

PCI 4-Channel Ultra3 SCSI RAID Adapter

Reference Guide

First Edition (October 2000)

Before using this information and the product it supports, read the information in "Appendix. Notices" on page 77.

International Business Machines Corporation, 2000. All rights reserved. Note to U.S. Government Users
Restricted Rights--Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About This Book	v
ISO 9000	v
Related Publications	v
Trademarks	v
 Chapter 1. PCI 4-Channel Ultra3 SCSI RAID Adapter Overview	 1
Disk Arrays	1
RAID Levels	3
RAID Level 0	4
RAID Level 1	5
RAID Level 5	6
RAID Level 5 Enhanced (5E)	8
PCI 4-Channel Ultra3 SCSI RAID Adapter Features	10
PCI 4-Channel Ultra3 SCSI RAID Adapter	10
Disk Array	10
SCSI Disk Drives	10
128 MB Non-volatile on Board Write Cache	11
PCI SCSI Disk Array Manager (PDAM)	11
Disk Array Overview	12
Disk Array Parameters	12
Identify Drives and Disk Arrays	16
Disk Array ODM Information	17
 Chapter 2. Installing the PCI 4-Channel Ultra3 SCSI RAID Adapter Software	 21
Installing an PCI 4-Channel Ultra3 SCSI RAID Adapter on AIX Version 4	21
Prerequisite	22
Installing a PCI 4-Channel Ultra3 SCSI RAID Adapter Software Update	24
Procedure	24
Verifying the PCI 4-Channel Ultra3 SCSI RAID Adapter Software Installation	25
 Chapter 3. Managing the PCI 4-Channel Ultra3 SCSI RAID Adapter	 27
Procedure	27
Starting the PCI SCSI Disk Array Manager	28
Creating a Disk Array	29
Modifying and Displaying Drive Status	31
Disk Array States	31
Physical Disk States	32
View the Current Disk Array Status	33
Add a Spare Drive	33
Delete a Spare or Failed Drive	34
Fail a Drive	34
Removing a Disk Array	35
Delete a PCI SCSI Disk Array	35
Reviving a Failed Drive	36
Assigning Hot Spare Drives	36
Transitioning a Hot Spare Drive to a Spare Drive	36
Performing a Consistency Check on a Disk Array	37

Reconstructing a Failed Drive	37
Change/Show PCI SCSI Disk Array Parameters	38
Configuration Synchronization	39
Clearing a Configuration.	40
Restoring Configuration on Replacement Adapter	40
Adding Disk Arrays to the AIX Operating System	41
Deleting Disk Arrays from the Operating System.	43
Extending Logical Volumes and Disk Arrays	44
 Chapter 4. PCI 4-Channel Ultra3 SCSI RAID Adapter Problem Determination and Recovery	45
Drive Failures and Disk Arrays	45
Drive Status.	45
Disk Array Failures	47
Reconstructing Redundant Disk Arrays after Drive Failure	50
The Reconstruction Rate - Low, Medium, High	51
Checking Consistency (Check and Repair Array Parity)	52
Perform Check Consistency	52
Restoring and Recovering Disk Arrays	53
Restoring a Disk Array	58
Replacing a Dead or Optimal Array	59
Replacing the PCI 4-Channel Ultra3 SCSI RAID Adapter	60
Detecting and Resolving Array Configuration Conflicts	61
Resolving Disk Array Configuration Conflicts	63
Resolving Common Problems	65
Messages and Error Log Templates	66
Temporary RAID Adapter Error	66
Permanent RAID Adapter Error	67
Temporary Array SCSI Bus Error	68
Temporary Physical Volume Error	69
Permanent Physical Volume Error	70
Array Configuration Error	71
Array Component Failure	72
Temporary Unknown Software Failure	73
Recovered RAID Adapter Error	74
Alertable Messages	75
 Appendix. Notices	77
 Index	79
 Reader's Comments — We'd Like to Hear From You	81

About This Book

This book provides information on how to set up the system unit, install and remove options, and verify system operation.

ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

Related Publications

The following publications provide additional information about or related to the PCI 4-Channel Ultra3 SCSI RAID Adapter.

- The *Diagnostic Information for Multiple Bus Systems*, order number SA38-0509, contains diagnostic information, service request numbers (SRNs), and failing function codes (FFCs).
- The RAID book, *A Source Book for Disk Array Technology*, Edition 4. Published by: The RAID Advisory Board, ST Peter, Mn, 1994
- The *AIX Version 4 System Management Guide: Operating System and Devices*, order number SC23-2525.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

- AIX
- IBM

Other company, product, and service names may be trademarks or service marks of others.

Chapter 1. PCI 4-Channel Ultra3 SCSI RAID Adapter Overview

The primary features of Redundant Array of Inexpensive Disks (RAID) technology are:

- Large disk capacity
- Immediate availability and recovery of data
- Redundancy of data at a user selected level
- Enhanced performance
- 128 MB nonvolatile onboard write cache

RAID technology is used to store data across a series of disk drives known as a disk array. Depending on the RAID level selected, this storage technique provides the data redundancy required for a secure system without depleting memory. It may also provide faster retrieval through multiple channel access. If a hardware failure occurs, a single disk drive can usually be replaced without interrupting normal system operation.

Disk Arrays

Disk arrays are groups of disk drives that work together with a specialized array controller to achieve higher data transfer and input and output (I/O) rates than those provided by single large drives. The array controller keeps track of how the data is distributed across the drives.

Redundant disk arrays can also provide data redundancy, so that no data is lost if a single drive in the array fails. Two methods of writing data to the disk drives are used in a RAID subsystem:

striping

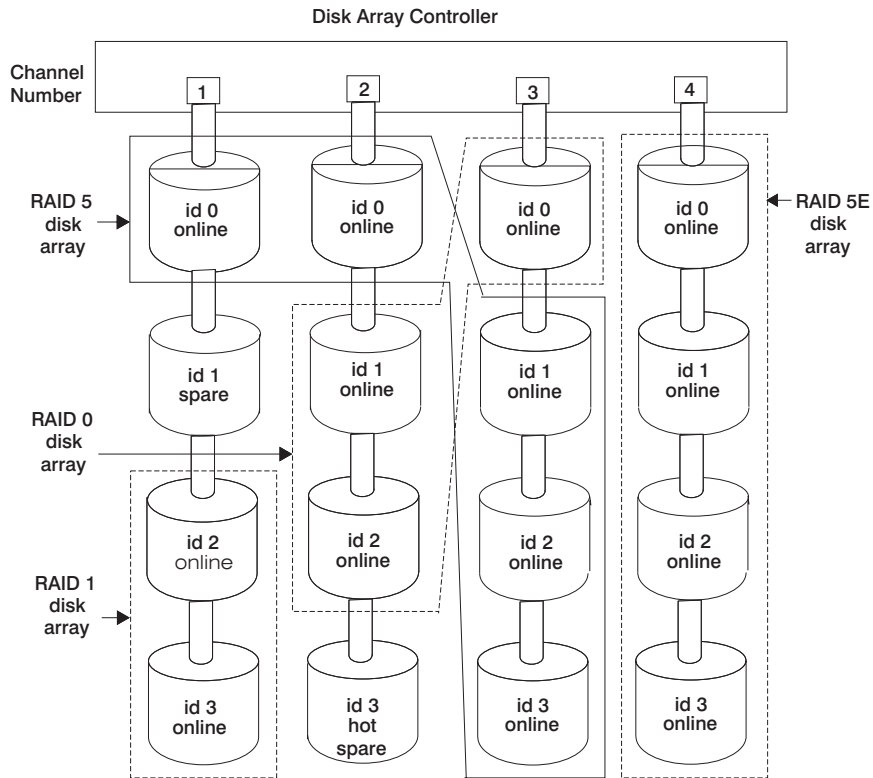
Data for a given file may be written in stripe units to different drives in the array, rather than being written to a single drive. By using multiple drives, the array can provide higher data transfer rates and higher I/O rates when compared to a single large drive.

mirroring

Data that is simultaneously written to two separate disks within the same array.

The method used for writing the data to an array is determined by the RAID level defined for that array and the number of drives used. "RAID Levels" on page 3 discusses the various RAID levels and their function in detail.

The Disk Arrays and Channels figure illustrates the terms used in describing disk arrays.



Disk Arrays and Channels

Each disk array is viewed as a single logical device with a unique disk name and location code.

Each disk array has its own array parameters (for example, RAID level and capacity). A disk array is treated as a single disk drive by the operating system. There are no special requirements for using it. AIX commands and utilities work on disk arrays just as they would with single non-RAID drives.

For example, after you configure a disk array with the PCI SCSI Disk Array Manager (PDAM), you can use AIX commands to make the disk array available to the system by adding it to a volume group, creating a file system on it, and so on. Raw I/O is also supported on disk arrays. When adding a disk array to your system, treat it in the same way you would treat a single disk drive.

Note: Diagnostic open and SCSI passthrough are not supported on disk arrays.

RAID Levels

For Each RAID level, a different error recovery is available when a part of the system fails. Depending on the RAID level chosen (with the exception of RAID level 0), if a single drive fails within an array, the data from that disk can be reconstructed from the data stored on other hard drives within the array. This data can be reconstructed with little or no impact to the users and programs that are using the system at the time of the failure.

Supported RAID levels vary depending on hardware and software platforms, but the most common supported levels are 0, 1, 3, and 5.

Note: The PCI 4-Channel Ultra3 SCSI RAID Adapter supports RAID levels 0, 1, 5, and 5E.

Each RAID level supported by the adapter provides a different method of writing data that has specific benefits.

RAID level 0

Does not provide data redundancy, but provides a potentially higher I/O rate.

RAID level 1

Maintains a duplicate copy of the data so that data can be reconstructed if it is lost when a drive fails.

RAID level 5

Creates array parity information so that data can be reconstructed if it is lost when a drive fails in the array.

RAID level 5E


Creates array parity information so that data can be reconstructed if it is lost when a drive fails. The Hot spare capacity is spread among all drives in the array.

The following sections provide detailed information about these RAID levels.

RAID Level 0

Attention: Regularly back up data on a RAID level 0 array. Only backed up data can be recovered if a disk drive fails.

RAID level 0 writes data across the drives in the array, one stripe unit at a time. In this example, a stripe unit is defined as 32 blocks of 512 bytes each (16K stripe unit size), as shown in the RAID Level 0 figure. For a 128K write, blocks 0 through 31 are written to drive 1, blocks 32 through 63 are written to drive 2, blocks 64 through 95 are written to drive 3 and so on as each drive has a single stripe unit written to it. Finally, blocks 224 through 255 are written to drive 3 to complete the write operation.



RAID Level 0

The host system treats a RAID level 0 array like a standard disk drive. RAID level 0 errors are reported in the same way as normal drive errors, and the recovery procedures are the same as those used on a standard disk drive. For example, if a drive fails, the array controller returns the same errors that occur during a read or write retry failure. All data on the array may be lost. However, unlike other RAID levels, the array controller never marks a RAID level 0 array as degraded as the result of a drive failure. If a physical drive fails in a RAID 0 disk array, the disk array is marked as dead.

RAID level 0 offers a higher potential I/O rate, but is a non-redundant configuration. That is, there is no array parity information generated for the purpose of reconstructing data in the event of a drive failure. Therefore, there is

no error recovery beyond what is normally provided on a single drive. All data in the array must be backed up regularly to protect against data loss.

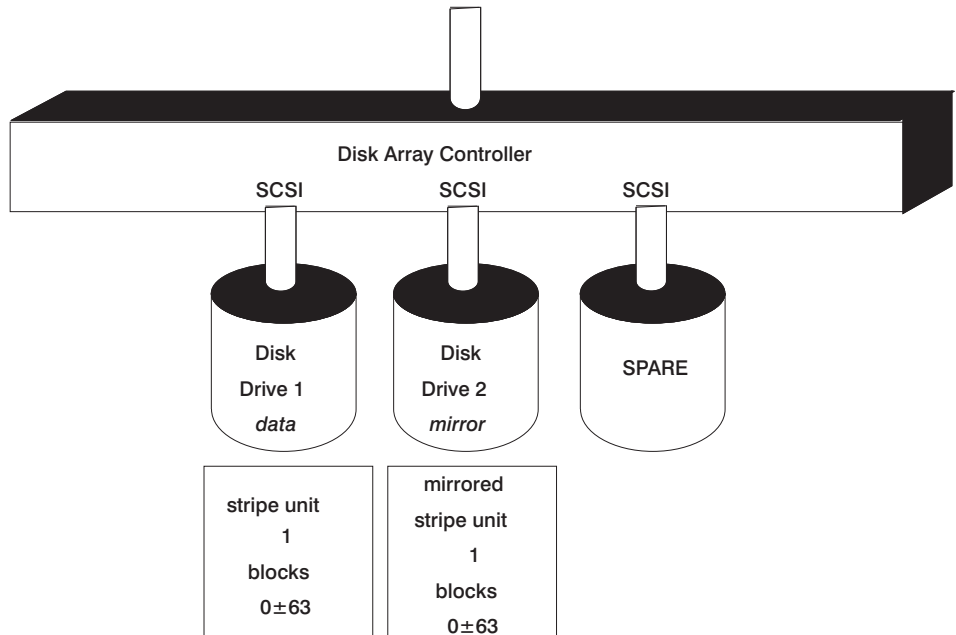
RAID Level 1

Attention: Although a RAID level 1 array has data redundancy, you should regularly back up data on the array. This is the only way to recover data in an event such as accidental file deletion or disaster recovery. You can continue to operate the array in degraded mode until you replace the drive. However, you should replace the drive as soon as possible. If you cannot replace the drive immediately, back up your data, file by file, to prevent potential data loss.

RAID level 1 transparently mirrors data by duplicating data stripes across drives. When data is written to a drive, it is also written to a mirrored stripe.

RAID level 1 has traditionally been used for critical fault-tolerant transaction processing. Mirrored data provides high reliability. When a small block size is used, mirrored data also provides a high I/O rate. However, RAID level 1 is a more costly RAID solution because it requires a mirrored data stripe for every physical drive in the array.

RAID level 1 writes data across the drives in the array, one stripe unit at a time. In this example, a stripe unit is defined as 64 blocks of 512 bytes each (32K stripe unit size) as shown in the RAID Level 1 figure. For a 32K write, blocks 0 through 63 are written to drive 1 and the mirrored data blocks 0-63 are written to the mirrored data drive .



RAID Level 1

If a single drive fails in a RAID level 1 array, you can continue to use the array. A RAID level 1 array operating with a single failed drive is said to be operating in degraded mode. Whenever you read or write to a disk array in degraded mode, the array controller retrieves the failed drive's data from its mirrored drive. Although you can continue to operate the RAID level 1 array with a failed drive, you should replace the drive and restore the array as soon as possible.

