

PCI-X SCSI RAID Controller

# Reference Guide for Linux



PCI-X SCSI RAID Controller

# Reference Guide for Linux

**Note**

Before using this information and the product it supports, be sure to read the general information under "Product Warranties and Notices" included with your system unit.

**Second Edition (August 2004)**

Before using this information and the product it supports, read the information in "Safety Information" on page ix and "Notices" on page 205.

A reader's comment form is provided at the back of this publication. If the form has been removed, address comments to Publications Department, Internal Zip 9561, 11501 Burnet Road, Austin, Texas 78758-3493. To send comments electronically, use this commercial internet address: [aix6kpub@austin.ibm.com](mailto:aix6kpub@austin.ibm.com). Any information that you supply may be used without incurring any obligation to you.

© **International Business Machines Corporation 2003, 2004. All rights reserved.**

Note to U.S. Government Users Restricted Rights--Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Safety Information</b> . . . . .	ix
<b>Handling Static Sensitive Devices</b> . . . . .	xi
<b>About This Book</b> . . . . .	xiii
ISO 9000 . . . . .	xiii
Highlighting . . . . .	xiii
References to Linux Operating System . . . . .	xiii
Related Publications . . . . .	xiii
Trademarks . . . . .	xiv
<b>Part 1. Linux 2.6</b> . . . . .	<b>1</b>
<b>Chapter 1. PCI-X SCSI RAID Controller Overview</b> . . . . .	<b>3</b>
Disk Arrays . . . . .	4
Supported RAID Levels. . . . .	5
RAID Level 0 . . . . .	5
RAID Level 5 . . . . .	6
RAID Level 6 . . . . .	7
RAID Level 10 . . . . .	7
Disk Array Capacities . . . . .	9
RAID Level Summary . . . . .	9
Stripe-Unit Size. . . . .	9
Disk Array Overview . . . . .	10
Disk Arrays States . . . . .	11
Physical Disk States . . . . .	11
I/O Adapter States . . . . .	11
<b>Chapter 2. PCI-X SCSI RAID Controller Software</b> . . . . .	<b>13</b>
Verifying the Installation of the PCI-X SCSI RAID Controller Software . . . . .	13
<b>Chapter 3. Common PCI-X SCSI RAID Controller Tasks</b> . . . . .	<b>15</b>
Using iprconfig . . . . .	15
Viewing the Status of Disks and Disk Arrays . . . . .	15
Viewing Disk Array Status . . . . .	17
Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays . . . . .	18
Formatting to Advanced Function. . . . .	18
Formatting to JBOD . . . . .	18
Deleting a PCI-X SCSI RAID Disk Array . . . . .	18
Adding Disks to an Existing Disk Array . . . . .	19
Using Hot Spare Disks . . . . .	19
Creating Hot Spare Disks . . . . .	19
Deleting Hot Spare Disks . . . . .	20
Viewing and Changing PCI-X SCSI RAID Controller Bus Settings. . . . .	20
Setting Bus Speed at Boot . . . . .	21
Creating a PCI-X SCSI RAID Disk Array . . . . .	22
<b>Chapter 4. PCI-X SCSI RAID Controller Maintenance</b> . . . . .	<b>23</b>
Do's and Don'ts . . . . .	23
Updating the PCI-X SCSI RAID Controller Microcode . . . . .	24
Separating a Removable Cache Card From the Base Card . . . . .	24
Replacing the Cache Directory Card . . . . .	29
Rechargeable Battery Maintenance . . . . .	32

Displaying Rechargeable Battery Information . . . . .	32
Forcing a Rechargeable Battery Error . . . . .	33
Replacing the Rechargeable Cache Battery Pack. . . . .	33
Replacing Physical Disks . . . . .	36
Removing a Failed Disk . . . . .	37
Installing a New Disk . . . . .	37
Recovery Procedures for Disk Failures . . . . .	38
RAID Level 0 . . . . .	38
RAID Level 5 . . . . .	38
RAID Level 6 . . . . .	39
RAID Level 10 . . . . .	39
Reclaiming IOA Cache Storage . . . . .	40
<b>Chapter 5. Problem Determination and Recovery.</b> . . . .	<b>43</b>
Error Log Analysis . . . . .	43
Basic vi Commands . . . . .	44
Sample Error Logs . . . . .	45
Adapter Dump . . . . .	49
Identifying the Disk Array Problem . . . . .	50
Unit Reference Code (URC) Tables . . . . .	50
Maintenance Analysis Procedures (MAPs) . . . . .	54
MAP 3300 . . . . .	54
MAP 3310 . . . . .	54
MAP 3311 . . . . .	55
MAP 3312 . . . . .	56
MAP 3313 . . . . .	56
MAP 3320 . . . . .	57
MAP 3321 . . . . .	58
MAP 3330 . . . . .	58
MAP 3331 . . . . .	59
MAP 3332 . . . . .	61
MAP 3333 . . . . .	62
MAP 3334 . . . . .	63
MAP 3335 . . . . .	64
MAP 3337 . . . . .	65
MAP 3350 . . . . .	66
MAP 3351 . . . . .	69
MAP 3390 . . . . .	70

---

## Part 2. RedHat Enterprise Linux 3.0. . . . . 71

<b>Chapter 6. PCI-X SCSI RAID Controller Overview.</b> . . . .	<b>73</b>
Disk Arrays . . . . .	73
Supported RAID Levels . . . . .	74
RAID Level 0 . . . . .	75
RAID Level 5 . . . . .	76
RAID Level 10 . . . . .	77
Disk Array Capacities . . . . .	78
RAID Level Summary . . . . .	78
Stripe-Unit Size . . . . .	78
Disk Array Overview . . . . .	78
Disk Array Models . . . . .	79
Physical Disk Models . . . . .	80
Disk Arrays States . . . . .	80
Physical Disk States . . . . .	80
I/O Adapter States . . . . .	81

<b>Chapter 7. PCI-X SCSI RAID Controller Software</b>	83
Verifying the Installation of the PCI-X SCSI RAID Controller Software	83
<b>Chapter 8. Common PCI-X SCSI RAID Controller Tasks</b>	85
Using iprconfig	85
Viewing the Status of Disks and Disk Arrays	85
Viewing Parity Status	88
Disk Array Parity Status	88
Disk Parity Status	89
Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays	89
Formatting to Advanced Function.	89
Formatting to JBOD	90
Creating a PCI-X SCSI RAID Disk Array	90
Deleting a PCI-X SCSI RAID Disk Array	91
Adding Disks to an Existing Disk Array	91
Using Hot Spare Disks	92
Creating Hot Spare Disks	92
Deleting Hot Spare Disks	92
Viewing and Changing PCI-X SCSI RAID Controller Bus Settings.	92
<b>Chapter 9. PCI-X SCSI RAID Controller Maintenance</b>	95
Do's and Don'ts	95
Updating the PCI-X SCSI RAID Controller Microcode	96
Replacing the Cache Directory Card	96
Rechargeable Battery Maintenance	98
Displaying Rechargeable Battery Information	98
Forcing a Rechargeable Battery Error	99
Replacing the Rechargeable Cache Battery Pack.	99
Replacing Physical Disks	101
Removing a Failed Disk	102
Installing a New Disk.	102
Recovery Procedures for Disk Failures	102
RAID Level 0	102
RAID Level 5	103
RAID Level 10	103
Reclaiming IOA Cache Storage	104
<b>Chapter 10. Problem Determination and Recovery</b>	105
Error Log Analysis.	105
Basic vi Commands	106
Sample Error Logs	107
Adapter Dump	111
Identifying the Disk Array Problem	112
Unit Reference Code (URC) Tables	112
Maintenance Analysis Procedures (MAPs)	116
MAP 3300.	116
MAP 3310.	117
MAP 3311.	118
MAP 3312.	118
MAP 3313.	119
MAP 3320	120
MAP 3321	121
MAP 3330	121
MAP 3331	122
MAP 3332	124
MAP 3333	125

MAP 3334 . . . . .	125
MAP 3335 . . . . .	126
MAP 3337 . . . . .	127
MAP 3350 . . . . .	128
MAP 3351 . . . . .	131
MAP 3390 . . . . .	132
 <b>Part 3. SuSE Linux Enterprise Server 8 . . . . .</b>	 <b>133</b>
 <b>Chapter 11. PCI-X SCSI RAID Controller Overview . . . . .</b>	 <b>135</b>
Disk Arrays . . . . .	135
Supported RAID Levels . . . . .	136
RAID Level 0 . . . . .	137
RAID Level 5 . . . . .	138
RAID Level 10 . . . . .	139
Disk Array Capacities . . . . .	140
RAID Level Summary . . . . .	140
Stripe-Unit Size . . . . .	140
Disk Array Overview . . . . .	140
Disk Array Models . . . . .	141
Physical Disk Models . . . . .	142
Disk Arrays States . . . . .	142
Physical Disk States . . . . .	142
I/O Adapter States . . . . .	143
 <b>Chapter 12. PCI-X SCSI RAID Controller Software . . . . .</b>	 <b>145</b>
Verifying the Installation of the PCI-X SCSI RAID Controller Software . . . . .	145
 <b>Chapter 13. Common PCI-X SCSI RAID Controller Tasks . . . . .</b>	 <b>147</b>
Using sisconfig . . . . .	147
Viewing the Status of Disks and Disk Arrays . . . . .	147
Viewing Parity Status . . . . .	150
Disk Array Parity Status . . . . .	150
Disk Parity Status . . . . .	151
Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays . . . . .	151
Formatting to Advanced Function . . . . .	151
Formatting to JBOD . . . . .	152
Creating a PCI-X SCSI RAID Disk Array . . . . .	152
Deleting a PCI-X SCSI RAID Disk Array . . . . .	153
Adding Disks to an Existing Disk Array . . . . .	153
Using Hot Spare Disks . . . . .	154
Creating Hot Spare Disks . . . . .	154
Deleting Hot Spare Disks . . . . .	154
Viewing and Changing PCI-X SCSI RAID Controller Bus Settings . . . . .	154
 <b>Chapter 14. PCI-X SCSI RAID Controller Maintenance . . . . .</b>	 <b>157</b>
Do's and Don'ts . . . . .	157
Updating the PCI-X SCSI RAID Controller Microcode . . . . .	158
Separating a Removable Cache Card From the Base Card . . . . .	158
Replacing the Cache Directory Card . . . . .	163
Rechargeable Battery Maintenance . . . . .	165
Displaying Rechargeable Battery Information . . . . .	165
Forcing a Rechargeable Battery Error . . . . .	166
Replacing the Rechargeable Cache Battery Pack . . . . .	166
Replacing Physical Disks . . . . .	168
Removing a Failed Disk . . . . .	169



Installing a New Disk. . . . .	169
Recovery Procedures for Disk Failures . . . . .	169
RAID Level 0 . . . . .	169
RAID Level 5 . . . . .	170
RAID Level 10 . . . . .	170
Reclaiming IOA Cache Storage . . . . .	171
<b>Chapter 15. Problem Determination and Recovery. . . . .</b>	<b>173</b>
Error Log Analysis. . . . .	173
Basic vi Commands . . . . .	174
Sample Error Logs . . . . .	175
Adapter Dump . . . . .	179
Identifying the Disk Array Problem . . . . .	180
Unit Reference Code (URC) Tables . . . . .	180
Maintenance Analysis Procedures (MAPs) . . . . .	184
MAP 3300 . . . . .	184
MAP 3310 . . . . .	185
MAP 3311. . . . .	186
MAP 3312 . . . . .	186
MAP 3313 . . . . .	187
MAP 3320 . . . . .	188
MAP 3321 . . . . .	189
MAP 3330 . . . . .	189
MAP 3331 . . . . .	190
MAP 3332 . . . . .	192
MAP 3333 . . . . .	193
MAP 3334 . . . . .	193
MAP 3335 . . . . .	195
MAP 3337 . . . . .	195
MAP 3350 . . . . .	196
MAP 3351 . . . . .	199
MAP 3390 . . . . .	200
<b>Part 4. Appendixes . . . . .</b>	<b>201</b>
<b>Appendixes . . . . .</b>	<b>203</b>
Communications Statements . . . . .	203
Federal Communications Commission (FCC) Statement. . . . .	203
European Union (EU) Statement . . . . .	203
International Electrotechnical Commission (IEC) Statement. . . . .	203
United Kingdom Telecommunications Safety Requirements. . . . .	203
Avis de conformité aux normes du ministère des Communications du Canada . . . . .	204
Canadian Department of Communications Compliance Statement . . . . .	204
VCCI Statement . . . . .	204
Electromagnetic Interference (EMI) Statement - Taiwan . . . . .	204
Radio Protection for Germany . . . . .	204
Notices. . . . .	205
<b>Index . . . . .</b>	<b>207</b>



---

## 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 for various versions of the Linux kernel. 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:

<b>Bold</b>	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.
<i>Italics</i>	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 Linux Operating System

This document may contain references to the Linux operating system, and is intended to be used with systems running the Linux kernel. Three different versions of Linux are described in this document: Linux 2.6, SuSE Linux Enterprise Server 8, and RedHat Enterprise Linux 3.0. Make sure you are consulting the appropriate section of this document for the operating system you are using.

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. Linux 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
- IPR Linux Device Driver Web site, available on the internet at <http://sourceforge.net/projects/iprdd/>
- *RS/6000 @server pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*, order number SA38-0516. Also available on the internet at <https://techsupport.services.ibm.com/server/library/>
- *Linux Documentation Project* Web site, available on the internet at <http://www.tldp.org/>
- *Linux for IBM @server pSeries* Web site, available on the internet at <http://www-1.ibm.com/servers/eserver/pseries/linux/>
- *RS/6000 @server pSeries Diagnostic Information for Multiple Bus Systems*, order number SA38-0509. Also 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
- *Penguinppc* Web site, dedicated to Linux on PowerPC, available on the internet at <http://penguinppc.org/>

---

## Trademarks

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

- @server
- PowerPC
- pSeries
- RS/6000

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



---

## Part 1. Linux 2.6



## Chapter 1. PCI-X SCSI RAID Controller Overview

The PCI-X SCSI RAID Controller is a 133 MHz 64-bit PCI-X adapter. This adapter has the following features:

- PCI-X system interface
- Dual or Quad Channel, Ultra320 (320 MB/s) 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

Card Type	PCI Form Factor	Channels	RAID Levels Supported	Write Cache Size	Read Cache Size	Multi Initiator Support	Adapter FFC	Recharge Battery Technology	Cache Battery Pack Concurrent Replace	Cache Battery Pack FFC
2780	Yes	4	0, 5, and 10	Up to 757 MB (comp)	Up to 1024 MB (comp)	No	2527	Lilon	Yes	2D01
5702, 1974	Yes	2	None	0 MB	0 MB	No	2522	None	N/A	N/A
5703, 1975	Yes	2	0, 5, and 10	40 MB	0 MB	No	2523	NiMH	No	2526
5709, 1976	No <sup>1</sup>	2	0, 5, and 10	16 MB	0 MB	No	2523	NiMH	No	2526
571A, 1912	Yes	2	None	0 MB	0 MB	No	252D	None	N/A	N/A
571B, 1913	Yes	2	0, 5, 6, and 10	90 MB	0 MB	No	2514	Lilon	No	2D03
572E	No <sup>2</sup>	2	0 and 10	0 MB	0 MB	No	252B	N/A	N/A	N/A
573D, 1908	No	2	0, 5, and 10	40 MB	0 MB	No	2529	NiMH	No	2526

<sup>1</sup> 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.

<sup>2</sup> This PCI-X SCSI RAID Controller is integrated onto the system board and attaches to Serial Attached SCSI (SAS) disks via 3 Gb connections.

---

## 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 Linux in the same way as it would a single SCSI disk. For example, after creating a disk array, you can use Linux commands to make the disk array available to the system by partitioning and creating a file system on it.

The PCI-X SCSI RAID Controller is managed by the **iprconfig** utility. The **iprconfig** utility is the interface to the RAID configuration, 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 into Rescue mode and creating the disk array before installing Linux. 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.



---

## 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 disks 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 1. 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:




Figure 2. 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 an *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:

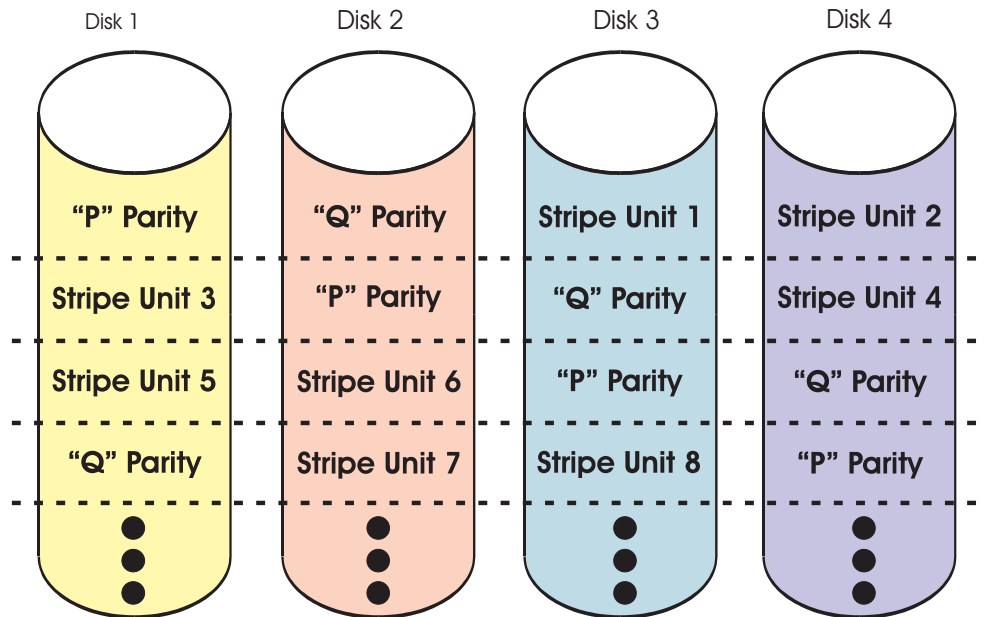


Figure 3. RAID Level 6

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:

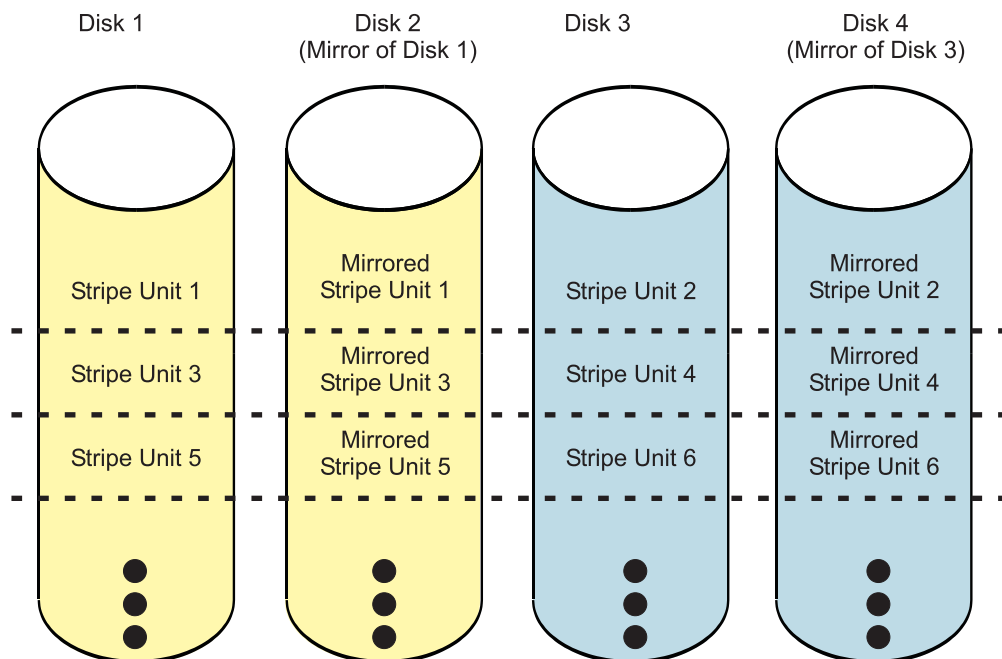


Figure 4. RAID Level 10

RAID level 10 can tolerate 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 *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 RAID Disk Array to 16 KB, 64 KB, or 256 KB. You may be able to maximize the performance of your PCI-X SCSI RAID 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.

**Note:** This guide and the **iprconfig** utility use common terminology for disk formats:

*JBOD* A *JBOD* disk is a disk that is formatted to 512 bytes/sector. JBOD stands for "Just a Bunch Of Disks".

*Advanced Function*

An *Advanced Function* disk is a disk that is formatted to 522 bytes/sector. This format allows disks to be used in disk arrays.

PCI-X SCSI RAID Disk Arrays are accessed in Linux as standard SCSI disk devices. These devices are automatically created when a disk array is created, and deleted whenever a disk array is deleted. The individual physical disks that comprise disk arrays or are candidates to be used in disk arrays, which are formatted for *Advanced Function*, are hidden from Linux and are accessible only through the **iprconfig** utility. Linux sees all *JBOD* disks. These disks must be formatted for *Advanced Function* before they can be used in disk arrays. For information on formatting *JBOD* disks to make them available for use in PCI-X SCSI RAID Disk Arrays, see "Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays" on page 18. The **Display Hardware Status** option in the **iprconfig** utility can be used to display these disks and the resource names associated with them. For details regarding how to view the disk information, see "Viewing the Status of Disks and Disk Arrays" on page 15. The following sample output is displayed when the **Display Hardware Status** option is invoked.

Display Hardware Status

Type option, press Enter.

1=Display hardware resource information details

OPT	Name	PCI/SCSI Location	Description	Status
		0000:02:04.0/2:	PCI-X SCSI RAID Adapter	Operational
		0000:02:04.0/2:1:14:0	Advanced Function Disk	Active
		0000:02:04.0/2:1:6:0	Advanced Function Disk	Active
		0000:02:04.0/2:1:8:0	Advanced Function Disk	Active
		0000:02:04.0/2:1:9:0	Advanced Function Disk	Active
sdt		0000:02:04.0/2:255:0:0	RAID 5 Disk Array	62% Rebuilt
		0000:02:04.0/2:1:10:0	RAID 5 Array Member	Active
		0000:02:04.0/2:1:11:0	RAID 5 Array Member	Active
		0000:02:04.0/2:1:12:0	RAID 5 Array Member	Active
		0000:02:04.0/2:1:13:0	RAID 5 Array Member	Active
sdb		0000:02:04.0/2:255:0:2	RAID 0 Disk Array	Active
		0000:02:04.0/2:1:4:0	RAID 0 Array Member	Active
sdc		0000:02:04.0/2:255:0:3	RAID 0 Disk Array	Active
		0000:02:04.0/2:1:3:0	RAID 0 Array Member	Active
sdd		0000:02:04.0/2:255:0:4	RAID 0 Disk Array	Active

More...

e=Exit    q=Cancel    r=Refresh    t=Toggle    f=PageDn    b=PageUp

Disk array, physical disk, and I/O adapter states are displayed in the fifth column of the **Display Hardware Status** screen. The following information provides a description for the possible states that may appear.

## Disk Arrays States

The fifth column in the preceding output displays the state of the disk array. The seven valid states for PCI-X SCSI RAID Disk Arrays are: *Active*, *Degraded*, *Rebuilding*, *R/W Protected*, *Missing*, *Offline*, and *Failed*.

**Active** The disk array is functional and fully protected (RAID level 5, 6, and 10) with all physical disks in the *Active* state.

### *Degraded*

The disk array's protection against disk failures is degraded or its performance is degraded. When one or more physical disks in the disk array are in the *Failed* state, the array is still functional but might no longer fully protected against disk failures. When all physical disks in the disk array are in the *Active* state, the array is not performing optimally because of a problem with the I/O adapter's nonvolatile write cache.

### *Rebuilding*

Data protection is being rebuilt on this disk array.

### *R/W Protected*

The disk array cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

### *Missing*

The disk array was not detected by the host operating system.

**Offline** The disk array has been placed offline due to unrecoverable errors.

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

## Physical Disk States

The fifth column in the preceding output displays the state of the device. The six possible states for physical disks are: *Active*, *Failed*, *Offline*, *Missing*, *R/W Protected*, and *Format Required*.

**Active** The disk is functioning properly.

**Failed** The IOA cannot communicate with the disk or the disk is the cause of the disk array being in the degraded state.

**Offline** The disk array has been placed offline due to unrecoverable errors.

### *Missing*

The disk was not detected by the host operating system.

### *R/W Protected*

The device cannot process a read nor write operation. A disk may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

### *Format Required*

The disk unit must be formatted to become usable on this IOA.

## I/O Adapter States

The fifth column in the preceding output displays the state of the I/O Adapter. The three possible states for I/O Adapters are: *Operational*, *Not Operational*, and *Not Ready*.

### *Operational*

The IOA is functional.

### *Not Operational*

The device driver cannot successfully communicate with this IOA.

*Not Ready*

The IOA requires a microcode download.

---

## Chapter 2. PCI-X SCSI RAID Controller Software

For the PCI-X SCSI RAID Controller to be identified and configured by Linux, the requisite device support software must be installed. Software for the PCI-X SCSI RAID Controller consists of a device driver and a set of utilities. The requisite software for the PCI-X SCSI RAID Controller is often preinstalled as part of the normal Linux installation. However, if the software package is not installed, software verification will fail. The missing package(s) can be installed from your Linux operating system CD-ROM. If you are missing components or need newer versions, obtain them from your Linux distributor or online at: <http://sourceforge.net/projects/iprdd/>

**Note:** Systems using the 2.4 Linux kernel require version 1.x of the **ipr** device driver and **iprutils**.  
Systems using the 2.6 Linux kernel require version 2.x of the **ipr** device driver and **iprutils**.

The PCI-X SCSI RAID Controller executes onboard microcode. Although a version of controller microcode may be distributed along with Linux, this does not necessarily represent the most recent version of microcode available for the controller. The **iprconfig** utility can be used to update the microcode being used by the controller. For more information regarding **iprconfig**, see “Updating the PCI-X SCSI RAID Controller Microcode” on page 24.

---

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

To verify that the **ipr** device driver for the PCI-X SCSI RAID Controller is installed, type:

```
modinfo ipr | grep version
```

The following is an example of the data that displays on your screen:

```
version:      2.0.9 DA9C6A0AA78C5D2B9D947A1
```

The following table describes the minimum ipr device driver version required for each supported adapter.

*Table 1. Minimum ipr device driver support*

Card Type	Minimum Supported kernel.org Version		Minimum Supported RedHat Enterprise Linux Version		Minimum Supported SuSE Enterprise Linux Version	
	ipr driver version	kernel version	ipr driver version	RHEL Version	ipr driver version	SLES Version
2780	2.0.10	2.6.8	2.0.11	RHEL 4	2.0.10	SLES 9
5702, 1974	2.0.10	2.6.8	2.0.11	RHEL 4	2.0.10	SLES 9
5703, 1975	2.0.10	2.6.8	2.0.11	RHEL 4	2.0.10	SLES 9
5709, 1976	2.0.10	2.6.8	2.0.11	RHEL 4	2.0.10	SLES 9
571A, 1912	2.0.13	2.6.12	2.0.11.1	RHEL 4 U2	2.0.10.3	SLES 9 SP2
571B, 1913	2.0.13	2.6.12	2.0.11.1	RHEL 4 U2	2.0.10.3	SLES 9 SP2
571B, 1913	2.0.13	2.6.12	2.0.11.1	RHEL 4 U2	2.0.10.3	SLES 9 SP2
572E	2.0.13	2.6.12	2.0.11.1	RHEL 4 U2	2.0.10.3	SLES 9 SP2
573D, 1908	2.0.13	2.6.12	2.0.11.1	RHEL 4 U2	2.0.10.3	SLES 9 SP2

To verify that the **iprconfig** utility is installed, type the command:

```
iprconfig --version
```

Output from this command will indicate if the utility is installed, and if so, version information will be displayed.

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 Linux distribution.

---

## 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 `iprconfig`

The interface for working with the PCI-X SCSI RAID Controller is **`iprconfig`**.

To start the **`iprconfig`** utility, type the command:

```
iprconfig
```

The main menu will display options for configuring and managing the PCI-X SCSI RAID Controller.

```
IBM Power RAID Configuration Utility

Select one of the following:

    1. Display hardware status
    2. Work with disk arrays
    3. Work with disk unit recovery
    4. Work with SCSI bus configuration
    5. Work with driver configuration
    6. Work with disk configuration
    7. Download microcode
    8. Analyze Log

Selection:

e=Exit
```

### Viewing the Status of Disks and Disk Arrays

The **`iprconfig`** utility offers an easy way to view the status of all devices controlled by the **`ipr`** device driver.

To view information about the disks and disk arrays on your system, do the following:

1. Run the **`iprconfig`** utility by typing `iprconfig`.
2. Select the **Display hardware status** option. The output displayed will be similar to the following:

## Display Hardware Status

Type option, press Enter.

