PCI-X SCSI RAID Controller

Reference Guide for AIX

ESCALA POWER6

ESCALA POWER6

PCI-X SCSI RAID Controller

Reference Guide for AIX

Hardware

Trademarks and Acknowledgements

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

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

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

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries

Contents

iii

Safety Information

DANGER

An electrical outlet that is not correctly wired could place hazardous voltage on metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock.

Before installing or removing signal cables, ensure that the power cables for the system unit and all attached devices are unplugged.

When adding or removing any additional devices to or from the system, ensure that the power cables for those devices are unplugged before the signal cables are connected. If possible, disconnect all power cables from the existing system before you add a device.

Use one hand, when possible, to connect or disconnect signal cables to prevent a possible shock from touching two surfaces with different electrical potentials.

During an electrical storm, do not connect cables for display stations, printers, telephones, or station protectors for communication lines.

Handling Static Sensitive Devices

Attention: Static electricity can damage this device and your system unit. To avoid damage, keep this device in its anti-static protective bag until you are ready to install it. To reduce the possibility of electrostatic discharge, follow the precautions listed below:

- Limit your movement. Movement can cause static electricity to build up around you.
- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or exposed printed circuitry.
- Do not leave the device where others can handle and possibly damage the device.
- While the device is still in its anti-static package, touch it to an unpainted metal part of the system unit for at least two seconds. (This drains static electricity from the package and from your body.)
- Remove the device from its package and install it directly into your system unit without setting it down. If it is necessary to set the device down, place it on its static-protective package. (If your device is an adapter, place it component-side up.) Do not place the device on your system unit cover or on a metal table.
- Take additional care when handling devices during cold weather, as heating reduces indoor humidity and increases static electricity.

About This Book

This book provides usage and maintenance information regarding the PCI-X SCSI RAID Controller. Chapters 1 through 4 of this book contain general information that is intended for all users of this product. Chapter 5 contains service information intended for a service representative specifically trained on the system unit and subsystem being serviced. Use this book in conjunction with your specific system unit and operating system documentation.

ISO 9000

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

Highlighting

The following highlighting conventions are used in this book:

Bold	Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.
Italics	Identifies parameters whose actual names or values are to be supplied by the user.
`Monospace`	Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

References to AIX Operating System

This document may contain references to the AIX operating system. If you are using another operating system, consult the appropriate documentation for that operating system.

This document may describe hardware features and functions. While the hardware supports them, the realization of these features and functions depends upon support from the operating system. AIX provides this support. If you are using another operating system, consult the appropriate documentation for that operating system regarding support for those features and functions.

Related Publications

The following publications contain related information:
- System unit documentation for information specific to your hardware configuration
- *RS/6000 @server pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*, order number SA38-0516, contains information about adapters, devices, and cables for your system.
- *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*, order number SA38-0509, contains diagnostic information, service request numbers (SRNs), and failing function codes (FFCs). Available on the internet at https://techsupport.services.ibm.com/server/library
- *AIX System Management Guide: Operating System and Devices* and *AIX Installation Guide and Reference*, available on the internet at https://techsupport.services.ibm.com/server/library
- *The RAIDbook: A Handbook of Storage Systems Technology*, Edition 6, Editor: Paul Massiglia

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:
- AIX
- AIX 5L
- @server
- PowerPC
- pSeries
- RS/6000

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

Chapter 1. PCI-X SCSI RAID Controller Overview

The PCI-X SCSI RAID Controller is a 64-bit PCI-X adapter. This adapter has the following features:

- PCI-X system interface
- Dual or Quad Channel, Ultra320 SCSI controllers
- Embedded PowerPC RISC Processor, hardware XOR DMA Engine, and hardware Finite Field Multiplier (FFM) DMA Engine (for RAID level 6)
- Non-volatile write cache
- Support for RAID (Redundant Array of Independent Disks) levels 0, 5, 6, and 10 disk arrays
- Disk array supported as a bootable device
- Supports attachment of other devices such as non-RAID disks, tape, and optical devices
- Advanced RAID features:
  - Hot spares for RAID level 5, 6, and 10 disk arrays
  - Ability to increase the capacity of an existing RAID level 5 or 6 disk array by adding disks
  - Background parity checking
  - Background data scrubbing
  - Disks formatted to 522 bytes/sector, providing Longitudinal Redundancy Checking (LRC) and Logically Bad Block checking

Table 1. RAID Controller cards by Custom Card Identification Number (CCIN)

CCIN (Custom Card Identification Number)	PCI Form Factor	Channels	RAID Levels Supported	Write Cache Size	Read Cache Size	Multi Initiator and High Availability Support	Auxillary Cache Support	Adapter LED/FFC	Cache Battery Pack Technology	Cache Battery Pack Concurrent Maintenance	Cache Battery Pack FFC
2780[9]	Yes	4	0, 5, and 10	Up to 757 MB (compressed)	Up to 1024 MB (compressed)	No	Yes	2527	LiIon	Yes	2D01
5703, 1975	Yes	2	0, 5, and 10	40 MB	0 MB	Yes[3]	No	2523	NiMH	No	2526
5709[1], 1976[1]	No [2]	2	0, 5, and 10	16 MB	0 MB	No	No	2523	NiMH	No	2526
571B, 1913	Yes	2	0, 5, 6 and 10	90 MB	0 MB	Yes[3]	No	2514	LiIon	No	2D03
571E[9]	Yes	4	0, 5, 6 and 10	Up to 1.5 GB (compressed)	Up to 1.6 GB (compressed)	No	Yes	2512	LiIon	Yes	2D01
572E	No [4]	2	0 and 10	0 MB	0 MB	No	No	252B	NA	NA	NA
573D[1], 1908[1]	No [2]	2	0, 5, and 10	40 MB	0 MB	No	No	2529	NiMH	No	2526
571F/575B[5, 9]	Yes[5]	4[6]	0, 5, 6 and 10	Up to 1.5 GB (compressed)	Up to 1.6 GB (compressed)	No	Yes	2513 / 252E	LiIon	Yes[7]	2D06
574F[8, 9]	Yes	1	NA	Up to 1.5 GB (compressed)	0 MB	No	Yes	2D07	LiIon	Yes	2D01

[1] The CCIN describes the RAID Enablement Card.

[2] Some PCI-X SCSI RAID Controllers do not have the form of a standard PCI adapter. The form of these controllers is a combination of SCSI and PCI-X bus interface logic integrated onto the system board and a RAID Enablement Card. The card plugs into the system board in order to connect to the SCSI and PCI-X bus interface logic.

[3] For Multi Initiator and High Availability information, see Chapter 4, "Multi-Initiator and High-Availability," on page 25.

[4] This PCI-X SCSI RAID Controller is integrated onto the system board and attaches to Serial Attached SCSI (SAS) disks via 3 Gb connections.

[5] Double-wide adapter which takes up two PCI-X slots. It is composed of both a storage I/O adapter (571F) and an Auxiliary Cache adapter (575B).

[6] Three channels available for external device attachment. One internal channel dedicated to Auxiliary Cache attachment. See the "Auxiliary Cache" on page 10 section for more information.

[7] The Cache Battery Pack for both adapters is contained in a single battery FRU which is physically located on the 575B Auxiliary Cache card.

[8] This is a standalone Auxiliary Cache adapter which should be connected to the last SCSI port of either a 2780 or 571E storage I/O adapter. See the "Auxiliary Cache" on page 10 section for more information.

[9] This adapter is only supported in an AIX partition on an iSeries or System i server.

Note: This guide refers to various hardware and software features and functions. The realization of these features and functions depends on the limitations of your hardware and software. AIX supports all functions mentioned. If you are using another operating system, consult the appropriate

documentation for that operating system regarding support for the mentioned features and functions. This guide also uses the terms *channel* and *bus* interchangeably.
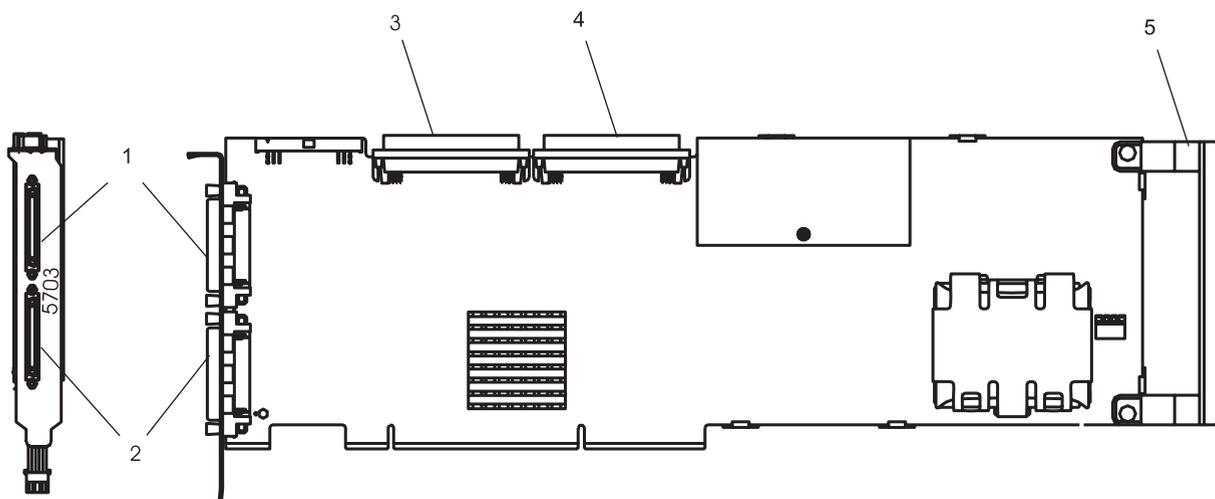
Figure 1. PCI-X SCSI RAID Controller

1   Channel 1 external, 68-position, VHDCI (very high density cable interconnect)

2   Channel 0 external, 68-position, VHDCI (very high density cable interconnect)

3   Channel 0 internal, 68-position, high-density plastic D-shell SCSI standard connector for 16-bit attachment

4   Channel 1 internal, 68-position, high-density plastic D-shell SCSI standard connector for 16-bit attachment

5   Adapter Handle

Disk Arrays

RAID technology is used to store data across a group of disks known as a *disk array*. Depending on the RAID level selected, this storage technique provides the data redundancy required to keep data secure and the system operational. If a disk failure occurs, the disk can usually be replaced without interrupting normal system operation. Disk arrays also have the potential to provide higher data transfer and input and output (I/O) rates than those provided by single large disks.

Each disk array can be used by AIX in the same way as it would a single non-RAID disk. For example, after creating a disk array, you can create a file system on the disk array or use AIX commands to make the disk array available to the system by adding the disk array to a volume group.

The PCI-X SCSI RAID Controller is managed by the PCI-X SCSI Disk Array Manager. The PCI-X SCSI Disk Array Manager serves as the interface to the RAID configuration. It is also responsible for the monitoring and recovery features of the controller.

If a disk array is to be used as the boot device, it may be required to prepare the disks by booting from the Diagnostic CD and creating the disk array before installing AIX. You might want to perform this procedure when the original boot drive is to be used as part of a disk array.

The following figure illustrates a possible disk array configuration.

Disk Array Controller

SCSI Bus 0   SCSI Bus 1

RAID 10 Disk Array

SCSI ID 0   SCSI ID 1   SCSI ID 0   SCSI ID 1

RAID 0 Disk Array   RAID 5 Disk Array

SCSI ID 2   SCSI ID 3   SCSI ID 2   SCSI ID 3

Hot Spare   Hot Spare

SCSI ID 4   SCSI ID 5   SCSI ID 4   SCSI ID 5
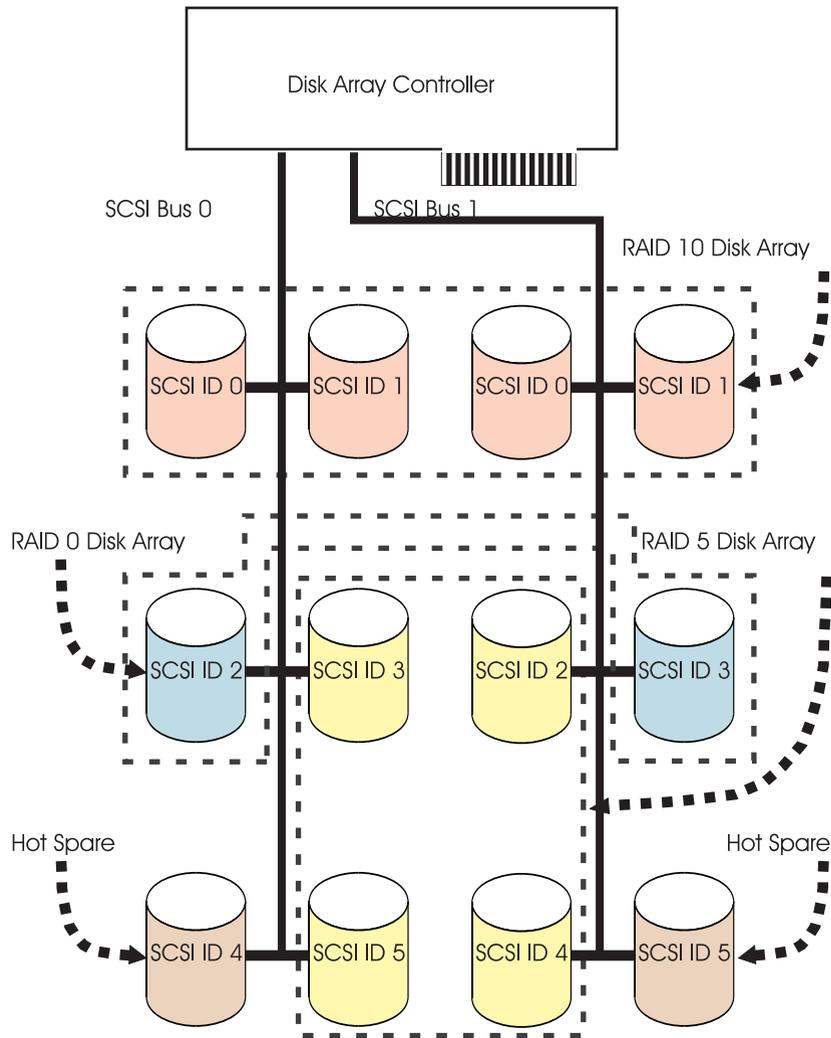
Supported RAID Levels

The RAID level of a disk array determines how data is stored on the disk array and the level of protection that is provided. When a part of the RAID system fails, different RAID levels help to recover lost data in different ways. With the exception of RAID level 0, if a single drive fails within an array, the array controller can reconstruct the data for the *Failed* disk by using the data stored on other hard drives within the array. This data reconstruction has little or no impact to current system programs and users. The PCI-X SCSI RAID Controller supports RAID levels 0, 5, 6, and 10. Each RAID level supported by the PCI-X SCSI RAID Controller has its own attributes and uses a different method of writing data. The following information details each supported RAID level.

RAID Level 0

RAID level 0 stripes data across the disks in the array, for optimal performance. For a RAID level 0 array of three disks, data would be written in the following pattern.
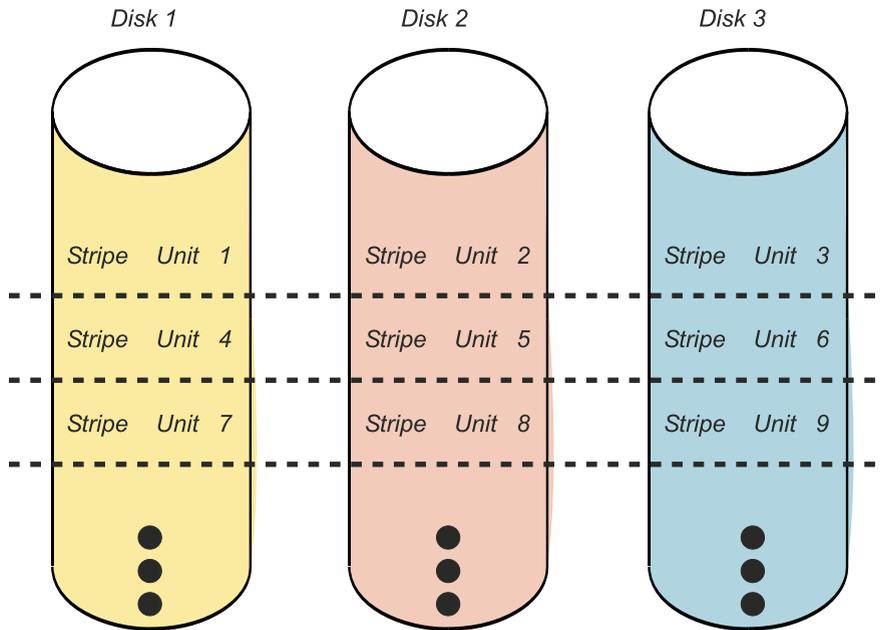
Figure 2. RAID Level 0

RAID level 0 offers a high potential I/O rate, but it is a nonredundant configuration. As a result, there is no data redundancy available for the purpose of reconstructing data in the event of a disk failure. There is no error recovery beyond what is normally provided on a single disk. Unlike other RAID levels, the array controller never marks a RAID level 0 array as *Degraded* as the result of a disk failure. If a physical disk fails in a RAID level 0 disk array, the disk array is marked as *Failed*. All data in the array must be backed up regularly to protect against data loss.

RAID Level 5

RAID level 5 stripes data across all disks in the array. RAID level 5 also writes array parity data. The parity data is spread across all the disks. For a RAID level 5 array of three disks, array data and parity information are written in the following pattern:

Disk 1   Disk 2   Disk 3

Parity	Stripe Unit 1	Stripe Unit 2
Stripe Unit 3	Parity	Stripe Unit 4
Stripe Unit 5	Stripe Unit 6	Parity
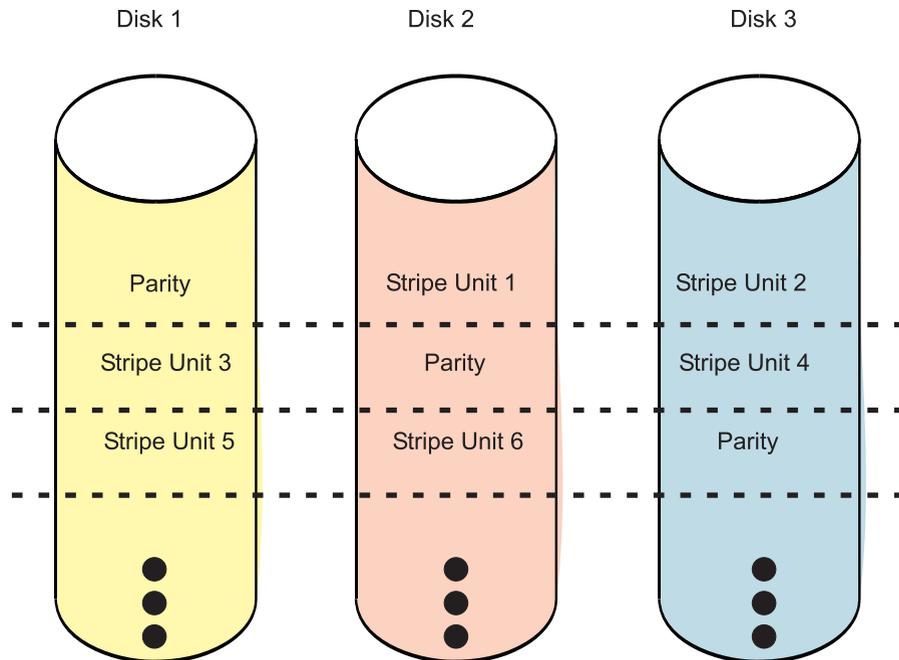
Figure 3. RAID Level 5

If a disk 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* disk 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* disk by using data and parity blocks on the operational disks. If a second disk fails, the array will be placed in the *Failed* state and will not be accessible.

RAID Level 6

RAID level 6 stripes data across all disks in the array. RAID level 6 also writes array "P" and "Q" parity data. The P and Q parity data, which is based on Reed Solomon algorithms, is spread across all the disks. For a RAID level 6 array of four disks, array data and parity information are written in the following pattern:

Disk 1	Disk 2	Disk 3	Disk 4
"P" Parity	"Q" Parity	Stripe Unit 1	Stripe Unit 2
Stripe Unit 3	"P" Parity	"Q" Parity	Stripe Unit 4
Stripe Unit 5	Stripe Unit 6	"P" Parity	"Q" Parity
"Q" Parity	Stripe Unit 7	Stripe Unit 8	"P" Parity
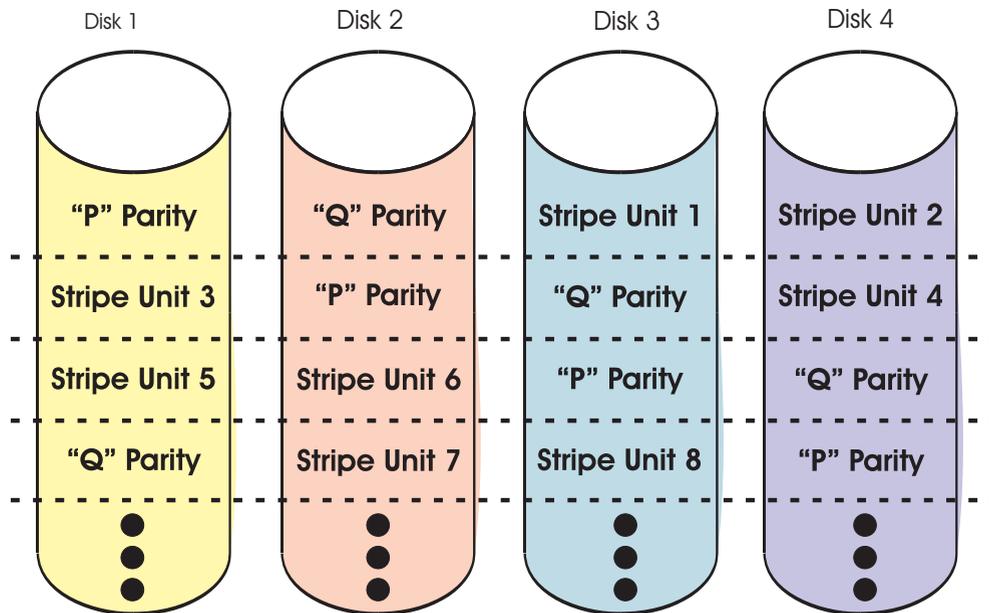
Figure 4. RAID Level 6

If one or two disks fail in a RAID level 6 array, you can continue to use the array normally. A RAID level 6 array operating with a one or two *Failed* disks 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 disk(s) by using data and parity blocks on the operational disks. A RAID level 6 array with a single *Failed* disk has similar protection to that of a RAID level 5 array with no disk failures. If a third disk fails, the array will be placed in the *Failed* state and will not be accessible.

RAID Level 10

RAID level 10 uses mirrored pairs to redundantly store data. The array must contain an even number of disks. Two is the minimum number of disks needed to create a RAID 10 array. The data is striped across the mirrored pairs. For example, a RAID level 10 array of four disks would have data written to it in the following pattern:

Disk 1	Disk 2 (Mirror of Disk 1)	Disk 3	Disk 4 (Mirror of Disk 3)
Stripe Unit 1	Mirrored Stripe Unit 1	Stripe Unit 2	Mirrored Stripe Unit 2
Stripe Unit 3	Mirrored Stripe Unit 3	Stripe Unit 4	Mirrored Stripe Unit 4
Stripe Unit 5	Mirrored Stripe Unit 5	Stripe Unit 6	Mirrored Stripe Unit 6
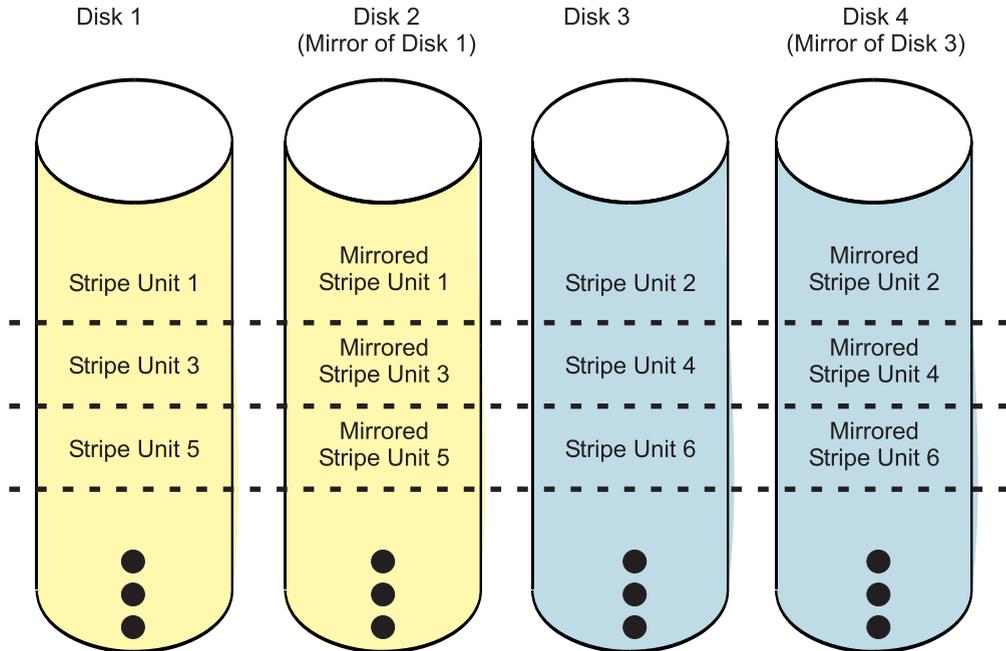
Figure 5. RAID Level 10

RAID level 10 tolerates multiple disk failures. If one disk in each mirrored pair fails, the array will still be functional, operating in *Degraded* mode. You can continue to use the array normally because for each *Failed* disk, the data is stored redundantly on its mirrored pair. However, if both members of a mirrored pair fail, the array will be placed in the *Failed* state and will not be accessible.

When a RAID level 10 disk array is created, the controller will automatically attempt to select the disks for each mirrored pair from a different SCSI bus. For example, if four disks selected for the disk array reside on one of the controller's SCSI buses and another four disks selected reside on another of the controller's SCSI buses, the controller will automatically attempt to create each mirrored pair from one disk on each SCSI bus. In the event of a SCSI bus failure, each mirrored pair will continue to operate in a *Degraded* mode.

Disk Array Capacities

The capacity of a disk array depends on the capacity of the disks used and the RAID Level of the array. To calculate the capacity of a disk array, do the following:

RAID Level 0

Multiply the number of disks by the disk capacity.

RAID Level 5

Multiply one fewer than the number of disks by the disk capacity.

RAID Level 6

Multiply two fewer than the number of disks by the disk capacity.

RAID Level 10

Multiply the number of disks by the disk capacity and divide by 2.

Note: If disks of different capacities are used in the same array, all disks are treated as if they have the capacity of the smallest disk.

RAID Level Summary

The following information provides data redundancy, usable disk capacity, read performance, and write performance for each RAID level.

RAID Level	Data Redundancy	Usable Disk Capacity	Read Performance	Write Performance
RAID 0	None	100%	Very Good	Excellent
RAID 5	Very Good	67% to 94%	Very Good	Good
RAID 6	Excellent	50% to 89%	Very Good	Fair to Good
RAID 10	Excellent	50%	Excellent	Very Good

RAID Level 0

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

RAID Level 5

Creates array parity information so that the data can be reconstructed if a disk in the array fails. Provides better capacity than RAID level 10 but possibly lower performance.

RAID Level 6

Creates array "P" and "Q" parity information, using Reed Solomon algorithms, so that the data can be reconstructed if one or two disks in the array fail. Provides better data redundancy than RAID level 5 but with slightly lower capacity and possibly lower performance. Provides better capacity than RAID level 10 but possibly lower performance.

RAID Level 10

Stores data redundantly on mirrored pairs to provide maximum protection against disk failures. Provides generally better performance than RAID level 5 or 6, but has lower capacity.

Stripe-Unit Size

With RAID technology, data is striped across an array of physical disks. This data distribution scheme complements the way the operating system requests data. The granularity at which data is stored on one disk of the array before subsequent data is stored on the next disk of the array is called the stripe-unit size. The collection of stripe units, from the first disk of the array to the last disk of the array, is called a *stripe*.

You can set the stripe-unit size of a PCI-X SCSI Disk Array to 16 KB, 64 KB, or 256 KB. You may be able to maximize the performance of your PCI-X SCSI Disk Array by setting the stripe-unit size to a value that

is slightly larger than the size of the average system I/O request. For large system I/O requests, use a stripe-unit size of 256 KB. The recommended stripe size for most applications is 64 KB.

Disk Array Overview

Disk arrays are groups of disks that work together with a specialized array controller to potentially achieve higher data transfer and input and output (I/O) rates than those provided by single large disks. The array controller keeps track of how the data is distributed across the disks. RAID level 5, 6, and 10 disk arrays also provide data redundancy, so that no data is lost if a single disk in the array fails.