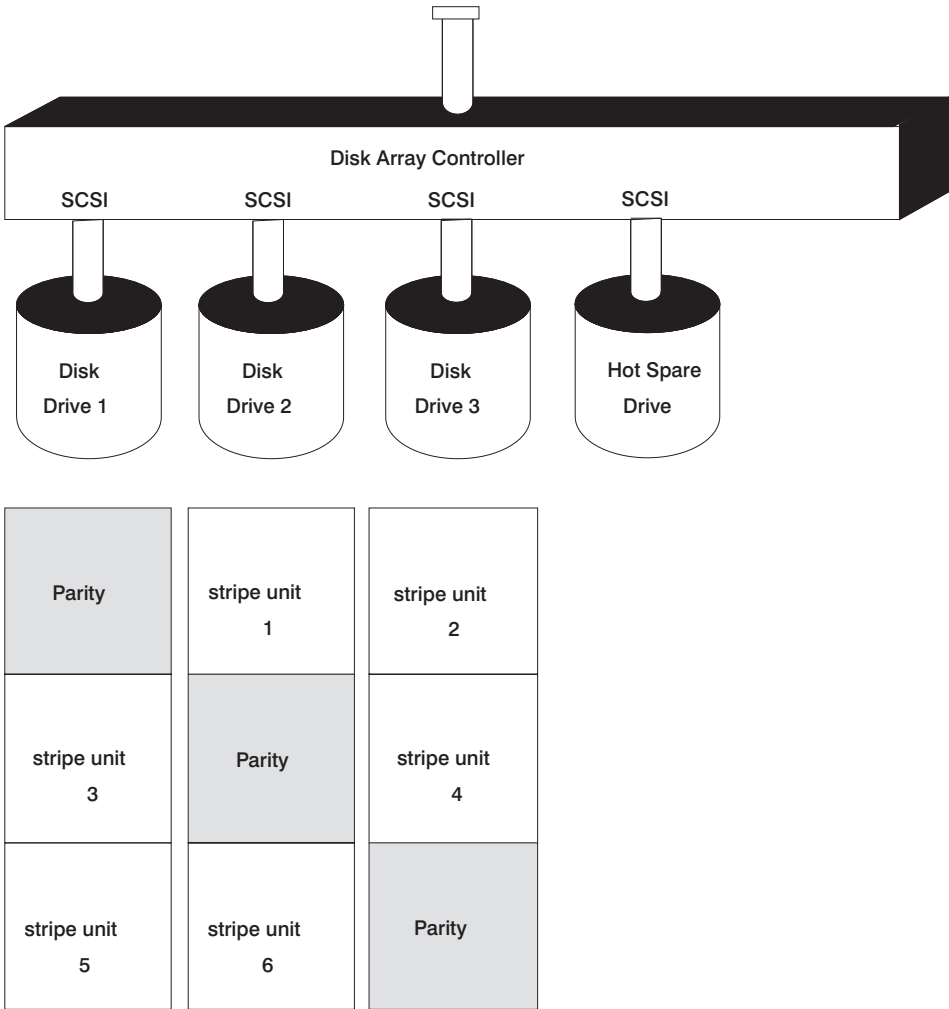
Using the Perform Consistency Check menu allows the array controller to compare the data drive and mirror to verify consistency after an abnormal system shutdown. See "Performing a Consistency Check on a Disk Array" on page 37 for more information.

Raid Level 5

Attention: Although a RAID level 5 array maintains parity information, you should regularly back up data on the array. This is the only way to recover data in an event such as accidental file deletion or disaster recovery. You can continue to operate the array in degraded mode until you replace the drive. However, you should replace the drive as soon as possible. If you cannot replace the drive immediately, back up your data, file by file, to prevent potential data loss.

RAID level 5 stripes data across all drives in the array, one stripe unit at a time (a stripe unit can contain multiple blocks). RAID level 5 also writes array parity data. The parity data is spread across all the drives.

In the RAID Level 5 example figure, a stripe unit is defined as 64 blocks of 512 bytes each (32K stripe unit size).



RAID Level 5

Stripes are written as follows:

- Stripe 1 is written in the first position to drive 2
- Stripe 2 is written to drive 3
- Parity data for the data in stripes 1 and 2 are written to drive 1

If a drive fails in a RAID level 5 array, you can continue to use the array normally. A RAID level 5 array operating with a single failed drive is said to be operating in degraded mode. Whenever data is read from a degraded disk array, the array controller recalculates the data on the failed drive by using data and parity blocks on the operational drives.

For example, to recalculate data in data stripe unit 2 in the RAID level 5 figure (the first position on drive 3), the array controller would use the parity information from drive 1 and the data from drive 2 (data stripe unit 1) to reconstruct the data. This process is repeated to reconstruct each block of the failed drive, as needed, so you can continue to operate the RAID level 5 array.

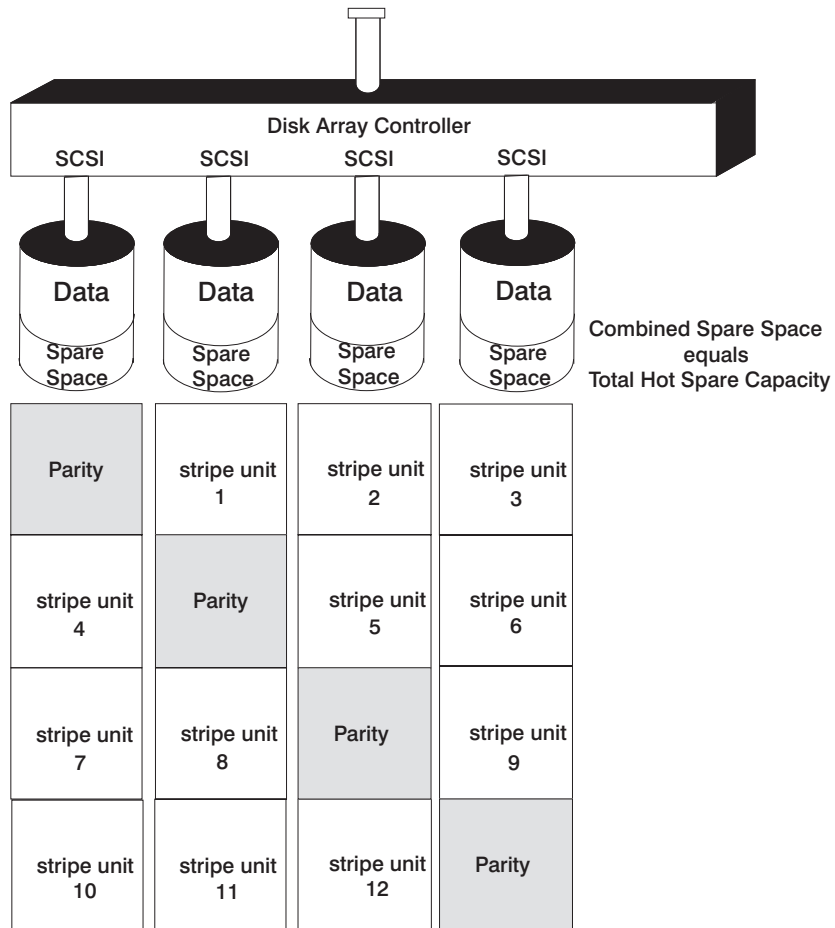
The Perform Consistency Check menu allows the array controller to check the integrity of the array parity. If the array subsystem has an abnormal system shutdown, you may check and repair array parity/mirror on the affected disk array(s). See “Performing a Consistency Check on a Disk Array” on page 37 for more information.

Raid Level 5 Enhanced (5E)

Attention: Although a RAID level 5E array maintains parity information, you should regularly back up data on the array. This is the only way to recover data in an event such as accidental file deletion or disaster recovery. You can continue to operate the array in degraded mode until you replace the drive. However, you should replace the drive as soon as possible. If you cannot replace the drive immediately, back up your data, file by file, to prevent potential data loss.

Note: RAID level 5E requires a minimum of four drives.

RAID level 5E is an IBM enhancement of RAID level 5. Traditional RAID 5 has active drives plus an inactive hot spare drive. In RAID 5E however, the hot spare drive is incorporated as an active element in the array. The data is laid out in such a way that the Hot spare capacity is spread among all the drives in the array. If a drive failure occurs, the data on the failed drive is reconstructed using the Hot spare capacity on the remaining drives in the array. This process of reconstruction in a RAID 5E array is referred to as compression. When compression is complete the disk array will have an Optimal status. A compressed RAID 5E array is identical to an optimal traditional RAID 5 array. If the failed drive in a RAID 5E array is then replaced with a new Hot spare drive, the array data on the array will migrate back to a RAID 5E layout with Hot spare capacity again spread among all the drives. See “Assigning Hot Spare Drives” on page 36. This process is referred to as decompression. Decompression will begin when a Hot spare, of equal or greater capacity to the failed drive is added.



RAID Level 5E

Note: If a Hot spare is added while the disk array has a status of Compressing, then the disk array will finish compressing before decompressing the disk array.

On a decompressed RAID level 5E array data stripes have a 3:1 ratio to data parity.

Stripes are written as follows:

- Stripe 1 is written in first position of drive 2
- Stripe 2 is written to drive 3
- Stripe 3 is written to drive 4
- After all three stripes are written the data parity is written to drive 1.

The data parity is calculated using stripes 1,2 and 3.

PCI 4-Channel Ultra3 SCSI RAID Adapter Features

The major components of the PCI 4-Channel Ultra3 SCSI RAID Adapter subsystem are:

- PCI 4-Channel Ultra3 SCSI RAID Adapter. The SCSI controllers and RAID logic are located on the adapter.
- 128 MB Non-volatile on Board Write Cache Module.
- Disk array. A collection of one or more SCSI disk drives.
- SCSI disk drives.
- PDAM (SMIT and command line interface).
- AIX device drivers.

The components of the RAID subsystem combine to provide an advanced storage technique designed to secure online information while speeding file access through multiple, independent transactions.

PCI 4-Channel Ultra3 SCSI RAID Adapter

The PCI 4-Channel Ultra3 SCSI RAID Adapter has built-in support for RAID levels 0, 1, 5, and 5E. It serves as an interface between the host PCI bus and the disk arrays attached to the SCSI busses. The adapter contains two dual channel Ultra3 SCSI controllers which manage the SCSI busses. It also supports background parity initialization which permits disk arrays to be used immediately by the system (without waiting for data parity/mirror initialization of all drives to complete).

The RAID functions of the adapter are controlled by firmware that is executed on an embedded PowerPC RISC processor. Performance is enhanced through a non-volatile write cache on the adapter.

Consistency of the data (RAID levels 1, 5, or 5E) contained on the RAID arrays is monitored continuously through the data scrubbing feature.

Note: The PCI 4-Channel Ultra3 SCSI RAID Adapter does not support boot devices.

Disk Array

The disk array is an abstraction that combines one or more SCSI disk drives and presents them as a logical drive to the host. SCSI disk drives can be grouped together to form a disk array (RAID disk array). See “Disk Array Overview” on page 12 for more information.

Note: The disk arrays do not support the root volume group (**rootvg**). No portion of **rootvg** should exist on a disk array.

SCSI Disk Drives

Each SCSI bus supports attachment of up to 15 SCSI disk drives. Each adapter supports attachment of up to 60 SCSI disk drives.

Note: SCSI devices other than disk drives are not supported on the PCI 4-Channel Ultra3 SCSI RAID Adapter.

128 MB Non-volatile on Board Write Cache

The PCI 4-Channel Ultra3 SCSI RAID Adapter has battery backed non-volatile on board write cache. Select Display Status of Adapter Write Cache in the PDAM to display the status of the write cache.

PCI SCSI Disk Array Manager (PDAM)

The PCI 4-Channel Ultra3 SCSI RAID Adapter is managed by the PDAM. The PDAM is the only interface to the RAID configuration, monitoring, and recovery features of the adapter. Some of the tasks performed using the PDAM include:

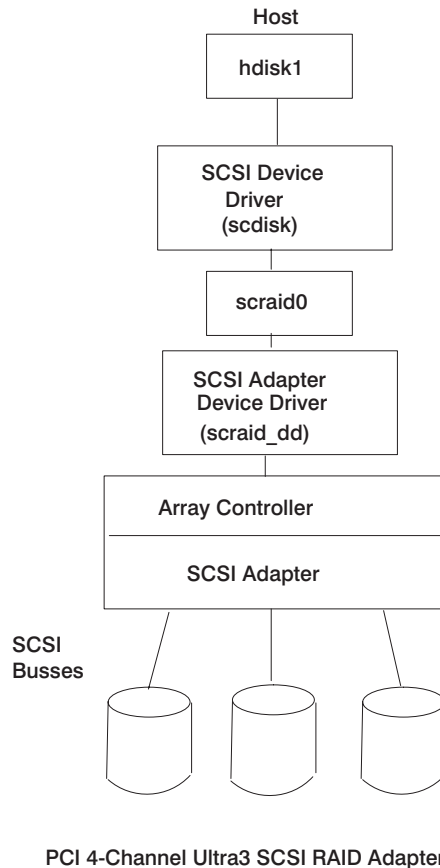
- Check device status for the disk array on your system.
- Create and delete disk arrays.
- Configure and reconfigure disk arrays.
- Change disk array configuration attributes.
- Check and repair array parity/mirror on disk array.
- Display information of physical drives and disk arrays.
- Reconstruct a RAID level 1, 5, or 5E disk array after a single drive failure.
- Delete/recreate a RAID level 0 disk array after a drive failure, or a RAID level 1, 5, or 5E disk array after multiple drive failures.
- Use adapter configuration and error recovery options.

The PDAM is supported through the SMIT interface. See “Chapter 2. Installing the PCI 4-Channel Ultra3 SCSI RAID Adapter Software” on page 21 for more information

Disk Array Overview

The disk array is an abstraction that combines one or more SCSI disk drives and presents them as an logical drive to the host.

The following figure illustrates a typical configuration of a RAID disk array connected to a host system through a PCI 4-Channel Ultra3 SCSI RAID Adapter. This disk array is made up of three physical disks located on three separate SCSI buses. This configuration results in making the disk array available as `hdisk1`. The figure includes the logical names assigned by the system during configuration.



Disk Array Parameters

Disk arrays are configured using the PDAM. See “Chapter 2. Installing the PCI 4-Channel Ultra3 SCSI RAID Adapter Software” on page 21 for additional information.

Each disk array has a set of parameters that determine how data is stored and accessed on it. Some of these parameters may differ for each array, while other parameters must be the same for all arrays created on the adapter. All disk arrays attached to a PCI 4-Channel Ultra3 SCSI RAID Adapter must share the same stripe unit

size, reconstruction rate, and data scrubbing option. Before creating the first disk array, a stripe unit size (16K, 32K, or 64K), reconstruction rate (low, med, or high) must be selected. Once the first disk array is created, all subsequent disk arrays must have the same stripe unit size, reconstruction rate, and data scrubbing option. These parameters cannot be changed until all disk arrays on the adapter have been deleted.

You can define up to:

- 8 disk arrays per PCI 4-Channel Ultra3 SCSI RAID Adapter
- 16 physical drives per disk array.

The following disk array parameters table summarizes the more detailed descriptions that follow:

Note: See the Create a PCI SCSI Disk Array and Change/Show PCI SCSI Drive Status menus in SMIT.

Parameter	Meaning
RAID Level	Determines how data is stored on the disk array, and whether there is data redundancy.
Size of Array (Mbs)	Size of the disk array.
Stripe Size**	Determines the amount of data written to a single drive in the disk array before the controller writes data on the next drive.
Reconstruction Rate**	Time between reconstruction operations.
Data Scrubbing**	Enables automatic perpetual background parity check for redundant RAID level (1, 5, 5E) arrays.
Select Drives	Determines what drives make up the disk array.
Command Queue Depth	Size of command queue.
Read-ahead Enable	Enables read-ahead caching on a disk array.

** Parameters can be changed only from the Change/Show PCI SCSI RAID Adapter SMIT menu when no disk arrays exist on an adapter.

RAID Level

Determines how data is stored on the disk array. RAID levels 1, 5, and 5E offer data redundancy; RAID level 0 does not. The levels you select depend on your storage and performance needs. See "RAID Levels" on page 3 for more information.

The RAID level dictates the minimum number of drives that can be included in the disk array and the realized capacity of the array. Refer to the following parameter descriptions.

Size of Array

Size is determined by the RAID level, number of drives, and the capacity of those drives. On the PCI 4-Channel Ultra3 SCSI RAID Adapter, when a physical drive is selected to be part of a disk array, its entire capacity is used. Physical drives cannot be fragmented or spread across multiple disk arrays.

Note: The capacity should be the same for each drive in the disk array. If drives differ in capacity, the array controller uses the smallest capacity in the selected drive list to determine the usable capacity of all the drives when creating the disk array.

To determine the size of a disk array you want to create on your array subsystem, you need to know the capacity of a single drive. To determine this value, select the Change/Show PCI SCSI Drive Status option from the PCI SCSI Disk Array Manager menu in SMIT.

Each time you create a new disk array from spare drives, use the following formulas to determine the maximum size of the disk array you can create. This calculation depends on both the RAID level of the disk array you want to create and the size of the drives:

RAID level 0

Multiply the number of drives by the drive capacity.

RAID level 1

Multiply the number of drives by the drive capacity and divide by 2.

RAID level 5

Multiply one less than the number of drives by the drive capacity.

RAID level 5E

Multiply two less than the number of drives by the drive capacity.

The resulting capacity for each of the preceding calculations yields the approximate size of the disk array.

