ”File” menu.)

If you wish to monitor the progress of the build operation, Storage Manager may be re–entered and the information window for the Array Group may be viewed. Information windows for Hardware Arrays which are members of the Array Group may also be viewed to monitor the progress for that component of the Array Group.

6. Once the array build process has completed, restore the backup data onto the modified Array Group.

Deleting an Array Group

To delete an Array Group, follow the steps below:

1. From the Logical Configuration View screen, select the LSU icon of the Array Group you wish to delete. Then select the ”Delete Array Group” button.
2. You will be warned that all data will be lost. Choose ”OK” or ”Cancel”. (An Array Group is not physically deleted until you exit Storage Manager and choose to save changes, or select ”Set System Configuration” from the ”File” menu.)

Logical Device Address

Every SCSI device and Array Group is assigned a ”Logical Device Address” by Storage Manager. This address is used by the OS to access the device or array. Logical Device Addresses appear in parenthesis under the logical device and LSU icons on the left of the Logical Configuration View screen, and are composed of four fields: (HBA, Bus, Device, LUN). Logical Device Addresses are assigned to SCSI devices as follows:

- **HBA** – Host Bus Adapter to which the device is attached. In PCI systems, bus slots are scanned, starting from the lowest slot, looking for SCSI RAID adapters. As SCSI RAID adapters are found, they are assigned numbers incrementally, starting with 1.
- **Bus** – SCSI bus to which the device is attached.
  The SCSI bus on a SCSI RAID adapter is always Bus 0.
  The SCSI bus on an SX4030/1 module is Bus 1.
- **Device** – The assigned SCSI ID for that device.
- **LUN** – SCSI Logical Unit Number for that device. (Normally 0)

Array Groups are automatically assigned the lowest Logical Device Address of any drive in that Array Group, see Figure 25.

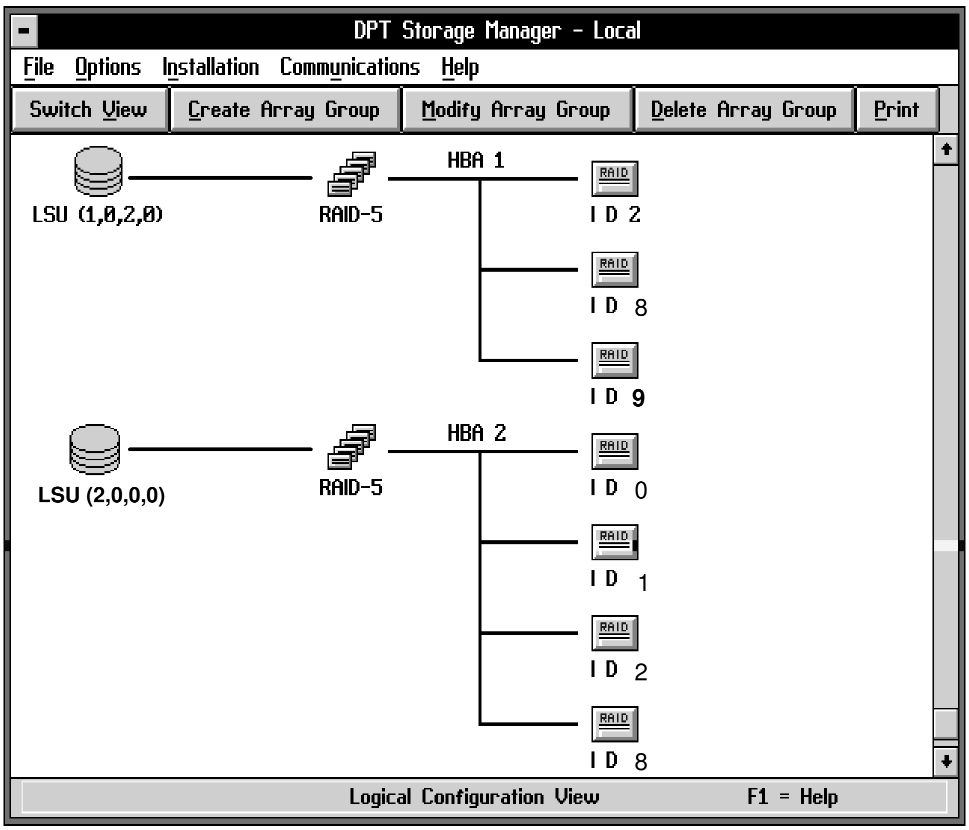

Figure 25. Storage Manager Window Showing Logical Device Addressing

(Once an Array Group has been created, its Logical Device Address will not change due to its lowest drive being replaced by a Hot Spare. When replaced, the failed drive will become the new Hot Spare. If you should choose to no longer make that drive a Hot Spare then Storage Manager will prompt you to choose an unused Logical Device Address for the drive.)

Status

Status is reported by SCSI RAID adapters for arrays and hard drives. Some status conditions are indicated by Storage Manager through status flags on the drive or array icons. More detailed status information can be obtained by viewing the information window for that drive or array. Changes in status conditions are logged and may be optionally broadcast. The following status conditions may be indicated for hard drives or arrays.

Optimal: **Status flag: none** – The drive or array is fully functional and is not running diagnostics which might impact performance.

Verifying: **Status flag: white** – This flag appears on drive or array icons to indicate that hardware or media diagnostics are being run on the drive or redundancy is being verified on the array. Performance may be affected.

Warning: **Status flag: yellow** – This flag appears on a drive icon (with S.M.A.R.T. failure prediction) to indicate that failure may be imminent or on a RAID–1 or RAID–5 array that has lost a drive.

Failed: **Status flag: red** – This flag appears on a hard drive icon when a drive failure has occurred. On an array icon, this flag indicates that a drive has failed in a RAID–0 array or two or more drives have failed in a RAID–1 or RAID–5 array.

Building: **Status flag: blue** – This flag appears on an array icon when an array which has been created is now being built.

Status flag: white – This flag appears on all hard drive icons that are members of an array that is being built.

Rebuilding: **Status flag: blue** – This flag appears on an array icon when the array is being rebuilt after a drive failure.

Status flag: white – This flag appears on a hard drive icon when data is being rebuilt onto that drive.

Missing: **Status flag: black** – This flag appears on a hard drive icon when the drive is physically missing or will not respond to commands on the SCSI bus. On an array icon, this flag indicates that the array has been created or modified within Storage Manager but the array has not yet been built. On a Hot Spare icon this flag indicates that the Hot Spare has not yet been created by the adapter.

Events

Events are generated for detected fault conditions as well as subsystem status changes. The events are grouped into four categories as follows:

- **Level 1: Soft Error** – An error occurred on a disk drive but upon retry, the operation was successful.
- **Level 2: Recoverable Hard Error** – A hard error on a disk drive, adapter or SCSI bus occurred, in which the data could be recovered from ECC or re–synthesized from redundant array information.
- **Level 3: Nonrecoverable Hard Error** – A hard error on a disk drive, adapter or SCSI bus occurred, in which the data could not be recovered from ECC or re–synthesized from redundant array information.
- **Level 4: Status Change** – The status of an array or hard drive has changed. Examples of this would be a drive or array failure or an array build or rebuild operation initiated or completed.

Event Logging

When events occur, they are automatically logged in the cache on the SCSI RAID adapter on which they occurred. In addition, you may specify that Storage Manager maintain an event log on disk. The combined contents of these event logs may be displayed by selecting the ”Event Log” button in any adapter, hard drive or array information window. Only the events pertaining to that adapter, hard drive or array will be displayed.

When the Event Log button is selected, an Event Log Display Window will appear, Figure 26. Within the Event Log Display Window, you may specify that the display be limited only to events of a specific level or higher. Click on the arrow buttons to adjust the ”Display Threshold” to the desired level. Then click the ”View Log” button to view the event messages whose levels are indicated by the bar.

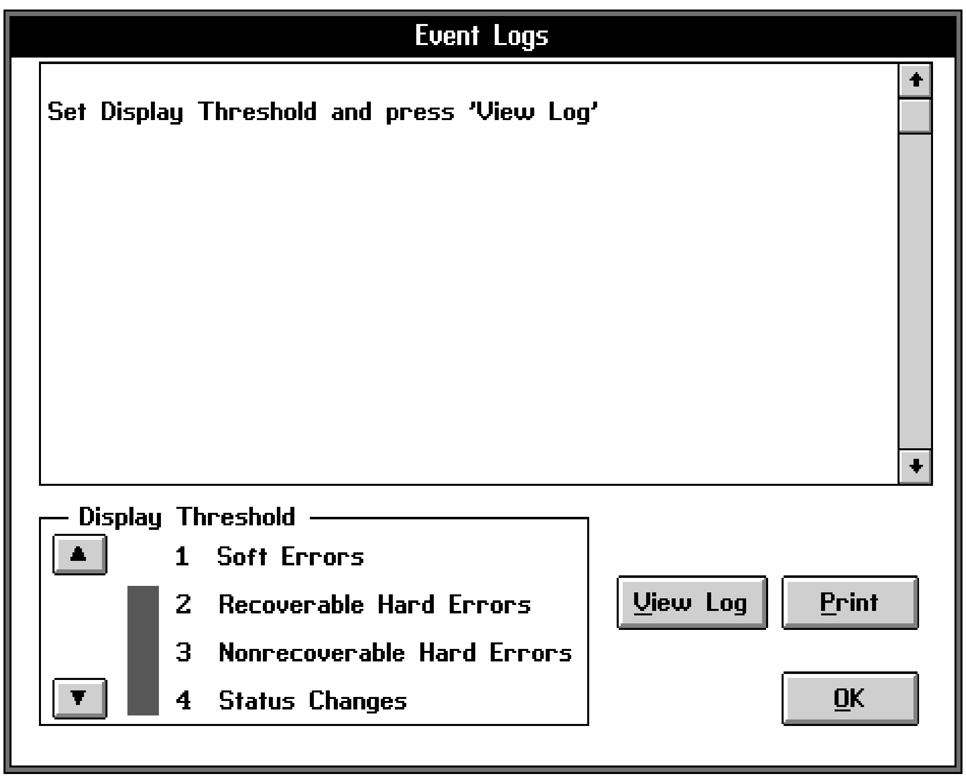

Figure 26. Event Log Display Window

Broadcaster

The Broadcaster collects events logged by SCSI RAID HBAs in the computer on which that broadcaster is running. It then records these events to disk files for each HBA for use by Storage Manager. Additionally, event whose level is greater than or equal to 3 are sent to AIX error log.

Broadcaster for AIX

A Broadcaster is provided for AIX versions 4.1.4 and higher. The Broadcaster collects events from SCSI RAID adapter and records them in the */var/adm/dpt* directory for use by Storage Manager.

Installing the Broadcaster

The Broadcaster is placed, by default, in */usr/lpp/dpt.* During Broadcaster installation, the Broadcaster is to be loaded automatically when the system is booted into multiuser mode. Doing so will allow events to be gathered and recorded without user intervention.

Viewing Events

Events may be viewed from the system console or a terminal through the "Event Log Display" window in Storage Manager, see **Event Logging**, on page 5-18. The Events are also logged in AIX system. The command `errpt -a` allows the Events to be shown from the AIX errlog file without loading in the Storage Manager.

The file */var/adm/sraidlog.log* contains also all events from SCSI RAID adapters, in ASCII format. This file can be viewed using "vi".

Hard Drive Formatting

SCSI RAID adapters are capable of performing format operations on SCSI hard drives and floptical media. This function is accessible through Storage Manager by viewing the device's information window.

Formatting Hard Drives

Formatting options for hard drives include the standard 512–byte format, or, if supported by the drive, a 528–byte format that provides SCSI Bus ECC protection. Storage Manager will present the current drive format as the default selection.

Note: Do not power–down the drive until the format operation has completed. Doing so may cause some drives to be left in an indeterminate state requiring a return to the manufacturer.

To perform a standard 512–byte format, follow the instructions below:

1. Select the "Format" button in the drive's information window. This will cause the "Format Options" window to appear, Figure 27.

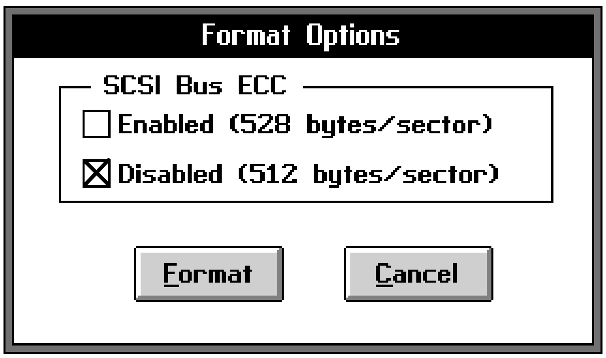

Figure 27. Format Options Window

2. Select the 512–byte format option. Select "OK" and confirm.
3. You may now exit Storage Manager. The format operation will continue even though Storage Manager is not running. To determine if the format has completed, run Storage Manager and view the drive's icon. A white flag indicates that the format is still in progress. If the drive icon contains no flag, then the format is complete. You may view the drive's event log to check that the format executed successfully.

To perform a 528–byte format supporting SCSI Bus ECC protection, follow the instructions below:

1. Select "View README Information" from the Storage Manager "Help" menu and review the list of drives which have been certified to support 528–byte sector formats. If your drive is not listed but the 528–byte format option is selectable in Step 2 below, you may continue. Perform Storage Manager drive diagnostics after the format has completed to assure that the format was successful.
2. Select the "Format" button in the drive's information window. This will cause the "Format Options" window to appear.
3. Select the 528–byte format option. Select "OK" and confirm.
4. You may now exit Storage Manager. The format operation will continue even though Storage Manager is not running. To determine if the format has completed, run Storage Manager and view the drive's icon. A white flag indicates that the format is still in progress. If the drive icon contains no flag, then the format is complete. View the drive's event log to check that the format executed successfully.
5. Ensure that the "ECC Enabled" box in the drive's information window is checked. Perform Storage Manager drive diagnostics to confirm that the format was successful.

Hard Drive Failures

If a hard drive belonging to an Array Group fails, the drive icon will appear marked with a red ”failed” flag on both the Physical and Logical Configuration View screens.

If the failed drive belongs to a RAID–0 Array Group, the icons for the drive, Hardware Array and Array Group will be marked with red failed flags, indicating that the array has suffered a catastrophic failure and that data on that Array Group has been lost.

If the failed drive belongs to a RAID–1 or RAID–5 Array Group, the Hardware Array icon will be marked with a yellow ”warning” flag, indicating that the array is currently running in degraded mode. If another drive belonging to the same Hardware Array fails before the failed drive has been replaced and the array Rebuilt, the yellow flag will turn to red, indicating that the array has failed and that data has been lost.

Audible Alarm

Failure of a hard drive which is a member of an Array Group will cause an audible alarm to sound on the adapter with the failed drive. This alarm will be automatically silenced when Storage Manager is entered. Alarms which occur while Storage Manager is running can be silenced by selecting ”Turn Off Audible Alarms” from the ”Options” menu.

Hot Swapping Drives

When a drive in a RAID–1 or RAID–5 Array Group fails (indicated by a solid amber LED) and that drive is not protected by a Hot Spare, follow the procedure below to replace the failed drive:

- Before removing that drive, select ”Rebuild” in the array’s information windows even if the drive has not been replaced to use the drive again. The success is indicated by a white flag on the storage manager and a flashing amber fault LED on the drive (Rebuild in progress).
- If that operation does not succeed, remove the failed drive as described in ”Replacing drives” on page 4-3. When a good drive is inserted in place of the failed drive, the array rebuild operation will start automatically.

The drive appears with a white flag to indicate that a rebuild operation is in process. The Hardware Array and LSU icons appear with yellow ”degraded” flags. The percentage completion of the rebuild operation is displayed by bar graphs in both the drive and Array Group information windows. Upon completion of the rebuild, the white flags will disappear.

Note: SCSI devices contain specially designed glitchless connectors on the SCSI bus so that failed drives may be removed or replaced without first quieting the bus. In addition, adapters detect when a failed drive has been physically replaced and will automatically start the rebuild operation. Consequently, when using RAID adapters, failed drives may be removed and replaced without running Storage Manager.

Hot Spares

Any non–arrayed hard drive can be assigned as a Hot Spare. Hot Spares are not accessible to the OS and are used to automatically replace failed drives in RAID–1 or RAID–5 arrays. Hot Spares can protect drives of equal or less capacity that are attached to any SCSI bus on the same adapter as the Hot Spare.

To make a drive a Hot Spare, select the ”Make HotSpare” button in the drive’s information window. Selecting the ”Remove HotSpare” button in the information window of a Hot Spare drive will cause that drive to be assigned as a normal hard drive, accessible by the O.S.

When an array which is protected by a Hot Spare has a drive failure, SCSI RAID adapter will automatically begin to rebuild the missing data onto the Hot Spare. During this process, Storage Manager will swap the positions of the failed drive and the Hot Spare in the Logical Configuration View. The failed drive will appear with a red ”failed” flag in the old position of the Hot Spare, and the Hot Spare will appear as a member of the Array Group with a white flag to indicate that a rebuild operation is in process. The array and LSU icons will appear with yellow degraded flags.

Upon completion of the rebuild, the Hot Spare icon and flags will disappear and the drive will be displayed as a normal member of the Array Group. The array rebuild operation will start automatically The red flag will remain on the failed drive until that drive is replaced. Follow the procedure below to replace the failed drive:

- Remove the failed drive as described in ”Replacing drives” on page 4-3.
- When a good drive is inserted in place of the failed drive, it will become the new Hot Spare, replacing the lost Hot Spare which is now a member of the rebuilt Array Group.