PCI-X SCSI Disk Arrays are assigned names using the *hdisk* form, the same as other disk storage units in AIX. These names are automatically assigned whenever you create a disk array. The names are deleted when you delete the disk array. The individual physical disks that comprise disk arrays or serve as candidates to be used in disk arrays are represented by *pdisk* names. A *pdisk* is a disk that is formatted to 522 bytes/sector. Disks that are formatted to 512 bytes/sector are assigned names using the *hdisk* form. These disks must be formatted to 522 bytes/sector before they can be used in disk arrays. For additional information about formatting 512 bytes/sector hdisks to make them available for use in PCI-X SCSI Disk Arrays, see "Preparing Disks for Use in PCI-X SCSI Disk Arrays" on page 16.

The **List PCI-X SCSI Disk Array Configuration** option in the PCI-X SCSI Disk Array Manager can be used to display these *pdisk* and *hdisk* names, along with their associated location codes. For details regarding how to view the disk array information, see "Viewing the Disk Array Configuration" on page 17. The following sample output is displayed when the **List PCI-X SCSI Disk Array Configuration** option is invoked.

```
--------------------------------------------------------------------
Name      Location     State       Description       Size
--------------------------------------------------------------------
sisioa1   30-60        Available   PCI-XDDR  Dual Channel U320 SCSI RAID Adapter

hdisk15   30-60-ff-0,1 Optimal     Raid 5 Array      54.3GB
 pdisk0   30-60-00-5,0 Active      Array Member      18.1GB
 pdisk1   30-60-00-6,0 Active      Array Member      18.1GB
 pdisk2   30-60-01-2,0 Active      Array Member      36.4GB
 pdisk4   30-60-01-3,0 Active      Array Member      18.1GB

hdisk16   30-60-ff-0,2 Degraded    Raid 6 Array      72.8GB
 pdisk3   30-60-00-10,0 Active     Array Member      36.4GB
 pdisk5   30-60-01-5,0 Failed      Array Member      36.4GB
 pdisk6   30-60-00-11,0 Active     Array Member      36.4GB
 pdisk9   30-60-00-1,0 Failed      Array Member      36.4GB

pdisk7    30-60-00-12,0 Active     Array Candidate   36.4GB
pdisk8    30-60-00-13,0 Active     Hot Spare         72.8GB

hdisk20   30-60-01-4,0 Available   16 Bit LVD SCSI    9.0GB
hdisk21   30-60-01-6,0 Defined     16 Bit LVD SCSI    9.0GB
```

The third column in the preceding output displays the state of the device.

States for Disk Arrays (hdisks)

The six valid states for PCI-X SCSI Disk Arrays are: *Optimal, Degraded, Rebuilding, Failed, Missing*, and *Unknown.*

Optimal
>   The array is functional and fully protected (RAID 5, 6, and 10) with all *Array Member pdisks* in the *Active* state.

Degraded
>   The array's protection against disk failures is degraded or its performance is degraded. When one

or more *Array Member pdisks* are in the *Failed* state, the array is still functional but might no longer be fully protected against disk failures. For more information about recovery procedures for disk failures, see "Recovery Procedures for Disk Failures" on page 59. When all *Array Member pdisks* are in the *Active* state, the array is not performing optimally because of a problem with the controller's nonvolatile write cache.

Rebuilding
> Redundancy data for the array is being reconstructed. After the rebuild process has completed, the array will return to the *Optimal* state. Until then, the array is not fully protected against disk failures.

Failed   The array is no longer accessible because of disk failures or configuration problems.

Missing
> A previously configured disk array no longer exists.

Unknown
> The state of the disk array could not be determined.

States for Physical Disks (pdisks)

The valid states for *pdisks* are: *Active, RWProtected, Failed, Missing,* and *Unknown*.

Active   The disk is functioning properly.

RWProtected
> The disk is unavailable because of a hardware or a configuration problem.

Failed   The controller can not communicate with the disk or the *pdisk* is the cause of the disk array being in a *Degraded* state.

Missing
> The disk was previously connected to the controller but is no longer detected.

Unknown
> The state of the disk could not be determined.

pdisk Descriptions

The fourth column in the preceding output is a description of the device. For an array, the description indicates the RAID level of the array. The description of a *pdisk* indicates whether the disk is configured as an *Array Member, Hot Spare*, or an *Array Candidate*.

Array Member
> A 522 bytes/sector *pdisk* that is configured as a member of an array.

Hot Spare
> A 522 bytes/sector *pdisk* that can be used by the controller to automatically replace a Failed disk in a *Degraded* RAID level 5, 6, or 10 disk array. A *Hot Spare* disk is useful only if its capacity is greater than or equal to the capacity of the smallest disk in an array that becomes *Degraded*. For more information about hot spare disks, see "Using Hot Spare Disks" on page 20.

Array Candidate
> A 522 bytes/sector *pdisk* that is a candidate for becoming an *Array Member* or a *Hot Spare*.

Auxiliary Cache

The Auxiliary Write Cache adapter (AWC) provides a duplicate, non-volatile copy of write cache data of the RAID controller to which it is connected. Protection of data is enhanced by having two battery backed (non-volatile) copies of write cache each stored on separate adapters. If a failure occurs to the write cache portion of the RAID controller, or the RAID controller itself fails in such a way that the write cache data is not recoverable, the AWC adapter provides a backup copy of the write cache data to prevent data loss during the recovery of the failed RAID controller. The cache data is recovered to the new replacement

RAID controller and then written out to disk before resuming normal operations. The AWC adapter is not a failover device that can keep the system operational by continuing disk operations when the attached RAID controller fails. The system cannot use the auxiliary copy of the cache for runtime operations even if only the cache on the RAID controller fails. The AWC adapter does not support any other device attachment and performs no other tasks than communicating with the attached RAID controller to receive back up write cache data. The purpose of the AWC is to minimize the length of an unplanned outage, due to a failure of a RAID controller, by preventing loss of critical data which may have otherwise required a system reload.

Should a failure of either the RAID controller or the Auxiliary Cache occur, it is extremely important that the Maintenance Analysis Procedures (MAPs) for the Service Request Numbers (SRNs) in the AIX error log be followed precisely. Needed service information can be found in Chapter 6, "Problem Determination and Recovery," on page 61.

The RAID controller and the AWC adapter each require a PCI-X slot and are required to be in the same partition. The two adapters are connected by an internal SCSI connection. This dedicated SCSI bus runs in Ultra4 mode (U320).

If the AWC adapter itself fails or the link between the two adapters fails, the RAID controller will stop caching operations, destage existing write cache data to disk, and run in a performance degraded mode. Once the AWC is replaced or the link is reestablished, the RAID controller will automatically recognize the AWC, synchronize the cache area, resume normal caching function, and resume writing the duplicate cache data to the AWC.

The AWC adapter is typically used in conjunction with RAID protection. RAID functions are not affected by the attachment of an AWC. Because the AWC does not control other devices over the bus and communicates directly with it's attached RAID controller directly over a dedicated SCSI bus, it has little, if any, performance impact on the system.

The AWC adapter may be directly mated with its associated storage I/O adapter (i.e. 571F/575B adapters) or the AWC adapter may be standalone (i.e. 574F) with a SCSI cable connecting it to its associated storage I/O adapter (i.e. 2780 or 571E).

RAID Controller / AWC Configuration

Installation

The following is the required Order of Installation when creating an Auxiliary Cache environment. It is critical to pay close attention to the information contained below and follow the steps exactly as written to avoid problems during installation.

1. Disk Arrays can be previously configured or new arrays can be created after the Auxiliary Cache environment configuration is set up.
2. Ensure both the storage I/O adapter and AWC adapter are installed in the same partiition and in the same enclosure.

   Attention: When working with sepearate storage I/O and AWC adapters (i.e. 2780 or 571E connected to a 574F AWC), do not install the connecting SCSI cable between the storage I/O adapter and AWC adapter until after the required code support is installed.
3. You must update to the latest adapter microcode from the code download website, and to the required levels of both the AIX version and the AIX driver package for your specific adapters. See Table 2 on page 13 for the required code levels and also refer to the installation information for the adapter.
4. When working with sepearate storage I/O and AWC adapters (i.e. 2780 or 571E connected to a 574F AWC), power down the system or partition and install the SCSI cable between the last port of the storage I/O adapter (i.e. farthest from the end where the Cache Battery Pack resides) and the AWC SCSI port.
5. Power on the system or partition and verify correct port cabling, and the function of adapters and disk arrays. Refer to "Viewing the Disk Array Configuration" on page 17. The output displayed will be similar to the following screen examples (the first example screen is for the storage I/O adapter while the second example screen is for the AWC adapter).
6. Verify that both adapters indicate they are Available and Linked to the other adapter, and that all arrays indicate Optimal.

```
+----------------------------------------------------------------------------
|                                 COMMAND STATUS
|
|Command: OK            stdout: yes           stderr: no
|
|Before command completion, additional instructions may appear below.
|
|
|-----------------------------------------------------------------------
|Name       Location      State       Description        Size
|-----------------------------------------------------------------------
|sisioa4    0I-08         Available   PCI-XDDR Quad Channel U320 SCSI RAID Adapter
| scsi19    0I-08-03-07,0 Linked      Target on remote adapter is 01-06,0
|
|hdisk4     0I-08-ff-0,0  Optimal     RAID 5 Array      175.6GB
| pdisk34   0I-08-02-0,0  Active      Array Member       35.1GB
| pdisk35   0I-08-02-1,0  Active      Array Member       35.1GB
| pdisk36   0I-08-02-2,0  Active      Array Member       35.1GB
| pdisk37   0I-08-02-3,0  Active      Array Member       35.1GB
| pdisk38   0I-08-02-4,0  Active      Array Member       35.1GB
| pdisk39   0I-08-02-5,0  Active      Array Member       35.1GB
|
|
+----------------------------------------------------------------------------
```

```
+----------------------------------------------------------------------------
|                                 COMMAND STATUS
|
|Command: OK            stdout: yes           stderr: no
|
|Before command completion, additional instructions may appear below.
|
|
|-----------------------------------------------------------------------
|Name       Location      State       Description        Size
|-----------------------------------------------------------------------
|sisioa5    0J-08         Available   PCI-XDDR Auxiliary Cache Adapter
| scsi20    0J-08-01-06,0 Linked      Target on remote adapter is 03-07,0
|
|
+----------------------------------------------------------------------------
```

Chapter 2. PCI-X SCSI RAID Controller Software

For the PCI-X SCSI RAID Controller to be identified and configured by AIX, the requisite device support software must be installed. The requisite software for the PCI-X SCSI RAID Controller is often preinstalled as part of the normal AIX installation. However, it may be necessary to perform operations related to the installation, verification, and maintenance of the AIX device software for the PCI-X SCSI RAID Controller.

Software for the PCI-X SCSI RAID Controller is packaged in installp format and distributed as part of the base AIX installation media, AIX update media, and through the web-based Fix Delivery Center for AIX. This information is an overview of the AIX software support required for the PCI-X SCSI RAID Controller. For complete information related to the installation and maintenance of AIX, refer to the *AIX Installation Guide and Reference*.

The PCI-X SCSI RAID Controller executes onboard microcode. The AIX command **lsmcode** can be used to determine the level of onboard microcode being used by the controller. Although a version of controller microcode may be distributed along with AIX, this does not necessarily represent the most recent version of microcode available for the controller. For more information regarding obtaining and downloading new PCI-X SCSI RAID Controller microcode, see "Updating the PCI-X SCSI RAID Controller Microcode" on page 32.

Verifying the Installation of the PCI-X SCSI RAID Controller Software

Support for the PCI-X SCSI RAID Controller is contained in the AIX package described by the table below. The device support package contains multiple filesets, each related to a different aspect of device support (for example, run-time environment, diagnostics support, and controller microcode).

These package can be downloaded from http://techsupport.services.ibm.com/server/mdownload2/adapter.html

Table 2. AIX software requirements by Custom Card Identification Number

CCIN (Custom Card Identification Number)	AIX Package	Minimum Required AIX Version
2780	`devices.pci.14106402`	5.2
5703, 1975	`devices.pci.14107802`	5.1 5.2 ML_07* or 5.3 ML_03*
5709, 1976	`devices.pci.14107802`	5.1
571B, 1913	`devices.pci.1410be02`	5.2
571E	`devices.pci.1410bf02`	5.2 ML_08 or 5.3 ML_04
572E	`devices.pci.14108d02`	5.2
573D, 1908	`devices.pci.1410d402`	5.2
571F/575B	devices.pci.1410d502	5.2 ML_10 or 5.3 ML_06
574F	devices.pci.14100c03	5.2 ML_10 or 5.3 ML_06
* Minium level required for Multi Initiator and High Availability support.		
Attention: Ensure the adapters are updated with the latest microcode from the code download website as part of the initial installation. Go to http://techsupport.services.ibm.com/server/mdownload2/adapter.html		

To verify that the device support package for the PCI-X SCSI RAID Controller is installed, type as an example:

```
lslpp -l devices.pci.14107802*
```

13

Output from this command will indicate if device support software for the PCI-X SCSI RAID Controller is installed, and if so, what the corresponding levels of each fileset are.

If the output indicates that no filesets of this name are installed, you must install the appropriate package in order for the PCI-X SCSI RAID Controller to be made available for use. This software package is available as part of the base AIX installation media, AIX update media, and through the web-based Fix Delivery Center for AIX.

Over time, it may become necessary to install software updates in order to have the very latest available level of device software support for the PCI-X SCSI RAID Controller. Updates to the device support software are packaged, distributed, and installed through the same mechanisms used for other portions of the AIX base operating system. The standard AIX technical support procedures can be used to determine the latest available level of device software support for the PCI-X SCSI RAID Controller.

Chapter 3. Common PCI-X SCSI RAID Controller Tasks

The instructions in this chapter pertain to the various tasks that can be performed in order to manage disk arrays.

Using the PCI-X SCSI Disk Array Manager

The PCI-X SCSI Disk Array Manager can be accessed either through the System Management Interface Tool (SMIT), or for some tasks, the AIX command line. For information about using the command line interface, see Appendix C, "AIX Command-Line Interface," on page 105. The PCI-X SCSI Disk Array Manager can also be started from Diagnostics.

To start the PCI-X SCSI Disk Array Manager, do the following:

1. At the command prompt, type `smit`, and press Enter.
2. Select **Devices**.
3. Select **Disk Array**.
4. Select **IBM PCI-X SCSI Disk Array**. From the menu with options for configuring and managing the PCI-X SCSI RAID Controller, select **PCI-X SCSI Disk Array Manager** to display the following menu for managing disk arrays.

```
                        PCI-X SCSI Disk Array Manager
 Move cursor to desired item and press Enter.

   List PCI-X SCSI Disk Array Configuration
   Create an Array Candidate pdisk and Format to 522 Byte Sectors
   Create a PCI-X SCSI Disk Array
   Delete a PCI-X SCSI Disk Array
   Add Disks to an Existing PCI-X SCSI Disk Array
   Configure a Defined PCI-X SCSI Disk Array
   Change/Show Characteristics of a PCI-X SCSI Disk Array
   Reconstruct a PCI-X SCSI Disk Array
   Change/Show PCI-X SCSI pdisk Status
   Diagnostics and Recovery Options



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

You can also use a SMIT fast path to start the PCI-X SCSI Disk Array Manager. On the AIX command line, type `smit pxdam`, and press Enter.

If a disk array is to be used as the boot device, you might need to prepare the disks by booting from the Diagnostic CD and creating the disk array before installing AIX. You might want to perform this procedure when the original boot drive is to be used as part of a disk array.

To start the PCI-X SCSI Disk Array Manager from Diagnostics, do the following:

1. Start Diagnostics and on the Function Selection screen, select **Task Selection**.
2. Select **RAID Array Manager** and press Enter.
3. Select **PCI-X SCSI Disk Array Manager** and press Enter.

Preparing Disks for Use in PCI-X SCSI Disk Arrays

Before a disk can be used in a PCI-X SCSI Disk Array, it must be a 522 bytes/sector *Array Candidate pdisk*. To create an *Array Candidate pdisk* and format it to 522 byte sectors, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Create an Array Candidate pdisk and Format to 522 Byte Sectors**.
3. Select the appropriate PCI-X SCSI RAID Controller.
4. Select the disks that you want to prepare for use in the PCI-X SCSI Disk Arrays.

   Attention: Continuing with this option will format the disks. All data on the disks will be lost. Some disks require that their microcode be updated to the latest level before being formatted to 522 byte sectors. Also, some disks do not support being formatted to 522 byte sectors. The PCI-X SCSI Disk Array Manager will inform the user of these situations when they exist.
5. A message will display asking if you want to continue. To proceed with the format, select **OK** or press Enter to continue. To return to the previous menu without formatting the disks, select **Cancel**.

After the formatting is complete, the disks will be *Array Candidate pdisks* and will be ready for use in PCI-X SCSI Disk Arrays. This operation will also zero all the data on the disks. The PCI-X SCSI RAID Controller keeps track of the disks that have their data zeroed. These *Zeroed Array Candidate pdisks* can be used to create a disk array that will be immediately protected against disk failures, and they are the only disks that can be added to an existing disk array. An *Array Candidate pdisk* will lose its *Zeroed* state after it has been used in an array or is unconfigured. It will also lose its *Zeroed* state after the system has been rebooted. To return an *Array Candidate pdisk* to the *Zeroed* state, follow the steps previously described in this section for preparing disks for use in disk arrays.

Creating a PCI-X SCSI Disk Array

A disk array is created using a set of *Active Array Candidate pdisks*. For disk arrays with data redundancy (RAID level 5, 6, and 10), if all of the *pdisks* are in the *Zeroed* state, the array will become immediately protected against failures. However, if one or more of the *pdisks* are not *Zeroed*, the newly created array will initially be in the *Rebuilding* state. It will be unprotected against disk failures until parity data on all of the disks has been recalculated. For disk arrays with no data redundancy (RAID level 0), it is of no importance whether the *pdisks* are in the *Zeroed* state. To create a PCI-X SCSI Disk Array, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Create a PCI-X SCSI Disk Array**.
3. Select the appropriate PCI-X SCSI RAID Controller on which you want to create an array.
4. Select the RAID level for the array. For more information about selecting an appropriate RAID level, see "Supported RAID Levels" on page 3.
5. Select the stripe size in kilobytes for the array. For more information about the stripe-size parameter, see "Stripe-Unit Size" on page 8. A selection screen similar to the following displays. It will display a list of *Array Candidate pdisks* and notes regarding array requirements. The minimum and maximum number of supported disks will be specified, along with any other specific requirements for the array.
6. Select the disks that you want to use in the array according to the requirements on this screen.

```
                     Create a PCI-X SCSI Disk Array
Move cursor to desired item and press F7. Use arrow keys to scroll.
    ONE OR MORE items can be selected.
Press Enter AFTER making all selections.

   # RAID 10 supports a minimum of 2 and a maximum of 18 disks
   # The total number of disks must be a multiple of 2

   pdisk10   30-60-01-1,0  Active      Array Candidate   72.8GB
   pdisk11   30-60-00-4,0  Active      Array Candidate   72.8GB
   pdisk12   30-60-01-8,0  Active      Array Candidate   72.8GB
   pdisk18   30-60-00-2,0  Active      Array Candidate   72.8GB



F1=Help                 F2=Refresh              F3=Cancel
F7=Select               F8=Image                F10=Exit
Enter=Do                /=Find                  n=Find Next
```

7. A SMIT Dialog Screen summarizes your selections. Press Enter to create the array.

You can now add the disk array to a volume group. Logical volumes and file systems can also be created. Use standard AIX procedures to perform these tasks, and treat the array in the same way that you would treat any *hdisk*.

Viewing the Disk Array Configuration

To view the configuration of arrays and disks associated with a particular PCI-X SCSI RAID Controller, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.

2. Select **List PCI-X SCSI Disk Array Configuration**.

3. Choose one or more PCI-X SCSI RAID Controllers. The output displayed will be similar to the following:

```
                        COMMAND STATUS
Command: OK              stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

-------------------------------------------------------------------
Name       Location     State       Description        Size
-------------------------------------------------------------------
sisioa1    30-60        Available   PCI-XDDR Dual Channel U320 SCSI RAID Adapter

hdisk11    30-60-ff-0,1 Optimal     Raid 5 Array       54.3GB
 pdisk0    30-60-00-5,0 Active      Array Member       18.1GB
 pdisk1    30-60-00-6,0 Active      Array Member       18.1GB
 pdisk2    30-60-01-2,0 Active      Array Member       36.4GB
 pdisk4    30-60-01-3,0 Active      Array Member       18.1GB

hdisk12    30-60-ff-0,2 Degraded    Raid 5 Array       72.8GB
 pdisk3    30-60-00-10,0 Active     Array Member       36.4GB
 pdisk5    30-60-01-5,0 Failed      Array Member       36.4GB
 pdisk6    30-60-00-11,0 Active     Array Member       36.4GB

hdisk13    30-60-ff-0,0 Degraded    Raid 6 Array       72.8GB
 pdisk18   30-60-01-10,0 Active     Array Member       36.4GB
 pdisk19   30-60-01-3,0 Failed      Array Member       36.4GB
 pdisk20   30-60-01-11,0 Active     Array Member       36.4GB
 pdisk21   30-60-01-5,0 Failed      Array Member       36.4GB

hdisk14    30-60-ff-0,3 Rebuilding  Raid 10 Array      72.8GB Create 55%
 pdisk7    30-60-00-12,0 Active     Array Member       72.8GB
 pdisk8    30-60-01-12,0 Active     Array Member       72.8GB

hdisk15    30-60-ff-0,6 Failed      Raid 5 Array       109.2GB
 pdisk14   30-60-00-0,0 RWProtected Array Member       36.4GB
 pdisk15   30-60-00-9,0 RWProtected Array Member       36.4GB
 pdisk16   30-60-00-2,0 Missing     Array Member       36.4GB
 pdisk17   30-60-00-3,0 Missing     Array Member       36.4GB

hdisk16    30-60-ff-0,8 Missing     Raid 0 Array       145.6GB

pdisk9     30-60-00-1,0 Active      Hot Spare          18.1GB
pdisk10    30-60-01-1,0 Active      Array Candidate    72.8GB
pdisk11    30-60-00-4,0 Active      Array Candidate    72.8GB Zeroed
pdisk12    30-60-01-13,0 RWProtected Array Candidate   72.8GB Format 35%

hdisk20    30-60-01-4,0 Available   16 Bit LVD SCSI    9.0GB
hdisk21    30-60-01-6,0 Defined     16 Bit LVD SCSI    9.0GB

F1=Help            F2=Refresh          F3=Cancel           F6=Command
F8=Image           F9=Shell            F10=Exit            /=Find
n=Find Next
```

The controller's name, location, status, and description are displayed first. Each PCI-X SCSI disk array *hdisk* is displayed with its *Array Member pdisks* directly underneath it.

- The first column of output is the name of the disk array (*hdisk*) or physical disk (*pdisk*).

- The second column of output is the device's location code. The format of the location code is *AB-CD-EF-G,H*. The *AB-CD* fields identify the location of the PCI-X SCSI RAID Controller. The *EF* value identifies the Controller's SCSI bus that the device is attached to. For a disk array this value will be ff which indicates a logical bus for RAID devices. The *G,H* values are the SCSI ID and Logical Unit Number (LUN) of the device.

- The third column of the above output is the state of the disk array or *pdisk*. For information about the possible disk array and *pdisk* states, see "Disk Array Overview" on page 9. For 512 byte/sector standalone disks (*hdisks*), this column is the AIX device state (for example, *Available* or *Defined*).

- The fourth column is a description of the device. For a disk array, the description is the RAID level of the array. For a *pdisk*, the description can be *Array Candidate, Hot Spare*, or *Array Member*.

- The fifth column is the capacity of the array or disk. For information about how the capacity of an array is calculated for each RAID level, see "Disk Array Capacities" on page 8.
- The sixth column is the status of a long-running command issued to a disk array or *pdisk*. This column is also used to indicate that an *Array Candidate pdisk* has had its data zeroed. The following values may be displayed:

Create nn%
> Disk array is in process of being created.

Delete nn%
> Disk array is in process of being deleted.

Rebuild nn%
> Disk array is in process of being reconstructed.

Resync nn%
> Disk array is in process of having it parity data resynchronized.

Adding nn%
> Disk array is in process of having one or more disks added to it.

Format nn%
> *pdisk* is in process of being formatted.

Zeroed
> *pdisk* has been zeroed.

Array Candidate pdisks and *Hot Spare pdisks* are displayed at the bottom of this screen. The *pdisk* names are displayed, along with location, state, description, capacity, and long-running command status. Any 512 bytes/sector standalone disks (*hdisks*) are displayed, along with location, state, description, and capacity.

Deleting a PCI-X SCSI Disk Array

Attention: After a disk array is deleted, it cannot be accessed. All data will be lost. A disk array that is currently in use or opened cannot be deleted. Also, if a disk array command (such as a disk creation command) is in progress, that disk array cannot be deleted.

To preserve the data on the disk array, you must first back up all files in the logical volumes and file systems on the disk array before removing the disk array from its volume group. To delete the array, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Delete a PCI-X SCSI Disk Array**.
3. Select the PCI-X SCSI RAID Controller.
4. Select the disk array to delete.

When the disk array has been deleted, any *Active Array Member pdisks* will become *Active Array Candidate pdisks*.

Adding Disks to an Existing Disk Array

The PCI-X SCSI RAID Controller supports adding disks to existing RAID level 5 or 6 disk arrays, which allows you to dynamically increase the capacity of a disk array while preserving its existing data. Extensive use of this feature, however, will result in a performance penalty; because the added disks will not contain parity and the data will not be restriped.

An *Array Candidate pdisk* is not necessarily a candidate that can be added to an existing array. In addition to being an Array Candidate, the *pdisk* must also be recognized by the adapter as having its data zeroed. This situation ensures that when the disks are added to the array, the parity data will be correct and the array will remain protected against disk failures.

To add disks to an existing array, do the following:

1. Ensure that the disks to be added are *Zeroed Array Candidate pdisks*. For assistance in viewing and changing the state of the disk, see "Preparing Disks for Use in PCI-X SCSI Disk Arrays" on page 16 and "Viewing the Disk Array Configuration" on page 17.

2. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.

3. Select **Add Disks to an Existing PCI-X SCSI Disk Array**.

4. Select the PCI-X SCSI Disk Array to which you want to add disks.

5. A screen will display similar to the following example. If a particular disk is not included in the list, it may not be a candidate that can be added to the array because of the following reasons:

   • The disk's capacity is less than that of the smallest disk already in the array.

   • The disk has not been formatted as a 522 bytes/sector *Array Candidate pdisk*.

   • The disk does not have its data zeroed.

   For the second and third cases, the disk can be added to an array if it is first formatted using the **Create an Array Candidate pdisk and Format to 522 Byte Sectors** option in the PCI-X SCSI Disk Array Manager.

```
              Add Disks to an Existing PCI-X SCSI Disk Array
  Move cursor to desired item and press F7. Use arrow keys to scroll.
      ONE OR MORE items can be selected.
  Press Enter AFTER making all selections.

    # Choose up to 14 of the following disks to add to hdisk2

    pdisk10   30-60-01-1,0  Active      Array Candidate   72.8GB Zeroed
    pdisk11   30-60-00-4,0  Active      Array Candidate   72.8GB Zeroed
    pdisk12   30-60-01-8,0  Active      Array Candidate   72.8GB Zeroed
    pdisk18   30-60-00-2,0  Active      Array Candidate   80.0GB Zeroed

    # Note: If a disk is not listed here it is either not a candidate
    # to be added to this array or it does not have its data zeroed
    # Use the Create an Array Candidate pdisk and Format to 522 Byte
    # Sectors option to format and zero the disk.

  F1=Help                 F2=Refresh              F3=Cancel
  F7=Select               F8=Image                F10=Exit
  Enter=Do                /=Find                  n=Find Next
```

   A screen summarizes your selections and indicates the performance penalty associated with this option.

6. Press Enter to add the disks to the array. To enable higher level components in the system to use the increased capacity of the disk array, additional steps might be needed.

Using Hot Spare Disks

Hot Spare disks are used to automatically replace Failed disks in a redundant RAID environment. Hot spare disks are useful only if their capacity is greater than or equal to that of the smallest capacity disk in an array that becomes *Degraded*.

Creating Hot Spare Disks

To create hot spare disks, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Change/Show PCI-X SCSI pdisk Status**.
3. Select **Create a Hot Spare**.
4. Select the appropriate PCI-X SCSI RAID Controller.
5. Select the *pdisks* that you want to designate as hot spares. A screen summarizes your selections.
6. Press Enter to create the hot spares.

The disk state changes to *Hot Spare*. On subsequent disk failures, reconstruction of *Failed* disks will occur automatically for RAID level 5, 6, and 10 disk arrays.

Note: If there is a *Degraded* disk array at the time that a *Hot Spare* is configured, reconstruction of the *Failed* disk begins automatically.

Deleting Hot Spare Disks

To delete a *Hot Spare* and change the disk to an *Array Candidate pdisk*, do the following:
1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Change/Show PCI-X SCSI pdisk Status**.
3. Select **Delete a Hot Spare**.
4. Select the appropriate PCI-X SCSI RAID Controller.
5. Select the hot spares to delete.

The *Hot Spare* becomes an *Array Candidate pdisk*.

Viewing and Changing PCI-X SCSI RAID Controller Bus Settings

To view and change settings for the PCI-X SCSI RAID Controller, do the following:
1. At the command prompt, type `smit`, and press Enter.
2. Select **Devices**.
3. Select **Disk Array**.
4. Select **IBM PCI-X SCSI Disk Array**.
5. Select **Change/Show PCI-X SCSI RAID Controller Bus**.
6. Select the PCI-X SCSI RAID Controller bus that you want to change settings on or view.