Stripe Size

Specifies the amount of data written on a single drive in the disk array before the controller continues writing the data on the next drive in the disk array. For example, if the stripe unit size of a RAID level 0 disk array is 16384 bytes (16K), the controller will write 16384 bytes of data on drive 1, the next 16384 bytes of data on drive 2, the next on drive 3, and so on.

The adapter supports stripe sizes of 16K (16384), 32K (32768) and 64K (65536).

Note: The stripe size parameter can only be changed if there are no disk arrays associated with the adapter.

Reconstruction Rate

Controls the rate of data reconstruction on RAID level 1, 5, and 5E disk arrays (data on a RAID level 0 array cannot be reconstructed). This parameter controls the amount of resources the adapter allocates for the reconstruction. This is a PCI 4-Channel Ultra3 SCSI RAID Adapter attribute and can only be modified before the first disk array is created. Low, medium and high rates are supported.

Note: The reconstruction rate parameter can only be changed if there are no disk arrays associated with the adapter.

Data Scrubbing

Determines if data scrubbing is enabled on the redundant RAID level (1, 5, and 5E) disk arrays on the adapter. Data scrubbing refers to an automatic consistency check performed by the adapter on a small number of data stripes at regular periodic intervals. This feature monitors the accessibility and consistency of the data on redundant RAID level arrays and can identify problems in RAID arrays independent of any host access to the data contained on those arrays. Data scrubbing is an attribute of the PCI 4-Channel Ultra3 SCSI RAID Adapter and its value can only be modified when no disk arrays exist on the adapter.

Note: The data scrubbing parameter can only be changed if there are no disk arrays associated with the adapter.

Select Drives: Channel SCSI ID

Defines the drives included in the disk array. Individual drives are identified by channel number and SCSI ID

The RAID level of the disk array sets some restrictions on drive selection:

- 0** Each disk array is allowed 1 to 16 drives.
- 1** Each disk array is allowed 2 to 16 drives.
- 5** Each disk array is allowed 3 to 16 drives.
- 5E** Each disk array is allowed 4 to 16 drives.

Note: The drives you select for a disk array depend on your storage and performance needs. Disperse drives among separate SCSI channels, as much as possible.

Command Queue Depth

Allows the user to customize the depth of the adapter command queue utilized for a particular disk array. Supported queue depths range from 8 to 126 (default value for this parameter is 8). Customizing this parameter for your system may improve performance.

Read Ahead Enable

Enables and disables the adapter's read ahead data caching. When enabled, the adapter always reads to the end of a stripe on read operations.

Identify Drives and Disk Arrays

There are times when you need to determine the AIX logical name of a disk array, which volume groups your disk array are in, or which logical volumes are on a given disk array. For example, to back up data after a drive in an array fails, you need to know the drives and logical volumes or file systems affected by that failure.

The following table explains how to make a connection between physical disks, disk arrays, and AIX.

Task	Tools and Procedures
Find the physical disk(s) that make up a disk array.	The PCI Disk Array Manager (PDAM): The List all PCI SCSI RAID Drives option in the Change/Show PCI SCSI RAID Drive Status displays all disk arrays as well as the SCSI disks that make up the disk array.
Determine the logical volumes and file systems on a disk array.	SMIT or the lspv command: From the console, enter lspv -l <i>hdisk</i> , where <i>hdisk</i> is the disk name of the disk array you want to check. The resulting display shows all the logical volumes and file systems on the disk array.
Determine which disk arrays are in what volume groups.	Either SMIT or the lspv command: From the console, enter lspv . The resulting display shows the disk contents of all the volume groups on your system.

Disk Array Device Names and Location Codes

Disk arrays are assigned names using the *hdisk* form, the same as any other disk storage unit in AIX. These names are automatically assigned whenever you create a disk array with the PDAM. The names are deleted when you delete the disk array. You can display these names and the location codes associated with them by using the **List IBM PCI SCSI Disk Arrays** option using PDAM.

The location codes used for disk arrays differ slightly from those used by most other disks in this operating system. In addition to identifying the SCSI controller, as all AIX location codes do (in the second from the last digit), a disk array location code also contains the disk array number (in the last digit of the location code) allocated by the PCI 4-Channel Ultra3 SCSI RAID Adapter. For example, the standard AIX Version 4 location code has the following format:

00-02-01-5,0

Each field helps identify the location of the disk controller. The second digit of the second field (2) identifies the slot location of the adapter. The fourth field (5) identifies the id of the disk array. On most SCSI disks this field refers to the SCSI id. Disk arrays are typically made up of several physical SCSI disks with several SCSI ids. For the disk array, this field refers to the index into the adapters list of disk arrays. The last field (0) identifies the logical unit number of the disk array. This field is always 0 for disk arrays.

For example, you may see a display similar to the following when you select the **List IBM PCI SCSI Disk Arrays** option using PDAM.

In the following sample display, there are four disk arrays (disk array 00 through 03) attached to one PCI 4-Channel Ultra3 SCSI RAID Adapter (slot 1).

```
hdisk1  Available Raid 5 04-01-00-0,0 2064 MB Status OPTIMAL
        hdisk1  11 Channel 1 ID 1  ONLINE
        hdisk1  22 Channel 2 ID 2  ONLINE
        hdisk1  33 Channel 3 ID 3  ONLINE

hdisk2  Available Raid 1 04-01-00-1,0 2048 MB Status OPTIMAL
        hdisk2  12 Channel 1 ID 2  ONLINE
        hdisk2  20 Channel 2 ID 0  ONLINE

hdisk3  Available Raid 0 04-01-00-2,0 3072 MB Status OPTIMAL
        hdisk3  13 Channel 1 ID 3  ONLINE
        hdisk3  21 Channel 2 ID 1  ONLINE
        hdisk3  30 Channel 3 ID 0  ONLINE

hdisk4  Defined Raid 0 04-01-00-3,0 2144 MB Status DEAD
        hdisk4  10 Channel 1 ID 0  FAILED DRIVE
```

Disk Array ODM Information

The operating system stores information about each device that can be connected to it in a database maintained by the Object Data Manager (ODM). The information stored about each device includes the following:

- Device class, subclass, and type
- Device parent
- Device location code
- Connection point
- Attributes.

The following sections describe the values stored for each of these categories for the PCI 4-Channel Ultra3 SCSI RAID Adapter.

Device Class, Subclass, and Type

A device's class, subclass, and type unambiguously identify it to the operating system. With this information, the operating system can locate the device's record in the Object Data Manager (ODM) and obtain other information stored in the database about the device.

Together, a device's class, subclass, and type make up the unique type. The PCI 4-Channel Ultra3 SCSI RAID Adapter unique type is **adapter/pci/14102e00**.

Device Class

Describes the general category of devices for a device. For example, all printers are in the *printer class*. The PCI 4-Channel Ultra3 SCSI RAID Adapter is in the *adapter class* device class.

Device Subclass

Describes how the device is connected to the host. This is usually the type of adapter used to communicate between the host and the device. For example, standard SCSI disks are in the *scsi subclass*. The PCI 4-Channel Ultra3 SCSI RAID Adapter is in the *pci subclass* because it uses a PCI bus to communicate to the host.

Device Type

Describes the characteristics of a device within a class. For example, for a standard SCSI disk, the type can be a description of its storage capacity, such as 2gb. To uniquely identify the PCI 4-Channel Ultra3 SCSI RAID Adapter within its class, the PCI configuration signature (device id/vendor id) is used. The PCI 4-Channel Ultra3 SCSI RAID Adapter signature is 0x14102e00.

Device Parent

A device parent is the device that must be configured or defined before the device itself can be configured or defined. A device can have only one parent. For example, the parent device of a standard SCSI disk is the SCSI adapter, because the system cannot communicate with the disk drive without going through the adapter. The parent device is specified by a system assigned name, such as *scsi0*, *scsi1*, or *vscsi*. The PCI bus attached to the adapter is the device parent of the PCI 4-Channel Ultra3 SCSI RAID Adapter. This is normally *bus0* or *bus1*.

Device Location Code

On AIX systems, all devices have device location codes that uniquely identify them to the system. The PCI 4-Channel Ultra3 SCSI RAID Adapter location code conforms to existing PCI adapters on the AIX platform. The location code is of the form: 04-0N where *N* is the slot number. Disk arrays attached to the PCI 4-Channel Ultra3 SCSI RAID Adapter follow the location code convention of existing SCSI disks and are of the form: 04-0N-00-00. The last “nibble” indicates the device ID and LUN. For a PCI 4-Channel Ultra3 SCSI RAID Adapter attached disk array, the LUN field will always be zero and the ID field will be the index into the adapter’s list of disk arrays. These location codes follow AIX convention for adapters and devices. See *AIX Version 4 System Management Guide: Operating System and Devices* for a thorough description of all of the location code fields.

PCI 4-Channel Ultra3 SCSI RAID Adapter AIX Configuration

As with any device connected to an AIX system, the PCI 4-Channel Ultra3 SCSI RAID Adapter must be defined and configured before it can be accessed.

- A device is defined when the system creates a record for it in the Customized Database portion of the Object Data Manager (ODM). The ODM database contains information about all the devices that can be connected to the system. In the ODM database, device information is stored in an object relationship. The ODM database stores generic descriptions of the various supported devices in its Predefined database. For each instance of a particular device type, the system creates a separate copy of the predefined record.
- A device is configured when all necessary device drivers are loaded into the kernel and all other components necessary for communication with the device are configured. When configuration is completed, the device is marked as available.

At boot time, or whenever you invoke the **cfgmgr** command, the AIX operating system configures the PCI 4-Channel Ultra3 SCSI RAID Adapter and any previously created disk arrays automatically.

Chapter 2. Installing the PCI 4-Channel Ultra3 SCSI RAID Adapter Software

In order to correctly install the PCI 4-Channel Ultra3 SCSI RAID Adapter software, complete the following steps carefully. Failure to follow these steps may result in errors that will not allow further configuration of the PCI 4-Channel Ultra3 SCSI RAID Adapter.

This chapter contains information about:

- Installing the PCI 4-Channel Ultra3 SCSI RAID Adapter Software
- Installing a PCI 4-Channel Ultra3 SCSI RAID Adapter Software Update
- Verifying the PCI 4-Channel Ultra3 SCSI RAID Adapter Software Installation.

Installing an PCI 4-Channel Ultra3 SCSI RAID Adapter on AIX Version 4

This section describes the steps for installing the PCI 4-Channel Ultra3 SCSI RAID Adapter software for the first time on AIX Version 4. Note that it is often not necessary to install the PCI 4-Channel Ultra3 SCSI RAID Adapter software as a separate step. If the PCI 4-Channel Ultra3 SCSI RAID Adapter was installed in the system when the required level of AIX version 4 was installed, then all the necessary software for the adapter was installed.

Note: AIX Version 4.3.3 or later is necessary for PCI 4-Channel Ultra3 SCSI RAID Adapter

1. The PCI RAID Adapter requires AIX Version 4.3.3 or later.
2. Do not run the diagnostics on the PCI 4-Channel Ultra3 SCSI RAID Adapter at any stage during the installation of the PCI 4-Channel Ultra3 SCSI RAID Adapter software. Wait until installation is complete before running diagnostics. In most cases, it is unnecessary to run diagnostics following installation. See the section on PCI 4-Channel Ultra3 SCSI RAID Adapter Service Aids to determine if diagnostics are required.
3. Array controller microcode is distributed on the PCI 4-Channel Ultra3 SCSI RAID. Array controller microcode updates can be downloaded at:
<http://www.rs6000.ibm.com/support/micro/download.html>

Note: Downloading of array controller microcode is not usually required because those files are already installed on the adapter. Read and follow any instructions shipped with the PCI 4-Channel Ultra3 SCSI RAID Adapter for situations in which download is required.

Prerequisite

Ensure that the PCI 4-Channel Ultra3 SCSI RAID Adapter software is not already installed by using the command:

```
lspp -l devices.pci.14102e00.rte
```

If the software is not installed, the following messages is displayed:

```
lspp: 0504-132 Fileset devices.pci.14102e00.rte not installed
```

Procedure

1. Insert the install media into the machine.
At this point the software can be installed using SMIT or on the command line using the **cfgmgr** command.
To install the software using the **cfgmgr** command, the PCI 4-Channel Ultra3 SCSI RAID Adapter must be in the machine.
2. To install the software from the command line, type **cfgmgr -i devicename**

Note: **devicename** can be a directory or a device (ie. /dev/rmt0).

To install the software using SMIT, follow steps 3 through 7.

3. Run the **smit** command and select:
 - a. **Software Installation and Maintenance**
 - b. **Install and Update Software**
 - c. **Install and Update Software by Package Name**
4. Select **Input Device** using F4.
Select **Devices** to Install using F4 to display the list of install images on the tape or CD-ROM.
Then select **devices.pci.14102e00 ALL**. Press Enter to begin the installation process.
During the installation process, the screen will display messages while the PCI 4-Channel Ultra3 SCSI RAID Adapter software is loaded onto the system. If the messages indicate that the installation was not successful, repeat the installation procedure or contact your service representative.

5. The screen will look similar to the following:

Install and Update Software by Package Name (includes devices and printers)

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

[Entry Fields]

```
* INPUT device / directory for software      .
* SOFTWARE to install                        [devices.pci.14102e00 ]
PREVIEW only? (install operation will NOT occur) no
COMMIT software updates?                    yes
SAVE replaced files?                        no
AUTOMATICALLY install requisite software?   yes
EXTEND file systems if space needed?        yes
OVERWRITE same or newer versions?          no
VERIFY install and check file sizes?        no
Include corresponding LANGUAGE filesets     yes
DETAILED output?                            no
Process multiple volumes?                   Yes
```

```
F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Cmand       F7=Edit        F8=Image
F9=Shell     F10=Exit      Enter=Do
```

6. To make the PCI 4-Channel Ultra3 SCSI RAID Adapter available, type **cfgmgr** at the command line and press Enter. Alternatively, you may shutdown the system and restart it.
7. Verify the PCI 4-Channel Ultra3 SCSI RAID Adapter software and controller connections (see “Verifying the PCI 4-Channel Ultra3 SCSI RAID Adapter Software Installation” on page 25).
8. Ensure that all PCI 4-Channel Ultra3 SCSI RAID Adapters are in the available state.

Note: If the instructions shipped with the PCI 4-Channel Ultra3 SCSI RAID Adapter require downloading of microcode, those actions will begin at this point.

The installation of the PCI 4-Channel Ultra3 SCSI RAID Adapter hardware and software is now complete.

Installing a PCI 4-Channel Ultra3 SCSI RAID Adapter Software Update

This checklist describes the steps required to update the PCI 4-Channel Ultra3 SCSI RAID Adapter software on a host system that is already running the PCI 4-Channel Ultra3 SCSI RAID Adapter software.

Procedure

1. Insert the new PCI 4-Channel Ultra3 SCSI RAID Adapter software media into the media device.
2. From SMIT, select:
 - a. **Software Installation and Maintenance**
 - b. **Install and Update Software**
 - c. **Install Software by Fix (APAR)**
3. Select **Input Device** using F4.
4. Select Apar number that you would like to install. Toggle COMMIT software to *no* and SAVE replaced files to *yes*
5. Press Enter to begin the installation process.

The screen will look similar to the following:

Update Software by Fix

Type or select values in entry fields.

Press Enter AFTER making all desired changes.

[Entry Fields]

```
* INPUT device / directory for software      .
* FIXES to install                          [IY05058]
PREVIEW only? (install operation will NOT occur) no
COMMIT software updates?                    yes
SAVE replaced files?                        no
EXTEND file systems if space needed?        yes
VERIFY install and check file sizes?        no
DETAILED output?                           no
Process multiple volumes                    yes
```

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Cmand	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

6. Press Enter to begin the installation process.

During the installation process, the screen will display messages while the PCI 4-Channel Ultra3 SCSI RAID Adapter software is loaded onto the system. If the messages indicate that the installation was not successful, repeat the installation procedure or contact your service representative
7. Type the `cfgmgr` command to make the PCI 4-Channel Ultra3 SCSI RAID Adapter available and press Enter. Alternatively, you may shutdown and restart the system.
8. Verify the PCI 4-Channel Ultra3 SCSI RAID Adapter software and controller connections.