Diagnostics for a B4–4 Adapter

On–line diagnostics may be performed on any hard drive or array by selecting the Diagnostics button in the information window of the hard drive or array. Diagnostics are performed automatically by the SCSI RAID adapter, concurrent with normal system operation and require no host computer intervention. Although there is no impact on the host CPU, system performance may be impacted due to increased drive utilization. If this is a concern, you may want to use the diagnostics scheduling features discussed later in this section or assign a low level to background task priority.

SCSI Device Diagnostics

Selecting the Diagnostics button from a hard drive information window will display the SCSI Device Diagnostics window, Figure 28:

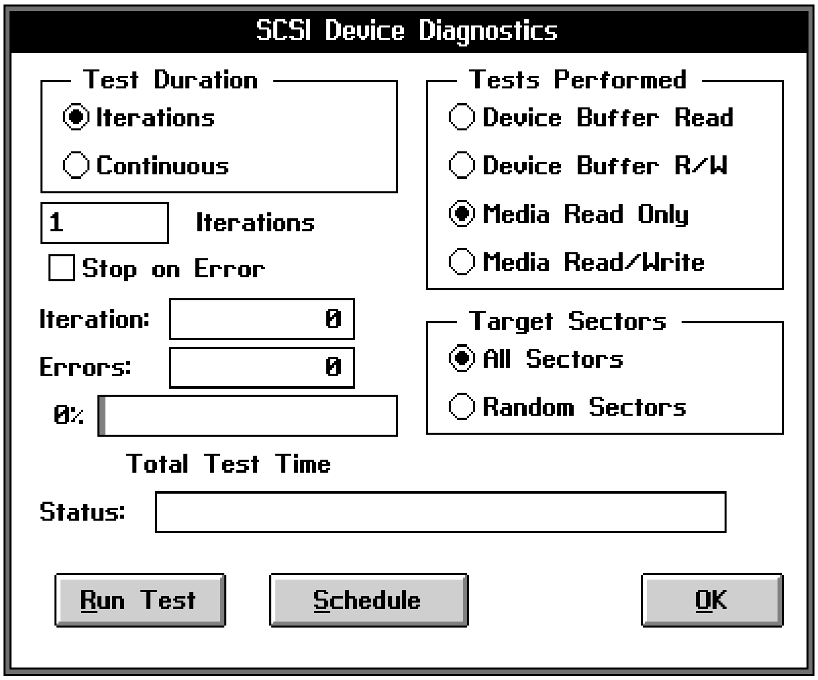

Figure 28. SCSI Device Diagnostics Window for a B4–4 Adapter

Diagnostics are non–destructive and may be performed without interrupting normal usage of the device. Diagnostics include device buffer tests, which test the integrity of the SCSI bus by sending and receiving data to a test buffer on the hard drives, and media read tests which test that each sector on the hard drives contains error–free data. Media read/write tests momentarily hold user I/O to each sector while a test pattern is written, read and compared. The original data is restored for each sector after it is tested. Diagnostics may be performed either for a specified number of iterations or in a continuous loop. Diagnostics may also be performed on multiple drives or arrays simultaneously. The current diagnostics status and percentage complete are displayed at the bottom of the diagnostics window. The statuses are: ”Running”, ”Completed”, ”Stopped on Error”, ”Stopped by User” and ”Scheduled”.

Array Diagnostics for a B4–4 Adapter

Data verification diagnostics can be performed for RAID–1 and RAID–5 arrays to assure that the redundant information contained in the array is consistent. Selecting the Diagnostics button from a Hardware Array or Array Group information window will display the Array Diagnostics window, Figure 29:

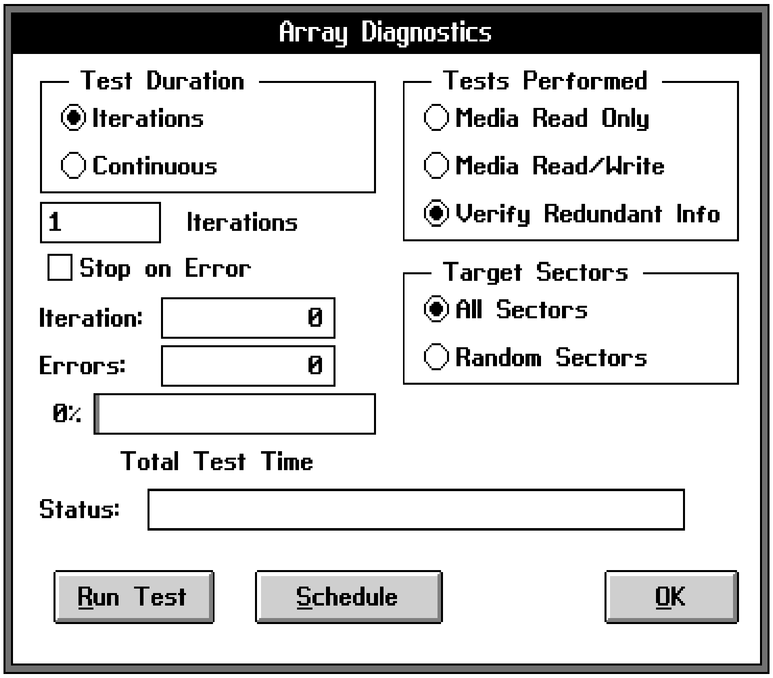

Figure 29. Array Diagnostics Window

If any inconsistencies in the data redundancy are found, they will be made consistent. In the case of RAID–1 arrays, the mirrored drive pairs are compared sector by sector to assure that both drives contain identical data. In the case of RAID–5 arrays, parity is recalculated and checked against the stored parity information.

Note: Under normal conditions, data inconsistencies should never occur. However, a power failure which interrupts an array write operation can possibly cause inconsistencies to occur. **Making the data consistent again through the Verify function does not necessarily ensure that the new consistent data will be the correct data.**

Scheduling Diagnostics for a B4–4 Adapter

Diagnostics for SCSI devices or arrays may be scheduled to begin up to six days in the future and repeat periodically with an interval of up to thirty days. For example, an array verification could be scheduled to occur automatically every weekend or every night.

To schedule a diagnostic, first set the desired diagnostic parameters. Then select the ”Schedule” button to display the ”Schedule Diagnostic” window, Figure 30.

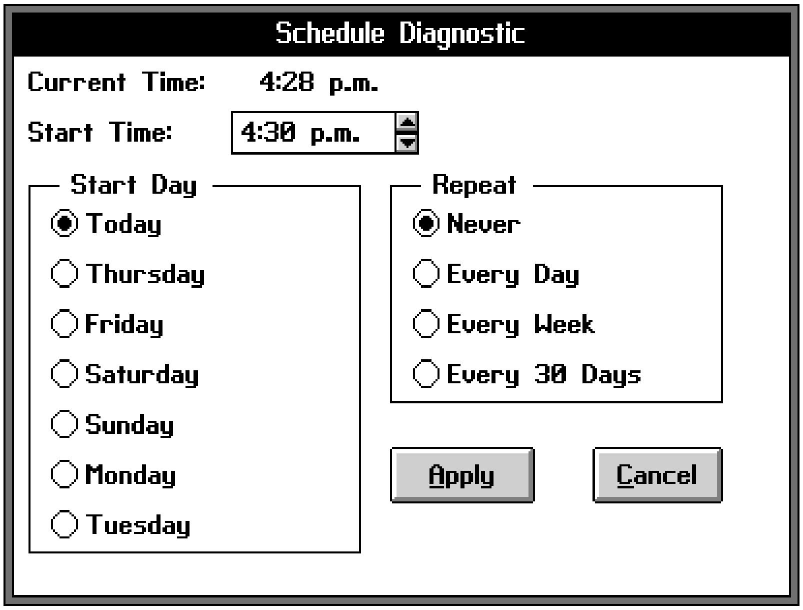

Figure 30. Schedule Diagnostics Window

You may now specify the day and time the diagnostic is to begin and the frequency at which it will be repeated. The defaults are to start the diagnostic at the current day and time and to not repeat. ”Start Day” determines when the diagnostic will begin and may be set to any day from the current day to six days in the future. For example, if it is Wednesday then the displayed options will be ”Today”, ”Thursday”, ”Friday”, ”Saturday”, ”Sunday”, ”Monday” and ”Tuesday”. ”Start Time” is selectable starting from the current time to 11:30PM or any half hour increment in between. If a day other than ”Today” is chosen, the ”Start Time” options will be from 12:00AM to 11:30PM. ”Repeat Period” determines when the diagnostic will repeat and may be set to ”Never”, ”Every Day”, ”Every Week” or ”Every 30 days”.

Background Task Priority

Array rebuilds, array verification and drive diagnostics may be executed without taking the drives or arrays off–line. These functions are performed entirely by the SCSI RAID adapter, as background tasks transparent to the OS, allowing for normal system operation. However, since the adapter interleaves I/O from the OS with I/O from the background task, this may sometimes impact system performance.

The relative priority given to I/O from the OS as compared to the background task can be controlled by selecting ”Background Task Priority” from the ”Options” menu. The ”Task Priority” section of the displayed window, Figure 31, assigns background task priority for all SCSI RAID adapters in the computer. Priority is set by the slide–bar which may be moved between ”Background” and ”Foreground” in increments. ”Background” specifies that background tasks will be processed only when there is no disk I/O from the computer for a period of at least second. As the bar is moved toward ”Foreground”, background tasks will consume more of the bandwidth. Regardless of this parameter setting, 100% of the bandwidth is allocated to background tasks during periods of no disk I/O from the OS.

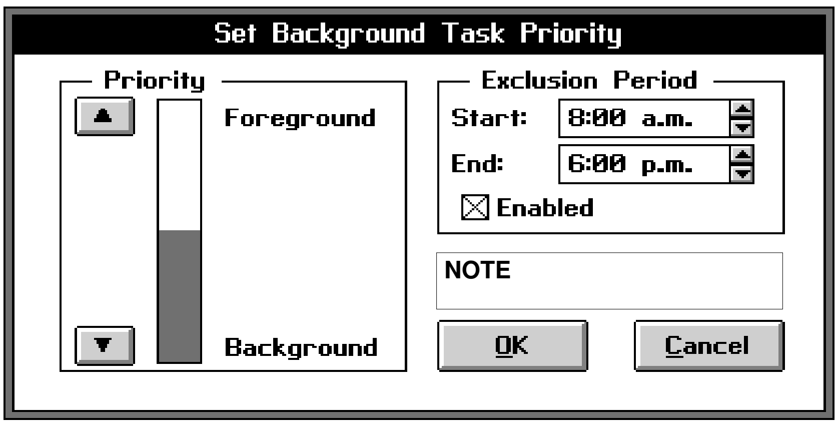

Figure 31. Background Task Priority Window

Note: “Exclusion Period” is only applicable to the B4–4 adapters

The ”Exclusion Period” section of the window specifies a time period during which array rebuilds will not be initiated. For example, this option could be used to exclude rebuilds during working hours so that maximum disk bandwidth would be available. If a drive failure occurred during working hours, the subsequent automatic rebuild onto a Hot Spare would be delayed until after hours.

Note: Tasks started prior to the exclusion period will run to completion, even if they run into the exclusion period.

The exclusion period is established by checking the ”Exclusion Period Active” check box and defining the ”Start Time” and ”Stop Time” of the exclusion period. Start and stop times may be set in 15 minute increments from 12:00AM (midnight) to 11:45PM.

I/O Monitor

SCSI RAID adapters automatically tally I/O operations in cache RAM for analysis of SCSI subsystem I/O loading. These numbers can be viewed by selecting the ”I/O Monitor” button in the Information Window for any hard drive, Hardware Array or Array Group, Figure 32. By analyzing these statistics, the array architecture, cache and stripe size can be optimized for your system.

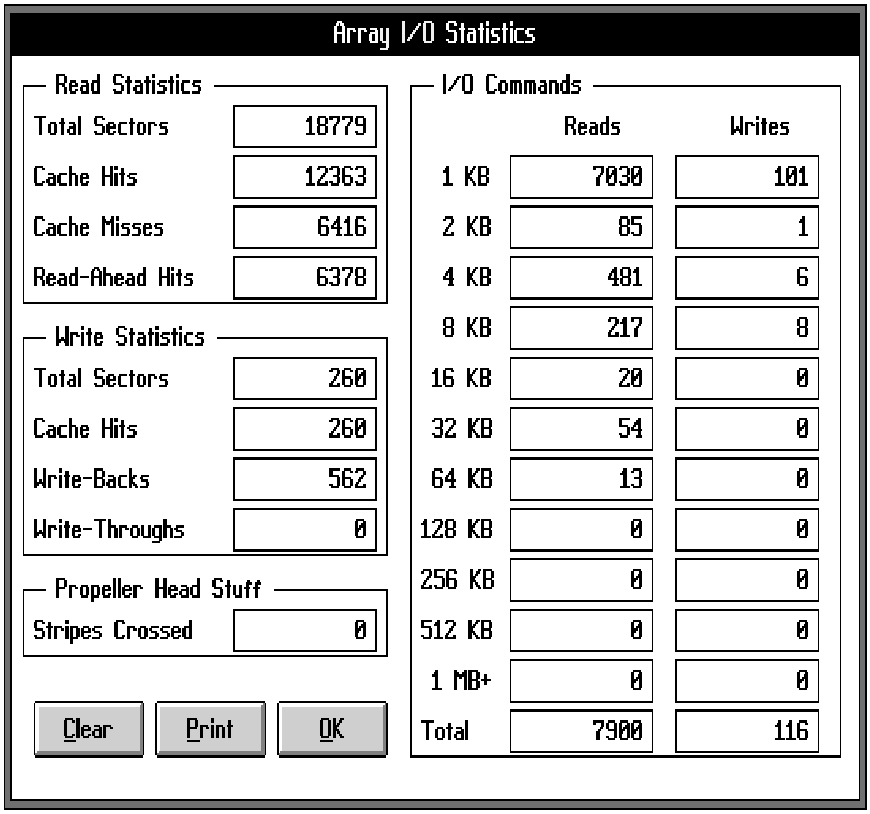

Figure 32. I/O Monitor Statistics Window

- **I/O Commands** counts the read and write commands issued by the computer to the adapter. With RAID and caching, the number of commands issued to the hard drives may significantly differ from this value.
- **Total Sectors** reports the total number of sectors read and written from the computer.
- **Cache Hits** counts the total number of sectors read or written from the computer that were found in the adapter’s cache, thus not requiring a disk access.
- **Read–Ahead Hits** is the number of the read cache hits due to data residing in cache from previous disk read–ahead operations.
- **Write–Backs** is the number of sectors written to disk that were held in the adapter cache and written some time after the computer’s write command terminated.
- **Write–Throughs** is the number of sectors written to disk before write command termination.

”Write–Backs” plus ”Write–Throughs” equals the total number of sectors written to disk by the adapter. In RAID–1 arrays, ”Write–Backs” plus ”Write–Throughs” will be up to twice the value of ”Total Sectors” since each sector written from the computer results in a write to each mirrored disk. In RAID–5 arrays, each write from the computer will generate up to two disk reads and two disk writes. Because of the adapter’s cache, the number of sectors read or written to disk may be less than this value.

Environments with a large number of sequential reads should generate a high number of ”Read–Ahead Hits” relative to ”Total Sectors”. These hits reduce the number of seeks and increase performance. Read–ahead hits can be increased by adding more cache RAM to the adapter or by increasing the ”Maximum % Read Ahead” in the ”HBA Caching Configuration” window. A high percentage of 4KB or smaller I/O operations is another indication that the adapter cache will be beneficial. If cache hits are low, adding more cache RAM may increase the hit rate. Systems that have a large number of disk writes typically derive significant performance benefits from the adapter cache.

For RAID–0 and RAID–5 arrays, a default stripe size of 32KB has been chosen for optimal performance in a variety of environments. However, some specific environments may benefit from smaller or larger stripes. In RAID–5 environments, if a write falls across all drives, thereby causing a large number of stripe crossings, the RAID–5 parity can be generated in a much more efficient manner, increasing performance. However, if the write crosses one or more stripes but does not involve all of the drives, the performance will be negatively impacted. Therefore, the stripe size should be chosen relative to the I/O size and number of drives in the array so that most I/O operations either cross no stripes and involve a single drive or cross many stripes and involve all the drives in the array.

Battery Backup Configuration for a B4–C Adapter

Note: During the initial calibration cycle for the Battery Backup module, the adapter disables automatic, low–battery write–through mode. After the calibration, use the Battery Backup Configuration dialog to set a threshold for entering write–through mode when the battery charge drops below a predetermined level.

This option lets you view the status of the Battery Backup module and set operating parameters when the battery reaches a predetermined level. The battery Status and available Backup capacity (in hours) is displayed. The Backup value is monitored periodically by software and changes whenever the battery pack is charging or discharging.