```
     Change/Show Characteristics of a PCI-X SCSI RAID Controller Bus
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                      [Entry Fields]
  SCSI Adapter                        scsi8
  Description                         PCI-X U320 SCSI RAID >
  Status                              Available
  Location                            30-60-00
  Adapter card SCSI ID                [7]                   +#
  Wide bus enabled                    yes                   +
  Maximum SCSI Bus Speed in MB/s      320                   +
  QAS (Quick Arbitration) Capability  Enable                +
  Apply change to DATABASE only       no                    +


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

A screen displays information about the controller bus. The last five attributes on the screen are changeable. The following is information regarding each of the attributes:

Adapter card SCSI ID

Specifies the SCSI ID value that the controller will use to send commands as an initiator device. If multi-initiators are supported, the SCSI ID that the controller responds to when acting as a target device will be specified. This value is usually changed to avoid a conflict or duplication of SCSI IDs among the controller and other devices on the SCSI bus. Before cabling any new device to the SCSI bus, you should verify that the device's SCSI ID is unique. If the controller does not support multi-initiators, 7 is the only valid value for the SCSI ID.

Wide bus enabled

Specifies whether the controller will negotiate with the device for a 16-bit data transfer width on the SCSI bus. If *Yes* is specified, which is the default value, the controller negotiates for 16-bit data transfers. *Yes* can be specified when an 8-bit device is connected through an 8-bit cable to the SCSI connector. The *No* option sets the data transfer width to 8 bits. Use the *No* option when older devices are connected that may not handle the wide data transfer request, WDTR, messages properly. The WDTR messages are used to detect the data transfer width capabilities of a device.

Maximum SCSI Bus Speed in MB/s

This attribute is used to limit the maximum speed in megabytes per second (MB/s) that the adapter SCSI bus will use. Speed is negotiated between the adapter and SCSI devices. The maximum negotiated speed cannot be greater than the Maximum SCSI Bus Speed in MB/s value.

QAS (Quick Arbitration) Capability

This attribute is used to enable or disable the quick arbitration capability. Enabling QAS capability does not guarantee that QAS will actually occur for any device on a SCSI bus.

Apply Change to Database only

Indicates whether the configuration changes should be applied only to the database or to both the database and the current device operation. For devices that are in use and cannot be changed, this allows the database to be changed for the device and the changes will take effect the next time the system is rebooted.

7. After changing any PCI-X SCSI RAID Controller attribute, press Enter to make the desired changes.

Viewing PCI-X SCSI Disk Array Settings

To view the settings for a PCI-X SCSI Disk Array, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select the **Change/Show Characteristics of a PCI-X SCSI Disk Array** option.
3. Select the desired PCI-X SCSI Disk Array.
4. A SMIT dialog screen displays the attributes of the selected array. The output displayed will be similar to the following:

```
   Change/Show Characteristics of a PCI-X SCSI Disk Array

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                     [Entry Fields]
  RAID Array                         hdisk15
  Description                        SCSI RAID 5 Disk Arra>
  Status                             Available
  Location                           30-60-ff-0,1
  Serial Number                      12D0E934
  Physical volume identifier         none
  Queue DEPTH                        16
  Size in Megabytes                  109200
  RAID Level                         5
  Stripe Size in KB                  64




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

- The **Physical volume identifier** field is a unique value assigned to the *hdisk* if the disk array is a member of a volume group. If the disk array is not a member of a volume group, this field value is none.
- The **Queue DEPTH** field is the depth of the command queue used for this disk array.
- The **Size in Megabytes** field represents the usable capacity of the disk array. For information about calculating capacities for each RAID level, see "Supported RAID Levels" on page 3.
- The **RAID Level** field is the level of protection chosen for this array.
- The **Stripe Size in KB** field is the number of contiguous kilobytes which will be written to a single disk before switching to the next disk in the disk array. It provides the host with a method to tune data striping according to the typical I/O request size.

You cannot change any of the attributes on this screen. The RAID level and stripe size must be specified when the array is created.

Viewing PCI-X SCSI pdisk Settings

To view the PCI-X SCSI *pdisk* settings, do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Change/Show PCI-X SCSI pdisk Status**.
3. Select **Change/Show PCI-X SCSI pdisk**.
4. Select a *pdisk* from the list.
5. The following attributes are displayed:

```
                  Change/Show PCI-X SCSI pdisk
Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                        [Entry Fields]
  Disk                                 pdisk3
  Description                          Physical SCSI Disk Drive
  Status                               Available
  Location                             30-60-00-10,0
  Serial Number                        E3V1W93B
  Vendor and Product ID                IBM     IC35L036UCDY10>
  Service Level
  Size in Megabytes                    60400
  Format Timeout in minutes            [120]                        +#


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

The **Size in Megabytes** field represents the capacity of the *pdisk*.

To display a *pdisk's* Vital Product Data (VPD), do the following:

1. Start the PCI-X SCSI Disk Array Manager by following the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.

2. Select **Change/Show PCI-X SCSI pdisk Status**.

3. Select **Display pdisk Vital Product Data**.

4. Select the appropriate PCI-X SCSI RAID Controller.

5. Select the desired *pdisk*.

Chapter 4. Multi-Initiator and High-Availability

The Multi-Initiator and High-Availability option provides a high-availability environment for system storage by enabling two systems to have access to the same set of disk arrays. This feature is typically used with the IBM High-Availability Cluster Multiprocessing application (HACMP). The IBM HACMP software provides a low-cost commercial computing environment that ensures the mission-critical applications can recover quickly from hardware and software failures.

Use of the High Availability feature requires adapter and AIX software support. The adapters which support this High Availability feature are identified in Table 1 on page 1. The AIX software support levels required for the High Availability feature are identified in Table 2 on page 13. These adapters support being either in a single adapter configuration or as part of this dual adapter (High Availability) configuration. The adapter firmware will determine the correct operating mode based upon the presence or absence of another compatible adapter, and will automatically change the operating mode as needed. No mode jumpers or special configuration settings are needed to configure the adapter.

This function is only supported for connections to disk arrays between two systems or partitions. It is not supported for connections to disk arrays from two adapters in the same partition or system.

One adapter functions as the primary adapter and performs direct management of the physical devices. The second adapter serves as a client of the primary adapter and does not perform reads or writes to devices directly. Commands directed to disk arrays through the secondary adapter are forwarded through the SCSI bus to the primary adapter for execution. The secondary adapter detects if the primary adapter goes offline and transitions itself to become the primary adapter. When the original primary adapter comes back online, it now serves as the secondary adapter, unless it was previously designated as the preferred primary adapter.

When in a High Availability configuration, the adapter behavior is altered to enable easier transition of the disk arrays between adapters in the event of an adapter failure (or for other causes). Write cache on the adapters is disabled in a High Availability configuration which impacts performance. Also, in a High Availability configuration, parity footprints are kept in reserved sectors of the device in addition to nonvolatile adapter memory. Both of these are different than the typical adapter behavior and are applicable only if two adapters are present. The adapter will automatically switch its behavior as needed based on the current configuration.

Multi Initiator and High Availability Configuration Specifics

- Connections to disk arrays must be between two systems or partitions and not between two adapters in the same partition or system.
- All (pdisks) in the disk arrays must be shared between both adapters in a High Availability configuration.
- Non-pdisk devices, including tape and optical, can continue to be connected non-shared to a single adapter.
- A disk array in a High Availability configuration cannot be used as a boot device.
- Attachment to the shared disk enclosures must be from the same SCSI bus port connections on each adapter. See Figure 1 on page 2 for reference.
- Unlike adapter types can be connected in a High Availability configuration as long as both adapter types are IBM PCI-X SCSI Raid Adapters and the same Advanced Functions are supported on each adapter (for example, RAID levels).
- The end user and their applications are responsible to ensure orderly read and write operations to the shared disk arrays, for example, when using device reservation commands (persistent reservation is not supported).

Comparison of JBOD (512 byte/sector standalone disks (hdisks)) versus Disk Arrays (522 bytes/sector disks (pdisks)) in High Availability Configurations

JBOD (512 byte/sector standalone disks (hdisks))

External Low-voltage Differential Connection

The interconnection cable, which has VHDCI connectors, is provided with the external subsystem.
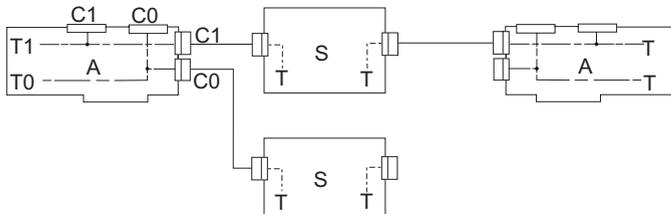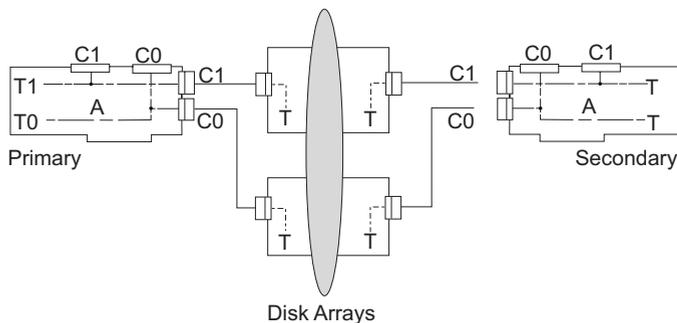

Figure 6. JBOD (512 byte/sector standalone disks (hdisks))

Disk Arrays (522 byte/sector disks (pdisks))

External Low-voltage Differential Connection

The interconnection cable, which has VHDCI connectors, is provided with the external subsystem.

Figure 7. Disk Arrays (522 byte/sector disks (pdisks))

Similarities between JBOD and Disk Arrays in High Availability Configurations:

- Both adapters in a High Availability configuration connect to shared disk enclosures.
- Both adapters in a High Availability configuration can attach 512 byte/sector standalone disks (hdisks).
- Each adapter's SCSI ID must be set to a different value when connected to shared disk enclosures.
- In the event of an adapter failure, the system containing the functional adapter can still access the shared disks.

Differences between JBOD and Disk Arrays in High Availability Configurations

- In a JBOD High Availability environment, there is a direct data path for either adapter to the shared disks.
- In a disk array High Availability environment, there is an indirect data path for the secondary adapter:
  - The Primary adapter controls the disk arrays and has direct access to the pdisks.
  - The Secondary adapter does not communicate directly to the disk arrays, but forwards commands/data to and from the disks arrays via the Primary adapter.
  - The Primary adapter will log errors related to problems with the disk array. These array errors will not be forwarded to the secondary adapter.
- In a disk array High Availability environment, SCSI bus connections from each adapter to the share disk enclosure must be from the same port on each adapter.
- In a disk array High Availability environment, adapters will automatically switch roles (primary to secondary) as needed to maintain access to the disk arrays and to satisfy the system's desire for preferred primary adapter.

  Typical reasons for the Primary and Secondary to switch roles from what was expected or preferred are as follows:
  - Adapters will switch roles for asymmetric reasons. For example one adapter detects more disk drives than the other. If the Secondary adapter is able to find devices that are not found by the Primary adapter, an automatic transition (failover) will occur. The adapters will communicate with each other, compare device information, and make the decision to switch roles.
  - Powering off the Primary adapter or the system/partition which contains the Primary adapter will cause an automatic transition (failover) to occur.
  - Failure of Primary Adapter or the system/partition which contains the Primary adapter will cause an automatic transition (failover) to occur.
  - If the preferred Primary adapter is delayed in coming active, automatic transition (failover) will occur. Once the preferred Primary adapter does come active, the roles will switch back.
  - If the Primary adapter loses contact to the disks that are also accessible by the secondary, an automatic transition (failover) will occur.
- In a disk array High Availability environment, the automatic transition (failover) will not occur for the following:
  - Microcode download of the Primary adapter.
  - Recoverable errors on the Primary adapter.

Performance Considerations

Write cache on the PCI-X SCSI Raid adapters is disabled in an High Availability configuration. The performance impacts from the lack of write cache function will be greater when running Raid-5 or Raid-6 arrays. Performance impacts will be somewhat less if running Raid-0 or Raid-10 arrays. Disk array access via the secondary adapter must be routed through the primary adapter which will also affect read/write performance. Additionally, Raid parity footprints being kept/updated on disk for Raid 5,6,10 arrays, will cause performance impacts.

Installation

The following is the required Order of Installation when creating a Multi-Initiator High-Availability environment. It is critical to pay close attention to the information contained below and follow the steps exactly as written to avoid problems during installation.

1. Before proceeding with physically installing the high availability configuration, you must update to the latest adapter microcode from the code download website, and to the required levels of both the AIX

version and the AIX driver package for your specific adapters. See Table 2 on page 13 for the required code levels and also refer to the installation information for the adapter.

Attention: Do not install any adapters or attach cables in a Multi-Initiator High-Availability configuration until after the required code support is installed.

2. Set the SCSI IDs of the adapters before they are installed.

Attention: Any attached devices with SCSI ID's of 5 and 6 will need to be moved or removed before proceeding with this step.
When installing a Multi-Initiator configuration, each adapter must have a unique SCSI ID and must not duplicate the ID of other devices on the same SCSI bus. Recommended SCSI IDs for the two adapters are 5 and 6. Refer to "Viewing and Changing PCI-X SCSI RAID Controller Bus Settings" on page 21 for setting the SCSI ID.

3. Disk Arrays can be previously configured or new arrays can be created after the high-availability configuration is set up. If new disk arrays are created afterwards, the configuration manager must be run on the system or partition containing the secondary adapter in order for it to recognize the new disk arrays (see 8 on page 29).

   If there are existing disk arrays configured on the adapters that are being used in the high-availability configuration, you must perform a normal power down of the system or partition prior to setting up the high-availability environment.

4. Perform a normal shutdown of the system or partition.

5. Refer to the Figure 6 on page 26 and Figure 7 on page 26 and attach cables from the shared disk enclosures to the same SCSI bus port on each adapter.

   Some disk enclosures might require a change to their configuration to support a dual adapter environment so that both adapters will detect all the disk units in the enclosure and also the other adapter. Refer to specific disk enclosure documentation for details.

   Note: Some disk enclosures may require a power cycle for the changes to take affect.

6. Power on system or partition and verify proper port cabling, and the function of adapters and disk arrays. Refer to "Viewing the Disk Array Configuration" on page 17. The output displayed will be similar to the following. - Verify that like bus numbered busses are connected on the primary and remote (secondary) adapters and the SCSI IDs are set to 05 and 06 on each adapter as highlighted below.

```
                              COMMAND STATUS

Command: OK            stdout: yes            stderr: no

Before command completion, additional instructions may appear below.


----------------------------------------------------------------------
Name       Location     State    Description        Size
----------------------------------------------------------------------
sisioa0    1V-08        Primary  PCI-X Dual Channel U320 SCSI RAID Adapter
 scsi4     1V-08-00-06,0 Linked  Target on remote adapter is 00-05,0
 scsi5     1V-08-01-06,0 Linked  Target on remote adapter is 01-05,0

hdisk4     1V-08-ff-0,0 Optimal  RAID 0 Array       35.3GB
 pdisk1    1V-08-00-1,0 Active   Array Member       35.3GB

hdisk5     1V-08-ff-0,1 Optimal  RAID 5 Array      141.1GB
 pdisk0    1V-08-00-0,0 Active   Array Member       35.3GB
 pdisk2    1V-08-00-2,0 Active   Array Member       35.3GB
 pdisk3    1V-08-00-3,0 Active   Array Member       35.3GB
 pdisk4    1V-08-00-4,0 Active   Array Member       35.3GB
 pdisk5    1V-08-00-8,0 Active   Array Member       35.3GB

 pdisk6    1V-08-00-10,0 Active   Array Candidate   71.4GB
 pdisk7    1V-08-00-11,0 Active   Array Candidate   71.4GB
 pdisk8    1V-08-00-12,0 Active   Array Candidate   71.4GB
 pdisk9    1V-08-00-13,0 Active   Array Candidate   35.3GB
 pdisk10   1V-08-00-14,0 Active   Array Candidate   71.4GB
```

7. You can decide to select one of the adapters in the high availability configuration to be the preferred primary adapter. This may be done for performance and usability reasons such as disk configuration changes. If you do not select one of the adapters to be the preferred primary adapter, the adapters will default to primary or secondary through a negotiation process during boot.

Things to consider when determining the preferred primary adapter:

- Because all Disk Array access must go through the Primary adapter, performance will be better for disk I/O operations from the system or partition containing the Primary adapter.
- All Disk Array configuration changes must be done on the system or partition containing the Primary adapter.
- Most disk service including error log analysis will be performed from the system or partition containing the Primary adapter. However, errors may also be presented by the Secondary adapter which may require actions on the system or partition containing the Secondary adapter.

To set the preferred primary mode:

a. At the command prompt, type `smit` and press Enter.
b. Select **Devices**.
c. Select **Disk Array**.
d. Select **IBM PCI-X SCSI Disk Array**.
e. Select **Change/Show PCI-X SCSI RAID Controller**.
f. Select the desired PCI-X Raid controller to be the preferred primary controller and press Enter.
g. The output displayed will be similar to the following. Use the F4 key to change the **Preferred operating mode**.

```
Change/Show Characteristics of a PCI-X SCSI RAID Controller

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                        [Entry Fields]
SCSI Adapter                            sisioa0
Description                             PCI-XDDR Quad Channel>
Status                                  Available
Location                                1A-08
Operating mode                          Primary Adapter
Preferred operating mode                [No Preference]
Link Operational                        No
Remote Adapter Vendor and Product ID
Remote Adapter Serial Number
```

8. Any new disk arrays can now be created from the primary adapter. See "Creating a PCI-X SCSI Disk Array" on page 16. If new disk arrays are created, the configuration manager must be run on the system or partition that contains the secondary adapter in order for it to recognize the new disk arrays.

9. If any configuration changes were made during the installation, for example in the previous step, or are made at a later time, the configuration manager must be run on the system or partition containing the secondary adapter in order for it to recognize the new configuration.

Chapter 5. PCI-X SCSI RAID Controller Maintenance

This chapter provides information to assist in ensuring optimal performance of your PCI-X SCSI RAID Controller.

Do's and Don'ts

To help avoid PCI-X SCSI RAID Controller and disk array problems, use the following tips :

- Always perform a normal system shutdown before physically replacing or moving the RAID adapter or members of disk arrays. A normal shutdown of the system will flush the adapter's write cache and remove dependencies between the adapter and the *pdisks*. Using the **rmdev** command on the controller (for example, `rmdev -Rl sisioa3`) has the same effect on the controller as it would on a single adapter when the system shutdown command is used.

   Note: *pdisks* that are a *Failed* member of a *Degraded* disk array can be replaced and the disk array reconstructed while the system continues to run.

- You can physically move *pdisks* from one adapter to another. However, if the *pdisks* are members of a disk array, be sure to move all the disks as a group. Prior to attempting a disk movement, ensure that the disk array is not in a *Degraded* state because of a disk failure.

- When physically removing *pdisks* that are members of a disk array and there is no need to preserve data and no intent to use the disk array again, delete the disk array before removing the disks. This action avoids disk-array related problems the next time that these disks are used.

- Always use the SCSI and SCSI RAID Hot Plug Manager to remove and replace a *pdisk*. For instructions on how to remove and replace a disk, see "Replacing pdisks" on page 58.

- If a disk array is being used as a boot device and the system fails to boot because of a suspected disk array problem, boot using the Standalone Diagnostic media. Error Log Analysis, AIX error logs, the PCI-X SCSI Disk Array Manager, and other tools are available on the Standalone Diagnostics to help determine and resolve the problem with the disk array.

- Do not attempt to correct problems by swapping adapters and disks unless you are directed to do so by the service procedures. Use Error Log Analysis to determine what actions to perform, and when appropriate, follow the appropriate MAPs for problem determination. If multiple errors occur at approximately the same time, look at them as a whole to determine if there is a common cause. For additional information regarding problem determination, see Chapter 6, "Problem Determination and Recovery," on page 61.

- Do not confuse the Cache Directory Card, which is a small rectangular card with a round button-shaped battery, for a cache card. The nonvolatile write cache memory is integrated into the main adapter card or integrated into the removable cache card on the newest adapters. The memory itself is not removable, and is battery-backed by the large, rechargeable Cache Battery Pack. The Cache Directory Card contains only a secondary copy of the write cache directory, no cache data, and should never be removed except under very specific cases as described in the MAPs.

- Do not unplug or exchange a Cache Battery Pack without following the procedures as outlined in this section or in the MAPs.

- When invoking diagnostic routines for a PCI-X SCSI RAID Controller, it is recommended that Problem Determination (PD) mode be used instead of System Verification (SV) mode unless there is a specific reason to use SV mode (for example, you were directed to run SV mode by a MAP).

- After diagnostic routines for a PCI-X SCSI RAID Controller are run in SV mode, it is recommended to run diagnostics in PD mode to ensure that new errors are analyzed. These actions should be performed especially when using Standalone Diagnostic media.

Updating the PCI-X SCSI RAID Controller Microcode

To determine if an update is needed for your PCI-X SCSI RAID controller, follow the directions at: http://techsupport.services.ibm.com/server/mdownload/download.html. If updates are needed, download instructions are also located at that Web address.

Note: For either an Auxiliary Cache adapter or a storage I/O adapter which attaches to an Auxiliary Cache adapter, it is recommended that both the storage I/O adapter and Auxiliary Cache adapter be updated.

To install the update to a PCI-X SCSI RAID Controller, do the following:
1. Type `smit` and press Enter.
2. Select **Devices**.
3. Select **Disk Array.**
4. Select **IBM PCI-X SCSI Disk Array.**
5. Select **Download Microcode to a PCI-X SCSI RAID Controller.**
6. Select your PCI-X SCSI RAID Controller.
7. Follow the instructions to complete the update.

Changing pdisks to hdisks

To change Array Candidate *pdisks* (522 bytes/sector) to standalone *hdisks* (512 bytes/sector), you must delete and format the *pdisks*.

Note: *pdisks* which are members of a disk array or are hot spares cannot be changed to standalone *hdisks*.

To change the *pdisks* to standalone *hdisks*, do the following:
1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Change/Show PCI-X SCSI pdisk Status**.
3. Select **Delete an Array Candidate pdisk and Format to 512 Byte Sectors**.
4. Select the appropriate PCI-X SCSI RAID Controller.
5. Select the 522 bytes/sector *pdisks* to be formatted to 512 bytes/sector standalone *hdisks*.

   Attention: Continuing with this option will format the disks. All data on the disks will be lost. When the format completes, the *pdisk* will be deleted and replaced by an *hdisk*.

Rechargeable Battery Maintenance

Attention: This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).

Rechargeable battery maintenance tasks described in this section include displaying rechargeable battery information, forcing a rechargeable battery error, and replacing the rechargeable Cache Battery Pack.

Displaying Rechargeable Battery Information

To display information about the controller's rechargeable battery, do the following:

1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps found in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Diagnostics and Recovery Options**.
3. Select **Controller Rechargeable Battery Maintenance**.
4. Select **Display Controller Rechargeable Battery Information**.
5. Select the PCI-X SCSI RAID Controller.
6. The screen displayed will be look similar to the following example:

```
                       COMMAND STATUS
Command: OK              stdout: yes            stderr: no

Before command completion, additional instructions may appear below.

RAID Adapter  . . . . . . . . . . . . : sisioa1
Battery Type  . . . . . . . . . . . . : Lithium Ion (LiIon)
Battery State . . . . . . . . . . . . : No battery warning/error
Power-on time (days)  . . . . . . . . : 256
Adjusted power-on time (days) . . . . : 288
Estimated time to warning (days)  . . : 512
Estimated time to error (days)  . . . : 768
Concurrently maintainable battery pack: Yes
Battery pack can be safely replaced . : No


F1=Help            F2=Refresh         F3=Cancel          F6=Command
F8=Image           F9=Shell           F10=Exit           /=Find
n=Find Next
```

The following are the fields displayed on the rechargeable battery information screen:

RAID Adapter
>   The name of the selected controller.

Battery Type
>   The type of rechargeable Cache Battery Pack.

Battery State
>   Indicates if an error condition currently exists related to the rechargeable Cache Battery Pack. The possible values for this field are:
>
>   **No battery warning/error**
>   >   No warning or error condition currently exists.
>
>   **Warning condition**
>   >   A warning condition currently exists and an error has been logged.
>
>   **Error condition**
>   >   An error condition currently exists and an error has been logged.

Unknown

Information is not available to determine whether a warning or error condition currently exists.

Power-on time (days)

Indicates the raw power-on time, in units of days, of the rechargeable Cache Battery Pack.

Adjusted power-on time (days)

Indicates the adjusted (prorated) power-on time, in units of days, of the rechargeable Cache Battery Pack.

Note: Some rechargeable Cache Battery Packs are negatively affected by higher temperatures and thus are prorated based on the amount of time that they spend at various ambient temperatures.

Estimated time to warning (days)

Estimated time, in units of days, until a message is issued indicating that the replacement of the rechargeable Cache Battery Pack should be scheduled.

Estimated time to error (days)

Estimated time, in units of days, until an error is reported indicating that the rechargeable Cache Battery Pack must be replaced.

Concurrently maintainable battery pack

Indicates that the rechargeable Cache Battery Pack can be replaced while the adapter continues to operate.

Battery pack can be safely replaced

Indicates that the adapter's write cache has been disabled and the rechargeable Cache Battery Pack can be safely replaced.

Forcing a Rechargeable Battery Error

This option allows you to force your PCI-X SCSI RAID Controller to put its Cache Battery Pack into an error state. This state requires replacement of the battery. Typically, performing this task ensures that write caching is stopped prior to replacing a Cache Battery Pack.

Forcing the battery error will result in the following:

- The system logs an error.
- Data caching becomes disabled on the selected controller.
- System performance could become significantly degraded until the Cache Battery Pack is replaced on the selected controller.

Ensure that you have the correct type and quantity of Cache Battery Packs to do the replacement. To resume normal operations, replace the Cache Battery Pack.

The Cache Battery Pack for 571F storage I/O adapter and 575B Auxiliary Cache adapter is contained in a single battery FRU that is physically located on the 575B Auxiliary Cache adapter. The functions of forcing a battery pack error and starting adapter cache on either adapter in the card set will result in the same function automatically being performed on the other adapter in the card set.

To force the battery error, do the following:

1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Diagnostics and Recovery Options**.
3. Select **Controller Rechargeable Battery Maintenance**.
4. Select **Force Controller Rechargeable Battery Error**.

5. Select your PCI-X SCSI RAID Controller whose battery you want to replace.

   Note: Using this option places the battery into the error state, which requires it to be replaced.
6. Press Enter.
7. Replace the Cache Battery Pack as described in "Replacing the Rechargeable Cache Battery Pack."

Replacing the Rechargeable Cache Battery Pack

This section provides instructions for replacing the rechargeable Cache Battery Pack.

Note: When replacing the Cache Battery Pack, the battery must be disconnected for at least 60 seconds before connecting the new battery. This is the minimum amount of time needed for the card to recognize that the battery has been replaced.

CAUTION:
The battery is a nickel metal hydride or lithium ion battery. To avoid possible explosion, do not burn the battery. Exchange it with the IBM approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call.