Verifying the PCI 4-Channel Ultra3 SCSI RAID Adapter Software Installation

1. For AIX Version 4.3.3 or later, verify the PCI 4-Channel Ultra3 SCSI RAID Adapter software installation by entering the following command:

```
lspp -l devices.pci.14102e00.rte
```

The displayed text should indicate that the PCI 4-Channel Ultra3 SCSI RAID Adapter Software is in the Committed state.

2. After you have verified the installation, go to “Chapter 3. Managing the PCI 4-Channel Ultra3 SCSI RAID Adapter” on page 27.

Chapter 3. Managing the PCI 4-Channel Ultra3 SCSI RAID Adapter

You can manage the PCI 4-Channel Ultra3 SCSI RAID Adapter by using the PCI SCSI Disk Array Manager (PDAM), System Management Interface Tool (SMIT), or for some tasks, the AIX command line. In most instances, disk arrays are managed by using SMIT.

Procedure

1. To start the PDAM, take the following path through SMIT:

- a. Run the **smit** command.
- b. Select **Devices**
- c. Select **Disk Array**
- d. Select **IBM PCI SCSI Disk Array**

From this panel you can change attributes of the PCI 4-Channel Ultra3 SCSI RAID Adapter, configure the adapter, or start the PDAM:

IBM PCI SCSI Disk Array

PCI SCSI Disk Array Manager
Configure a Defined PCI SCSI RAID Adapter
Change/Show PCI SCSI RAID Adapter

2. To change adapter parameters, select the **Change/Show PCI SCSI RAID Adapter** option of the PDAM panel.
3. Select the PCI 4-Channel Ultra3 SCSI RAID Adapter you want to work with and press Enter. The following panel will be displayed:

Change/Show PCI SCSI RAID Adapter

PCI SCSI RAID Adapter	scraid1
Battery Backup Enabled	no
Stripe Size (in KBytes)	16
Reconstruction Rate	med
Data Scrubbing Enabled	yes

4. Use the arrow keys to select the field to be altered. Change the values and press Enter.

Stripe Size

This parameter specifies the amount of data that will be accessed (striped) across a single physical disk in a disk array before accessing data on the next disk in the disk array. Supported stripe sizes are 16K, 32K and 64K.

Battery Backup

You should change this to yes **only** if the system is equipped with an Uninterrupted Power Supply (UPS). This parameter may be changed at any time, but will not take effect until the next system reboot.

Reconstruction Rate

This parameter specifies the rate at which a drive reconstruction (data rebuild) will occur. If the rate is set to low, the reconstruction will be performed at a slower rate but overall system performance should not be impacted during reconstructs. The opposite is true if the rate is set to high.

Data Scrubbing

This parameter specifies whether data scrubbing will be enabled on redundant level disk arrays created on this adapter. A value of yes indicates that a regular periodic interval, a consistency check will be performed by the adapter on a small number of data stripes.

Note: The Stripe Size, Reconstruction Rate, and Data Scrubbing parameters may only be changed if **no** disk arrays are currently configured on this PCI 4-Channel Ultra3 SCSI RAID Adapter.

Starting the PCI SCSI Disk Array Manager

1. To start the PDAM, take the following SMIT fast path:
`smit pdam`
2. From the PDAM panel, select the PCI SCSI Disk Array Manager option. The screen then displays the PCI SCSI Disk Array Manager panel.

PCI SCSI Disk Array Manager

List PCI SCSI Disk Arrays
Create a PCI SCSI Disk Array
Delete a PCI SCSI Disk Array
Configure a Defined PCI SCSI Disk Array
Change/Show a PCI SCSI Disk Array
Reconstruct a PCI SCSI Disk Array
Revive a FAILED Drive in a PCI SCSI Disk Array
Fail a Drive in a PCI SCSI Disk Array
Change/Show PCI SCSI RAID Drive Status
Perform Consistency Check
Display Status of Adapter Write Cache
Recovery Options

Creating a Disk Array

Use the following procedure to create or recreate a disk array:

- 1. Select the **Create a PCI SCSI Disk Array** option.
- 2. When prompted, select a **PCI SCSI RAID Adapter**.
- 3. When prompted, select the **RAID level**.
- 4. When prompted, select each drive you want in the disk array.

The following considerations apply when assigning drives:

RAID LEVEL

- 0** Each disk array is allowed 1 to 16 physical SCSI disk drives.
- 1** Each disk array is allowed 2 to 16 drives.
- 5** Each disk array is allowed 3 to 16 drives.
- 5E** Each disk array is allowed 4 to 16 drives.

Note: The drives you select for a disk array depend on your storage and performance needs.

- 5. After selecting the drives for the disk array, the **Create a PCI SCSI Disk Array** panel containing the disk array parameters is displayed. Use the Down Arrow or Enter key to select the value you want to change, and then enter the new value.

Note: You cannot change the RAID Level parameter or selected drives without restarting the creation procedure.

An example Create a PCI SCSI Disk Array SMIT panel follows

Create a PCI SCSI Disk Array	
	Entry Fields]
PCI SCSI RAID Adapter	scraid0
RAID Level	0
Stripe Size (KBytes)	64
Size of Array (MBytes)	1024
Select Drives: Channel-SCSI ID	01
Command Queue Depth	8
Read Ahead Enabled	yes
Initialize Parity	yes

Refer to the following information to determine the values you may use for the disk array parameters:

RAID Level

This parameter was selected on a previous panel and cannot be changed.

Stripe Size

This parameter can be changed but not from this panel. See "Change/Show PCI SCSI Disk Array Parameters" on page 38.

Size of Array

This parameter is calculated as the maximum size for the drives and the RAID level specified and may not be changed from this panel.

Command Queue Depth

This parameter can be changed from this panel. This parameter allows you to customize the depth of the adapter command queue for this disk array.

Read Ahead Enabled

This parameter can be changed. See “Change/Show PCI SCSI Disk Array Parameters” on page 38.

Initialize Parity

If parity is not initialized at disk array creation, data reconstruction capabilities may be compromised.

Select Drives: Channel-SCSI ID

These are the physical drives that were selected to be part of the disk array.

Note: The AIX Logical Volume Manager has a limit of 1016 physical partitions per physical volume. For the default physical partition size of 4 MB, this translates to a maximum physical volume size of approximately 4 GB. Larger physical volumes can be configured by increasing the physical partition size. You will know you have exceeded the 1016 physical partitions per physical volume (hdisk) if you receive this message while trying to create a volume group on a disk array:

```
0516-862 mkvg: Unable to create volume group
```

6. After you have set all the parameters you want to set, press Enter.

After you press Enter, if the Initialize Parity parameter was set to yes, the PCI 4-Channel Ultra3 SCSI RAID Adapter automatically initializes the parity/mirror on the new disk array in the background. The disk array is available for use while the parity/mirror is being initialized. Progress of the parity/mirror initialization can be monitored by using the List PCI SCSI Disk Arrays option of PDAM.

Note: Parity/mirror initialization is not applicable to RAID level 0 disk arrays.

7. If you need to configure more disk arrays, return to step 1. If you do not need to configure more disk arrays, exit PDAM.
8. Create AIX file systems on the disk arrays. See “Adding Disk Arrays to the AIX Operating System” on page 41 for more information.

Modifying and Displaying Drive Status

Using the Change/Show PCI SCSI RAID Drive Status option, you can display the status of the physical drives and disk arrays, delete a spare drive, add a spare drive, add a hot spare drive, identify a drive, remove a failed drive, or display a physical drive's Vital Product Data (VPD).

Change/Show PCI SCSI RAID Drive Status

Move cursor to desired item and press Enter

```
List all PCI SCSI RAID Drives
Delete a Spare Drive
Add a Spare Drive
Add a Hot Spare Drive
Identify a Drive
Remove a FAILED Drive
Display Vital Product Data
```

F1=Help	F2=Refresh	F3=Cancel	F8=Image
F9=Shell	F10=Exit	Enter=Do	

There are five possible values that may be displayed for the status of a disk array and seven possible values that may be displayed for the status of a physical drive.

Disk Array States

The state of a disk array is uniquely determined by the states of the physical drives that make up the disk array. Depending on the number of Failed physical drives, replacement of a Failed drive with a good drive may change the state of the disk array from Dead to Degraded or Optimal after data rebuild. When a Failed physical drive is replaced in a Degraded disk array, the data is rebuilt onto the replaced drive before the disk array state transitions to the Optimal. The figure below depicts the valid disk array state transitions.

Optimal

All physical drives in the disk array are Online.

Compressing

This is a special state for RAID level 5E disk arrays only. This state indicates that one physical drive in the array is in the Failed state. The data from the failed drive is compressed onto the Hot-spare by being distributed across the spare space of the remaining drives. When compression is complete the disk array indicates an Optimal state.

Decompressing

This is a special state for RAID level 5E disk arrays only. Decompressing occurs on compressed RAID level 5E disk arrays. A RAID level 5E array is compressed when the disk array is Optimal and there is one physical drive in the Failed state. When a Hot-spare is added to the disk array, compressed

data on the disk array is decompressed onto the newly added Hot-spare. When decompressing completes the disk array will be Optimal with all physical arrays in the Online state.

Degraded

One of the physical drives that is part of a redundant disk array is in the Failed state. For RAID level 5E disk array, one of the physical drives in a compressed disk array is in the Failed state.

Dead

For a non-redundant disk array, this state indicates that one or more physical drives are in the Failed state. For a redundant disk array, this state indicates that two or more physical drives are in the Failed state. For RAID level 5E disk arrays, this state indicates that three or more physical drives are in the Failed state. A Dead disk array is not accessible.

Physical Disk States

Online	The drive is part of a disk array and is functioning properly.
Spare	The drive is connected to and spun up by the PCI 4-Channel Ultra3 SCSI RAID Adapter but not configured into a disk array. This drive is available for the creation of a disk array.
Failed	<p>The drive was Failed by the PCI 4-Channel Ultra3 SCSI RAID Adapter or the user and must be replaced. A drive may go to the Failed state if one of the following conditions occur:</p> <ul style="list-style-type: none"> • Drive does not respond to selection. • Drive failed to spin-up. • Fail in inquiry or read capacity. • Fail to read or write reserve area. • Drive fails to respond to SCSI commands. • Inquiry, capacity, serial number and SCSI ID do not match configuration data stored on the adapter NVRAM. • User failed the drive using PDAM.
Reconstruct	The drive is undergoing a data reconstruction operation.
Warning	<p>The drive has been put into a Warning state as the result of Preventative Failure Analysis (PFA) error reported by the disk. The severity of this status depends on the RAID level of the disk array.</p> <p>In all cases, the Warning drive should be replaced as soon as possible.</p>
Hot Spare	The drive automatically replaces an equivalent Failed drive in a redundant disk array.
Non Existent	A drive is marked as non existent if no device is detected at a particular channel/SCSI ID location.

When the **Change/Show Drive Status** option is used to change the status of a drive, the state or status of the drive is changed. The following table describes the various actions performed and the new status of the drive after the action is complete.

Modifying Drive Status	
How To	Status Becomes
Add a Spare drive	Spare
Delete a Spare/Failed drive	Non Existent
Add a Hot-spare	Hot-spare

View the Current Disk Array Status

To see the current status of the drive, select the **List All PCI SCSI Disk Arrays** option.

A display similar to the following appears:

```

hdisk1 Available Raid 5 04-01-00-0,0 2064 MB Status OPTIMAL
    hdisk1 11 Channel 1 ID 1 ONLINE
    hdisk1 22 Channel 2 ID 2 ONLINE
    hdisk1 33 Channel 3 ID 3 ONLINE

hdisk2 Available Raid 1 04-01-00-1,0 2048 MB Status OPTIMAL
    hdisk2 12 Channel 1 ID 2 ONLINE
    hdisk2 20 Channel 2 ID 0 ONLINE

hdisk3 Available Raid 0 04-01-00-2,0 3072 MB Status OPTIMAL
    hdisk3 13 Channel 1 ID 3 ONLINE
    hdisk3 21 Channel 2 ID 1 ONLINE
    hdisk3 30 Channel 3 ID 0 ONLINE

```

Add a Spare Drive

This option is used to have the PCI 4-Channel Ultra3 SCSI RAID Adapter recognize a physical drive that has been added to the system subsequent to the last reboot (hot-plugged). If the drive was attached to the adapter before the last reboot, no user intervention is required as the adapter will recognize the new device during the boot process. Adding a drive through software causes the PCI 4-Channel Ultra3 SCSI RAID Adapter to change the drive status from Non-Existent to Spare, thus making the drive available to create a disk array.

1. Attach a physical drive with a unique SCSI ID to one of the available SCSI channels on the PCI 4-Channel Ultra3 SCSI RAID Adapter.
2. Select the **Change/Show PCI SCSI RAID Drive Status** option of the PCI SCSI Disk Array Manager panel. Each disk array and its associated disks are displayed.
3. Select the **Add a Spare Drive** option.
4. Select the drive with the Non-Existent status which is located at the channel/SCSI ID location desired and press Enter.

The drive state should now be Spare and the drive can now be used to create a disk array.

Delete a Spare or Failed Drive

A drive must be deleted if it is to be removed from the PCI 4-Channel Ultra3 SCSI RAID Adapter and not subsequently replaced. Deleting a drive through software causes the PCI 4-Channel Ultra3 SCSI RAID Adapter to change the drive status from Spare or Failed to Non-Existent. You cannot delete a drive that is part of a disk array. See “Deleting Disk Arrays from the Operating System” on page 43 to delete the disk array. Next, delete the drive and physically remove the drive from the system.

1. Select the **Change/Show PCI SCSI RAID Drive Status** option of the PCI SCSI Disk Array Manager panel. All disks, Ultra SCSI PCI RAID Adapters, and associated disks are displayed.
2. Select the **Remove a Spare Drive or Remove a Failed Drive** option of the PCI SCSI Disk Array Manager panel. For the **Remove a Spare Drive** option, all spare drives are displayed. For the **Remove a Failed Drive** option, all Failed drives are displayed.
3. Select the drive to be deleted and press Enter.

Fail a Drive

Attention: Do not fail a drive in a RAID level 0 disk array unless you want to replace the drive. RAID level 0 has no redundancy. Once a drive is Failed in a RAID level 0 disk array, it is unlikely that the data can be recovered. Do not fail a drive in a RAID level 1, 5, or 5E disk array if the disk array is already degraded. See “Restoring and Recovering Disk Arrays” on page 53 for instructions on which drives to fail and when to fail them.

You can place a drive in the Failed state when you want to replace a drive in Warning state and restore or reconstruct a disk array. Failing a drive through software causes the PCI 4-Channel Ultra3 SCSI RAID Adapter to change the drive status from Warning to Failed. Once a drive is Failed, the PCI 4-Channel Ultra3 SCSI RAID Adapter cannot access that drive's data until you reconstruct the drive data.

1. Select the **Fail a Drive in a PCI SCSI Disk Array** option. Each disk array and its associated disks are displayed.
2. Select the drive to be Failed and press enter.

Attention: If you fail a drive, you may lose data redundancy or data.

The status of the selected drive changes from Online/Warning to Failed

Repeat steps 1 and 2 until you have failed all the drives you want to fail.

Attention: Do not fail more than one Online drive in a single disk array. Failing two drives with this procedure could result in the loss of all data on the disk array.

Note: A drive that has been Failed manually using the Fail a Drive in a PCI SCSI Disk Array option will be called out by Diagnostics as needing replacement. This is not necessarily true as the drive may have been functioning properly when the drive was Failed. See “Reviving a Failed Drive” on page 36.

Removing a Disk Array

To change the RAID level, drive selection, or size of an existing disk array, see “Delete a PCI SCSI Disk Array”.

Delete a PCI SCSI Disk Array

Delete the disk array first in order to complete any of the following tasks:

- Change the RAID level of a disk array
- Change the drives that make up a disk array
- Change the stripe size of the adapter

If the disk array is part of a volume group, you must first back up all files in the logical volumes and file systems on the disk array and remove the disk array from its volume group. See “Deleting Disk Arrays from the Operating System” on page 43.