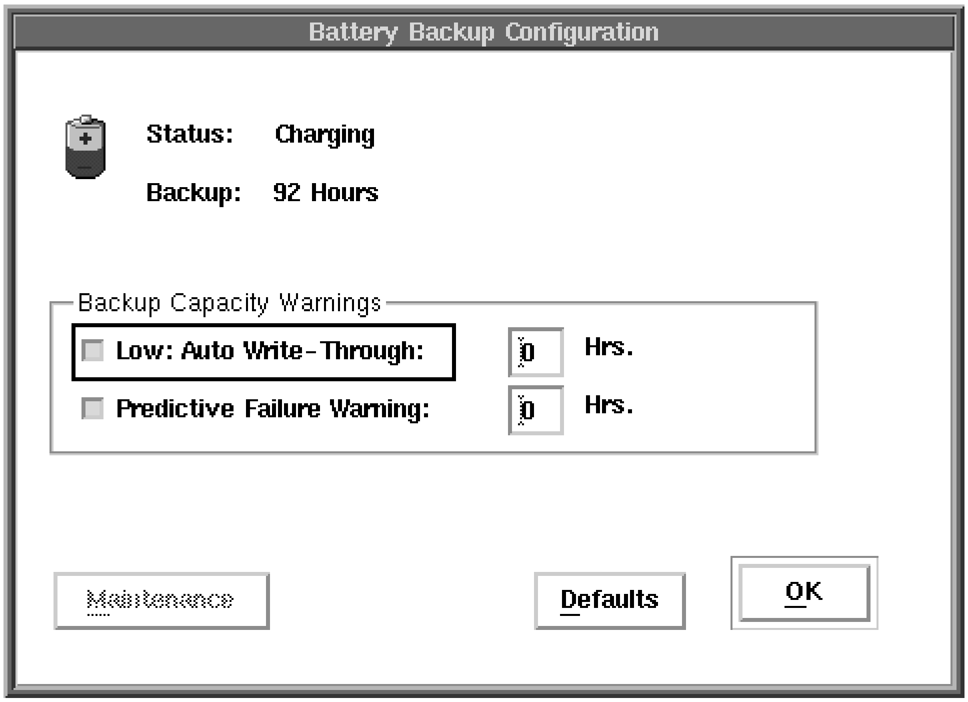

The Backup Capacity Warnings parameters let you activate the following options when the module reaches a predetermined level of remaining backup capacity (in hours). The capacity warning thresholds can change depending on the capacity and number of memory modules. Generally, more cache will result in lower battery backup capacity.

- **Low: Auto Write–Through**: when checked this option automatically sets the cache to write–through mode when the specified number of hours (Hrs.) of battery capacity remain. The number of hours must be greater than or equal to the number of hours for a Predictive Failure Warning.
- **Predictive Failure Warning**: when this option is checked a warning message will be issued when the battery backup capacity is almost depleted. Enter the number of hours (Hrs.) of remaining backup capacity when this message should be issued.

Click the Defaults button to set the Battery Capacity Warnings to their default values. Actual default values are calculated by the controller based on the size of the cache memory.

Battery Status Messages

Battery Module status is reported as one of the following:

Full

Charging

Initial calibration charge

Initial calibration discharge

Initial calibration recharge

Maintenance calibration discharge

Maintenance calibration charge

Initial Calibration

The Initial Calibration operation is started when a Battery Backup Module is first installed on a controller. This function ensures the battery is fully charged for subsequent backup operation. The calibration cycle requires approximately 24 hours to complete. The module will not be able to provide backup capability during this operation. When the process is completed, the module is ready for normal operation.

This function has three phases:

1. Initial charge: The battery is charged to its full capacity.
2. Discharge: The battery is discharged until no backup capacity remains.
3. Recharge: The battery is recharged to full capacity.

Note: The Battery Backup module will not be able to provide backup capability during this operation.

The Maintenance operation can be used to ensure the battery remains capable of accepting a full charge. The Maintenance cycle completely discharges and then recharges the battery pack. This helps to eliminate any voltage–depression effects resulting from the battery pack being partially discharged and then recharged repeatedly during normal operation.

When this option is selected, the date of the most recent Maintenance cycle is displayed.

Flash Configuration for a B4–C Adapter

This dialog is displayed when you click on the flash button in the Configure Host Bus Adapter window.

Note: It is easier to perform an update if the image files are placed in the root directory of a diskette. Each component must be upgraded as a separate operation.

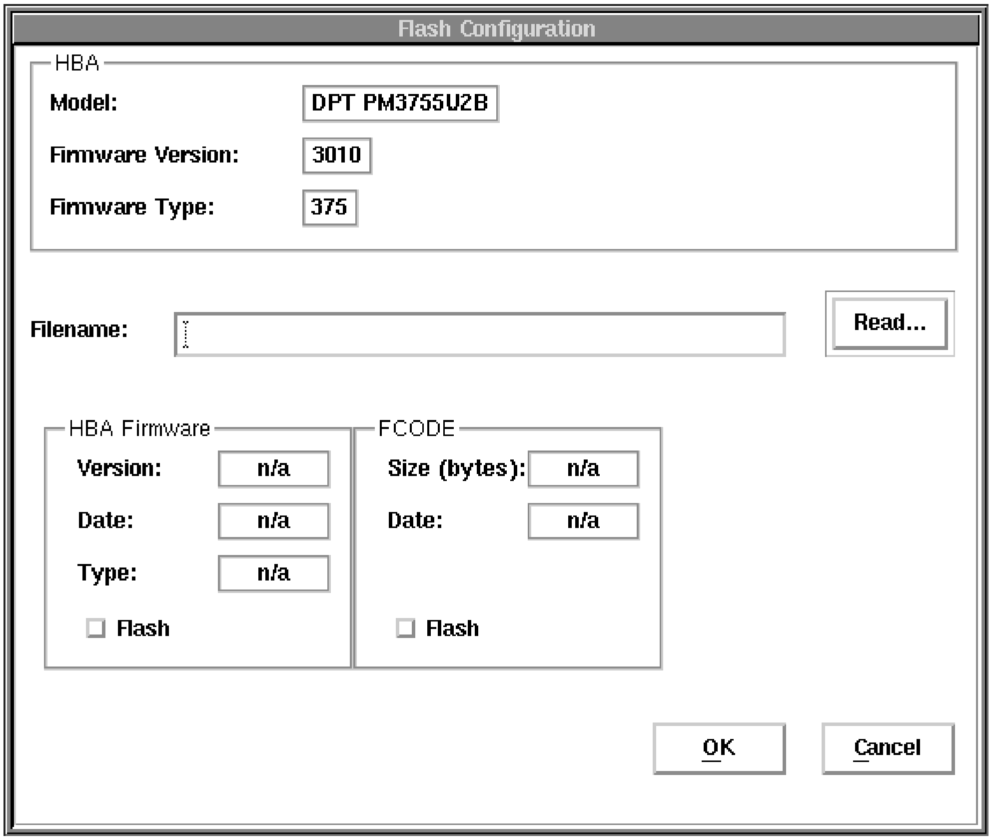

The current controller Model, firmware Version and Type are displayed. To specify an image file for the flash operation, you can type a path and filename in the ilename field or click on Browse to use the standard file selection dialog.

The firmware image for a B4–C adapter is located in */etc/microcode/firm375.fwi.* This file will typically be updated during installation or update of fileset “devices.pci.441001a5.rte”. The adapter firmware will be automatically flashed during reboot following the software installation.

When you select an image file, Storage Manager reads the file to determine the type of image selected, firmware or Fcode. The Version, Date and Type are displayed in the corresponding section of the dialog.

Click OK to begin the flash operation. Click Cancel to return to the Configure Host Bus Adapter window.

Setting the Configuration

Under the File menu are four options:

- **Read System Configuration** causes Storage Manager to re–read the current hardware configuration. Any changes which have been made and not saved will be lost. ”Read System Configuration” happens automatically when Storage Manager is run.
- **Set System Configuration** causes Storage Manager to save to hardware, any changes that have been made to the SCSI subsystem configuration. If any Array Groups have been created or modified, then setting the system configuration will cause the adapters to initiate array build operations.
- **Load Configuration File** allows a saved configuration to be loaded into Storage Manager and applied to the current hardware.
- **Save Configuration File** allows the current configuration, or any changes to that configuration, to be saved to a file for later use. This feature allows SCSI subsystems to be configured for other machines and later loaded from the configuration file.

Chapter 6. SCSI RAID Adapter – Theory of Operation

Describes the theory of operation.

Disk Caching

SCSI RAID adapters may be configured with 1 to 64MB of hardware cache. Adding more cache RAM to the SCSI RAID adapter will typically result in better RAID performance – especially in RAID–5 environments.

Software and Hardware Caching

The SCSI RAID adapter cache is specifically designed to enhance performance in environments like UNIX and Novel which are already software cached. Although software caches are almost always beneficial, in larger, more heavily loaded systems, hardware caching offers a significant advantage when combined with the software cache in the operating system.

The primary purpose of the operating system’s cache is to eliminate disk reads. If data requested by an application is found in the OS cache then a disk read will be eliminated, see Figure 33. Software caches can also improve application response time by immediately accepting data to be written to disk from the application programs but postponing the actual writes until the disk is otherwise idle. Write caching in the OS provides benefits for lightly loaded systems but experience degradation in larger multiuser environments with heavier loading. This is because frequent periods of system idle time are required to flush the OS cache buffers to disk.

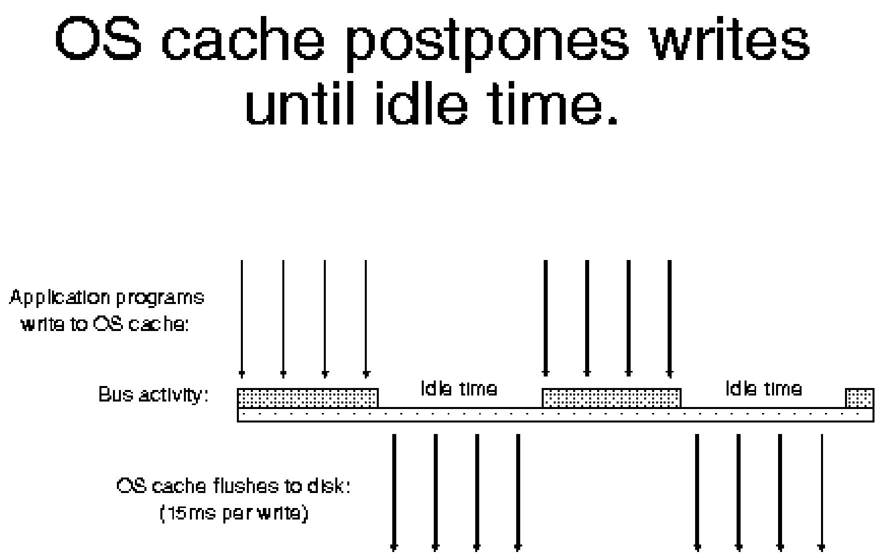

Figure 33. OS Cache

As disk activity becomes more intense, periods of system idle time become shorter and less frequent, Figure 34, causing the OS cache to fill up with dirty sectors waiting to be written to disk. These sectors take up space which could be used for more current data. When cache is added to a SCSI RAID adapter, the operating system's cache buffers must still flush but the adapter receives the flushed data in a fraction of the time.

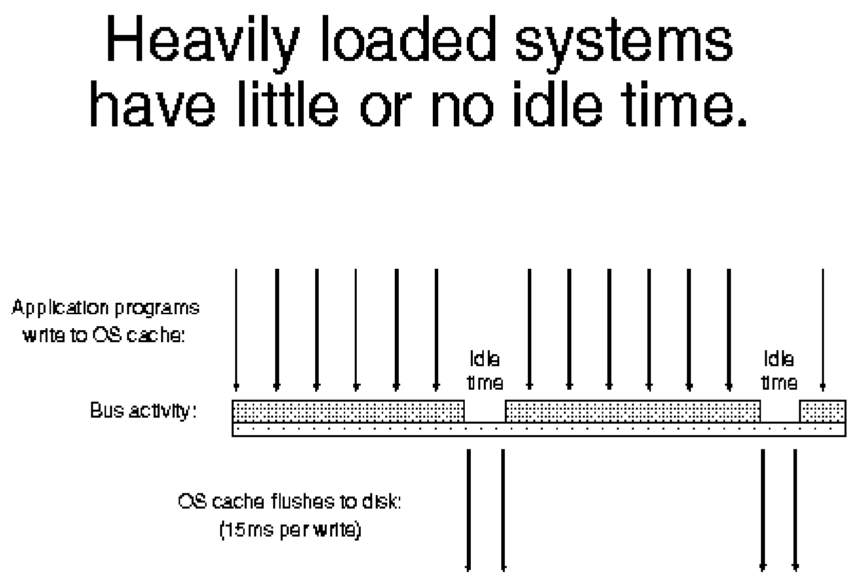

Figure 34. System Idle Time

The adapter later copies the data back to disk in elevator–sorted order without interfering with other system activity. As an example, consider an operating system with 4MB of dirty data in its cache buffers, Figure 35. With a 4KB record size and an average disk I/O time of 15ms per record, this flush operation would take 15 seconds. However, with the adapter cache installed the entire flush operation from the software cache would take approximately 1 second instead of 15 seconds. The SCSI RAID adapter would then flush its cache to disk in elevator–sorted order using its separate I/O bus, concurrently with other system activity.

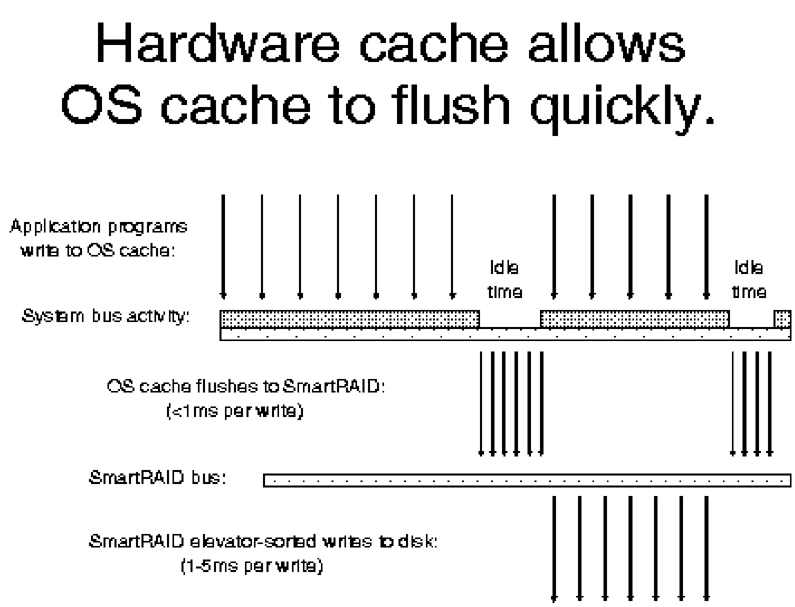

Figure 35. Flushing of OS Cache

Elevator Sorting

Flush operations from the SCSI RAID adapter cache occur much faster than from the OS cache, since all data in the adapter cache is kept in elevator sorted order. Data is written to disk in order of increasing cylinder, head, and sector number, minimizing physical disk seeks and rotational latency. This feature, called ”Elevator Sorted Write–Back”, eliminates much of the head thrashing commonly associated with disk–intensive operating systems like Windows NT and UNIX. By elevator sorting the sectors in cache, the average write access time to the drive may be reduced from the normal 15ms (seek plus rotation) time, to between 1 and 5ms, Figure 36.

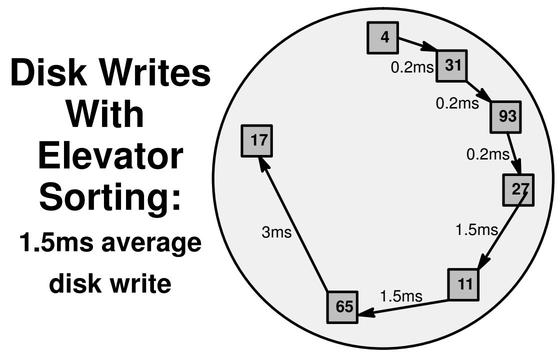

Figure 36. Disk Writes With Elevator Sorting

A good analogy is an elevator in a multi–story building. Consider how much slower the elevator would be if it went to the requested floors in the same order the buttons were pressed, instead of in order of floor number!, Figure 37. An additional benefit of elevator sorting is that multiple short disk write operations can often be combined within the adapter cache into one large write operation, thereby significantly reducing SCSI overhead.

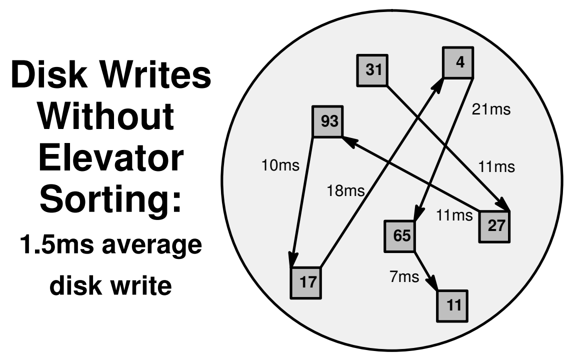

Figure 37. Disk Writes Without Elevator Sorting

Write–back Algorithm

Whenever a SCSI RAID adapter has not received an I/O request from the computer for a particular hard disk drive for 250ms, it will begin to flush the dirty pages for that disk. (Flush operations copy data from cache to disk but do not remove the data from cache.) Since the data is written to disk in elevator sorted order, several seconds of idle time is normally sufficient for all dirty pages to be flushed. Flush operations are also initiated whenever more than 80% of the cache available for write caching becomes dirty. The flush will continue and the execution of further write commands from the computer will be delayed until less than 80% of the cache is dirty.

The 80% ”Maximum Percentage Dirty” limit and the 250ms ”Write Back Delay” may be modified through the ”HBA Caching Configuration” window accessible from the adapter’s Information Window in Storage Manager.

Read–ahead Algorithm

Another way that average disk access time is reduced is through disk read–ahead caching, Figure 38. Since disk data is often grouped in clumps of contiguous sectors on the disk, it is sometimes beneficial to continue to read sequential sectors into cache following a sector that has been requested by the computer. This is known as the principle of ”locality of reference”. SCSI RAID adapter's algorithms analyze the pattern of disk access, looking for cases where sequential I/O may be occurring. If it is determined that a read command from the computer is part of a pattern of sequential reads, the adapter will bring additional sequential sectors into cache so that future read commands will result in cache hits.