Attention: If the cache battery pack is not already in the error state, follow the steps described in "Forcing a Rechargeable Battery Error" on page 34 before proceeding to prevent data loss. After the battery is in the error state, replace it with a new battery by doing the following:

Procedure for non-Concurrently Maintainable Battery Packs

Use the following procedure for replacing non-concurrently maintainable battery packs

1. Determine that it is safe to replace the Cache Battery Pack. Refer to "Displaying Rechargeable Battery Information" on page 33. It is safe to replace the Cache Battery Pack when the display lists **Yes** next to **Battery pack can be safely replaced**.
2. Remove the adapter from the system, see your system documentation for instructions.
3. Use the following illustration to locate the type of adapter on which you are working.
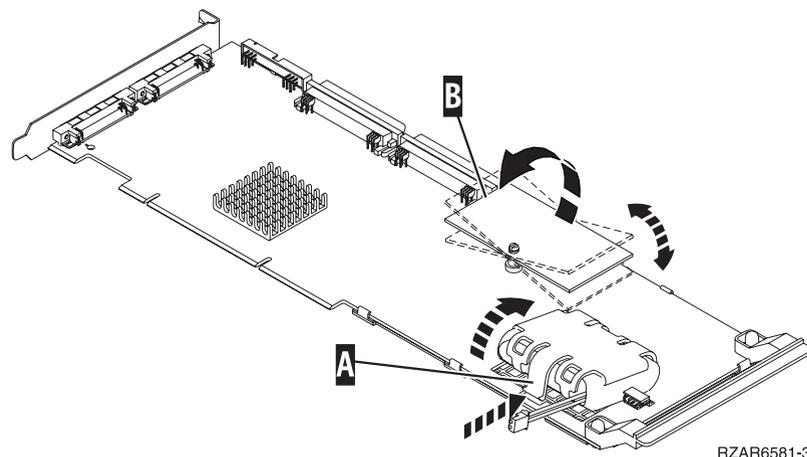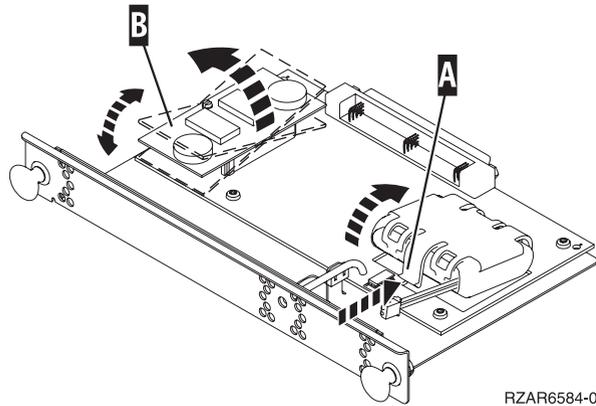
RZAR6581-3

RZAR6584-0
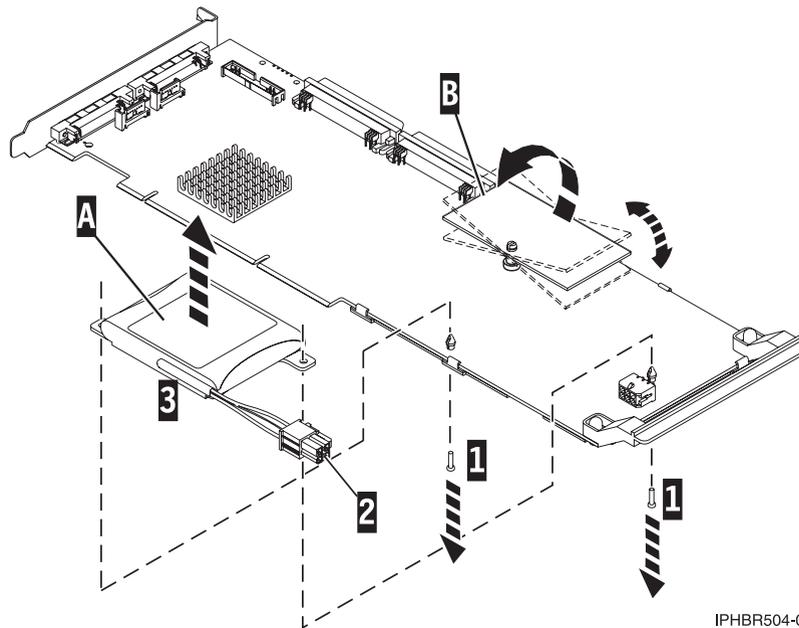
IPHBR504-0

A Battery Pack Casing
B Cache Directory Card (do not disturb, shown here for identification purposes only)
1 Plastic rivet pins

4. Remove the plug that connects the battery unit and the card.

   Note: The plug fits in the board only one way so that it cannot be inserted incorrectly during the replacement steps of this procedure.

5. Do one of the following:
   - For cards that have a hard plastic casing around the battery, squeeze the casing **A** to remove the battery unit from the card. Remove the battery unit from the battery casing. Save the battery casing to reuse on the replacement battery pack.
   - For batteries that are attached to the card with rivets, work from the back side of the card to remove the two pins (**1**) from the inside of the rivets. Working from the front side of the card, release the plastic rivets. A hint in releasing the plastic rivets is to use a medium retractable ball point pen, in the retracted position, to push the rivets through the tab of the battery pack.

     Note: Ensure that the Cache Battery Pack is disconnected for at least 60 seconds before connecting the new battery. This is the minimum amount of time needed for the card to recognize that the battery has been replaced.

6. Install the new Cache Battery Pack by reversing this procedure.

Procedure for Concurrently Maintainable Battery Packs

Use the following procedure for concurrently maintainable battery packs.

1. Determine that it is safe to replace the Cache Battery Pack. Refer to "Displaying Rechargeable Battery Information" on page 33. It is safe to replace the Cache Battery Pack when the display lists **Yes** next to **Battery pack can be safely replaced**.

2. Use the following illustration to locate the metal cover **A** that holds the battery pack. Pull out on the push-rivet to release the metal cover **A** .

RZAR6585-2

IPHBR612-1
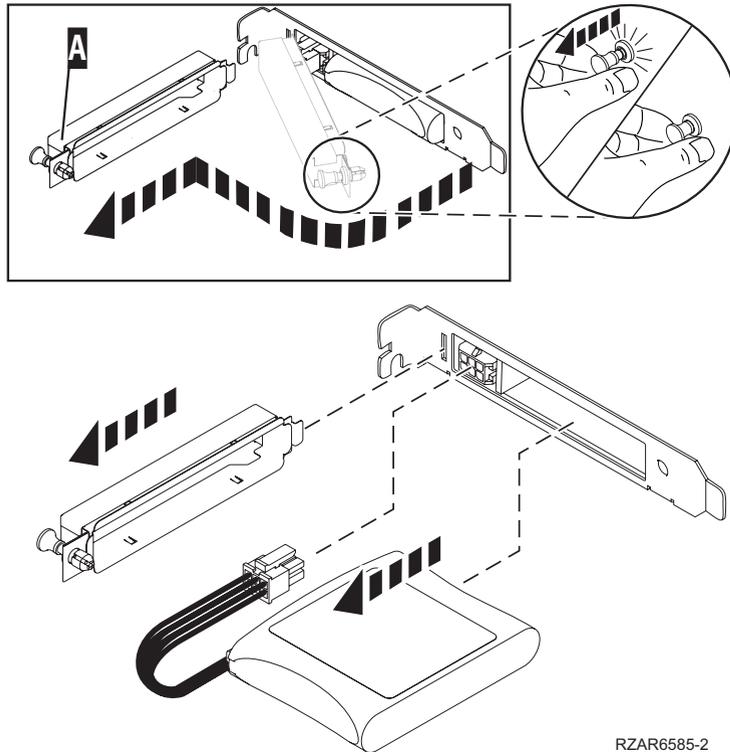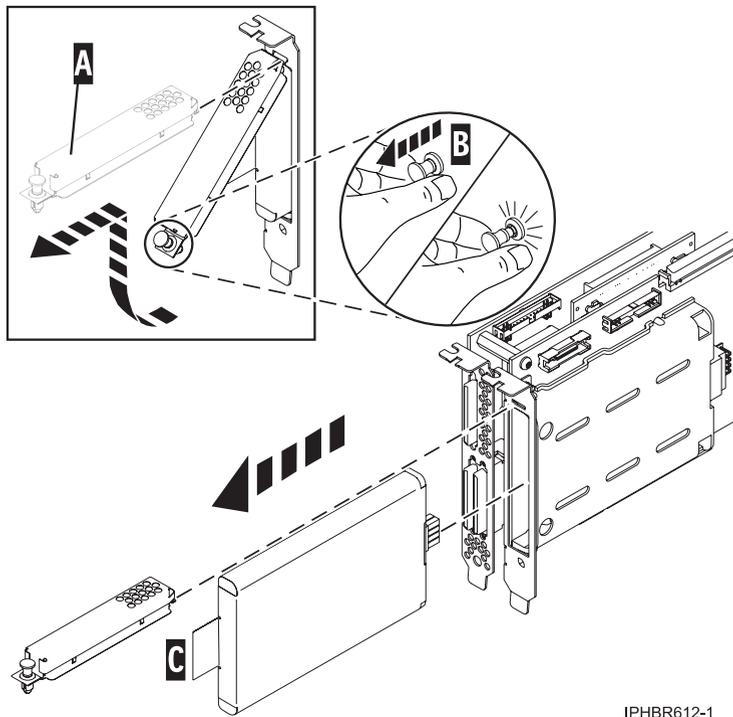
A Battery Pack Cover

3. For all except the 571F/575B card set, remove the plug that connects the battery unit and card by squeezing the retaining latch while gently pulling on the plug.

   Note: The plug fits in the board only one way so that it cannot be inserted incorrectly during the replacement phase.

4. Remove the battery unit by gently pulling it out.

5. Install the new Cache Battery Pack by reversing this procedure.

6. Restart the adapter's write cache by doing the following:
   • Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
   • Select **Diagnostics and Recovery Options**.
   • Select **Controller Rechargeable Battery Maintenance**.
   • Select **Start Adapter Cache**.
   • Select your PCI-X SCSI RAID Controller whose battery you just replaced.
   • Press Enter.

Separating a Removable Cache Card From the Base Card on Type 2780 and 571E Cards

To complete this procedure, you will need a T-10 TORX tool.

Attention: To avoid loss of cache data, do not disconnect the cache battery during this procedure.

All cards are sensitive to electrostatic discharge. See "Handling Static Sensitive Devices" on page ix.

Attention: This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).

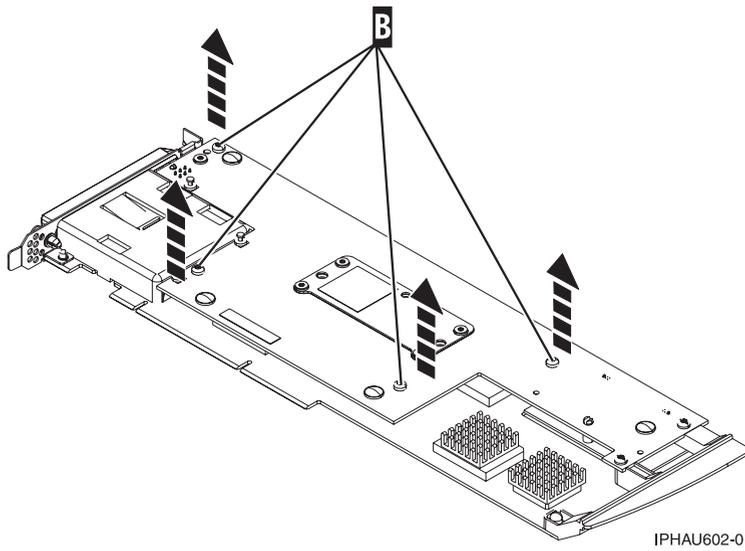To separate the removable cache card from the base card, do the following:

1. Label both sides of the card before separating the cache card from the base card.

2. Are you servicing a 2780 adapter or a 571E adapter?

   - If you are servicing a 2780 adapter, go to step 3.
   - If you are servicing a 571E adapter, go to step 9 on page 41.

3. Place the 2780 adapter on an ESD protective surface and orientate it as shown in step 4. Note the position of the card handle.

4. Remove three screws **A** from the base card as shown. Note the thread differences of the screws for correct replacement in the following steps. The two screws that were secured into the plastic handle have a different thread pattern from the machine thread screw that was removed from the front of the card.

IPHAU600-1

5. Turn the card over.

IPHAU601-1
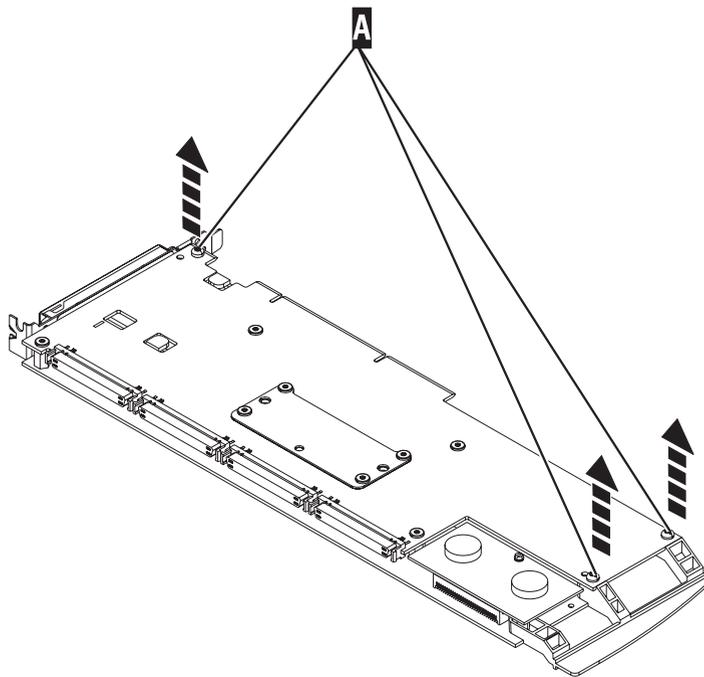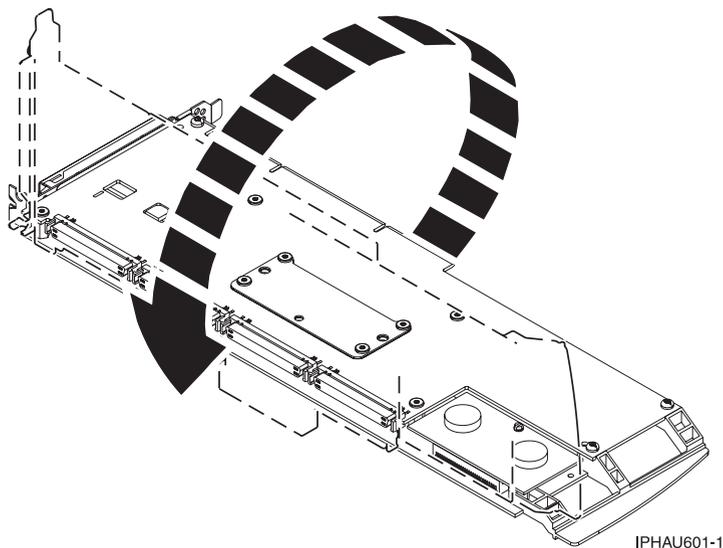
6. Remove four screws [B] from the removable cache card as shown.

IPHAU602-0

7. To separate the cards, turn the jack screw counterclockwise until you feel the cards separate.

IPHAU603-0
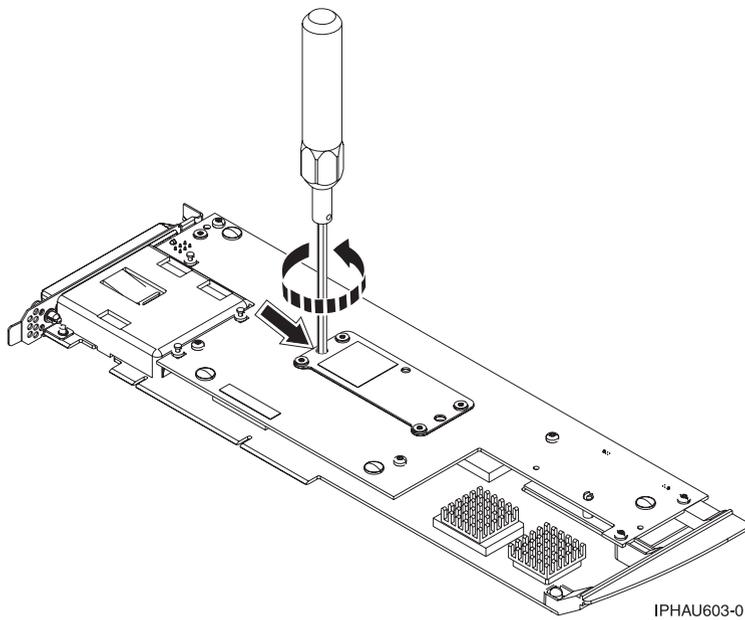
8. Separate the base card from the removable cache card.

   Important: After separating the base card from the cache card, turn the jack screw clockwise to reseat it. then go to "Attach the Removable Cache Card to the Base Card" on page 45.
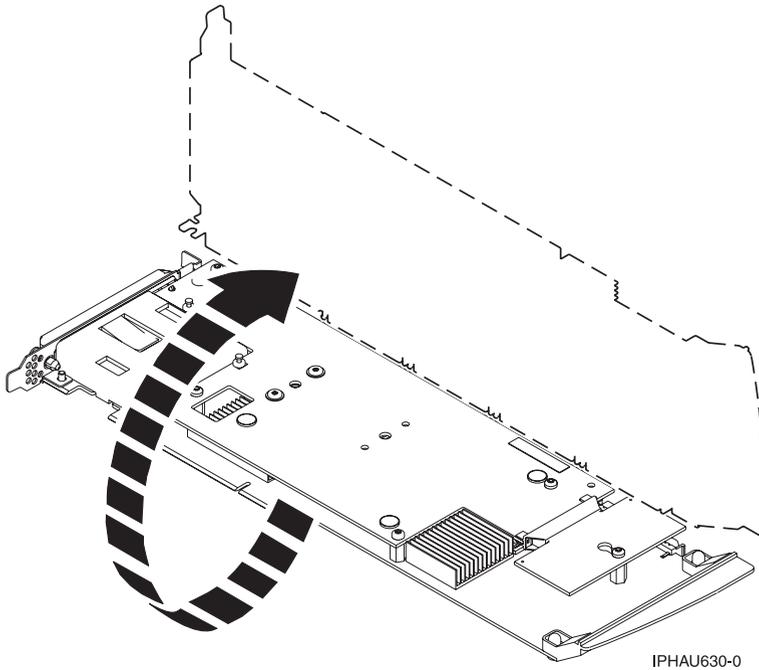

IPHAU604-0
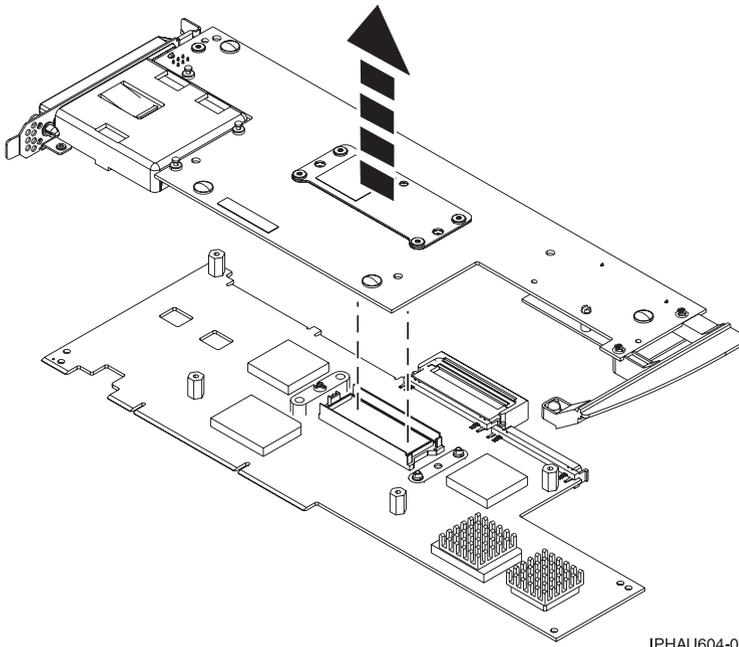
   This ends this procedure.

9. Place the 571E adapter on an ESD protective surface and orientate it as shown in step 10. Note the position of the card handle.

10. Remove screw **A** from the base card as shown.
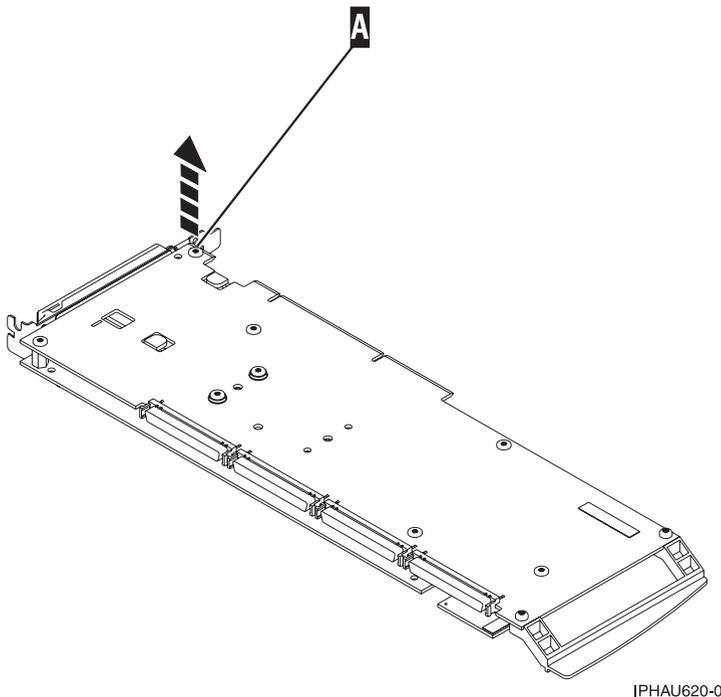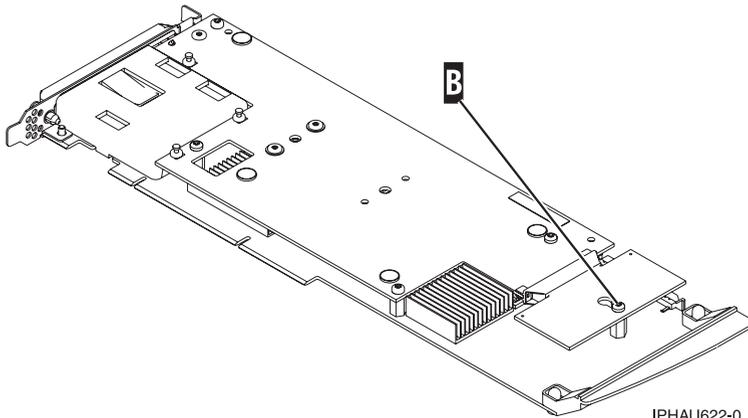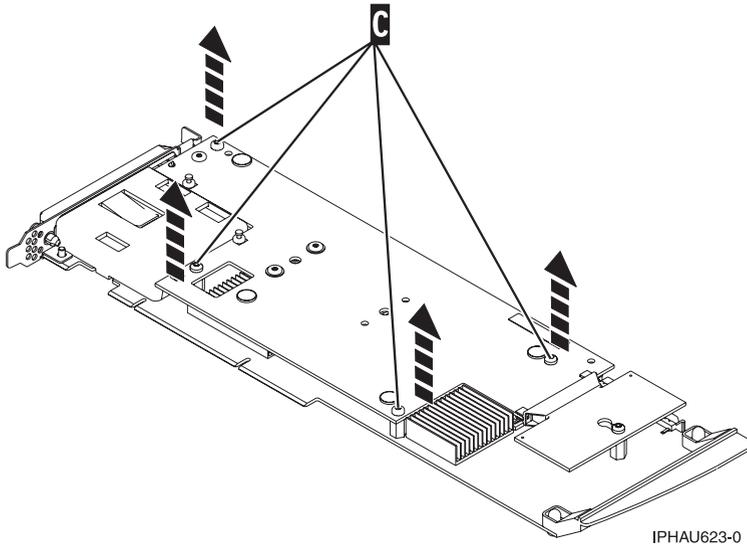

IPHAU620-0

11. Turn the card over.

IPHAU630-0

12. Loosen screw **B** one full turn. Do not remove this screw.

IPHAU622-0

13. Remove the four screws **C** from the removable cache card as shown.

> **Note:** Observe the keyhole slot located below screw **B** . In step 15 on page 44 when asked to separate the cache card, you will need to have the keyhole slot clear the retaining screw.

IPHAU623-0

14. Turn the jack screw counterclockwise to begin separating the cards. Continue turning the jack screw until you feel the cards separate.

IPHAU624-0
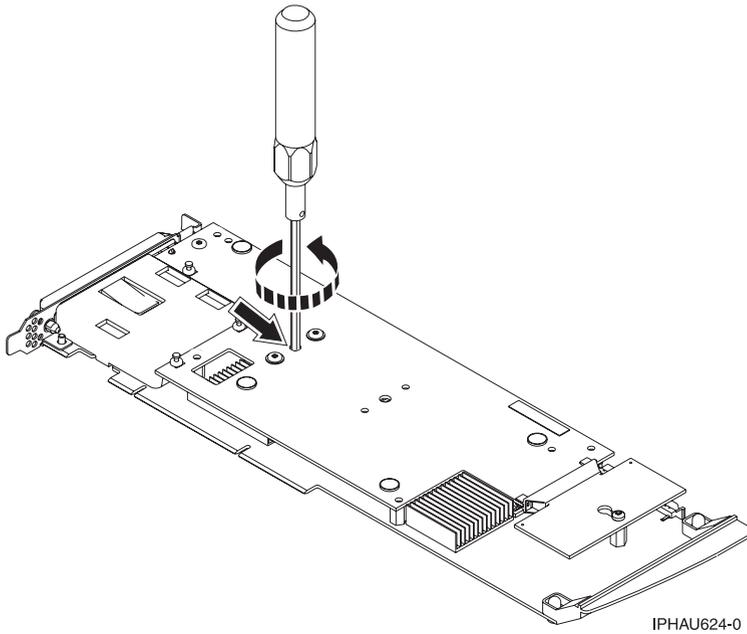
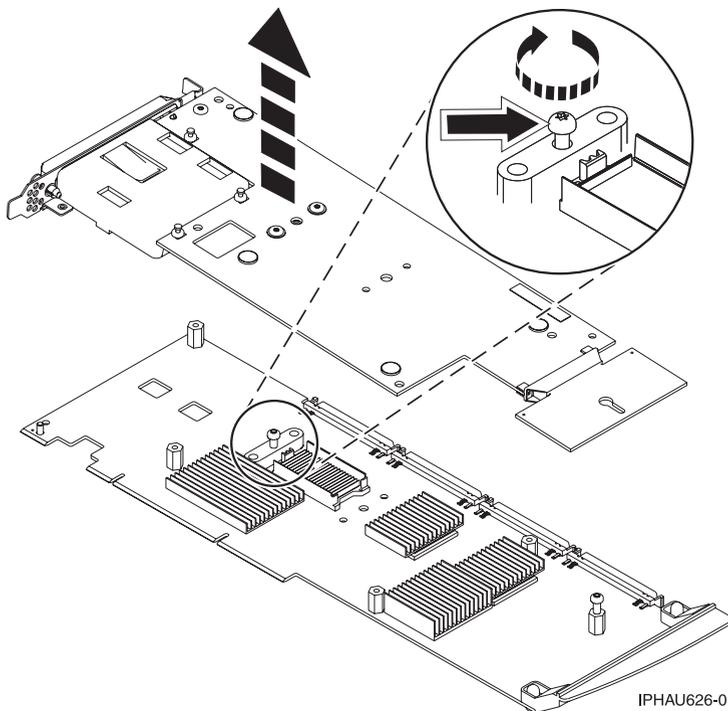15. Carefully slide the removable cache card towards the plastic handle. Slightly lift up the back portion of the cache card freeing the cache directory card's keyhole slot from its retaining screw.

IPHAU625-0
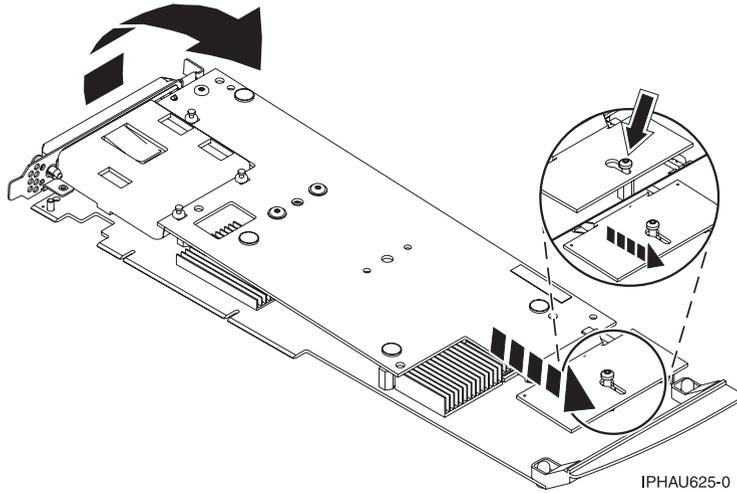
16. Separate the base card from the removable cache card.

   Note: After separating the base card from the cache card, turn the jack screw clockwise to reseat it.
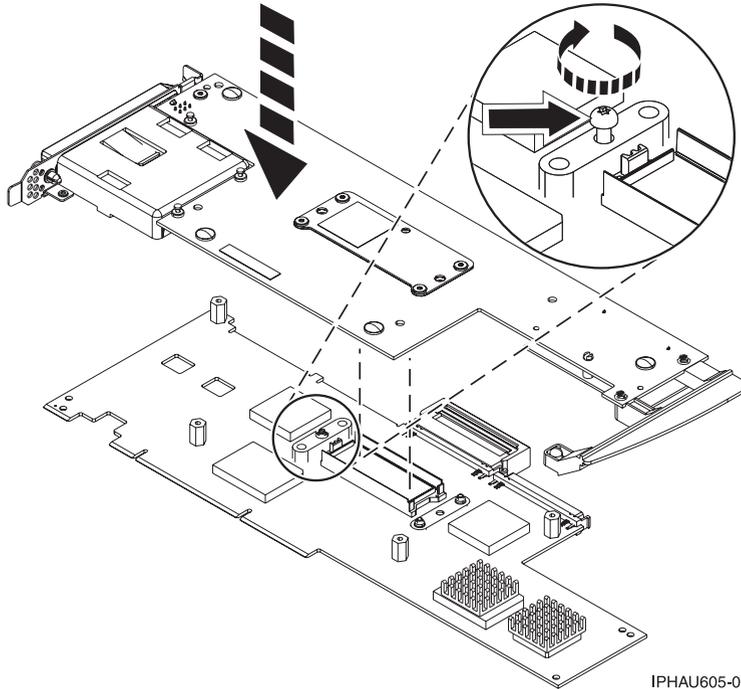
IPHAU626-0

   This ends this procedure.