1=Display hardware resource information details

OPT	Name	PCI/SCSI Location	Description	Status
		0000:01:01.0/0:	PCI-X SCSI Adapter	Operational
		0000:41:01.0/1:	PCI-X SCSI RAID Adapter	Operational
	sda	0000:41:01.0/1:0:3:0	Physical Disk	Active
	sdb	0000:41:01.0/1:0:4:0	Physical Disk	Active
	sdc	0000:41:01.0/1:1:3:0	Physical Disk	Active
	sdd	0000:41:01.0/1:1:4:0	Physical Disk	Active
	sde	0000:41:01.0/1:1:5:0	Physical Disk	Active
		0001:61:01.0/2:	PCI-X SCSI RAID Adapter	Operational
	sdf	0001:61:01.0/2:0:3:0	Physical Disk	Active
	sdg	0001:61:01.0/2:0:9:0	Physical Disk	Active
	sdh	0001:61:01.0/2:255:0:0	RAID 5 Disk Array	Active
		0001:61:01.0/2:0:4:0	RAID 5 Array Member	Active
		0001:61:01.0/2:0:5:0	RAID 5 Array Member	Active
		0001:61:01.0/2:0:6:0	RAID 5 Array Member	Active
		0001:61:01.0/2:0:8:0	RAID 5 Array Member	Active

e=Exit q=Cancel r=Refresh t=Toggle

- The first column is an input field used to select a device to display additional information.
- The second column of output is the device's resource name.
- The third column of output is the device's PCI/SCSI Location code. The format of this field is PCI Location/SCSI Host:SCSI Bus:SCSI Target:SCSI Lun.
- The fourth column of output is the device's description. Typing a 't' will toggle this field to be the device's Vendor/Product ID.
- The fifth column of output is the device's hardware status. For an overview of the possible disk and disk array hardware statuses, see "Disk Array Overview" on page 10.

To view information on a specific device, select the desired device with a 1 and press Enter. Output similar to the following will be displayed:



```

Disk Unit Hardware Resource Information Details

Manufacturer. . . . . : IBM
Product ID. . . . . : IC35L036UCDY10-0
Firmware Version. . . . . : 53323846 (S28F)
Serial Number . . . . . : E3V0E77B
Capacity. . . . . : 36.40 GB
Resource Name. . . . . : /dev/sda

Physical location:
PCI Address . . . . . : 0000:41:01.0
SCSI Host Number . . . . . : 1
SCSI Channel . . . . . : 0
SCSI Id. . . . . : 3
SCSI Lun . . . . . : 0

Extended Details
FRU Number . . . . . : 08K0293
EC Level . . . . . : H32224

Press Enter to continue.

e=Exit      q=Cancel    f=PageDn    b=PageUp
```

The previous screen shows an overview of a particular piece of hardware on your system. Multiple pages of information may be available.

## Viewing Disk Array Status

To view the disk array status, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Display disk array status**. The **Display Disk Array Status** screen will appear.

The following is sample output that will display:

```

Display Disk Array Status

Type option, press Enter.
1=Display hardware resource information details

OPT Name  PCI/SCSI Location      Description              Status
-----
sdh       0001:61:01.0/2:255:0:0  RAID 5 Disk Array       Active
          0001:61:01.0/2:0:4:0    RAID 5 Array Member     Active
          0001:61:01.0/2:0:5:0    RAID 5 Array Member     Active
          0001:61:01.0/2:0:6:0    RAID 5 Array Member     Active
          0001:61:01.0/2:0:8:0    RAID 5 Array Member     Active

e=Exit    q=Cancel  r=Refresh  t=Toggle
```

## Disk Array Status

The last column in the preceding output displays the status of each disk array. These states are described in “Disk Arrays States” on page 11.

## Physical Disk Status

The last column in the preceding output displays the status of each physical disk. These states are described in “Physical Disk States” on page 11.

---

## Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays

Before a disk can be used in a PCI-X SCSI RAID Disk Array, it must be formatted for *Advanced Function*. Before a disk is recognized as a standalone disk, it must be formatted to *JBOD*. Steps for performing both of these actions are contained in this section.

### Formatting to Advanced Function

To format a disk for *Advanced Function*, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Format device for RAID function**.
4. From the list of eligible disk units, choose the disks you want to format for *Advanced Function* and press Enter.

**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 for *Advanced Function*. These disks will not show up on the list of choices. In some cases, errors may be logged in the `/var/log/messages` file. For more detailed information, view that log file.

5. A message will display asking if you want to continue. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), resource names (such as `/dev/sdb`) may change when the system is rebooted. This could affect kernel command line entries and `fstab` entries.

After the formatting is complete, the disks will be ready for use in PCI-X SCSI RAID Disk Arrays.

### Formatting to JBOD

To format a disk back to *JBOD* format and allow Linux to use it as a standalone disk, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Format device for JBOD function (512)**.
4. From the list of eligible disk units, choose the disks you want to format to *JBOD* and press Enter.

**Attention:** Continuing with this option will format the disks. All data on the disks will be lost.

5. A message will display asking if you want to continue. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), it is advised that you reboot the system now. Resource names (such as `/dev/sdb`) may change as a result of reformatting. This could affect kernel command line entries and `fstab` entries.

---

## Deleting a PCI-X SCSI RAID Disk Array

**Attention:** After a disk array is deleted, it cannot be accessed. All data will be lost.

**Note:** If a disk array is currently being rebuilt or synched, that disk array cannot be deleted. A disk array that is currently in use can be deleted. Make sure the disk array is not being used before deleting it.

To preserve the data on the disk array, you must first back up all data that you wish to save. To delete the array, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Delete a disk array**.
4. From the listed disk arrays, select the one you wish to delete then press Enter.

**Attention:** All data on the selected drives will be lost when the disk array is deleted. If you are sure you want to delete the disk array, press Enter. If you do not want to delete the disk array, type `q` to cancel.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), resource names (such as `/dev/sdb`) may change when the system is rebooted. This could affect kernel command line entries and `fstab` entries.

---

## 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. This feature can be used to dynamically increase the capacity of a disk array while preserving existing data in the disk array. Extensive use of this feature, however, will result in a performance penalty because the data will not be restriped.

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Add a device to a disk array**.
4. Select the PCI-X SCSI RAID Disk Array to which you want to add the disk(s).
5. Select the disk(s) to be included into the disk array then press Enter.

**Attention:** All data on the selected drives will be lost when the disk(s) are added to the existing array.

6. If you are sure you want to include the disk(s), press Enter. If you do not want to include the disk(s) in the disk array, type `q` to cancel.

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 one or more 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 for *Advanced Function*.
- The disk belongs to another disk array or is configured as a *Hot Spare*.

---

## Using Hot Spare Disks

*Hot Spare* disks are used to automatically replace failed disks in a 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*. In order to assign a disk for use as a hot spare, it must be formatted for *Advanced Function*.

## Creating Hot Spare Disks

To create *Hot Spare* disks, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.

2. Select **Work with disk arrays**.
3. Select **Create a hot spare**.
4. Select the adapter for which you want to create hot spares and press Enter.
5. Select the disks that you want to designate as hot spares and press Enter.

**Attention:** All data on the selected drives will be lost when the disks are configured as hot spares.

6. If you are sure you want to configure the disks as hot spares, press Enter. If you do not want to configure the disks as hot spares, type q to cancel.

## Deleting Hot Spare Disks

To delete *Hot Spare* disks and free them for use in other disk arrays, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Delete a hot spare device**.
4. Select the *Hot Spare* disks to be deleted and press Enter.
5. If you are sure you want to reconfigure the disks, press Enter. If you do not want to reconfigure the disks, type q to cancel

---

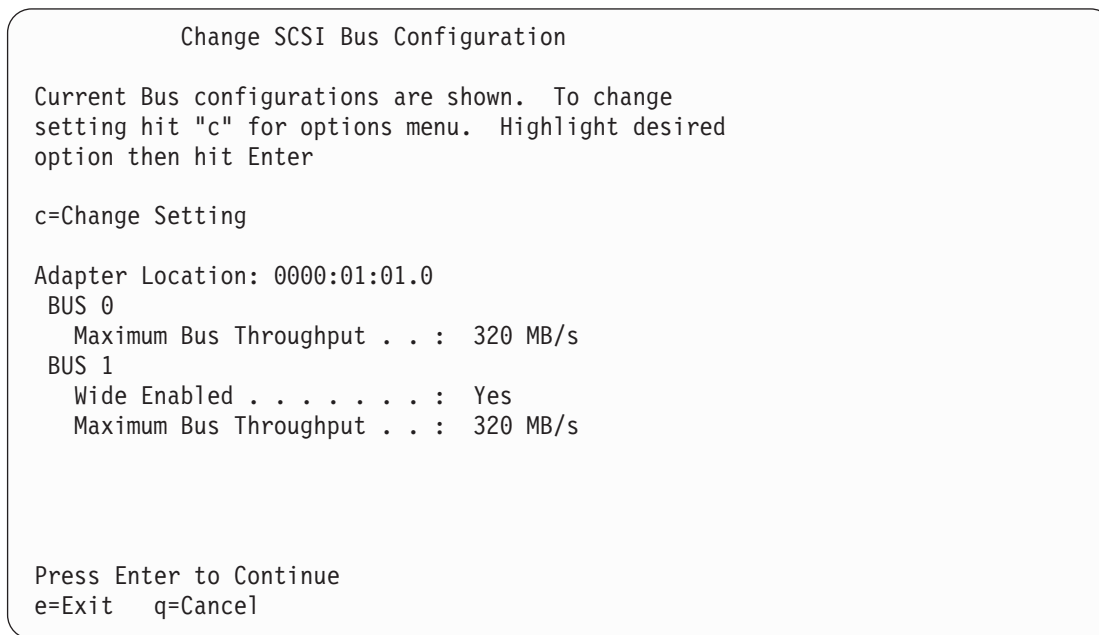
## Viewing and Changing PCI-X SCSI RAID Controller Bus Settings

**Note:** The maximum speed and bus width are automatically negotiated based on the characteristics of the attached devices. If you are not successful running at these speeds, you can reduce the adapter speed by changing the **Maximum bus throughput**.

**Note:** The driver will default to a maximum of 160 MB/s until the **iprinit** utility runs successfully. After **iprinit** completes, maximum configured bus speed is enabled. Each time the system is booted, **iprinit** is run automatically.

To view and change settings for the PCI-X SCSI RAID Controller, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with SCSI bus configuration**.
3. Select the adapter that you want to change settings on or view by typing 1 then pressing Enter. The **Change SCSI Bus Configuration** screen similar to the following will appear:



4. Follow the directions on screen to change settings, if desired.

The preceding screen displays information about the controller bus. The following is information regarding attributes displayed:

### Maximum Bus Throughput

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 options available will differ depending on your hardware configuration.

### Wide 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.

**Note:** If a device that requires 16-bit data transfer is attached, you will not have the option to disable wide transfer mode.

### Host SCSI ID

Specifies the SCSI ID to use for the controller. The default value is 7. Change this when in multi-initiator configurations.

## Setting Bus Speed at Boot

If running a 2.6 kernel, a module parameter is provided to allow for setting the maximum bus speed at boot time before the `iprinit` utility runs. Normally this should not be needed, but if you are having difficulty booting due to running 160 MB/s, the bus speed can be lowered to 80 MB/s by using `ipr.max_speed=0` on the kernel command line.

---

## Creating a PCI-X SCSI RAID Disk Array

A disk array is created using a set of disks that are formatted for *Advanced Function*. Disk arrays with data redundancy (RAID level 5, RAID level 6, and RAID level 10) will be unprotected until parity is finished being built. For disk arrays with no data redundancy (RAID 0), no parity is calculated and the array is built immediately. To create a PCI-X SCSI RAID Disk Array, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Create a disk array**.
4. Select the controller under which you would like to create a disk array and then press Enter.
5. Select the disk units to be included in the disk array then press Enter. A screen similar to the following will appear:

```
                Select Protection Level and Stripe Size

Default array configurations are shown.  To change
setting hit "c" for options menu.  Highlight desired
option then hit Enter

c=Change Setting

Protection Level . . . . . : RAID 5
Stripe Size . . . . . : 64 k
Queue Depth (default = 12). . . . . : 12

Press Enter to Continue

e=Exit   q=Cancel
```

6. Select the protection level desired for the array. For more information about selecting an appropriate RAID level, see “Supported RAID Levels” on page 5.
7. Select the stripe size desired in kilobytes for the array. For more information about selecting an appropriate stripe size, see “Stripe-Unit Size” on page 9.
8. Press Enter to continue.

**Attention:** All data on the selected drives will be lost when the disk array is created. If you are sure you want to create the disk array, press Enter. If you do not want to create the disk array, type q to cancel.

If you choose to create the disk array, a status screen displays until the operation is completed. If you wish to create additional disk arrays, you may leave the status screen by typing e. When a disk array has been built, it will be available to Linux as would any normal SCSI disk. To find the resource name that has been assigned to the disk array, refer to the **Display Disk Array Status** screen.

**Note:** If you are not using a device mapper (for example, LVM, md, or scsidev), it is advised that you reboot the system now. Resource names (such as `/dev/sdb`) may change as a result of creating a new disk array. This could affect kernel command line entries and `fstab` entries.

---

## Chapter 4. PCI-X SCSI RAID Controller Maintenance

This chapter provides information to assist in ensuring optimal performance of your PCI-X SCSI RAID Controller under Linux 2.6.

---

### 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 physical disks. Using the **modprobe -r ipr** command will have the same effect as a system shutdown. For systems that support PCI hotplug, PCI hotplug will also have the same effect as a system shutdown.

**Note:** A disk that is a failed member of a *Degraded* disk array can be replaced and the disk array rebuilt while the system continues to run.

- You can physically move disks from one adapter to another. However, if the physical disks 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 an *Degraded* state because of a disk failure.
- When physically removing disks 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 these disks are used.
- Always use the **Device Concurrent Maintenance** screen to remove and replace a physical disk. For instructions on how to remove and replace a disk, see "Replacing Physical Disks" on page 36.
- 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 into *Rescue* mode. Linux error logs, the **iprconfig** utility, and other tools are available 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. This is likely to make the problems worse. For additional information regarding problem determination, see Chapter 5, "Problem Determination and Recovery," on page 43.
- 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.
- If multiple errors occur at approximately the same time, look at them as a whole to determine if there may be a common cause.

---

## Updating the PCI-X SCSI RAID Controller Microcode

Disk drives attached to this RAID adapter must contain the latest microcode level.

The latest microcode level for your drive model and PCI-X SCSI RAID Controller can be downloaded from <http://techsupport.services.ibm.com/server/mdownload/download.html>.

Microcode can be updated using the **iprconfig** utility. The **iprconfig** utility allows for multiple levels of adapter and device microcode to exist on the system at the same time. After the adapter or device to update microcode is selected, all available microcode levels are presented to choose to download.

If you downloaded new microcode from the Web site, follow the instructions in the Readme on the Web site to install it. If there are no instructions available, install it as follows:

1. Install the package by typing:

```
rpm -ivh --ignoreos scsi.rpm
```

Press Enter.

2. Update the microcode by typing:

```
iprconfig
```

Press Enter.

3. Select **Download microcode** and press Enter.
4. Select the device or adapter you wish to update by typing a 1 then pressing Enter. Multiple devices may be selected.
5. Select the microcode image to be downloaded to the device or adapter by typing a 1 then pressing Enter.
6. Press Enter again to confirm the download.
7. If multiple devices were selected, repeat from Step 5 for the next device.

---

## Separating a Removable Cache Card From the Base Card

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


**Attention:** Do not disconnect the cache battery during this procedure or all cache data will be lost.

All cards are sensitive to electrostatic discharge. See Working with electrostatic discharge-sensitive parts.


**Note:** You will need a T-10 TORX tool to complete this procedure.

1. To separate the removable cache card from the base card, do the following:
  - a. Label both sides of the card before separating.
  - b. Place the adapter on an ESD protective surface and orientate it as shown in step 1c. Note the position of the card handle.
  - c. Remove three screws **A** from the base card as shown. Note the thread differences of the screws for correct replacement in the following steps.







d. Turn the card over.



e. Remove four screws **B** from the removable cache card as shown.




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




- g. Separate the base card from the removable cache card.

**Important:** Turn the jack screw back down fully clockwise.




- h. Place the adapter on an ESD protective surface and orientate it as shown in step 1c on page 24. Note the position of the card handle.
2. To attach the removable cache card to the base card, do the following:
  - a. 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.




- b. Attach the base card to the removable cache card by doing the following:


- 1) Line up the card connector (view the jack screw through the hole to help align).
- 2) Press on the stiffener plate to fully seat the connector.
- c. Attach four screws **A** to the removable cache card as shown.



- d. Turn the card over.



- e. Attach the three screws **B** to the base card as shown. Note the thread differences in the screws as pointed out in step 1c on page 24 of the separation procedure.



---

## Replacing the Cache Directory Card

**Attention:** All cards are sensitive to electrostatic discharge. See “Handling Static Sensitive Devices” on page xi. This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).

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 Figure 5 on page 30, Figure 6 on page 30, Figure 8 on page 31 or Figure 9 on page 32 below).
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.




Figure 5.




Figure 6.




Figure 7.




Figure 8.




Figure 9.

## Rechargeable Battery Maintenance

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. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Work with resources containing cache battery packs**.
4. Select your adapter and type 1. Then press Enter to display battery information.
5. The **Battery Information** screen displayed will look similar to the following example:

```

Battery Information

Resource Name . . . . . : /dev/sg10
Serial Number . . . . . : 03000001
Type . . . . . : 2780-001
PCI Address . . . . . : 0001:61:01.0
SCSI Host Number . . . . . : 2
Battery type . . . . . : Lithium Ion (LiIon)
Battery state . . . . . : No battery warning
Power on time (days) . . . . . : 56
Adjusted power on time (days) . . . . . : 61
Estimated time to warning (days) . . . . . : 865
Estimated time to error (days) . . . . . : 942
Concurrently maintainable battery pack . . : Yes
Battery pack can be safely replaced . . . : No

Press Enter to continue.

e=Exit      q=Cancel

```



The following information details each battery state that may be displayed on the **Battery Information** screen:

#### **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**

No warning 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**

The battery state cannot be determined.

## **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. 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, you will have to replace the Cache Battery Pack.

To force the battery error, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Work with resources containing cache battery packs**.
4. Select your adapter and type 2. Then press Enter to force the battery error.

**Attention:** This procedure should be done only under the guidance of IBM Hardware Service Support and poses serious risks to performance until the battery pack is replaced.

5. If you are sure you want to force a battery error, type `c` to confirm. If you do not want to force a battery error, type `q` to cancel.

## **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 15 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 33 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**

1. Determine that it is safe to replace the Cache Battery Pack. Refer to “Displaying Rechargeable Battery Information” on page 32. 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.





- 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

1. Determine that it is safe to replace the Cache Battery Pack. Refer to “Displaying Rechargeable Battery Information” on page 32. 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**.



**A** Battery Pack Cover  
**B** Cache Directory Card

3. Remove the plug that connects the battery unit and the card. It is necessary to squeeze 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:
  - Run the **iprconfig** utility by typing `iprconfig`.
  - Select **Work with disk unit recovery**.
  - Select **Work with resources containing cache battery packs**.
  - Start the IOA cache on the adapter whose battery you just replaced by typing a 3.
  - Press Enter.

## Replacing Physical Disks

Replace failed disks as soon as possible, even if a rebuild was initiated on a *Hot Spare* by the controller. Always use the **Device Concurrent Maintenance** option in the **iprconfig** utility to replace physical disks attached to a PCI-X SCSI RAID Controller.

## Removing a Failed Disk

Concurrent Device Remove

Choose a single location for remove operations  
1=Select

OPT	Name	PCI/SCSI Location	Description	Status
	sdc	0000:41:01.0/1:0:8:0	Physical Disk	Active
		0000:41:01.0/1:0:5:		Empty
	sdb	0000:41:01.0/1:0:4:0	Physical Disk	Active
	sda	0000:41:01.0/1:0:3:0	Physical Disk	Active
		0000:41:01.0/1:1:8:		Empty
	sdf	0000:41:01.0/1:1:5:0	Physical Disk	Active
	sde	0000:41:01.0/1:1:4:0	Physical Disk	Active
	sdd	0000:41:01.0/1:1:3:0	Physical Disk	Active
		0001:61:01.0/2:0:8:		Empty
	sdh	0001:61:01.0/2:0:9:0	Physical Disk	Active
	sdg	0001:61:01.0/2:0:3:0	Physical Disk	Active

e=Exit    q=Cancel    t=Toggle

To remove a failed disk, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Concurrent remove device**.
4. Select the device to be removed by typing a 1, then press Enter.
5. Verify the selected device is the device to be removed. The identify indicator should now be set. Press Enter.

**Attention:** **iprconfig** will allow removal of devices that are currently being used by the system. Ensure that the device is not in use by the system prior to removal to avoid data loss.

6. Remove the physical disk from the system.

**Attention:** Do not install the replacement disk at this time.

## Installing a New Disk

Concurrent Device Add

Choose a single location for add operations  
1=Select

OPT	Name	PCI/SCSI Location	Description	Status
		0000:41:01.0/1:0:5:		Empty
		0000:41:01.0/1:1:8:		Empty
		0001:61:01.0/2:0:8:		Empty

e=Exit    q=Cancel    t=Toggle

Once the failed disk has been removed, do the following to install a new disk:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Concurrent add device**.
3. Select the location to install the device by typing a 1. Press Enter.
4. Verify the location of the device to install. The identify indicator should now be set. Press Enter.
5. Install the physical disk into the system.
6. Press Enter.

If a *Hot Spare* disk was available and kept your disk array protected throughout the failure, no additional steps are required because the hot spare has replaced the failed disk unit in the disk array, and the newly installed disk is available to be configured as a hot spare. If a *Hot Spare* rebuild was not initiated by the controller, a rebuild needs to be initiated on the newly installed disk. To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk you wish to rebuild then press Enter.  
**Attention:** Data currently on the disk will be overwritten.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

---

## 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 and physical disk states, see “Disk Array Overview” on page 10. For a description of RAID levels, see “Supported RAID Levels” on page 5. This section describes recovery procedures for each RAID level.

### RAID Level 0

A single disk failure will cause a RAID level 0 array to transition to the *Failed* state. You will need to delete the disk array, replace the disk that is *Failed*, and recreate the disk array. You must then recreate the file systems on the disk array, and copy 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 fails, the disk array status changes to *Degraded*. The disk array remains functional because the data on the failed disk can be rebuilt using parity and data on the remaining disks.

Whenever a disk fails, replace it as soon as possible. If a *Hot Spare* disk was available, the controller may rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

## Multiple-Disk Failures

If a second disk in a RAID level 5 disk array fails, the array will also fail and the data will not be accessible. You will have to replace the failed disks, then delete and recreate the disk array. You must then recreate the file systems on the disk array and copy data 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 rebuilt 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 rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

### 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. You will have to replace the failed disks, then delete and recreate the disk array. You must then recreate the file systems on the disk array and copy data 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 fails, 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 fails, replace it as soon as possible. If a *Hot Spare* disk was available, the controller may rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

## 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 (Initiate the rebuild operation after you replace the failed disks.)

When both members of a mirrored pair fail, the disk array status changes to *Failed*. You will have to delete the disk array, replace the failed disks, and then recreate the disk array. You must then recreate the file systems on the disk array and copy data to the restored disk array from your backup media.

---

## Reclaiming IOA Cache Storage

IOA cache storage should only be reclaimed when indicated by a Maintenance Analysis Procedure (MAP) in Chapter 5, “Problem Determination and Recovery,” on page 43. This process is used to resolve cache problems.

**Attention:** Data may be lost.

To reclaim IOA cache storage, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Reclaim IOA cache storage**.
4. Select the desired adapter, then press Enter.
5. A screen will display showing the devices that may be affected. If you are sure you want to reclaim IOA cache storage, type `c` to confirm. If you do not want to reclaim IOA cache storage, type `q` to cancel.

After the process is complete, one of the following screens will display:

```
Reclaim IOA Cache Storage Results

IOA cache storage reclamation has completed. The number of
lost sectors could not be determined.

Press Enter to continue.

e=Exit  q=Cancel
```



### Reclaim IOA Cache Storage Results

IOA cache storage reclamation has completed. Use the number of lost sectors to decide whether to restore data from the most recent save media or to continue with possible data loss.

Number of lost sectors . . . . . : 1572608

Press Enter to continue.

e=Exit q=Cancel



---

## Chapter 5. 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. For additional sources of related service procedures required to service this subsystem, see “Related Publications” on page xiii.

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

- **ipr** error log entries associated with the problem, viewed using the **Analyze Log** option in **iprconfig** or directly from **/var/log/messages**
- Disk array and physical disk status, viewed using the **iprconfig** utility

Entries in the error log will contain a Unit Reference Code (URC) and a description of the problem. Each URC corresponds to an action that should be taken. 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 Linux or system documentation.

---

### Error Log Analysis

Errors logged by the adapter and devices are stored in the **/var/log/messages** file. While the experienced user can search through this file with commands such as **grep** and **sed**, the **iprconfig** utility also offers a menu for analyzing the error log.

To analyze the error log, select the **Analyze Log** option from the main menu in **iprconfig**. A screen will display similar to the following:

Kernel Messages Log

Select one of the following:

1. View most recent ipr error messages
2. View ipr error messages
3. View all kernel error messages
4. View iprconfig error messages
5. Set root kernel message log directory
6. Set default editor
7. Restore defaults
8. View IBM boot time messages

Selection:

e=Exit q=Cancel

The **Set root kernel message log directory** option allows you to change the directory where the error log is stored. Normally, this does not need to be used. It may be useful when booting from an alternate root file system if you wish to access **/var/log/messages** on another file system.

If you are more familiar with a text editor other than **vi**, choose the **Set default editor** option to change the text editor used when viewing the error log. Invoking **iprconfig** with the **-e** command line option will also change the default editor. Execute **iprconfig -h** for more information.

If at any time you wish to return the location of the error log to **/var/log/messages** and set the default editor back to **vi**, choose the **Restore defaults** option.

All other options on this screen gather certain information from the error log and present it in the default editor. This is used mostly during Maintenance Analysis Procedures to search for Unit Reference Codes.

## Basic vi Commands

Several frequently used **vi** commands follow:

**h** Move cursor to the right (in case arrow keys have no effect)

**j** Move cursor down (in case arrow keys have no effect)

**k** Move cursor up (in case arrow keys have no effect)

**l** Move cursor to the left (in case arrow keys have no effect)

**CTRL-f**

Page down (forward)

**CTRL-b**

Page up (backward)

**nG** Go to line number *n*. If *n* is omitted, go to the last line in file.

**/pattern**

Search forward for *pattern*. After searching, type *n* to repeat the search.

**?pattern**

Search backward for *pattern*. After searching, type *n* to repeat the search.

**n** Repeat the previous search

**:q** Quit **vi** and return to the **iprconfig** utility

## Example Using vi

To search through the most recent errors for the URC 3400, do the following:

1. From the **Kernel Messages Log** screen, select the **Use vi to view most recent ipr error messages** option.
2. Type **G** to move to the last line of the file.
3. Type **?3400** and press Enter to search backward for the pattern 3400.
4. If the first instance found is not part of an error log entry header, repeat the command by typing *n* as necessary.

## Sample Error Logs

When viewed through **iprconfig**, **ipr** error messages are displayed in the following forms:

Generic IOA or device error log entry

```
2:0:9:0: FFFE: Soft device bus error recovered by the IOA
00000000: 01080000 00000900 FFFFFFFF 1104E092
00000010: 00000000 00000002 00000000 00000000
00000020: 00000000 00000000 00000000 00000000
00000030: 00000000 00000000 0034D780 00000900
00000040: 00000311 041F1CC7 0000E092 0000E092
00000050: 0000E092 0000E092 0000E092 0000E092
00000060: 28282828 28282828 D000100F 00052800
00000070: 041F1CC7 00000100 00000000 90280000
00000080: 00000500 00000000 0C804000 00CCDDDD
00000090: 00000000 00000000 49000000 00315208
000000A0: 83000000 40010000 80000000 28E30000
000000B0: 000001F4 00000000 00000000 00000000
000000C0: 00000000 00000000 00000000 00000000
000000D0: 00000000 00000000 00000000 00000000
000000E0: 00000000 00000000 00000000 00000000
000000F0: 00000000 00000000 00000000 0034D780
00000100: 00000000 FFFFFFFF 00000000 00000000
00000110: 00000000 53544154 E00000A2 0034D780
```

## Device configuration error log entry

```
0001:61:01.0: 9002: IOA reserved area LRC error
Device Errors Detected/Logged: 1/1
-----
Device 1: 2:0:9:0
Vendor/Product ID: IBM      ST336605LC
Serial Number: 000227CE
-----New Device Information-----
Vendor/Product ID:
Serial Number: 00000000
Cache Directory Card Information:
Vendor/Product ID:
Serial Number: 00000000
Adapter Card Information:
Vendor/Product ID:
Serial Number: 00000000
Additional IOA Data: 00000000 00000000 00000000 00000000 00000000
```

#### Array error log entry

0001:61:01.0: 9030: Array no longer protected due to missing or failed disk unit

-----  
RAID 10 Array Configuration: 2:255:0:0  
-----

Exposed Array Member 0:

Vendor/Product ID: IBM IC35L036UCDY10-0

Serial Number: E3V1WEAB

Current Location: 2:0:4:0

Expected Location: 2:0:4:0  
-----

Array Member 1:

Vendor/Product ID: IBM IC35L036UCDY10-0

Serial Number: E3V0J55B

Current Location: 2:0:5:0

Expected Location: 2:0:5:0

## Cache error log entry

```
0001:61:01.0: 9010: Cache data associated with attached devices cannot be found
-----Current Configuration-----
Cache Directory Card Information:
Vendor/Product ID: IBM      5703001
Serial Number: 03060038
Adapter Card Information:
Vendor/Product ID: IBM      5703001
Serial Number: 03060038
-----Expected Configuration-----
Cache Directory Card Information:
Vendor/Product ID:
Serial Number:
Adapter Card Information:
Vendor/Product ID: IBM      5703001
Serial Number: 03060038
Additional IOA Data: 00000000 00000000 152101D0
```



The first portion of the log entry is the header, which contains important information such as the Unit Reference Code, the error class, the IOA under which the error occurred, the error description text, and the hardware location. The last block of data in the entry is the IOA Error Data, which is data recorded by the IOA to assist with recovery, if the recommended Service Action did not solve the problem.