1. Start SMIT by entering:

```
smit pdam
```

2. Select the **Delete a PCI SCSI Disk Array** option. A display similar to the following appears:

Move cursor to desired item and press Enter.

```
hdisk1 Available Raid 5 04-01-00-0,0 2064 MB Status OPTIMAL
      hdisk1 11 Channel 1 ID 1  ONLINE
      hdisk1 22 Channel 2 ID 2  ONLINE
      hdisk1 33 Channel 3 ID 3  ONLINE
```

```
hdisk2 Available Raid 1 04-01-00-1,0 2048 MB Status OPTIMAL
      hdisk2 12 Channel 1 ID 2  ONLINE
      hdisk2 20 Channel 2 ID 0  ONLINE
```

```
hdisk3 Available Raid 0 04-01-00-2,0 3072 MB Status OPTIMAL
      hdisk3 13 Channel 1 ID 3  ONLINE
      hdisk3 21 Channel 2 ID 1  ONLINE
      hdisk3 30 Channel 3 ID 0  ONLINE
```

3. Select the disk array you want to delete and press Enter.

When the disk array has been deleted, the physical drives may change to one of the following states:

- Online drives go to Spare
- Warning drives go to Spare
- Failed drives remain Failed

Reviving a Failed Drive

Attention: Caution should be used when using the **Revive a Failed Drive** option as it may lead to corrupted data. As an example, consider a RAID level 0 disk array that is in the Dead state because one physical drive went to the Failed state. If the Failed physical drive was replaced with a good drive and the **Revive a Failed Drive** option was used, the status of the disk array would change to Optimal because all of the drives in the disk array are online. The disk array is now accessible but the data on the replaced drive is not valid. The proper recovery in this example would have been to delete and recreate the disk array (after replacing the Failed drive) and restoring the data from the backup copy. The Revive a Failed Drive option cannot be used on a failed drive which is part of a Degraded RAID array. A Reconstruct operation must be performed if you wish to revive a failed drive which is part of a Degraded RAID array.

Assigning Hot Spare Drives

Hot Spare drives are used to automatically replace Failed drives with drives of equivalent capacity in a redundant RAID environment.

1. Select the **Change/Show PCI SCSI RAID Drive Status** option.
2. Select the **Add a Hot Spare Drive** option. Select the PCI 4-Channel Ultra3 SCSI RAID Adapter you want to work with and press Enter.
3. Select the drive to be the hot spare and press Enter.

The drive state will transition to Hot Spare. On subsequent drive failures, reconstruction of Failed drives occur automatically for redundant disk arrays.

Transitioning a Hot Spare Drive to a Spare Drive

1. Select the **Change/Show PCI SCSI RAID Drive Status** option.
2. Select the **Add a Spare Drive** option.
3. Select the PCI 4-Channel Ultra3 SCSI RAID Adapter you want to work with and press Enter.
4. Select the hot spare drive you want to change to spare drive and press Enter.

The hot spare drive will transition to a spare drive.

Performing a Consistency Check on a Disk Array

The consistency check of a RAID level 5 or 5E disk array determines if the computed parity matches the actual parity on the disk array. Performing a consistency check on a RAID level 1 disk array compares the data contained on the primary copy of the data with the secondary copy of the data. If *no* is specified for the **Enable Automatic Consistency Repair** option, then an error will be returned if an inconsistency is found. At this point, the user can either restore the data using the most recent backup or rerun check consistency with the **Enable Automatic Consistency Repair** option set to *yes*. In the latter case, the PCI 4-Channel Ultra3 SCSI RAID Adapter attempts to repair the parity/mirror. Note that restoring consistency may result in loss of valid data for the blocks found to be inconsistent.

To perform a consistency check on a disk array, follow these steps:

1. Select the **Perform Consistency Check** option.
2. Select the disk array to perform the consistency check upon. Select yes/no for the **Enable Automatic Consistency Repair** option and press Enter.

Note: The **Check Consistency** option can only be run on disk arrays that are in the Optimal state.

Reconstructing a Failed Drive

You can reconstruct a drive after replacing a Failed drive in a degraded RAID level 1, 5, or 5E disk array.

Note: Only one reconstruct at a time is supported on the PCI 4-Channel Ultra3 SCSI RAID Adapter.

To reconstruct a Failed drive, follow these steps:

1. Replace the Failed drive with a working drive.
To transition this drive to the **Reconstructing** state and begin the rebuild:
2. Select the **Reconstruct a PCI SCSI Disk Array** option.
3. Select any Spare drive you wish to rebuild onto or leave the default value of the Failed drive if it has been replaced

Note: PDAM defaults to the Failed drive when selecting a drive to rebuild. The user can choose the default or any spare drive of equal (or greater) capacity on the adapter.

Change/Show PCI SCSI Disk Array Parameters

This section describes how to change the RAID level, drive selection, or size of array parameters on an existing disk array. If you want to change any of these values of a disk array, you must first back up the data in all logical volumes and file systems on the disk array, delete the logical volumes and file systems from the disk array, and then remove the disk array from its volume group before continuing. You must then delete the disk array and recreate using the new parameters.

Attention: Failure to back up data before deleting the disk array will result in a loss of data.

Use the following procedure to change the RAID level, drive selection, or size of the array of an existing disk array:

1. Back up the data if there is any data in the logical volumes or file systems on the disk array.
2. Unmount all file systems on the disk array you want to reconfigure.
3. Delete all logical volumes and file systems from the disk array you want to reconfigure, and remove the disk array from its volume group.
4. Delete the disk array. See “Removing a Disk Array” on page 35 for more information.
5. Recreate the disk array using the new parameters. See “Creating a Disk Array” on page 29 for more information.

Configuration Synchronization

Configuration synchronization makes the configuration on the PCI 4-Channel Ultra3 SCSI RAID Adapter consistent with the configuration on the disk arrays. This is useful in cases where the configurations are out of sync. For example, when an adapter is replaced but the user would like to continue to use the drives configured in the same way, configuration synchronization should be used.

1. Select the **Recovery Options** option.
2. Select **Resolve PCI SCSI RAID Adapter Configuration**.

From this panel you can select:

Display/Accept Configuration Changes
Accept the Configuration on Drives
Retry the Current Configuration

To synchronize to the configuration stored on the adapter, select the **Display/Accept Configuration Changes** option. To synchronize to the configuration on the drives (for example, adapter replacement), select the **Accept the Configuration on Drives** option.

The following options include:

Display/Accept Configuration Changes

If a PCI SCSI RAID configuration conflict is detected during a boot sequence, an error will be posted in the system error log. This option will allow you to view any drive state change and/or any unidentified drives within an existing Disk Array. Once you have viewed the changes, you may choose to accept them. Alternatively, you may modify the existing hardware to your needs and execute the Retry Current Configuration option.

Accept Configuration on Drives

Allows you to synchronize the PCI SCSI RAID configuration with the configuration currently stored on the majority of the physical drives. This would typically be done only in the case of a PCI SCSI RAID adapter replacement.

Note: Following an Accept Configuration on Drives option, the PCI 4-Channel Ultra3 SCSI RAID Adapter automatically performs a Consistency Check with Auto-Repair on all redundant level RAID arrays. Progress of the Consistency Check can be monitored using the List PCI SCSI Disk Arrays option of PDAM.

Retry Current Configuration

Allows you to perform a configuration reset that will cause the adapter to verify that all of the physical disk drives are detected and responding. If the adapter is unable to detect all of the physical disks which were previously attached to the adapter, an error is returned and a configuration conflict will exist.

Clearing a Configuration

Attention: This option should be used with extreme CAUTION. All disk arrays and data associated with the selected adapter may be *DESTROYED*! Selecting this option is equivalent to deleting all disk arrays configured on the adapter. This is an error recovery option normally used by your service representatives.

1. Select the **Recovery Options** option of the PCI SCSI Disk Array Manager panel.
2. Select **Clear PCI SCSI RAID Adapter Configuration**.
3. Select the adapter to clear the configuration on and press Enter.

Restoring Configuration on Replacement Adapter

Attention: If the PCI 4-Channel Ultra3 SCSI RAID Adapter being replaced experienced an abrupt failure, there may be valid data in the adapter's non volatile write cache.

If an existing disk array configuration will be imported from the attached drives, then the write cache module from the adapter being replaced should be used with the new adapter. This will ensure any data which may exist in the non volatile write cache module is correctly written to the attached disk drives.

A configuration conflict may result when a PCI 4-Channel Ultra3 SCSI RAID Adapter within a configured system is replaced (possibly due to a defective adapter). The configuration stored on the adapter may be out of sync with the configuration stored on the drives. The configuration can be restored by synchronizing the adapter configuration with the drive configuration. This can be accomplished using the **Accept Configuration on Drives** option under PDAM's Recovery Options. Use the following to replace the PCI 4-Channel Ultra3 SCSI RAID Adapter:

1. Start the SMIT PDAM by entering the following command:

```
smit pdam
```
2. Select the **Recovery Options** option.
3. Select **Resolve PCI SCSI RAID Adapter Configuration** option.
4. Select **Accept Configuration on Drives** option.
5. Select the adapter that was replaced.

Note: Following an Accept Configuration on Drives option, the PCI 4-Channel Ultra3 SCSI RAID Adapter automatically performs a Consistency Check with Auto-Repair on all redundant level RAID arrays. Progress of the Consistency Check can be monitored using the List PCI SCSI Disk Arrays option of PDAM.

Adding Disk Arrays to the AIX Operating System

After you configure a disk array using the PDAM, you can add the disk array to a volume group and create logical volumes and file systems on it before using it. Use standard AIX procedures to do this and treat the array in the same way you would treat a single disk drive. You can also access the disk array using raw IO.

Notes:

1. PCI 4-Channel Ultra3 SCSI RAID Adapter is not supported as a boot device.
2. PCI SCSI disk arrays do not support the root volume group.

Attention: The AIX Logical Volume Manager has a limit of 1016 physical partitions in a physical volume. When creating AIX volume groups, select a physical partition size sufficiently large so that you don't have more than 1016 physical partitions per physical volume. For example, if you are creating a volume group composed of a RAID 5 disk array composed of 5 physical disks of 2 GB each, use a physical partition size of 8 MB minimum; the default size of 4 MB will exceed the 1016 physical partition limit.

You will know you have exceeded the 1016 physical partitions per physical volume (hdisk) if you receive this message while trying to create a volume group on a disk array (hdisk):

```
0516-862 mkvg: Unable to create volume group
```

Do not mix hdisks with different RAID characteristics in the same volume group. If you mix RAID 0 and RAID 5 hdisks, and a file system has partitions on both hdisks, a drive failure in the RAID 0 disk array will cause you to lose all your data in the file system.

If performance is important, you should plan your volume groups and logical volume placement before you create them because recreating them and restoring your data can be time-consuming. Contact your point of sale to get performance planning assistance or more information.

The following procedure gives the basic steps required to create a file system on a configured disk array so that you can use it. This procedure makes the disk array its own volume group and creates a single logical volume or file system on it. To add the disk array to an existing volume group, extend a logical volume or file system onto it, or to change other file system options, see the procedures covered in your AIX system documentation.

1. From the SMIT System Management panel, select **Physical & Logical Storage**.
2. From the Physical & Logical Storage panel, select **Logical Volume Manager**.
3. From the Logical Volume Manager panel, select **Volume Group**.
4. Select the **Add A Volume Group** option from the Volume Group screen.
5. Enter values for the Add A Volume Group panel fields, as follows:

VOLUME GROUP name

Enter a volume group name.

PHYSICAL VOLUME names

Enter the hard disk name of the configured disk array. Press F4 to display a list of options.

Activate volume group AUTOMATICALLY at system restart?

Indicate whether you want the volume to be automatically activated whenever you restart the system.

ACTIVATE volume group after it is created?

Indicate whether you want the volume to be automatically activated when you are finished creating it (that is, when you exit SMIT).

Volume group MAJOR NUMBER

Give the volume a volume number or leave this option blank.

6. Press Enter to save these values. Then press the F3 key until you return to the Physical & Logical Storage panel.
7. Select the **File Systems** option from the Physical & Logical Storage panel.
8. Select the **Add / Change / Show / Delete File Systems** option from the File Systems panel.
9. Select the **Journaled File System** option from the Add / Change / Show / Delete File Systems panel. A list of volume groups is displayed.
10. Select **Add a Journaled File System** from the Journaled File System panel.
11. Select the volume group that you just created from the list of names displayed.
12. Enter values for the Journaled File System fields, as follows:

SIZE of file system

Enter the size, in number of 512-byte blocks, that you want to assign to the file system. Any size up to the limit of the volume group or 2 GB can be entered, although the size will be set to a multiple of the partition size.

MOUNT POINT Enter a mount point for the file system. This is the directory you will enter to access the file system.

Mount AUTOMATICALLY at system restart?

Select *yes* or *no*, depending on whether you want this file system mounted automatically after the system is started.

PERMISSIONS Leave this at **read/write**.

Mount OPTIONS

Enter the mount options or leave this field blank.

13. Press Enter to create the journaled file system.

After completing this procedure, you can mount and use the newly created file system, using the specified mount point.

Deleting Disk Arrays from the Operating System

Attention: Failure to back up any data before performing this procedure could result in the loss of all data on the disk arrays.

It is necessary to delete a disk array whenever you want to change or modify disk array parameters. The procedures required are basically the reverse of those used to add disk arrays.

For each disk array you want to delete, perform the following procedure:

1. Back up any data on the disk array. There may be more than one logical volume and file system on the drive. To find out which logical volumes or file systems are on the disk array, see “Identify Drives and Disk Arrays” on page 16.
2. Unmount any file systems on the disk array.

Delete any file systems on the disk array using the following steps:

1. Return to the Logical Volume Manager panel.
2. Select the **File Systems** option from the Logical Volume Manager panel.
3. Select the **Add / Change / Show / Delete File Systems** option from the File Systems panel.
4. Select the **Journaled File System** option from the Add / Change / Show / Delete File Systems panel. A list of volume groups is displayed.
5. Select **Delete a Journaled File System** from the Journaled File System panel.
6. Select the file system to be removed.

Delete any logical volumes on the disk array using the following steps:

1. From the SMIT System Management panel, select **Physical & Logical Storage**.
2. From the Physical & Logical Storage panel, select **Logical Volume Manager**.
3. From the Logical Volume Manager panel, select **Logical Volumes**.
4. Select the **Remove a Logical Volume** option from the Logical Volumes panel.
5. Select the logical volume to be removed.

Remove the disk array from its volume group using the following steps:

1. Return to the **Logical Volume Manager** panel.
2. Select the **Volumes Group** option from the Logical Volume Manager panel.
3. Select either the **Remove a Volume Group** or **Set Characteristics of a Volume Group** option, depending on whether there are other physical volumes in the volume group.

To find out which volume group the disk array is in, see “Identify Drives and Disk Arrays” on page 16. Then go to “Delete a PCI SCSI Disk Array” on page 35 to delete the disk array using PDAM.

Extending Logical Volumes and Disk Arrays

You can extend an existing logical volume onto a disk array just as you can with any other disk storage device. However, you should remember that the data reliability of a combined logical volume is only as great as the reliability of its weakest part. For example, if you combine a RAID level 5 disk array and a RAID level 0 disk array into the same logical volume, the resulting volume has the data redundancy protection of a RAID level 0 disk array (that is, none). Even though the RAID level 5 portion of the logical volume has data redundancy, if the RAID level 0 portion fails, the whole volume may become unusable, and all data on the volume will be lost. However, a RAID level 1 and a RAID level 5 logical volume would offer data redundancy.

Chapter 4. PCI 4-Channel Ultra3 SCSI RAID Adapter Problem Determination and Recovery

This chapter provides information to assist you in determining and recovering from error conditions. The following topics are discussed in detail:

- Identifying and recovering from drive failures. If a drive attached to the PCI 4-Channel Ultra3 SCSI RAID Adapter fails, procedures are provided by the PCI SCSI Disk Array Manager (PDAM) for removing and replacing the failed drive.
- Instructions on how to recover the data from the failed drive (in redundant environment) or restoring the data from the drive (on non-redundant environments) are included.
- Information describing the various RAID levels and how drive failures are handled by each of the levels.
- Instructions on how to replace a PCI 4-Channel Ultra3 SCSI RAID Adapter.
- Information describing how to detect and resolve array configuration conflicts.

You may also refer to the informational error messages returned by the PCI 4-Channel Ultra3 SCSI RAID Adapter device driver. These include alertable error messages and AIX error log entries in “Messages and Error Log Templates” on page 66.