Attach the Removable Cache Card to the Base Card

To attach the removable cache card to the base card, do the following:

1. Make sure the jack screw is turned clockwise until it is seated. If the jack screw is not fully seated, the cards will not connect together properly and damage may occur.

IPHAU605-0

2. Are you servicing a 2780 adapter or a 571E adapter?
   - If you are servicing a 2780 adapter, go to step 3.
   - If you are servicing a 571E adapter, go to step 7 on page 47.
3. Align then attach the 2780 base card to the removable cache card by doing the following:
   a. Align the card connector.

      Note: To help align the card connector, view the jack screw through the hole.
   b. Press down on the stiffener plate to fully seat the connector.

4. Attach four screws A to the removable cache card as shown.

IPHAU606-0

5. Turn the card over.

IPHAU607-0
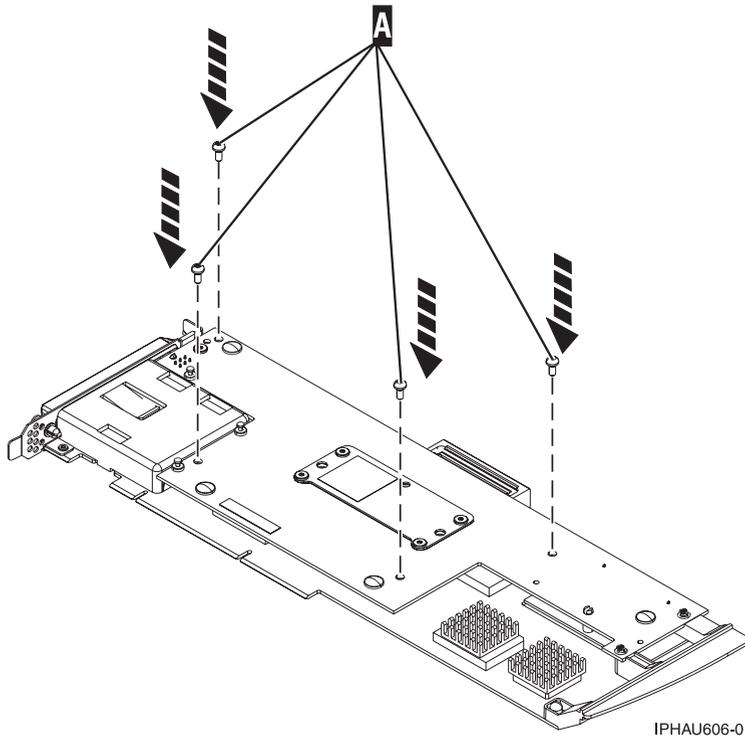
6. Attach the three screws **B** to the base card as shown. Note the thread differences in the screws as pointed out in step 4 on page 39 of the previous separation procedure.
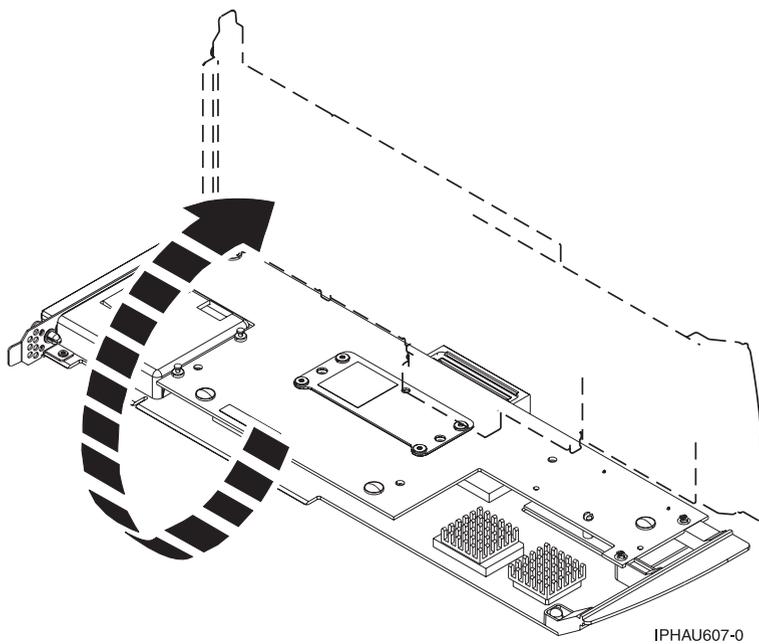

IPHAU608-1
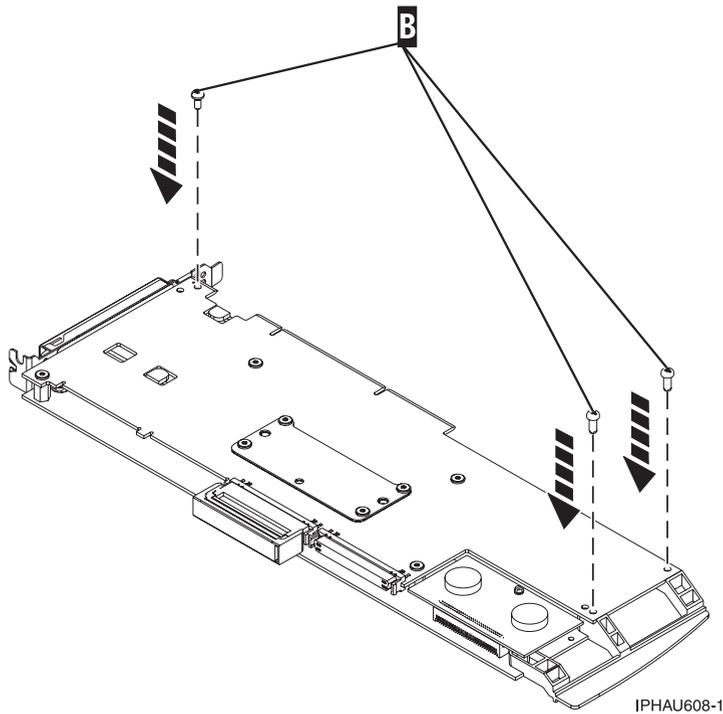
This ends this procedure.

7. Attach the 571E base card to the removable cache card by doing the following:

   a. Align the removable cache card with the base card and place the hole in the cache directory card over the screw **B** .

   b. Slide the removable cache card away from the plastic handle to line up the card connector. (view the jack screw through the hole to help alignment).

IPHAU628-0
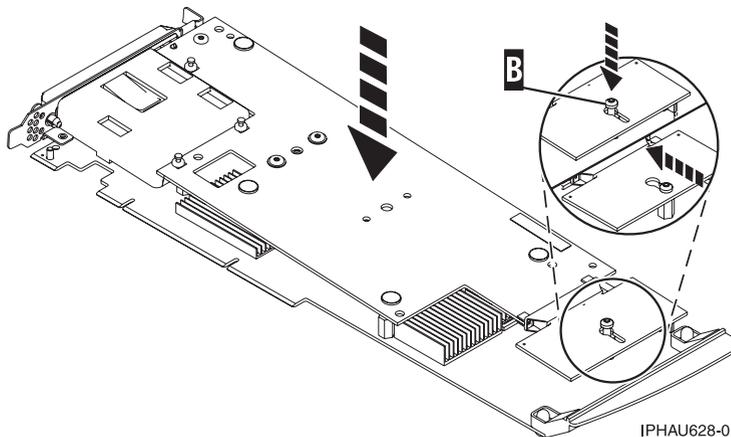
   c. Press down on the area of the card over the connector to fully seat the connector.

8. Attach the four screws $\boxed{C}$ to the removable cache card as shown.

IPHAU629-0
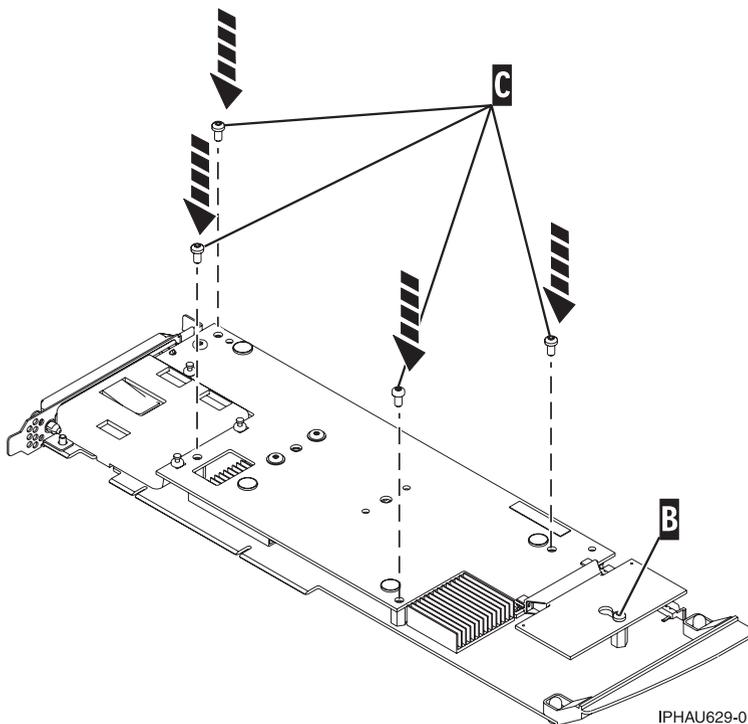
9. Tighten screw $\boxed{B}$ on the cache directory card.

IPHAU628-0
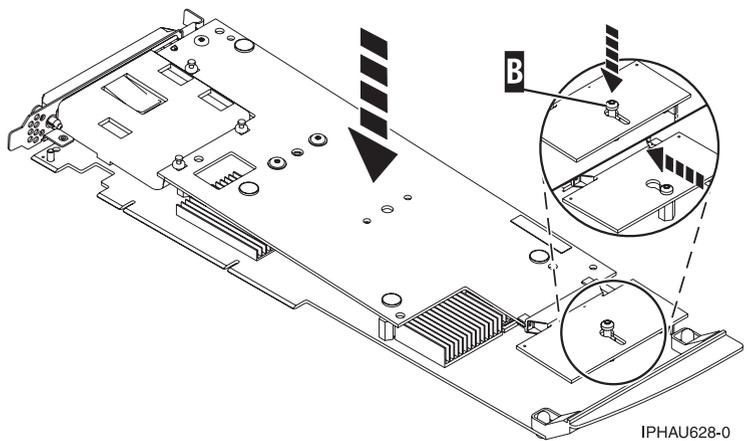
10. Turn the card over.

IPHAU630-0
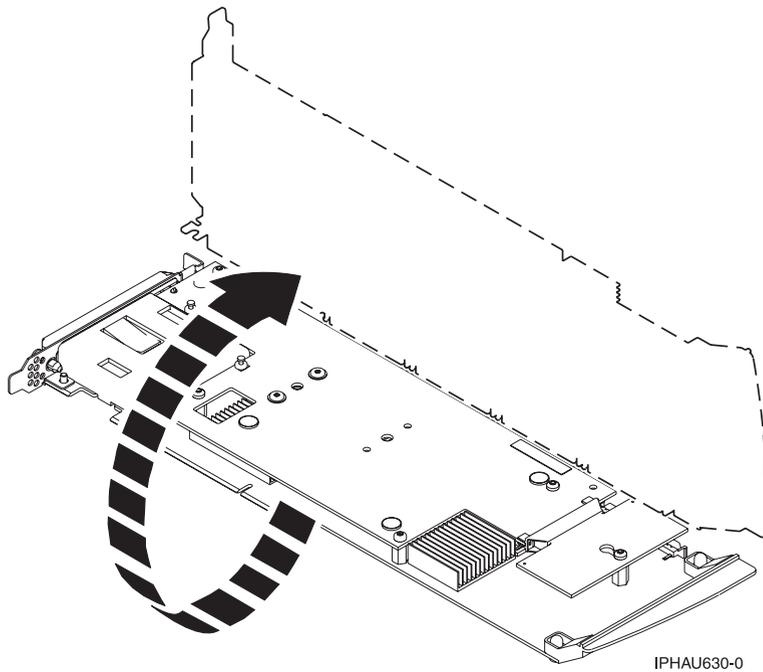
11. Attach the screw A to the base card.

A

IPHAU631-0
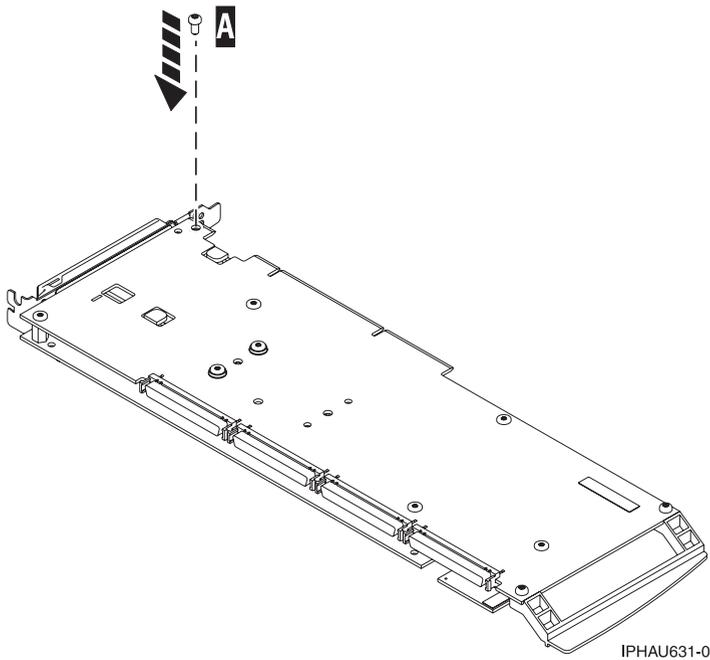
This ends this procedure.

Replacing the Cache Directory Card

Attention: This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).

To replace the Cache Directory card, do one of the following:
- For non-571E cards, go to "To replace a non-571E cache directory card" on page 50.

- For 571E cards, go to "To replace the 571E Cache Directory card" on page 52.

Attention: All cards are sensitive to electrostatic discharge. See "Handling Static Sensitive Devices" on page ix.

- To replace a non-571E cache directory card, do the following:
  1. Remove the card using the card remove and replace procedure for the model or expansion unit on which you are working.
  2. Locate the cache directory card **B** . It is a small rectangular card mounted on the I/O card (see the following figures for your adapter).
  3. Unseat the connector on the cache directory card by wiggling the two corners furthest from the mounting peg using a rocking motion. Then, pivot the cache directory card back over the mounting peg to disengage the cache directory card from the mounting peg.
  4. Install the replacement card by seating it on the connector and mounting peg.

RZAR6581-3
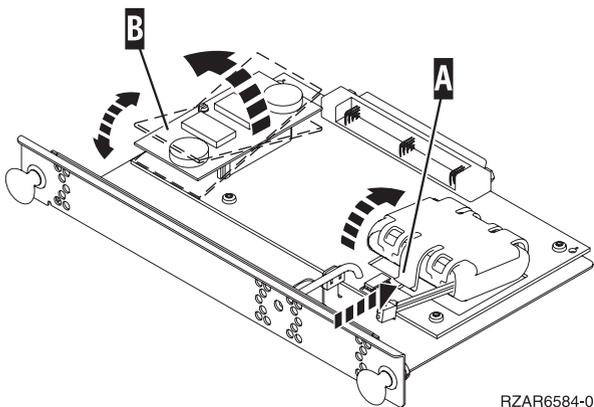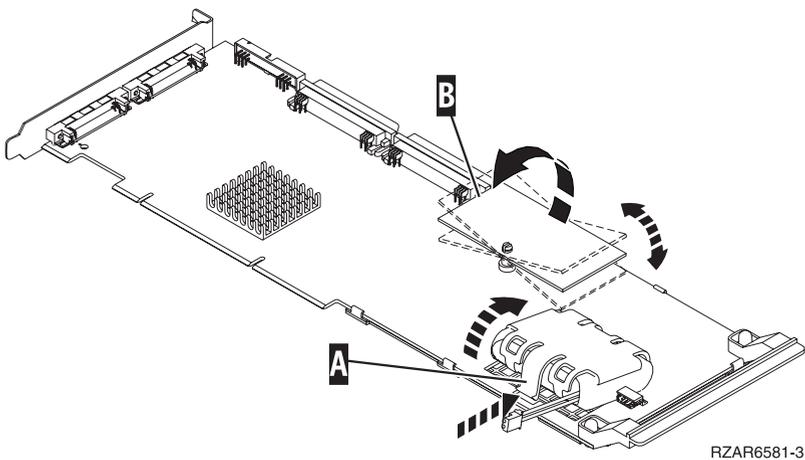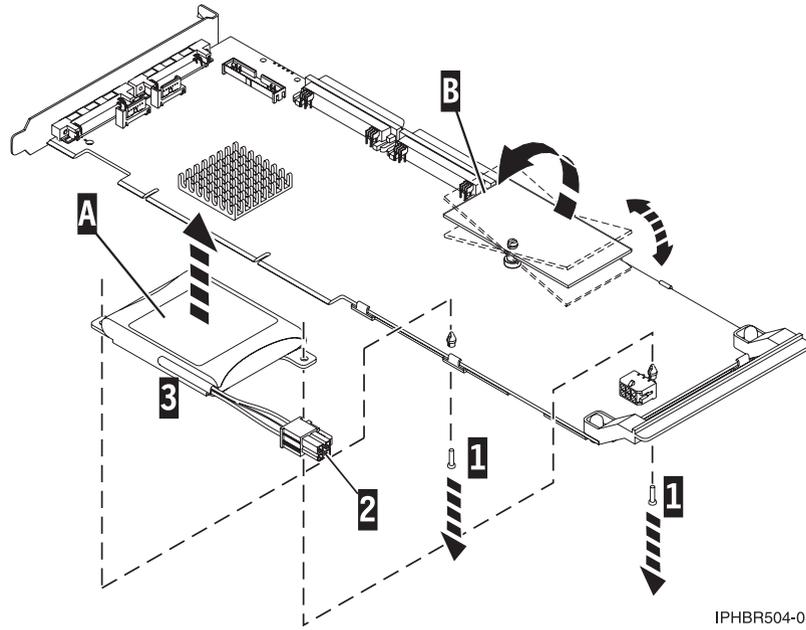
RZAR6584-0

IPHBR504-0
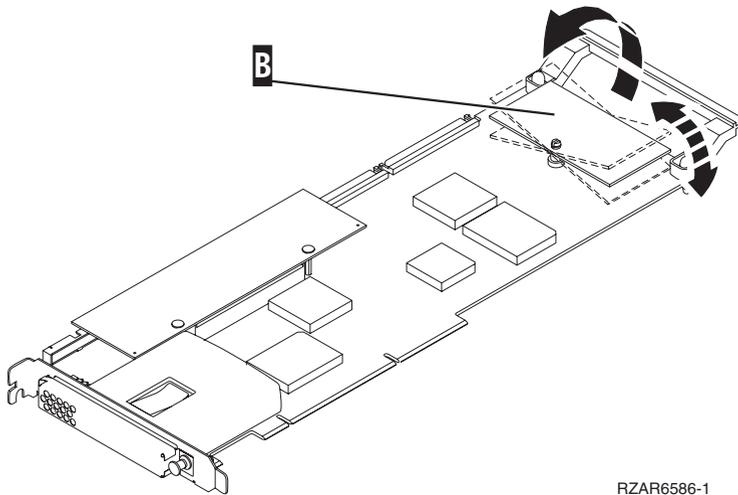
RZAR6586-1
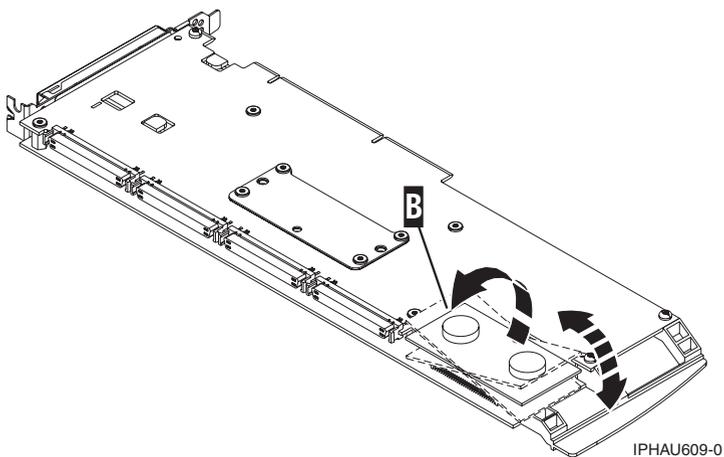
IPHAU609-0

5. **This ends the procedure.**
- To replace the 571E Cache Directory card, do the following:
   1. Locate the retaining screw for the cache directory card **B** .

IPHAU622-0
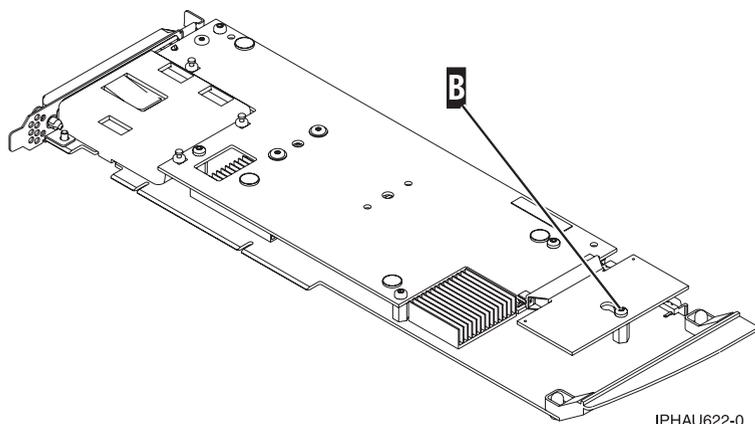
   2. Loosen the retaining screw located on the cache directory card one full turn. Do not remove the retaining screw.
   3. Unseat the cache directory card from the connector. To unseat and remove the cache directory card, pull the cache directory card in the direction of the plastic card handle. You might need to pull the card on one side then the other to help disengage the connector. When you have disengaged the connector, lift the cache directory card up and off the retaining screw.
   4. Install the replacement cache directory card. To install the replacement cache directory card, place it down over the retaining screw, then slide it away from the plastic card handle to fully seat it in the connector.

IPHAU628-0
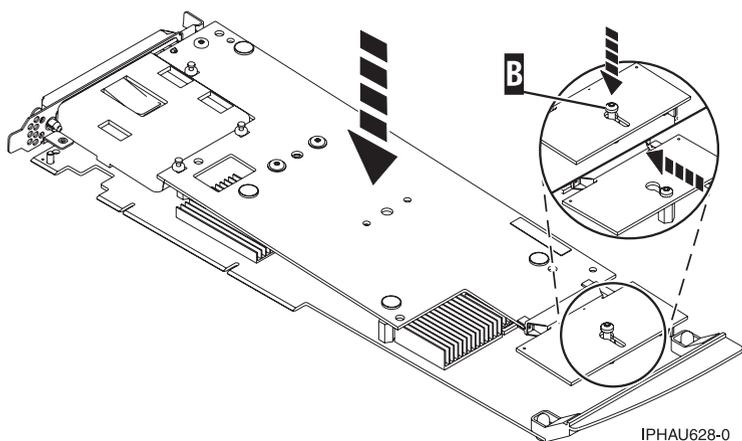
   5. Tighten the retaining screw.
   6. **This ends the procedure**.

Separating the 571F/575B Card Set and Moving the Cache Directory Card

Detailed instructions for separating the card set and moving the cache directory card. To avoid loss of cache data, you *must not* remove the cache battery during this procedure.

Note:

1. This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).
2. If you are removing the adapter from a double-wide cassette, go to the procedures in your system unit's service information for removing a double-wide adapter from a double-wide cassette.

To complete this procedure, you will need a T-10 TORX bit or driver.

Attention: To avoid loss of cache data, you *must not* remove the cache battery during this procedure.

To separate the 571F/575B card set and move the cache directory card, do the following:

Attention: All cards are sensitive to electrostatic discharge. See "Handling Static Sensitive Devices" on page ix before beginning this procedure.

1. Label both sides of the card before separating.
2. Go to the next step if you are not servicing a 571F/575B card set containing a light pipe assembly. If you are servicing a 571F/575B card set that contains a light pipe assembly, you will need to remove the light pipes. To remove the light pipes, do the following:
   a. Place the 571F/575B card set adapter on an ESD protective surface.
   b. Remove the light pipe retaining screw **C** from the 571F/575B card set.
   c. Slide the light pipe assembly **D** from between the 571F/575B cards.
   d. Put both the screw and the light pipe assembly in a safe place.
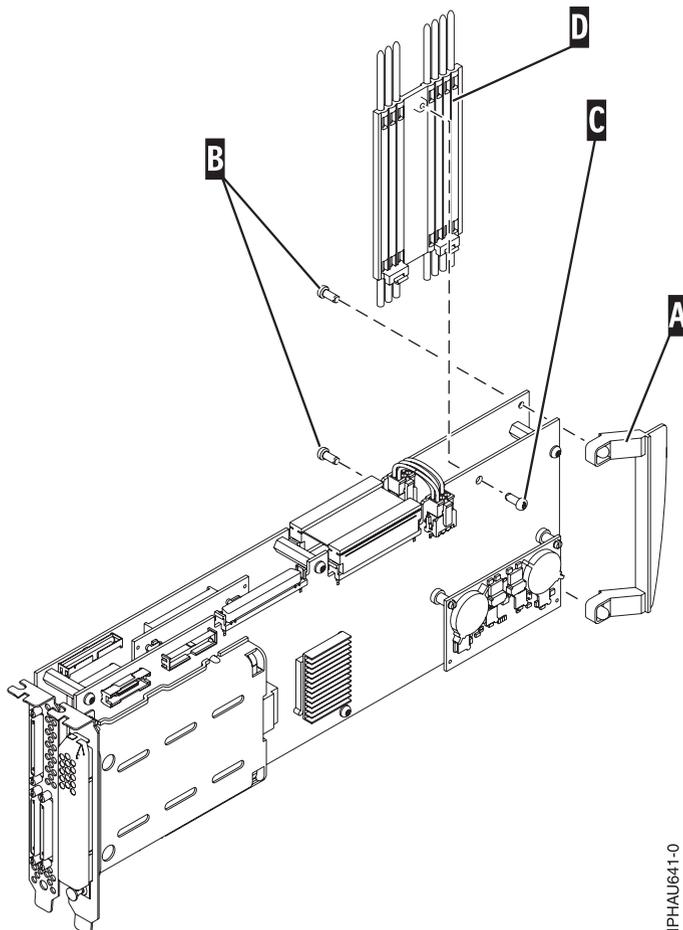
IPHAU641-0

3. Place the 571F/575B card set adapter on an ESD protective surface and orient it as shown.

IPHAU635-1
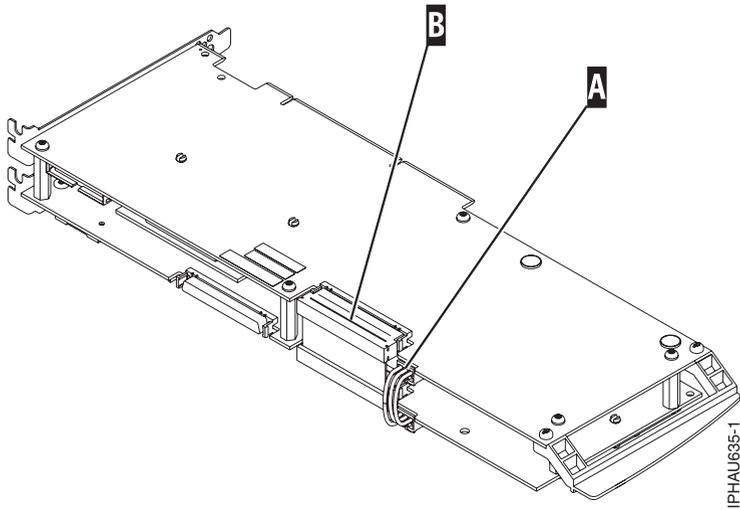
4. Disconnect battery cable **A** and SCSI cable B from the 571F storage adapter. Leave the other end of the cables attached to the 575B auxiliary cache adapter.