---

## Adapter Dump

Should the adapter ever experience a critical error that requires an adapter reset to recover, the **iprdump** utility will dump the contents of the adapter's memory into a dump file. This dump file contains information from the adapter and the device driver. This dump file is approximately 4 MB. **iprdump** is usually configured to start at boot time and run in the background until such an error occurs. Whenever a critical error occurs, **iprdump** logs a message in the system messages log containing the location and name of the dump file that corresponds to that error and the location of the affected adapter.

Critical errors require that IBM Hardware Service Support be contacted. They will need your error log in order to make a complete analysis and determine the cause of the critical error.

---

## Identifying the Disk Array Problem

A disk array problem is uniquely identified by a Unit Reference Code (URC). The URC indicates the specific problem that has occurred and must be obtained in order to determine which MAP to use.

An URC is provided in the **ipr** error log with each error. This URC describes the exact problem that has been detected and should be considered the primary means of identifying a problem. However, the **Display Hardware Status** screen within the **iprconfig** utility is also a useful tool in identifying a problem or confirming a problem described by the error log. For additional information, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Obtain the URC and proceed to the next section to obtain a more detailed description of the problem and to determine which Maintenance Analysis Procedure to use.

## Unit Reference Code (URC) Tables

Using the URC obtained from the **ipr** error log, use the following table to determine which MAP to use. For additional details regarding the error log, see “Identifying the Disk Array Problem.” For MAPs, see “Maintenance Analysis Procedures (MAPs)” on page 54.

**Note:** The following table includes only URCs that are associated with Maintenance Analysis Procedures contained in this document.

Table 2. URCs

URC	Description Text	Service Action	Failing Items
102E	Out of alternate sectors for disk storage	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 69.	Disk drive
3002	Addressed device failed to respond to selection	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 69.	Disk drive
3010	Disk device returned wrong response to IOA	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 69.	Disk drive
3020	Storage subsystem configuration error IOA detected a SCSI bus configuration error	Perform “MAP 3350” on page 66.	
3029	A device replacement has occurred	No action required.	
3100	Device bus error	A SCSI bus error occurred. Perform “MAP 3350” on page 66.	
3109	IOA timed out a device command	Perform “MAP 3350” on page 66.	
3110	Device bus error, message or command phase	Perform “MAP 3350” on page 66.	
3150	SCSI bus configuration error	Internal and external SCSI cables are connected to a SCSI bus at the same time. Perform “MAP 3350” on page 66.	
3400	Device, I/O Adapter, Any device on I/O bus, Signal cables	Exchange the failing items in the Failing Items list one at a time.	Device, I/O Adapter, Any device on I/O bus, Signal cables

Table 2. URCs (continued)

URC	Description Text	Service Action	Failing Items
7001	IOA sector reassignment successful	If three 7001 messages have occurred for the same disk drive location, then exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 69.	Disk drive
8008	A permanent cache battery pack failure occurred	Perform "MAP 3300" on page 54. If the problem is not resolved, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8009	Impending Cache Battery Pack failure	Perform "MAP 3300" on page 54.	
8150	Permanent IOA failure	Exchange the failing items in the Failing Items list one at a time.	I/O Adapter Any device on I/O bus
8150	PCI bus error	If two errors have occurred for the same I/O Adapter in 24 hours, exchange the failing items in the Failing Items list one at a time.	I/O Adapter Any device on I/O bus
8151	IOA microcode error	Update adapter microcode. (See Updating the PCI-X SCSI RAID Controller Microcode.) If the problem is not resolved, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8157	IOA error requiring IOA reset to recover	If two 8157 messages have occurred for the same I/O Adapter location, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
9000	IOA reserved area data check	Perform MAP 3390.	
9001	IOA reserved area invalid data pattern	Perform MAP 3337.	
9002	IOA reserved area LRC error	Perform MAP 3390.	
9008	IOA does not support functions expected by devices	Perform MAP 3330.	
9010	Cache data associated with attached devices cannot be found	Perform MAP 3320.	
9011	Cache data belongs to devices other than those attached	Perform MAP 3390.	
9020	Array missing 2 or more devices with only 1 device present	Perform MAP 3311.	
9021	Array missing 2 or more devices with 2 or more devices present	Perform MAP 3311.	
9022	Exposed array is missing a required device	Perform MAP 3311.	
9023	Array member(s) not at required physical locations	Perform MAP 3312.	
9024	Array not functional due to present hardware configuration	Perform MAP 3390.	
9025	Disk unit is not supported at its physical location	Perform MAP 3310.	

Table 2. URCs (continued)

URC	Description Text	Service Action	Failing Items
9026	Array not functional due to present hardware configuration	Perform MAP 3390.	
9027	Array is missing a device and parity is out of sync	Perform MAP 3313.	
9028	Maximum number of arrays already exist	Perform MAP 3390.	
9029	Incorrect hardware configuration change has been detected	Perform MAP 3390.	
9030	Array no longer protected due to missing or failed disk unit	Perform MAP 3310.	
9031	Array protection temporarily suspended, protection resuming	Perform MAP 3310.	
9032	Array exposed but still protected	Perform MAP 3310.	
9040	Array protection temporarily suspended, protection resuming	No action required. The array is synching. Wait until the synch is complete.	
9041	Array protection temporarily suspended	Background array parity checking detected and corrected errors Perform MAP 3390.	
9042	Corrupt array parity detected on specified device	Perform MAP 3390.	
9050	Required cache data cannot be located for a disk unit	Perform MAP 3331.	
9051	IOA cache data exists for a missing or failed device	Perform MAP 3332.	
9052	Cache data exists for a device that has been modified	Perform MAP 3390.	
9054	IOA resources not available due to previous problems	Perform MAP 3321.	
9060	One or more disk pairs are missing from an array	Perform MAP 3311.	
9061 9062	One or more disks are missing from an array	Perform MAP 3311.	
9063	Maximum number of functional arrays has been exceeded	Perform MAP 3390.	
9081	IOA detected device error	Exchange the failing items in the Failing Items list one at a time.	Disk drive, I/O adapter
9082	IOA detected device error	Exchange the failing items in the Failing Items list one at a time.	Disk drive I/O adapter signal cables backplane any device on I/O bus
9090	Disk unit has been modified after the last known status	Perform MAP 3333.	

Table 2. URCs (continued)

URC	Description Text	Service Action	Failing Items
9091	Incorrect hardware configuration change has been detected	Perform MAP 3333.	
9092	Disk unit requires initialization before use	Perform MAP 3334.	
FF3D	Soft IOA error recovered by the IOA	If 10 FF3D messages have occurred for the same I/O Adapter physical location within a week, then exchange the failing items in the Failing Items list one at a time.	I/O adapter
FFF3	Disk media format bad	Perform "MAP 3335" on page 64.	
FFF4	Disk device problem Command to logical unit failed Device microcode is corrupt Data transfer overlength error Data transfer underlength error	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 69.	Disk drive
FFF6	Failure prediction threshold exceeded	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform MAP 3351.	Disk drive
FFF6	Device bus error, message or command phase  Device hardware error recovered by the device  Device hardware error recovered by the IOA	No action required.	
FFF7	Media error recovered by IOA rewrite procedures  Media error recovered by device rewrite procedures	No action required.	
FFF9	Soft media error. Sector reassignment recommended  Device sector reassign successful	No action required.	
FFFA	Undefined device response recovered by the IOA	If 10 FFFA messages have occurred for the same disk drive location in a one-week time period, then exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 69.	Disk drive
FFFB	SCSI bus was reset  SCSI bus was reset by another initiator	No action required.	
FFFE	Soft device bus error recovered by the IOA	If 10 Statistical FFFE messages have occurred for the same disk drive location in a one-week time period, perform "MAP 3350" on page 66.	

---

## 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 “Unit Reference Code (URC) Tables” on page 50 to identify which MAP to use.

### MAP 3300

A permanent cache battery pack failure occurred (URC 8008) or Impending Cache Battery Pack failure (URC 8009).

#### Step 3300-1

Are you here because of a permanent Cache Battery Pack failure (URC 8008)?

**NO** Go to Step 3300-2.

**YES** Go to Step 3300-3.

#### Step 3300-2

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.

**Note:** Ensure that you have a replacement Cache Battery Pack before proceeding.  
Force the Cache Battery Pack into an error state by following the instructions outlined in “Forcing a Rechargeable Battery Error” on page 33.

Go to Step 3300-3.

#### Step 3300-3

Replace the rechargeable battery pack by following the instructions outlined in “Replacing the Rechargeable Cache Battery Pack” on page 33.

**Attention:** All cards are sensitive to electrostatic discharge.

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

### MAP 3310

Disk unit is not supported at its physical location (URC 9025), Array no longer protected due to missing or failed disk unit (URC - 9030), Array protection temporarily suspended (URC - 9031), Disk Array is degraded due to missing or failed disk (URC – 9032).

#### Step 3310-1

Identify the disk array by examining the error log and the **Display Disk Array Status** screen. For information regarding accessing the **Display Disk Array Status** screen, see “Viewing Disk Array Status” on page 17.

Go to Step 3310-2.

#### Step 3310-2

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

**NO** Go to Step 3310-3.

**YES** Go to Step 3310-4.

### Step 3310-3

The affected disk array should have a status of either *Rebuilt* or *Active* due to the use of a *Hot Spare* disk. Create a new *Hot Spare* disk for the disk array by doing the following:

1. Identify the failed disk using the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15. The failed disk should have a status of *Failed*.
2. Remove the failed disk and replace it with a new disk that will be used as the *Hot Spare*. For further details, see “Replacing Physical Disks” on page 36.
3. The new disk must be formatted for *Advanced Function* in order to be used as a hot spare. If the new disk needs to be formatted, see “Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays” on page 18.
4. Assign the new disk as a hot spare for the disk. For further details, see “Using Hot Spare Disks” on page 19.

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

### Step 3310-4

The *Failed* disk should be replaced by using the following procedure:

1. Identify the failed disk by using the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15. The failed disk should have a status of *Failed*.
2. Remove the failed disk, replace it with a new disk, and rebuild the disk unit data. For further details, see “Replacing Physical Disks” on page 36.

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

## MAP 3311

Array not functional due to present hardware configuration (URC 9020 / 9021 / 9022), One or more disk pairs are missing from an array (URC 9060), or One or more disks are missing from an array (URC 9061 / 9062).

### Step 3311-1

Identify the disks missing from the disk array by examining the error log and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3311-2.

### Step 3311-2

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 missing disks and install them in the correct physical locations in the system. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reset the adapter by performing the following:
  1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.
  2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 18.

**Attention:** All data on the disk array will be lost.
- Format the remaining members of the disk array, as follows:

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3312

Array member(s) not at required resource address (URC 9023).

### Step 3312-1

Identify the disks that are not at their required physical locations by using the error log entry and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3312-2.

### Step 3312-2

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 missing disks and install them in the correct physical locations in the system. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reset the adapter by performing the following:
  1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.
  2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 18.

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

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

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3313

Array not functional due to present hardware configuration (URC 9027).

### Step 3313-1

Identify the adapter and disks related to the error by examining the error log and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3313-2.



## Step 3313-2

Have the disks or adapter card been physically moved recently?

- NO**                      Contact your service support organization.
- YES**                     Go to Step 3313-3.

## Step 3313-3

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. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reset the adapter by performing the following:
  1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.
  2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 18.  
**Attention:** All data on the disk array will be lost.
- Format the remaining members of the disk array, as follows:  
**Attention:** All data on the disk array will be lost.
  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

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

## MAP 3320

Use this MAP to resolve the following problem:

Cache data associated with attached devices cannot be found (URC 9010).

### Step 3320-1

Has the system been powered off for several days?

- NO**                      Go to Step 3320-2.
- YES**                     Go to Step 3320-5.

### Step 3320-2

**Note:** Label all parts (*original* 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 29.
- The removable cache card from the *original* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

Go to step 3320-3.

### Step 3320-3

Has a new URC 9010 or URC 9050 occurred?

- NO**            Go to Step 3320-6.  
**YES**            Go to Step 3320-4.

### Step 3320-4

Was the new URC 9050?

- NO**            The new URC 9010.  
                 Reclaim the controller cache storage. See “Reclaiming IOA Cache Storage” on page 40.  
**YES**            Contact your Service Support organization

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

### Step 3320-5

If the system has been powered off for several days after an abnormal power-down, the Cache Battery Pack may be depleted. Do *not* replace the adapter nor the Cache Battery Pack.

Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 40.

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

### Step 3320-6

1. 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 29.
  - The removable cache card from the *new* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

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

## MAP 3321

IOA resources not available due to previous problems (URC 9054).

### Step 3321-1

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 3330

IOA does not support functions expected by devices (URC 9008).

### Step 3330-1

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3330-2.

### Step 3330-2

Have the adapter card or disks been physically moved recently?

**NO** Contact your service support organization.

**YES** Go to Step 3330-3.

### Step 3330-3

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. Then either IPL the system or logical partition (it may be necessary to boot into *Rescue* mode) or reset the adapter by performing the following:
  1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.
  2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.
- Format the remaining members of the disk array, as follows:

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

  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

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

## MAP 3331

Use this MAP to resolve the following problem:

Required cache data cannot be located for a disk unit (URC 9050).

### Step 3331-1

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

**NO** Go to Step 3331-3.

**YES** Go to Step 3331-2.

### Step 3331-2

**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 29.
- The removable cache card from the original storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

Go to Step 3331-8.

### Step 3331-3

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3331-4.

### Step 3331-4

Have the adapter card or disks been physically moved recently?

**NO** Contact your service support organization.

**YES** Go to Step 3331-5.

### Step 3331-5

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

**NO** Go to Step 3331-6.

**YES** Go to Step 3331-7.

### Step 3331-6

There are three possible ways to correct the problem. Perform only *one* of the following options that are listed by preference.

- Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 40.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 18.
- Format the remaining members of the disk array, as follows:

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

### Step 3331-7

The adapter and disks, identified in the preceding steps, 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 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 3331-8

Has a new URC 9010 or URC 9050 occurred?

**NO** Go to Step 3331-10.

**YES** Go to Step 3331-9.

### Step 3331-9

Was the new URC-9050?

**NO** The new URC was 9010. Reclaim IOA cache storage. See "Reclaiming IOA Cache Storage" on page 40.

**YES** Contact your service support organization.

### Step 3331-10

**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 new storage I/O adapter. (Refer to "Replacing the Cache Directory Card" on page 29.
- The removable cache card from the NEW storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to "Separating a Removable Cache Card From the Base Card" on page 24.

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

## MAP 3332

IOA cache data exists for a missing or failed device (URC 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 3332-1

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For further details, see "Viewing the Status of Disks and Disk Arrays" on page 15.

Go to Step 3332-2.

### Step 3332-2

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

**NO** Go to Step 3332-4.

**YES** Go to Step 3332-3.

### Step 3332-3

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 3332-4

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

**NO** Go to Step 3332-5.

**YES** Go to Step 3332-6.

### Step 3332-5

Reclaim IOA cache storage. See "Reclaiming IOA Cache Storage" on page 40.

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

### Step 3332-6

Have the adapter card or disks been physically moved recently?

**NO** Contact your service support organization.

**YES** Go to Step 3332-7.

### Step 3332-7

The adapter and disks, identified in the preceding steps, 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 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*.

## MAP 3333

Disk unit has been modified after the last known status (URC 9090) or Incorrect hardware configuration change has been detected (URC 9091).

### Step 3333-1

IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reset the adapter by performing the following:

1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.
2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.

Take action on any new errors which surface.

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

## MAP 3334

Disk unit requires initialization before use (URC 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 **Device Concurrent Maintenance** screen in **iprconfig** when concurrently removing and installing disks (see “Replacing Physical Disks” on page 36) 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 3334-1

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3334-2.

### Step 3334-2

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

**NO** Go to Step 3334-4.

**YES** Go to Step 3334-3.

### Step 3334-3

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 3334-4

Have the disks or adapter card been physically moved recently?

**NO** Go to Step 3334-5.

**YES** Go to Step 3334-6.

### Step 3334-5

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

**NO** Go to Step 3334-7.

**YES** Go to Step 3334-6.

### Step 3334-6

There are three possible ways to correct this problem. Perform only *one* of the following three options:

- IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reset the adapter by performing the following:
  1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.

2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.
- Restore the adapter and disks back to their original configuration. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reset the adapter by performing the following:
    1. Find the scsi host number associated with the adapter by using the Display Hardware Status screen. For more details see Viewing the Status of Disks and Disk Arrays.
    2. Use the `echo 1 > /sys/class/scsi_host/hostX/reset_host` command to reset the adapter, where X is the scsi host number from the previous step.
  - 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 3334-7

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

- Format the disks, as follows:
 

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

  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.
- If the disks are members of a disk array, delete the disk array. See “Deleting a PCI-X SCSI RAID Disk Array” on page 18.

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

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

## MAP 3335

Disk media format bad (URC 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 3335-1

Identify the affected disk by examining the error log and the **Display Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3335-2.

## Step 3335-2

Format the disks, as follows:

- Attention:** All data on the disks will be lost.
1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.



4. Select the drives you wish to format then press Enter.

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

## MAP 3337

IOA detected device error (URC 9001).

### Step 3337-1

A device configuration error has been detected. The configuration sectors on the device may be incompatible with the current I/O adapter.

Has the I/O adapter been replaced with a different type of I/O adapter, or have the devices been moved from a different type of I/O adapter to this one recently?

**NO**                      Go to Step 3337-2.

**YES**                     Go to Step 3337-3.

### Step 3337-2

Contact your service support organization.

Exit this procedure.

### Step 3337-3

Identify the affected disk(s) by examining the error log and the **Display Hardware Status** screen. For further details, see "Viewing the Status of Disks and Disk Arrays" on page 15.

Go to Step 3337-4.

### Step 3337-4

If the I/O Adapter has been replaced by a different type of I/O Adapter, reinstall the original adapter.

Go to Step 3337-5.

### Step 3337-5

If the disks involved have been moved from a different type of I/O Adapter to this one, return them to their original I/O Adapter.

Go to Step 3337-6.

### Step 3337-6

**Attention:** The next step will cause data loss on the disks involved. If the data on these disks is important, perform back up procedures now.

Delete the disk array(s). For further details, see "Deleting a PCI-X SCSI RAID Disk Array" on page 18.

Go to Step 3337-7.

### Step 3337-7

Restore the initial hardware configuration that caused the error.

Go to Step 3337-8.

## Step 3337-8

Create new disk arrays, then restore the saved data. For further details, see “Creating a PCI-X SCSI RAID Disk Array” on page 22.

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

## MAP 3350

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 Card Type 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.

**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 an 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 fail if functioning disks are removed. Also, additional problems may be created.

## Step 3350-1

Identify the SCSI bus on which the problem is occurring by examining the error log entry.

Go to Step 3350-2.

## Step 3350-2

Have changes been made to the SCSI configuration recently?

- |            |                    |
|------------|--------------------|
| <b>NO</b>  | Go to Step 3350-5. |
| <b>YES</b> | Go to Step 3350-3. |

## Step 3350-3

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)
- Ensure the SCSI bus does not have multi-initiators (for example, set up for a high-availability configuration)

**Note:** Multi-initiator support is not provided at this time.

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 3350-5.

**YES** Go to Step 3350-4.

### Step 3350-4

1. Power off the system or logical partition.
2. Correct the problem.
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus-related failure occur?

**NO** Go to Step 3350-14.

**YES** Go to Step 3350-5.

### Step 3350-5

Determine if any of the disk arrays on the adapter are in an *Degraded* state. For further details, see "Viewing Disk Array Status" on page 17.

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

**NO** Go to Step 3350-7.

**YES** Go to Step 3350-6.

### Step 3350-6

1. Identify the failed disk(s) by first finding disk arrays with a state of *Degraded* , then finding disks on those arrays with a state of *Failed*.
2. Remove the failed disk(s) from each *Degraded* disk array. For further details, see "Replacing Physical Disks" on page 36.
3. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO** Go to Step 3350-14.

**YES** Go to Step 3350-7.

### Step 3350-7

Are there any non-essential removable media devices (such as Tape, CDROM, and DVDROM) on the SCSI bus?

**NO** Go to Step 3350-10.

**YES** Go to Step 3350-8.

### Step 3350-8

1. Power off the system or logical partition.
2. Remove one of the non-essential removable media devices.
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

- NO** Go to Step 3350-9.  
**YES** Go to Step 3350-7.

### Step 3350-9

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 3350-14.

### Step 3350-10

Are there any non-essential disks that are not disk array members (such as 512 byte/sector *JBOD* disks, *Hot Spare* disks, or Advanced Function disks) on the SCSI bus?

- NO** Go to Step 3350-13.  
**YES** Go to Step 3350-11.

### Step 3350-11

1. Remove one of the non-essential disk devices. For further details, see “Replacing Physical Disks” on page 36.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

- NO** Go to Step 3350-12.  
**YES** Go to Step 3350-10.

### Step 3350-12

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 3350-14.

### Step 3350-13

Contact your service support organization.

Exit this procedure

### Step 3350-14

1. Reinstall any good devices that were removed during the course of this MAP.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.
3. Take action on any other non-SCSI bus related errors if they exist.

## MAP 3351

Use the following to determine what other Field Replaceable Units (FRUs) besides a disk, may need to be replaced in order to solve a problem.

You are sent here when an *Advanced Function* disk, a physical disk in 522 bytes/sector format, was identified as the primary FRU to replace in order to resolve 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.
- 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 Card Type 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 fail or become *Degraded* if functioning disks are removed. Also, additional problems may be created.

### Step 3351-1

Identify the SCSI bus on which the problem is occurring by examining the error log entry.

Go to Step 3351-2.

### Step 3351-2

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 or logical partition
2. Replace a component listed above
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

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

## MAP 3390

The problem that occurred is uncommon or complex to resolve. Information should be gathered and assistance obtained from your service support organization.

### Step 3390-1

Make a copy of the entire `/var/log/messages` and `/var/log/boot.msg` files.

Go to Step 3390-2.

### Step 3390-2

Collect the current disk array configuration. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3390-3.

### Step 3390-3

Collect any `ipr` dump files that may be applicable to the problem. They will be located in `/var/log/`.

Go to Step 3390-4

### Step 3390-4

Contact your service support organization.

Exit this procedure.

---

## Part 2. RedHat Enterprise Linux 3.0





---

## Chapter 6. PCI-X SCSI RAID Controller Overview

The PCI-X SCSI RAID Controller is a 133 MHz 64-bit PCI-X adapter. This adapter has the following features:

- PCI-X system interface
- Dual or Quad Channel, Ultra320 (320 MB/s) SCSI controllers
- Embedded PowerPC RISC Processor and Hardware XOR DMA Engine
- Non-volatile write cache
- Support for RAID (Redundant Array of Independent Disks) levels 0, 5, 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 and 10 disk arrays
  - Ability to increase the capacity of an existing RAID level 5 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

**Note:** Multi-initiator support is not provided at this time.

Card Type	PCI Form Factor	Channels	RAID Levels Supported	Write Cache Size	Read Cache Size	Recharge Battery Technology	Cache Battery Pack Concurrent Replace	Cache Battery Pack FFC
2780	Yes	4	0, 5, and 10	Up to 757 MB (comp)	Up to 1024 MB (comp)	Lilon	Yes	2D01
5702, 1974	Yes	2	None	0 MB	0 MB	None	N/A	N/A
5703, 1975	Yes	2	0, 5, and 10	40 MB	0 MB	NiMH	No	2526
5709, 1976	No *	2	0, 5, and 10	16 MB	0 MB	NiMH	No	2526

\* 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.

---

### 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 Linux in the same way as it would a single SCSI disk. For example, after creating a disk array, you can use Linux commands to make the disk array available to the system by partitioning and creating a file system on it.

The PCI-X SCSI RAID Controller is managed by the **iprconfig** utility. The **iprconfig** utility is the interface to the RAID configuration, 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 into Rescue mode and creating the disk array before installing Linux. 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.



## 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 disks 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, 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 10. 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 *Unprotected* 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:




Figure 11. 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 *Unprotected* mode. Whenever data is read from an *Unprotected* 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 *R/W Protected* 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:

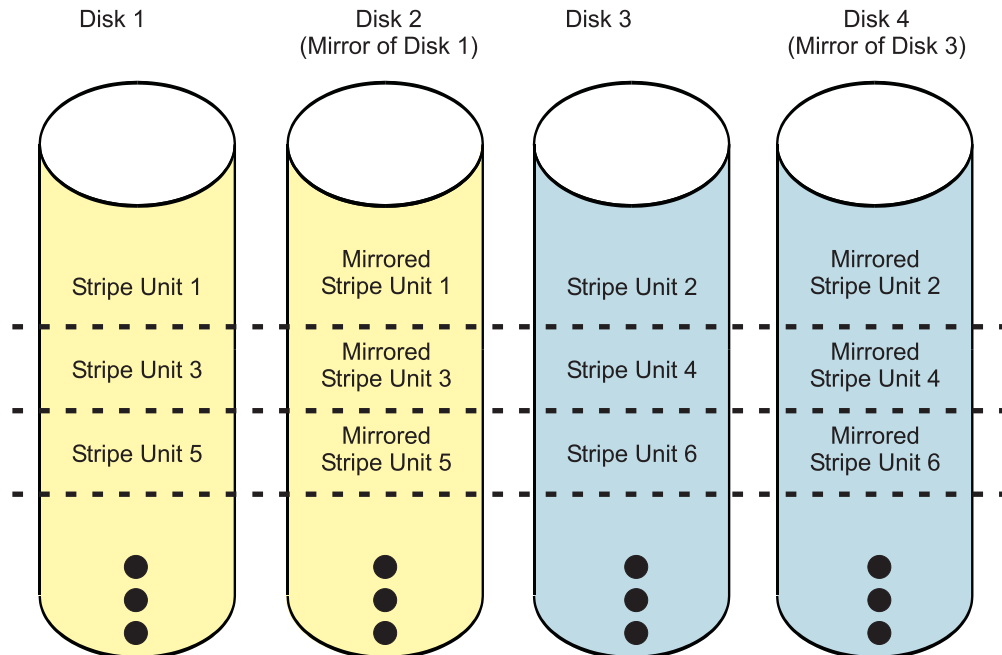


Figure 12. RAID Level 10

RAID level 10 can tolerate multiple disk failures. If one disk in each mirrored pair fails, the array will still be functional, operating in *Unprotected* 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 *R/W Protected* 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 *Unprotected* 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 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 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 10

Stores data redundantly on mirrored pairs to provide maximum protection against disk failures. Provides generally better performance than RAID level 5, 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 RAID Disk Array to 16 KB, 64 KB, or 256 KB. You may be able to maximize the performance of your PCI-X SCSI RAID 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 and 10 disk arrays also provide data redundancy, so that no data is lost if a single disk in the array fails.

**Note:** This guide and the **iprconfig** utility use common terminology for disk formats:

**JBOD** A **JBOD** disk is a disk that is formatted to 512 bytes/sector. **JBOD** stands for "Just a Bunch Of Disks".

**Advanced Function**

An **Advanced Function** disk is a disk that is formatted to 522 bytes/sector. This format allows disks to be used in disk arrays.

PCI-X SCSI RAID Disk Arrays are accessed in Linux as standard SCSI disk devices. These devices are automatically created when a disk array is created, and deleted whenever a disk array is deleted. The individual physical disks that comprise disk arrays or are candidates to be used in disk arrays, which are formatted for **Advanced Function**, are hidden from Linux and are accessible only through the **iprconfig** utility. Linux sees all **JBOD** disks. These disks must be formatted for **Advanced Function** before they can be used in disk arrays. For information on formatting **JBOD** disks to make them available for use in PCI-X SCSI RAID Disk Arrays, see "Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays" on page 89. The **Display Disk Hardware Status** option in the **iprconfig** utility can be used to display these disks and the resource names associated with them. For details regarding how to view the disk information, see "Viewing the Status of Disks and Disk Arrays" on page 85. The following sample output is displayed when the **Display Disk Hardware Status** option is invoked.

Display Disk Hardware Status					
Serial Number	Vendor ID	Product ID	Model	Resource Name	Hardware Status
03000001	IBM	5703001	001	/dev/iplr0	Operational
00009DB1	IBM	ST3146807LC	020	/dev/sdd	Operational
000012F0	IBM	ST3146807LC	090		HS/Active
00001315	IBM	ST3146807LC	079		DPY/Active
000012F5	IBM	ST3146807LC	070		DPY/Active
00001304	IBM	ST3146807LC	079		DPY/Active
E7800377	IBM	01C933F1	205	/dev/sdc	Operational

Press Enter to continue.

e=Exit                      q=Cancel                      r=Refresh    f=PageDn    b=PageUp  
d=Display disk unit details    p=Display device parity status

Disk array and physical disk model numbers are displayed in the fourth column of the **Display Disk Hardware Status** screen. The following information provides a description for the possible model numbers that may appear.