Drive Failures and Disk Arrays

The PCI 4-Channel Ultra3 SCSI RAID Adapter handles drive failures differently, depending on the RAID level involved.

Drive Status

There are seven possible values that may be displayed for the status of a physical drive and five possible values that may be displayed for the status of a disk array.

Disk Array States

The state of a disk array is uniquely determined by the states of the physical drives that make up the disk array. Depending on the number of Failed physical drives, replacement of a Failed drive with a good drive may change the state of the disk array from Dead to Degraded or Optimal after data rebuild. When a Failed physical drive is replaced in a Degraded disk array, the data is rebuilt onto the replaced drive before the disk array state transitions to the Optimal.

Optimal The disk array is in good functional state.

Compressing This is a special state for RAID level 5E disk arrays only. This state indicates that one physical drive in the array is in the Failed state. The data from the failed drive is compressed onto the Hot-spare by being distributed across the spare space of the remaining drives. When compression is complete the disk array indicates an Optimal state.

Decompressing

This is a special state for RAID level 5E disk arrays only. Decompressing occurs on compressed RAID level 5E disk arrays. A RAID level 5E array is compressed when the disk array is Optimal and there is one physical drive in the Failed state. When a Hot-spare is added to the disk array, compressed data on the disk array is decompressed onto the newly added Hot-spare. When decompressing completes the disk array will be Optimal with all physical arrays in the Online state.

Degraded

One of the physical drives that is part of a disk array is in the Failed state.

Dead

For a non-redundant disk array, this state indicates that one or more physical drives are in the Failed state. For a redundant disk array, this state indicates that two or more physical drives are in the Failed state. A Dead disk array is not accessible.

Physical Disk States**Online**

The drive is part of a disk array and is in good functional state.

Spare

The drive is connected to the adapter, but not configured into a disk array. This drive is available to for the creation of a disk array.

Failed

The drive was Failed by the PCI 4-Channel Ultra3 SCSI RAID Adapter or by the user and must be replaced. A drive goes to Failed state when one of the following conditions occur:

- Drive does not respond to selection.
- Drive failed to spin-up.
- Fail in inquiry or read capacity.
- Fail to read or write reserve area.
- Drive fails to respond to SCSI commands.
- Inquiry, capacity, serial number and SCSI ID do not match configuration data stored on the adapter NVRAM.
- User failed the drive using PDAM.

Reconstruct

The drive is undergoing a data reconstruction operation.

Warning

The drive has set its Predictive Failure Analysis (PFA) indication. Although the drive may remain functional, this an indication that a serious failure of the drive is imminent. The severity of this status depends to some extent on the RAID level of the disk array.

The Warning drive should be replaced as soon as possible.

Hot Spare

The drive automatically replaces an equivalent Failed drive in a redundant disk array.

Non Existent

A drive location is marked as non existent if no disk was detected at that channel/SCSI ID.

Disk Array Failures

This section contains information on disk array drive failures:

- RAID Level 0
- RAID Level 1
- RAID Level 5
- RAID Level 5E

Raid Level 0

When read or write errors occur on one or more drives in a RAID level 0 disk array, the drive state changes from Online to Failed or Warning. Follow the steps below:

1. You should attempt to back up the files on all logical volumes and file systems on the disk array. If the error is not serious, you may be able to recover the data.
2. Regardless of whether the backup succeeds, use the PDAM to identify the disk drive on which the error occurred.
3. Delete all the logical volumes and file systems on the disk array.
4. Remove the disk array from the volume group.
5. Replace the Failed or Warning drive and use the PDAM to revive the replaced drive.
6. Add the disk array back to a volume group.
7. Recreate the logical volumes and file systems on it.
8. Copy data back to the restored disk array from your backup media.

Raid Level 1

This section describes drive failures on a RAID level 1 disk array, and the disk array and drive status values that result from such failures.

Single-Drive Failures: Whenever a single drive in a RAID level 1 disk array transitions to the Failed state, the disk array status changes to Degraded. The disk array remains functional because the data on the Failed drive can be accessed from the mirrored stripe.

Whenever a drive is marked as Failed, you should replace it as soon as possible. If a Hot-spare drive was defined, the adapter may reconstruct the data on the drive automatically. If a Hot-spare reconstruct was not initiated by the adapter, then using the PDAM, initiate a reconstruct on the failed drive following replacement. If you cannot replace the drive immediately, back up the files on all logical volumes and file systems on the disk array and replace the drive when you can.

Multiple-Drive Failures: Whenever the second drive in a RAID level 1 disk array has experienced read or write errors, the adapter takes the following actions:

- If the error is a recoverable read error, then the disk array status remains Degraded and an entry is posted in the error log.
- If the error is an unrecoverable error, and the drive is the mirrored pair of a drive that is already Failed, the drive status changes to Failed. The disk array status changes to Dead.

Note: Only one Failed drive per RAID level 1 disk array is supported on the PCI 4-Channel Ultra3 SCSI RAID Adapter.

Raid Level 5

This section describes drive failures on a RAID level 5 disk array, and the disk array and drive statuses that result from such failures.

Single-Drive Failure: Whenever a single drive in a RAID level 5 disk array transitions to the Failed state, the disk array status changes to Degraded. The disk array remains functional because the data on the Failed drive can be reconstructed using parity and data on the remaining drives.

Whenever a drive is marked as Failed, you should replace it as soon as possible. If a Hot-spare drive was defined, the adapter may reconstruct the data on the drive automatically. If a Hot-spare reconstruct was not initiated by the adapter, then using the PDAM, initiate a reconstruct on the failed drive following replacement. If you cannot replace the drive immediately, back up the files on all logical volumes and file systems on the disk array and replace the drive as soon possible.

Multiple-Drive Failures: If a second drive in a RAID level 5 disk array experiences read or write errors, the adapter takes the following actions:

- If the error is a recoverable read error, then the drive status remains Degraded and an entry is posted in the error log.
- If the errors are unrecoverable, the drive status changes to Failed. The disk array status changes to Dead.

If the disk array status is Degraded, you will be able to reconstruct the data on the disk array by replacing the Failed drive, and reconstructing the data on the drive. If the disk array status is Dead, you will have to replace the Failed drives and then delete and recreate the disk array (after deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group). You must then add the disk array back to a volume group, recreate the logical volumes and file systems on it, and copy data back to the restored disk array from your backup media.

Raid Level 5E

This section describes drive failures on a RAID level 5E disk array, and the disk array and drive statuses that result from such failures.

Single-Drive Failure: Whenever a single drive in a RAID level 5E disk array transitions to the Failed state, the disk array status changes to Compressing. After the drive has successfully compressed, the disk array status changes to Optimal. The disk array remains functional because the data from the Failed drive is compressed onto the Hot-spare space on the remaining drives.

Whenever a disk array has been compressed, you should decompress it as soon as possible. If a Hot-spare drive was defined, the adapter may decompress the data on the drive automatically. If a Hot-spare was not defined, then you should add a Hot-spare.

When the Hot-spare is added, the disk array begins decompressing the data onto the Hot-spare. When the disk array has successfully decompressed, the disk array status changes to Optimal.

Note: Whenever the disk array is compressed and one drive fails, then the disk array status changes to Degraded. The disk array remains functional because the data can be reconstructed using parity and data remaining on the drives. If a second drive fails before compression completes the disk array will have a status of Dead.

Multiple-Drive Failures: Two situations may occur for multiple drive failures.

- If a second drive fails while the disk array is Compressing then the disk array status will change to Dead. If three or more drives fail at any time the disk array status will change to Dead.
- If a second drive fails after the disk array has completed Compressing and has an Optimal status then the status will change to Degraded.

If the disk array status is Degraded, you will be able to reconstruct the data on the disk array by adding a Hot-spare drive to the array. If the disk array status is Dead, you will have to replace the Failed drives and then delete and recreate the disk array (after deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group). You must then add the disk array back to a volume group, recreate the logical volumes and file systems on it, and copy data back to the restored disk array from your backup media.

Reconstructing Redundant Disk Arrays after Drive Failure

Reconstruction is a process used to restore a degraded RAID level 1, 5, or 5E disk array to its Optimal state after a single drive has been replaced. During reconstruction, the adapter recalculates the data on the drive that was replaced, using data and parity from the other drives in the disk array. The controller then writes this data to the replaced drive. Although RAID level 1 does not have parity, the adapter can reconstruct data on a RAID level 1 disk array by copying data from the mirrored stripe.

If a Hot-spare has been defined, the adapter automatically initiates the reconstruction process after a drive status changes to Failed. If a Hot-spare of appropriate capacity had not previously been defined, the reconstruct must be initiated by the user after replacing the Failed drive. Once reconstruction is initiated (either by the user or the adapter), the adapter completes the following actions:

- Copies special array configuration information files to the new drive
- Recalculates the data and/or parity from the data and parity on the other disk array drives
- Writes the recalculated data and parity to the new drive.

The disk array remains accessible while the Reconstruction is in progress. However, in some cases (for example, if you have two Warning drives in a RAID level 5 disk array) you should stop using the logical volumes and file systems on the disk array to prevent subsequent drive failure and data loss.

Drive reconstruction may have a varying impact on system I/O performance. The disk array continues to operate in degraded mode. The reconstruction rate is a tunable attribute that must be selected before any disk arrays are configured on the adapter. See the section on Reconstruction Rate for additional information.

If there are multiple-drive failures in a RAID level 1 or 5 disk array (at least two Failed drives), you may not be able to reconstruct the drive data. To restore a RAID level 1 or 5 disk array to a working state in this case, you need to replace the drives and recreate the disk array (after deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group). You must then add the disk array back to a volume group, recreate the logical volumes and file systems on it, and copy data back to the restored disk array from your backup media.

Notes:

1. Only one reconstruct can be in progress at any given time on the PCI 4-Channel Ultra3 SCSI RAID Adapter. Concurrent reconstructs are not supported.
2. Allow reconstruction to complete before running diagnostics. Running diagnostics on the PCI 4-Channel Ultra3 SCSI RAID Adapter during a reconstruct kills the reconstruction. If a reconstruction fails, it may be restarted using PDAM.

The Reconstruction Rate - Low, Medium, High

The rate at which a reconstruction occurs depends on the value of the reconstruction rate parameter. This parameter determines how much resource the adapter allocates for reconstructions. When the adapter reconstructs data on a drive, it divides the adapter's internal resources between reconstruction and normal I/O operations. As reconstruction rate increases, system I/O performance and reconstruction time may decrease. These parameters control only the rate of data reconstruction, not total reconstruction time.

Change Reconstruction Rate Parameters

Use this procedure to change the Reconstruction Rate parameters.

Note: The **Reconstruction Rate** attribute may only be changed prior to the creation of any disk arrays on the adapter. Once a disk array has been configured on an adapter, the **Reconstruction Rate** cannot be changed. In order to modify the **Reconstruction Rate**, all disk arrays on the adapter must first be removed.

1. Start the SMIT PDAM:
`smit pdam`
2. Select the device (hdisk) name of the disk array you want.
3. Select the **Change/Show IBM PCI SCSI RAID Adapter** option.
4. Select an adapter when prompted.
5. Select the **Reconstruction Rate** option and enter the new value for this field.
6. Press Enter to make the changes and exit SMIT.

Checking Consistency (Check and Repair Array Parity)

Note: The **Check Consistency** option applies only to RAID levels 1, 5, and 5E. RAID level 0 is a non-redundant array and therefore cannot be checked and repaired. RAID level 1 does not have parity either, but a parity check compares data on the mirrored drives. In addition, the RAID level 1, 5, or 5E disk array must be at Optimal status in order to run check consistency.

Check consistency performs the following functions:

1. Scans the disk array and checks the array parity for each block in the disk array. On a RAID level 1 disk array, Parity Check compares the data on each mirror, stripe by stripe.
2. Repairs any array parity errors found during the parity check. On a RAID level 1 disk array, the adapter changes the data on the mirror stripe to make it match the data on the data disk. On a RAID level 5 and 5E disk arrays, the controller changes the parity so that it is consistent with the data.

The data is not repaired if array parity errors result from corrupted data. Only the array parity is repaired. Array parity is used in RAID Level 5 and 5E disk arrays to enable data to be reconstructed if a single drive fails. Checking and repairing parity helps ensure that you can recover after a drive failure.

Note: Upon creation of a disk array, if the parity/mirror is not initialized (Initialize Parity=no), then the **Check Consistency** option encounters inconsistencies. It is recommended that the parity/mirror be initialized at disk array creation. This applies to only RAID levels 1, 5, and 5E. RAID level 0 is a non-redundant array and therefore cannot be initialized.

Perform Check Consistency

Always maintain backup files, even with RAID level 1, 5, and 5E disk arrays.

1. Start the SMIT PCI SCSI Disk Array Manager
2. Select the **Perform Check Consistency** option.
3. Select the device (hdisk) name of the disk array you want to check and repair.

Note: Setting the **Enable Automatic Consistency Repair** option to *yes* enables the adapter to attempt to repair any inconsistencies encountered.

You must run Check Consistency on each disk array separately.

4. When all the entries are correct, press Enter to perform the check on the selected disk arrays.
5. Repeat steps 3 through 4 as needed to check and repair additional disk arrays.

Attention: It is assumed the data is good and the parity is bad if an inconsistency is found when repairing a RAID level 5 or 5E array. Always maintain backup files also for RAID level 1, 5, and 5E disk arrays. Backups are the most reliable method of restoring data.

Restoring and Recovering Disk Arrays

Disk arrays may be restored to a working state after drive failure using several methods. Procedures for disk array restoration are provided for each of the RAID levels because each level may have a unique recovery process. To restore a disk array, refer to “Restoring a Disk Array” on page 58.

Error recovery procedures are also provided for various disk array status values. Refer to the following information:

- “Delete a Spare or Failed Drive” on page 34
- “Add a Spare Drive” on page 33
- “Reviving a Failed Drive” on page 36
- “Reconstructing a Failed Drive” on page 37

Restoring a disk array is the process of returning the disk array to the Optimal state after a drive failure. The tasks involved depend on the RAID level of the disk array and on how many drives have Failed or have experienced errors:

- Single-drive failures in RAID levels 1, 5, and 5E are easily restored without losing data and without shutting down the disk array.
- Multiple-drive failures, or single-drive failures on RAID level 0, are treated like non-array drive failures in AIX. Delete the logical volumes, file systems, and volume groups on the disk array. Replace the drive, recreate the disk array and AIX file structures, and restore the data from the backup copy.

Note: Single-drive failures on redundant RAID Levels 1, 5, and 5E are recoverable. The user should identify the Failed drive, physically remove it from the subsystem, and replace it with another drive. The user may then initiate a reconstruct on the Failed drive. Alternatively the user may reconstruct the data using an available Spare drive. The current Failed drive remains Failed but no longer exists as part of the disk array. For single-drive failures on RAID level 5E disk arrays, the disk array will compress and the disk array status changes to Compressing. After a successful compression the disk array status will change to Optimal. A Hot-spare drive should be added to decompress the disk array. When the Hot-spare drive is added the adapter will automatically decompress the array. After a successful decompression, the current Failed drive remains Failed but no longer exists as part of the array. This task is intended to be performed by the system administrator, and involves no reorganization or deletion of file systems or volume groups.

You need to replace a drive when one of the following events occur:

- The AIX error log facility sends a message indicating a disk array component failure (a drive state transitioned to Failed).
- You see a disk array status other than Optimal displayed by the **List all PCI SCSI RAID Arrays** option.
- You see a drive status other than Online or Spare displayed by the **List all PCI SCSI RAID Arrays** option.

The recovery tasks are summarized in the RAID Recovery Tasks Summary tables that follow. See “Restoring and Recovering Disk Arrays” on page 53 for more information.

RAID Level 0 Recovery Tasks Summary	
Disk Array and Drive Status	Recovery Steps
Disk array status is Dead	<ol style="list-style-type: none">1. Delete the file systems and logical volumes on the affected disk array and remove the disk array from the volume groups (see “Deleting Disk Arrays from the Operating System” on page 43.2. Replace failed drives.3. Delete and recreate the disk array.4. Create new logical volumes and file systems on restored units (see “Adding Disk Arrays to the AIX Operating System” on page 415. Copy backed-up data to the restored logical volumes or file systems.