IPHAU636-1
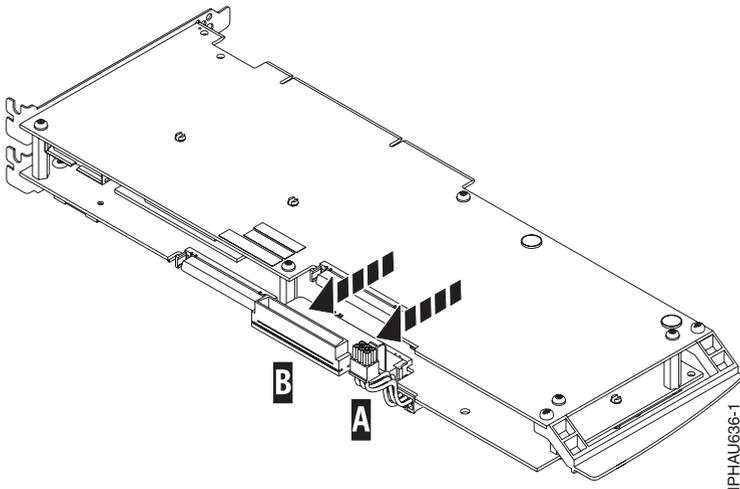
5. To prevent possible card damage, first loosen all 5 retaining screws **C** BEFORE removing them. After all five retaining screws have been loosened, remove the screws **C** from the 571F storage adapter.

IPHAU637-1
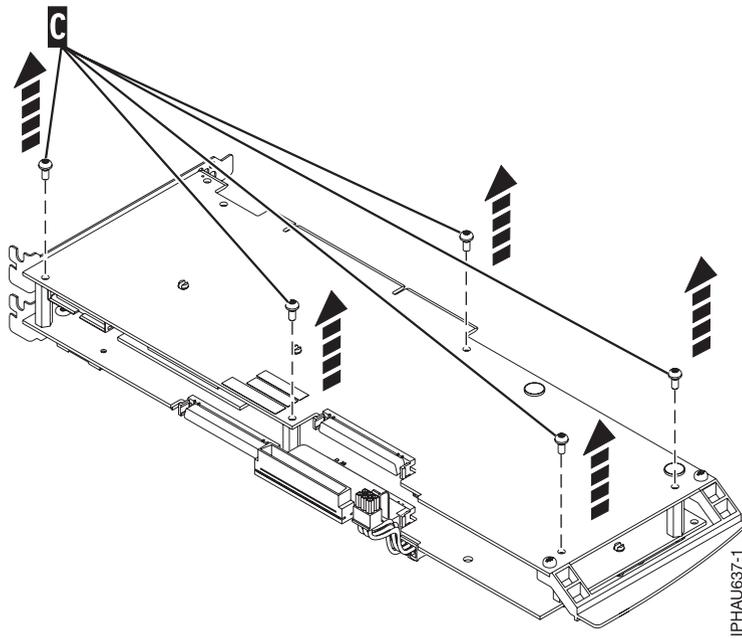
6. Carefully lift the 571F storage adapter off the standoffs and set it on the ESD protective surface.

IPHAU638-2
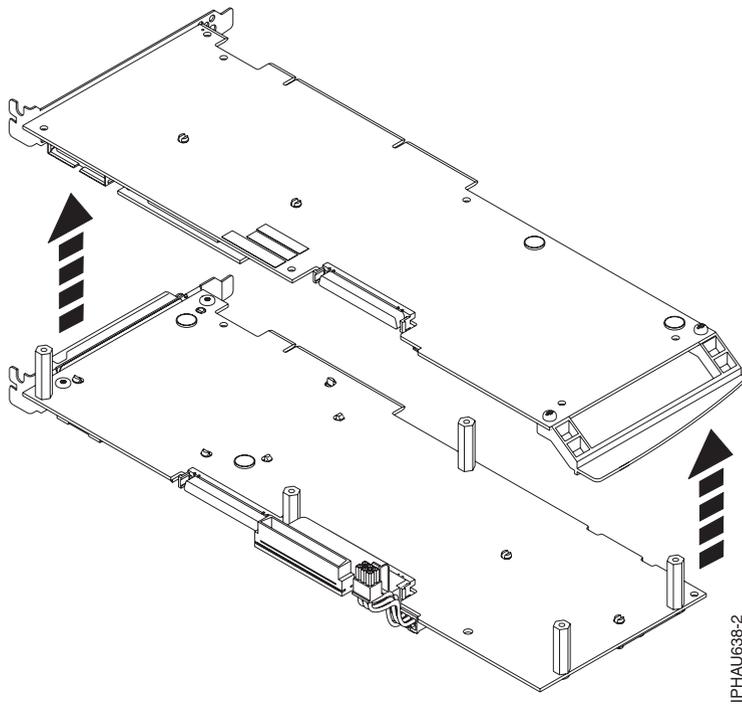
7. Turn the 571F storage adapter over so the components are facing up and locate the cache directory card **D** on the 571F storage adapter. It is the small rectangular card mounted on the I/O card.

IPHAU640-0
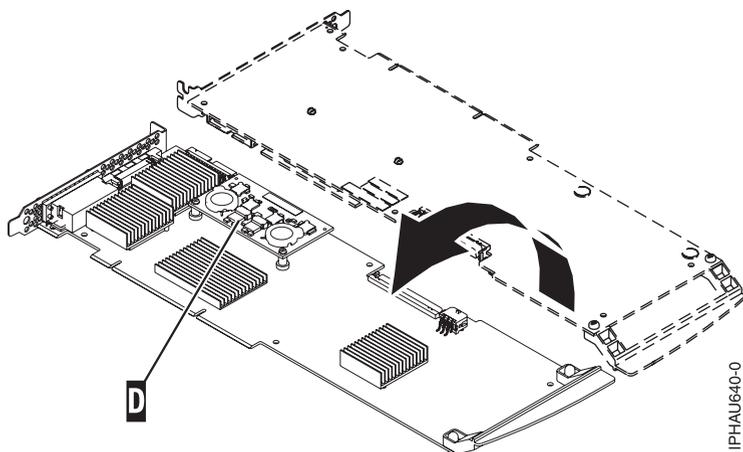
8. Unseat the connector on the cache directory card by wiggling the two corners that are farthest from the mounting pegs. To disengage the mounting pegs, pivot the cache directory card back over the mounting peg.

IPHAU639-2
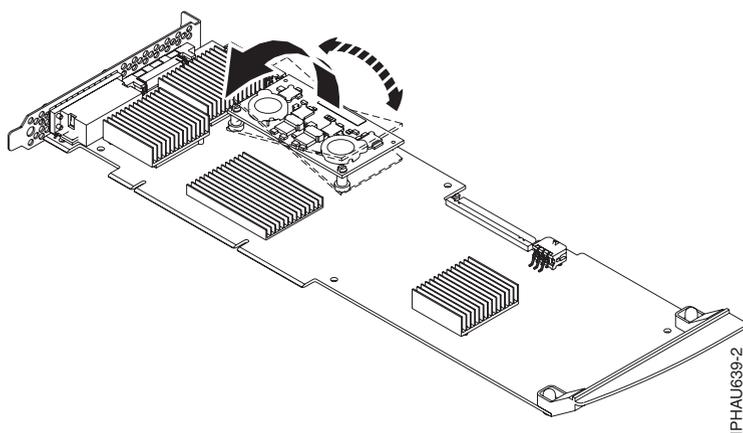
9. Move the cache directory card to the replacement 571F storage adapter and seat it on the connector and mounting pegs.

10. To reassemble the cards do the above procedure in reverse order. To prevent possible card damage, insert all 5 screws **C** before tightening any of them.

> **Note:** If you are installing the 571F/575B card set adapter into a cassette, do the following:
>
> a. Remove the adapter handle **B** as shown in the following illustration.
>
> b. If applicable, remove the light-pipe assembly.
>
> c. Remove the two plastic covers from the SCSI jumper cable **A** as shown in the following figure:
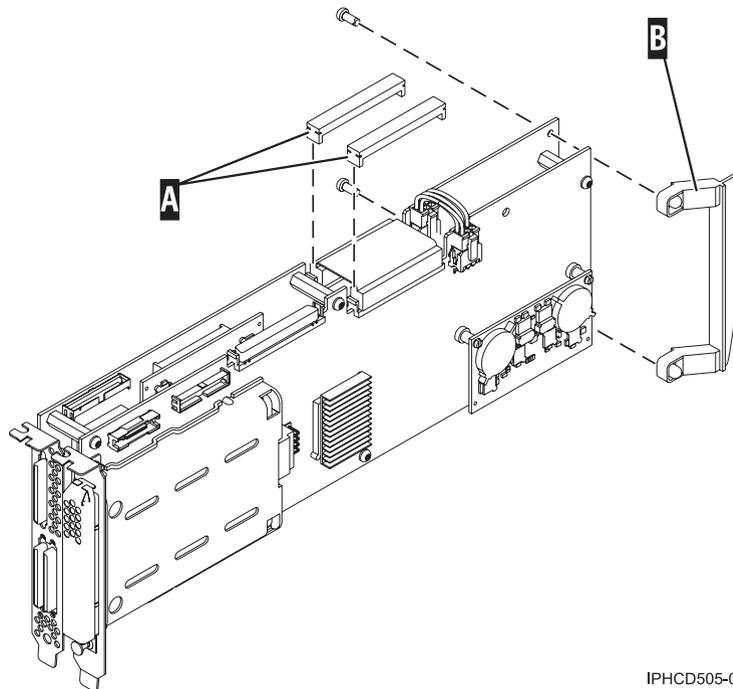
IPHCD505-0

    d. Ensure that battery cable **A** and SCSI cable **B** are reconnected to the 575B auxiliary cache adapter.

    e. If you removed the double-wide PCI adapter from a cassette in the beginning of this procedure, reinstall the adapter into the double-wide cassette to complete the installation. Refer to the procedures in your system unit's service information for installing a double-wide adapter in a double-wide cassette.

11. If applicable, to replace the light pipes, do the following:

    a. Ensure that the 571F/575B cards are correctly assembled.

    b. If necessary, remove the 571F/575B card handle **A** from the back of the card set. If you misplace the light-pipe retaining screw **C**, use either of the screws that secure the card handle **B** to the card.

    c. Slide the light pipe assembly **D** between the 571F/575B cards. Ensure that the threaded screw hole located on the light pipe assembly is aligned with the screw clearance hole located on the card.

IPHAU641-0
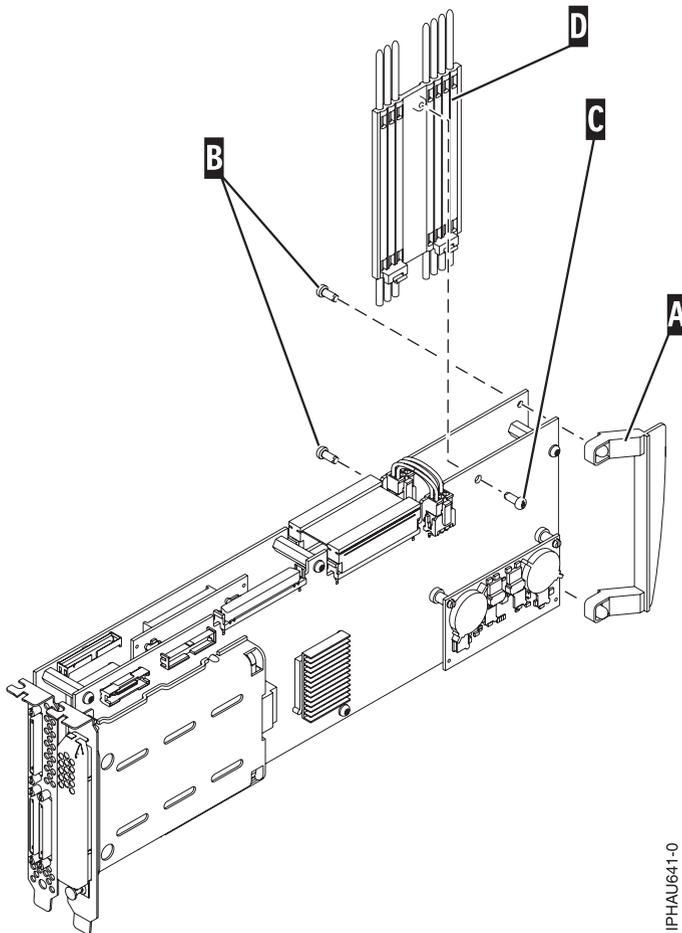
12. Return to the procedure that sent you here. **This ends this procedure.**

Replacing pdisks

Replace *Failed pdisks* as soon as possible, even if a reconstruction was initiated with a hot spare by the controller. The **Replace/Remove a Device Attached to an SCSI Hot Swap Enclosure Device** option in the SCSI and SCSI RAID Hot Plug Manager can be used to replace *Failed* pdisks. The PCI-X SCSI Disk Array Manager provides a shortcut to the SCSI and SCSI RAID Hot Plug Manager. The SCSI and SCSI RAID Hot Plug Manager allows you to perform the following tasks:

* List all SCSI hot swap slots and their contents
* Set the identify indication for a device attached to a SCSI hot swap enclosure device
* Add a device to a SCSI hot swap enclosure device
* Replace and remove a hot swap device
* Configure added and replaced devices

To start the SCSI and SCSI RAID Hot Plug Manager, do the following:

1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Diagnostics and Recovery Options**.
3. Select **SCSI and SCSI RAID Hot Plug Manager**.
4. Select the action you want to take, and follow the instructions on the screen to complete the selected operation.

Attention: Always use the SCSI and SCSI RAID Hot Plug Manager for devices attached to a PCI-X SCSI RAID Controller. *Do not* use utilities intended for other RAID products, such as RAID Hot Plug Devices.

Recovery Procedures for Disk Failures

The PCI-X SCSI RAID Controller handles disk failures differently, depending on the RAID level involved. For a description of disk array states, *pdisk* states, and RAID levels, see "Supported RAID Levels" on page 3. This section describes recovery procedures for each RAID level.

RAID Level 0

A single disk failure will cause a RAID level 0 array to change to the *Failed* state. After deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group, delete the disk array, replace the *Failed* disk, then re-create the disk array. You must then add the disk array back to a volume group, re-create the disk array logical volumes and file systems, and copy the data back to the restored disk array from your backup media.

RAID Level 5

Single-Disk Failure
Whenever a single disk 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* disk can be reconstructed using parity and data on the remaining disks.

Whenever a disk is marked as *Failed*, replace it as soon as possible. If a hot spare disk was defined, the controller may reconstruct the data on the disk automatically. If a hot spare reconstruct was not initiated by the controller, replace the *Failed* disk and then initiate a reconstruct.

To initiate a reconstruct, do the following:
1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Reconstruct a PCI-X SCSI Disk Array**.
3. Before you proceed, ensure that you have replaced the *Failed* disk.
4. Select the *pdisk* that you want to reconstruct.

Multiple-Disk Failures
If a second disk in a RAID level 5 disk array fails, the array will be marked as *Failed* and the data will not be accessible. After deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group, delete the disk array, replace the *Failed* disks, then re-create the disk array. You must then add the disk array back to a volume group, re-create the disk array logical volumes and file systems, and copy the data back to the restored disk array from your backup media.

RAID Level 6

Single-Disk or Dual-Disk Failure
Whenever one or two disks in a RAID level 6 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 disk can be reconstructed using "P" and "Q" parity and data on the remaining disks.

Whenever a disk is marked as Failed, replace it as soon as possible. If a hot spare disk was defined, the controller may reconstruct the data on the disk automatically. If a hot spare reconstruct was not initiated by the controller, replace the Failed disk and then initiate a reconstruct.

To initiate a reconstruct, do the following:

1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Reconstruct a PCI-X SCSI Disk Array**.
3. Before you proceed, ensure that you have replaced the Failed disk.
4. Select the *pdisk* that you want to reconstruct.

Three-Disk Failures

If a third disk in a RAID level 6 disk array fails, the array will be marked as Failed and the data will not be accessible. After deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group, delete the disk array, replace the Failed disks, then re-create the disk array. You must then add the disk array back to a volume group, re-create the disk array logical volumes and file systems, and copy the data back to the restored disk array from your backup media.

RAID Level 10

Single-Disk Failure

Whenever a single disk in a RAID level 10 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* disk is also stored on the other member of its mirrored pair.

Whenever a disk is marked as *Failed*, replace it as soon as possible. If a hot spare disk was *Defined*, the controller may reconstruct the data on the disk automatically. If a hot spare reconstruct was not initiated by the controller, then initiate a reconstruct on the *Failed* disk following replacement.

To initiate a reconstruct, do the following:

1. Navigate to the PCI-X SCSI Disk Array Manager by using the steps in "Using the PCI-X SCSI Disk Array Manager" on page 15.
2. Select **Reconstruct a PCI-X SCSI Disk Array**.
3. Before you proceed, ensure that you have replaced the *Failed* disk.
4. Select the *Failed pdisk* that you want to reconstruct.

Multiple-Disk Failures

If each *Failed* disk belongs to a different mirrored pair, then the array will be in the *Degraded* state because the data on each *Failed* disk is still available from its mirrored pair. The recovery procedure for this case is the same as the single-disk failure case. After you replace the *Failed* disks, initiate reconstruct operations.

When both members of a mirrored pair fail, the disk array status changes to *Failed*. After deleting the logical volumes and file systems on the disk array and removing the disk array from the volume group, delete the disk array, replace the *Failed* disks, then re-create the disk array. You must then add the disk array back to a volume group, re-create the disk array logical volumes and file systems, and copy the data back to the restored disk array from your backup media.

Chapter 6. Problem Determination and Recovery

Note: The procedures contained in this chapter are intended for service representatives specifically trained on the system unit and subsystem that is being serviced. Additionally, some of the service actions in this chapter might require involvement of the system administrator.

If a problem arises related to disk arrays and associated *pdisks*, use the following to identify the problem:

- Information presented by the Error Log Analysis (ELA)
- Hardware error logs viewed using the Display Hardware Error Report diagnostic task
- Disk array *hdisk* and *pdisk* status, viewed using the PCI-X SCSI Disk Array Manager

Error Log Analysis (ELA) analyzes errors presented by the adapter and recommends actions that need to be performed to correct the errors. It is sometimes recommended that you perform a Maintenance Analysis Procedure (MAP) to further determine what actions should be taken to resolve the problem. This chapter provides many of these MAPs.

The MAPs contained in this chapter are intended to address only problems directly related to disk arrays and SCSI bus-problem isolation. MAPs related to other device or adapter problems, when applicable, are located in other system documentation.

Considerations

Read the following before using these problem determination and recovery procedures:

- If a disk array is being used as a boot device and the system fails to boot because of a suspected disk-array problem, boot using the Standalone Diagnostic media. Error Log Analysis, AIX error logs, the PCI-X SCSI Disk Array Manager, and other tools are available on the Standalone Diagnostics to help determine and resolve the problem with the disk array.
- When invoking diagnostic routines for a PCI-X SCSI RAID Controller, it is recommended that Problem Determination (PD) mode be used instead of System Verification (SV) mode unless there is a specific reason to use SV mode (for example, you were directed to run SV mode by a MAP).
- After diagnostic routines for a PCI-X SCSI RAID Controller are run in SV mode, it is recommended to run diagnostics in PD mode to ensure that new errors are analyzed. These actions should be performed especially when using Standalone Diagnostic media.

Location Codes

Much of the location information contained in the hardware error logs is in the format of AIX location codes.

The AIX location code format for a SCSI device is: *AB-CD-EF-G,H* where

- *AB-CD* identifies the PCI-X SCSI RAID Controller.
- *EF* identifies the Controller's SCSI bus that the device is attached to.

   Note: For a disk array this value will be *ff*, which indicates a logical bus for RAID devices.
- *G,H* are the SCSI ID and logical unit number (LUN) of the device, respectively.

Often only a portion of the AIX location code is shown. For example:

- *AB-CD* identifies a PCI-X SCSI RAID Controller.
- *AB-CD-EF* identifies a PCI-X SCSI RAID Controller's SCSI bus.
- *EF-G,H* identifies the SCSI bus, SCSI ID, and LUN of an attached disk.

Identifying the Disk Array Problem

A disk array problem is uniquely identified by a Service Request Number (SRN). An SRN is in the format `nnnn - rrrr`, where the first four digits of the SRN preceding the dash (-) is known as the Failing Function Code (FFC, for example 2523) and the last four digits of the SRN following the dash (-) is known as the reason code. The reason code indicates the specific problem that has occurred and must be obtained in order to determine which MAP to use.

An SRN is provided by Error Log Analysis, which directs you to the MAPs contained in this chapter. To obtain the reason code (last four digits of the SRN) from an AIX error log, see "Finding an SRN Given an AIX Error Log" on page 97.

The SRN describes the problem that has been detected and should be considered the primary means of identifying a problem. However, the **List PCI-X SCSI Disk Array Configuration** option within the PCI-X SCSI Disk Array Manager is also useful in identifying a problem or confirming a problem described by Error Log Analysis. For additional information about the PCI-X Disk Array Manager, see "Using the PCI-X SCSI Disk Array Manager" on page 15.

Obtain the SRN and proceed to the next section to obtain a more detailed description of the problem and to determine which MAP to use.

Service Request Number (SRN) Table

Using the SRN obtained from Error Log Analysis or from the AIX error log, use the following table to determine which MAP to use. For MAPs, see "Maintenance Analysis Procedures (MAPs)" on page 64.

Note: The following table includes only SRNs that are associated with Maintenance Analysis Procedures contained in this document. A complete list of SRNs can be found in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Table 3.

SRN	Description	MAP
nnnn - 8008	Permanent Cache Battery Pack failure	MAP 3000
nnnn - 8009	Impending Cache Battery Pack failure	MAP 3000
nnnn - 9000 nnnn - 9001 nnnn - 9002	Controller detected device error during configuration discovery	MAP 3090
nnnn - 9008	Controller does not support function expected for one or more disks	MAP 3030
nnnn - 9010	Cache data associated with attached disks cannot be found	MAP 3020
nnnn - 9011	Cache data belongs to disks other than those attached	MAP 3090
nnnn - 9020 nnnn - 9021 nnnn - 9022	Two or more disks are missing from a RAID-5 or RAID-6 Disk Array	MAP 3011
nnnn - 9023	One or more Disk Array members are not at required physical locations	MAP 3012
nnnn - 9024	Physical location of Disk Array members conflict with another Disk Array	MAP 3090
nnnn - 9025	Incompatible disk installed at degraded disk location in Disk Array	MAP 3010
nnnn - 9026	Previously degraded disk in Disk Array not found at required physical location	MAP 3090

Table 3. (continued)

SRN	Description	MAP
nnnn - 9027	Disk Array is or would become degraded and parity data is out of synchronization	MAP 3013
nnnn - 9028	Maximum number of functional Disk Arrays has been exceeded	MAP 3090
nnnn - 9029	Maximum number of functional Disk Array disks has been exceeded	MAP 3090
nnnn - 9030	Disk Array is degraded due to missing or failed disk	MAP 3010
nnnn - 9031	Automatic reconstruction initiated for Disk Array	MAP 3010
nnnn - 9032	Disk Array is degraded due to a missing or failed disk	MAP 3010
nnnn - 9041	Background Disk Array parity checking detected and corrected errors	MAP 3090
nnnn - 9042	Background Disk Array parity checking detected and corrected errors on specified disk	MAP 3090
nnnn - 9050	Required cache data can not be located for one or more disks	MAP 3031
nnnn - 9051	Cache data exists for one or more missing or failed disks	MAP 3032
nnnn - 9052	Cache data exists for one or more modified disks	MAP 3090
nnnn - 9054	RAID controller resources not available due to previous problems	MAP 3021
nnnn - 9060	One or more disk pairs are missing from a RAID-10 Disk Array	MAP 3011
nnnn - 9061 nnnn - 9062	One or more disks are missing from a RAID-0 Disk Array	MAP 3011
nnnn - 9063	Maximum number of functional Disk Arrays has been exceeded	MAP 3090
nnnn - 9073	Multiple controllers connected in an invalid configuration	MAP 3040
nnnn - 9074	Multiple controllers not capable of similar functions or controlling the same set of devices	MAP 3041
nnnn - 9081 nnnn - 9082	Controller detected device error during internal media recovery	MAP 3090
nnnn - 9090	Disk has been modified after last known status	MAP 3033
nnnn - 9091	Incorrect disk configuration change has been detected	MAP 3033
nnnn - 9092	Disk requires Format before use	MAP 3034
nnnn - 9200	Identify disk to be replaced	MAP 3036
nnnn - FFF3	Disk media format bad	MAP 3035
nnnn - 719 nnnn - 720 nnnn - 721 nnnn - 3020 nnnn - 3100 nnnn - 3101 nnnn - 3109 nnnn - 3110 nnnn - FFFE	Various errors requiring SCSI bus problem isolation	MAP 3050

Table 3. (continued)

SRN	Description	MAP
nnnn - 102E nnnn - 3002 nnnn - 3010 nnnn - 7001 nnnn - FFF4 nnnn - FFF6 nnnn - FFFA	Various errors which call for disk replacement, however, replacing disk did not solve problem	MAP 3051

Maintenance Analysis Procedures (MAPs)

Use the following procedures to resolve adapter, cache, or disk array problems associated with a PCI-X SCSI RAID Controller. See "Service Request Number (SRN) Table" on page 62 to identify which MAP to use.

MAP 3000

Use this MAP to resolve the following problems:
- Permanent Cache Battery Pack failure (SRN nnnn - 8008)
- Impending Cache Battery Pack failure (SRN nnnn - 8009)

Step 3000-1

Prior to replacing the Cache Battery Pack, it must be forced into an error state. This will ensure that write caching is stopped prior to replacing the battery pack thus preventing possible data loss.

1. Follow the steps described in "Forcing a Rechargeable Battery Error" on page 34.

2. Go to Step 3000-2

Step 3000-2

Follow the steps described in "Replacing the Rechargeable Cache Battery Pack" on page 35.

When the problem is resolved, then go to MAP 0410: Repair Checkout, in RS/6000 pSeries Diagnostic Information for Multiple Bus Systems.

MAP 3010

Use this MAP to resolve the following problems:
- Incompatible disk installed at the degraded disk location in Disk Array (SRN nnnn - 9025)
- Disk Array is degraded due to a missing or failed disk (SRN nnnn - 9030)
- Automatic reconstruction initiated for a Disk Array (SRN nnnn - 9031)
- Disk Array is degraded due to a missing or failed disk (SRN nnnn - 9032)

Step 3010-1

Identify the disk array by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.

2. Select **Display Hardware Error Report**.

3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.

4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.

5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem that sent you here and select it.

> **Note:** If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. This error log displays the following disk array information under the *Array Information* heading: *Location*, *S/N*(serial number), and *RAID Level*.
7. Go to Step 3010-2.

Step 3010-2

View the current disk array configuration as follows:

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **List PCI-X SCSI Disk Array Configuration**.
3. Select the PCI-X SCSI RAID Controller identified in the hardware error log.
4. Go to Step 3010-3.

Step 3010-3

Does a disk array have a state of *Degraded*?

NO   Go to Step 3010-4.

YES   Go to Step 3010-5.

Step 3010-4

The affected disk array should have a state of either *Rebuilding* or *Optimal* due to the use of a hot spare disk.

Identify the *Failed* disk, which is no longer a part of the disk array, by finding the *pdisk* listed at the bottom of the screen that has a state of either *Failed* or *RWProtected*. Using appropriate service procedures, such as use of the SCSI and SCSI RAID Hot Plug Manager, remove the failed disk and replace it with a new disk to use as a hot spare:

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options** .
3. Select **SCSI and SCSI RAID Hot Plug Manager**.
4. Select **Identify a Device Attached to an SCSI Hot Swap Enclosure Device**.
5. Choose the location for the device you want to remove or install.

   > **Note:** The visual indicator on the Device will blink at the Identify rate.

6. If you are removing a device:
   a. Select **Replace/Remove a Device Attached to an SCSI Hot Swap Enclosure Device**.

      > **Note:** The visual indicator on the device will blink at the Remove rate.

   b. Remove the device.

   else, if installing a device:

   a. Select **Attach a Device to an SCSI Hot Swap Enclosure Device**.

> **Note:** The visual indicator on the device will blink at the Remove rate.

    b. Insert the device.

If a new disk is not listed as a *pdisk*, it may first need to be prepared for use in a disk array. Do the following:

1. Start the PCI-X SCSI Disk Array Manager

    a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

    b. Select **RAID Array Manager**.

    c. Select **PCI-X SCSI Disk Array Manager**.

2. Select **Create an Array Candidate pdisk and Format to 522 Byte Sectors**.

3. Select the PCI-X SCSI RAID Controller.

4. Select the disk(s) from the list that you want to prepare for use in the disk arrays.

In order to make the new disk usable as a hot spare, do the following:

1. Start the PCI-X SCSI Disk Array Manager.

    a. Start Diagnostics and select **Task Selection** on the Function Selection screen .

    b. Select **RAID Array Manager**.

    c. Select **PCI-X SCSI Disk Array Manager**.

2. Select **Change/Show PCI-X SCSI pdisk Status**.

3. Select **Create a Hot Spare**.

4. Select the PCI-X SCSI RAID Controller.

5. Select the *pdisk* that you want to designate as a hot spare.

> **Note:** Hot spare disks are useful only if their capacity is greater than or equal to that of the smallest capacity disk in a disk array that becomes *Degraded*.

When the problem is resolved, go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3010-5

Identify the failed disk by finding the *pdisk* listed for the degraded disk array that has a state of *Failed*. Using appropriate service procedures, such as use of the SCSI and SCSI RAID Hot Plug Manager, remove the failed disk and replace it with a new disk to use in the disk array. The SCSI and SCSI RAID Hot Plug Manager can be invoked as follows:

1. Start the PCI-X SCSI Disk Array Manager.

    a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

    b. Select **RAID Array Manager**.

    c. Select **PCI-X SCSI Disk Array Manager**.