## Disk Array Models

The fourth column in the preceding output displays the model number of the disk array. The model number indicates the parity configuration.

200      RAID level 0

205      RAID level 5

210      RAID level 10

## Physical Disk Models

The fourth column in the preceding output displays the model number of the device. The model number indicates the device configuration.

- 020 JBOD disk unit resource. Linux sees this disk.
- 050 Advanced Function disk unit resource. This disk is not currently configured in a disk array.
- 07x Parity protected disk unit. The x describes the amount of the device that is used as parity.
- 9 unknown parity
  - 4 1/4 parity
  - 2 1/8 parity
  - 1 1/16 parity
  - 0 no parity
- 090 Hot Spare disk unit.

Disk array, physical disk, and I/O adapter states are displayed in the sixth column of the **Display Disk Hardware Status** screen. The following information provides a description for the possible states that may appear.

## Disk Arrays States

The sixth column in the preceding output displays the state of the disk array. The five valid states for PCI-X SCSI RAID Disk Arrays are: *Operational*, *DPY/Unprotected*, *DPY/Rebuilding*, *R/W Protected*, and *Perf Degraded*.

### *Operational*

The disk array is functional and fully protected (RAID level 5 and 10) with all physical disks in the *DPY/Active* state.

### *DPY/Unprotected*

The disk array is no longer protected by data redundancy due to a disk failure.

### *DPY/Rebuilding*

Data protection is being rebuilt on this disk array.

### *R/W Protected*

The disk array cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

### *Perf Degraded*

The device is functional, but performance may be affected because of other hardware problems (such as a problem with the IOA cache.)

## Physical Disk States

The sixth column in the preceding output displays the state of the device. The 11 possible states for physical disks are: *Operational*, *Not Operational*, *Not Ready*, *R/W Protected*, *DPY/Failed*, *DPY/Active*, *Perf Degraded*, *DPY/Unprotected*, *DPY/Rebuilding*, *Format Required*, and *HS/Active*.

### *Operational*

The disk unit is operational and ready to accept input or output operations.

### *Not Operational*

This device cannot communicate with the IOA. Verify that the unit is powered on.

### *Not Ready*

The device cannot perform media-related functions, but it can still communicate with the IOA.



#### *R/W Protected*

The device cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

#### *DPY/Failed*

This disk unit is part of a disk array that has device parity protection. The disk unit failed within its device parity set which caused the loss of data protection for the device parity set.

#### *DPY/Active*

This disk unit is part of a disk array that has device parity protection. The disk unit is operational and ready to accept input or output operations.

#### *Perf Degraded*

The device is functional, but performance may be affected because of other hardware problems (such as a problem with the IOA cache.)

#### *DPY/Unprotected*

This disk unit is part of a disk array that has device parity protection. Data protection is no longer in effect because of a failure of another resource.

#### *DPY/Rebuilding*

This disk unit is part of a disk array that has device parity protection. Data protection is being rebuilt.

#### *Format Required*

The disk unit must be formatted to become usable on this IOA.

#### *HS/Active*

The disk unit is configured to be a *Hot Spare*.

## **I/O Adapter States**

The sixth column in the preceding output displays the state of the I/O Adapter. The three possible states for I/O Adapters are: *Operational*, *Not Operational*, and *Not Ready*.

#### *Operational*

The IOA is functional.

#### *Not Operational*

The device driver cannot successfully communicate with this IOA.

#### *Not Ready*

The IOA requires a microcode download.



---

## Chapter 7. PCI-X SCSI RAID Controller Software

For the PCI-X SCSI RAID Controller to be identified and configured by Linux, the requisite device support software must be installed. Software for the PCI-X SCSI RAID Controller consists of a device driver and a set of utilities. The requisite software for the PCI-X SCSI RAID Controller is often preinstalled as part of the normal Linux installation. However, if the software package is not installed, software verification will fail. The missing package(s) can be installed from your Linux operating system CD-ROM. If you are missing components or need newer versions, obtain them from your Linux distributor or online at: <http://sourceforge.net/projects/iprdd/>

**Note:** Systems using the 2.4 Linux kernel require version 1.x of the **ipr** device driver and **iprutils**.  
Systems using the 2.6 Linux kernel require version 2.x of the **ipr** device driver and **iprutils**.

The PCI-X SCSI RAID Controller executes onboard microcode. Although a version of controller microcode may be distributed along with Linux, this does not necessarily represent the most recent version of microcode available for the controller. The **iprupdate** utility can be used to update the microcode being used by the controller. For more information regarding **iprupdate**, see “Updating the PCI-X SCSI RAID Controller Microcode” on page 96.

---

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

To verify that the **ipr** device driver for the PCI-X SCSI RAID Controller is installed, type:

```
modinfo ipr
```

Output from this command will indicate that the module is not found if the supporting software has not been installed. If information about **ipr** is displayed, type the following command to obtain version information:

```
strings filename | grep "ipr version"
```

where *filename* is the file name displayed after the **modinfo** command was typed in the previous step.

To verify that the **iprconfig** utility is installed, type the command:

```
iprconfig --version
```

Output from this command will indicate if the utility is installed, and if so, version information will be displayed.

**Attention:** To ensure proper functionality, the version numbers of **ipr** and **iprconfig** should be the same. Both the **ipr** device driver and the **iprconfig** utility should be Version 1.x or later.

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 Linux distribution.



---

## Chapter 8. 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 `iprconfig`

The interface for working with the PCI-X SCSI RAID Controller is **`iprconfig`**.

To start the **`iprconfig`** utility, type the command:

```
iprconfig
```

The main menu will display options for configuring and managing the PCI-X SCSI RAID Controller.

```
Work with Disk Units

Select one of the following:

    1. Display disk hardware status
    2. Work with device parity protection
    3. Work with disk unit recovery
    4. Work with configuration
    5. Analyze Log

Selection:

e=Exit
```

---

### Viewing the Status of Disks and Disk Arrays

The **`iprconfig`** utility offers an easy way to view the status of all devices controlled by the **`ipr`** device driver.

To view information about the disks and disk arrays on your system, do the following:

1. Run the **`iprconfig`** utility by typing `iprconfig`.
2. Select the **Display disk hardware status** option. The output displayed will be similar to the following:

### Display Disk Hardware Status

Serial Number	Vendor ID	Product ID	Model	Resource Name	Hardware Status
03000001	IBM	5703001	001	/dev/ipr0	Operational
00009DB1	IBM	ST3146807LC	020	/dev/sdd	Operational
000012F0	IBM	ST3146807LC	090		HS/Active
00001315	IBM	ST3146807LC	079		DPY/Active
000012F5	IBM	ST3146807LC	070		DPY/Active
00001304	IBM	ST3146807LC	079		DPY/Active
E7800377	IBM	01C933F1	205	/dev/sdc	Operational

Press Enter to continue.

e=Exit                      q=Cancel                      r=Refresh    f=PageDn    b=PageUp  
d=Display disk unit details      p=Display device parity status

- The first column of output is the device's 8-digit hexadecimal serial number.
- The second column of output is the device's vendor ID.
- The third column of output is the device's product ID. Different types of devices may have very dissimilar product IDs.
- The fourth column of output is the device's model. For an overview of the possible disk and disk array models, see "Disk Array Overview" on page 78.
- The fifth column of output is the device's Linux resource name. Devices that are hidden from Linux are not assigned resource names, and are only visible through **iprconfig**.
- The sixth column of output is the device's hardware status. For an overview of the possible disk and disk array hardware statuses, see "Disk Array Overview" on page 78.

For more information on the physical devices shown, type d to display disk unit details. Output similar to the following will be displayed:

### Display Disk Unit Details

Type option, press Enter.

5=Display hardware resource information details

OPT	Serial Number	Location	PCI Bus	PCI Dev	SCSI Host	SCSI Bus	SCSI ID	SCSI Lun
	03000001	P1-I2	64	12	2			
	000012F5	P1-I2/Z1-A4	64	12	2	0	4	0
	000012F0	P1-I2/Z1-A5	64	12	2	0	5	0
	00001315	P1-I2/Z1-A6	64	12	2	0	6	0
	00001304	P1-I2/Z1-A9	64	12	2	0	9	0

e=Exit                      q=Cancel                      f=PageDn    b=PageUp

The previous screen shows information pertaining to the location of the hardware on your system. To view information on a specific device, select the desired device with a 5 and press Enter. Output similar to the following will be displayed:

```
Disk Unit Hardware Resource Information Details

Manufacturer. . . . . : IBM
Machine Type and Model. . . : ST3146807LC
Firmware Version. . . . . : 43353046
Model . . . . . : 079
Serial Number . . . . . : 00001315
Capacity. . . . . : 142.87 GB

Physical location:
Location . . . . . : P1-I2/Z1-A6
PCI Bus. . . . . : 64
PCI Device . . . . . : 12
SCSI Host Number . . . . . : 2
SCSI Channel . . . . . : 0
SCSI Id. . . . . : 6
SCSI Lun . . . . . : 0

Press Enter to continue.

e=Exit      q=Cancel      r=Remove      v=Display Extended VPD
```

The previous screen shows an overview of a particular piece of hardware on your system. If extended information is available, press v to view it. Output similar to the following will be displayed:

```
Disk Unit Hardware Resource Extended Details

Manufacturer. . . . . : IBM
Machine Type and Model. . . : ST3146807LC
FRU Number. . . . . : 55P4081
Serial Number . . . . . : 000012F0
EC Level . . . . . : H31325
Part Number. . . . . : 55P4083
Device Specific (Z0). . . . : 000003129F00013E
Device Specific (Z1). . . . : 0626C50F
Device Specific (Z2). . . . : 0002
Device Specific (Z3). . . . : 02291
Device Specific (Z4). . . . : 0001
Device Specific (Z5). . . . : 22
Device Specific (Z6). . . . : H31325

Press Enter to continue.

e=Exit      q=Cancel
```

---

## Viewing Parity Status

The parity status can be viewed in two ways.

- To view the parity status from the main menu, do the following:
  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with device parity protection**.
  3. Select **Display device parity status**. The **Display Device Parity** screen will appear. Sample output follows.
- To view the parity status from the **Display Disk Hardware Status** screen, type `p`. The **Display Device Parity** screen will appear.

The following is sample output that will display:

Display Device Parity Status										
Parity Set	Serial Number	Model	Resource Name	PCI Bus	PCI Dev	SCSI Bus	SCSI ID	SCSI Lun	Status	
1	FF3BA7E4	210	/dev/sdc	64	12	255	0	0	Active	
	00249E55	070		64	12	1	0	0	Active	
	00C02D9E	070		64	12	1	1	0	Active	
2	DCC7BB45	200	/dev/sdd	64	12	255	0	1	Active	
	001DA3BA	070		64	12	1	2	0	Active	

e=Exit    q=Cancel    r=Refresh    f=PageDn    b=PageUp

## Disk Array Parity Status

The tenth column in the preceding output displays the parity status of each disk array. The five valid parity states for PCI-X SCSI RAID Disk Arrays are: *Active*, *R/W Prot*, *Synched*, *Rebuilt*, and *Unprot*.

**Active** This disk array is functioning properly.

### *R/W Prot*

The disk array cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

### *Synched*

The parity for this disk array is being built. Accompanying this state is a % complete.

### *Rebuilt*

The data on this disk array is being rebuilt from other units in the disk array. Accompanying this state is a % complete.

**Unprot** Due to one or more disk failures, this disk array is no longer fully protected by parity.



## Disk Parity Status

The tenth column in the preceding output displays the parity status of each disk. The seven valid parity states for disks are: *Active*, *Failed*, *Not Ready*, *R/W Prot*, *Synched*, *Rebuilt*, and *Unprot*.

*Active* This disk unit is part of a disk array and is fully operational.

*Failed* This disk unit is part of a disk array and has failed.

*Not Ready*

The device cannot perform media-related functions, but it can still communicate with the IOA.

*R/W Prot*

The device cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

*Synched*

The redundant data for the disk array is being built from the data within the disk array.

Accompanying this state is a % complete.

*Rebuilt*

This disk unit is part of a disk array that has device parity protection. The data on this disk unit is being rebuilt from other units in the disk array. Accompanying this state is a % complete.

*Unprot* This disk unit is part of a disk array that has device parity protection. This disk unit is operational. However, another disk unit in the disk array has failed or is being rebuilt. If another disk unit in the disk array fails, data could be lost.

---

## Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays

Before a disk can be used in a PCI-X SCSI RAID Disk Array, it must be formatted for *Advanced Function*. Before a disk is recognized as a standalone disk, it must be formatted to *JBOD*. Steps for performing both of these actions are contained in this section.

### Formatting to Advanced Function

To format a disk for *Advanced Function*, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Format device for advanced function**.
5. From the list of eligible disk units, choose the disks you want to format for *Advanced Function* and press Enter.

**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 for *Advanced Function*. These disks will not show up on the list of choices. In some cases, errors may be logged in the `/var/log/messages` file. For more detailed information, view that log file.

6. A message will display asking if you want to continue. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), resource names (such as `/dev/sdb`) may change when the system is rebooted. This could affect kernel command line entries and `fstab` entries.

After the formatting is complete, the disks will be ready for use in PCI-X SCSI RAID Disk Arrays.

## Formatting to JBOD

To format a disk back to *JBOD* format and allow Linux to use it as a standalone disk, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Format device for JBOD function**.
5. From the list of eligible disk units, choose the disks you want to format to *JBOD* and press Enter.  
**Attention:** Continuing with this option will format the disks. All data on the disks will be lost.
6. A message will display asking if you want to continue. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), it is advised that you reboot the system now. Resource names (such as `/dev/sdb`) may change as a result of reformatting. This could affect kernel command line entries and `fstab` entries.

---

## Creating a PCI-X SCSI RAID Disk Array

A disk array is created using a set of disks that are formatted for *Advanced Function*. Disk arrays with data redundancy (RAID level 5 and RAID level 10) will be unprotected until parity is finished being built. For disk arrays with no data redundancy (RAID 0), no parity is calculated and the array is built immediately. To create a PCI-X SCSI RAID Disk Array, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with device parity protection**.
3. Select **Start device parity protection**.
4. Select the controller under which you would like to start device parity protection then press Enter.
5. Select the disk units to be included in the disk array then press Enter. A screen similar to the following will appear:

```
Select Protection Level and Stripe Size

Default array configurations are shown. To change
setting hit "c" for options menu. Highlight desired
option then hit Enter

c=Change Setting

Protection Level . . . . . : RAID 5
Stripe Size . . . . . : 64 k

Press Enter to Continue

e=Exit   q=Cancel
```

6. Select the protection level desired for the array. For more information about selecting an appropriate RAID level, see “Supported RAID Levels” on page 74.
7. Select the stripe size desired in kilobytes for the array. For more information about selecting an appropriate stripe size, see “Stripe-Unit Size” on page 78.

8. Press Enter to continue. If you wish to create additional disk arrays, type a and go back to Step 5 on page 90. Otherwise, type c to continue.

**Attention:** All data on the selected drives will be lost when the disk array is created. If you are sure you want to create the disk array, press Enter. If you do not want to create the disk array, type q to cancel.

If you choose to create the disk array, a status screen displays until the operation is completed. When a disk array has been built, it will be available to Linux as would any normal SCSI disk. To find the resource name that has been assigned to the disk array, refer to the **Display Device Parity Status** screen.

**Note:** If you are not using a device mapper (for example, LVM, md, or scsidev), it is advised that you reboot the system now. Resource names (such as /dev/sdb) may change as a result of creating a new disk array. This could affect kernel command line entries and fstab entries.

---

## Deleting a PCI-X SCSI RAID Disk Array

**Attention:** After a disk array is deleted, it cannot be accessed. All data will be lost.

**Note:** A disk array that is currently in use or opened (that is busy) cannot be deleted. Also, if a disk array is currently being rebuilt or synched, that disk array cannot be deleted.

To preserve the data on the disk array, you must first back up all data that you wish to save. To delete the array, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with device parity protection**.
3. Select **Stop device parity protection**.
4. From the listed disk arrays, select the one you wish to delete then press Enter.

**Attention:** All data on the selected drives will be lost when the disk array is deleted. If you are sure you want to delete the disk array, press Enter. If you do not want to delete the disk array, type q to cancel.

**Note:** If you are not using a device mapper (for example, LVM, md, or scsidev), resource names (such as /dev/sdb) may change when the system is rebooted. This could affect kernel command line entries and fstab entries.

---

## Adding Disks to an Existing Disk Array

The PCI-X SCSI RAID Controller supports adding disks to existing RAID level 5 disk arrays. This feature can be used to dynamically increase the capacity of a disk array while preserving existing data in the disk array. Extensive use of this feature, however, will result in a performance penalty because the data will not be restriped.

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with device parity protection**.
3. Select **Include unit in device parity protection**.
4. Select the PCI-X SCSI RAID Disk Array to which you want to add the disk(s).
5. Select the disk(s) to be included into the disk array then press Enter.

**Attention:** All data on the selected drives will be lost when the disk(s) are added to the existing array.

6. If you are sure you want to include the disk(s), press Enter. If you do not want to include the disk(s) in the disk array, type q to cancel.

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 one or more 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 for *Advanced Function*.
- The disk belongs to another disk array or is configured as a *Hot Spare*.

---

## Using Hot Spare Disks

*Hot Spare* disks are used to automatically replace failed disks in a 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 *Unprotected*. In order to assign a disk for use as a hot spare, it must be formatted for *Advanced Function*.

### Creating Hot Spare Disks

To create *Hot Spare* disks, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Configure a hot spare device**.
5. Select the adapter for which you want to create hot spares and press Enter.
6. Select the disks that you want to designate as hot spares and press Enter.

**Attention:** All data on the selected drives will be lost when the disks are configured as hot spares.

7. If you are sure you want to configure the disks as hot spares, press Enter. If you do not want to configure the disks as hot spares, type `q` to cancel.

### Deleting Hot Spare Disks

To delete *Hot Spare* disks and free them for use in other disk arrays, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Unconfigure a hot spare device**.
5. Select the adapter from which you wish to remove a hot spare and press Enter.
6. Select the *Hot Spare* disks to be deleted and press Enter.
7. If you are sure you want to reconfigure the disks, press Enter. If you do not want to reconfigure the disks, type `q` to cancel.

---

## Viewing and Changing PCI-X SCSI RAID Controller Bus Settings

The maximum speed and bus width are automatically negotiated based on the characteristics of the attached devices. If you are not successful running at these speeds, you can reduce the adapter speed by changing the **Maximum bus throughput**.

**Note:** The driver will default to a maximum of 80 MB/s until the **iprupdate** utility runs successfully. After **iprupdate** completes, maximum configured bus speed is enabled. Each time the system is booted, **iprupdate** is run automatically. In order to achieve bus speeds of more than 80 MB/s, you must download the latest microcode for your adapter and install it on your system. For instructions, see "Updating the PCI-X SCSI RAID Controller Microcode" on page 96.

**Note:**

To view and change settings for the PCI-X SCSI RAID Controller, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with configuration**.
3. Select **Work with SCSI bus configuration**.
4. Select the adapter that you want to change settings on or view by typing 1 then pressing Enter. The **Change SCSI Bus Configuration** screen similar to the following will appear:

```
Change SCSI Bus Configuration

Current Bus configurations are shown. To change
setting hit "c" for options menu. Highlight desired
option then hit Enter

c=Change Setting

/dev/ipr0
BUS 0
  Maximum Bus Throughput . . : 320 MB/s
BUS 1
  Wide Enabled . . . . . : Yes
  Maximum Bus Throughput . . : 320 MB/s

Press Enter to Continue
e=Exit   q=Cancel
```

5. Follow the directions on screen to change settings, if desired.

The preceding screen displays information about the controller bus. The following is information regarding attributes displayed:

#### Maximum Bus Throughput

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 options available will differ depending on your hardware configuration.

#### Wide 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.

**Note:** If a device that requires 16-bit data transfer is attached, you will not have the option to disable wide transfer mode.



---

## Chapter 9. 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 physical disks. Using the **modprobe -r ipr** command will have the same effect as a system shutdown.

**Note:** A disk that is a failed member of an *Unprotected* disk array can be replaced and the disk array rebuilt while the system continues to run.

- You can physically move disks from one adapter to another. However, if the physical disks 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 an *Unprotected* state because of a disk failure.
- When physically removing disks 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 these disks are used.
- Always use the **Device Concurrent Maintenance** screen to remove and replace a physical disk. For instructions on how to remove and replace a disk, see "Replacing Physical Disks" on page 101.
- 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 into *Rescue* mode. Linux error logs, the **iprconfig** utility, and other tools are available 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. This is likely to make the problems worse. For additional information regarding problem determination, see Chapter 10, "Problem Determination and Recovery," on page 105.
- 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.
- If multiple errors occur at approximately the same time, look at them as a whole to determine if there may be a common cause.

---

## Updating the PCI-X SCSI RAID Controller Microcode

Disk drives attached to this RAID adapter must contain the latest microcode level.

The latest microcode level for your drive model and PCI-X SCSI RAID Controller can be downloaded from <http://techsupport.services.ibm.com/server/mdownload/download.html>.

Microcode can be updated using the **iprupdate** utility. If **iprupdate** finds newer microcode in `/etc/microcode` or `/usr/lib/microcode`, it will be downloaded to the adapter. The **iprupdate** utility should be set to run during boot time, after the **ipr** device driver is loaded.

If you downloaded new microcode from the Web site, follow the instructions in the Readme on the Web site to install it. If there are no instructions available, install it as follows:

1. Install the package by typing:

```
rpm -ivh --ignoreos scsi.rpm
```

Press Enter.

2. Update the microcode by typing:

```
iprupdate
```

Press Enter.

---

## Replacing the Cache Directory Card

**Attention:** All cards are sensitive to electrostatic discharge. See “Handling Static Sensitive Devices” on page xi. This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).

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 Figure 13, Figure 14 on page 97, Figure 15 on page 97 or Figure 16 on page 97 below).
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.




Figure 13. Cache battery pack and cache directory card for type 2782, 5703






Figure 14. Cache battery pack and cache directory card for type 5709




Figure 15. Cache battery pack and cache directory card for type 2780 (style A)




Figure 16. Cache battery pack and cache directory card for type 2780 (style B)

---

## Rechargeable Battery Maintenance

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. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Work with resources containing cache battery packs**.
4. Select your adapter and type 5. Then press Enter to display battery information.
5. The **Battery Information** screen displayed will look similar to the following example:

```
Battery Information

Resource Name . . . . . : /dev/ipr0
Serial Number . . . . . : 03000001
Type-model . . . . . : 5703-001
Location . . . . . : P1-I2
PCI Bus . . . . . : 64
PCI Device . . . . . : 12
SCSI Host Number . . . . . : 2
Battery type . . . . . : Nickel Metal Hydride (NiMH)
Battery state . . . . . : No battery warning
Power on time (days) . . . . . : 56
Adjusted power on time (days) . . . : 61
Estimated time to warning (days) . : 865
Estimated time to error (days) . . : 942

Press Enter to continue.

e=Exit      q=Cancel
```

The following information details each battery state that may be displayed on the **Battery Information** screen:

#### 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

No warning 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

The battery state cannot be determined.

## 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. Performing this task ensures that write caching is stopped prior to replacing a Cache Battery Pack. Forcing the battery error can be performed when you want to replace a Cache Battery Pack on one controller and you also want to replace other Cache Battery Packs on controllers in the same system that have not yet issued the cache battery warning nor cache battery error.

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, you will have to replace the Cache Battery Pack.

To force the battery error, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Work with resources containing cache battery packs**.
4. Select your adapter and type 2. Then press Enter to force the battery error.  
**Attention:** This procedure should be done only under the guidance of IBM Hardware Service Support and poses serious risks to performance until the battery pack is replaced.
5. If you are sure you want to force a battery error, type `c` to confirm. If you do not want to force a battery error, type `q` to cancel.

## 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 15 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 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.**

If the Cache Battery Pack is not already in the error state, follow the steps described in “Forcing a Rechargeable Battery Error” before proceeding. After the battery is in the error state, replace it with a new battery by taking the following steps:


1. Remove the adapter from the system by referring to your system documentation.
2. Use the following illustration to locate the Battery Pack Casing that holds the battery pack.
3. Squeeze the Battery Pack Casing to remove the battery unit from the card.
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 phase.

5. Remove the battery unit from the Battery Pack Casing. Save the Battery Pack Casing. The replacement battery pack does not provide a casing.

**Note:** Ensure that the Cache Battery Pack is disconnected for at least 15 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.



- A** Battery Pack Casing  
**B** Cache Directory Card

---

## Replacing Physical Disks

Replace failed disks as soon as possible, even if a rebuild was initiated on a *Hot Spare* by the controller. Always use the **Device Concurrent Maintenance** option in the **iprconfig** utility to replace physical disks attached to a PCI-X SCSI RAID Controller.

### Device Concurrent Maintenance

Type the choices, then press Enter.

Specify either Physical Location or PCI/SCSI location.

Physical Location . . . :

OR

PCI/SCSI . . . PCI Bus: PCI Device:

SCSI Channel: SCSI Id: SCSI Lun:

Specify action as 1=Remove device 2=Install device

Action to be performed . . . . . :

Enter a time value between 01 and 19.

Time needed in minutes . . . . . :

Or leave blank and press Enter to cancel

## Removing a Failed Disk

To remove a failed disk, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Device Concurrent Maintenance**.
4. Enter either the physical location code or the PCI/SCSI location of the device to be removed.
5. Specify **Remove device** as the action to be performed.
6. Enter a time value in minutes. This should be enough time for you to remove the specified hardware. During this time interval, **iprconfig** will wait and the LED(s) at the specified physical location will indicate that the device is ready for removal.
7. Remove the physical disk from the system.

**Attention:** Do not install the replacement disk at this time.

## Installing a New Disk

Once the failed disk has been removed, do the following to install a new disk:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Device Concurrent Maintenance**.
4. Enter either the physical location code or the PCI/SCSI location where the new device will be installed.
5. Specify **Install device** as the action to be performed.
6. Enter a time value in minutes. This should be enough time for you to install the specified hardware. During this time interval, **iprconfig** will wait and the LED(s) at the specified physical location will indicate that the new device can be installed.
7. Install the physical disk into the system.

If a *Hot Spare* disk was available and kept your disk array protected throughout the failure, no additional steps are required because the hot spare has replaced the failed disk unit in the disk array, and the newly installed disk is available to be configured as a hot spare. If a *Hot Spare* rebuild was not initiated by the controller, a rebuild needs to be initiated on the newly installed disk. To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Rebuild disk unit data**.
  4. Select the disk you wish to rebuild then press Enter.
- Attention:** Data currently on the disk will be overwritten.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

---

## 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 and physical disk states, see “Disk Array Overview” on page 10. For a description of RAID levels, see “Supported RAID Levels” on page 5. This section describes recovery procedures for each RAID level.

### RAID Level 0

A single disk failure will cause a RAID level 0 array to transition to the *R/W Protected* state. You will need to delete the disk array, replace the disk that is *Not Operational*, and recreate the disk array. You must then recreate the file systems on the disk array, and copy 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 fails, the disk array status changes to *Unprotected*. The disk array remains functional because the data on the failed disk can be rebuilt using parity and data on the remaining disks.

Whenever a disk fails, replace it as soon as possible. If a *Hot Spare* disk was available, the controller may rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

### Multiple-Disk Failures

If a second disk in a RAID level 5 disk array fails, the array will also fail and the data will not be accessible. You will have to replace the failed disks, then delete and recreate the disk array. You must then recreate the file systems on the disk array and copy data 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 fails, the disk array status changes to *Unprotected*. 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 fails, replace it as soon as possible. If a *Hot Spare* disk was available, the controller may rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

### Multiple-Disk Failures

If each failed disk belongs to a different mirrored pair, then the array will be in the *Unprotected* 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 (Initiate the rebuild operation after you replace the failed disks.)

When both members of a mirrored pair fail, the disk array status changes to *R/W Protected*. You will have to delete the disk array, replace the failed disks, and then recreate the disk array. You must then recreate the file systems on the disk array and copy data to the restored disk array from your backup media.

---

## Reclaiming IOA Cache Storage

IOA cache storage should only be reclaimed when indicated by a Maintenance Analysis Procedure (MAP) in Chapter 10, “Problem Determination and Recovery,” on page 105. This process is used to resolve cache problems.

**Attention:** Data may be lost.

To reclaim IOA cache storage, do the following:

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Reclaim IOA cache storage**.
4. Select the desired adapter, then press Enter.
5. A screen will display showing the devices that may be affected. If you are sure you want to reclaim IOA cache storage, type `c` to confirm. If you do not want to reclaim IOA cache storage, type `q` to cancel.

After the process is complete, one of the following screens will display:

```
Reclaim IOA Cache Storage Results

IOA cache storage reclamation has completed. The number of
lost sectors could not be determined.

Press Enter to continue.

e=Exit  q=Cancel
```

```
Reclaim IOA Cache Storage Results

IOA cache storage reclamation has completed. Use the number
of lost sectors to decide whether to restore data from the
most recent save media or to continue with possible data loss.

Number of lost sectors . . . . . : 1572608

Press Enter to continue.

e=Exit  q=Cancel
```



---

## Chapter 10. 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 physical disks, use the following to identify the problem:

- **ipr** error log entries associated with the problem, viewed using the **Analyze Log** option in **iprconfig** or directly from **/var/log/messages**
- Disk array and physical disk status, viewed using the **iprconfig** utility

Entries in the error log will contain a System Reference Code (SRC) and a description of the problem. Each SRC corresponds to an action that should be taken. 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 Linux or system documentation.