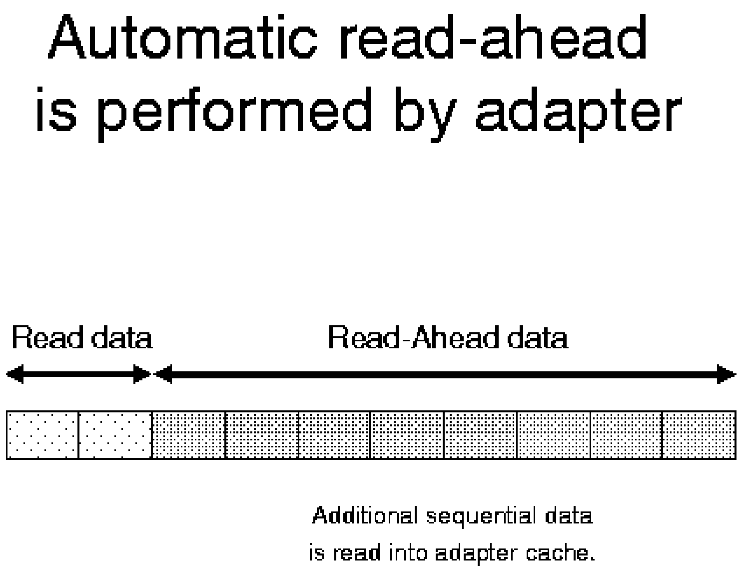

Figure 38. Read–ahead Disk Access

In order for disk read–ahead to be effective in a multiuser environment, it is necessary to cache a large number of read–ahead sectors from many different areas of the disk. This is because multiuser and networking operating systems contain highly fragmented files and time–slice I/O requests from many different users or tasks. Data which has been cached by a read–ahead operation may not be accessed until many commands in the future when the same user or task has had another chance to access the disk.

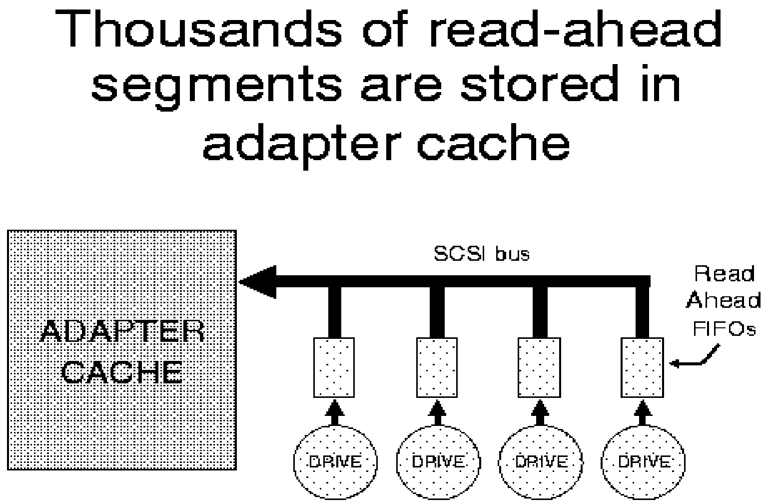

Figure 39. Read–ahead Caching

In addition, patterns of sequential I/O are harder to detect in large multiuser systems since they are typically interleaved with I/O from many other users. This is why a SCSI drive's read–ahead FIFO, which can efficiently store only a limited number of data segments, may provide acceptable performance in small systems, but must be augmented by the much larger cache on the SCSI RAID adapter in systems with more users.

By default, SCSI RAID adapter limits the amount of cache which can hold read–ahead data to 30%. This "Maximum Percentage Read Ahead" limit may be modified through the "HBA Caching Configuration" window accessible from the adapter's Information Window in Storage Manager.

Optimizing OS Cache Size

When configuring a server, it is important to take into account the effect of the operating system's cache on the disk subsystem. Although many disk read operations may be eliminated by the OS cache, all writes must get through to the disk sooner or later. In effect, the OS cache acts as a read filter for the disk subsystem.

As the OS cache size is increased, more reads hits are serviced from the OS cache resulting in fewer reads issued to the disk subsystem, Figure 40. However, the number of disk writes stays relatively constant. At some point, adding more cache to the OS will not result in significantly reduced disk I/O. At this point it is better to add more cache to the adapter to improve write throughput.

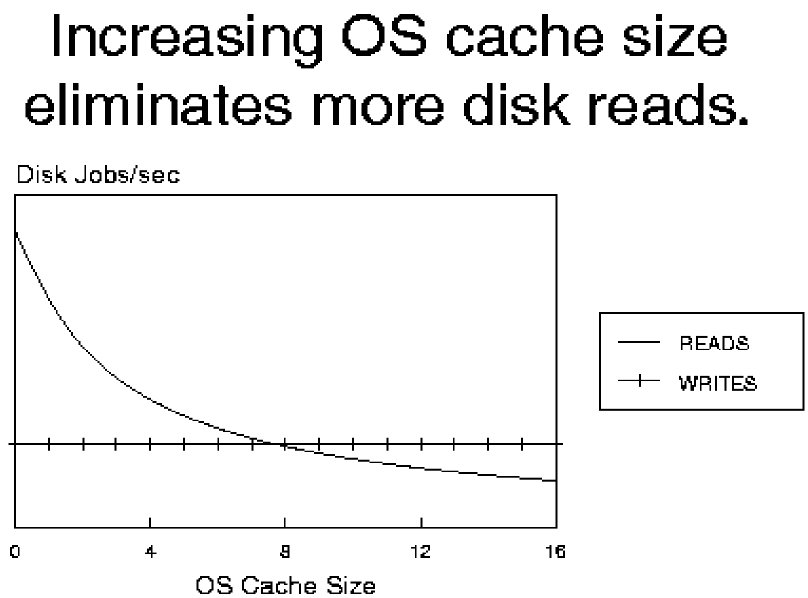

Figure 40. Optimizing OS Cache Size

Note: If too large a percentage of the system RAM is allocated to disk cache or I/O buffers, not enough will be left to hold operating system code and application programs. The computer may excessively swap or page application and OS code to and from the disk, resulting in system performance degradation. As a rule of thumb, it is best to limit the operating system cache and I/O buffers to no more than 10% of total system RAM.

Optimizing Adapter Cache Size

The SCSI RAID adapter cache is used primarily to process elevator sorted write–back and read–ahead segments. Increasing the size of this cache makes these features more efficient. More cache RAM allows the adapter to hold and elevator–sort more records, resulting in closer spacing between consecutive write–back segments. This results in lower average access time for disk write operations and thus higher disk throughput. The SCSI RAID adapter cache may be expanded incrementally as needed to maintain optimum system performance. More cache is typically required for systems with many active users than for single–user systems. The same holds true for systems with large data files as opposed to small files. If a system has a large number of users who all access the same data, less cache may be needed.

Note: As a rule of thumb, the adapter should be configured with 0.5MB of cache for each active user on the system. With RAID–5 arrays, it is always recommended to configure the adapter with at least 4MB of cache.

RAID

In 1987, Patterson, Gibson and Katz at the University of California Berkeley, published a paper entitled ”A Case for Redundant Arrays of Inexpensive Disks (RAID)”. This paper described various types of disk arrays, referred to by the acronym ”RAID”. The basic idea of RAID was to combine multiple small, inexpensive disk drives into an array of disk drives which yields performance exceeding that of a Single Large Expensive Drive (SLED). This array of drives appears to the computer as a single logical storage unit or drive.

The Mean Time Between Failure (MTBF) of the array will be equal to the MTBF of an individual drive, divided by the number of drives in the array. Because of this, the MTBF of a non–redundant array would be too low for mission–critical systems. However, disk arrays can be made fault–tolerant by redundantly storing information in various ways.

Five types of array architectures, RAID–1 through RAID–5, were defined by the Berkeley paper, each providing disk fault–tolerance and each offering different trade–offs in features and performance. In addition to these five redundant array architectures, it has become popular to refer to a non–redundant array of disk drives as a RAID–0 array.

Disk Drive Striping

Fundamental to RAID is ”striping”, a method of combining multiple drives into one logical storage unit. Striping involves partitioning each drive’s storage space into stripes which may be as small as one sector (512 bytes) or as large as several megabytes, Figure 41. These stripes are then interleaved round–robin, so that the combined space is composed alternately of stripes from each drive. In effect, the storage space of the drives is shuffled like a deck of cards. The type of operating environment determines whether large or small stripes should be used.

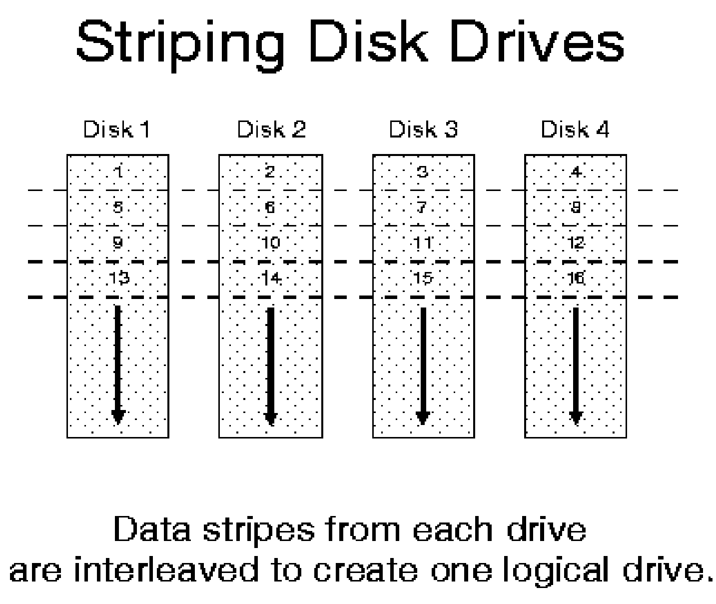

Figure 41. Striping Disk Drives

Most operating systems today support concurrent disk I/O operations across multiple drives. However, in order to maximize throughput for the disk subsystem, the I/O load must be balanced across all the drives so that each drive can be kept busy as much as possible. In a multiple drive system without striping, the disk I/O load is never perfectly balanced. Some drives will contain data files which are frequently accessed and some drives will only rarely be accessed. By striping the drives in the array with stripes large enough so that each record falls entirely within one stripe, the records will be evenly distributed across all drives and the I/O load will be balanced. All drives in the array will thus be kept busy during heavy load situations. This allows all drives to concurrently work on different I/O operations, and thus maximizes the number of simultaneous I/O operations which can be performed by the array.

Definition of RAID Levels

RAID–0 is typically defined as a non–redundant group of striped disk drives without parity. RAID–0 arrays may be configured with large stripes for multi–user environments or small stripes for single–user systems which access long sequential records. If one drive in a RAID–0 array fails, the entire array fails. However, RAID–0 arrays deliver the best performance and data storage efficiency of any array type, Figure 42.

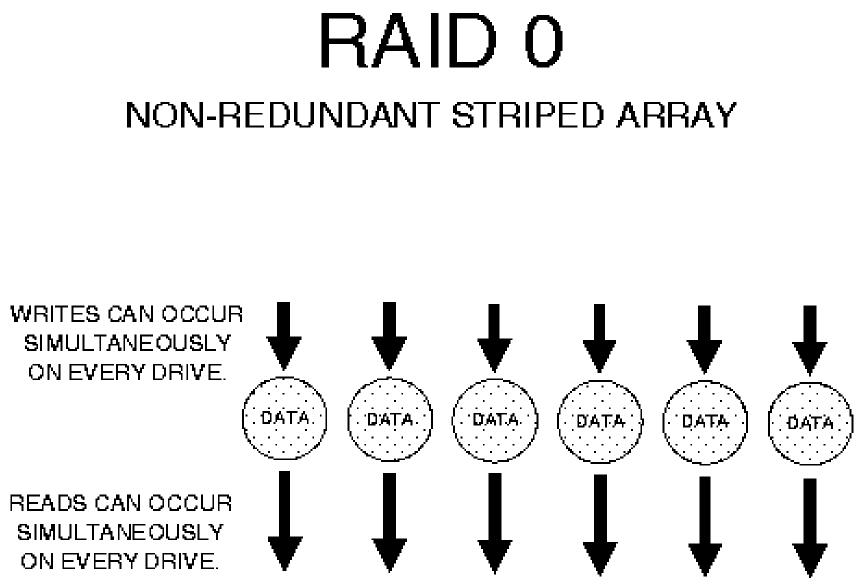

Figure 42. Non–redundant Striped Array

RAID–1, also known as disk mirroring, is simply a pair of disk drives which store duplicate data but appear to the computer as a single drive. Although striping is not used within a single mirrored drive pair, multiple RAID–1 arrays may be striped together to create a single large array consisting of pairs of mirrored drives. All writes must go to both drives of a mirrored pair so that the information on the drives is kept identical. Each individual drive, however, can perform simultaneous read operations. Mirroring thus doubles the read performance of a single un-mirrored drive and leaves the write performance unchanged. RAID–1 delivers the best performance of any redundant array type. In addition, there is less performance degradation during drive failure than in RAID–5 arrays, Figure 43.

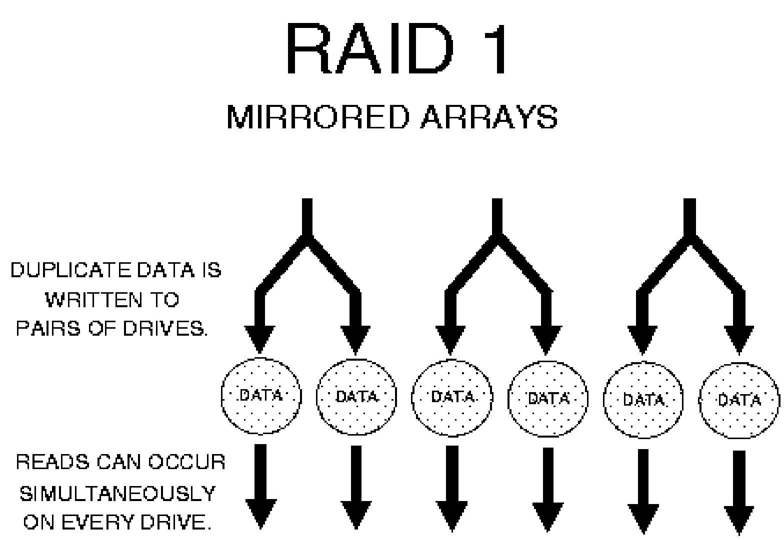

Figure 43. Mirrored Arrays

RAID–2 arrays sector–stripe data across groups of drives, with some drives relegated to storing ECC information. Since all disk drives today embed ECC information within each sector, RAID–2 offers no significant advantages over RAID–3 architecture and is not supported by SCSI RAID adapter, Figure 44.

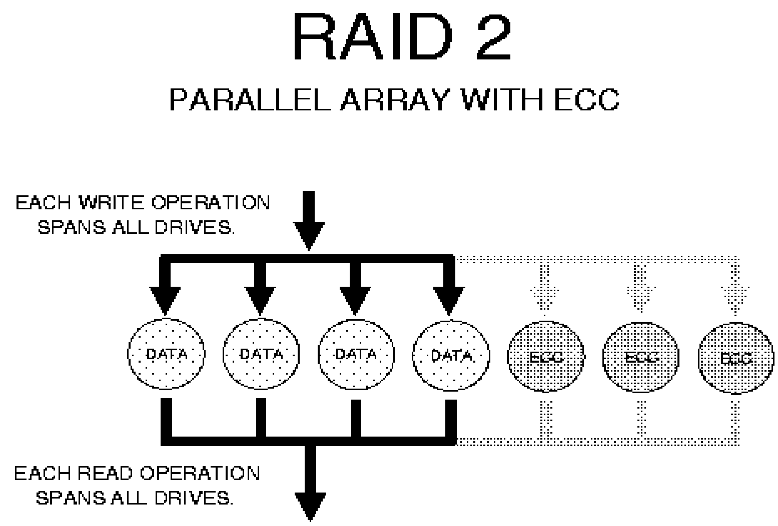

Figure 44. Parallel Array with ECC

RAID–3, as with RAID–2, sector–stripes data across groups of drives, but one drive in the group is dedicated to storing parity information. RAID–3 relies on the embedded ECC in each sector for error detection. In the case of drive failure, data recovery is accomplished by calculating the exclusive OR (XOR) of the information recorded on the remaining drives. Records typically span all drives, thereby optimizing disk transfer rate. Since each I/O accesses every drive in the array, RAID–3 arrays perform only one I/O at a time and thus deliver best performance in single–user, single–tasking environments with long records. Synchronized–spindle drives are required for RAID–3 arrays in order to avoid performance degradation with short records. Since RAID–5 arrays with small stripes can yield similar performance to RAID–3 arrays, RAID–3 is not supported by SCSI RAID adapter, Figure 45.