RAID Level 1 Recovery Tasks Summary	
Disk Array and Drive Status	Recovery Steps
Disk array status is Optimal, one or more drives have status of Warning.	<ol style="list-style-type: none"> 1. If the PDAM displays two or more drives within a disk array that have the Warning state, stop using all file systems and logical volumes on the affected disk arrays and attempt to back up all files in those logical volumes before continuing. 2. One at a time, fail each drive (using the PDAM) that is in the Warning state. Replace the Warning drive, and reconstruct the data on it.
Disk array status is Degraded. Drive status shows a Failed drive.	<ol style="list-style-type: none"> 1. Replace the Failed drive. 2. Using PDAM, reconstruct the data on the replaced drive.
Disk array status is Dead. Drive status shows two or more Failed drives.	<ol style="list-style-type: none"> 1. Delete the file systems and logical volumes on the affected disk arrays, and remove the disk arrays from their volume groups (see "Deleting Disk Arrays from the Operating System" on page 43). 2. Delete the Dead disk array. The Failed drives will remain Failed drives outside of a disk array. Delete the Failed drives from the system using PDAM, then physically replace the Failed drives. Add new drives back into the system as Spare drives using PDAM. Recreate the disk array. 3. Create new logical volumes and file systems on the restored units (see "Adding Disk Arrays to the AIX Operating System" on page 41). 4. Copy backed-up data to the restored file systems or logical volumes.

RAID Level 5 Recovery Tasks Summary	
Disk Array and Drive Status	Recovery Steps
Disk array status is Optimal.	<ol style="list-style-type: none"> 1. If the PDAM displays two or more drives within a disk array that have the Warning state, stop using all file systems and logical volumes on the affected disk arrays and attempt to back up all files in those logical volumes before continuing. 2. One at a time, fail each drive (using the PDAM) that is in the Warning state. Replace the Warning drive, and reconstruct the data on it.
Disk array status is Degraded. Drive status shows one Failed drive.	<ol style="list-style-type: none"> 1. Replace the Failed drive. 2. Reconstruct the data on the Failed drive by using the PDAM.
Disk array status is Dead. Drive status shows two or more Failed drives.	<ol style="list-style-type: none"> 1. Delete the file systems and logical volumes on the affected disk arrays, and remove the disk arrays from their volume groups (see "Deleting Disk Arrays from the Operating System" on page 43). 2. Delete the Dead disk array. Delete the Failed drives from the system using PDAM, then physically replace the Failed drives. Add the drives back into the system as Spare drives using PDAM. Recreate the disk array. 3. Create new logical volumes and file systems on the restored units (see "Adding Disk Arrays to the AIX Operating System" on page 41). 4. Copy backed-up data to the restored file systems or logical volumes.

RAID Level 5E Recovery Tasks Summary	
Disk Array and Drive Status	Recovery Steps
Disk array status is Optimal.	<ol style="list-style-type: none"> 1. If the PDAM displays two or more drives within a disk array that have the Warning state, stop using all file systems and logical volumes on the affected disk arrays and attempt to back up all files in those logical volumes before continuing. 2. Add a Hot-spare drive to the disk array for each drive in the Warning state. The Hot-spare drive must have a capacity greater than or equal to the smallest disk capacity in the disk array. 3. One at a time, fail each drive (using the PDAM) that is in the Warning state. The adapter will initiate the reconstruct of the data. When the reconstruct completes you may repeat this step for each Warning drive.
Disk array status is Degraded. Drive status shows two Failed drives.	<ol style="list-style-type: none"> 1. Add a Hot-spare drive to the disk array for each drive in the Warning state. The Hot-spare drive must have a capacity greater than or equal to the smallest disk capacity in the disk array. 2. The adapter may begin the reconstruct of data after the Hot-spare is added.
Disk array status is Dead. Drive status shows two or more Failed drives.	<ol style="list-style-type: none"> 1. Delete the file systems and logical volumes on the affected disk arrays, and remove the disk arrays from their volume groups (see "Deleting Disk Arrays from the Operating System" on page 43). 2. Delete the Dead disk array. Delete the Failed drives from the system using PDAM, then physically replace the Failed drives. Add the drives back into the system as Spare drives using PDAM. Recreate the disk array. 3. Create new logical volumes and file systems on the restored units (see "Adding Disk Arrays to the AIX Operating System" on page 41). 4. Copy backed-up data to the restored file systems or logical volumes.

Restoring a Disk Array

Use the following procedure to restore a disk array:

1. Determine the degree to which the drive fault has affected the disk array. This can be accomplished with the PDAM. This will report any drives that have had errors.
2. Start the SMIT PCI SCSI Disk Array Manager by entering the following command:

```
smit pdam
```
3. Select the **List all PCI SCSI Disk Arrays** option from the PCI disk array Manager menu. This will return the status of each disk array and the status of each drive within the disk array.
4. If a disk array status is Optimal, but one or more drives are in the Warning state, then the drive(s) has exceeded its soft error threshold. Go to “Replacing Failed or Warning Drives (RAID Levels 1, 5, and 5E)”.
5. If a disk array status is Degraded and the **List all PCI SCSI RAID ARRAY** option shows the Failed drive, go to “Replacing Failed or Warning Drives (RAID Levels 1, 5, and 5E)” and complete the procedure.
If a disk array status is Optimal and compressed, go to “Replacing Failed or Warning Drives (RAID Levels 1, 5, and 5E)” and complete the procedure.
6. If the disk array status is Dead, then go to “Replacing a Dead or Optimal Array” on page 59.

Replacing Failed or Warning Drives (RAID Levels 1, 5, and 5E)

Use the following procedure to replace a Failed or Warning drive in a RAID level 1, 5, or 5E disk array. If a drive in a disk array has a status of Failed, the disk array will have a status of Degraded. For RAID level 5E disk arrays, compressed disk arrays with one Failed drive will have a status of Degraded while compressed disk arrays with no Failed drives will have a status of Optimal. If one or more drives in a disk array have a status of Warning, the disk array status will remain Optimal.

1. Start the PDAM by entering the following command:

```
smit pdam
```
2. Select the **Change/Show PCI SCSI Drive Status** option.
3. Locate the Failed or Warning drive.
4. Replace the drive with a good drive of equal or greater capacity. For RAID level 5E disk arrays, add a Hot-spare drive with a good drive of equal or greater capacity.

Attention: If the replaced drive was a Warning drive, it is necessary to fail the drive before continuing to the next step. See the Fail a Drive section in Chapter 4. Do not fail more than one drive at a time, even if you need to replace more than one drive. Fail and reconstruct the drives one at a time.

Note: If a service request number (SRN) is needed for the device, it will be necessary to run hardware diagnostics. See the *Diagnostic Information for Multiple Bus Systems* for complete instructions on running system diagnostics.

5. After you have physically replaced the drive, initiate a reconstruct on the replaced drive. Alternatively, the user may reconstruct onto any Spare drive. The Failed drive will remain Failed but will no longer be part of the disk array. The status of the replaced drive should change to Reconstructing.

For RAID level 5E disk arrays, after adding the Hot-spare drive the disk array will do one of the following:

- Compressed disk arrays with an Optimal status will decompress. After successfully decompressing the disk array will be decompressed with a status of Optimal.
- Compressed disk arrays with one failed drive will have a status of Degraded. The disk array will begin reconstructing. After successfully reconstructing the disk array will be compressed with a status of Optimal.

Note: If reconstruction does not start automatically (a Hot-spare was not in place), select the **Reconstruct an IBM PCI SCSI Disk Array** option from the PDAM main menu to initiate reconstruction manually.

Wait for reconstruction to finish. The drive status of the replaced drive changes to Online, and the status of the affected disk array changes to Optimal.

6. After you have replaced and reconstructed all the Failed drives, exit PDAM.

Attention: Do not use this procedure to replace drives in a RAID level 0 array. For the RAID level 0 procedure, see “Replacing a Dead or Optimal Array”. If there are two or more drives that need to be replaced, stop using all logical volumes and file systems on the disk array immediately. Back up all logical volumes on the disk array before continuing.

Replacing a Dead or Optimal Array

This procedure is essentially the same procedure you would use to replace a Failed non-array disk drive in the AIX operating system. Use this procedure to restore disk arrays in the following circumstances:

- All RAID level 0 disk arrays with a status of Dead or Optimal.

Note: A RAID level 0 disk array that has a drive in the Warning state will be in the Optimal state. A RAID level 0 disk array that has a drive in the Failed state will be in the Dead state.

- A RAID level 1 disk array with a status of Dead.
 - A RAID level 5 disk array with a status of Dead.
 - A RAID level 5E disk array with a status of Dead.
1. For a RAID level 0 disk array that is Optimal, attempt to back up the files in all the logical volumes and file systems on the disk array.
 2. Remove all logical volumes and file systems on the disk array affected by the bad drive. See “Deleting Disk Arrays from the Operating System” on page 43.
 3. Start the PCI SCSI Disk Array Manager by entering the following command:

```
smit pdam
```
 4. Select the **Change/Show PCI SCSI Drive Status** option.
 5. Locate the drive that the PDAM displays as Failed or Warning.

Note: If the drive is in the Warning state, fail the drive manually before continuing. See “Reviving a Failed Drive” on page 36.

6. Delete the disk array. Select the **Delete a PCI SCSI Disk Array** option from the PDAM main menu. Select the appropriate disk array and press Enter.
7. The Failed drives will remain Failed outside of the disk array. Select the **Remove a Failed Drive** option from PDAM and remove the Failed drive from the system. Physically replace the drives, adding them back to the system using the **Add a Spare Drive** option.
8. Recreate the disk array. Use the **Create a PCI SCSI Disk Array** option from the PDAM main menu.
9. Create new logical volumes and file systems on the restored disk arrays (see "Adding Disk Arrays to the AIX Operating System" on page 41)
10. Copy backed-up data back to the restored logical volumes and file systems.

Replacing the PCI 4-Channel Ultra3 SCSI RAID Adapter

Attention: If the PCI 4-Channel Ultra3 SCSI RAID Adapter being replaced experienced an abrupt failure, there may be valid data in the adapter's non volatile write cache.

If an existing disk array configuration will be imported from the attached drives, then the write cache module from the adapter being replaced should be used with the new adapter. This will ensure any data which may exist in the non volatile write cache module is correctly written to the attached disk drives.

A configuration conflict may result when a PCI 4-Channel Ultra3 SCSI RAID Adapter within a configured system is replaced (possibly due to a defective adapter). The configuration stored on the adapter may be out of sync with the configuration stored on the drives. The configuration can be restored by synchronizing the adapter configuration with the drive configuration. This can be accomplished using the **Accept Configuration on Drives** option under **PDAM's Recovery Options**. Use the following to replace the PCI 4-Channel Ultra3 SCSI RAID Adapter:

1. Start the SMIT PDAM by entering the following command:

```
smit pdam
```
2. Select the **Recovery Options** option.
3. Select **Resolve PCI SCSI RAID Adapter Configuration** option.
4. Select **Accept Configuration on Drives** option.
5. Select the adapter that was replaced.

Note: Following an Accept Configuration on Drives option, the PCI 4-Channel Ultra3 SCSI RAID Adapter will automatically perform a Consistency Check with Auto-Repair on all redundant level RAID arrays. Progress of the Consistency Check can be monitored using the List PCI SCSI Disk Arrays option of PDAM.

Detecting and Resolving Array Configuration Conflicts

An array configuration conflict exists when, during adapter configuration, the adapter is unable to detect an Online, Warning, Spare, Hot-spare disk drive, or when a drive in a disk array is moved to another location. All disk arrays (hdisks) associated with the adapter will remain in the Defined state when this condition occurs. An array configuration conflict can only exist following adapter configuration. It is, therefore, most likely to occur during system reboot, but can also be seen if the adapter is configured during runtime (for example, changed from the Defined to Available state). No PDAM operations are supported until the configuration conflict is resolved. Attempting to perform any PDAM operation will result in the following error:

```
scraid0 A conflicting Disk Array configuration has been
detected. Please proceed to the 'Display/Accept
Configuration Changes' option under the Recovery Options
panel to resolve.
<#0106>0016 - RAID Adapter command has failed.
```

Also, the following message is displayed and mail is sent to the root user when the conflict is first detected:

```
A condition has been detected that requires immediate
attention. A physical disk drive previously configured on
the IBM PCI SCSI RAID adapter scraid0 was not detected
during adapter configuration. The physical disk has
either been removed or is not detectable. Proceed to the
IBM PCI SCSI Disk Array Manager Recovery Options where
you can either accept that the drive is missing
(Display/Accept Configuration Changes), or attempt again
to detect the disk (Retry Current Configuration). Until
one of these options is successfully executed, no logical
RAID disks will be made available on this adapter.
```

An entry is also posted to the system error log indicating an array configuration conflict exists. The error log entry appears as follows:

```
9A1495A2    0308113896 U H scraid0    ARRAY CONFIGURATION
ERROR
```

The error log entry in its expanded form appears as follows:

```
LABEL:          SCSI_ARRAY_ERR6
IDENTIFIER:      9A1495A2
```

Date/Time:

Fri Mar 8 11:38:38

Sequence Number: 9

Machine Id: 00300001C400

Node Id: drool

Class: H

Type: UNKN

Resource Name: scraid0

Resource Class: adapter

Resource Type: 14102e00

Location: 04-07

VPD:

```
Part Number.....12345678
EC Level.....1234567890
Serial Number.....12345678
FRU Number.....12345678
Manufacturer.....IBM000
Displayable Message.....PCIRAI
Diagnostic Level.....01
Device Driver Level.....01
Loadable Microcode Level....96064!
ROS Level and ID..... P10809
```

Description

ARRAY CONFIGURATION ERROR

Probable Causes

ARRAY COMPONENT UNDETECTABLE

Failure Causes

DISK DRIVE

CABLE LOOSE OR DEFECTIVE

SCSI BACKPLANE

SCSI REPEATER

Recommended Actions

PERFORM DISK ARRAY MANAGER RECOVERY ACTIONS

Detail Data

SENSE DATA

```
0000 0052 0000 0708 0000 0000 0000 0000 0000 0000
0000 0000 0000 0000 0000
```

Any of the following indicate the presence of an array configuration conflict: error from PDAM, alertable message, or error log entry.

Resolving Disk Array Configuration Conflicts

An action must be taken to resolve the conflict when an array configuration conflict occurs. No disk arrays configured on the adapter are available to the system until the conflict is resolved.

Note: The conflict should be handled as described in the “Replacing the PCI 4-Channel Ultra3 SCSI RAID Adapter” on page 60 if an array configuration conflict occurs subsequent to an adapter replacement. An Array configuration conflict resulting from an adapter replacement should not be resolved by synchronizing to the configuration on the adapter.

If the disk drive(s) which was not detected was part of an array, the array(s) associated with the disk drive(s) and its contents are displayed along with the status of the array prior to the conflict being detected. The location of the undetected disk drive and the state previous to being undetected are included.

The following information is displayed:

```
hdisk2  Raid 5 04-07-00-0,0 2144 MB Status OPTIMAL
        hdisk2  20 Channel 2 ID 0  ONLINE
        hdisk2  21 Channel 2 ID 1  ONLINE
        hdisk2  25 Channel 2 ID 5  ONLINE
Channel 2 ID 0 indicates change from status ONLINE to
status FAILED DRIVE
```

Once it has been determined which specific disk drives are involved in the array configuration conflict, take the following actions to resolve the conflict:

Retry Current Configuration

1. Select the **Recovery Options** option of the PDAM
2. Select the **Resolve PCI SCSI RAID Adapter Configuration** option.
3. Select **Retry Current Configuration**.

This action causes the adapter to attempt to detect the missing disk drives. This option would be selected after making some adjustment to the hardware that may affect the adapter's ability to detect the disk drive; for example, powering on an external disk enclosure or reattaching a loose cable connection. If the adapter is successful in detecting the missing disk drives, then a good completion is indicated by this option, all disk arrays associated with the adapter are configured (for example, made Available), and the conflict is resolved. If the adapter is not successful in detecting the missing disk drives, then the array configuration conflict has not been resolved. See “Messages and Error Log Templates” on page 66 or contact your service representative.