---

### Error Log Analysis

Errors logged by the adapter and devices are stored in the **/var/log/messages** file. While the experienced user can search through this file with commands such as **grep** and **sed**, the **iprconfig** utility also offers a menu for analyzing the error log.

To analyze the error log, select the **Analyze Log** option from the main menu in **iprconfig**. A screen will display similar to the following:

Kernel Messages Log

Select one of the following:

1. Use vi to view most recent IBM Storage error messages
2. Use vi to view IBM Storage error messages
3. Use vi to view all kernel error messages
4. Use vi to view iprconfig error messages
5. Set root kernel message log directory
6. Set default editor
7. Restore defaults
8. Use vi to view IBM Storage boot time messages

Selection:

e=Exit q=Cancel

The **Set root kernel message log directory** option allows you to change the directory where the error log is stored. Normally, this does not need to be used. It may be useful when booting from an alternate root file system if you wish to access **/var/log/messages** on another file system.

If you are more familiar with a text editor other than **vi**, choose the **Set default editor** option to change the text editor used when viewing the error log. Invoking **iprconfig** with the **-e** command line option will also change the default editor. Execute **iprconfig -h** for more information.

If at any time you wish to return the location of the error log to **/var/log/messages** and set the default editor back to **vi**, choose the **Restore defaults** option.

All other options on this screen gather certain information from the error log and present it in the default editor. This is used mostly during Maintenance Analysis Procedures to search for System Reference Codes.

## Basic vi Commands

Several frequently used **vi** commands follow:

**h** Move cursor to the right (in case arrow keys have no effect)

**j** Move cursor down (in case arrow keys have no effect)

**k** Move cursor up (in case arrow keys have no effect)

**l** Move cursor to the left (in case arrow keys have no effect)

**CTRL-f**

Page down (forward)

**CTRL-b**

Page up (backward)

**nG** Go to line number **n**. If **n** is omitted, go to the last line in file.

**/pattern**

Search forward for **pattern**. After searching, type **n** to repeat the search.

**?pattern**

Search backward for **pattern**. After searching, type **n** to repeat the search.

**n** Repeat the previous search

**:q** Quit **vi** and return to the **iprconfig** utility

## Example Using vi

To search through the most recent errors for the SRC ending in 3400, do the following:

1. From the **Kernel Messages Log** screen, select the **Use vi to view most recent IBM Storage error messages** option.
2. Type **G** to move to the last line of the file.
3. Type **?3400** and press Enter to search backward for the pattern 3400.
4. If the first instance found is not part of an error log entry header, repeat the command by typing **n** as necessary.

## Sample Error Logs

When viewed through **iprconfig**, IBM Storage error messages are displayed in the following forms:

### Device error log entry

```
Sep 22 03:48:56: begin-entry*****
Sep 22 03:48:56: SRC: 6600 FFF3
Sep 22 03:48:56: Class: Permanent
Sep 22 03:48:56: Disk media format bad
Sep 22 03:48:56: IOA Location: PCI Bus: 64, Device: 12 Location P1-I2
Sep 22 03:48:56: Device Location: P1-I2/Z1-A5
Sep 22 03:48:56: Device Serial Number: 000012F0
Sep 22 03:48:56: Device Manufacturer: IBM
Sep 22 03:48:56: Device Machine Type and Model: ST3146807LC
Sep 22 03:48:56: Device FRU Number: 55P4081
Sep 22 03:48:56: Device EC Level: H31325
Sep 22 03:48:56: Device Part Number: 55P4083
Sep 22 03:48:56: Device Specific (Z1): 0626C50F
Sep 22 03:48:56: Device Specific (Z2): 0002
Sep 22 03:48:56: Device Specific (Z3): 02291
Sep 22 03:48:56: Device Specific (Z4): 0001
Sep 22 03:48:56: Device Specific (Z5): 22
Sep 22 03:48:56: Device Specific (Z6): H31325
Sep 22 03:48:56: IOASC: 0x03310000
Sep 22 03:48:56: PRC: 0x1104E275
Sep 22 03:48:56: Driver version: Ver. 1 Rev. 0.5 SMP
Sep 22 03:48:56: IOA Firmware version: 0309003B
Sep 22 03:48:56: Predictive Analysis Seeks/256 counter: 0
Sep 22 03:48:56: Predictive Analysis Sectors Read/256 counter: 0
Sep 22 03:48:56: SCSI Sense Data: F0 00 03 00 00 00 00 0A 00 00 00 00 31 00 0C
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Sep 22 03:48:56: IOA Error Data:
Sep 22 03:48:56: Offset          0 1 2 3 4 5 6 7 8 9 A B C D E F
Sep 22 03:48:56: 00000000      00297C00 00000500 00000101 00000000
Sep 22 03:48:56: 00000010      3100E275 00000000 00000000 00000000
Sep 22 03:48:56: 00000020      00000000 00000000 00FFFFFF FFFFFFFF
```

## IOA error log entry

```

Sep 23 00:36:32: begin-entry*****
Sep 23 00:36:32: SRC: 5703 3400
Sep 23 00:36:32: Class: Informational
Sep 23 00:36:32: IOA detected device error
Sep 23 00:36:32: IOA Serial Number: 03000001
Sep 23 00:36:32: PCI Bus: 64, Device: 12 Location P1-I2
Sep 23 00:36:32: IOASC: 0x06290000
Sep 23 00:36:32: PRC: 0x10253380
Sep 23 00:36:32: Driver version: Ver. 1 Rev. 0.5 SMP
Sep 23 00:36:32: IOA Firmware version: 0309003B
Sep 23 00:36:32: IOA Error Data:
Sep 23 00:36:32: Offset          0 1 2 3 4 5 6 7 8 9 A B C D E F
Sep 23 00:36:32: 00000000      00000000 09A2C000 00000003 83000000
Sep 23 00:36:32: 00000010      80000000 80000000 09050435 30280000
Sep 23 00:36:32: 00000020      60008400 00000000 00000000 0006DDDD
Sep 23 00:36:32: 00000030      AABBBBBB 00000210 00BB0000 00000000
Sep 23 00:36:32: 00000040      00000000 0028B000 10000800 49000210
.
.
.
Sep 23 00:36:32: 000003B0      6A000050 005F82EC 00000030 A4E60F44
Sep 23 00:36:32: 000003C0      605C0050 01060409 003C0102 00000000
Sep 23 00:36:32: 000003D0      00000000 00000000
Sep 23 00:36:32: end-entry*****

```

## Device configuration error log entry

```
May 25 11:17:57: begin-entry*****
May 25 11:17:57: SRC: 5703 9092
May 25 11:17:57: Class: Permanent
May 25 11:17:57: Disk unit requires initialization before use
May 25 11:17:57: IOA Serial Number: 03060044
May 25 11:17:57: PCI Bus: 64, Device: 12 Location P1-I2
May 25 11:17:57: IOASC: 0x07279100
May 25 11:17:57: PRC: 0x1400B704
May 25 11:17:57: Driver version: Ver. 1 Rev. 0.5 SMP
May 25 11:17:57: IOA Firmware version: 0309001A
May 25 11:17:57: Device Errors Detected: 1
May 25 11:17:57: Device Errors Logged: 1
May 25 11:17:57: -----
May 25 11:17:57: Device 1:
May 25 11:17:57: IOA Location: PCI Bus: 64, Device: 12 Location P1-I2
May 25 11:17:57: Device Location: P1-I2/Z1-A1
May 25 11:17:57: Vendor ID: IBMAS400
May 25 11:17:57: Product ID: DDYS36M
May 25 11:17:57: Serial Number: 000003F8
May 25 11:17:57: New Device Information:
May 25 11:17:57: Vendor ID: IBMAS400
May 25 11:17:57: Product ID: DDYS36M
May 25 11:17:57: Serial Number: 000003F8
May 25 11:17:57: I/O Processor Information:
May 25 11:17:57: Vendor ID: IBM
May 25 11:17:57: Product ID: 5703001
May 25 11:17:57: Serial Number: 03060044
May 25 11:17:57: Cache Adapter Card Information:
May 25 11:17:57: Vendor ID: IBM
May 25 11:17:57: Product ID: 5703001
May 25 11:17:57: Serial Number: 03060044
May 25 11:17:57: Additional IOA Data: . . .
```

## Array error log entry

```
May 23 14:47:28: begin-entry*****
May 23 14:47:28: SRC: 5703 9030
May 23 14:47:28: Class: Permanent
May 23 14:47:28: Array no longer protected due to missing or failed disk unit
May 23 14:47:28: IOA Serial Number: 03060044
May 23 14:47:28: PCI Bus: 64, Device: 12 Location P1-I2
May 23 14:47:28: IOASC: 0x066B0200
May 23 14:47:28: PRC: 0x1705B004
May 23 14:47:28: Driver version: Ver. 1 Rev. 0.5 SMP
May 23 14:47:28: IOA Firmware version: 0309001A
May 23 14:47:28: -----
May 23 14:47:28: Array Member 0:
May 23 14:47:28:     Vendor ID: IBMAS400
May 23 14:47:28:     Product ID: DDYS36M
May 23 14:47:28:     Serial Number: 000003F8
May 23 14:47:28:     PCI/SCSI Address:
May 23 14:47:28:         Current: 40:0C/000100
May 23 14:47:28:         Expected: 40:0C/000100
May 23 14:47:28:     Physical Location:
May 23 14:47:28:         Current: P1-I2/Z1-A1
May 23 14:47:28:         Expected: P1-I2/Z1-A1
May 23 14:47:28: -----
May 23 14:47:28: Array Member 1:
May 23 14:47:28:     Vendor ID: IBMAS400
May 23 14:47:28:     Product ID: DGHS18U
May 23 14:47:28:     Serial Number: 00C1260A
May 23 14:47:28:     PCI/SCSI Address:
May 23 14:47:28:         Current: 40:0C/000200
May 23 14:47:28:         Expected: 40:0C/000200
May 23 14:47:28:     Physical Location:
May 23 14:47:28:         Current: P1-I2/Z1-A2
May 23 14:47:28:         Expected: P1-I2/Z1-A2
May 23 14:47:28: -----
May 23 14:47:28: Array Member 2:
May 23 14:47:28:     Vendor ID: IBMAS400
May 23 14:47:28:     Product ID: DGHS18U
May 23 14:47:28:     Serial Number: 00C124A7
May 23 14:47:28:     PCI/SCSI Address:
May 23 14:47:28:         Current: 40:0C/000400
May 23 14:47:28:         Expected: 40:0C/000400
May 23 14:47:28:     Physical Location:
May 23 14:47:28:         Current: P1-I2/Z1-A4
May 23 14:47:28:         Expected: P1-I2/Z1-A4
May 23 14:47:28: -----
May 23 14:47:28: end-entry*****
```

## Cache error log entry

```
May 27 15:00:27: begin-entry*****
May 27 15:00:27: SRC: 5703 9010
May 27 15:00:27: Class: Permanent
May 27 15:00:27: Cache data associated with attached devices cannot be found
May 27 15:00:27: IOA Serial Number: 03060044
May 27 15:00:27: PCI Bus: 64, Device: 12 Location P1-I2
May 27 15:00:27: IOASC: 0x07278100
May 27 15:00:27: PRC: 0x1400B23C
May 27 15:00:27: Driver version: Ver. 1 Rev. 0.5 SMP
May 27 15:00:27: IOA Firmware version: 0309001A
May 27 15:00:27: Current Configuration:
May 27 15:00:27:   I/O Processor Information:
May 27 15:00:27:       Vendor ID: IBM
May 27 15:00:27:       Product ID: 5703001
May 27 15:00:27:       Serial Number: 03060044
May 27 15:00:27:   Cache Adapter Card Information:
May 27 15:00:27:       Vendor ID: IBM
May 27 15:00:27:       Product ID: 5703001
May 27 15:00:27:       Serial Number: 03060044
May 27 15:00:27: Expected Configuration:
May 27 15:00:27:   I/O Processor Information:
May 27 15:00:27:       Vendor ID:
May 27 15:00:27:       Product ID:
May 27 15:00:27:       Serial Number:
May 27 15:00:27:   Cache Adapter Card Information:
May 27 15:00:27:       Vendor ID: IBM
May 27 15:00:27:       Product ID: 5703001
May 27 15:00:27:       Serial Number: 03060044
May 27 15:00:27: Additional IOA Data: 00000000 00000000 152101D0
```

The first portion of the log entry is the header, which contains important information such as the System Reference Code, the error class, the IOA under which the error occurred, and the hardware location. The last block of data in the entry is the IOA Error Data, which is data recorded by the IOA to assist with recovery, if the recommended Service Action did not solve the problem.

---

## Adapter Dump

Should the adapter ever experience a critical error that requires an adapter reset to recover, the **iprdump** utility will dump the contents of the adapter's memory into a dump file. This dump file contains information from the adapter and the device driver. This dump file is approximately 4 MB. **iprdump** is usually configured to start at boot time and run in the background until such an error occurs. Whenever a critical error occurs, **iprdump** logs a message in the system messages log containing the location and name of the dump file that corresponds to that error and the location of the affected adapter.

Critical errors require that IBM Hardware Service Support be contacted. They will need your error log in order to make a complete analysis and determine the cause of the critical error.

---

## Identifying the Disk Array Problem

A disk array problem is uniquely identified by a System Reference Code (SRC). An SRC is in the format *cccc uuuu*, where the first four digits of the SRC is the Card Type (for example 5703) and the last four digits of the SRC is known as the Unit Reference Code (URC). The URC indicates the specific problem that has occurred and must be obtained in order to determine which MAP to use.

An URC is provided in the **ipr** error log with each error. This SRC describes the exact problem that has been detected and should be considered the primary means of identifying a problem. However, the **Display Disk Hardware Status** screen within the **iprconfig** utility is also a useful tool in identifying a problem or confirming a problem described by the error log. For additional information, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Obtain the SRC and proceed to the next section to obtain a more detailed description of the problem and to determine which Maintenance Analysis Procedure to use.

## Unit Reference Code (URC) Tables

Using the SRC obtained from the **ipr** error log, use the following table to determine which MAP to use. For additional details regarding the error log, see “Identifying the Disk Array Problem.” For MAPs, see “Maintenance Analysis Procedures (MAPs)” on page 116.

**Note:** The following table includes only URCs that are associated with Maintenance Analysis Procedures contained in this document.

If the Card Type part of the SRC (the first four digits) is 432x, 660x, 671x, or 673x, the error pertains to an attached device and you should refer to the table of Device URCs. Otherwise, the error pertains to the adapter and you should refer to the table of Adapter URCs.

*Table 3. Device URCs*

URC	Description Text	Error Class	Service Action	Failing Items
102E	Out of alternate sectors for disk storage	Permanent	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 131.	Disk drive
3002	Addressed device failed to respond to selection	Recoverable	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 131.	Disk drive
3010	Disk device returned wrong response to IOA	Permanent	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 131.	Disk drive
3020	Storage subsystem configuration error	Permanent	Perform “MAP 3350” on page 128.	
3029	A device replacement has occurred	Informational	No action required.	
3109	IOA timed out a disk command	Recoverable	Perform “MAP 3350” on page 128.	
3110	Disk bus interface error occurred	Permanent	Perform “MAP 3350” on page 128.	
7000	Disk sector read error	Permanent	No action required.	



Table 3. Device URCs (continued)

URC	Description Text	Error Class	Service Action	Failing Items
7001	Temporary disk data error	Statistical	If two 7001 messages have occurred for the same disk drive location, then exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 131.	Disk drive
FFF3	Disk media format bad	Permanent	Perform "MAP 3335" on page 126.	
FFF4	Disk device problem	Permanent	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 131.	Disk drive
FFF5	Disk sector read error	Permanent	If two FFF5 messages have occurred for the same disk drive location in a one-day time period, then exchange the failing items in the Failing Items list one at a time.	Disk drive
FFF6	Disk device detected recoverable error	Statistical	No action required.	
FFF6	Disk device detected recoverable error	Threshold	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 131.	Disk drive
FFF7	Temporary disk data error	Temporary	No action required.	
FFF9	Temporary disk data error	Temporary	No action required.	
FFF9	Temporary disk data error	Statistical	No action required.	
FFFA	Temporary disk bus error	Statistical	If 10 Statistical FFFA messages have occurred for the same disk drive location in a one-week time period, then exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 131.	Disk drive
FFFB	SCSI bus reset occurred	Statistical	No action required.	
FFFE	Temporary disk bus error	Statistical	If 10 Statistical FFFE messages have occurred for the same disk drive location in a one-week time period, perform "MAP 3350" on page 128.	
FFFE	Temporary disk bus error	Informational	No action required.	

Table 4. Adapter URCs

URC	Description	Error Class	Service Action	Failing Items
302x	IOA detected a SCSI bus configuration error	Permanent	An error occurred on I/O bus x. Exchange the failing items in the Failing Items list one at a time.	Any device on I/O bus x I/O Adapter
310x	IOA detected interface error	Permanent	An error occurred on I/O bus x. Perform "MAP 3350" on page 128.	

Table 4. Adapter URCs (continued)

URC	Description	Error Class	Service Action	Failing Items
315x	SCSI bus configuration error	Permanent	Internal and external SCSI cables are connected to SCSI bus x at the same time. Perform “MAP 3350” on page 128.	
3400	IOA detected device error	Permanent, Recoverable	Exchange the failing items in the Failing Items list one at a time.	Device
				I/O Adapter
				Any device on I/O bus
				Signal cables
3400	IOA detected device error	Informational	No action required.	Device, I/O Adapter, Any device on I/O bus, Signal Cables
34FF	Disk device format in progress	Informational	No action required.	
8008	A permanent cache battery pack failure occurred	Permanent	Perform “MAP 3300” on page 116. If the problem is not resolved, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8009	Impending Cache Battery Pack failure		Perform “MAP 3300” on page 116.	
813x	IOA detected recoverable device bus error	Statistical	No action required.	
8140	IOA detected recoverable device bus error	Statistical	No action required.	
8141	IOA detected recoverable device bus error	Temporary, Statistical	No action required.	
8145	A recoverable IOA error occurred	Threshold	Exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8146	Disk device detected recoverable error	Threshold	Exchange the failing items in the Failing Items list one at a time.	Disk Drive
8150	Permanent IOA failure	Permanent	Exchange the failing items in the Failing Items list one at a time.	I/O Adapter
				Any device on I/O bus
8151	IOA Licensed Internal Code error	Permanent	Update adapter microcode. (See “Updating the PCI-X SCSI RAID Controller Microcode” on page 96.) If the problem is not resolved, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8157	Hardware Error, IOA error requiring IOA reset to recover	Permanent	I/O Adapter was reset. If the error persists, replace the I/O Adapter.	I/O Adapter
9000	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device
				I/O Adapter
9001	IOA detected device error	Permanent	Perform “MAP 3337” on page 127.	

Table 4. Adapter URCs (continued)

URC	Description	Error Class	Service Action	Failing Items
9002	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device
				I/O Adapter
				Signal cables
				Backplane
				Any device on I/O bus
9008	IOA does not support functions expected by devices	Permanent	Perform "MAP 3330" on page 121.	
9010	Cache data associated with attached devices cannot be found	Permanent	Perform "MAP 3320" on page 120.	
9011	Cache data belongs to devices other than those attached	Permanent	Perform "MAP 3390" on page 132.	
9020 9021 9022	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3311" on page 118.	
9023	Array member(s) not at required resource address	Permanent	Perform "MAP 3312" on page 118.	
9024	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3390" on page 132.	
9025	Disk unit is not supported at its physical location	Permanent	Perform "MAP 3310" on page 117.	
9026	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3390" on page 132.	
9027	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3313" on page 119.	
9028	Incorrect hardware configuration change has been detected	Permanent	Perform "MAP 3390" on page 132.	
9029	Incorrect hardware configuration change has been detected	Permanent	Perform "MAP 3390" on page 132.	
9030	Array no longer protected due to missing or failed disk unit	Permanent	Perform "MAP 3310" on page 117.	
9031	Array protection temporarily suspended	Temporary	Perform "MAP 3310" on page 117.	
9040	Array protection temporarily suspended	Temporary	No action required.  The array is synching. Wait until the synch is complete.	
9041	Background array parity checking detected and corrected errors	Temporary	Perform "MAP 3390" on page 132.	
9050	Required cache data cannot be located for a disk unit	Permanent	Perform "MAP 3331" on page 122.	
9051	IOA cache data exists for a missing or failed device	Permanent	Perform "MAP 3332" on page 124.	
9052	Cache data exists for a device that has been modified	Permanent	Perform "MAP 3390" on page 132.	

Table 4. Adapter URCs (continued)

URC	Description	Error Class	Service Action	Failing Items
9054	IOA resources not available due to previous problems	Permanent	Perform “MAP 3321” on page 121.	
9060	One or more disk pairs are missing from an array	Permanent	Perform “MAP 3311” on page 118.	
9061 9062	One or more disks are missing from an array	Permanent	Perform “MAP 3311” on page 118.	
9063	Maximum number of functional arrays has been exceeded	Permanent	Perform “MAP 3390” on page 132.	
9081	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device I/O Adapter
9082	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device I/O Adapter Signal cables Backplane Any device on I/O bus
9090	Disk unit has been modified after the last known status	Permanent	Perform “MAP 3333” on page 125.	
9091	Incorrect hardware configuration change has been detected	Permanent	Perform “MAP 3333” on page 125.	
9092	Disk unit requires initialization before use	Permanent	Perform “MAP 3334” on page 125.	
FF3D	Soft IOA error recovered by the IOA	Statistical	If 16 Statistical FF3D messages have occurred for the same I/O Adapter physical location, then exchange the failing items in the Failing Items list one at a time.	I/O Adapter

## 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 “Unit Reference Code (URC) Tables” on page 112 to identify which MAP to use.

### MAP 3300

A permanent cache battery pack failure occurred (SRC *cccc* 8008) or Impending Cache Battery Pack failure (SRC *cccc* 8009).

#### Step 3300-1

Are you here because of a permanent Cache Battery Pack failure (SRC *cccc* 8008)?

**NO** Go to Step 3300-2.

**YES** Go to Step 3300-3.

#### Step 3300-2

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.

**Note:** Ensure that you have a replacement Cache Battery Pack before proceeding.  
Force the Cache Battery Pack into an error state by following the instructions outlined in “Forcing a Rechargeable Battery Error” on page 99.

Go to Step 3300-3.

### Step 3300-3

Replace the rechargeable battery pack by following the instructions outlined in “Replacing the Rechargeable Cache Battery Pack” on page 99.

**Attention:** All cards are sensitive to electrostatic discharge.

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

## MAP 3310

Disk unit is not supported at its physical location (SRC cccc 9025), Array no longer protected due to missing or failed disk unit (SRC cccc - 9030), or Array protection temporarily suspended (SRC cccc - 9031).

### Step 3310-1

Identify the disk array by examining the error log and the **Display Device Parity Status** screen. For information regarding accessing the **Display Device Parity Status** screen, see “Viewing Parity Status” on page 88.

Go to Step 3310-2.

### Step 3310-2

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

**NO** Go to Step 3310-3.

**YES** Go to Step 3310-4.

### Step 3310-3

The affected disk array should have a status of either *Rebuilt* or *Active* due to the use of a *Hot Spare* disk. Create a new *Hot Spare* disk for the disk array by doing the following:

1. Identify the failed disk using the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85. The failed disk should have a status of *Not Operational* or *DPY/Failed*.
2. Remove the failed disk and replace it with a new disk that will be used as the *Hot Spare*. For further details, see “Replacing Physical Disks” on page 101
3. The new disk must be formatted for *Advanced Function* in order to be used as a hot spare. If the new disk needs to be formatted, see “Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays” on page 89.
4. Assign the new disk as a hot spare for the disk. For further details, see “Using Hot Spare Disks” on page 92.

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

### Step 3310-4

The *Failed* disk should be replaced by using the following procedure:

1. Identify the failed disk by using the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85. The failed disk should have a status of *Not Operational* or *DPY/Failed*.
2. Remove the failed disk, replace it with a new disk, and rebuild the disk unit data. For further details, see “Replacing Physical Disks” on page 101.

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

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

## MAP 3311

Array not functional due to present hardware configuration (SRC cccc 9020 / cccc 9021 / cccc 9022), One or more disk pairs are missing from an array (SRC cccc 9060), or One or more disks are missing from an array (SRC cccc 9061 / cccc 9062).

### Step 3311-1

Identify the disks missing from the disk array by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3311-2.

### Step 3311-2

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 missing disks and install them in the correct physical locations in the system. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ipr` command to remove the device driver.
  2. Use the `modprobe ipr` command to reload the device driver.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 91.

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

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

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3312

Array member(s) not at required resource address (SRC cccc 9023).

### Step 3312-1

Identify the disks that are not at their required physical locations by using the error log entry and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3312-2.

## Step 3312-2

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 missing disks and install them in the correct physical locations in the system. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ipr` command to remove the device driver.
  2. Use the `modprobe ipr` command to reload the device driver.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 91.  
**Attention:** All data on the disk array will be lost.
- Format the remaining members of the disk array, as follows:  
**Attention:** All data on the disk array will be lost.
  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

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

## MAP 3313

Array not functional due to present hardware configuration (SRC cccc 9027).

### Step 3313-1

Identify the adapter and disks related to the error by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3313-2.

### Step 3313-2

Have the disks or adapter card been physically moved recently?

- |            |  |
|------------|--|
| <b>NO</b>  | Contact your service support organization. |
| <b>YES</b> | Go to Step 3313-3.                         |

### Step 3313-3

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. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ipr` command to remove the device driver.
  2. Use the `modprobe ipr` command to reload the device driver.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 91.  
**Attention:** All data on the disk array will be lost.
- Format the remaining members of the disk array, as follows:  
**Attention:** All data on the disk array will be lost.
  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.

3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3320

Cache data associated with attached devices cannot be found (SRC cccc 9010).

### Step 3320-1

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

- NO**                      Go to Step 3320-3.  
**YES**                     Go to Step 3320-2.

### Step 3320-2

**Note:** Label all parts (*original* 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 29.
- The removable cache card from the *original* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

Go to step 3320-3.

### Step 3320-3

Has a new SRC cccc 9010 or SRC cccc 9050 occurred?

- NO**                      Go to Step 3320-6.  
**YES**                     Go to Step 3320-4.

### Step 3320-4

Was the new SRC cccc 9050?

- NO**                      The new SRC cccc 9010.  
                             Reclaim the controller cache storage. See “Reclaiming IOA Cache Storage” on page 40.  
**YES**                     Contact your Service Support organization

### Step 3320-5

If the system has been powered off for several days after an abnormal power-down, the Cache Battery Pack may be depleted. Do *not* replace the adapter nor the Cache Battery Pack.

Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 104.

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



## Step 3320-6

1. 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 29.
  - The removable cache card from the *new* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

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

## MAP 3321

IOA resources not available due to previous problems (SRC cccc 9054).

### Step 3321-1

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 3330

IOA does not support functions expected by devices (SRC cccc 9008).

### Step 3330-1

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3330-2.

### Step 3330-2

Have the adapter card or disks been physically moved recently?

**NO** Contact your service support organization.

**YES** Go to Step 3330-3.

### Step 3330-3

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. Then either IPL the system or logical partition (it may be necessary to boot into *Rescue* mode) or reload the device driver by performing the following:
  1. Use the `modprobe -r ipr` command to remove the device driver.
  2. Use the `modprobe ipr` command to reload the device driver.
- Format the remaining members of the disk array, as follows:

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

  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

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

## MAP 3331

Required cache data cannot be located for a disk unit (SRC cccc 9050).

### Step 3331-1

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

- NO** Go to Step 3331-3.  
**YES** Go to Step 3331-2.

### Step 3331-2

**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.

When continuing with the new adapter, perform the following:

**Note:** Label all parts (*original* 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 29.
- The removable cache card from the *original* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

Go to step 3331-8.

### Step 3331-3

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3331-4.

### Step 3331-4

Have the adapter card or disks been physically moved recently?

- NO** Contact your service support organization.  
**YES** Go to Step 3331-5.

### Step 3331-5

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

- NO** Go to Step 3331-6.  
**YES** Go to Step 3331-7.

### Step 3331-6

There are three possible ways to correct the problem. Perform only *one* of the following options that are listed by preference.

- Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 104.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 91.
- Format the remaining members of the disk array, as follows:

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

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

### Step 3331-7

The adapter and disks, identified in the preceding steps, 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 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 3331-8

Has a new SRC *cccc* 9010 or SRC *cccc* 9050 occurred?

**NO** Go to Step 3331-10.

**YES** Go to Step 3331-9.

### Step 3331-9

Was the new SRC *cccc* 9050?

**NO** The new SRC *cccc* 9010.

Reclaim the controller cache storage. See “Reclaiming IOA Cache Storage” on page 40.

**YES** Contact your Service Support organization

### Step 3331-10

1. 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 29.
- The removable cache card from the *new* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to “Separating a Removable Cache Card From the Base Card” on page 24.

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

## MAP 3332

IOA cache data exists for a missing or failed device (SRC *cccc* 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 3332-1

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3332-2.

### Step 3332-2

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

**NO** Go to Step 3332-4.

**YES** Go to Step 3332-3.

### Step 3332-3

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 3332-4

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

**NO** Go to Step 3332-5.

**YES** Go to Step 3332-6.

### Step 3332-5

Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 104.

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

### Step 3332-6

Have the adapter card or disks been physically moved recently?