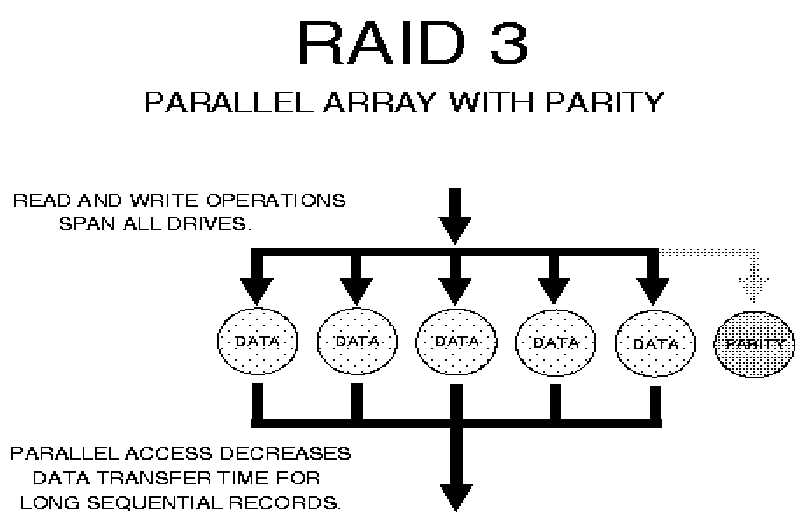

Figure 45. Parallel Array with Parity

RAID–4 is identical to RAID–3 except that large stripes are used, so that records can be read from any individual drive in the array (except the parity drive), allowing read operations to be overlapped. However, since all write operations must update the parity drive, they cannot be overlapped. This architecture offers no significant advantages over RAID–5 and is not supported by SCSI RAID adapter, Figure 46.

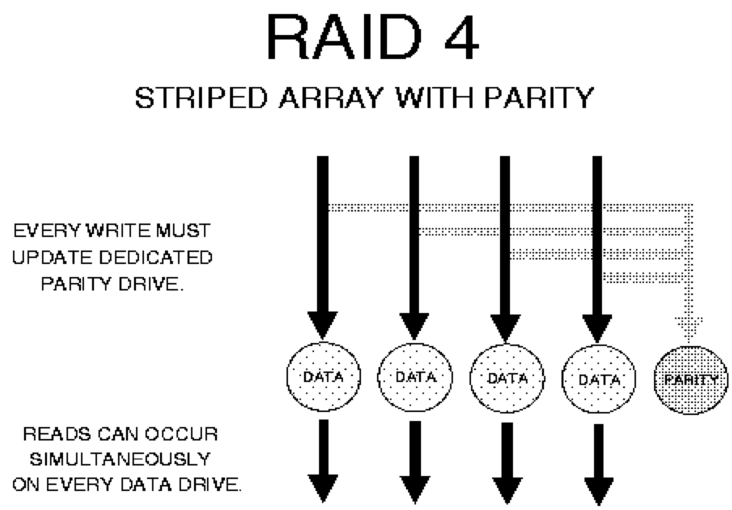

Figure 46. Striped Array with Parity

RAID–5, sometimes called a Rotating Parity Array, avoids the write bottleneck caused by the single dedicated parity drive of RAID–4. Like RAID–4, large stripes are used so that multiple I/O operations can be overlapped. However, each drive takes turns storing parity information for a different series of stripes. Since there is no dedicated parity drive, all drives contain data and read operations can be overlapped on every drive in the array. Write operations will typically access one data drive and one parity drive. Since different records store their parity on different drives, write operations can usually be overlapped, Figure 47.

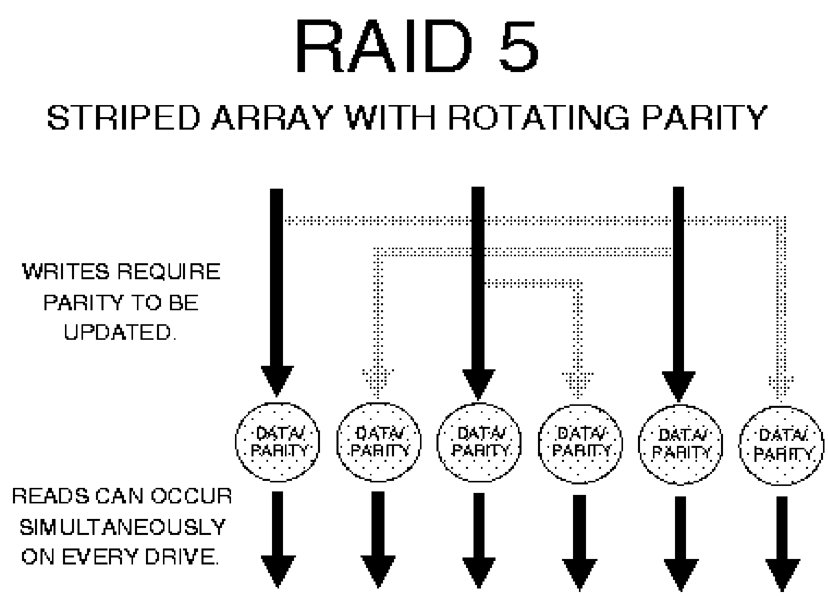

Figure 47. Striped Array with Rotating Parity

RAID Summary

- **RAID–0** is the fastest and most efficient array type but offers no fault–tolerance.
- **RAID–1** is the array of choice for performance–critical, fault–tolerant environments. In addition, RAID–1 is the only choice for fault–tolerance if no more than two drives are desired.
- **RAID–2** is seldom used today since ECC is embedded in all hard drives. RAID–2 is not supported by SCSI RAID adapter.
- **RAID–3** can be used in single–user environments which access long sequential records to speed up data transfer and provide fault–tolerance. However, RAID–3 does not allow multiple I/O operations to be overlapped and requires synchronized–spindle drives in order to avoid performance degradation with short records. Since RAID–5 with short stripes offers similar performance, RAID–3 is not supported by SCSI RAID adapter.
- **RAID–4** offers no advantages over RAID–5 and does not support multiple simultaneous write operations. RAID–4 is not supported by SCSI RAID adapter.
- **RAID–5** combines efficient, fault–tolerant data storage with good performance characteristics. However, write performance and performance during drive failure are slower than with RAID–1. Rebuild operations also require more time than with RAID–1. At least three drives are required for RAID–5 arrays.

Creating Data Redundancy

RAID–5 offers improved storage efficiency over RAID–1 since parity information is stored, rather than a complete redundant copy of all data. The result is that three or more drives can be combined into a RAID–5 array, with the effective storage capacity of only one drive sacrificed to store the parity information. Therefore, RAID–5 arrays provide greater storage efficiency than RAID–1 arrays. However, this comes at the cost of a corresponding loss in performance.

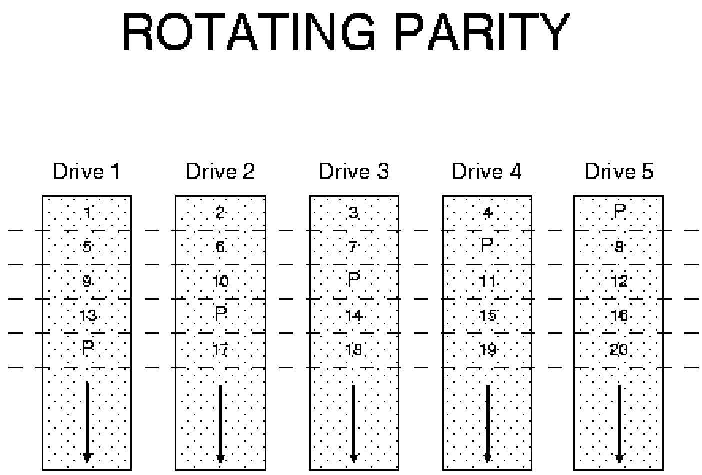

Figure 48. Data Redundancy – Rotating Parity

When data is written to a RAID–5 array, the parity information must be updated. There are two ways to accomplish this. The first way is straightforward but very slow. The parity information is the XOR of the data on every drive in the array. Therefore, whenever one drive's data is changed, the other drives in the array which hold data are read and XORed to create the new parity. This requires accessing every drive in the array for each write operation.

The second method of updating parity, which is usually more efficient, is to find out which data bits were changed by the write operation and then change the corresponding parity bits. This is accomplished by first reading the old data to be overwritten. This data is then XORed with the new data which is to be written. The result is a bit mask which has a one in the position of every bit which has changed. This bit mask is then XORed with the old parity information which is read from the parity drive. This results in the corresponding bits being changed in the parity information. The new updated parity is then written back to the parity drive. This results in only two reads, two writes and two XOR operations, rather than a read or write and XOR for every drive in the array.

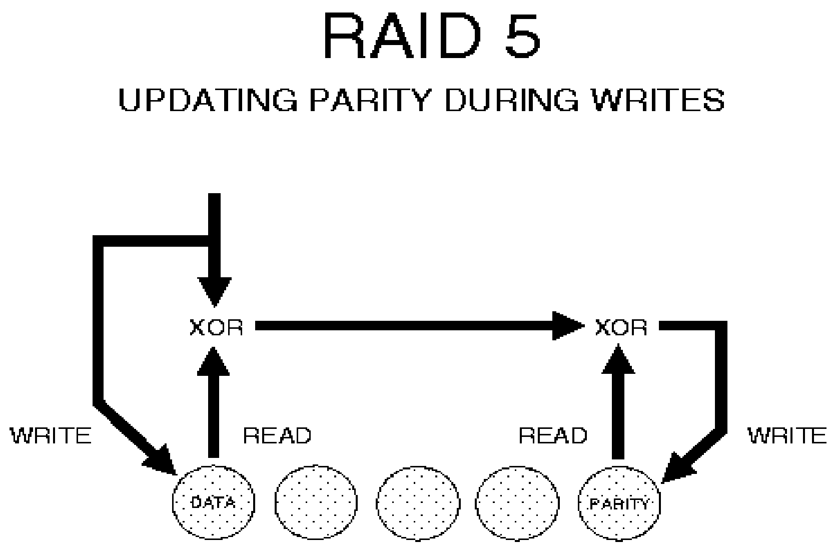

Figure 49. Updating Parity During Writes

The cost of storing parity, rather than redundant data as in RAID–1, is the extra time taken during write operations to regenerate the parity information. This additional time results in a degradation of write performance for RAID–5 arrays over RAID–1.

Degraded Mode

When a drive fails in a RAID–0 array, the entire array fails. In a RAID–1 array, a failed drive has little impact on performance since data can be read from the duplicate mirrored drive. However, in a RAID–5 array, the data must be synthesized by reading and XORing the corresponding data stripes from the remaining drives in the array. This process is relatively slow and is thus called ”degraded mode”. The more drives contained in a RAID–5 array, the slower degraded mode operation becomes.

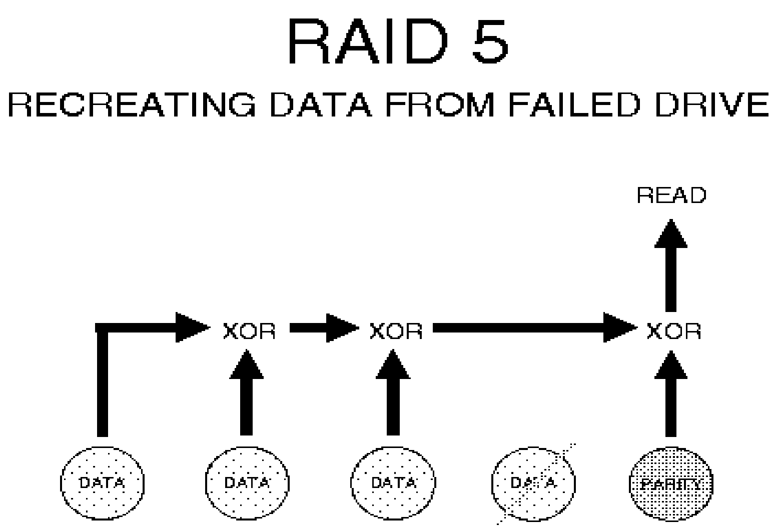

Figure 50. Recreating Data From Failed Drive

Rebuilding a Failed Hard Drive

When a failed drive is replaced in a RAID–1 or RAID–5 array, or if a Hot Spare is present, the SCSI RAID adapter will rebuild the lost data onto the new drive or Hot Spare. This rebuild operation occurs on–line while other normal reads and writes are being processed by the array. In the case of RAID–1 arrays, the rebuild will occur relatively quickly since it simply involves copying data from the duplicate mirrored drive to the replacement drive. In the case of RAID–5 arrays, the new data to be written to the replacement drive must be synthesized by reading and XORing the corresponding data stripes from the remaining drives in the array. RAID–5 arrays containing more drives will require more time for rebuilds.

SCSI RAID Adapter Architecture

SCSI RAID adapters can process commands from the OS using two different protocols:

- **ASPI Protocol** provides compatibility with DOS, OS/2 and NetWare tape and CD–ROM application software. It is not used on this machine.
- **EATA Protocol** utilizes the full performance capabilities and functionality of the PCI, EISA and ISA busses.

EATA Protocol

EATA protocol was designed to utilize the full performance capabilities and functionality of the PCI, EISA and ISA busses. It supports PCI, EISA and ISA bus mastering, overlapped I/O processing, command queuing and scatter/gather memory access.

Command Overlapping

Peripheral controllers which cannot perform command overlapping can only execute one command from the computer at a time. EATA protocol allows SCSI RAID adapter to receive another command before the first command has finished executing. By allowing multiple outstanding commands, SCSI RAID adapters can optimize performance by overlapping commands for multiple SCSI devices, thereby keeping all devices busy as much as possible.

Command Queuing

EATA protocol allows the adapter to queue multiple commands from the computer and execute them out of order. Once a command has completed, SCSI RAID adapter transfers a Status Packet into the computer's memory, containing status information plus information indicating which command was just completed. This allows commands to be executed in a different order than they were issued by the computer. SCSI RAID adapter can thus queue commands and rearrange them so that they may be executed in a more efficient order for the particular peripheral configuration.

In addition, SCSI RAID adapter supports Tagged Command Queuing on the SCSI bus. This feature allows the adapter to transfer up to 64 SCSI commands to a single device if that device will accept queued commands. SCSI Command Queuing may be disabled through the Configuration Utility or from the "Configure Host Bus Adapter" window accessible from the adapter's Information Window in Storage Manager. However, this should not be necessary as the adapter will negotiate with each SCSI device to determine if it is capable of utilizing the command queuing feature.

Scatter/Gather

When using virtual memory addressing schemes, system memory may appear contiguous to the user but is actually fragmented into many widely scattered physical address locations. Because of this, it is often necessary when accessing a large amount of contiguous data from a peripheral device, to break up the transfer into many different locations in system memory. SCSI RAID adapters support Scatter/Gather, a method of providing multiple memory addresses for data transfer in one command packet. This greatly increases performance in environments such as UNIX, Novell NetWare, Windows and OS/2.

Byte/Word Alignment

Data to be read or written by the SCSI RAID adapters need not be aligned to word or long–word boundaries in system memory. If the data starts on an odd memory address, the necessary bus control signals will be generated by the SCSI RAID adapter to access 8–bit, 16–bit, 24–bit and 32–bit data as required. Data blocks of any byte length will be correctly handled by SCSI RAID adapter.

Command Processing

Using EATA protocol, SCSI RAID adapters process a command from the computer as follows:

1. The CPU creates a Command Packet in system RAM and writes the address of this packet into a register in the adapter. The CPU can then perform other tasks. All remaining operations will be performed by the adapter. The Command Packet contains parameters, specifying the type of operation, the amount of data and the location on the peripheral device. The Command Packet may also contain pointers to a Scatter/Gather Table and a Status Block. The Scatter/Gather Table describes a list of data segments to be read or written to sequential locations on the peripheral device. These data segments may be scattered in many different areas of system RAM. The Status Packet is a place in system RAM for SCSI RAID adapter to save status information concerning the successful or unsuccessful completion of the command.
2. RAID becomes the bus master (takes control of the system bus) and DMAs the Command Packet and Scatter/Gather Table into its local memory.
3. When SCSI RAID adapter is ready to transfer the data, it again becomes the bus master and DMAs the data to the areas specified by the Scatter/Gather Table.
4. When the command is complete, SCSI RAID adapter becomes the bus master and DMAs the Status Packet into the proper RAM location.
5. RAID interrupts the CPU to tell the computer that the command has completed.

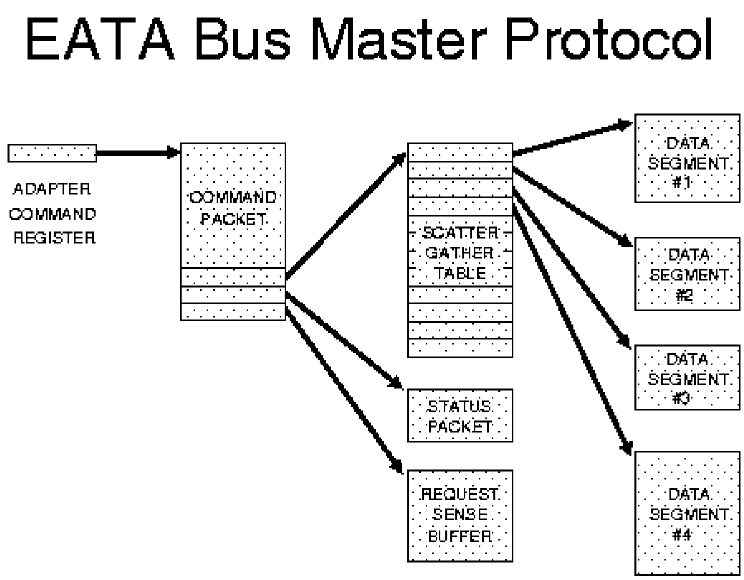

Figure 51. EATA Bus Master Protocol

PCI Data Transfer

Through bus mastering, SCSI RAID adapter PM3334UW adapters can transfer data at up to 132MB/s across the PCI bus. In addition, several other data transfer protocols are supported as described in the following sections:

PCI Bus Mastering

PCI achieves 132MB/s bandwidth by using a 33MHz high speed clock rate as opposed to the 8.33MHz clock used by ISA and EISA busses. In order to maintain data integrity at that high clock rate, the PCI bus is limited to four card slots and six devices integral to the motherboard. However, motherboard designs can include multiple PCI busses interconnected through bridges. PCI allows two modes of data transfer: Bus mastering and PIO. Third party DMA is not supported by PCI.