2. Select **Diagnostics and Recovery Options**.

3. Select **SCSI and SCSI RAID Hot Plug Manager**.

4. Select **Identify a Device Attached to an SCSI Hot Swap Enclosure Device**.

5. Choose the location for the device you wish to remove/install.

> **Note:** The visual indicator on the device will blink at the Identify rate.

6. If removing a device:

    a. Select **Replace/Remove a Device Attached to an SCSI Hot Swap Enclosure Device**.

>   **Note:** The visual indicator on the device will blink at the Remove rate.

    b. Remove the device.

    else, if installing a device:

a. Select **Attach a Device to an SCSI Hot Swap Enclosure Device**.

> **Note:** The visual indicator on the device will blink at the Remove rate.

b. Insert the device.

To bring the disk array back to a state of *Optimal*, do the following:

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Reconstruct a PCI-X SCSI Disk Array**.
3. Select the failed *pdisk* to reconstruct.

> **Note:** The replacement disk should have a capacity that is greater than or equal to that of the smallest capacity disk in the degraded disk array.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3011

Use this MAP to resolve the following problems:
- Two or more disks are missing from a RAID-5 or RAID-6 Disk Array (SRN nnnn - 9020 / nnnn - 9021 / nnnn - 9022)
- One or more disk pairs are missing from a RAID-10 Disk Array (SRN nnnn - 9060)
- One or more disks are missing from a RAID-0 Disk Array (SRN nnnn - 9061 / nnnn - 9062)

Step 3011-1

Identify the disks missing from the disk array by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   > **Note:** If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, the missing disks are those listed under *Array Member Information* with an *Actual Location* of *unkwn*.
7. Go to Step 3011-2.

Step 3011-2

There are three possible ways to correct the problem. Perform only *one* of the following three options, listed in the order of preference:

- Locate the identified disks and install them in the correct physical locations (that is the *Expected Locations*) in the system. Perform only *one* of the following two options:
  - IPL the system or logical partition
  - Unconfigure and reconfigure the adapter by performing the following:
    1. Unconfigure the adapter.

a. Start the PCI-X SCSI Disk Array Manager.

      1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

      2) Select **RAID Array Manager**.

      3) Select **PCI-X SCSI Disk Array Manager**.

   b. Select **Diagnostics and Recovery Options**.

   c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.

2. Configure the adapter.

   a. Start the PCI-X SCSI Disk Array Manager

      1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

      2) Select **RAID Array Manager**.

      3) Select **PCI-X SCSI Disk Array Manager**.

   b. Select **Diagnostics and Recovery Options**.

   c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @*server *pSeries Diagnostic Information for Multiple Bus Systems*.

- Delete the disk array, as follows:

   Attention: All data on the disk array will be lost.

   1. Start the PCI-X SCSI Disk Array Manager.

      a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

      b. Select **RAID Array Manager**.

      c. Select **PCI-X SCSI Disk Array Manager**.

   2. Select **Delete a PCI-X SCSI Disk Array**.

   3. Select the PCI-X SCSI RAID Controller.

   4. Select the disk array to delete.

   When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @*server *pSeries Diagnostic Information for Multiple Bus Systems*.

- Format the remaining members of the disk array, as follows

   Attention: All data on the disk array will be lost.

   1. Start the PCI-X SCSI Disk Array Manager.

      a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

      b. Select **RAID Array Manager**.

      c. Select **PCI-X SCSI Disk Array Manager**.

   2. Select **Diagnostics and Recovery Options**.

   3. Select **Format Physical Disk Media (pdisk)**.

   When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @*server *pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3012

Use this MAP to resolve the following problem:

One or more Disk Array Members are not at required physical locations (SRN nnnn - 9023)

Step 3012-1

Identify the disks which are not at their required physical locations by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view.

   Viewing the hardware error log, the disks which are not at their required locations are those listed under *Array Member Information* with an *Expected Location* and *Actual Location* which do not match.

   Note: An *Actual Location* of *unkwn* is acceptable, and no action is needed to correct it. This *unkwn* location should only occur for the disk array member that corresponds to the *Degraded Disk S/N*.

7. Go to Step 3012-2.

Step 3012-2

There are three possible ways to correct the problem. Perform only *one* of the following three options, listed in the order of preference:

- Locate the identified disks and install them in the correct physical locations (that is the *Expected Locations*) in the system. Perform only *one* of the following two options:
  - IPL the system or logical partition
  - Unconfigure and reconfigure the adapter by performing the following:
    1. Unconfigure the adapter.
       a. Start the PCI-X SCSI Disk Array Manager.
          1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
          2) Select **RAID Array Manager**.
          3) Select **PCI-X SCSI Disk Array Manager**.
       b. Select **Diagnostics and Recovery Options**.
       c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.
    2. Configure the adapter.
       a. Start the PCI-X SCSI Disk Array Manager
          1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
          2) Select **RAID Array Manager**.
          3) Select **PCI-X SCSI Disk Array Manager**.
       b. Select **Diagnostics and Recovery Options**.
       c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Delete the disk array, as follows:

  Attention: All data on the disk array will be lost.

  1. Start the PCI-X SCSI Disk Array Manager.
     a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
     b. Select **RAID Array Manager**.
     c. Select **PCI-X SCSI Disk Array Manager**.
  2. Select **Delete a PCI-X SCSI Disk Array**.
  3. Select the PCI-X SCSI RAID Controller.
  4. Select the disk array to delete.

  When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Format the remaining members of the disk array, as follows

  Attention: All data on the disk array will be lost.

  1. Start the PCI-X SCSI Disk Array Manager.
     a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
     b. Select **RAID Array Manager**.
     c. Select **PCI-X SCSI Disk Array Manager**.
  2. Select **Diagnostics and Recovery Options**.
  3. Select **Format Physical Disk Media (pdisk)**.

  When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3013

Use this MAP to resolve the following problem:

Disk array is or would become degraded and parity data is out of synchronization (SRN nnnn - 9027)

Step 3013-1

Identify the adapter and disks by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, if the disk array member which corresponds to the *Degraded Disk S/N* has an *Actual Location* of *unkwn* and is not physically present, it may be helpful to find this disk.
7. Go to Step 3013-2.

Step 3013-2

Have the adapter card or disks been physically moved recently?

NO   Contact your service support organization.

YES   Go to Step 3013-3.

Step 3013-3

There are three possible ways to correct the problem. Perform only *one* of the following three options, listed in the order of preference:

- Restore the adapter and disks back to their original configuration. Perform only*one* of the following two options:
  - IPL the system or logical partition
  - Unconfigure and reconfigure the adapter by performing the following:
    1. Unconfigure the adapter.
        a. Start the PCI-X SCSI Disk Array Manager.
            1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
            2) Select **RAID Array Manager**.
            3) Select **PCI-X SCSI Disk Array Manager**.
        b. Select **Diagnostics and Recovery Options**.
        c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.
    2. Configure the adapter.
        a. Start the PCI-X SCSI Disk Array Manager
            1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
            2) Select **RAID Array Manager**.
            3) Select **PCI-X SCSI Disk Array Manager**.
        b. Select **Diagnostics and Recovery Options**.
        c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

  When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Delete the disk array, as follows:

  Attention: All data on the disk array will be lost.
    1. Start the PCI-X SCSI Disk Array Manager.
        a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
        b. Select **RAID Array Manager**.
        c. Select **PCI-X SCSI Disk Array Manager**.
    2. Select **Delete a PCI-X SCSI Disk Array**.
    3. Select the PCI-X SCSI RAID Controller.
    4. Select the disk array to delete.

  When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Format the remaining members of the disk array, as follows

  Attention: All data on the disk array will be lost.
    1. Start the PCI-X SCSI Disk Array Manager.
        a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
        b. Select **RAID Array Manager**.
        c. Select **PCI-X SCSI Disk Array Manager**.

2. Select **Diagnostics and Recovery Options**.
3. Select **Format Physical Disk Media (pdisk)**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3020

Use this MAP to resolve the following problem:

Cache data associated with attached disks cannot be found (SRN nnnn - 9010)

Step 3020-1
Has the server been powered off for several days?

NO   Go to Step 3020-2.

YES   Go to Step 3020-7.

Step 3020-2
Are you working with a 571F/575B card set?

NO   Go to Step 3020-4.

YES   Go to Step 3020-3.

Step 3020-3
Using the appropriate service procedures, remove the 571F/575B card set. Create and install new card set with the following parts installed on it:

Note: Label all parts (original and new) before moving them from one place to another.
* The new replacement 571F storage I/O adapter.
* The cache directory card from the original 571F storage I/O adapter.
* The original 575B Auxiliary Cache adapter.

Notes:
1. Refer to "Separating the 571F/575B Card Set and Moving the Cache Directory Card" on page 52.
2. Ensure that the SCSI cable and the battery power cable on the top edge of the storage I/O side of the card set are connected to the top edge of the Auxiliary Cache side of the card set.

Go to Step 3020-5..

Step 3020-4

Note: Label all parts (*original* and *new*) before moving them from one place to another.

Using the appropriate service procedures, remove the I/O adapter. Install the *new* replacement storage I/O adapter with the following parts installed on it:
* The cache directory card from the *original* storage I/O adapter. (Refer to "Replacing the Cache Directory Card" on page 49.
* The removable cache card from the *original* storage I/O adapter (This only applies to 571E and some 2780 I/O adapters. Refer to "Separating a Removable Cache Card From the Base Card on Type 2780 and 571E Cards" on page 38.
* If the storage I/O adapter is attached to a 574F Auxiliary Cache adapter, ensure that the SCSI cable on the last port of the storage I/O adapter is connected to the Auxiliary Cache adapter.

Step 3020-5

Has a new SRN of *nnnn*-9010 or *nnnn*-9050 occurred?

NO   Go to Step 3020-8.

YES   Go to Step 3020-6.

Step 3020-6

Was the new SRN *nnnn*-9050?

NO   The new SRN was *nnnn*-9010.

Reclaim the controller cache storage as follows:

Attention: Data might be lost. When an Auxiliary Cache adapter connected to the RAID Controller logs a nnnn - 9055 SRN in the hardware error log, the Reclaim does not result in lost sectors. Otherwise, the Reclaim does result in lost sectors.

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **Reclaim Controller Cache Storage**.
4. Select the PCI-X SCSI RAID Controller.
5. Confirm that you wish to proceed.
   Note: On the Reclaim Controller Cache Storage results screen, the number of lost sectors is displayed. If the number is 0, there is no data loss. If the number is not 0, data has been lost and the system operator may want to restore data after this procedure is completed.
6. Go to Step 3020-8.

YES   Contact your Service Support organization

Step 3020-7

If the server has been powered off for several days after an abnormal power-down, the cache battery pack may be depleted. Do not replace the adapter or the cache battery pack. Reclaim the controller cache storage as follows:

Attention: Data might be lost. When an Auxiliary Cache adapter connected to the RAID Controller logs a nnnn - 9055 SRN in the hardware error log, the Reclaim does not result in lost sectors. Otherwise, the Reclaim does result in lost sectors.

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **Reclaim Controller Cache Storage**.
4. Select the PCI-X SCSI RAID Controller.
5. Confirm that you wish to proceed.

   Note: On the Reclaim Controller Cache Storage results screen, the number of lost sectors is displayed. If the number is 0, there is no data loss. If the number is not 0, data has been lost and the system operator may want to restore data after this procedure is completed.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*

Step 3020-8
Are you working with a 571F/575B card set?

NO　　　Go to Step 3020-10.

YES　　　Go to Step 3020-9.

Step 3020-9
Using the appropriate service procedures, remove the 571F/575B card set. Create and install new card set with the following parts installed on it:

Note: Label all parts (original and new) before moving them around.

- The new 571F storage I/O adapter.
- The cache directory card from the new 571F storage I/O adapter.
- The new 575B Auxiliary Cache adapter.

Notes:

1. Refer to "Separating the 571F/575B Card Set and Moving the Cache Directory Card" on page 52.
2. Ensure that the SCSI cable and the battery power cable on the top edge of the storage I/O side of the card set are connected to the top edge of the Auxiliary Cache side of the card set.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*

Step 3020-10
Using the appropriate service procedures, remove the I/O adapter. Install the *new* replacement storage I/O adapter with the following parts installed on it:

- The cache directory card from the *new* storage I/O adapter. (Refer to "Replacing the Cache Directory Card" on page 49.
- The removable cache card from the *new* storage I/O adapter (This only applies to 571E and some 2780 I/O adapters. Refer to "Separating a Removable Cache Card From the Base Card on Type 2780 and 571E Cards" on page 38.
- If the storage I/O adapter is attached to a 574F Auxiliary Cache adapter, ensure that the SCSI cable on the last port of the storage I/O adapter is connected to the Auxiliary Cache adapter.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*

MAP 3021

Use this MAP to resolve the following problem:

RAID controller resources not available due to previous problems (SRN nnnn - 9054)

Step 3021-1
Perform the following:

1. Remove any new or replacement disks which have been attached to the adapter.
2. Take action on the other errors which have occurred at the same time as this error.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3030

Use this MAP to resolve the following problem:

Controller does not support function expected by one or more disks (SRN nnnn - 9008)

Step 3030-1

Identify the affected disks by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, the *Device Errors Detected* field indicates the total number of disks which are affected. The *Device Errors Logged* field indicates the number of disks for which detailed information is provided. Under the *Device* heading, the *Location*, *Vendor/Product ID*, and *S/N* are provided for up to three disks. Additionally, the *Controller Type* and *S/N* for each of these disks indicates the adapter to which the disk was last attached when it was operational.
7. Go to Step 3030-2.

Step 3030-2

Have the adapter card or disks been physically moved recently?

NO   Contact your service support organization.

YES   Go to Step 3030-3.

Step 3030-3

There are two possible ways to correct the problem. Perform only *one* of the following two options, listed in the order of preference:

- Restore the adapter and disks back to their original configuration. Perform only*one* of the following two options:
  - IPL the system or logical partition
  - Unconfigure and reconfigure the adapter by performing the following:
    1. Unconfigure the adapter.
       a. Start the PCI-X SCSI Disk Array Manager.
          1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
          2) Select **RAID Array Manager**.
          3) Select **PCI-X SCSI Disk Array Manager**.
       b. Select **Diagnostics and Recovery Options**.
       c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.
    2. Configure the adapter.
       a. Start the PCI-X SCSI Disk Array Manager
          1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

2) Select **RAID Array Manager**.

3) Select **PCI-X SCSI Disk Array Manager**.

   b. Select **Diagnostics and Recovery Options**.

   c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Format the disks, as follows:

   Attention: All data on the disk array will be lost.

   1. Start the PCI-X SCSI Disk Array Manager.

      a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

      b. Select **RAID Array Manager**.

      c. Select **PCI-X SCSI Disk Array Manager**.

   2. Select **Diagnostics and Recovery Options**.

   3. Select **Format Physical Disk Media (pdisk)**.

   When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3031

Use this MAP to resolve the following problem:

   Required cache data cannot be located for one or more disks (SRN nnnn - 9050)

Step 3031-1

Did you just exchange the adapter as the result of a failure?

NO   Go to Step 3031-5.

YES   Go to Step 3031-2.

Step 3031-2

Are you working with a 571F/575B card set?

NO   Go to Step 3031-4.

YES   Go to Step 3031-3.

Step 3031-3

Using the appropriate service procedures, remove the 571F/575B card set. Create and install new card set with the following parts installed on it:

Note: Label all parts (original and new) before moving them from one place to another.

- The new replacement 571F storage I/O adapter.
- The cache directory card from the original 571F storage I/O adapter.
- The original 575B Auxiliary Cache adapter.

Notes:

1. Refer to "Separating the 571F/575B Card Set and Moving the Cache Directory Card" on page 52.
2. Ensure that the SCSI cable and the battery power cable on the top edge of the storage I/O side of the card set are connected to the top edge of the Auxiliary Cache side of the card set.

Go to Step 3031-10.

Step 3031-4

Note: The failed adapter that you have just exchanged contains cache data that is required by the disks that were attached to that adapter. If the adapter that you just exchanged is failing intermittently, reinstalling it and IPLing the system may allow the data to be successfully written to the disks. After the cache data is written to the disks and the system is powered off normally, the adapter can be replaced without data being lost. Otherwise, continue with this procedure.

Note: Label all parts (old and new) before moving them around.

Using the appropriate service procedures, remove the I/O adapter. Install the new replacement storage I/O adapter with the following parts installed on it:

- The cache directory card from the original storage I/O adapter. (Refer to "Replacing the Cache Directory Card" on page 49.
- The removable cache card from the original storage I/O adapter (This only applies to 571E and some 2780 I/O adapters. Refer to "Separating a Removable Cache Card From the Base Card on Type 2780 and 571E Cards" on page 38.
- If the storage I/O adapter is attached to a 574F Auxiliary Cache adapter, ensure that the SCSI cable on the last port of the storage I/O adapter is connected to the Auxiliary Cache adapter.

Go to Step 3031-10.

Step 3031-5

Identify the affected disks by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, the *Device Errors Detected* field indicates the total number of disks which are affected. The *Device Errors Logged* field indicates the number of disks for which detailed information is provided. Under the *Device* heading, the *Location, Vendor/Product ID*, and *S/N* are provided for up to three disks. Additionally, the *Controller Type* and *S/N* for each of these disks indicates the adapter to which the disk was last attached when it was operational.
7. Go to Step 3031-6.

Step 3031-6

Have the adapter card or disks been physically moved recently?

NO   Contact your service support organization.

YES   Go to Step 3031-7

Step 3031-7

Is the data on the disks needed for this or any other system?

NO   Go to Step 3031-9.

YES   Go to Step 3031-8.

Step 3031-8

The adapter and disks, identified above, must be reunited so that the cache data can be written to the disks.

Restore the adapter and disks back to their original configuration.

Once the cache data is written to the disks and the system is powered off normally, the adapter and/or disks may be moved to another location.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3031-9

There are three possible ways to correct the problem. Perform only *one* of the following three options, listed in the order of preference:

* Reclaim Controller Cache Storage by performing the following:

   Attention: All data on the disk array will be lost.

   1.   Start the PCI-X SCSI Disk Array Manager.

      a.   Start Diagnostics and select **Task Selection** on the Function Selection screen.

      b.   Select **RAID Array Manager**.

      c.   Select **PCI-X SCSI Disk Array Manager**.

   2.   Select **Diagnostics and Recovery Options**.

   3.   Select **Reclaim Controller Cache Storage**.

   4.   Select the PCI-X SCSI RAID Controller.

   5.   Confirm that you will Allow Unknown Data Loss.

   6.   Confirm that you wish to proceed.

   When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

* If the disks are members of a disk array, delete the disk array.

   Attention: All data on the disk array will be lost.
   This may be performed as follows:

   1.   Start the PCI-X SCSI Disk Array Manager.

      a.   Start Diagnostics and select **Task Selection** on the Function Selection screen.

      b.   Select **RAID Array Manager**.

      c.   Select **PCI-X SCSI Disk Array Manager**.

   2.   Select **Delete a PCI-X SCSI Disk Array**.

   3.   Select the PCI-X SCSI RAID Controller.

   4.   Select the disk array to delete.

   When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

* Format the disks, as follows:

   Attention: All data on the disks will be lost.

   1.   Start the PCI-X SCSI Disk Array Manager.

      a.   Start Diagnostics and select **Task Selection** on the Function Selection screen.

      b.   Select **RAID Array Manager**.

      c.   Select **PCI-X SCSI Disk Array Manager**.

2. Select **Diagnostics and Recovery Options**.
3. Select **Format Physical Disk Media (pdisk)**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3031-10
Has a new SRN nnnn-9010 or nnnn-9050 occurred?

NO Go to Step 3031-12.

YES Go to Step 3031-11.

Step 3031-11
Was the new SRN *nnnn*-9050?

NO The new SRN was *nnnn*-9010.

Reclaim the controller cache storage as follows:

Attention: Data might be lost. When an Auxiliary Cache adapter connected to the RAID Controller logs a nnnn - 9055 SRN in the hardware error log, the Reclaim does not result in lost sectors. Otherwise, the Reclaim does result in lost sectors.
1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **Reclaim Controller Cache Storage**.
4. Select the PCI-X SCSI RAID Controller.
5. Confirm that you wish to proceed.
   Note: On the Reclaim controller cache storage results screen, the number of lost sectors is displayed. If the number is 0, there is no data loss. If the number is not 0, data has been lost and the system operator may want to restore data after this procedure is completed.
6. Go to Step 3031-12.

YES Contact your Service Support organization

Step 3031-12
Are you working with a 571F/575B card set?

NO Go to Step 3031-14.

YES Go to Step 3031-13.

Step 3031-13
Using the appropriate service procedures, remove the 571F/575B card set. Create and install new card set with the following parts installed on it:

Note: Label all parts (original and new) before moving them from one place to another.
- The new 571F storage I/O adapter.
- The cache directory card from the new 571F storage I/O adapter.
- The new 575B Auxiliary Cache adapter.

Notes:
1. Refer to "Separating the 571F/575B Card Set and Moving the Cache Directory Card" on page 52.

2. Ensure that the SCSI cable and the battery power cable on the top edge of the storage I/O side of the card set are connected to the top edge of the Auxiliary Cache side of the card set.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000* @server *pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3031-14
Using the appropriate service procedures, remove the I/O adapter. Install the new replacement storage I/O adapter with the following parts installed on it:

- The cache directory card from the new storage I/O adapter. (Refer to "Replacing the Cache Directory Card" on page 49.
- The removable cache card from the NEW storage I/O adapter (This only applies to 571E and some 2780 I/O adapters. Refer to "Separating a Removable Cache Card From the Base Card on Type 2780 and 571E Cards" on page 38.
- If the storage I/O adapter is attached to a 574F Auxiliary Cache adapter, ensure that the SCSI cable on the last port of the storage I/O adapter is connected to the Auxiliary Cache adapter.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000* @server *pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3032
Use this MAP to resolve the following problem:

Cache data exists for one or more missing or failed disks (SRN nnnn - 9051)

The possible causes are:

- One or more disks have failed on the adapter.
- One or more disks were either moved concurrently or were removed after an abnormal power off.
- The adapter was moved from a different system or a different location on this system after an abnormal power off.
- The cache of the adapter was not cleared before it was shipped to the customer.

Step 3032-1
Identify the affected disks by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, the *Device Errors Detected* field indicates the total number of disks which are affected. The *Device Errors Logged* field indicates the number of disks for which detailed information is provided. Under the *Device* heading, the *Location*, *Vendor/Product ID*, and *S/N* are provided for up to three disks. Additionally, the *Controller Type* and *S/N* for each of these disks indicates the adapter to which the disk was last attached when it was operational.
7. Go to Step 3030-2.

Step 3032-2

Are there other disk or adapter errors which have occurred at about the same time as this error?

NO   Go to Step 3032-3.

YES   Go to Step 3032-6.

Step 3032-3

Is the data on the disks (and thus the cache data for the disks) needed for this or any other system?

NO   Go to Step 3032-7.

YES   Go to Step 3032-4.

Step 3032-4

Have the adapter card or disks been physically moved recently?

NO   Contact your service support organization.

YES   Go to Step 3032-5.

Step 3032-5

The adapter and disks must be reunited so that the cache data can be written to the disks.

Restore the adapter and disks back to their original configuration.

After the cache data is written to the disks and the system is powered off normally, the adapter and/or disks may be moved to another location.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000* @server *pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3032-6

Take action on the other errors that have occurred at the same time as this error.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000* @server *pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3032-7

Attention: Data will be lost.

Reclaim the Controller Cache Storage by performing the following:
1. Start the PCI-X SCSI Disk Array Manager.
    a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
    b. Select **RAID Array Manager**.
    c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **Reclaim Controller Cache Storage**.
4. Select the PCI-X SCSI RAID Controller.
5. Confirm that you will Allow Unknown Data Loss.
6. Confirm that you wish to proceed.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3033

Use this MAP to resolve the following problems:
* Disk has been modified after last known status (SRN nnnn - 9090)
* Incorrect disk configuration change has been detected (SRN nnnn - 9091)

Step 3033-1

Perform only *one* of the following two options:
* IPL the system or logical partition
* Unconfigure and reconfigure the adapter by performing the following:
    1. Unconfigure the adapter.
        a. Start the PCI-X SCSI Disk Array Manager.
            1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
            2) Select **RAID Array Manager**.
            3) Select **PCI-X SCSI Disk Array Manager**.
        b. Select **Diagnostics and Recovery Options**.
        c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.
    2. Configure the adapter.
        a. Start the PCI-X SCSI Disk Array Manager
            1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
            2) Select **RAID Array Manager**.
            3) Select **PCI-X SCSI Disk Array Manager**.
        b. Select **Diagnostics and Recovery Options**.
        c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

Take action on the other errors which have occurred at the same time as this error.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3034

Use this MAP to resolve the following problem:

Disk requires Format before use (SRN nnnn - 9092)

The possible causes are:
* Disk is a previously failed disk from a disk array and was automatically replaced by a hot spare disk.
* Disk is a previously failed disk from a disk array and was removed and later reinstalled on a different adapter or different location on this adapter.
* Appropriate service procedures were not followed when replacing disks or reconfiguring the adapter, such as not using the SCSI and SCSI RAID Hot Plug Manager when concurrently removing/installing disks or not performing a normal power down of the system prior to reconfiguring disks and adapters.
* Disk is member of a disk array, but was detected subsequent to the adapter being configured.
* Disk has multiple or complex configuration problems.

Step 3034-1

Identify the affected disks by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.
6. Select the hardware error log to view. Viewing the hardware error log, the *Device Errors Detected* field indicates the total number of disks which are affected. The *Device Errors Logged* field indicates the number of disks for which detailed information is provided. Under the *Device* heading, the *Location*, *Vendor/Product ID*, and *S/N* are provided for up to three disks. Additionally, the *Controller Type* and *S/N* for each of these disks indicates the adapter to which the disk was last attached when it was operational.
7. Go to Step 3030-2.

Step 3034-2
Are there other disk or adapter errors which have occurred at about the same time as this error?

NO   Go to Step 3034-3.

YES   Go to Step 3034-5.

Step 3034-3
Have the adapter card or disks been physically moved recently?

NO   Go to Step 3034-4.

YES   Go to Step 3034-6.

Step 3034-4
Is the data on the disks not needed for this or any other system and you wish to continue to use them with this adapter?

NO   Go to Step 3034-6.

YES   Go to Step 3034-7.

Step 3034-5
Take action on the other errors which have occurred at the same time as this error.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3034-6
Perform only *one* of the following actions that is most applicable to your situation:
- Perform only *one* of the following two options:
  - IPL the system or logical partition
  - Unconfigure and reconfigure the adapter by performing the following:
    1. Unconfigure the adapter.
       a. Start the PCI-X SCSI Disk Array Manager.

1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

2) Select **RAID Array Manager**.

3) Select **PCI-X SCSI Disk Array Manager**.

    b. Select **Diagnostics and Recovery Options**.

    c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.

2. Configure the adapter.

    a. Start the PCI-X SCSI Disk Array Manager

1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

2) Select **RAID Array Manager**.

3) Select **PCI-X SCSI Disk Array Manager**.

    b. Select **Diagnostics and Recovery Options**.

    c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

Take action on the other errors which have occurred at the same time as this error.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Restore the adapter and disks to their original configuration. Once this has been done, perform only*one* of the following two options:

  – IPL the system or logical partition

  – Unconfigure and reconfigure the adapter by performing the following:

1. Unconfigure the adapter.

    a. Start the PCI-X SCSI Disk Array Manager.

1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

2) Select **RAID Array Manager**.

3) Select **PCI-X SCSI Disk Array Manager**.

    b. Select **Diagnostics and Recovery Options**.

    c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.

2. Configure the adapter.

    a. Start the PCI-X SCSI Disk Array Manager

1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

2) Select **RAID Array Manager**.

3) Select **PCI-X SCSI Disk Array Manager**.

    b. Select **Diagnostics and Recovery Options**.

    c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

- Remove the disks from this adapter.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3034-7

There are two possible ways to correct the problem. Perform only *one* of these options.

- Format the disks.

Attention: All data on the disks will be lost.

This may be performed as follows:

1. Start the PCI-X SCSI Disk Array Manager.

    a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **Format Physical Disk Media (pdisk)**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

* If the disks are members of a disk array, delete the disk array by doing the following.

   Attention: All data on the disk array will be lost.

   Note: In some rare scenarios, deleting the disk array will not have no effect on a disk and the disk must be formatted instead.

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Delete a PCI-X SCSI Disk Array**.
3. Select the PCI-X SCSI RAID Controller.
4. Select the disk array to delete.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3035

Use this MAP to resolve the following problem:

   Disk media format bad (SRN nnnn - FFF3)

The possible causes are:
* Disk was being formatted and was powered off during this process.
* Disk was being formatted and was reset during this process.

Step 3035-1

Identify the affected disk by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapter resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, under the *Disk Information* heading, the *Location*, *Vendor/Product ID*, and *S/N* are provided for the disk.
7. Go to Step 3035-2.

Step 3035-2

Format the disk by performing the following.

Attention: All data on the disks will be lost.

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **Format Physical Disk Media (pdisk)**.