**NO** Contact your service support organization.

**YES** Go to Step 3332-7.

### Step 3332-7

The adapter and disks, identified in the preceding steps, 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 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*.

## MAP 3333

Disk unit has been modified after the last known status (SRC cccc 9090) or Incorrect hardware configuration change has been detected (SRC cccc 9091).

### Step 3333-1

IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:

1. Use the `modprobe -r ipr` command to remove the device driver.
2. Use the `modprobe ipr` command to reload the device driver.

Take action on any new errors which surface.

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

## MAP 3334

Disk unit requires initialization before use (SRC cccc 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 **Device Concurrent Maintenance** screen in **iprconfig** when concurrently removing and installing disks (see “Replacing Physical Disks” on page 101) 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 3334-1

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3334-2.

### Step 3334-2

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

- |            |                    |
|------------|--------------------|
| <b>NO</b>  | Go to Step 3334-4. |
| <b>YES</b> | Go to Step 3334-3. |

### Step 3334-3

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 3334-4

Have the disks or adapter card been physically moved recently?

- NO Go to Step 3334-5.
- YES Go to Step 3334-6.

## Step 3334-5

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

- NO Go to Step 3334-7.
- YES Go to Step 3334-6.

## Step 3334-6

There are three possible ways to correct this problem. Perform only *one* of the following three options:

- IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ipr` command to remove the device driver.
  2. Use the `modprobe ipr` command to reload the device driver.
- Restore the adapter and disks back to their original configuration. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ipr` command to remove the device driver.
  2. Use the `modprobe ipr` command to reload the device driver.
- 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 3334-7

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

- Format the disks, as follows:

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

  1. Run the **iprconfig** utility by typing `iprconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.
- If the disks are members of a disk array, delete the disk array. See “Deleting a PCI-X SCSI RAID Disk Array” on page 91.

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

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

## MAP 3335

Disk media format bad (SRC cccc 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 3335-1

Identify the affected disk by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3335-2.

### Step 3335-2

Format the disks, as follows:

**Attention:** All data on the disks will be lost.

1. Run the **iprconfig** utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3337

IOA detected device error (SRC cccc 9001).

### Step 3337-1

A device configuration error has been detected. The configuration sectors on the device may be incompatible with the current I/O adapter.

Has the I/O adapter been replaced with a different type of I/O adapter, or have the devices been moved from a different type of I/O adapter to this one recently?

- NO**                      Go to Step 3337-2.  
**YES**                     Go to Step 3337-3.

### Step 3337-2

Contact your service support organization.

Exit this procedure.

### Step 3337-3

Identify the affected disk(s) by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3337-4.

### Step 3337-4

If the I/O Adapter has been replaced by a different type of I/O Adapter, reinstall the original adapter.

Go to Step 3337-5.

### Step 3337-5

If the disks involved have been moved from a different type of I/O Adapter to this one, return them to their original I/O Adapter.

Go to Step 3337-6.

## Step 3337-6

**Attention:** The next step will cause data loss on the disks involved. If the data on these disks is important, perform back up procedures now.

Stop parity protection on the disk array(s) involved. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 91.

Go to Step 3337-7.

## Step 3337-7

Restore the initial hardware configuration that caused the error.

Go to Step 3337-8.

## Step 3337-8

Start parity protection on the disks then restore the saved data. For further details, see “Creating a PCI-X SCSI RAID Disk Array” on page 90.

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

## MAP 3350

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 Card Type 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.

**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 an 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 *Unprotected* or fail if functioning disks are removed. Also, additional problems may be created.

## Step 3350-1

Identify the SCSI bus on which the problem is occurring by examining the error log entry.

Go to Step 3350-2.



## Step 3350-2

Have changes been made to the SCSI configuration recently?

- NO** Go to Step 3350-5.  
**YES** Go to Step 3350-3.

## Step 3350-3

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)
- Ensure the SCSI bus does not have multi-initiators (for example, set up for a high-availability configuration)

**Note:** Multi-initiator support is not provided at this time.

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 3350-5.  
**YES** Go to Step 3350-4.

## Step 3350-4

1. Power off the system or logical partition.
2. Correct the problem.
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus-related failure occur?

- NO** Go to Step 3350-14.  
**YES** Go to Step 3350-5.

## Step 3350-5

Determine if any of the disk arrays on the adapter are in an *Unprotected* state. For further details, see "Viewing Parity Status" on page 88.

Does any disk array have a state of *Unprotected*?

- NO** Go to Step 3350-7.  
**YES** Go to Step 3350-6.

## Step 3350-6

1. Identify the failed disk(s) by first finding disk arrays with a state of *Unprotected*, then finding disks on those arrays with a state of *Failed*.
2. Remove the failed disk(s) from each *Unprotected* disk array. For further details, see "Replacing Physical Disks" on page 101.

3. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO** Go to Step 3350-14.

**YES** Go to Step 3350-7.

### Step 3350-7

Are there any non-essential removable media devices (such as Tape, CDROM, and DVDROM) on the SCSI bus?

**NO** Go to Step 3350-10.

**YES** Go to Step 3350-8.

### Step 3350-8

1. Power off the system or logical partition.
2. Remove one of the non-essential removable media devices.
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO** Go to Step 3350-9.

**YES** Go to Step 3350-7.

### Step 3350-9

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 3350-14.

### Step 3350-10

Are there any non-essential disks that are not disk array members (such as 512 byte/sector *JBOD* disks, *Hot Spare* disks, or array candidates) on the SCSI bus?

**NO** Go to Step 3350-13.

**YES** Go to Step 3350-11.

### Step 3350-11

1. Remove one of the non-essential disk devices. For further details, see “Replacing Physical Disks” on page 101.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO** Go to Step 3350-12.

**YES** Go to Step 3350-10.

## Step 3350-12

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 3350-14.

## Step 3350-13

Contact your service support organization.

Exit this procedure

## Step 3350-14

1. Reinstall any good devices that were removed during the course of this MAP.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.
3. Take action on any other non-SCSI bus related errors if they exist.

## MAP 3351

Use the following to determine what other Field Replaceable Units (FRUs) besides a disk, may need to be replaced in order to solve a problem.

You are sent here when an *Advanced Function* disk, a physical disk in 522 bytes/sector format, was identified as the primary FRU to replace in order to resolve 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.
- 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 Card Type 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 fail or become *Unprotected* if functioning disks are removed. Also, additional problems may be created.

## Step 3351-1

Identify the SCSI bus on which the problem is occurring by examining the error log entry.

Go to Step 3351-2.

## Step 3351-2

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 or logical partition
2. Replace a component listed above
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

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

## MAP 3390

The problem that occurred is uncommon or complex to resolve. Information should be gathered and assistance obtained from your service support organization.

### Step 3390-1

Make a copy of the error log entry, as well as other error log entries that occurred around the same time. These entries are all in **/var/log/messages** file if the default error log path has not been changed.

Go to Step 3390-2.

### Step 3390-2

Collect the current disk array configuration. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 85.

Go to Step 3390-3.

### Step 3390-3

Contact your service support organization.

Exit this procedure.

---

## Part 3. SuSE Linux Enterprise Server 8



---

## Chapter 11. PCI-X SCSI RAID Controller Overview

The PCI-X SCSI RAID Controller is a 133 MHz 64-bit PCI-X adapter. This adapter has the following features:

- PCI-X system interface
- Dual Channel, Ultra320 (320 MB/s) SCSI controllers
- Embedded PowerPC RISC Processor and Hardware XOR DMA Engine
- Non-volatile write cache
- Support for RAID (Redundant Array of Independent Disks) levels 0, 5, 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 and 10 disk arrays
  - Ability to increase the capacity of an existing RAID level 5 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

**Note:** Multi-initiator support is not provided at this time.

Card Type	PCI Form Factor	Write Cache Size	Adapter FFC	Rechargeable Battery Technology	Cache Battery Pack FFC
5702, 1974 <sup>1</sup>	Yes	0 MB	2522	None	Not Applicable
5703, 1975	Yes	40 MB	2523	NiMH	2526
5709, 1976	No <sup>2</sup>	16 MB	2523	NiMH	2526
<sup>1</sup> Card type 5702 is not RAID capable.					
<sup>2</sup> 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. The Card Type describes the RAID Enablement Card.					

---

### 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 Linux in the same way as it would a single SCSI disk. For example, after creating a disk array, you can use Linux commands to make the disk array available to the system by partitioning and creating a file system on it.

The PCI-X SCSI RAID Controller is managed by the **sisconfig** utility. The **sisconfig** utility is the interface to the RAID configuration, 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 into Rescue mode and creating the disk array before installing Linux. 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.



## 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 disks 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, 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 17. 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 *Unprotected* 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:




Figure 18. 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 *Unprotected* mode. Whenever data is read from an *Unprotected* 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 *R/W Protected* 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:

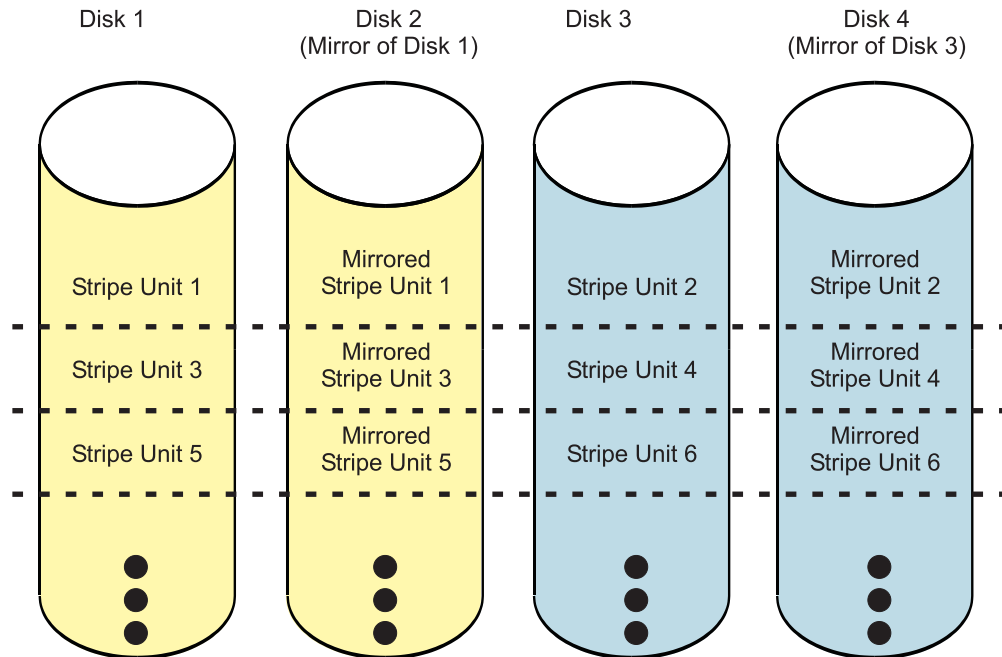


Figure 19. RAID Level 10

RAID level 10 can tolerate multiple disk failures. If one disk in each mirrored pair fails, the array will still be functional, operating in *Unprotected* 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 *R/W Protected* 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 *Unprotected* 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 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 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 10

Stores data redundantly on mirrored pairs to provide maximum protection against disk failures. Provides generally better performance than RAID level 5, 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 RAID Disk Array to 16 KB, 64 KB, or 256 KB. You may be able to maximize the performance of your PCI-X SCSI RAID 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 and 10 disk arrays also provide data redundancy, so that no data is lost if a single disk in the array fails.

**Note:** This guide and the **sisconfig** utility use common terminology for disk formats:

*JBOD* A *JBOD* disk is a disk that is formatted to 512 bytes/sector. *JBOD* stands for "Just a Bunch Of Disks".

*Advanced Function*

An *Advanced Function* disk is a disk that is formatted to 522 bytes/sector. This format allows disks to be used in disk arrays.

PCI-X SCSI RAID Disk Arrays are accessed in Linux as standard SCSI disk devices. These devices are automatically created when a disk array is created, and deleted whenever a disk array is deleted. The individual physical disks that comprise disk arrays or are candidates to be used in disk arrays, which are formatted for *Advanced Function*, are hidden from Linux and are accessible only through the **sisconfig** utility. Linux sees all *JBOD* disks. These disks must be formatted for *Advanced Function* before they can be used in disk arrays. For information on formatting *JBOD* disks to make them available for use in PCI-X SCSI RAID Disk Arrays, see "Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays" on page 151. The **Display Disk Hardware Status** option in the **sisconfig** utility can be used to display these disks and the resource names associated with them. For details regarding how to view the disk information, see "Viewing the Status of Disks and Disk Arrays" on page 147. The following sample output is displayed when the **Display Disk Hardware Status** option is invoked.

Display Disk Hardware Status					
Serial Number	Vendor ID	Product ID	Model	Resource Name	Hardware Status
03000001	IBM	5703001	001	/dev/ibmsis0	Operational
00009DB1	IBM	ST3146807LC	020	/dev/sdd	Operational
000012F0	IBM	ST3146807LC	090		HS/Active
00001315	IBM	ST3146807LC	079		DPY/Active
000012F5	IBM	ST3146807LC	070		DPY/Active
00001304	IBM	ST3146807LC	079		DPY/Active
E7800377	IBM	01C933F1	205	/dev/sdc	Operational

Press Enter to continue.

e=Exit                      q=Cancel                      r=Refresh    f=PageDn    b=PageUp  
d=Display disk unit details    p=Display device parity status

Disk array and physical disk model numbers are displayed in the fourth column of the **Display Disk Hardware Status** screen. The following information provides a description for the possible model numbers that may appear.

## Disk Array Models

The fourth column in the preceding output displays the model number of the disk array. The model number indicates the parity configuration.

200    RAID level 0

205    RAID level 5

210    RAID level 10

## Physical Disk Models

The fourth column in the preceding output displays the model number of the device. The model number indicates the device configuration.

- 020 JBOD disk unit resource. Linux sees this disk.
- 050 Advanced Function disk unit resource. This disk is not currently configured in a disk array.
- 07x Parity protected disk unit. The x describes the amount of the device that is used as parity.
  - 9 unknown parity
  - 4 1/4 parity
  - 2 1/8 parity
  - 1 1/16 parity
  - 0 no parity
- 090 Hot Spare disk unit.

Disk array, physical disk, and I/O adapter states are displayed in the sixth column of the **Display Disk Hardware Status** screen. The following information provides a description for the possible states that may appear.

## Disk Arrays States

The sixth column in the preceding output displays the state of the disk array. The five valid states for PCI-X SCSI RAID Disk Arrays are: *Operational*, *DPY/Unprotected*, *DPY/Rebuilding*, *R/W Protected*, and *Perf Degraded*.

### *Operational*

The disk array is functional and fully protected (RAID level 5 and 10) with all physical disks in the *DPY/Active* state.

### *DPY/Unprotected*

The disk array is no longer protected by data redundancy due to a disk failure.

### *DPY/Rebuilding*

Data protection is being rebuilt on this disk array.

### *R/W Protected*

The disk array cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

### *Perf Degraded*

The device is functional, but performance may be affected because of other hardware problems (such as a problem with the IOA cache.)

## Physical Disk States

The sixth column in the preceding output displays the state of the device. The 11 possible states for physical disks are: *Operational*, *Not Operational*, *Not Ready*, *R/W Protected*, *DPY/Failed*, *DPY/Active*, *Perf Degraded*, *DPY/Unprotected*, *DPY/Rebuilding*, *Format Required*, and *HS/Active*.

### *Operational*

The disk unit is operational and ready to accept input or output operations.

### *Not Operational*

This device cannot communicate with the IOA. Verify that the unit is powered on.

### *Not Ready*

The device cannot perform media-related functions, but it can still communicate with the IOA.

#### *R/W Protected*

The device cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

#### *DPY/Failed*

This disk unit is part of a disk array that has device parity protection. The disk unit failed within its device parity set which caused the loss of data protection for the device parity set.

#### *DPY/Active*

This disk unit is part of a disk array that has device parity protection. The disk unit is operational and ready to accept input or output operations.

#### *Perf Degraded*

The device is functional, but performance may be affected because of other hardware problems (such as a problem with the IOA cache.)

#### *DPY/Unprotected*

This disk unit is part of a disk array that has device parity protection. Data protection is no longer in effect because of a failure of another resource.

#### *DPY/Rebuilding*

This disk unit is part of a disk array that has device parity protection. Data protection is being rebuilt.

#### *Format Required*

The disk unit must be formatted to become usable on this IOA.

#### *HS/Active*

The disk unit is configured to be a *Hot Spare*.

## **I/O Adapter States**

The sixth column in the preceding output displays the state of the I/O Adapter. The three possible states for I/O Adapters are: *Operational*, *Not Operational*, and *Not Ready*.

#### *Operational*

The IOA is functional.

#### *Not Operational*

The device driver cannot successfully communicate with this IOA.

#### *Not Ready*

The IOA requires a microcode download.





---

## Chapter 12. PCI-X SCSI RAID Controller Software

For the PCI-X SCSI RAID Controller to be identified and configured by Linux, the requisite device support software must be installed. Software for the PCI-X SCSI RAID Controller consists of a device driver and a set of utilities. The requisite software for the PCI-X SCSI RAID Controller is often preinstalled as part of the normal Linux installation. However, if the software package is not installed, software verification will fail. The missing package(s) can be installed from your Linux operating system CD-ROM. If you are missing components or need newer versions, obtain them from your Linux distributor or online at: <http://www6.software.ibm.com/dl/linuxsdd/linuxsdd-p>

**Note:** Support of the PCI-X SCSI RAID Controller requires Version 1.19 or later of both the **ibmsis** device driver and the **sisconfig** utility.

The PCI-X SCSI RAID Controller executes onboard microcode. Although a version of controller microcode may be distributed along with Linux, this does not necessarily represent the most recent version of microcode available for the controller. The **sisupdate** utility can be used to update the microcode being used by the controller. For more information regarding **sisupdate**, see “Updating the PCI-X SCSI RAID Controller Microcode” on page 158.

---

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

To verify that the **ibmsis** device driver for the PCI-X SCSI RAID Controller is installed, type:

```
modinfo ibmsis
```

Output from this command will indicate that the module is not found if the supporting software has not been installed. If information about **ibmsis** is displayed, type the following command to obtain version information:

```
strings filename | grep "ibmsis version"
```

where *filename* is the file name displayed after the **modinfo** command was typed in the previous step.

To verify that the **sisconfig** utility is installed, type the command:

```
sisconfig --version
```

Output from this command will indicate if the utility is installed, and if so, version information will be displayed.

**Attention:** To ensure proper functionality, the version numbers of **ibmsis** and **sisconfig** should be the same. Both the **ibmsis** device driver and the **sisconfig** utility should be Version 1.19 or later.

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 Linux distribution.



---

## Chapter 13. 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 **sisconfig**

The interface for working with the PCI-X SCSI RAID Controller is **sisconfig**.

To start the **sisconfig** utility, type the command:

```
sisconfig
```

The main menu will display options for configuring and managing the PCI-X SCSI RAID Controller.

```
Work with Disk Units

Select one of the following:

    1. Display disk hardware status
    2. Work with device parity protection
    3. Work with disk unit recovery
    4. Work with configuration
    5. Analyze Log

Selection:

e=Exit
```

---

### Viewing the Status of Disks and Disk Arrays

The **sisconfig** utility offers an easy way to view the status of all devices controlled by the **ibmsis** device driver.

To view information about the disks and disk arrays on your system, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select the **Display disk hardware status** option. The output displayed will be similar to the following:

### Display Disk Hardware Status

Serial Number	Vendor ID	Product ID	Model	Resource Name	Hardware Status
03000001	IBM	5703001	001	/dev/ibmsis0	Operational
00009DB1	IBM	ST3146807LC	020	/dev/sdd	Operational
000012F0	IBM	ST3146807LC	090		HS/Active
00001315	IBM	ST3146807LC	079		DPY/Active
000012F5	IBM	ST3146807LC	070		DPY/Active
00001304	IBM	ST3146807LC	079		DPY/Active
E7800377	IBM	01C933F1	205	/dev/sdc	Operational

Press Enter to continue.

e=Exit                      q=Cancel                      r=Refresh    f=PageDn    b=PageUp  
d=Display disk unit details      p=Display device parity status

- The first column of output is the device's 8-digit hexadecimal serial number.
- The second column of output is the device's vendor ID.
- The third column of output is the device's product ID. Different types of devices may have very dissimilar product IDs.
- The fourth column of output is the device's model. For an overview of the possible disk and disk array models, see "Disk Array Overview" on page 140.
- The fifth column of output is the device's Linux resource name. Devices that are hidden from Linux are not assigned resource names, and are only visible through **sisconfig**.
- The sixth column of output is the device's hardware status. For an overview of the possible disk and disk array hardware statuses, see "Disk Array Overview" on page 140.

For more information on the physical devices shown, type d to display disk unit details. Output similar to the following will be displayed:

### Display Disk Unit Details

Type option, press Enter.

5=Display hardware resource information details

Serial Number	Location	PCI Bus	PCI Dev	SCSI Host	SCSI Bus	SCSI ID	SCSI Lun
03000001	P1-I2	64	12	2			
000012F5	P1-I2/Z1-A4	64	12	2	0	4	0
000012F0	P1-I2/Z1-A5	64	12	2	0	5	0
00001315	P1-I2/Z1-A6	64	12	2	0	6	0
00001304	P1-I2/Z1-A9	64	12	2	0	9	0

e=Exit                      q=Cancel                      f=PageDn    b=PageUp

The previous screen shows information pertaining to the location of the hardware on your system. To view information on a specific device, select the desired device with a 5 and press Enter. Output similar to the following will be displayed:

```
Disk Unit Hardware Resource Information Details

Manufacturer. . . . . : IBM
Machine Type and Model. . . : ST3146807LC
Firmware Version. . . . . : 43353046
Model . . . . . : 079
Serial Number . . . . . : 00001315
Capacity. . . . . : 142.87 GB

Physical location:
Location . . . . . : P1-I2/Z1-A6
PCI Bus. . . . . : 64
PCI Device . . . . . : 12
SCSI Host Number . . . . . : 2
SCSI Channel . . . . . : 0
SCSI Id. . . . . : 6
SCSI Lun . . . . . : 0

Press Enter to continue.

e=Exit      q=Cancel      r=Remove      v=Display Extended VPD
```

The previous screen shows an overview of a particular piece of hardware on your system. If extended information is available, press v to view it. Output similar to the following will be displayed:

```
Disk Unit Hardware Resource Extended Details

Manufacturer. . . . . : IBM
Machine Type and Model. . . : ST3146807LC
FRU Number. . . . . : 55P4081
Serial Number . . . . . : 000012F0
EC Level . . . . . : H31325
Part Number. . . . . : 55P4083
Device Specific (Z0). . . . : 000003129F00013E
Device Specific (Z1). . . . : 0626C50F
Device Specific (Z2). . . . : 0002
Device Specific (Z3). . . . : 02291
Device Specific (Z4). . . . : 0001
Device Specific (Z5). . . . : 22
Device Specific (Z6). . . . : H31325

Press Enter to continue.

e=Exit      q=Cancel
```

---

## Viewing Parity Status

The parity status can be viewed in two ways.

- To view the parity status from the main menu, do the following:
  1. Run the **sisconfig** utility by typing `sisconfig`.
  2. Select **Work with device parity protection**.
  3. Select **Display device parity status**. The **Display Device Parity** screen will appear. Sample output follows.
- To view the parity status from the **Display Disk Hardware Status** screen, type `p`. The **Display Device Parity** screen will appear.

The following is sample output that will display:

Display Device Parity Status										
Parity Set	Serial Number	Model	Resource Name	PCI Bus	PCI Dev	SCSI Bus	SCSI ID	SCSI Lun	Status	
1	FF3BA7E4	210	/dev/sdc	64	12	255	0	0	Active	
	00249E55	070		64	12	1	0	0	Active	
	00C02D9E	070		64	12	1	1	0	Active	
2	DCC7BB45	200	/dev/sdd	64	12	255	0	1	Active	
	001DA3BA	070		64	12	1	2	0	Active	

e=Exit    q=Cancel    r=Refresh    f=PageDn    b=PageUp

## Disk Array Parity Status

The tenth column in the preceding output displays the parity status of each disk array. The five valid parity states for PCI-X SCSI RAID Disk Arrays are: *Active*, *R/W Prot*, *Synched*, *Rebuilt*, and *Unprot*.

**Active** This disk array is functioning properly.

### *R/W Prot*

The disk array cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

### *Synched*

The parity for this disk array is being built. Accompanying this state is a % complete.

### *Rebuilt*

The data on this disk array is being rebuilt from other units in the disk array. Accompanying this state is a % complete.

**Unprot** Due to one or more disk failures, this disk array is no longer fully protected by parity.

## Disk Parity Status

The tenth column in the preceding output displays the parity status of each disk. The seven valid parity states for disks are: *Active*, *Failed*, *Not Ready*, *R/W Prot*, *Synched*, *Rebuilt*, and *Unprot*.

*Active* This disk unit is part of a disk array and is fully operational.

*Failed* This disk unit is part of a disk array and has failed.

*Not Ready*

The device cannot perform media-related functions, but it can still communicate with the IOA.

*R/W Prot*

The device cannot process a read nor write operation. A disk array may be in this state because of a cache, device configuration, or any other problem that could cause a data integrity exposure.

*Synched*

The redundant data for the disk array is being built from the data within the disk array. Accompanying this state is a % complete.

*Rebuilt*

This disk unit is part of a disk array that has device parity protection. The data on this disk unit is being rebuilt from other units in the disk array. Accompanying this state is a % complete.

*Unprot* This disk unit is part of a disk array that has device parity protection. This disk unit is operational. However, another disk unit in the disk array has failed or is being rebuilt. If another disk unit in the disk array fails, data could be lost.

---

## Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays

Before a disk can be used in a PCI-X SCSI RAID Disk Array, it must be formatted for *Advanced Function*. Before a disk is recognized as a standalone disk, it must be formatted to *JBOD*. Steps for performing both of these actions are contained in this section.

### Formatting to Advanced Function

To format a disk for *Advanced Function*, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Format device for advanced function**.
5. From the list of eligible disk units, choose the disks you want to format for *Advanced Function* and press Enter.

**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 for *Advanced Function*. These disks will not show up on the list of choices. In some cases, errors may be logged in the `/var/log/messages` file. For more detailed information, view that log file.

6. A message will display asking if you want to continue. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), resource names (such as `/dev/sdb`) may change when the system is rebooted. This could affect kernel command line entries and `fstab` entries.

After the formatting is complete, the disks will be ready for use in PCI-X SCSI RAID Disk Arrays.

## Formatting to JBOD

To format a disk back to *JBOD* format and allow Linux to use it as a standalone disk, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Format device for JBOD function**.
5. From the list of eligible disk units, choose the disks you want to format to *JBOD* and press Enter.  
**Attention:** Continuing with this option will format the disks. All data on the disks will be lost.
6. A message will display asking if you want to continue. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

**Note:** If you are not using a device mapper (for example, LVM, md, or `scsidev`), it is advised that you reboot the system now. Resource names (such as `/dev/sdb`) may change as a result of reformatting. This could affect kernel command line entries and `fstab` entries.

---

## Creating a PCI-X SCSI RAID Disk Array

A disk array is created using a set of disks that are formatted for *Advanced Function*. Disk arrays with data redundancy (RAID level 5 and RAID level 10) will be unprotected until parity is finished being built. For disk arrays with no data redundancy (RAID 0), no parity is calculated and the array is built immediately. To create a PCI-X SCSI RAID Disk Array, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with device parity protection**.
3. Select **Start device parity protection**.
4. Select the controller under which you would like to start device parity protection then press Enter.
5. Select the disk units to be included in the disk array then press Enter. A screen similar to the following will appear:

```
Select Protection Level and Stripe Size

Default array configurations are shown. To change
setting hit "c" for options menu. Highlight desired
option then hit Enter

c=Change Setting

Protection Level . . . . . : RAID 5
Stripe Size . . . . . : 64 k

Press Enter to Continue

e=Exit   q=Cancel
```

6. Select the protection level desired for the array. For more information about selecting an appropriate RAID level, see “Supported RAID Levels” on page 136.
7. Select the stripe size desired in kilobytes for the array. For more information about selecting an appropriate stripe size, see “Stripe-Unit Size” on page 140.



8. Press Enter to continue. If you wish to create additional disk arrays, type a and go back to Step 5 on page 152. Otherwise, type c to continue.

**Attention:** All data on the selected drives will be lost when the disk array is created. If you are sure you want to create the disk array, press Enter. If you do not want to create the disk array, type q to cancel.

If you choose to create the disk array, a status screen displays until the operation is completed. When a disk array has been built, it will be available to Linux as would any normal SCSI disk. To find the resource name that has been assigned to the disk array, refer to the **Display Device Parity Status** screen.

**Note:** If you are not using a device mapper (for example, LVM, md, or scsidev), it is advised that you reboot the system now. Resource names (such as /dev/sdb) may change as a result of creating a new disk array. This could affect kernel command line entries and fstab entries.

---

## Deleting a PCI-X SCSI RAID Disk Array

**Attention:** After a disk array is deleted, it cannot be accessed. All data will be lost.

**Note:** A disk array that is currently in use or opened (that is busy) cannot be deleted. Also, if a disk array is currently being rebuilt or synched, that disk array cannot be deleted.

To preserve the data on the disk array, you must first back up all data that you wish to save. To delete the array, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with device parity protection**.
3. Select **Stop device parity protection**.
4. From the listed disk arrays, select the one you wish to delete then press Enter.