Each PCI device contains configuration registers which allow the system to be configured for maximized throughput. Parameters such as Latency Timer, Cache Line Size, Minimum Bus Grant (MIN_GNT), Maximum Latency (MAX_LAT), as well as configurable I/O Address and IRQs can be read or written by the PCI system BIOS. SCSI RAID adapters support the full range of required and optional features as described by the PCI Revision 2.1 specification.

SCSI RAID adapters are fully compliant with the PCI Revision 2.1 specification for bus drivers and receivers. These transceivers are unique in their ability to vary the drive current in order to minimize signal reflections.

One option supported by SCSI RAID adapter PCI adapters is Dual Address Cycles. This feature allows the adapter to address the full 64–bit memory address space supported by the PCI specification.

All mandatory and optional memory read/write commands are supported by SCSI RAID adapters. Mandatory commands include Memory Read and Memory Write which are used by all PCI bus mastering devices. Optional commands include Memory Read Line, Memory Read Multiple and Memory Write and Invalidate. These commands maximize performance in PCI systems with intelligent bridges which perform CPU cache line optimization.

SCSI Interface

The Small Computer Systems Interface is an ANSI standard parallel interface designed to communicate with intelligent peripheral devices. SCSI RAID adapters conform to the SCSI specification, defined in ANSI document number X3.131–1986. In addition, SCSI RAID adapters conform to the ANSI SCSI–2 and SCSI–3 specifications and CCS (Common Command Set) documents.

SCSI Device IDs

The SCSI ID on SCSI RAID adapters is set by default to ID 7 and may be changed through Storage Manager or the Configuration Utility to any ID from 0 to 7 (on B4–4 adapters only). It cannot be changed on a B4–C adapter.

In 8–bit SCSI systems, the device with ID 7 has the highest priority and ID 0 has the lowest priority when arbitrating for use of the SCSI bus. In 16–bit Wide SCSI systems, IDs are prioritized as follows:

Highest: 7, 6, 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, 8 :Lowest

Since IDs 8–15 will always have a lower priority during arbitration than 8–bit devices, this allows 8–bit devices which do not recognize IDs 8–15 to co–exist on a SCSI bus with Wide SCSI devices which may be assigned those IDs. When controlling 8–bit devices, however, a Wide SCSI adapter may not be configured at IDs 8–15 since 8–bit devices would then not recognize the adapter.

LUNs

Any SCSI device may contain up to eight sub–devices or Logical Units, assigned Logical Unit Numbers (LUN) 0–7. However, most SCSI devices other than bridge controllers consist of only one LUN and will be accessed as LUN 0 by SCSI RAID adapter.

Commands and Messages

All SCSI commands are supported by SCSI RAID adapter. Using EATA protocol, SCSI RAID adapters will pass through any SCSI command from the computer to the SCSI peripheral device.

RAID SCSI Adapter supports one byte and extended messages, such as Abort, SCSI Reset, Initiator Detected Error, Parity Error, Wide Negotiation and Synchronous Negotiation. SCSI message protocol is normally handled automatically by SCSI RAID adapters without requiring intervention by the host computer. This allows much faster SCSI protocol handling and puts less load on the host CPU, freeing it for other tasks.

Disconnect/Reconnect

Full support for disconnection and re-connection by SCSI peripheral devices is provided. All Disconnect/Reconnect message protocol is handled automatically by SCSI RAID adapters, without host intervention.

Disconnection allows a SCSI device and adapter to temporarily release control of the SCSI bus during periods of non–activity while a SCSI command is in process, so that another SCSI device can share the bus. However, disconnection and re-connection protocol adds unnecessary overhead when only one device is utilizing the SCSI bus. SCSI RAID adapters optimize performance by disabling disconnection during periods of single device activity.

Transfer Padding and Residue Reporting

In unusual cases, a SCSI peripheral device may read or write more data than requested by the computer. If this occurs during a peripheral write operation, SCSI RAID adapters will automatically transfer ”pad–data” by adding additional bytes of FFH to the end of the data to be written to the SCSI device. If this occurs during a peripheral read operation, the adapter will automatically strip off additional bytes of data read from the SCSI device, thereby preventing host memory from being over–written. It is also possible in some cases, for a SCSI peripheral device to read or write less data than requested by the computer. If this occurs, SCSI RAID adapters will report the number of additional bytes requested by the computer which were not transferred by the SCSI peripheral, as the ”residue” for the transfer.

Single–ended/Differential SCSI

The SCSI electrical specification has two mutually exclusive transceiver specifications:

1. Differential RS–485 transceivers that allow up to 20MHz data transfer at a maximum cable length of 25 meters.
2. Single–Ended TTL transceivers which allow:

- Asynchronous data transfer at a maximum cable length of 6 meters. (No maximum transfer rate is given but typical rates range from 1.5MHz to 3MHz.)
- Slow synchronous data transfer up to 5MHz at a maximum cable length of 6 meters.
- Fast 10MHz synchronous data transfer at a maximum cable length of 3 meters.
- Fast–20 20MHz synchronous data transfer at a maximum cable length of 3 meters with up to four devices and 1.5 meters with up to eight devices.

Most SCSI RAID adapter products support Single–Ended SCSI only.

Wide SCSI

The SCSI–2 specification defines 8–bit, 16–bit and 32–bit SCSI bus widths. All SCSI RAID adapters support both 8–bit SCSI and 16–bit Wide SCSI.

SCSI Transfer Rate

SCSI RAID adapters are capable of handling mixed SCSI devices with different transfer speeds, as well as both synchronous and asynchronous protocol, on the same SCSI Bus. Using SCSI message protocol, SCSI RAID adapters automatically negotiate with each SCSI device during the power–on sequence to determine the fastest transfer speed and method possible for that device.

SCSI RAID adapters can optionally detect the presence of an external SCSI cable and automatically lower the maximum SCSI bus transfer rate when an external cable is present. (Both the external and internal cable data rates will be lowered.) Since some external SCSI cabling methods are not capable of handling maximum transfer rates without causing data errors this safety feature protects the user against data integrity problems. This external cable detection feature may be enabled and the default maximum SCSI bus transfer rate may be modified by using the Configuration Utility or from the ”Configure Host Bus Adapter” window accessible from the adapter’s information window in Storage Manager.

The original SCSI specification allowed synchronous transfer rates of up to 5MHz. The SCSI–2 specification increased the maximum synchronous rate to 10MHz. Devices which support the faster 10MHz rate are called Fast SCSI devices. The SCSI–3 specification defines faster rates up to 20MHz, called Fast–20 SCSI.

All SCSI RAID adapters support 10MHz Fast SCSI and slower devices.

Cabling Single–ended SCSI

When using 10MHz and higher SCSI data rates, a common problem is poor data reliability, especially as the number of devices and length of the cable grow. Most data reliability problems stem from signal reflections and noise which are read by SCSI receivers as incorrect data or false SCSI bus phases. The SCSI cable is a transmission line that has a characteristic impedance, the value of which depends upon the type and configuration of cable used. Discontinuities in this impedance can cause signal reflections. These impedance variations can be the result of extra capacitance internal to SCSI devices, connectors, improper terminators, mixing of different cable types, cable stubs, etc. At 10MHz Fast and 20MHz Fast–20 SCSI rates, these reflections are much more problematic than at the slower 5MHz SCSI rates. Additional noise picked up from external devices, as well as crosstalk from other signals on the SCSI cable, can add to SCSI bus signal integrity problems.

A well–configured Single–Ended SCSI bus can reliably transfer data at up to 20MHz without a problem. However, good engineering practices should be followed in order to guarantee success:

- **Use as short a cable length as possible.** For 10MHz data transfers the total SCSI bus length should not exceed 3 meters. For 20MHz data transfers, total SCSI bus length should not exceed 3 meters with up to 4 devices on the bus, or 1.5 meters with between 5 and 8 devices on the bus.

  Note: When an external SCSI storage cabinet is used, the cabinet's internal SCSI backplane will add additional length to the bus.

- **Avoid stub clustering.** Avoid spacing SCSI devices on the cable closer than 0.3 meters apart. When devices are clustered closely together on the SCSI cable, their capacitances add together to create an impedance discontinuity and thus reflections. SCSI devices should be spaced as evenly as possible.
- **Cable stub length should not exceed 0.1 meters.** Some SCSI devices may create stubs internal to the device which exceed this value, resulting in excessive capacitive loading and signal reflections. This parameter is under the control of the SCSI device manufacturer. The SCSI cabling itself should include no stubs.
- **Watch out for capacitance.** As devices are added to a SCSI bus, capacitance is introduced to each signal from the connectors, receivers, and PC board traces. The SCSI–3 working specification limits this capacitance to 25pF since added capacitance lowers the impedance of the cable near the device and adds delay. Look for input filters that may be attached to the SCSI front–end of the device's printed circuit board. These filters add capacitance.
- **Avoid unnecessary connector converters.** They cause impedance discontinuities and signal reflections.
- **Route cable with care.** Avoid rolling the cable up on itself, running the cable alongside of metal for long lengths or routing the cable past noise generators (i.e. power supplies). Placing the cable near ground planes created by grounded metal cabinetry reduces its impedance. The SCSI–3 working committee suggests that in order to minimize discontinuities due to local impedance variation, a flat cable should be spaced at least 1.27 mm (0.050 in.) from other cables, any other conductor, or the cable itself when the cable is folded.
- **Follow the SCSI specification for cable impedance.** Cable impedance for 10MHz Fast SCSI systems should be limited to 84+/–12W. Cable impedance for 20MHz Fast–20 SCSI systems should be limited to 90+/–12W for REQ and ACK signals and 90+/–10W for all other signals.
- **Avoid mixing cable types.** Select either flat or round, shielded or non–shielded. Typically mixing cables mixes impedances. Cable impedance mis–match is a common problem resulting in signal reflections. Internal cables are typically flat ribbon cables, while external cables should be shielded. Where they offer easier routing, size advantages, and better air flow, round cables can be used internally as well. This in fact, may be desirable if it allows for better impedance matching to the external cable.
- **When round cable is used, select a cable that uses a wise placement of signals within the cable.** Ribbon cable shows fairly good crosstalk rejection characteristics due to the GND–Signal–GND layout. However, more care needs to be taken to insure adequate performance when round cable is employed. With standard 25–pair round cable, pairs are arranged in three layers. The closer a pair is to the cable center, the

higher the impedance. Using centrally located high impedance pairs for speed–critical signals such as REQ and ACK is desired. By locating data pairs in the outermost layer of the cable, crosstalk between REQ, ACK, and the data lines is minimized. The middle layer might contain status lines such as C/D, I/O, MSG, ATN, etc. Make sure that the lowest impedance wire in the cable is used for TERMPWR to minimize transmission line effects on this voltage supply line. Some SCSI cables have a low–impedance conductor for this purpose. Typically, a larger wire gauge and an insulator with a higher dielectric constant are used on this conductor.

SCSI Termination

SCSI termination for SCSI RAID adapters is enabled through Storage Manager or the Configuration Utility. Termination may be set to ”Off” or ”On”, causing termination for all SCSI signals to be enabled or disabled respectively. An additional setting of ”High Only” enables termination for signals which are present on Wide SCSI devices but not on 8–bit SCSI devices as shown below:

SCSI RAID adapters contain active SCSI terminators (called Alternative 2 in the SCSI–2 specification) for maximum reliability with Fast SCSI devices. Active termination uses one 110W resistor per signal, pulled up to a locally supplied voltage which is regulated at 2.85V. Features of active termination include:

- Better immunity to fluctuations in TERMPWR.
- Closer match to the characteristic impedance of the cable minimizes reflections.
- Increased noise margin.

Passive termination (Alternative–1) is still utilized on many SCSI devices today. Passive termination employs one 220W pull–up and one 330W pull–down resistor per signal. For better signal integrity, especially in Fast SCSI environments, SCSI devices with active termination should be used whenever possible. However, active and passive termination can be mixed on one SCSI bus and results in better reliability than using passive termination alone. Fast–20 does not allow passive termination in the Single–Ended configuration.

TERMPWR

SCSI terminators need power in order to operate. SCSI RAID adapters always supply power for their own on–board SCSI terminators and are configured by default to supply termination power for other devices via the TERMPWR line on the SCSI cable. Most SCSI devices allow the manner in which TERMPWR is supplied to their on–board SCSI terminators to be selected through jumpers on the device. It is best that devices supplying SCSI bus termination be configured to supply their own isolated TERMPWR. This prevents loss of receiver noise margin due to TERMPWR DC voltage drop across the cable. (TERMPWR, however, should remain enabled if there is a possibility of power–loss on an external SCSI device which supplies its own TERMPWR.)

If all devices which are supplying termination for the SCSI cable are also capable of supplying their own termination power then TERMPWR supplied to the SCSI cable can be disabled on SCSI RAID adapters through the Configuration Utility or through Storage Manager.

SCSI RAID adapters supply termination power onto the SCSI cable through a thermistor and a Schottky diode. A resetable fuse, or polyswitch, on the SCSI RAID adapter protects it from a short to ground on the TERMPWR line. (Some SCSI devices use a fuse for protection, necessitating the replacement of the part if a short occurs or if the SCSI cable is inserted incorrectly.)

The ANSI XT3T9.3 SCSI–3 Parallel Interface working group recommends that all TERMPWR lines be decoupled at each terminator to minimize TERMPWR glitch coupling. The minimum recommended values are a 2.2mF solid tantalum capacitor along with .01mF ceramic capacitor in parallel to help with high frequency, low voltage noise. These capacitors, when utilized, will supply the high frequency, low impedance path to ground necessary to filter out glitches. SCSI RAID adapters provide power to the on–board termination directly from a highly decoupled power plane, thus insuring minimal TERMPWR glitch coupling.

SCSI Bus ECC Protection

Hard drives that support 528–byte sector formats have the ability to store ECC (Error Correcting Code) generated by SCSI RAID adapters. From the time that data first enters the SCSI RAID adapter from the computer, 16 bytes of ECC are appended to each 512–byte block of data. This ECC protects the data against most errors which occur during transmission through the adapter, data buffers, cache, SCSI cable, storage cabinet backplane, and drive electronics. This extra level of protection is particularly important in systems which employ hot–pluggable devices since hot–plugging causes electrical glitches on the SCSI bus which could otherwise cause data corruption.

When data is written to a SCSI RAID adapter hard drive, the 16–byte ECC plus 512 bytes of data are written on the drive as a 528–byte sector. (Additional ECC is added to the sector by the drive electronics in order to correct errors caused by flaws in the media.) Later, when the 528–byte sector is read from the drive, the 16–byte ECC is separated from the data by the SCSI RAID adapter before it is transferred to the computer. Corrupted data is automatically corrected by the adapter.

If supported by the hard drive, Storage Manager may be used to format the drive with 528–byte sectors for SCSI bus ECC support. ECC SIMMs are not supported on Estrella yet.

Appendix A. SCSI RAID B4–4 Adapter – Troubleshooting

This Troubleshooting Guide provides answers to many commonly asked questions. If a situation occurs which is not covered here, or if the recommendations here do not correct the problem, contact Technical Support.

SCSI RAID B4–4 Adapter – Troubleshooting

Note: LEDs and other components are shown in Figure 1, on page 2-2.

Problem: When the SCSI BIOS displays the SCSI devices at system bootup, one device appears at all seven SCSI IDs.

Solution: One of the SCSI devices has been set to the same ID as the adapter. The SCSI RAID adapter is set by default to SCSI ID 7. Assure that all SCSI devices have a unique ID.

Problem: The adapter does not respond during bootup and LEDs 1 and 2 or 1, 2 and 10 remain lit.

Solution: Either the SCSI bus is improperly terminated or the SCSI cable is on backwards. See **Configuring SCSI Termination**, on page 2-5.

If auto termination is enabled and the controller is in the middle of a Wide SCSI bus, turn off termination.

Problem: LEDs 4 and 8 flash once per second at power–up and the adapter does not respond.

Solution: This pattern indicates that the adapter has detected a problem with the SCSI cable termination. Check that the cable termination is configured correctly. See **Configuring SCSI Termination**, on page 2-5**.**

Problem: LEDs 1 and 3, or LEDs 1, 3 and 6 remain lit at power–up and the SCSI BIOS ROM does not display any attached devices.

Solution: This pattern indicates that the adapter has detected a severe problem with the SCSI cable termination. Check that the cable termination is configured correctly. See **Configuring SCSI Termination**, on page 2-5.

Problem: LEDs 1, 2 and 8 flash once per second at power–up and the adapter does not respond.

Solution: This pattern indicates that there is no SIMM in socket #1 of the SCSI RAID adapter. Socket #1 must always contain the largest capacity SIMM.

Problem: LEDs 2 and 8 flash once per second at power–up and the adapter does not respond.

Solution: This pattern indicates that SIMM Socket #1 of the SCSI RAID adapter does not contain the largest SIMM. Rearrange the SIMMs so that the largest capacity SIMM is in Socket #1.

Problem: Pressing <Ctrl–D> to access the Configuration Utility either does not work, or the information displayed is garbled.