When the problem is resolved then go to MAP 0410: Repair Checkout, in*RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3036

Use this MAP to resolve the following problem:

Identify disk to be replaced (SRN nnnn - 9200)

You are sent here when a *pdisk* (that is physical disk in 522 bytes/sector format) is to be replaced, however, the location of this disk was not provided.

Step 3036-1

Identify the failing disk by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, under the *Disk Information* heading, the *Location*, *Vendor/Product ID*, and *S/N* are provided for the disk.
7. Go to Step 3036-2.

Step 3036-2

Using appropriate service procedures, such as use of the SCSI and SCSI RAID Hot Plug Manager, remove the failed disk and replace it with a new disk. The SCSI and SCSI RAID Hot Plug Manager can be invoked as follows:

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **Diagnostics and Recovery Options**.
3. Select **SCSI and SCSI RAID Hot Plug Manager**.
4. Select **Identify a Device Attached to an SCSI Hot Swap Enclosure Device**.

5. Choose the location for the device you wish to remove or install

   Note: The visual indicator on the Device will blink at the Identify rate.
6. If removing a device:
   a. Select **Replace/Remove a Device Attached to an SCSI Hot Swap Enclosure Device**

      Note: The visual indicator on the Device will blink at the Remove rate.
   b. Remove the device

   else if installing a device:
   a. Select **Attach a Device to an SCSI Hot Swap Enclosure Device**

      Note: The visual indicator on the Device will blink at the Remove rate.
   b. Insert the device
7. Go to Step 3036-3.

Step 3036-3

1. Run diagnostics in system verification mode on the adapter
2. Take action on any other errors which may have surfaced due to removing the disk, if any, such as for degraded disk arrays

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3040

Use this MAP to resolve the following problem:

   Multiple controllers connected in an invalid configuration (SRN nnnn - 9073)

Step 3040-1

Was the error logged by a 574F Auxiliary Cache adapter, a storage I/O adapter which attaches an Auxiliary Cache adapter (for example, 2780 or 571E), or by a 571F/575B card set?

NO   Go to Step 3040-5.

YES   Go to Step 3040-2.

Step 3040-2

Are you working with a 571F/575B card set?

NO   Go to Step 3040-4.

YES   Go to Step 3040-3.

Step 3040-3

You must configure both the 571F and the 575B in the same partition.

Are both the 571F side of the card set and the 575B side of the card set configured in the same partition?

NO   Change the configuration so that the same parition controls both cards in the card set.

YES   Replace the entire card set.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3040-4

There are five primary reasons for receiving this error (SRN nnnn - 9073).

- The cable is not properly connected to the storage I/O adapter and the Auxiliary Cache adapter. Ensure that the SCSI cable on the last port of the storage I/O adapter is connected to the Auxiliary Cache adapter.
- Both adapters are not in the same system/partition and powered on. Ensure that both the storage I/O adapter and Auxiliary Cache adapter are assigned to the same system/partition and they are powered on.
- The adapter microcode levels are not up to date. Ensure that the microcode for both the storage I/O adapter and Auxiliary Cache adapter are at the latest level.
- An unsupported Auxiliary Cache adapter is connected to the adapter that is logging this error. Verify that the connected Auxiliary Cache adapter is listed in Table 1. If the Auxiliary Cache adapter is not listed in the table, and thus is not supported, it should be removed.

  Note: The adapter which is logging this error will run in a performance degraded mode, without caching, until the unsupported Auxiliary Cache adapter is removed.

- A storage I/O adapter which does not support Auxiliary Cache is attached to an Auxiliary Cache adapter. Verify that the storage I/O adapter is listed in Table 1 with Auxiliary Cache support marked as YES. If the storage I/O adapter does not support Auxiliary Cache then the Auxiliary Cache must be disconnected from the storage I/O adapter.

Determine which of these is the cause of your specific error and take the appropriate actions listed. If this does not correct the error, contact your next level of support.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3040-5

There are two primary reasons for receiving this error (SRN nnnn - 9073).

- An adapter in a multi-initiator and high-availability environment sees more than one other adapter connected in the configuration. Only two adapters are supported connected together in the multi-initiator and high-availability configuration. Work with the customer to identify and correct the invalid configuration.
- Incompatible adapters are connected in a multi-initiator and high-availability environment. One or both of the adapters do not support attachment in a multi-initiator configuration. Verify the adapters logging this error are listed in Table 1 with multi-initiator support mark as YES.

Determine which of these is the cause of your specific error and take the appropriate actions listed. If this does not correct the error, contact your next level of support.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3041

Use this MAP to resolve the following problem:

Multiple controllers not capable of similar functions or controlling the same set of devices (SRN nnnn - 9074).

Step 3041-1

To obtain the reason/description for this failure, you must find the formatted error information in the AIX error log. This should also contain information about the other connected adapter (Remote Adapter Fields).

Display the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select Task Selection on the Function Selection screen.
2. Select Display Hardware Error Report.
3. Select Display Hardware Errors for PCI-X SCSI RAID Adapters.
4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. When viewing the hardware error log, the Detail Data section contains the "Reason for Failure" and "Remote Adapter" information.

Step 3041-2

Find the Reason for Failure and information for the other attached adapter (Remote Adapter) shown in the error log, and perform the action listed for the Reason in the following table:

Table 4.

Reason for Failure	Description	Action	Adapter on which to perform the action
Secondary does not support RAID level being used by the Primary.	Logged by the secondary adapter. For example, the Primary has a Raid-6 array and the Secondary does not support Raid-6.	Customer needs to upgrade the type of secondary adapter or change the RAID level of the array on the Primary to a level that is supported by the secondary.	Physically change the type of adapter which logged the error. Change RAID level on Primary adapter (Remote adapter indicated in the error log).
Secondary does not support disk function being used by the Primary.	Secondary detected some device function/attribute that it does not support.	Customer may need to upgrade the Adapter micro code or upgrade the type of secondary adapter.	Adapter which logged the error.
Secondary is unable to find devices found by the Primary	Secondary logging the error, cannot see all the devices that the Primary sees.	Verify the connections to the devices from the adapter logging the error. View the Disk Array Configuration screens to determine the SCSI bus with the problem.	Adapter which logged the error.
Secondary found devices not found by the Primary	Secondary adapter sees more devices than the Primary. After this error is logged, automatic failover will occur.	Verify the connections to the devices from the Remote Adapter as indicated in the error log.	Remote adapter indicated in the error log.
Secondary bus not connected to the same numbered bus on Primary	SCSI connections from the adapter to the devices are incorrect. Both adapters must be connected to the same devices on the same adapter SCSI bus	Verify connections and re-cable SCSI connections as necessary.	Either adapter.

Table 4. (continued)

Reason for Failure	Description	Action	Adapter on which to perform the action
Primary lost contact with disks accessible by Secondary	Link failure from Primary adapter to devices. Failover will occur.	Verify cable connections from the adapter which logged the error. Possible disk enclosure failure.	Adapter which logged the error.
Other	Not currently defined	Contact Next Level Support	

When the problem is resolved then go to MAP 0410: Repair Checkout, in RS/6000 eServer pSeries Diagnostic Information for Multiple Bus Systems.

MAP 3050

Use the following to perform SCSI bus problem isolation.

Considerations:

- Remove power from the system before connecting and disconnecting cables or devices, as appropriate, to prevent hardware damage or erroneous diagnostic results.
- Some systems have SCSI and PCI-X bus interface logic integrated onto the system boards and use a pluggable RAID Enablement Card (a non-PCI form factor card) for these SCSI/PCI-X buses. An example of such a RAID Enablement Card is FC 5709. For these configurations, replacement of the RAID Enablement Card is unlikely to solve a SCSI bus-related problem because the SCSI bus interface logic is on the system board.
- Some adapters provide two connectors, one internal and one external, for each SCSI bus. For this type of adapter, it is not acceptable to use both connectors for the same SCSI bus at the same time. SCSI bus problems are likely to occur if this is done. However, it is acceptable to use an internal connector for one SCSI bus and an external connector for another SCSI bus. The internal and external connectors are labeled to indicate which SCSI bus they correspond to.
- When two adapters are connected in a Multi-Initiator and High-Availability configuration, as described in Chapter 4, "Multi-Initiator and High-Availability," on page 25, each adapter's SCSI ID must be set to a different value when connected to shared disk enclosures. If the SCSI IDs are not set properly, many SCSI bus problems can occur.

Attention: Replacing RAID adapters is not recommended without assistance from your service support organization when SCSI bus problems exist. Because the adapter may contain non-volatile write cache data and configuration data for the attached disk arrays, additional problems can be created by replacing a adapter when SCSI bus problems exist.

Attention: Removing functioning disks in a disk array is not recommended without assistance from your service support organization. A disk array may become degraded or failed if functioning disks are removed and additional problems may be created.

Common Device Removal and Installation Procedure

When this MAP calls for a device (that is disk, tape, CD-ROM, or DVD-ROM) to be removed or installed, use this common device removal and installation procedure:

If the boot device is not on this adapter *and* the device to be removed or installed is a hot-swap device, follow this procedure. If these conditions do not apply to your situation, see the *else* section located at the end of this procedure.

1. Remove or install the device by invoking the SCSI and SCSI RAID Hot Plug Manager.
   a. Start the PCI-X Disk Array Manager.
      1) Start Diagnostics and select **Task Selection** on the Function Selection screen.
      2) Select **RAID Array Manager**.

3) Select **PCI-X SCSI Disk Array Manager**.

b. Select **Diagnostics and Recovery Options**.

c. Select **SCSI and SCSI RAID Hot Plug Manager**.

d. Select **Identify a Device Attached to an SCSI Hot Swap Enclosure Device**.

e. Choose the location for the device you wish to remove or install.

   Note: The visual indicator on the device will blink at the Identify rate.

f. If removing a device:

   1) Select **Replace/Remove a Device Attached to an SCSI Hot Swap Enclosure Device**.

      Note: The visual indicator on the device will blink at the Remove rate.

   2) Remove the device.

   3) Label the device with the slot it was removed from to ensure it can be reinstalled in the same location.

   else if installing a device:

   1) Select **Attach a Device to an SCSI Hot Swap Enclosure Device**.

      Note: The visual indicator on the device will blink at the Remove rate.

   2) Insert the device.

2. Unconfigure the adapter

   a. Start the PCI-X SCSI Disk Array Manager.

      1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

      2) Select **RAID Array Manager**.

      3) Select **PCI-X SCSI Disk Array Manager**.

   b. Select **Diagnostics and Recovery Options**.

   c. Select **Unconfigure an Available PCI-X SCSI RAID Controller**.

3. Configure the adapter.

   a. Start the PCI-X SCSI Disk Array Manager.

      1) Start Diagnostics and select **Task Selection** on the Function Selection screen.

      2) Select **RAID Array Manager**.

      3) Select **PCI-X SCSI Disk Array Manager**.

   b. Select **Diagnostics and Recovery Options**.

   c. Select **Configure a Defined PCI-X SCSI RAID Controller**.

else:

1. Power off the system/logical partition

2. Remove/install the device(s)

3. Power on the system/logical partition

Step 3050-1

Identify the SCSI bus which the problem is occurring on by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.

2. Select **Display Hardware Error Report**.

3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.

4. Select the adapter resource, or select all adapter resources if the adapter resource is not known.

5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

> **Note:** If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, under the *Disk Information* heading, the *Location* field can be used to identify which SCSI bus the error is associated with (that is the SCSI bus is the value *EF* given a location of *EF-G,H*).

7. Go to Step 3050-2.

Step 3050-2
Have recent changes been made to the SCSI configuration?

NO Go to Step 3050-5.

YES Go to Step 3050-3.

Step 3050-3
Check for the following problems:
- Address conflicts between devices
- Cabling problems such as; configurations that exceed the maximum cable lengths, missing termination, or excessive termination
- Both internal and external connectors for this SCSI bus are being used at the same time (only one should have a cable attached)
- If in a Multi-Initiator High-Availability configuration, verify correct cabling and SCSI ID settings as described in Chapter 4, "Multi-Initiator and High-Availability," on page 25.

For more details about supported SCSI cabling, refer to *RS/6000 @server pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*.

Did you find a problem?

NO Go to Step 3050-5.

YES Go to Step 3050-4.

Step 3050-4
1. Power off the system or logical partition.
2. Correct the problem.
3. Power on the system or logical partition, and run diagnostics in system verification mode on the adapter.

Did a SCSI bus-related failure occur?

NO Go to Step 3050-16.

YES Go to Step 3050-5.

Step 3050-5
Is problem related to the thermal fuse (that is SRN nnnn-719)?

NO Go to Step 3050-7.

YES Go to Step 3050-6.

Step 3050-6

The thermal fuse protects the SCSI bus from high currents due to shorts on the terminator, cable, or device. It is unlikely that the thermal fuse can be tripped by a defective adapter. A fault (short-circuit) causes an increase in resistance and temperature of the thermal fuse. The increase in temperature causes the thermal fuse to halt current flow. The thermal fuse returns to a low resistive and low temperature state when the fault is removed from the SCSI bus or when the system is turned off. Wait 10 seconds for the thermal fuse to reset itself and recover, then retest.

If the same error persists, replace the components of the failing SCSI bus in the following order. Wait 10 seconds for the thermal fuse to reset itself between steps.

1. Cable (if present)
2. DASD backplane (if present)
3. System board (if SCSI bus interface logic is on the system board)

To replace a component and verify that the problem was corrected, do the following:

1. Power off the system or logical partition.
2. Replace a component listed above.
3. Power on the system or logical partition, and run diagnostics in system verification mode on the adapter.

Did you correct the problem?

NO   Go to Step 3050-7.

YES   Go to Step 3050-16.

Step 3050-7

Determine if any of the disk arrays on the adapter are in a *Degraded* state as follows:

1. Start the PCI-X SCSI Disk Array Manager.
   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.
   b. Select **RAID Array Manager**.
   c. Select **PCI-X SCSI Disk Array Manager**.
2. Select **List PCI-X SCSI Disk Array Configuration**.
3. Select the PCI-X SCSI RAID Controller identified in the hardware error log.

Does any disk array have a State of *Degraded*?

NO   Go to Step 3050-9.

YES   Go to Step 3050-8.

Step 3050-8

1. Identify the failed disk(s) by first finding disk arrays with a state of *Degraded* and then a *pdisk* for that disk array which has a state of *Failed*.
2. Remove the failed disk from each degraded disk array by using the "Common Device Removal and Installation Procedure" on page 90.
3. Run diagnostics in system verification mode on the adapter.

Did a SCSI bus related failure occur?

NO   Go to Step 3050-16.

YES   Go to Step 3050-9.

Step 3050-9

Are there any non-essential removable media devices (for example tape, CD-ROM, or DVD-ROM) on the SCSI bus?

NO   Go to Step 3050-12.

YES   Go to Step 3050-10.

Step 3050-10

1. Remove one of the non-essential removable media devices by using the "Common Device Removal and Installation Procedure" on page 90.
2. Run diagnostics in system verification mode on the adapter.

Did a SCSI bus related failure occur?

NO   Go to Step 3050-11.

YES   Go to Step 3050-9.

Step 3050-11

The last removable media device removed from the SCSI bus may be the cause of the SCSI bus problems. Follow the repair procedures for that device.

Go to Step 3050-16.

Step 3050-12

Are there any non-essential disks which are not disk array members (for example, 512 byte/sector standalone disks, hot spare disks, or Array Candidates) on the SCSI bus?

NO   Go to Step 3050-15.

YES   Go to Step 3050-13.

Step 3050-13

1. Remove one of the non-essential disks devices by using the "Common Device Removal and Installation Procedure" on page 90.
2. Run diagnostics in system verification mode on the adapter.

Did a SCSI bus related failure occur?

NO   Go to Step 3050-14.

YES   Go to Step 3050-12.

Step 3050-14

The last disk removed from the SCSI bus may be the cause of the SCSI bus problems. Follow the repair procedures for that device.

Go to Step 3050-16.

Step 3050-15

Contact your service support organization.

Exit this procedure

Step 3050-16

1. Reinstall any good devices that were removed by using the "Common Device Removal and Installation Procedure" on page 90.
2. Run diagnostics in system verification mode on the adapter.
3. Perform the appropriate problem determination steps to resolve any other non-SCSI bus related errors that appear, if any, such as for degraded disk arrays.

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3051

Use the following to determine what other FRUs, besides a disk, may need to be replaced in order to solve a problem.

You are sent here when a *pdisk* (that is physical disk in 522 bytes/sector format) was identified as the primary FRU to replace in order to solve a problem. However, if replacing the disk did not resolve the problem then other FRUs may need to be replaced.

Considerations:

- Remove power from the system before connecting and disconnecting cables or devices, as appropriate, to prevent hardware damage or erroneous diagnostic results.
- Keep in mind that some systems have SCSI and PCI-X bus interface logic integrated onto the system boards and use a pluggable RAID Enablement Card (a non-PCI form factor card) for these SCSI/PCI-X busses. An example of such a RAID Enablement Card is FC 5709. For these configurations, replacement of the RAID Enablement Card is unlikely to solve a SCSI bus related problem since the SCSI bus interface logic is on the system board.
- Some adapters provide two connectors, one internal and one external, for each SCSI bus. For this type of adapter, it is not acceptable to use both connectors for the same SCSI bus at the same time. SCSI bus problems are likely to occur if this is done. However, it is acceptable to use an internal connector for one SCSI bus and an external connector for another SCSI bus. The internal and external connectors are labeled to indicate which SCSI bus they correspond to.

Attention: Replacing RAID adapters is not recommended without assistance from your service support organization when SCSI bus problems exist. Because the adapter may contain non-volatile write cache data and configuration data for the attached disk arrays, additional problems can be created by replacing a adapter when SCSI bus problems exist.

Attention: Removing functioning disks in a disk array is not recommended without assistance from your service support organization. A disk array may become degraded or failed if functioning disks are removed and additional problems may be created.

Step 3051-1

Identify the SCSI bus which the problem is occurring on by examining the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.
2. Select **Display Hardware Error Report**.
3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.
4. Select the adapter resource, or select all adapter resources if the adapter resource is not known.
5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view. Viewing the hardware error log, under the *Disk Information* heading, the *Location* field can be used to identify which SCSI bus the error is associated with (that is the SCSI bus is the value *EF* given a location of *EF-G,H*).

7. Go to Step 3051-2.

Step 3051-2

While the error persists, replace the components of the failing SCSI bus in the following order.

1. Cable (if present)

2. Adapter (if SCSI bus interface logic is on the adapter) or system board (if SCSI bus interface logic is on the system board)

3. DASD backplane (if present)

To replace a component and see if the problem was corrected, do the following:

1. Power off the system/logical partition

2. Replace a component listed above

3. Power on the system/logical partition, and run diagnostics in system verification mode on the adapter

When the problem is resolved then go to MAP 0410: Repair Checkout, in *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3090

The problem that occurred is uncommon or complex to resolve. Information should be gathered and assistance obtained from your service support organization.

Step 3090-1

Record the hardware error log. The hardware error log may be viewed as follows:

1. Start Diagnostics and select **Task Selection** on the Function Selection screen.

2. Select **Display Hardware Error Report**.

3. Select **Display Hardware Errors for PCI-X SCSI RAID Adapters**.

4. Select the adapter resource, or select all adapters resources if the adapter resource is not known.

5. On the Error Summary screen, look for an entry with a SRN corresponding to the problem which sent you here and select it.

   Note: If multiple entries exist for the SRN, some entries could be older versions or a problem has occurred on multiple entities (such as adapters, disk arrays, and devices). Older entries can be ignored, however, this MAP may need to be used multiple times if the same problem has occurred on multiple entities.

6. Select the hardware error log to view.

7. Go to Step 3090-2.

Step 3090-2

Collect any hardware error logged about the same time for the adapter.

Go to Step 3090-3.

Step 3090-3

Collect the current disk array configuration. The disk array configuration may be viewed as follows:

1. Start the PCI-X SCSI Disk Array Manager.

   a. Start Diagnostics and select **Task Selection** on the Function Selection screen.

b. Select **RAID Array Manager**.

    c. Select **PCI-X SCSI Disk Array Manager**.

2. Select **List PCI-X SCSI Disk Array Configuration**.

3. Select the PCI-X SCSI RAID Controller identified in the hardware error log.

4. Go to Step 3090-4.

Step 3090-4
Contact your service support organization.

Exit this procedure.

Finding an SRN Given an AIX Error Log

Normally Error Log Analysis (ELA) will examine the error logs and present a Service Request Number (SRN) to the user as appropriate. If you need to determine an SRN given an AIX error log, perform the following steps:

1. Display the error log using the AIX **errpt** command (for example **errpt** for a summary followed by `errpt -a -s timestamp` or `errpt -a -N resource_name`).

2. Ensure that the Error ID is of the form SISIOA_xxxx (for example SISIOA_ARY_DEGRADED).

    Note: Only Error IDs of the form SISIOA_xxxx are potentially related to disk arrays.

3. Locate the SENSE DATA in the Detail Data.

4. Identify the bytes 24-27 of the SENSE DATA from the 32 bytes shown.

    Note: Use the following example AIX Error Log to help you identify bytes 24-27.

5. The first four digits of the SRN, known as the Failing Function Code (FFC), can be found in the following table:

CCIN (bytes 24-25) of SENSE DATA	Failing Function Code (FFC)
2780	2527
5703 , 1975	2523
5709 , 1976	2523
571B, 1913	2514
571E	2512
571F	2513
572E	252B
573D, 1908	2529
574F	2D07
575B	252E

6. The second four digits of the SRN, known as the Reason Code is equal to bytes 26-27 of the SENSE DATA

For the example error log:
- Bytes 24-27 of the SENSE DATA are **5703 9030**
- The first four digits of the SRN, using **5703** in the table above, is **2523**
- The second four digits of the SRN is **9030**
- Thus the SRN would be **2323 - 9030**

Example AIX Error Log (Error ID = SISIOA_ARY_DEGRADED)

```
LABEL:          SISIOA_ARY_DEGRADED
IDENTIFIER:     8267403B
Date/Time:      Wed Dec 11 15:19:54 CST
Sequence Number: 2600
Machine Id:     000B51FD4C00
Node Id:        rokon
Class:          H
Type:           TEMP
Resource Name:  sisioa0
Resource Class: adapter
Resource Type:  14107802
Location:       17-08
VPD:
        Product Specific.(  ).......PCI-X Dual Channel Ultra320 SCSI RAID
                                    Adapter
        Part Number.................21P6477
        FRU Number..................21P6477
        Serial Number...............YL3382127011
        Manufacture ID..............0038
        EC Level....................3
        ROM Level.(alterable).......03090025
        Product Specific.(Z0).......58
        Device Specific.(YL)........U0.2-P1-I1


Description
DISK ARRAY PROTECTION SUSPENDED


        Recommended Actions
        PERFORM PROBLEM DETERMINATION PROCEDURES


Detail Data
SENSE DATA
00FF FFFF 1900 00F0 066B 0200 0201 0000 0309 0025 0000 0004 5703 9030 0000 0000
                                                            ^    ^
                                                            |    |
                                                            |    Last 4-digits of SRN (bytes 26-27)
                                  CCIN of Controller (bytes 24-25)
ARRAY INFORMATION
Location  S/N       RAID Level
ff-0,0    23897FFD 5

DEGRADED DISK S/N
0023ABA7

ARRAY MEMBER INFORMATION
Expected  Actual     Vendor/
Location  Location   Product ID        S/N
00-2,0    00-2,0     IBM    DFHSS4W  00661BF3
00-1,0    00-1,0     IBM    DFHSS4W  002395F4
00-3,0    00-3,0     IBM    DFHSS4W  0023ABA7
00-4,0    00-4,0     IBM    DFHSS4W  000F4089
00-5,0    00-5,0     IBM    DFHSS4W  001B10F2



ADDITIONAL HEX DATA
E210 0080 0400 0000 06F0 0017 00BF 632B 0000 0000 0000 036C 066B 0200 00FF FFFF
FFFF FFFF 1705 B004 0000 0000 0000 0000 0000 0000 0000 0000 2020 2020 2020 2020
2020 2020 2020 2020 2020 2020 2020 2020 3030 3030 3030 3030 0000 0000
```

Appendix A. Communications Statements

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

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

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

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

European Union (EU) Statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. The manufacturer cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of option cards supplied by third parties. Consult with your dealer or sales representative for details on your specific hardware.

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

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

International Electrotechnical Commission (IEC) Statement

This product has been designed and built to comply with IEC 60950.

United Kingdom Telecommunications Safety Requirements

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

The network adapter interfaces housed within this equipment are approved separately, each one having its own independent approval number. These interface adapters, supplied by the manufacturer, do not use or contain excessive voltages. An excessive voltage is one which exceeds 70.7 V peak ac or 120 V dc. They

interface with this equipment using Safe Extra Low Voltages only. In order to maintain the separate (independent) approval of the manufacturer's adapters, it is essential that other optional cards, not supplied by the manufacturer, do not use main voltages or any other excessive voltages. Seek advice from a competent engineer before installing other adapters not supplied by the manufacturer.

Avis de conformité aux normes du ministère des Communications du Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Canadian Department of Communications Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003..

VCCI Statement

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。　　　　ＶＣＣＩ－Ａ

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

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.

Electromagnetic Interference (EMI) Statement - Taiwan

警告使用者：
這是甲類的資訊產品，在
居住的環境中使用時，可
能會造成射頻干擾，在這
種情況下，使用者會被要
求採取某些適當的對策。

The following is a summary of the EMI Taiwan statement above.

Warning: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user will be required to take adequate measures.

Radio Protection for Germany

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

Der Aussteller der Konformitätserklärung ist die IBM Germany.

Dieses Gerät erfüllt die Bedingungen der EN 55022 Klasse A. Für diese von Geräten gilt folgende Bestimmung nach dem EMVG:

Geräte dürfen an Orten, für die sie nicht ausreichend entstört sind, nur mit besonderer Genehmigung des Bundesministers für Post und Telekommunikation oder des Bundesamtes für Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Störungen zu erwarten sind.

(Auszug aus dem EMVG vom 9.Nov.92, Para.3, Abs.4)

Hinweis

Dieses Genehmigungsverfahren ist von der Deutschen Bundespost noch nicht veröffentlicht worden.

Appendix B. 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.

Appendix C. AIX Command-Line Interface

Many tasks used to manage the PCI-X SCSI RAID Controller can be performed by using the AIX command line instead of using the PCI-X SCSI Disk Array Manager as described throughout this guide. The following table summarizes these commands.

General Help	sisraidmgr -h
"Viewing the Disk Array Configuration" on page 17	sisraidmgr -Ll '*controller name*' -j1
"Preparing Disks for Use in PCI-X SCSI Disk Arrays" on page 16	sisraidmgr -P -z *disk list* (For example, sisraidmgr -P -z `hdisk1 hdisk2 pdisk3 pdisk4')
"Changing pdisks to hdisks" on page 32	sisraidmgr -U -z `*pdisk list*'
"Creating a PCI-X SCSI Disk Array" on page 16	sisraidmgr -C -r `raid level' -s `stripe size' -z `pdisk list'
"Deleting a PCI-X SCSI Disk Array" on page 19	sisraidmgr -D -l '*controller name*' -d `*array name*'
"Adding Disks to an Existing Disk Array" on page 19	sisraidmgr -A -l `*array name*' -z `*pdisk list*'
"Creating Hot Spare Disks" on page 20	sisraidmgr -H -z `*pdisk list*'
"Deleting Hot Spare Disks" on page 21	sisraidmgr -I -z `*pdisk list*'
"Displaying Rechargeable Battery Information" on page 33	sisraidmgr -M -o0 -l '*adapter name*'
"Forcing a Rechargeable Battery Error" on page 34	sisraidmgr -M -o1 -l '*adapter name*'
"Recovery Procedures for Disk Failures" on page 59	sisraidmgr -R -z '*pdisk list*'

Index

S

V

Technical publication remarks form

Title :	ESCALA POWER6 PCI-X SCSI RAID Controller Reference Guide for AIX

Reference:	86 A1 59EV 00	Date:	November 2007

ERRORS IN PUBLICATION

SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

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

NAME : _____ Date : _____

COMPANY : _____

ADDRESS : _____

_____

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

Bull - Documentation D^{ept.}
1 Rue de Provence
BP 208
38432 ECHIROLLES CEDEX
FRANCE
info@frec.bull.fr

Technical publications ordering form

To order additional publications, please fill in a copy of this form and send it via mail to:

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

Phone: +33 (0) 2 41 73 72 66
FAX: +33 (0) 2 41 73 70 66
E-Mail: srv.Duplicopy@bull.net

Reference	Designation	Qty
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
_ _ _ _ _ _ _ _ _ [_ _]		
[_ _] : The latest revision will be provided if no revision number is given.		

NAME: _____   Date:_____

COMPANY:_____

ADDRESS: _____

_____

PHONE: _____   FAX: _____

E-MAIL: _____

For Bull Subsidiaries:

Identification: _____

For Bull Affiliated Customers:

Customer Code: _____

For Bull Internal Customers:

Budgetary Section: _____

For Others: Please ask your Bull representative.

BULL CEDOC

357 AVENUE PATTON

B.P.20845

49008 ANGERS CEDEX 01

FRANCE

REFERENCE
86 A1 59EV 00