**Attention:** All data on the selected drives will be lost when the disk array is deleted. If you are sure you want to delete the disk array, press Enter. If you do not want to delete the disk array, type q to cancel.

**Note:** If you are not using a device mapper (for example, LVM, md, or scsidev), resource names (such as /dev/sdb) may change when the system is rebooted. This could affect kernel command line entries and fstab entries.

---

## Adding Disks to an Existing Disk Array

The PCI-X SCSI RAID Controller supports adding disks to existing RAID level 5 disk arrays. This feature can be used to dynamically increase the capacity of a disk array while preserving existing data in the disk array. Extensive use of this feature, however, will result in a performance penalty because the data will not be restriped.

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

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with device parity protection**.
3. Select **Include unit in device parity protection**.
4. Select the PCI-X SCSI RAID Disk Array to which you want to add the disk(s).
5. Select the disk(s) to be included into the disk array then press Enter.

**Attention:** All data on the selected drives will be lost when the disk(s) are added to the existing array.

6. If you are sure you want to include the disk(s), press Enter. If you do not want to include the disk(s) in the disk array, type q to cancel.

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 one or more 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 for *Advanced Function*.
- The disk belongs to another disk array or is configured as a *Hot Spare*.

---

## Using Hot Spare Disks

*Hot Spare* disks are used to automatically replace failed disks in a 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 *Unprotected*. In order to assign a disk for use as a hot spare, it must be formatted for *Advanced Function*.

### Creating Hot Spare Disks

To create *Hot Spare* disks, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Configure a hot spare device**.
5. Select the adapter for which you want to create hot spares and press Enter.
6. Select the disks that you want to designate as hot spares and press Enter.  
**Attention:** All data on the selected drives will be lost when the disks are configured as hot spares.
7. If you are sure you want to configure the disks as hot spares, press Enter. If you do not want to configure the disks as hot spares, type `q` to cancel.

### Deleting Hot Spare Disks

To delete *Hot Spare* disks and free them for use in other disk arrays, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with configuration**.
3. Select **Work with disk unit configuration**.
4. Select **Unconfigure a hot spare device**.
5. Select the adapter from which you wish to remove a hot spare and press Enter.
6. Select the *Hot Spare* disks to be deleted and press Enter.
7. If you are sure you want to reconfigure the disks, press Enter. If you do not want to reconfigure the disks, type `q` to cancel

---

## Viewing and Changing PCI-X SCSI RAID Controller Bus Settings

**Note:** The maximum speed and bus width are automatically negotiated based on the characteristics of the attached devices. If you are not successful running at these speeds, you can reduce the adapter speed by changing the **Maximum bus throughput**.

**Note:** The driver will default to a maximum of 80 MB/s until the **sisupdate** utility runs successfully. After **sisupdate** completes, maximum configured bus speed is enabled. Each time the system is booted, **sisupdate** is run automatically.

To view and change settings for the PCI-X SCSI RAID Controller, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.

2. Select **Work with configuration**.
3. Select **Work with SCSI bus configuration**.
4. Select the adapter that you want to change settings on or view by typing 1 then pressing Enter. The **Change SCSI Bus Configuration** screen similar to the following will appear:

```
Change SCSI Bus Configuration

Current Bus configurations are shown. To change
setting hit "c" for options menu. Highlight desired
option then hit Enter

c=Change Setting

/dev/ibmsis0
BUS 0
  Maximum Bus Throughput . . : 320 MB/s
BUS 1
  Wide Enabled . . . . . : Yes
  Maximum Bus Throughput . . : 320 MB/s

Press Enter to Continue
e=Exit   q=Cancel
```

5. Follow the directions on screen to change settings, if desired.

The preceding screen displays information about the controller bus. The following is information regarding attributes displayed:

#### **Maximum Bus Throughput**

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 options available will differ depending on your hardware configuration.

#### **Wide 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.

**Note:** If a device that requires 16-bit data transfer is attached, you will not have the option to disable wide transfer mode.



---

## Chapter 14. PCI-X SCSI RAID Controller Maintenance

This chapter provides information to assist in ensuring optimal performance of your PCI-X SCSI RAID Controller under SuSE Linux Enterprise Server 8.

---

### 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 physical disks. Using the **modprobe -r ibmsis** command will have the same effect as a system shutdown.

**Note:** A disk that is a failed member of an *Unprotected* disk array can be replaced and the disk array rebuilt while the system continues to run.

- You can physically move disks from one adapter to another. However, if the physical disks 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 an *Unprotected* state because of a disk failure.
- When physically removing disks 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 these disks are used.
- Always use the **Device Concurrent Maintenance** screen to remove and replace a physical disk. For instructions on how to remove and replace a disk, see "Replacing Physical Disks" on page 168.
- 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 into *Rescue* mode. Linux error logs, the **sisconfig** utility, and other tools are available 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. This is likely to make the problems worse. For additional information regarding problem determination, see Chapter 15, "Problem Determination and Recovery," on page 173.
- 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.
- If multiple errors occur at approximately the same time, look at them as a whole to determine if there may be a common cause.

---

## Updating the PCI-X SCSI RAID Controller Microcode

Disk drives attached to this RAID adapter must contain the latest microcode level.

The latest microcode level for your drive model and PCI-X SCSI RAID Controller is shipped as part of the Linux distribution. If an update or replacement is needed, the latest microcode can be obtained from your Linux distributor, or found as part of the device driver package at the following Web site:

<http://www6.software.ibm.com/dl/linuxsdd/linuxsdd-p>.

Microcode can be updated using the **sisupdate** utility. If **sisupdate** finds newer microcode in `/etc/microcode`, it will be downloaded to the adapter. The **sisupdate** utility should be set to run during boot time, after the **ibmsis** device driver is loaded.

If you downloaded new microcode from the Web site, install it as follows:

1. Untar the package, type:

```
tar -xzvf ibmsis-$ver.tgz
```

Press Enter.

2. Change directories to **ibmsis/obj/firmware/**. Type:

```
cd ibmsis/obj/firmware/
```

3. Install the microcode. Type:

```
make install
```

Press Enter.

4. Update the microcode, type:

```
sisupdate
```

Press Enter.

---

## Separating a Removable Cache Card From the Base Card


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

**Attention:** Do not disconnect the cache battery during this procedure or all cache data will be lost.

All cards are sensitive to electrostatic discharge. See Working with electrostatic discharge-sensitive parts.


**Note:** You will need a T-10 TORX tool to complete this procedure.

1. To separate the removable cache card from the base card, do the following:
  - a. Label both sides of the card before separating.
  - b. Place the adapter on an ESD protective surface and orientate it as shown in step 1c. Note the position of the card handle.
  - c. Remove three screws **A** from the base card as shown. Note the thread differences of the screws for correct replacement in the following steps.




IPHAU600-1

- d. Turn the card over.




IPHAU601-1

- e. Remove four screws **B** from the removable cache card as shown.




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




- g. Separate the base card from the removable cache card.



**Important:** Turn the jack screw back down fully clockwise.




- h. Place the adapter on an ESD protective surface and orientate it as shown in step 1c on page 158. Note the position of the card handle.
2. To attach the removable cache card to the base card, do the following:
  - a. 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.




- b. Attach the base card to the removable cache card by doing the following:


- 1) Line up the card connector (view the jack screw through the hole to help align).
- 2) Press on the stiffener plate to fully seat the connector.
- c. Attach four screws **A** to the removable cache card as shown.



- d. Turn the card over.



- e. Attach the three screws **B** to the base card as shown. Note the thread differences in the screws as pointed out in step 1c on page 158 of the separation procedure.



## Replacing the Cache Directory Card

**Attention:** All cards are sensitive to electrostatic discharge. See “Handling Static Sensitive Devices” on page xi. This procedure should only be performed if directed from an isolation procedure or a maintenance analysis procedure (MAP).

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 Figure 20 on page 164, Figure 21 on page 164, Figure 22 on page 164 or Figure 23 on page 165 below).
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.




Figure 20.




Figure 21.




Figure 22.




Figure 23.

## Rechargeable Battery Maintenance

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. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Work with resources containing cache battery packs**.
4. Select your adapter and type 5. Then press Enter to display battery information.
5. The **Battery Information** screen displayed will look similar to the following example:

```

Battery Information

Resource Name . . . . . : /dev/ibmsis0
Serial Number . . . . . : 03000001
Type-model . . . . . : 5703-001
Location . . . . . : P1-I2
PCI Bus . . . . . : 64
PCI Device . . . . . : 12
SCSI Host Number . . . . . : 2
Battery type . . . . . : Nickel Metal Hydride (NiMH)
Battery state . . . . . : No battery warning
Power on time (days) . . . . . : 56
Adjusted power on time (days) . . . : 61
Estimated time to warning (days) . : 865
Estimated time to error (days) . . : 942

```

Press Enter to continue.

e=Exit            q=Cancel

The following information details each battery state that may be displayed on the **Battery Information** screen:

#### **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**

No warning 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**

The battery state cannot be determined.

## **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. Performing this task ensures that write caching is stopped prior to replacing a Cache Battery Pack. Forcing the battery error can be performed when you want to replace a Cache Battery Pack on one controller and you also want to replace other Cache Battery Packs on controllers in the same system that have not yet issued the cache battery warning nor cache battery error.

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, you will have to replace the Cache Battery Pack.

To force the battery error, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Work with resources containing cache battery packs**.
4. Select your adapter and type 2. Then press Enter to force the battery error.

**Attention:** This procedure should be done only under the guidance of IBM Hardware Service Support and poses serious risks to performance until the battery pack is replaced.

5. If you are sure you want to force a battery error, type `c` to confirm. If you do not want to force a battery error, type `q` to cancel.

## **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 15 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 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.

If the Cache Battery Pack is not already in the error state, follow the steps described in “Forcing a Rechargeable Battery Error” on page 33 before proceeding. After the battery is in the error state, replace it with a new battery by taking the following steps:


1. Remove the adapter from the system by referring to your system documentation.
2. Use the following illustration to locate the Battery Pack Casing that holds the battery pack.
3. Squeeze the Battery Pack Casing to remove the battery unit from the card.
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 phase.

5. Remove the battery unit from the Battery Pack Casing. Save the Battery Pack Casing. The replacement battery pack does not provide a casing.

**Note:** Ensure that the Cache Battery Pack is disconnected for at least 15 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.



**A** Battery Pack Casing  
**B** Cache Directory Card

---

## Replacing Physical Disks

Replace failed disks as soon as possible, even if a rebuild was initiated on a *Hot Spare* by the controller. Always use the **Device Concurrent Maintenance** option in the **sisconfig** utility to replace physical disks attached to a PCI-X SCSI RAID Controller.

### Device Concurrent Maintenance

Type the choices, then press Enter.

Specify either Physical Location or PCI/SCSI location.

Physical Location . . . :

OR

PCI/SCSI . . . PCI Bus: PCI Device:

SCSI Channel: SCSI Id: SCSI Lun:

Specify action as 1=Remove device 2=Install device

Action to be performed . . . . . :

Enter a time value between 01 and 19.

Time needed in minutes . . . . . :

Or leave blank and press Enter to cancel



## Removing a Failed Disk

To remove a failed disk, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Device Concurrent Maintenance**.
4. Enter either the physical location code or the PCI/SCSI location of the device to be removed.
5. Specify **Remove device** as the action to be performed.
6. Enter a time value in minutes. This should be enough time for you to remove the specified hardware. During this time interval, **sisconfig** will wait and the LED(s) at the specified physical location will indicate that the device is ready for removal.
7. Remove the physical disk from the system.

**Attention:** Do not install the replacement disk at this time.

## Installing a New Disk

Once the failed disk has been removed, do the following to install a new disk:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Device Concurrent Maintenance**.
4. Enter either the physical location code or the PCI/SCSI location where the new device will be installed.
5. Specify **Install device** as the action to be performed.
6. Enter a time value in minutes. This should be enough time for you to install the specified hardware. During this time interval, **sisconfig** will wait and the LED(s) at the specified physical location will indicate that the new device can be installed.
7. Install the physical disk into the system.

If a *Hot Spare* disk was available and kept your disk array protected throughout the failure, no additional steps are required because the hot spare has replaced the failed disk unit in the disk array, and the newly installed disk is available to be configured as a hot spare. If a *Hot Spare* rebuild was not initiated by the controller, a rebuild needs to be initiated on the newly installed disk. To initiate a rebuild, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk you wish to rebuild then press Enter.  
**Attention:** Data currently on the disk will be overwritten.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

---

## 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 and physical disk states, see “Disk Array Overview” on page 140. For a description of RAID levels, see “Supported RAID Levels” on page 136. This section describes recovery procedures for each RAID level.

### RAID Level 0

A single disk failure will cause a RAID level 0 array to transition to the *R/W Protected* state. You will need to delete the disk array, replace the disk that is *Not Operational*, and recreate the disk array. You must then recreate the file systems on the disk array, and copy 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 fails, the disk array status changes to *Unprotected*. The disk array remains functional because the data on the failed disk can be rebuilt using parity and data on the remaining disks.

Whenever a disk fails, replace it as soon as possible. If a *Hot Spare* disk was available, the controller may rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

### Multiple-Disk Failures

If a second disk in a RAID level 5 disk array fails, the array will also fail and the data will not be accessible. You will have to replace the failed disks, then delete and recreate the disk array. You must then recreate the file systems on the disk array and copy data 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 fails, the disk array status changes to *Unprotected*. 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 fails, replace it as soon as possible. If a *Hot Spare* disk was available, the controller may rebuild the data on the disk automatically. If a *Hot Spare* rebuild was not initiated by the controller, replace the failed disk then initiate a rebuild.

To initiate a rebuild, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disk(s) you wish to rebuild then press Enter.
5. If you are sure you want to rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

**Attention:** Data currently on the disk will be overwritten.

### Multiple-Disk Failures

If each failed disk belongs to a different mirrored pair, then the array will be in the *Unprotected* 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 (Initiate the rebuild operation after you replace the failed disks.)

When both members of a mirrored pair fail, the disk array status changes to *R/W Protected*. You will have to delete the disk array, replace the failed disks, and then recreate the disk array. You must then recreate the file systems on the disk array and copy data to the restored disk array from your backup media.

---

## Reclaiming IOA Cache Storage

IOA cache storage should only be reclaimed when indicated by a Maintenance Analysis Procedure (MAP) in Chapter 15, “Problem Determination and Recovery,” on page 173. This process is used to resolve cache problems.

**Attention:** Data may be lost.

To reclaim IOA cache storage, do the following:

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Reclaim IOA cache storage**.
4. Select the desired adapter, then press Enter.
5. A screen will display showing the devices that may be affected. If you are sure you want to reclaim IOA cache storage, type `c` to confirm. If you do not want to reclaim IOA cache storage, type `q` to cancel.

After the process is complete, one of the following screens will display:

```
Reclaim IOA Cache Storage Results

IOA cache storage reclamation has completed. The number of
lost sectors could not be determined.

Press Enter to continue.

e=Exit  q=Cancel
```

```
Reclaim IOA Cache Storage Results

IOA cache storage reclamation has completed. Use the number
of lost sectors to decide whether to restore data from the
most recent save media or to continue with possible data loss.

Number of lost sectors . . . . . :      1572608

Press Enter to continue.

e=Exit  q=Cancel
```



---

## Chapter 15. 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. For additional sources of related service procedures required to service this subsystem, see “Related Publications” on page xiii.

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

- **ibmsis** error log entries associated with the problem, viewed using the **Analyze Log** option in **sisconfig** or directly from **/var/log/messages**
- Disk array and physical disk status, viewed using the **sisconfig** utility

Entries in the error log will contain a System Reference Code (SRC) and a description of the problem. Each SRC corresponds to an action that should be taken. 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 Linux or system documentation.

---

### Error Log Analysis

Errors logged by the adapter and devices are stored in the **/var/log/messages** file. While the experienced user can search through this file with commands such as **grep** and **sed**, the **sisconfig** utility also offers a menu for analyzing the error log.

To analyze the error log, select the **Analyze Log** option from the main menu in **sisconfig**. A screen will display similar to the following:

Kernel Messages Log

Select one of the following:

1. Use vi to view most recent IBM Storage error messages
2. Use vi to view IBM Storage error messages
3. Use vi to view all kernel error messages
4. Use vi to view sisconfig error messages
5. Set root kernel message log directory
6. Set default editor
7. Restore defaults
8. Use vi to view IBM Storage boot time messages

Selection:

e=Exit q=Cancel

The **Set root kernel message log directory** option allows you to change the directory where the error log is stored. Normally, this does not need to be used. It may be useful when booting from an alternate root file system if you wish to access **/var/log/messages** on another file system.

If you are more familiar with a text editor other than **vi**, choose the **Set default editor** option to change the text editor used when viewing the error log. Invoking **sisconfig** with the **-e** command line option will also change the default editor. Execute **sisconfig -h** for more information.

If at any time you wish to return the location of the error log to **/var/log/messages** and set the default editor back to **vi**, choose the **Restore defaults** option.

All other options on this screen gather certain information from the error log and present it in the default editor. This is used mostly during Maintenance Analysis Procedures to search for System Reference Codes.

## Basic vi Commands

Several frequently used **vi** commands follow:

**h** Move cursor to the right (in case arrow keys have no effect)

**j** Move cursor down (in case arrow keys have no effect)

**k** Move cursor up (in case arrow keys have no effect)

**l** Move cursor to the left (in case arrow keys have no effect)

**CTRL-f**

Page down (forward)

**CTRL-b**

Page up (backward)

**nG** Go to line number **n**. If **n** is omitted, go to the last line in file.

**/pattern**

Search forward for **pattern**. After searching, type **n** to repeat the search.

**?pattern**

Search backward for **pattern**. After searching, type **n** to repeat the search.

**n** Repeat the previous search

**:q** Quit **vi** and return to the **sisconfig** utility

## Example Using vi

To search through the most recent errors for the SRC ending in 3400, do the following:

1. From the **Kernel Messages Log** screen, select the **Use vi to view most recent IBM Storage error messages** option.
2. Type **G** to move to the last line of the file.
3. Type **?3400** and press Enter to search backward for the pattern 3400.
4. If the first instance found is not part of an error log entry header, repeat the command by typing **n** as necessary.

## Sample Error Logs

When viewed through **sisconfig**, IBM Storage error messages are displayed in the following forms:

### Device error log entry

```
Sep 22 03:48:56: begin-entry*****
Sep 22 03:48:56: SRC: 6600 FFF3
Sep 22 03:48:56: Class: Permanent
Sep 22 03:48:56: Disk media format bad
Sep 22 03:48:56: IOA Location: PCI Bus: 64, Device: 12 Location P1-I2
Sep 22 03:48:56: Device Location: P1-I2/Z1-A5
Sep 22 03:48:56: Device Serial Number: 000012F0
Sep 22 03:48:56: Device Manufacturer: IBM
Sep 22 03:48:56: Device Machine Type and Model: ST3146807LC
Sep 22 03:48:56: Device FRU Number: 55P4081
Sep 22 03:48:56: Device EC Level: H31325
Sep 22 03:48:56: Device Part Number: 55P4083
Sep 22 03:48:56: Device Specific (Z1): 0626C50F
Sep 22 03:48:56: Device Specific (Z2): 0002
Sep 22 03:48:56: Device Specific (Z3): 02291
Sep 22 03:48:56: Device Specific (Z4): 0001
Sep 22 03:48:56: Device Specific (Z5): 22
Sep 22 03:48:56: Device Specific (Z6): H31325
Sep 22 03:48:56: IOASC: 0x03310000
Sep 22 03:48:56: PRC: 0x1104E275
Sep 22 03:48:56: Driver version: Ver. 1 Rev. 19.12 SMP
Sep 22 03:48:56: IOA Firmware version: 0309003B
Sep 22 03:48:56: Predictive Analysis Seeks/256 counter: 0
Sep 22 03:48:56: Predictive Analysis Sectors Read/256 counter: 0
Sep 22 03:48:56: SCSI Sense Data: F0 00 03 00 00 00 00 0A 00 00 00 00 31 00 0C
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
Sep 22 03:48:56: IOA Error Data:
Sep 22 03:48:56: Offset          0 1 2 3 4 5 6 7 8 9 A B C D E F
Sep 22 03:48:56: 00000000      00297C00 00000500 00000101 00000000
Sep 22 03:48:56: 00000010      3100E275 00000000 00000000 00000000
Sep 22 03:48:56: 00000020      00000000 00000000 00FFFFFF FFFFFFFF
```

## IOA error log entry

```

Sep 23 00:36:32: begin-entry*****
Sep 23 00:36:32: SRC: 5703 3400
Sep 23 00:36:32: Class: Informational
Sep 23 00:36:32: IOA detected device error
Sep 23 00:36:32: IOA Serial Number: 03000001
Sep 23 00:36:32: PCI Bus: 64, Device: 12 Location P1-I2
Sep 23 00:36:32: IOASC: 0x06290000
Sep 23 00:36:32: PRC: 0x10253380
Sep 23 00:36:32: Driver version: Ver. 1 Rev. 19.12 SMP
Sep 23 00:36:32: IOA Firmware version: 0309003B
Sep 23 00:36:32: IOA Error Data:
Sep 23 00:36:32: Offset          0 1 2 3 4 5 6 7 8 9 A B C D E F
Sep 23 00:36:32: 00000000      00000000 09A2C000 00000003 83000000
Sep 23 00:36:32: 00000010      80000000 80000000 09050435 30280000
Sep 23 00:36:32: 00000020      60008400 00000000 00000000 0006DDDD
Sep 23 00:36:32: 00000030      AABBBBBB 00000210 00BB0000 00000000
Sep 23 00:36:32: 00000040      00000000 0028B000 10000800 49000210
.
.
.
Sep 23 00:36:32: 000003B0      6A000050 005F82EC 00000030 A4E60F44
Sep 23 00:36:32: 000003C0      605C0050 01060409 003C0102 00000000
Sep 23 00:36:32: 000003D0      00000000 00000000
Sep 23 00:36:32: end-entry*****

```



## Device configuration error log entry

```
May 25 11:17:57: begin-entry*****
May 25 11:17:57: SRC: 5703 9092
May 25 11:17:57: Class: Permanent
May 25 11:17:57: Disk unit requires initialization before use
May 25 11:17:57: IOA Serial Number: 03060044
May 25 11:17:57: PCI Bus: 64, Device: 12 Location P1-I2
May 25 11:17:57: IOASC: 0x07279100
May 25 11:17:57: PRC: 0x1400B704
May 25 11:17:57: Driver version: Ver. 1 Rev. 19.2 SMP
May 25 11:17:57: IOA Firmware version: 0309001A
May 25 11:17:57: Device Errors Detected: 1
May 25 11:17:57: Device Errors Logged: 1
May 25 11:17:57: -----
May 25 11:17:57: Device 1:
May 25 11:17:57: IOA Location: PCI Bus: 64, Device: 12 Location P1-I2
May 25 11:17:57: Device Location: P1-I2/Z1-A1
May 25 11:17:57:      Vendor ID: IBMAS400
May 25 11:17:57:      Product ID: DDYS36M
May 25 11:17:57:      Serial Number: 000003F8
May 25 11:17:57: New Device Information:
May 25 11:17:57:      Vendor ID: IBMAS400
May 25 11:17:57:      Product ID: DDYS36M
May 25 11:17:57:      Serial Number: 000003F8
May 25 11:17:57: I/O Processor Information:
May 25 11:17:57:      Vendor ID: IBM
May 25 11:17:57:      Product ID: 5703001
May 25 11:17:57:      Serial Number: 03060044
May 25 11:17:57: Cache Adapter Card Information:
May 25 11:17:57:      Vendor ID: IBM
May 25 11:17:57:      Product ID: 5703001
May 25 11:17:57:      Serial Number: 03060044
May 25 11:17:57: Additional IOA Data: . . .
```

## Array error log entry

```
May 23 14:47:28: begin-entry*****
May 23 14:47:28: SRC: 5703 9030
May 23 14:47:28: Class: Permanent
May 23 14:47:28: Array no longer protected due to missing or failed disk unit
May 23 14:47:28: IOA Serial Number: 03060044
May 23 14:47:28: PCI Bus: 64, Device: 12 Location P1-I2
May 23 14:47:28: IOASC: 0x066B0200
May 23 14:47:28: PRC: 0x1705B004
May 23 14:47:28: Driver version: Ver. 1 Rev. 19.2 SMP
May 23 14:47:28: IOA Firmware version: 0309001A
May 23 14:47:28: -----
May 23 14:47:28: Array Member 0:
May 23 14:47:28:     Vendor ID: IBMAS400
May 23 14:47:28:     Product ID: DDYS36M
May 23 14:47:28:     Serial Number: 000003F8
May 23 14:47:28:     PCI/SCSI Address:
May 23 14:47:28:         Current: 40:0C/000100
May 23 14:47:28:         Expected: 40:0C/000100
May 23 14:47:28:     Physical Location:
May 23 14:47:28:         Current: P1-I2/Z1-A1
May 23 14:47:28:         Expected: P1-I2/Z1-A1
May 23 14:47:28: -----
May 23 14:47:28: Array Member 1:
May 23 14:47:28:     Vendor ID: IBMAS400
May 23 14:47:28:     Product ID: DGHS18U
May 23 14:47:28:     Serial Number: 00C1260A
May 23 14:47:28:     PCI/SCSI Address:
May 23 14:47:28:         Current: 40:0C/000200
May 23 14:47:28:         Expected: 40:0C/000200
May 23 14:47:28:     Physical Location:
May 23 14:47:28:         Current: P1-I2/Z1-A2
May 23 14:47:28:         Expected: P1-I2/Z1-A2
May 23 14:47:28: -----
May 23 14:47:28: Array Member 2:
May 23 14:47:28:     Vendor ID: IBMAS400
May 23 14:47:28:     Product ID: DGHS18U
May 23 14:47:28:     Serial Number: 00C124A7
May 23 14:47:28:     PCI/SCSI Address:
May 23 14:47:28:         Current: 40:0C/000400
May 23 14:47:28:         Expected: 40:0C/000400
May 23 14:47:28:     Physical Location:
May 23 14:47:28:         Current: P1-I2/Z1-A4
May 23 14:47:28:         Expected: P1-I2/Z1-A4
May 23 14:47:28: -----
May 23 14:47:28: end-entry*****
```

## Cache error log entry

```
May 27 15:00:27: begin-entry*****
May 27 15:00:27: SRC: 5703 9010
May 27 15:00:27: Class: Permanent
May 27 15:00:27: Cache data associated with attached devices cannot be found
May 27 15:00:27: IOA Serial Number: 03060044
May 27 15:00:27: PCI Bus: 64, Device: 12 Location P1-I2
May 27 15:00:27: IOASC: 0x07278100
May 27 15:00:27: PRC: 0x1400B23C
May 27 15:00:27: Driver version: Ver. 1 Rev. 19.2 SMP
May 27 15:00:27: IOA Firmware version: 0309001A
May 27 15:00:27: Current Configuration:
May 27 15:00:27:   I/O Processor Information:
May 27 15:00:27:       Vendor ID: IBM
May 27 15:00:27:       Product ID: 5703001
May 27 15:00:27:       Serial Number: 03060044
May 27 15:00:27:   Cache Adapter Card Information:
May 27 15:00:27:       Vendor ID: IBM
May 27 15:00:27:       Product ID: 5703001
May 27 15:00:27:       Serial Number: 03060044
May 27 15:00:27: Expected Configuration:
May 27 15:00:27:   I/O Processor Information:
May 27 15:00:27:       Vendor ID:
May 27 15:00:27:       Product ID:
May 27 15:00:27:       Serial Number:
May 27 15:00:27:   Cache Adapter Card Information:
May 27 15:00:27:       Vendor ID: IBM
May 27 15:00:27:       Product ID: 5703001
May 27 15:00:27:       Serial Number: 03060044
May 27 15:00:27: Additional IOA Data: 00000000 00000000 152101D0
```

The first portion of the log entry is the header, which contains important information such as the System Reference Code, the error class, the IOA under which the error occurred, and the hardware location. The last block of data in the entry is the IOA Error Data, which is data recorded by the IOA to assist with recovery, if the recommended Service Action did not solve the problem.

---

## Adapter Dump

Should the adapter ever experience a critical error that requires an adapter reset to recover, the **sisdump** utility will dump the contents of the adapter's memory into a dump file. This dump file contains information from the adapter and the device driver. This dump file is approximately 4 MB. **sisdump** is usually configured to start at boot time and run in the background until such an error occurs. Whenever a critical error occurs, **sisdump** logs a message in the system messages log containing the location and name of the dump file that corresponds to that error and the location of the affected adapter.

Critical errors require that IBM Hardware Service Support be contacted. They will need your error log in order to make a complete analysis and determine the cause of the critical error.

---

## Identifying the Disk Array Problem

A disk array problem is uniquely identified by a System Reference Code (SRC). An SRC is in the format *cccc uuuu*, where the first four digits of the SRC is the Card Type (for example 5703) and the last four digits of the SRC is known as the Unit Reference Code (URC). The URC indicates the specific problem that has occurred and must be obtained in order to determine which MAP to use.

An SRC is provided in the **ibmsis** error log with each error. This SRC describes the exact problem that has been detected and should be considered the primary means of identifying a problem. However, the **Display Disk Hardware Status** screen within the **sisconfig** utility is also a useful tool in identifying a problem or confirming a problem described by the error log. For additional information, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Obtain the SRC and proceed to the next section to obtain a more detailed description of the problem and to determine which Maintenance Analysis Procedure to use.