Solution: Clear the adapter's NVRAM. See **NVRAM**, on page 2-12. Then re–run the Configuration Utility and reconfigure the adapter.

Problem: LEDs 1, 2, 3, 4 and 8 flash once per second at power–up and the adapter does not respond.

Solution: This pattern indicates that a combination of ECC and parity SIMMs have been installed on the SCSI RAID adapter. Remove either the parity or the ECC SIMMs.

Problem: Various LEDs in the 1–4 range flash once per second and the adapter does not respond.

Solution: This pattern indicates an internal Microprocessor Trap in the adapter. Remove all SCSI devices, cables and option modules and retry. If the Trap Error disappears, reconnect everything, one device at a time, until the faulty device, cable or module is isolated. Microprocessor Trap can also occur because of a firmware EPROM error during normal operation. Try reseating the EPROMs in their sockets to assure good contact.

Problem: Some other combination of LEDs not listed above flash once per second at power–up and the adapter does not respond.

Solution: This indicates a firmware trap has occurred. To determine the cause of the trap, run Storage Manager (from floppy if necessary) while the LEDs are still flashing. A description of the problem and possible solutions will be displayed upon entering the program.

Problem: The SCSI RAID adapter cannot access a hard drive.

Solution: This typically happens when a SCSI RAID adapter hard drive which has been formatted to support SCSI bus ECC protection is attached to a SCSI RAID adapter which contains one or more non–ECC SIMMs.

Problem: After flash updating the SCSI RAID adapter firmware or SCSI BIOS and rebooting, the adapter does not respond.

Solution: The firmware update may have been unsuccessful. See **Upgrading Flash Firmware**, on page 2-12, to disable the new firmware and re–attempt the update.

Problem: After flash updating the SCSI RAID adapter firmware or SCSI BIOS and rebooting, LEDs 1 through 8 walk back–and–forth, a single LED at a time.

Solution: This pattern indicates that the adapter startup code detected a firmware checksum error. See **Upgrading Flash Firmware**, on page 2-12, to disable the new firmware and re–attempt the update.

Appendix B. SCSI RAID B4–C Adapter – Troubleshooting

This Troubleshooting Guide provides answers to many commonly asked questions. If a situation occurs which is not covered here, or if the recommendations here do not correct the problem, contact Technical Support.

SCSI RAID B4–C Adapter – Troubleshooting

Note: LEDs and other components are shown in Figure 2, on page 2-3.

Problem: The controller does not respond and the IRQ LED (and possibly other LEDs) remains lit. See Appendix A for the location of the LEDs on your controller.

Solution: The IRQ LED indicates that the controller IRQ assignment is pending. This usually indicates an IRQ conflict with another card. Ensure that each card is set to a unique IRQ.

Problem: The controller does not respond and one of the following patterns of LEDs flash once per second at power–up:

PATTERN	MEANING
7, 6, 5, 2, 1	None
7, 6, 5, 3, 1	High
7, 6, 5, 3, 2	Mismatch
7, 6, 5, 3, 2, 1	Invalid

Solution: These patterns indicate that there is a problem with the memory modules on the controller.

None: Either no memory modules were detected on the controller, or there is no module in socket 1.

Socket 1 must always have a 16MB or 64MB 60ns EDO memory module installed.

High: Too much memory has been detected on a Decade or Century controller. Remove memory so that the total is less than or equal to 64MB.

Mismatch: Memory modules of mixed sizes have been detected or a SIMM slot was skipped. All installed modules must be the same size and must be filled sequentially from socket 1 to socket 4.

Invalid: A memory module of a size other than 16MB or 64MB has been detected. Use only 16MB or 64MB memory modules.

NOTE: Do not install non–EDO SIMMs or DIMMs. This will cause data corruption.

Problem: The controller does not respond and various LEDs in the 1 4 range flash once per second.

Solution: This pattern indicates an internal microprocessor trap occurred in the controller. Remove all attached devices, cables and option modules and retry. If the trap error disappears, reconnect the cables and devices, one device at a time, until the faulty device, cable or module is isolated. If the error persists, contact Technical Support.

Problem: Although the SCSI devices can be accessed by the SmartRAID V controller, the fault LEDs on the devices in a RAIDstation storage cabinet do not flash during boot–up and the SmartRAID V controller does not detect drive swaps or cabinet failures.

Solution: These symptoms indicate that the RAIDstation storage cabinet status signals are not being properly received by the SmartRAID V controller.

For DEC Fault Bus subsystems: this can be caused by another SCSI device or non–DPT cabinet connected to the external SCSI cable along with the RAIDstation cabinet. Other devices will typically ground these signals.

For SAF–TE or SES: this can result from a failed enclosure monitoring module in the subsystem cabinet.

Problem: After updating the SmartRAID V controller firmware or BIOS and rebooting, LEDs 1 and 5 or 2 and 5 flash once per second.

Solution: These patterns indicate that the adapter startup code detected a firmware checksum error or a flash error. Attempt the firmware update procedure again by using the procedure in the following Problem description to recover from this condition.

Problem: A flash ROM upgrade is unsuccessful, causing the controller to hang.

Solution: The new firmware can be temporarily disabled and the upgrade attempted again by following the steps below:

1. Power–off the system.

2. PM1554, PM2554 and PM2654 controllers   remove any RA4050 or SX405x expansion modules from the controller.

3. PM1554, PM2554 and PM2654 controllers   place shorting jumpers across pins 1 and 2 and pins 3 and 4 of P9 on the controller.

PM375x controllers   move the shorting jumper from pins 3 and 4 of P9 to pins 1 and 2 (from RUN to LOAD) as shown in Appendix A, " Assembly Drawings" .

NOTE: You do not need to remove an SX405x expansion module from PM375x controllers to access the jumper pins at P9.

4. Insert the DPT SMOR Boot diskette and power–up the system. This will start SMOR.

NOTE: The SMOR Boot diskette image is available from the DPT Technical Support ftp site. The download file contains the diskette image and instructions for use.

5. Use SMOR to update the firmware. You must restore all three components of the flash ROM, firmware, I2O BIOS, and SMOR.

6. Power–off the system and return the jumpers to their original positions.

7. Reattach the expansion module to the controller card and insert the card in a host system PCI slot.

8. Remove the SMOR Boot diskette from your floppy disk drive and power–up the system.

Problem: The controller's audible alarm is sounding during normal operation.

Solution: This indicates a drive has failed. Start Storage Manager or restart the host system and run SMOR to identify the failed drive. The alarm will stop when Storage Manager or SMOR finish the system scan. Replace the failed drive and start a rebuild operation for the array.

For additional information about procedures for failed drives, refer to Chapter , " Storage Manager" .

Index

A

Adapter
- Features, 1-1
- Installation, 2-9
- Installation Overview, 2-1
- SCSI RAID, 1-2

Adapter Architecture, 6-16
Adapter Cache Size, Optimization, 6-6
Adapter LEDs, 2-10, 2-11
Addressing a Logical Device, 5-16
AIX
- Adding Hard drives, 3-3
- Packaging Installation, 3-1
- Recognition Hard drives, 3-3
- System Installation, 3-1

Algorithm
- Read–ahead, 6-4
- Write–back, 6-3

Architecture, Adapter, 6-16
Array
- Diagnostics, 5-24
- Mirrored (Figure), 6-9
- Non–redundant, Striped (Figure), 6-9
- Parallel with ECC (Figure), 6-10
- Parallel with Parity (Figure), 6-10
- Recreating Data From Failed Drive (Figure), 6-14
- Striped with Parity (Figure), 6-11
- Striped with Rotating Parity (Figure), 6-11
- Updating Parity During Writes (Figure), 6-14

Array Diagnostics Window, 5-24
Array Group
- Create, 5-13
- Delete, 5-15
- Modify, 5-15

Array Group Caching, Information Windows, 5-11
Array Group Caching Configuration Window, (Figure), 5-11
Array Group Information Window, (Figure), 5-13
Array Groups, 5-12
Audible Alarm, 5-21

B

Background Task Priority, 5-26
Background Task Priority Window, (Figure), 5-26
Battery backup, 5-29
Broadcaster, 5-19
Broadcasters
- AIX, 5-19
- Install (AIX), 5-19
- Viewing Events (AIX), 5-19

Bus
- ECC Protection. *See* SCSI
- EISA, 6-16
- ISA, 6-16
- PCI, 6-16

Byte/Word Alignment. *See* EATA Protocol

C

Cache memory, Minimum size, 4-1
Cache size, 1-2
Caching
- Disk, 6-1
- Hardware, 6-1
- Software, 6-1

Command, ls –Ccdisk, 3-3
Command Overlapping. *See* EATA Protocol
Command Processing. *See* EATA Protocol
Command Queuing. *See* EATA Protocol
Commands, SCSI, 6-19
Configuration, Setting, 5-32
configuration, device, SCSI-ID, hard disk drive, 2-7, 2-8
Creating Data Redundancy, 6-13

D

Data Redundancy
- Creation, 6-13
- Rotating Parity (Figure), 6-13

Data Transfer, PCI, 6-18
Data transfer modes
- Bus mastering, 6-18
- PIO, 6-18

Definition, RAID Levels, 6-9
Degraded Mode, 6-14
Device, 5-16
Device Caching, Information Windows, 5-10
Device Caching Configuration Window, (Figure), 5-10
Device IDs, SCSI, 6-19
Diagnostics, 5-23
- Array, 5-24
- Scheduling, 5-25
- SCSI Device, 5-23

Differential SCSI. *See* SCSI
DIMMs, Installation, 2-3
Disk Caching, 6-1
Disk Drives, Striping, 6-7
Disk Writes
- With Elevator Sorting, 6-3
- Without Elevator Sorting, 6-3

drive, replacement, 4-3

E

EATA Bus Master Protocol, (Figure), 6-17
EATA Protocol
- Byte/Word Alignment, 6-16
- Command Overlapping, 6-16
- Command Processing, 6-17
- Command Queuing, 6-16
- Scatter/Gather, 6-16

ECC Protection, 6-23
Elevator Sorting, 6-3
Error Correcting Code, 5-9

Event Log Display Window, (Figure), 5-18
Event Logging, 1-3, 5-18
Event Notification, 1-3
Events, 5-18
Ext. Cable Detect, 5-7
Extend PCI REQ, 5-7

F

Failed Hard Drive, Rebuilding, 6-15
Firmware, Upgrade, 2-12
Flash Configuration, 5-31
Flash Firmware, Upgrading, 2-12
Flash Firmware Upgrade, Jumper Installation, 2-12
Flushing of OS Cache, (Figure), 6-2
Format Options Window, (Figure), 5-20

G

GUI, Storage Manager, 5-1

H

hard disk drive, SCSI-ID, 2-7, 2-8
Hard drive, Selection, 4-2
Hard Drives
    Failures, 5-21
    Formatting, 5-20
    Hot spares, 5-21
    Hot swapping, 5-21
Hard drives, 1-2
Hardware Caching, 6-1
Hardware Inventory Management, 1-3
HBA
    Caching Configuration, 5-8
    Configuration, 5-6
    Information Windows, 5-5
HBA Caching Configuration Window, (Figure), 5-8
HBA Information Window, (Figure), 5-5
hdisk objects, 3-3
Host Bus Adapter, 5-16
Host Bus Adapter windows, 5-5
Hot
    Spares, 5-21
    Swapping, 5-21
Hot spares, 1-2
How to
    Add Hard drives (AIX), 3-3
    Create an Array Group, 5-13
    Delete an Array Group, 5-15
    Format Hard Drives, 5-20
    Install the Broadcaster (AIX), 5-19
    Modify an Array Group, 5-15
    Recognizing Hard drives (AIX), 3-3

I

I/O Address, 5-7
I/O Monitor, 5-27
I/O Monitor Statistics Window, 5-27
Information Windows, 5-5
Information windows
    Array Group Caching, 5-11
    Device Caching, 5-10
    HBA, 5-5
    SCSI Device, 5-9
Installation, Storage Manager, 5-1
Installing
    Adapter, 2-9
    Adapter, Overview, 2-1
    Additional DIMMs, 2-3
    Additional SIMMs, 2-2
Interface, SCSI, 6-19
IRQ, 5-7

K

Keyboard functions, Storage Manager, 5-2

L

LED Display
    Adapter Active, 2-10
    Adapter Idle, 2-10
    Cache Status, 2-11
    Power–up, 2-10
LED Display States, (Table), 2-10
LED Function Indicators, (Figure), 2-2, 2-3
LEDs, Adapter, 2-10, 2-11
Locality of reference, 6-4
Logging, Events, 5-18
Logical Configuration View, 5-4
Logical Configuration View Screen, (Figure), 5-4
Logical Device, Addressing, 5-16
Logical Unit Number, 5-16
Logical Units. *See* LUNs
ls –Ccdisk Command, 3-3
LUNs, SCSI, 6-19

M

Maximum Record Size
    Read, 5-10
    Write, 5-11
Messages, SCSI, 6-19
Mode, Degraded, 6-14
Monitor, I/O, 5-27

N

NVRAM, 2-13

O

On–line Scheduled Diagnostics, 1-3
Optimizing Cache OS Size, (Figure), 6-6
OS Cache, (Figure), 6-1
OS Cache Size, Optimization, 6-6

P

PCI
    Bus Mastering, 6-18
    Data Transfer, 6-18
PCI Data Transfer, 6-18
    Bus Mastering, 6-18
PCI Parity Detect, 5-7
PCI Transfer/Burst, 5-7
Physical Configuration View, 5-3
Physical Configuration View Screen, (Figure), 5-3
Power–up Sequence
    LED Dispaly (Figure), 2-10
    LED Display, 2-10
Protocols available
    ASPI (unused), 6-16

EATA, 6-16

R

RAID, 6-7
  connection example, 2-8
  Creating Data Redundancy, 6-13
  Degraded Mode, 6-14
  Determining the level, 4-1
  Disk Drive Striping, 6-7
  Level definition, 6-9
  Published Paper, 6-7
  Rebuilding a Failed Hard Drive, 6-15
  Selecting hard drives, 4-2
  Subsystem configuration, 4-1
  Summary, 6-12
RAID Configuration, 1-3
RAID level
  0, 4-1
  1, 4-2
  5, 4-2
RAID Levels, Definition, 6-9
RAID levels, 1-1
RAID–1 Array Group Window, (Figure), 5-12
RAM, Non–volatile, 2-13
Read Caching Enabled, 5-10
Read–ahead Algorithm, 6-4
Read–ahead Caching, (Figure), 6-5
Read–ahead Disk Access, (Figure), 6-4
Real–time I/O Analysis, 1-3
Rebuilding a Failed Hard Drive, 6-15
Reconnect, SCSI, 6-19
Replacing drives, 4-3
Residue Reporting. *See* SCSI

S

S.M.A.R.T. Emulation, 5-11
SCAM, 5-7
Scatter/Gather. *See* EATA Protocol
Schedule Diagnostics Window, 5-25
Scheduling, Diagnostics, 5-25
SCSI
  8-bit/Wide Cables Termination (Figure), 2-6
  Cabling, single–ended, 6-20
  Commands, 6-19
  Configuring Device IDs, 2-7
  Device IDs, 6-19
  Differential, 6-20
  Disconnect, 6-19
  ECC Protection, 6-23
  Interface, 6-19
  LUNs, 6-19
  Messages, 6-19
  Reconnect, 6-19
  Residue Reporting, 6-19
  Single Cable Center Termination (Figure), 2-5
  Single Cable End Termination (Figure), 2-5
  Single–ended, 6-20
  Termination, 2-5, 6-22
  TERMPWR, 6-22
  Transfer Padding, 6-19
  Transfer Rate, 6-20
  Two-Cable Termination (Figure), 2-6
  Wide, 6-20
SCSI Bus, 5-16
SCSI Command Queuing, 5-7
SCSI Commands & Messages, 6-19
SCSI Device
  Diagnostics, 5-23
  Information Windows, 5-9
SCSI Device Diagnostics Window, 5-23
SCSI Device Information Window, (Figure), 5-9
SCSI Devices, 1-2
  Hard drives, 1-2
SCSI Disconnect/Reconnect, 6-19
SCSI ID, 5-7
SCSI Interface, 6-19
SCSI RAID
  Adapter Architecture, 6-16
  Adapter LEDs, 2-10, 2-11
  Audible alarm, 2-11
  NVRAM, 2-13
SCSI Termination, 5-7
SCSI TERMPWR, 5-7
SCSI Transfer Rate, 5-7
Select Array Type Window, (Figure), 5-14
Set–up Parameters, Permanent memory, 2-13
Setting Configuration, 5-32
SIMM Installation, (Figure), 2-2, 2-3
SIMM sockets, 1-2, 2-2, 2-3
SIMMs, Installation, 2-2
Single–ended
  See also SCSI
  Cabling. *See* SCSI
SmartPower, 5-7
Software Caching, 6-1
Sorting, Elevator, 6-3
Start Storage Manager, 5-2
Status, 5-17
Status conditions, Optimal, 5-17
Storage Manager, 1-3
  GUI, 5-1
  Installation, 5-1
  Keyboard functions, 5-2
  Running, 5-2
  Start, 5-2
Storage Manager Window, Logical Device Addressing (Figure), 5-16
Striping, Disk Drives, 6-7
Striping Disk Drives, (Figure), 6-7
System Idle Time, (Figure), 6-2

T

Task Priority, 5-26
Termination, Cabling. *See* SCSI
Terminators, Power. *See* SCSI
TERMPWR, 6-22
Transfer Padding. *See* SCSI
Transfer Rate. *See* SCSI
Troubleshooting
  B4–4, A-1
  B4–C, B-1

W

Wide SCSI. *See* SCSI
Write caching, 5-11
Write Caching Enabled, 5-10
Write Caching Mode, 5-10
Write Trough, 5-7

Write–back Algorithm, 6-3

Glossary

Adapter
SCSI Host Bus Adapter or HBA.