Accept Configuration Changes

1. Select the **Recovery Options** option of the PDAM.
2. Select **Resolve PCI SCSI RAID Adapter Configuration**.
3. Select **Display/Accept Configuration Changes**.
4. Select the adapter reporting the conflict.
5. Change the **Display only** option to NO.

The array configuration is updated to reflect an Online, Spare, or Hot-spare disk drive is now in the Failed state. If any disk arrays are affected by failed drives, the state of that array is updated to reflect the disk failure. If any hard configured drives have been relocated within the channel then their new location is updated. All Optimal and Degraded disk arrays are brought to the Available state, and the conflict is resolved.

User action is required to resolve array configuration conflicts. This allows the user to prevent arrays from becoming Failed or Degraded for reasons such as neglecting to supply power to a disk enclosure or connecting an SCSI cable. In these instances, the failure can be remedied without an array having become Degraded or Failed, and an unnecessary reconstruct operation may be avoided.

Auto Rearrange

Attention: This option should be used with CAUTION. Incorrect use of this option can result in drives marked failed and disk arrays marked dead. If this occurs due to improper use of auto rearrange, put the moved drives back to their original locations and retry the current configuration.

Auto rearrange allows physical drives that are part of a disk array to be moved to other locations on the same channel when the adapter is in the Defined state. During adapter configuration, a configuration conflict occurs. See “Accept Configuration Changes” to accept the new drive locations.

The PCI 4-Channel Ultra3 SCSI RAID Adapter supports auto rearrange when the following rules are followed:

1. The physical drives part of a disk array are moved when the adapter is in the Defined state.
2. The new location for the drive moved is on the same channel as its old location.
3. Drives are not removed or added to any channels on the adapter during auto rearrange.
4. Only drives part of a disk array can be moved and recognized at adapter configuration time.

Resolving Common Problems

Symptom:

The PCI 4-Channel Ultra3 SCSI RAID Adapter is in the Defined state.

Action:

Attempt to configure the Defined adapter instance.

1. Use the `smitty` devices fast path.
2. Select **Disk Array** option.
3. Select the **IBM PCI SCSI Disk Array** option.
4. Select the **Configure a Defined PCI SCSI RAID Adapter** option.
5. Select the adapter instance.

If this action fails to configure the adapter, verify the adapter is still located in the described PCI slot, and run diagnostics on the adapter.

Symptom:

The PCI 4-Channel Ultra3 SCSI RAID Adapter is in the Available state, but all associated arrays (hdisks) are in the Defined state.

Action:

If the adapter was just replaced, follow the steps described in “Replacing the PCI 4-Channel Ultra3 SCSI RAID Adapter” on page 60.

If an array configuration conflict exists, follow the steps described in “Detecting and Resolving Array Configuration Conflicts” on page 61

If it is a Failed disk array that will not configure, it is left in the Defined state and is not accessible.

Symptom:

A disk drive attached to the PCI 4-Channel Ultra3 SCSI RAID Adapter is not recognized by the adapter.

Action:

Attempt to configure the device. See “Add a Spare Drive” on page 33.

Note: If the drive was attached to the adapter before the last reboot, no user intervention is required. The adapter should recognize the new device during the boot process.

Determine if there is a device or cabling problem.

1. If the new device introduces a SCSI ID conflict, more than one device may be affected. Check the ID cable that is part of the hot-plug drive carrier.
2. See *Diagnostic Information for Multiple Bus System* for procedures about SCSI problem isolation. Follow the steps for cabling and devices.

Messages and Error Log Templates

This section contains alertable messages and error log templates.

Temporary RAID Adapter Error

F0>IDENTIFIER 51F6CEBE

Label: SCSI_ARRAY_ERR1

Class: H

Type: TEMP

Loggable: YES Reportable: YES Alertable: NO

Description

ADAPTER ERROR

Probable Causes

ADAPTER MICROCODE

ADAPTER HARDWARE

Failure Causes

ADAPTER MICROCODE

ADAPTER HARDWARE

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Permanent RAID Adapter Error

IDENTIFIER 1EB75C3D

Label: SCSI_ARRAY_ERR2

Class: H

Type: PERM

Loggable: YES Reportable: YES Alertable: NO

Description

ADAPTER ERROR

Probable Causes

ADAPTER MICROCODE

ADAPTER HARDWARE

Failure Causes

ADAPTER MICROCODE

ADAPTER HARDWARE

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Temporary Array SCSI Bus Error

IDENTIFIER 1D57AC2F

Label: SCSI_ARRAY_ERR3

Class: H

Type: TEMP

Loggable: YES Reportable: YES Alertable: NO

Description

SCSI BUS ERROR

Probable Causes

SCSI BACKPLANE

SCSI REPEATER

CABLE

DISK DRIVE ELECTRONICS

Failure Causes

SCSI BACKPLANE

SCSI REPEATER

CABLE LOOSE OR DEFECTIVE

DISK DRIVE ELECTRONICS

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Temporary Physical Volume Error

IDENTIFIER 5ABCCC6B

Label: SCSI_ARRAY_ERR4

Class: H

Type: TEMP

Loggable: YES Reportable: YES Alertable: NO

Description

DISK OPERATION ERROR

Probable Causes

DASD MEDIA

DASD DEVICE

Failure Causes

DASD MEDIA

DISK DRIVE

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Permanent Physical Volume Error

IDENTIFIER 15FD5EE8

Label: SCSI_ARRAY_ERR5

Class: H

Type: PERM

Loggable: YES Reportable: YES Alertable: NO

Description

DISK OPERATION ERROR

Probable Causes

DASD MEDIA

DASD DEVICE

Failure Causes

DASD MEDIA

DISK DRIVE

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Array Configuration Error

IDENTIFIER 9A1495A2

Label: SCSI_ARRAY_ERR6

Class: H

Type: UNKN

Loggable: YES Reportable: YES Alertable: NO

Description

ARRAY CONFIGURATION ERROR

Probable Causes

ARRAY COMPONENT UNDETECTABLE

Failure Causes

DISK DRIVE

CABLE LOOSE OR DEFECTIVE

SCSI BACKPLANE

SCSI REPEATER

Recommended Actions

PERFORM DISK ARRAY MANAGER RECOVERY ACTIONS

Detail Data

SENSE DATA

Array Component Failure

IDENTIFIER F084AF4B

Label: SCSI_ARRAY_ERR7

Class: H

Type: UNKN

Loggable: YES Reportable: YES Alertable: NO

Description

ARRAY COMPONENT FAILURE

Probable Causes

DASD DEVICE

CABLE

SCSI BACKPLANE

SCSI REPEATER

Failure Causes

DISK DRIVE

CABLE LOOSE OR DEFECTIVE

SCSI BACKPLANE

SCSI REPEATER

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Temporary Unknown Software Failure

IDENTIFIER DF7F0DED

Label: SCSI_ARRAY_ERR8

Class: S

Type: TEMP

Loggable: YES Reportable: YES Alertable: NO

Description

SOFTWARE PROGRAM ERROR

Probable Causes

SOFTWARE PROGRAM

ADAPTER MICROCODE

Failure Causes

SOFTWARE PROGRAM

ADAPTER MICROCODE

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

SENSE DATA

Recovered RAID Adapter Error

IDENTIFIER 8AC43378

Label: SCSI_ARRAY_ERR9

Class: H

Type: TEMP

Loggable: YES Reportable: YES Alertable: NO

Description

ADAPTER ERROR

Probable Causes

DASD DEVICE

CABLE

SCSI BACKPLANE

SCSI REPEATER

Failure Causes

DISK DRIVE

CABLE LOOSE OR DEFECTIVE

SCSI BACKPLANE

SCSI REPEATER

Recommended Actions

PERFORM PROBLEM DETERMINATION PROCEDURES

Detail Data

SENSE DATA

Alertable Messages

SCSI_ARRAY_ERR6

A condition has been detected that requires immediate attention. A physical disk drive previously configured on the IBM PCI SCSI RAID adapter scraidN was not detected during adapter configuration.

The physical disk has either been removed or is not detectable. Proceed to the IBM PCI SCSI Disk Array Manager Recovery Options where you can either accept that the drive is missing (Display/Accept Configuration Changes), or attempt again to detect the disk (Retry Current Configuration). Until one of these options is successfully executed, no logical RAID disks will be made available on this adapter.

SCSI_ARRAY_ERR7

A condition has been detected that requires immediate attention. A physical disk drive that is part of a logical RAID disk configured on the IBM PCI SCSI RAID adapter scraidN has failed. The corresponding logical RAID disk is now either degraded or offline. Proceed to the List IBM PCI SCSI Disk Arrays option in the IBM PCI SCSI Disk Array Manager to determine which physical disk has failed and if a hotspare rebuild has been initiated.

See “Resolving Disk Array Configuration Conflicts” on page 63 for more information.

Appendix. Notices

This information was developed for products and services offered in the U.S.A.

The manufacturer may not offer the products, services, or features discussed in this document in other countries. Consult the manufacturer's representative for information on the products and services currently available in your area. Any reference to the manufacturer's product, program, or service is not intended to state or imply that only that product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any intellectual property right of the manufacturer may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any product, program, or service.

The manufacturer may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the manufacturer.

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: THIS MANUAL IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. The manufacturer may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Information concerning products made by other than the manufacturer was obtained from the suppliers of those products, their published announcements or other publicly available sources. The manufacturer has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to products made by other than the manufacturer. Questions on the capabilities of products made by other than the manufacturer should be addressed to the suppliers of those products.

Index

A

- about this book v
- Adapter, managing the PCI 4-Channel Ultra3 SCSI RAID 27
- adapter overview 1
- add a spare drive 33
- adding disk arrays to the AIX operating system 41
- alertable messages 75
- array, disk 10
- array, replacing a dead or optimal 59
- array component failure 72
- array configuration conflicts, detecting and resolving 61
- array configuration error 71
- array manager, PCI SCSI 11
- array manager, PCI SCSI disk 11
- array overview, disk 12
- array parameters, disk 12
- arrays, disk 1
- arrays, identify drives and disk 16
- assigning hot spare drives 36

C

- cache, non-volatile on board write 11
- change reconstruction rate parameters 51
- change/show PCI SCSI disk array parameters 38
- checking consistency (check and repair array parity) 52
- clearing a configuration 40
- common problems, resolving 65
- configuration, clearing a 40
- configuration synchronization 39
- consistency, perform check 52
- creating a disk array 29

D

- delete a PCI SCSI disk array 35
- delete a spare or failed drive 34
- deleting disk arrays from the operating system 43
- detecting and resolving array configuration conflicts 61
- device names and location codes, disk array 16
- disk array 10
- disk array, creating a 29
- disk array, delete a PCI SCSI 35
- disk array, performing a consistency check on a 37
- disk array, removing a 35
- disk array, restoring a 58
- disk array configuration conflicts, resolving 63
- disk array device names and location codes 16
- disk array failures 47
 - Raid level 0 47
 - Raid level 1 47
 - Raid level 5 48
 - Raid level 5E 48

- disk array manager, starting the PCI SCSI 28
- disk array ODM information 17
 - device class, subclass, and type 17
 - device location code 18
 - device parent 18
 - PCI 4-Channel Ultra3 SCSI RAID Adapter AIX configuration 18
- disk array overview 12
- disk array parameters 12
- disk array states 31, 45
- disk arrays 1
- disk arrays, identify drives and 16
- disk arrays, restoring and recovering 53
- disk drives, SCSI 10
- drive, reconstructing a failed 37
- drive failures and disk arrays 45
- drive status 45
- drive status, modifying and displaying 31
- drives, SCSI disk 10
- drives and disk arrays, identify 16

E

- extending logical volumes and disk arrays 44

F

- fail a drive 34
- failed drive, reviving a 36
- failures, disk array 47
- features
 - PCI 4-Channel Ultra3 SCSI RAID Adapter 10

H

- hot spare drives, assigning 36

I

- identify drives and disk arrays 16
- information, disk array ODM 17
- installing a PCI 4-Channel Ultra3 SCSI RAID Adapter
 - software update 24
 - procedure 24
- installing an PCI 4-Channel Ultra3 SCSI RAID Adapter on AIX Version 4 21
 - Prerequisite 22
 - procedure 22
- installing the PCI 4-Channel Ultra3 SCSI RAID Adapter software 21

L

- levels, RAID 3
- location codes, disk array device names and 16
- logical volumes and disk arrays, extending 44

M

- manager, PCI SCSI disk array 11
- managing the PCI 4-Channel Ultra3 SCSI RAID Adapter 27
 - Procedure 27
- messages and error log templates 66
 - alertable messages 75
 - array component failure 72
 - array configuration error 71
 - permanent physical volume error 70
 - permanent RAID adapter error 67
 - recovered RAID adapter error 74
 - temporary array SCSI bus error 68
 - temporary physical volume error 69
 - temporary RAID adapter error 66
 - temporary unknown software failure 73
- modifying and displaying drive status 31

N

- names and location codes, disk array device 16
- non-volatile on board write cache 11
- Notices 77

O

- overview, adapter 1
- overview, disk array 12

P

- parameters, disk array 12
- PCI 4-Channel Ultra3 SCSI RAID Adapter 10
 - features 10
- PCI 4-Channel Ultra3 SCSI RAID Adapter, replacing the 60
- PCI 4-Channel Ultra3 SCSI RAID Adapter problem determination and recovery 45
- PCI SCSI disk array manager 11
- PCI SCSI disk array parameters, change/show 38
- perform check consistency 52
- performing a consistency check on a disk array 37
- permanent physical volume error 70
- permanent RAID adapter error 67
- physical disk states 32, 46
- preface v
- problem determination and recovery, PCI 4-Channel Ultra3 SCSI RAID Adapter 45

R

- RAID levels 3
 - RAID level 0 4
 - RAID level 1 5
 - RAID level 5 6
 - RAID level 5E 8
- Reader's Comments 81
- reconstructing a failed drive 37
- reconstructing redundant disk arrays after drive failure 50

- reconstructing redundant disk arrays after drive failure 50 (*continued*)
 - the reconstruction rate - low, medium, high 51
- recovered RAID adapter error 74
- related publications v
- removing a disk array 35
- replacing a dead or optimal array 59
- replacing failed or warning drives (RAID levels 1, 5, and 5E) 58
- replacing the PCI 4-Channel Ultra3 SCSI RAID Adapter 60
- resolving common problems 65
- resolving disk array configuration conflicts 63
 - accept configuration changes 64
 - auto rearrange 64
 - retry current configuration 63
- restoring a disk array 58
- restoring and recovering disk arrays 53
- restoring configuration on replacement adapter 40
- reviving a failed drive 36

S

- SCSI disk array manager, PCI 11
- SCSI disk drives 10
- SCSI RAID Adapter on AIX Version 4, installing an PCI 4-Channel Ultra3 21
- software, installing the PCI 4-Channel Ultra3 SCSI RAID Adapter 21
- software installation, verifying the PCI 4-Channel Ultra3 SCSI RAID Adapter 25
- starting the PCI SCSI disk array manager 28
- synchronization, configuration 39

T

- templates, messages and error log 66
- temporary array SCSI bus error 68
- temporary physical volume error 69
- temporary RAID adapter error 66
- temporary unknown software failure 73
- trademarks v
- transitioning a hot spare drive to a spare drive 36

U

- update, Installing a PCI 4-Channel Ultra3 SCSI RAID Adapter software 24

V

- verifying the PCI 4-Channel Ultra3 SCSI RAID Adapter software installation 25
- view the current disk array status 33

W

- write cache, non-volatile on board 11

Reader's Comments — We'd Like to Hear From You

PCI 4-Channel Ultra3 SCSI RAID Adapter

Form Number: SA23-1239-00

Overall how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall Satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your response. May we contact you? ☐ Yes ☐ No

When you send comments to us, you grant us a nonexclusive right to use or distribute your comments in any way we believe appropriate without incurring any obligation to you.

Name	Address
Company or Organization	
Phone Number	

Cut or Fold
Along Line

Fold and Tape

Please do not Staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

POSTAGE WILL BE PAID BY ADDRESSEE

Information Development
Department H6DS-9561
11400 Burnet Road
Austin, TX 78758-3493



Fold and Tape

Please do not Staple

Fold and Tape

Cut or Fold
Along Line

Part Number: 09P1950



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

October 2000

SA23-1239-00



(1P) P/N: 09P1950