## Unit Reference Code (URC) Tables

Using the SRC obtained from the **ibmsis** error log, use the following table to determine which MAP to use. For additional details regarding the error log, see “Identifying the Disk Array Problem.” For MAPs, see “Maintenance Analysis Procedures (MAPs)” on page 184.

**Note:** The following table includes only URCs that are associated with Maintenance Analysis Procedures contained in this document.

If the Card Type part of the SRC (the first four digits) is 432x, 660x, 671x, or 673x, the error pertains to an attached device and you should refer to the table of Device URCs. Otherwise, the error pertains to the adapter and you should refer to the table of Adapter URCs.

Table 5. Device URCs

URC	Description Text	Error Class	Service Action	Failing Items
102E	Out of alternate sectors for disk storage	Permanent	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 199.	Disk drive
3002	Addressed device failed to respond to selection	Recoverable	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 199.	Disk drive
3010	Disk device returned wrong response to IOA	Permanent	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform “MAP 3351” on page 199.	Disk drive
3020	Storage subsystem configuration error	Permanent	Perform “MAP 3350” on page 196.	
3029	A device replacement has occurred	Informational	No action required.	
3109	IOA timed out a disk command	Recoverable	Perform “MAP 3350” on page 196.	
3110	Disk bus interface error occurred	Permanent	Perform “MAP 3350” on page 196.	
7000	Disk sector read error	Permanent	No action required.	

Table 5. Device URCs (continued)

URC	Description Text	Error Class	Service Action	Failing Items
7001	Temporary disk data error	Statistical	If two 7001 messages have occurred for the same disk drive location, then exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 199.	Disk drive
FFF3	Disk media format bad	Permanent	Perform "MAP 3335" on page 195.	
FFF4	Disk device problem	Permanent	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 199.	Disk drive
FFF5	Disk sector read error	Permanent	If two FFF5 messages have occurred for the same disk drive location in a one-day time period, then exchange the failing items in the Failing Items list one at a time.	Disk drive
FFF6	Disk device detected recoverable error	Statistical	No action required.	
FFF6	Disk device detected recoverable error	Threshold	Exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 199.	Disk drive
FFF7	Temporary disk data error	Temporary	No action required.	
FFF9	Temporary disk data error	Temporary	No action required.	
FFF9	Temporary disk data error	Statistical	No action required.	
FFFA	Temporary disk bus error	Statistical	If 10 Statistical FFFA messages have occurred for the same disk drive location in a one-week time period, then exchange the failing items in the Failing Items list one at a time. If the problem is not resolved, perform "MAP 3351" on page 199.	Disk drive
FFFB	SCSI bus reset occurred	Statistical	No action required.	
FFFE	Temporary disk bus error	Statistical	If 10 Statistical FFFE messages have occurred for the same disk drive location in a one-week time period, perform "MAP 3350" on page 196.	
FFFE	Temporary disk bus error	Informational	No action required.	

Table 6. Adapter URCs

URC	Description	Error Class	Service Action	Failing Items
302x	IOA detected a SCSI bus configuration error	Permanent	An error occurred on I/O bus x. Exchange the failing items in the Failing Items list one at a time.	Any device on I/O bus x I/O Adapter
310x	IOA detected interface error	Permanent	An error occurred on I/O bus x. Perform "MAP 3350" on page 196.	

Table 6. Adapter URCs (continued)

URC	Description	Error Class	Service Action	Failing Items
315x	SCSI bus configuration error	Permanent	Internal and external SCSI cables are connected to SCSI bus x at the same time. Perform “MAP 3350” on page 196.	
3400	IOA detected device error	Permanent, Recoverable	Exchange the failing items in the Failing Items list one at a time.	Device I/O Adapter Any device on I/O bus Signal cables
3400	IOA detected device error	Informational	No action required.	
34FF	Disk device format in progress	Informational	No action required.	
8008	A permanent cache battery pack failure occurred	Permanent	Perform “MAP 3300” on page 184. If the problem is not resolved, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8009	Impending Cache Battery Pack failure		Perform “MAP 3300” on page 184.	
813x	IOA detected recoverable device bus error	Statistical	No action required.	
8140	IOA detected recoverable device bus error	Statistical	No action required.	
8141	IOA detected recoverable device bus error	Temporary, Statistical	No action required.	
8145	A recoverable IOA error occurred	Threshold	Exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8146	Disk device detected recoverable error	Threshold	Exchange the failing items in the Failing Items list one at a time.	Disk Drive
8150	Permanent IOA failure	Permanent	Exchange the failing items in the Failing Items list one at a time.	I/O Adapter Any device on I/O bus
8151	IOA Licensed Internal Code error	Permanent	Update adapter microcode. (See “Updating the PCI-X SCSI RAID Controller Microcode” on page 158.) If the problem is not resolved, exchange the failing items in the Failing Items list one at a time.	I/O Adapter
8157	Hardware Error, IOA error requiring IOA reset to recover	Permanent	I/O Adapter was reset. If the error persists, replace the I/O Adapter.	I/O Adapter
9000	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device I/O Adapter
9001	IOA detected device error	Permanent	Perform “MAP 3337” on page 195.	

Table 6. Adapter URCs (continued)

URC	Description	Error Class	Service Action	Failing Items
9002	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device
				I/O Adapter
				Signal cables
				Backplane
				Any device on I/O bus
9008	IOA does not support functions expected by devices	Permanent	Perform "MAP 3330" on page 189.	
9010	Cache data associated with attached devices cannot be found	Permanent	Perform "MAP 3320" on page 188.	
9011	Cache data belongs to devices other than those attached	Permanent	Perform "MAP 3390" on page 200.	
9020 9021 9022	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3311" on page 186.	
9023	Array member(s) not at required resource address	Permanent	Perform "MAP 3312" on page 186.	
9024	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3390" on page 200.	
9025	Disk unit is not supported at its physical location	Permanent	Perform "MAP 3310" on page 185.	
9026	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3390" on page 200.	
9027	Array not functional due to present hardware configuration	Permanent	Perform "MAP 3313" on page 187.	
9028	Incorrect hardware configuration change has been detected	Permanent	Perform "MAP 3390" on page 200.	
9029	Incorrect hardware configuration change has been detected	Permanent	Perform "MAP 3390" on page 200.	
9030	Array no longer protected due to missing or failed disk unit	Permanent	Perform "MAP 3310" on page 185.	
9031	Array protection temporarily suspended	Temporary	Perform "MAP 3310" on page 185.	
9040	Array protection temporarily suspended	Temporary	No action required.  The array is synching. Wait until the synch is complete.	
9041	Background array parity checking detected and corrected errors	Temporary	Perform "MAP 3390" on page 200.	
9050	Required cache data cannot be located for a disk unit	Permanent	Perform "MAP 3331" on page 190.	
9051	IOA cache data exists for a missing or failed device	Permanent	Perform "MAP 3332" on page 192.	
9052	Cache data exists for a device that has been modified	Permanent	Perform "MAP 3390" on page 200.	

Table 6. Adapter URCs (continued)

URC	Description	Error Class	Service Action	Failing Items
9054	IOA resources not available due to previous problems	Permanent	Perform “MAP 3321” on page 189.	
9060	One or more disk pairs are missing from an array	Permanent	Perform “MAP 3311” on page 186.	
9061 9062	One or more disks are missing from an array	Permanent	Perform “MAP 3311” on page 186.	
9063	Maximum number of functional arrays has been exceeded	Permanent	Perform “MAP 3390” on page 200.	
9081	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device I/O Adapter
9082	IOA detected device error	Permanent	Exchange the failing items in the Failing Items list one at a time.	Device I/O Adapter Signal cables Backplane Any device on I/O bus
9090	Disk unit has been modified after the last known status	Permanent	Perform “MAP 3333” on page 193.	
9091	Incorrect hardware configuration change has been detected	Permanent	Perform “MAP 3333” on page 193.	
9092	Disk unit requires initialization before use	Permanent	Perform “MAP 3334” on page 193.	
FF3D	Soft IOA error recovered by the IOA	Statistical	If 16 Statistical FF3D messages have occurred for the same I/O Adapter physical location, then exchange the failing items in the Failing Items list one at a time.	I/O Adapter

## 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 “Unit Reference Code (URC) Tables” on page 180 to identify which MAP to use.

### MAP 3300

A permanent cache battery pack failure occurred (SRC *cccc*) or Impending Cache Battery Pack failure (SRC *cccc* 8009).

#### Step 3300-1

Are you here because of a permanent Cache Battery Pack failure (SRC *cccc* 8008)?

**NO** Go to Step 3300-2.

**YES** Go to Step 3300-3.

#### Step 3300-2

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.



**Note:** Ensure that you have a replacement Cache Battery Pack before proceeding.  
Force the Cache Battery Pack into an error state by following the instructions outlined in “Forcing a Rechargeable Battery Error” on page 166.

Go to Step 3300-3.

### Step 3300-3

Replace the rechargeable battery pack by following the instructions outlined in “Replacing the Rechargeable Cache Battery Pack” on page 166.

**Attention:** All cards are sensitive to electrostatic discharge.

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

## MAP 3310

Disk unit is not supported at its physical location (SRC cccc 9025), Array no longer protected due to missing or failed disk unit (SRC cccc - 9030), or Array protection temporarily suspended (SRC cccc - 9031).

### Step 3310-1

Identify the disk array by examining the error log and the **Display Device Parity Status** screen. For information regarding accessing the **Display Device Parity Status** screen, see “Viewing Parity Status” on page 150.

Go to Step 3310-2.

### Step 3310-2

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

**NO** Go to Step 3310-3.

**YES** Go to Step 3310-4.

### Step 3310-3

The affected disk array should have a status of either *Rebuilt* or *Active* due to the use of a *Hot Spare* disk. Create a new *Hot Spare* disk for the disk array by doing the following:

1. Identify the failed disk using the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147. The failed disk should have a status of *Not Operational* or *DPY/Failed*.
2. Remove the failed disk and replace it with a new disk that will be used as the *Hot Spare*. For further details, see “Replacing Physical Disks” on page 168.
3. The new disk must be formatted for *Advanced Function* in order to be used as a hot spare. If the new disk needs to be formatted, see “Formatting Disks for Use in PCI-X SCSI RAID Disk Arrays” on page 151.
4. Assign the new disk as a hot spare for the disk. For further details, see “Using Hot Spare Disks” on page 154.

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

### Step 3310-4

The *Failed* disk should be replaced by using the following procedure:

1. Identify the failed disk by using the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147. The failed disk should have a status of *Not Operational* or *DPY/Failed*.
2. Remove the failed disk, replace it with a new disk, and rebuild the disk unit data. For further details, see “Replacing Physical Disks” on page 168.

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

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

## MAP 3311

Array not functional due to present hardware configuration (SRC cccc 9020 / cccc 9021 / cccc 9022), One or more disk pairs are missing from an array (SRC cccc 9060), or One or more disks are missing from an array (SRC cccc 9061 / cccc 9062).

### Step 3311-1

Identify the disks missing from the disk array by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3311-2.

### Step 3311-2

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 missing disks and install them in the correct physical locations in the system. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ibmsis` command to remove the device driver.
  2. Use the `modprobe ibmsis` command to reload the device driver.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 18.

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

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

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

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3312

Array member(s) not at required resource address (SRC cccc 9023).

### Step 3312-1

Identify the disks that are not at their required physical locations by using the error log entry and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3312-2.

## Step 3312-2

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 missing disks and install them in the correct physical locations in the system. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ibmsis` command to remove the device driver.
  2. Use the `modprobe ibmsis` command to reload the device driver.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 153.  
**Attention:** All data on the disk array will be lost.
- Format the remaining members of the disk array, as follows:  
**Attention:** All data on the disk array will be lost.
  1. Run the **sisconfig** utility by typing `sisconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

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

## MAP 3313

Array not functional due to present hardware configuration (SRC cccc 9027).

### Step 3313-1

Identify the adapter and disks related to the error by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3313-2.

### Step 3313-2

Have the disks or adapter card been physically moved recently?

- |            |  |
|------------|--|
| <b>NO</b>  | Contact your service support organization. |
| <b>YES</b> | Go to Step 3313-3.                         |

### Step 3313-3

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. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ibmsis` command to remove the device driver.
  2. Use the `modprobe ibmsis` command to reload the device driver.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 153.  
**Attention:** All data on the disk array will be lost.
- Format the remaining members of the disk array, as follows:  
**Attention:** All data on the disk array will be lost.
  1. Run the **sisconfig** utility by typing `sisconfig`.
  2. Select **Work with disk unit recovery**.

3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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


## MAP 3320

Cache data associated with attached devices cannot be found (SRC cccc 9010).

### Step 3320-1

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

- NO**                      Go to Step 3320-3.
- YES**                     Go to Step 3320-2.



**A** Battery Pack Casing  
**B** Cache Directory Card

### Step 3320-2

Perform the following:

1. Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 171.
2. Using appropriate service procedures, remove the new adapter, reinstall the Cache Directory Card (see illustration in “Step 3320-1” on page 57) that came with the new adapter, and reinstall the new adapter. The original Cache Directory Card from the failed adapter should be kept with the failed 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 3320-3

Has the Cache Directory Card been physically moved recently (see illustration in “Step 3320-1” on page 57)?

- NO**                      Go to Step 3320-5.

**YES**                      Go to Step 3320-4.

### **Step 3320-4**

Restore Cache Directory Card(s) and adapter(s) back to their original configuration.

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

### **Step 3320-5**

If the system has been powered off for several days after an abnormal power-down, the Cache Battery Pack may be depleted. Do *not* replace the adapter nor the Cache Battery Pack.

Reclaim IOA cache storage. See "Reclaiming IOA Cache Storage" on page 40.

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

### **Step 3320-6**

1. 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 29.
  - The removable cache card from the *new* storage I/O adapter (This only applies to some 2780 I/O adapters. Refer to "Separating a Removable Cache Card From the Base Card" on page 24.

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

## **MAP 3321**

IOA resources not available due to previous problems (SRC cccc 9054).

### **Step 3321-1**

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 3330**

IOA does not support functions expected by devices (SRC cccc 9008).

### **Step 3330-1**

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see "Viewing the Status of Disks and Disk Arrays" on page 15.

Go to Step 3330-2.

### **Step 3330-2**

Have the adapter card or disks been physically moved recently?

**NO**                      Contact your service support organization.

**YES**                      Go to Step 3330-3.

### Step 3330-3

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. Then either IPL the system or logical partition (it may be necessary to boot into *Rescue* mode) or reload the device driver by performing the following:
  1. Use the `modprobe -r ibmsis` command to remove the device driver.
  2. Use the `modprobe ibmsis` command to reload the device driver.
- Format the remaining members of the disk array, as follows:

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

  1. Run the **sisconfig** utility by typing `sisconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

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


## MAP 3331

Required cache data cannot be located for a disk unit (SRC *cccc* 9050).

### Step 3331-1

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

- NO**                      Go to Step 3331-3.
- YES**                     Go to Step 3331-2.



- A** Battery Pack Casing  
**B** Cache Directory Card

## Step 3331-2

**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.

When continuing with the new adapter, perform the following:

1. Using appropriate service procedures, remove the new adapter, replace the Cache Directory Card (see illustration in “Step 3331-1” on page 59) with the original Cache Directory Card from the failed adapter, and reinstall the new adapter.
2. An error with SRC *cccc* 9010 will occur. Ignore this and continue with the next steps.
3. Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 171.
4. Using appropriate service procedures, remove the new adapter, reinstall the Cache Directory Card (see illustration in “Step 3331-1” on page 59) that came with the new adapter, and reinstall the new adapter. The original Cache Directory Card from the failed adapter should be kept with the failed 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 3331-3

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3331-4.

## Step 3331-4

Have the adapter card or disks been physically moved recently?

- NO**                      Contact your service support organization.
- YES**                     Go to Step 3331-5.

## Step 3331-5

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

- NO**                      Go to Step 3331-6.
- YES**                     Go to Step 3331-7.

## Step 3331-6

There are three possible ways to correct the problem. Perform only *one* of the following options that are listed by preference.

- Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 171.
- Delete the disk array. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 153.
- Format the remaining members of the disk array, as follows:

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

1. Run the **sisconfig** utility by typing **sisconfig**.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

### Step 3331-7

The adapter and disks, identified in the preceding steps, 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 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*.

## MAP 3332

IOA cache data exists for a missing or failed device (SRC cccc 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 3332-1

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3332-2.

### Step 3332-2

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

- |            |                    |
|------------|--------------------|
| <b>NO</b>  | Go to Step 3332-4. |
| <b>YES</b> | Go to Step 3332-3. |

### Step 3332-3

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 3332-4

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

- |            |                    |
|------------|--------------------|
| <b>NO</b>  | Go to Step 3332-5. |
| <b>YES</b> | Go to Step 3332-6. |

### Step 3332-5

Reclaim IOA cache storage. See “Reclaiming IOA Cache Storage” on page 171.



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

### Step 3332-6

Have the adapter card or disks been physically moved recently?

**NO** Contact your service support organization.

**YES** Go to Step 3332-7.

### Step 3332-7

The adapter and disks, identified in the preceding steps, 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 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*.

## MAP 3333

Disk unit has been modified after the last known status (SRC cccc 9090) or Incorrect hardware configuration change has been detected (SRC cccc 9091).

### Step 3333-1

IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:

1. Use the `modprobe -r ibmsis` command to remove the device driver.
2. Use the `modprobe ibmsis` command to reload the device driver.

Take action on any new errors which surface.

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

## MAP 3334

Disk unit requires initialization before use (SRC cccc 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 **Device Concurrent Maintenance** screen in **sisconfig** when concurrently removing and installing disks (see “Replacing Physical Disks” on page 168) 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 3334-1

Identify the affected disks by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 15.

Go to Step 3334-2.

### Step 3334-2

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

**NO** Go to Step 3334-4.

**YES** Go to Step 3334-3.

### Step 3334-3

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 3334-4

Have the disks or adapter card been physically moved recently?

**NO** Go to Step 3334-5.

**YES** Go to Step 3334-6.

### Step 3334-5

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

**NO** Go to Step 3334-7.

**YES** Go to Step 3334-6.

### Step 3334-6

There are three possible ways to correct this problem. Perform only *one* of the following three options:

- IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ibmsis` command to remove the device driver.
  2. Use the `modprobe ibmsis` command to reload the device driver.
- Restore the adapter and disks back to their original configuration. Then, IPL the system or logical partition (it may be necessary to boot into *Rescue* mode), or reload the device driver by performing the following:
  1. Use the `modprobe -r ibmsis` command to remove the device driver.
  2. Use the `modprobe ibmsis` command to reload the device driver.
- 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 3334-7

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

- Format the disks, as follows:

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

  1. Run the **sisconfig** utility by typing `sisconfig`.
  2. Select **Work with disk unit recovery**.
  3. Select **Initialize and format disk unit**.
  4. Select the drives you wish to format then press Enter.

- If the disks are members of a disk array, delete the disk array. See “Deleting a PCI-X SCSI RAID Disk Array” on page 18.

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

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

## MAP 3335

Disk media format bad (SRC cccc 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 3335-1

Identify the affected disk by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3335-2.

### Step 3335-2

Format the disks, as follows:

**Attention:** All data on the disks will be lost.

1. Run the **sisconfig** utility by typing `sisconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Initialize and format disk unit**.
4. Select the drives you wish to format then press Enter.

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

## MAP 3337

IOA detected device error (SRC cccc 9001).

### Step 3337-1

A device configuration error has been detected. The configuration sectors on the device may be incompatible with the current I/O adapter.

Has the I/O adapter been replaced with a different type of I/O adapter, or have the devices been moved from a different type of I/O adapter to this one recently?

**NO** Go to Step 3337-2.

**YES** Go to Step 3337-3.

### Step 3337-2

Contact your service support organization.

Exit this procedure.

### Step 3337-3

Identify the affected disk(s) by examining the error log and the **Display Disk Hardware Status** screen. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3337-4.

### Step 3337-4

If the I/O Adapter has been replaced by a different type of I/O Adapter, reinstall the original adapter.

Go to Step 3337-5.

### Step 3337-5

If the disks involved have been moved from a different type of I/O Adapter to this one, return them to their original I/O Adapter.

Go to Step 3337-6.

### Step 3337-6

**Attention:** The next step will cause data loss on the disks involved. If the data on these disks is important, perform back up procedures now.

Stop parity protection on the disk array(s) involved. For further details, see “Deleting a PCI-X SCSI RAID Disk Array” on page 153.

Go to Step 3337-7.

### Step 3337-7

Restore the initial hardware configuration that caused the error.

Go to Step 3337-8.

### Step 3337-8

Start parity protection on the disks then restore the saved data. For further details, see “Creating a PCI-X SCSI RAID Disk Array” on page 152.

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

## MAP 3350

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 Card Type 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.

**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 an 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 *Unprotected* or fail if functioning disks are removed. Also, additional problems may be created.

### Step 3350-1

Identify the SCSI bus on which the problem is occurring by examining the error log entry.

Go to Step 3350-2.

### Step 3350-2

Have changes been made to the SCSI configuration recently?

**NO** Go to Step 3350-5.

**YES** Go to Step 3350-3.

### Step 3350-3

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)
- Ensure the SCSI bus does not have multi-initiators (for example, set up for a high-availability configuration)

**Note:** Multi-initiator support is not provided at this time.

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 3350-5.

**YES** Go to Step 3350-4.

### Step 3350-4

1. Power off the system or logical partition.
2. Correct the problem.
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus-related failure occur?

**NO** Go to Step 3350-14.

**YES** Go to Step 3350-5.

### Step 3350-5

Determine if any of the disk arrays on the adapter are in an *Unprotected* state. For further details, see “Viewing Parity Status” on page 150.

Does any disk array have a state of *Unprotected*?

**NO** Go to Step 3350-7.

**YES** Go to Step 3350-6.

### Step 3350-6

1. Identify the failed disk(s) by first finding disk arrays with a state of *Unprotected*, then finding disks on those arrays with a state of *Failed*.
2. Remove the failed disk(s) from each *Unprotected* disk array. For further details, see “Replacing Physical Disks” on page 168.
3. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO** Go to Step 3350-14.

**YES** Go to Step 3350-7.

### Step 3350-7

Are there any non-essential removable media devices (such as Tape, CDROM, and DVDROM) on the SCSI bus?

**NO** Go to Step 3350-10.

**YES** Go to Step 3350-8.

### Step 3350-8

1. Power off the system or logical partition.
2. Remove one of the non-essential removable media devices.
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO** Go to Step 3350-9.

**YES** Go to Step 3350-7.

### Step 3350-9

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 3350-14.

### Step 3350-10

Are there any non-essential disks that are not disk array members (such as 512 byte/sector *JBOD* disks, *Hot Spare* disks, or array candidates) on the SCSI bus?

**NO** Go to Step 3350-13.

**YES**                      Go to Step 3350-11.

### **Step 3350-11**

1. Remove one of the non-essential disk devices. For further details, see “Replacing Physical Disks” on page 168.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SCSI bus related failure occur?

**NO**                      Go to Step 3350-12.

**YES**                      Go to Step 3350-10.

### **Step 3350-12**

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 3350-14.

### **Step 3350-13**

Contact your service support organization.

Exit this procedure

### **Step 3350-14**

1. Reinstall any good devices that were removed during the course of this MAP.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.
3. Take action on any other non-SCSI bus related errors if they exist.

## **MAP 3351**

Use the following to determine what other Field Replaceable Units (FRUs) besides a disk, may need to be replaced in order to solve a problem.

You are sent here when an *Advanced Function* disk, a physical disk in 522 bytes/sector format, was identified as the primary FRU to replace in order to resolve 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.
- 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 Card Type 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 fail or become *Unprotected* if functioning disks are removed. Also, additional problems may be created.

### Step 3351-1

Identify the SCSI bus on which the problem is occurring by examining the error log entry.

Go to Step 3351-2.

### Step 3351-2

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 or logical partition
2. Replace a component listed above
3. Power on the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

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

## MAP 3390

The problem that occurred is uncommon or complex to resolve. Information should be gathered and assistance obtained from your service support organization.

### Step 3390-1

Make a copy of the error log entry, as well as other error log entries that occurred around the same time. These entries are all in **/var/log/messages** file if the default error log path has not been changed.

Go to Step 3390-2.

### Step 3390-2

Collect the current disk array configuration. For further details, see “Viewing the Status of Disks and Disk Arrays” on page 147.

Go to Step 3390-3.

### Step 3390-3

Contact your service support organization.

Exit this procedure.



---

## Part 4. Appendixes



---

## Appendixes

---

### 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

<p>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。</p> <p>VCCI-A</p>
---

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

<p>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。</p>
---

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.

---

## Notices

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

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

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

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

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

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



---

# Index

## A

adding disk to existing array 19, 91, 153  
avoiding shock hazard ix

## B

battery  
    displaying information 32, 98, 165  
    error, forcing 33, 99, 166  
bus settings  
    viewing and changing 20, 92, 154

## C

Cache Battery Pack  
    replacing 33, 99, 166  
capacities, disk array 9, 78, 140  
checkpoints 54, 116, 184  
common tasks 85, 147  
controller microcode, updating 24, 96, 158  
creating disk array 22, 90, 152  
creating disks, hot spare 19, 92, 154

## D

deleting disks hot spare 20, 92, 154  
deleting, disk array 18, 91, 153  
description  
    RAID Level 0 5, 75, 137  
    RAID Level 10 7, 77, 139  
    RAID Level 5 6, 7, 76, 138  
disk array  
    adding disk to an existing 19, 91, 153  
    deleting 18, 91, 153  
disk array problems  
    identifying 50, 112, 180  
disk array settings  
    viewing 15, 85, 147  
disk array settings, viewing 15, 85, 147  
disk array states 11, 80, 142  
disk array, capacities 9, 78, 140  
disk arrays, definitions 4, 73, 135  
disk arrays, models 79, 141  
disk arrays, overview 10, 78, 140

## E

error log analysis 43, 105, 173  
error logs, samples 45, 107, 175

## F

formatting disks  
    *Advanced Function* 18, 89, 151  
    JBOD 18, 90, 152

formatting disks for disk array use  
    *Advanced Function* 18, 89, 151

## H

highlighting xiii  
hot spare  
    creating disks 19, 92, 154  
    deleting disks 20, 92, 154  
Hot Spare Disks 19, 92, 154

## I

I/O Adapter States 11, 81, 143  
IOA cache storage, reclaiming 40, 104, 171  
iprconfig, using 15, 85

## M

Maintenance Analysis Procedures (MAPs) 54, 116, 184  
maintenance tips 23, 95, 157  
MAP 3300 54, 116, 184  
MAP 3310 54, 117, 185  
MAP 3311 55, 118, 186  
MAP 3312 56, 118, 186  
MAP 3313 56, 119, 187  
MAP 3320 57, 120, 188  
MAP 3321 58, 121, 189  
MAP 3330 58, 121, 189  
MAP 3331 59, 122, 190  
MAP 3332 61, 124, 192  
MAP 3333 62, 125, 193  
MAP 3334 63, 125, 193  
MAP 3335 64, 126, 195  
MAP 3337 65, 127, 195  
MAP 3350 66, 128, 196  
MAP 3351 69, 131, 199  
MAP 3390 70, 132, 200  
models  
    disk arrays 79, 141  
    physical disk 80, 142

## O

overview, controller 73, 135  
overview, disk arrays 10, 78, 140

## P

parity settings 17, 88, 150  
physical disk, models 80, 142  
physical disks  
    installing 37, 102, 169  
    removing 37, 102, 169  
    replacing 36, 101, 168  
physical disks states 11, 80, 142

problem determination 43, 105, 173

## R

RAID Level

recovery procedures 38, 39, 103, 170

RAID level 0 5, 75, 137

RAID Level 0

recovery procedures 38, 102, 169

RAID level 10 7, 77, 139

RAID Level 10

recovery procedures 39, 103, 170

RAID level 5 6, 76, 138

RAID Level 5 7

RAID levels 5, 9, 74, 78, 136, 140

rechargeable battery error, forcing 33, 99, 166

rechargeable battery information, displaying 32, 98, 165

Rechargeable Battery Maintenance 32, 98, 165

rechargeable cache battery pack, replacing 33, 99, 166

recovery 43, 105, 173

recovery procedures

RAID Level 0 38, 102, 169

RAID Level 10 39, 103, 170

RAID Level 5 38, 39, 103, 170

related publications xiii

## S

shock hazard, avoiding ix

sisconfig, using 147

software, verifying installation 13, 83, 145

states

disk arrays 11, 80, 142

I/O Adapter 11, 81, 143

physical disks 11, 80, 142

stripe-unit size 9, 78, 140

summary, RAID level 9, 78, 140

## U

Unit Reference Code (URC) Table 50, 112, 180

URC

table 50, 112, 180

## V

verifying installation of controller software 13, 83, 145

vi commands 44, 106, 174

viewing

parity settings 17, 88, 150

viewing and changing, bus settings 20, 92, 154



---

# Readers' Comments — We'd Like to Hear from You

PCI-X SCSI RAID Controller  
Reference Guide for Linux

Publication No. SA23-1327-01

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

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

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

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

Please tell us how we can improve this book:

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

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

---

Name

---

Address

---

Company or Organization

---

Phone No.

Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

**BUSINESS REPLY MAIL**

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

Information Development  
Department H6DS-9561  
11501 Burnet Road  
Austin, TX 78758-3493



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

SA23-1327-01