ANSI
American National Standards Institute.

Array Group
A group of disk drives which appear to the computer as a single LSU. RAID–1 and RAID–5 Array Groups can be composed of any number of Hardware Arrays. RAID–0 Array Groups can be composed of any combination of individual disk drives.

ASPI
Advanced SCSI Programming Interface – A protocol used by some SCSI application programs to communicate with SCSI adapters under DOS, OS/2 and Netware.

Asynchronous
Data transfer protocol which is not synchronized to a set timing interval. Asynchronous SCSI data transmitting devices must wait after each byte for acknowledgement from the receiving device. Either device can take as long as it wishes to send or acknowledge data. Asynchronous SCSI has no defined maximum transfer rate but is typically limited to 1.5 to 3MHz.

BIOS
Basic Input Output System – a ROM–based collection of device drivers and system boot–up routines which is provided as an integral part of every PC computer. The BIOS provide enough intelligence to enable the computer to understand some simple keyboard commands and load the operating system from disk upon power–up. In order to do this, the BIOS must contain drivers for the CRT display, keyboard and disk controller.

Bridge Controller
A SCSI device which appears as a single ID on the SCSI bus, but which bridges to multiple devices. Each of these devices is assigned a Logical Unit Number (LUN). Bridge controllers are typically used to control remote RAID subsystems.

Build
The operation of initializing a Redundant Array by creating consistent redundant information. In the case of a RAID–1 array, one drive must be copied to the other. In the case of RAID–5 arrays, the parity information must be generated.

Burst
A term used to describe data words which are transmitted as a single group across a bus without interruption by another device.

Bus Mastering
A method of data transfer which allows data to be moved between a peripheral controller and system memory without interaction with the host CPU or a third party DMA controller. This technique allows the peripheral controller to take control of the system bus and move data at up to 10MB/s, 33MB/s or 132MB/s for ISA, EISA and PCI systems respectively.

Cache
A temporary fast storage area for data which would normally be accessed from a slower storage device. A cache management algorithm monitors the data access patterns and selects which data from the slower device is to be kept in the cache for quick access. Caches are normally transparent or hidden from the accessing device.

CAD
Computer Aided Design – Highly graphically oriented software applications which aid in engineering or other types of design.

CD–ROM
Compact Disk Read–Only Memory. A read–only storage device which retrieves up to 640 Mbytes of information from a removable laser disk similar to an audio compact disk.

Command Overlapping
A feature which allows the simultaneous execution of more than one I/O command by a peripheral controller.

Command Packet
An area of host computer memory which is used to store an I/O command and associated parameters for a peripheral controller which uses EATA protocol. Commands for SCSI RAID adapters are stored in Command Packets.

Command Queuing
A feature which allows multiple I/O commands to be executed by a peripheral controller, in a more efficient order.

Degraded Mode
The mode of operation of a Redundant Array in which it can continue to be accessed after a component drive has failed. For read accesses, data is synthesized from the remaining good drives in the array. For write accesses, data is stored on the remaining good drives in such a way that is can be restored to the failed drive once it is replaced.

Device
According to the SCSI specification, up to fifteen SCSI devices may be connected to a single SCSI bus. SCSI devices include peripherals such as disk and tape drives, optical devices, scanners, printers, and host adapters. Each SCSI device is assigned a SCSI ID number from 0–7 for 8–bit SCSI, or 0–15 for Wide SCSI.

Differential
An electrical signal protocol which transmits information through a current loop rather than by changes in voltage, thereby reducing the susceptibility to electrical interference. Differential SCSI uses RS–485 transceivers to transfer 10MHz data at distances up to 25 meters (82 feet).

DIMM
Dual In–Line Memory Module – A standard way of packaging RAM on a small circuit board with a defined edge connector.

Dirty
A cache page in which data has been written or modified but which has not yet been copied to the storage device. Once the data has been copied to disk, the page is said to be ”clean”.

DMA
Direct Memory Access – A method of data transfer which allows data to be moved between a peripheral controller and system memory without interaction with the host CPU. The data may be moved by the peripheral controller itself, or by a separate 3rd–party DMA controller.

Driver
A software routine which receives I/O requests from higher levels within the operating system and converts those requests to the protocol required by a specific hardware device such as a SCSI adapter.

Duty Cycle
The percentage of time in which an operation is active.

EATA
Extended AT Attachment – The hardware–level high performance protocol used to interface to the host computer. EATA supports Bus Mastering, Scatter/Gather data transfers, and overlapped and queued commands.

ECC
Error Correcting Code – A method of generating redundant information which can be used to detect and correct errors in stored or transmitted data. ECC is stored on disk drives at the end of every sector to correct errors caused by media flaws. ECC is used in SCSI RAID adapters to detect and correct errors caused by cache RAM defects or electrical interference over the SCSI bus.

ECU
EISA Configuration Utility – A software utility program provided by the manufacturer of an EISA computer. It is used to configure EISA option cards instead of manually setting jumpers.

EISA
Extended Industry Standard Architecture – An enhanced version of the PC/AT architecture which provides backward compatibility with existing 8 and 16–bit hardware cards. In addition, EISA supports 32–bit data paths, 33 Mbyte/sec data transfers from Bus Mastering peripheral cards, automatic configuration, and a more sophisticated I/O addressing scheme.

Elevator Sorting
A method of sorting records or cache pages by physical location on disk so that the information may be written to disk with less seek and rotational latency.

Failed
The mode of operation of a drive or array in which the drive or array, because of a malfunction, can no longer be accessed.

Fast SCSI
The original SCSI specification defined synchronous data transmission rates of up to 5MHz. By assuming transceivers which provide tighter timing margins, the SCSI–2 standard allows synchronous transfers of up to 10MHz. This provides a transfer speed of 10MB/s for an 8–bit bus, and 20MB/s for a 16–bit bus. Devices which utilize these faster timings are called Fast SCSI devices.

Fast–20 SCSI
The SCSI–3 specification defined synchronous data transmission rates of up to 20MHz. The 20MHz transfer rate provides speeds up to 20MB/s on an 8–bit bus and 40MB/s on a 16–bit bus. Devices that support these timings are called Fast–20 SCSI devices.

Flash ROM
A ROM on the adapter containing firmware that can be reprogrammed via software without removing it from the board.

Flush
Copying all dirty data in the cache to disk.

Function Button
An object used in a graphical user interface, on which the user can point and click with a mouse, to execute a specific function.

GUI
Graphical User Interface. A software interface which interacts with the user through color graphics and a mouse.

Hardware Array
A group of disk drives which are all members of the same RAID–1 or 5 array implemented in host adapter hardware. Multiple Hardware Arrays can be combined into a common Array Group.

HBA
Host Bus Adapter – Refers to the SCSI RAID adapter, SmartCache or other peripheral controller cards which provide host computer access to the SCSI bus.

Hit
A data access in which the requested data is found in the cache.

Hit Ratio
The ratio of cache hits to total disk accesses. A hit ratio of 100% means that all disk accesses can be serviced from the cache.

Hot Plug
The operation of adding or removing a device from a bus while transactions involving other devices are occurring over the bus.

Hot Spare
A spare disk drive which, upon failure of a member of a redundant disk array, will automatically be used to replace the failed disk drive.

Hot Swap
The operation of removing a failed disk drive which is a member of a Redundant Array and replacing it with a good drive, while transactions involving other devices are occurring over the bus.

ID
A number from 0 to 7 (or 0 to 15 for 16–bit Wide SCSI) which SCSI devices use to address one another. SCSI IDs are typically selected by setting jumpers on the SCSI device.

Icon
A graphical representation of an object or function used by a GUI.

Icon Information Window
A window, displayed by Storage Manager upon selection of an icon, which contains information and functions pertaining to that device.

ISA
Industry Standard Architecture – the name given to the original IBM PC/AT 16–bit bus architecture.

Latency
The time required by a device to access stored data, excluding the data transfer time. Reducing disk latency results in more I/O operations per second being performed on a disk drive.

LED
Light Emitting Diode.

LRU
Least Recently Used – A cache management algorithm employed by SCSI RAID adapter to determine the next cache page to delete and re–use once all pages have been filled with data. By deleting the page in cache which has gone the longest without an access, the LRU algorithm typically results in the most frequently accessed data being cache resident.

LSU
Logical Storage Unit – A logical device on which the computer can store and retrieve information. An LSU may represent an individual disk drive or an Array Group.

LUN
Logical Unit Number – Each SCSI device may contain up to eight sub–devices or Logical Units, assigned Logical Unit Numbers 0–7. Typically, SCSI devices such as a disk or tape drives contain only LUN 0.

Menu
An item at the top of a GUI display which, when selected by clicking with a mouse, will cause a list of related functions to be displayed.

Mirroring
A popular term for RAID–1. A method of creating disk–fault tolerance by redundantly storing information on pairs of drives.

Miss
A data access in which the requested data is not found in cache. Cache misses result in the normal average disk latency of 20ms when accessing data.

MTBF
Mean Time Between Failure. The average time between expected failures of a device.

Non–Redundant Array
An Array Group with no fault tolerance (RAID–0). If one drive in a Non–Redundant Array fails, the entire Array Group will fail.

NVRAM
Non–Volatile Random Access Memory.

Optimal
The mode of operation of a disk array in which no drive failures have occurred.

OS
Operating System – Software which manages the computer's resources and provides the operating environment for application programs.

Page
The smallest region of cache which may be allocated to storing data. The SCSI RAID adapter can contain up to 128,000 pages, each storing one 512–byte sector from disk.

Parity
A method of generating redundant information which can be used to detect errors in stored or transmitted data. Parity is used in standard RAM SIMMs and over the SCSI bus to detect data errors. Parity is used in RAID–5 disk arrays to reconstruct flawed or missing data sectors.

PCI
Peripheral Component Interconnect – An intelligent computer bus specification which supports 32–bit data paths, 132 Mbyte/sec data transfers from Bus Mastering devices and plug–and–play automatic configuration of peripheral cards.

PIO
Programmed Input/Output – The method of data transfer used whereby data is moved between the peripheral controller and system memory via the CPU. PIO transfers typically occur at about 2.5MB/s.

Plug–and–Play
The ability to install peripheral cards or devices without requiring manual configuration by the user.

Plug–and–Play BIOS
The system BIOS in PCI computers which has the ability to automatically configure PCI peripheral cards so as to avoid contention with other installed devices.

RAID
Redundant Array of Inexpensive Disks – A method of combining hard disks into one logical storage unit which offers disk–fault tolerance and can operate at higher throughput levels than a single hard disk.

Rebuild
The operation of recreating data belonging to a failed member of a redundant disk array and writing that data onto a replacement disk drive.

Redundant Array
A fault tolerant Array Group. (RAID–1 or RAID–5).

Scatter/Gather
A feature which allows data to be transferred to or from multiple discontiguous areas of host computer memory with a single I/O command.

SCSI
Small Computer Systems Interface – an ANSI standard parallel interface designed to communicate with intelligent peripheral devices. The SCSI definition, created by ANSI is defined in document number X3.131–1986. Copies of this document can be obtained from the ANSI X3 Secretariat.

SIMM
Single In–line Memory Module – A standard way of packaging RAM on a small circuit board with a defined edge connector. SCSI RAID adapters can use any combination of 4MB or 16MB 36–bit single–sided low–profile RAM SIMMs.

Single–Ended
An electrical signal protocol which transmits information through changes in voltage. Single–Ended SCSI uses standard TTL signal–and–ground pairs to transmit information over the SCSI bus.

Status Packet
An area of host computer memory which is used to store the completion status for an I/O command for a peripheral controller which uses EATA protocol. SCSI RAID adapters return command completion status in Status Packets.

Stripe
A contiguous region of disk space. Stripes may be as small as one sector or may be composed of many contiguous sectors.

Striping
Also called RAID–0. A method of distributing data evenly across all drives in an array by concatenating interleaved stripes from each drive.

Synchronous
Data transmission protocol which is synchronized to a defined time interval. Synchronous SCSI can transmit data faster than asynchronous SCSI because the transmitting device does not wait for acknowledgement of each byte from the receiving device. Instead, it continues to transmit data at the rate negotiated by both devices (up to 20MHz).

Tagged Command Queuing
A feature of SCSI–2 and SCSI–3 protocols that allows SCSI commands to be executed out of order.

Termination
A method of matching the transmission impedance of a electrical bus so as to eliminate signal reflections from the physical ends of the bus.

TERMPWR
A signal on the SCSI bus which provides power for SCSI bus terminators on remote devices.

Throughput
A term used to describe the amount of data which can be processed by a system in a given amount of time.

Tool Bar
A bar, near the top of the screen used by a GUI, which contains Function Buttons.

Ultra SCSI
Another name used for Fast–20 SCSI.

VGA
Video Graphics Adapter.

Write–Back
A method postponing data to be written to a slow device such as a disk drive by temporarily saving the data in a cache. The data can then be written at a time when the device would otherwise be idle. In addition, the data can be processed through techniques such as elevator sorting, so that the write operations will go faster.

Write–Through
The opposite of write–back. In write–through mode, data must be written to the final destination before a write operation is completed.

Wide SCSI
A SCSI protocol and signal definition which provides a greater than 8–bit wide data path. Wide SCSI devices may support either 8 and 16–bit, or 8, 16 and 32–bit data transfers.

WORM
Write Once Read Multiple – An optical storage device similar to a CD–ROM which can write data one time only to any location on a removable laser disk. WORMs are primarily used as data archive devices because once written, the data can never be erased.

XOR
Exclusive OR – A logical operation performed on two binary operands which yields a 0 in every bit position where the operands are both 1 or both 0, and a 1 in every bit position when the operands are dissimilar. Performing XOR operations on every bit in a word generates a parity bit for that word.

Vos remarques sur ce document / Technical publication remark form

Titre / Title : Bull ESCALA SCSI RAID Adapter Installation & Configuration Guide

Nº Reférence / Reference Nº : 86 A1 44HX 02

Daté / Dated : April 2000

ERREURS DETECTEES / ERRORS IN PUBLICATION

AMELIORATIONS SUGGEREES / SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

Vos remarques et suggestions seront examinées attentivement.
Si vous désirez une réponse écrite, veuillez indiquer ci-après votre adresse postale complète.

Your comments will be promptly investigated by qualified technical personnel and action will be taken as required.
If you require a written reply, please furnish your complete mailing address below.

NOM / NAME : ______________________________ Date : __________

SOCIETE / COMPANY : ______________________________

ADRESSE / ADDRESS : ______________________________

____

Remettez cet imprimé à un responsable BULL ou envoyez-le directement à :

Please give this technical publication remark form to your BULL representative or mail to:

BULL ELECTRONICS ANGERS
CEDOC
34 Rue du Nid de Pie – BP 428
49004 ANGERS CEDEX 01
FRANCE

Technical Publications Ordering Form

Bon de Commande de Documents Techniques

To order additional publications, please fill up a copy of this form and send it via mail to:
Pour commander des documents techniques, remplissez une copie de ce formulaire et envoyez-la à :

BULL ELECTRONICS ANGERS
CEDOC
ATTN / MME DUMOULIN
34 Rue du Nid de Pie – BP 428
49004 ANGERS CEDEX 01
FRANCE

Managers / Gestionnaires :
Mrs. / Mme : **C. DUMOULIN** +33 (0) 2 41 73 76 65
Mr. / M : **L. CHERUBIN** +33 (0) 2 41 73 63 96

FAX : +33 (0) 2 41 73 60 19
E–Mail / Courrier Electronique : srv.Cedoc@franp.bull.fr

Or visit our web site at: / Ou visitez notre site web à:
`http://www-frec.bull.com` (PUBLICATIONS, Technical Literature, Ordering Form)

CEDOC Reference # Nº Référence CEDOC	**Qty** Qté	**CEDOC Reference #** Nº Référence CEDOC	**Qty** Qté	**CEDOC Reference #** Nº Référence CEDOC	**Qty** Qté
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
__ __ ____ _ [__]		__ __ ____ _ [__]		__ __ ____ _ [__]	
[_ _] : **no revision number means latest revision** / pas de numéro de révision signifie révision la plus récente					

NOM / NAME : ____________________ Date : __________

SOCIETE / COMPANY : ____________________

ADRESSE / ADDRESS : ____________________

____

PHONE / TELEPHONE : ____________________ FAX : ____________________

E–MAIL : ____________________

For Bull Subsidiaries / Pour les Filiales Bull :
Identification: ____________________

For Bull Affiliated Customers / Pour les Clients Affiliés Bull :
Customer Code / Code Client : ____________________

For Bull Internal Customers / Pour les Clients Internes Bull :
Budgetary Section / Section Budgétaire : ____________________

For Others / Pour les Autres :
Please ask your Bull representative. / Merci de demander à votre contact Bull.

PLACE BAR CODE IN LOWER
LEFT CORNER

BULL ELECTRONICS ANGERS
CEDOC
34 Rue du Nid de Pie – BP 428
49004 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A1 44HX 02

Bull

Utiliser les marques de découpe pour obtenir les étiquettes.
Use the cut marks to get the labels.

ESCALA

SCSI RAID Adapter
Installation &
Configuration
Guide

86 A1 44HX 02

ESCALA

SCSI RAID Adapter
Installation &
Configuration
Guide

86 A1 44HX 02

ESCALA

SCSI RAID Adapter
Installation &
Configuration
Guide

86 A1 44HX 02